

Published in final edited form as:

Nature. 2008 November 13; 456(7219): 245–249. doi:10.1038/nature07538.

Associative learning of social value

Timothy EJ Behrens^{1,2,*}, Laurence T Hunt^{1,2,*}, Mark W Woolrich¹, and Matthew FS Rushworth^{1,2}

¹FMRIB Centre, University of Oxford

²Dept Experimental Psychology, University of Oxford

Abstract

Our decisions are guided by information learnt from our environment. This information may come via personal experiences of reward, but also from the behaviour of social partners^{1, 2}. Social learning is widely held to be *distinct* from other forms of learning in its mechanism and neural implementation; it is often assumed to compete with simpler mechanisms, such as reward-based associative learning, to drive behaviour³. Recently however, neural signals have been observed during social exchange reminiscent of signals seen in associative paradigms⁴. Here, we demonstrate that social information may be acquired using the *same* associative processes assumed to underlie reward-based learning. We find that key computational variables for learning in the social and reward domains are processed in a similar fashion, but in parallel neural processing streams. Two neighbouring divisions of the anterior cingulate cortex were central to learning about social and reward-based information, and for determining the extent to which each source of information guides behaviour. When making a decision, however, the information learnt using these parallel streams was combined within ventromedial prefrontal cortex. These findings suggest that human social valuation can be realised via the same associative processes previously established for learning other, simpler, features of the environment.

In order to compare learning strategies for social and reward-based information, we constructed a task in which each outcome revealed information both about likely future outcomes (reward-based information) and about the trust that should be assigned to future advice from a confederate (social information).

24 subjects performed a decision-making task requiring the combination of information from three sources (fig 1, methods and supplementary information): (i) the *reward magnitude* of each option (generated randomly at each trial); (ii) the likely correct response (blue or green) based on *their own experience of rewards* on each option; and (iii) the *confederate's advice*, and how trustworthy the confederate currently was. When a new outcome was witnessed, subjects could use this single outcome to learn in parallel about the likely correct action, and the trustworthiness of the confederate.

The investigation resembles previous experiments that have compared animate and inanimate conditions in different trials or experiments^{5,6}. Here, however, both sources of information were present on each trial outcome but the relevance of each was manipulated

Corresponding Author: Timothy Behrens, FMRIB Centre, John Radcliffe Hospital, Oxford, OX3 9DU, Email: behrens@fmrib.ox.ac.uk, Phone +441865 222782 .

*These authors contributed equally to this work

Author contributions. All four authors contributed to generating the hypothesis and designing the experiment. Where specific roles can be assigned, LTH collected the data, TEJB and LTH analyzed the data, TEJB and MWW built the model, MFSR supervised the project.

continuously allowing determination of both the fMRI signal and the behavioural influence associated with each source of information.

Optimal behaviour in this task requires the subject to track the probability of the correct action and the probability of correct advice independently, and to combine these two probabilities into an overall probability of the correct response (supplementary information). Computational models of *reinforcement learning* (RL) have had considerable success in predicting how such probabilities are tracked in learning tasks outside the social domain⁷. The simplest RL models integrate information over trials by maintaining and updating the expected value of each option. When new information is observed this value is updated by the product of the *prediction error* and the *learning rate*⁷. In our task, there are two dissociable prediction errors; the *reward prediction error* (actual reward - expected value), for learning about the correct option; and the *confederate prediction error* (actual - expected fidelity), for learning about the trustworthiness of the confederate. The optimal learning rate depends on the *volatility* of the underlying information source⁸⁻¹⁰. In volatile conditions, subjects should give more weight to recent information, using a *fast* learning rate. In stable conditions, subjects should weigh recent and historical information almost equally, using a *slow* learning rate. By ensuring that the correct option and the confederate's advice became volatile at different times, we ensured that the learning rate for these two sources of information varied independently. We used a Bayesian reinforcement learning (RL)⁸ model (supplementary info) to generate the optimal estimates of prediction error, volatility and outcome probability separately for each source of information (fig 1b,c,d).

We first sought to establish whether human behaviour matched predictions from the RL model. We used logistic regression to determine the degree to which subject choices were influenced by the optimally-tracked confederate and outcome probabilities, and by the difference in reward magnitudes between options. Parameter estimates for all three information sources were significantly greater than zero, and there was no significant difference in the degree to which subjects used reward and social information to determine their behaviour (fig 1e). Furthermore there was no significant effect either of subjects blindly following confederate advice without learning its value, or of subjects assuming that the confederate would behave in the same way as the previous trial (fig 1e). Hence subjects were able to integrate the fidelity of the confederate over many trials in an RL-like fashion.

We then investigated whether the fMRI signal reflected the model's estimates of prediction error and volatility, for both social and reward information, when subjects witnessed new outcomes. In the reward domain, neural responses have been identified that encode these key parameters^{8, 11-16}. Dopamine neurons in the ventral tegmental area (VTA) code reward prediction errors^{12, 13, 17}. Similar signals are reported in the dopaminergic striatum^{11, 18} and even in the VTA itself, when specialized strategies are used in human fMRI studies¹⁹. fMRI correlates of the learning-rate in the reward domain have been reported in anterior cingulate sulcus (ACCs). If humans can learn from social information in a similar fashion, it should be possible to detect signals that co-vary with the same computational parameters, but in the social domain.

We observed BOLD correlates of the confederate prediction error in dorsomedial prefrontal cortex (DMPFC) in the vicinity of the paracingulate sulcus, right middle temporal gyrus (MTG), and in the right superior temporal sulcus at the temporoparietal junction (STS/TPJ) (figure 2a). Equivalent signals were present in the left hemisphere at the same threshold, but did not pass the cluster extent criterion; similar effects were also found bilaterally in the cerebellum (supplementary information). Notably, these regions showed a pattern of activation similar to known dopaminergic activity in reward learning¹³, but for social information. Activity correlated with the probability of a confederate lie after the subject

decision but before the outcome was revealed (a prediction signal). When the subjects observed the trial outcome, activity correlated negatively with this same probability, but positively with the event of a confederate lie (Figure 2b). This signal reflects both components of a prediction error signal for social information: The outcome (lie or truth) minus the expectation (Figure 2b). These signals cannot be influenced by reward prediction errors as the two types of prediction error were decorrelated in the task design. The presence of this prediction error signal in the brain is a prerequisite for any theory of an RL-like strategy for social valuation.

We performed a similar analysis for prediction errors on reward information (reward minus expected reward). We found a significant effect of reward prediction error in the ventral striatum (figure 2c), the ventromedial prefrontal cortex, and anterior cingulate sulcus (see supplementary information). As in the social domain, we observed significant effects of all three elements of the reward prediction error (Figure 2d) (see supplementary information for discussion).

As previously demonstrated⁸, the volatility of action-outcome associations predicted BOLD signal in a circumscribed region of ACCs (figure 3a). This effect varied across people such that those whose behaviour relied more on their own experiences (supplementary information) showed a greater volatility related signal in this region (figure 3b). The volatility of confederate advice correlated with BOLD signal in a circumscribed region in the adjacent ACC gyrus (ACCg) (figure 3a). Subjects whose behaviour relied more on this advice showed greater signal change in this region (figure 3c). Notably, this double dissociation [reflected in a three way interaction between area (ACCs versus ACCg), volatility type (social versus outcome) and degree of reliance on social ($F_{1,20}=7.145$, $p=0.015$) or experiential information ($F_{1,20}=5.379$, $p=0.031$)] can be understood by reference to a dissociation in macaque monkeys. Selective lesions to ACCs but not ACCg impair reward-guided decision-making in the reward domain²⁰. In the social domain, male macaques will forego food to acquire information about other individuals^{21, 22}. Selective lesions to ACCg but not ACCs abolish this effect²³. We found that BOLD signals in these two regions reflect the respective values of the same outcome for learning about the two different sources of information.

Learning about reward probability from vicarious and personal experiences recruits distinct neural systems, but subjects combine information across both sources when making decisions (figure 1e). A ventromedial portion of the prefrontal cortex (VMPFC) has been shown to code such an expected value signal for the chosen action^{24, 25} during decision-making.

We computed two probabilities of reward on the subject's chosen option; one based only on experience and one based only on confederate advice. BOLD Signal in the VMPFC was significantly correlated with both probabilities (figures 4a and S4). However, there was subject variability in whether the VMPFC signal better reflected the reward probability based on outcome history or on social information. The extent to which the VMPFC data reflected each source of information (at the time of the decision) was predicted by the ACCs/ACCg response to outcome/social volatility (at the time when the outcomes were witnessed) (figure 4b,c).

Here, we have shown that the weighting assigned to social information is subject to learning and continual update via associative mechanisms. We use techniques that predict behaviour when learning from personal experiences to show that similar mechanisms explain behaviour in a social context. Furthermore, we demonstrate fundamental similarities between the neural encoding of key parameters for reward-based and social learning.

Despite employing similar mechanisms, distinct anatomical structures code learning parameters in the two domains. However, information from both is combined in ventromedial prefrontal cortex when making a decision.

By comparing the two sources of information, we find that social prediction error signals similar to those reported in dopamine neurons for reward-based learning are coded in the MTG, STS/TPJ and DMPFC. BOLD signal fluctuations in these regions are often seen in social tasks^{26, 27}, and in tasks which involve the attribution of *motive* to stimuli²⁸. Such activations have been thought critical in studies of theory of mind²⁸. That these regions should code quantitative prediction and prediction error signals about a confederate, lends more weight to the argument that social evaluation mechanisms are able to rely on simple associative processes.

A second crucial parameter in reinforcement learning models is the *learning rate*, reflecting the value of each new piece of information. In the context of reward-based learning, this parameter predicts BOLD signal fluctuations in ACC *sulcus* at the crucial time for learning⁸ - a finding that is replicated here. We further demonstrate that the exact same computational parameter, in the context of social learning, predicts BOLD fluctuations in the neighbouring ACC *gyrus*. This functional dissociation is mirrored by differences in the regions' anatomical connectivity. In the macaque monkey, connections with motor regions lie predominantly in ACCs²⁹, giving access to information about the monkey's own actions. Connections with visceral and social regions, including the STS, lie predominantly in ACCg²⁹, giving access to information about other agents. Nevertheless, that it is the same computational parameter that is represented in ACCs and ACCg, suggests that *parallel* streams of learning occur within ACC for social and non-social information.

It has been suggested that VMPFC activity might represent a common currency in which the value of different types of items might be encoded^{25, 30}. Here we show that the same portion of the VMPFC represents the expected value of a decision based on the combination of information from social and experiential sources. However, the extent to which the VMPFC signal reflects each source of information during a decision is predicted by the extent to which the ACCs and ACCg modulate their activity at the point when information is learnt. If, as is suggested, the VMPFC response codes the expected value of a decision, then the ACCs response to each new outcome predicts the extent that this outcome will determine future valuation of an action; the ACCg response predicts the extent to which this outcome will determine future valuation of an individual.

Methods Summary

(For detailed methods see supplementary information). Short Description of task (Figure 1a).

Subjects performed a decision-making task whilst undergoing fMRI, repeatedly choosing between blue and green rectangles, each of which had a different reward magnitude available on each trial. The chance of the correct colour being blue or green depended on the recent outcome history. Prior to the experiment, subjects were introduced to a confederate. At each trial, the confederate would choose between supplying the subject with the correct or incorrect option, unaware of the number of points available. The subject's goal was to *maximise* the number of points gained during the experiment. In contrast, the confederate's goal was to ensure that the eventual score would lie within one of two pre-defined ranges, known to the confederate but not the subject. The confederate might therefore reasonably give consistently helpful or unhelpful advice, but this advice might change as the game progressed (supplementary information). During the experiment, the confederate was

replaced by a computer that gave correct advice on a prescribed set of trials. Subjects knew that the trial outcomes were determined by an inanimate computer program, but believed that the social advice came from an animate agent's decision.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We would like to acknowledge funding from the UK MRC (TEJB,MFSR), the Wellcome Trust (LTH) and the UK EPSRC (MWW). Thanks to Steven Knight for helping with data acquisition, and Kate Watkins for help with figure preparation.

Appendix

Methods

Detailed analysis of the task, the learning model, the behavioural analysis, the data acquisition and pre-processing, and several further results and discussion can be found in the supplementary information. Here, we describe aspects of the fMRI modelling that may be relevant to the interpretation of our results. Further technical details can also be found in the supplementary information.

fMRI single subject modelling

We performed two fMRI GLM analyses using FMRIB's Software library (FSL³¹). The first looked for learning-related activity (figures 2, 3 and S3), the second for decision-related activity (figure 4 and S4). In each case a general linear model was fit in pre-whitened data space (to account for autocorrelation in the fMRI residuals)³². Regressors were convolved and filtered according to FSL defaults (see supplement).

Timeseries model (learning-related activity)

The following regressors (plus their temporal derivatives) were included in the model: 4 regressors defining the different times during the task (see figure 1 and supplement): CUE, SUGGEST, INTERVAL, MONITOR; 4 regressors defining key learning parameters when the outcomes are presented (see supplement): [MONITOR x REWARD HISTORY VOLATILITY], [MONITOR x CONFEDERATE VOLATILITY], [MONITOR x REWARD PREDICTION ERROR], [MONITOR x CONFEDERATE PREDICTION ERROR].

Timeseries model (decision-related activity)

The following regressors (plus their temporal derivatives) were included in the model: 4 regressors defining the different times during the task (see figure 1 and supplement): CUE, SUGGEST, INTERVAL, MONITOR; 7 regressors defining key decision parameters at the times when they were available during the decision (see supplement): [CUE x EXPERIENCE-BASED PROBABILITY], [SUGGEST x EXPERIENCE-BASED PROBABILITY], [SUGGEST x CONFEDERATE -BASED PROBABILITY], [CUE x CHOSEN REWARD MAGNITUDE], [SUGGEST x CHOSEN REWARD MAGNITUDE], [CUE x UNCHOSEN REWARD MAGNITUDE], [SUGGEST x UNCHOSEN REWARD MAGNITUDE]. Note that probabilities were log-transformed such that their linear combination in the GLM would approximate the optimal combination for behaviour (see

supplement). Figure 4a was generated using the mean ([1 1 1]) contrast of all probability-related regressors.

FMRI group modelling

FMRI group analyses were carried out using a GLM with 3 regressors: A group mean, the weight for reward history information based on each subject's behaviour (see supplement), the weight for confederate information based on each subject's behaviour (see supplement).

FMRI region of interest analyses (figure 2)

The following processing steps are illustrated schematically and described in more detail in the supplement (figure S2). Individual subject data were taken from ROIs defined by the group clusters. Data from each trial were upsampled and re-aligned to points in the trial corresponding to the onset of the 4 trial stages. Data were Z-normalised across trials at each time point in the trial. We then performed 2 general linear models *across trials* for both reward, and confederate prediction errors. This allowed us (i) to test at which points in the trial the data correlated with the prediction of reward, or the prediction of confederate fidelity, and (ii) to test at which points after the outcome the data correlated with the trial outcome, or actual confederate fidelity. A prediction error signal should comprise 3 parts. (i) a positive correlation with the prediction after the decision; (ii) a positive correlation with the trial outcome at the time of this outcome; (iii) a negative correlation with the prediction at the time of the outcome (as a prediction error is defined as the outcome minus the prediction).

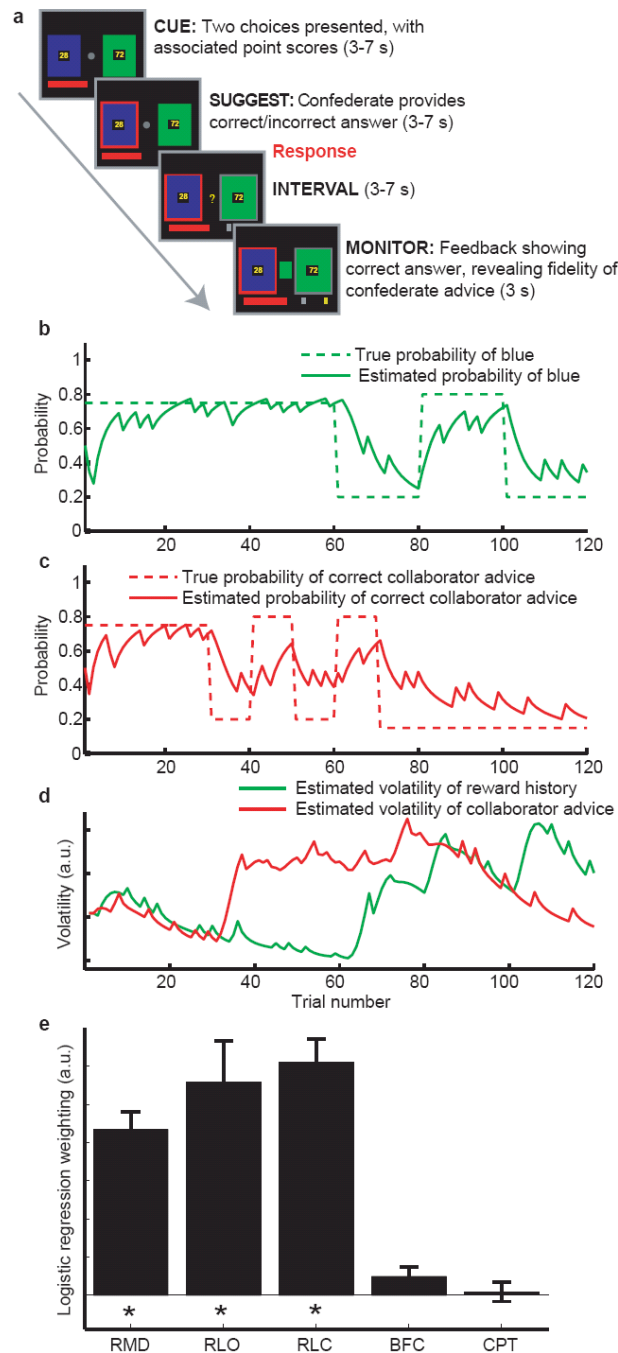
Statistical tests

We witnessed all 3 parts of the confederate prediction error as deflections in BOLD correlations at the relevant times. However, due to the nature of the haemodynamic response, it is difficult to test significance from just these deflections. We therefore fit a haemodynamic model to these correlation profiles in each subject (see supplement). The key test was whether the timecourse of correlations with the prediction could be accounted for by a positive haemodynamic impulse at the time of the decision and a negative haemodynamic impulse at the time of the outcome; and whether the timecourse of correlations with the outcome could be accounted for by a positive haemodynamic impulse at the time of the outcome. By fitting the haemodynamic model we were able to measure three parameter estimates for each of these three haemodynamic impulses in each subject, and perform random effects t-tests to measure statistical significance of each.

References

1. Fehr E, Fischbacher U. The nature of human altruism. *Nature*. 2003; 425:785–91. [PubMed: 14574401]
2. Maynard Smith, J. *Evolution and the Theory of Games*. Cambridge Univ. Press; 1982.
3. Delgado MR, Frank RH, Phelps EA. Perceptions of moral character modulate the neural systems of reward during the trust game. *Nat Neurosci*. 2005; 8:1611–8. [PubMed: 16222226]
4. King-Casas B, et al. Getting to know you: reputation and trust in a two-person economic exchange. *Science*. 2005; 308:78–83. [PubMed: 15802598]
5. Rilling J, et al. A neural basis for social cooperation. *Neuron*. 2002; 35:395–405. [PubMed: 12160756]
6. Gallagher HL, Jack AI, Roepstorff A, Frith CD. Imaging the intentional stance in a competitive game. *Neuroimage*. 2002; 16:814–21. [PubMed: 12169265]
7. Sutton, RS.; Barto, AG. *Reinforcement Learning: An introduction*. Cambridge: MIT Press; 1998.

8. Behrens TE, Woolrich MW, Walton ME, Rushworth MF. Learning the value of information in an uncertain world. *Nat Neurosci.* 2007; 10:1214–21. [PubMed: 17676057]
9. Courville AC, Daw ND, Touretzky DS. Bayesian theories of conditioning in a changing world. *Trends Cogn Sci.* 2006; 10:294–300. [PubMed: 16793323]
10. Dayan P, Kakade S, Montague PR. Learning and selective attention. *Nat Neurosci.* 2000; 3(Suppl): 1218–23. [PubMed: 11127841]
11. O'Doherty J, et al. Dissociable roles of ventral and dorsal striatum in instrumental conditioning. *Science.* 2004; 304:452–4. [PubMed: 15087550]
12. Schultz W, Dayan P, Montague PR. A neural substrate of prediction and reward. *Science.* 1997; 275:1593–9. [PubMed: 9054347]
13. Waelti P, Dickinson A, Schultz W. Dopamine responses comply with basic assumptions of formal learning theory. *Nature.* 2001; 412:43–8. [PubMed: 11452299]
14. Matsumoto M, Matsumoto K, Abe H, Tanaka K. Medial prefrontal cell activity signaling prediction errors of action values. *Nat Neurosci.* 2007; 10:647–56. [PubMed: 17450137]
15. Tanaka SC, et al. Prediction of immediate and future rewards differentially recruits cortico-basal ganglia loops. *Nat Neurosci.* 2004; 7:887–93. [PubMed: 15235607]
16. Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ. Cortical substrates for exploratory decisions in humans. *Nature.* 2006; 441:876–9. [PubMed: 16778890]
17. Bayer HM, Glimcher PW. Midbrain dopamine neurons encode a quantitative reward prediction error signal. *Neuron.* 2005; 47:129–41. [PubMed: 15996553]
18. Haruno M, Kawato M. Different neural correlates of reward expectation and reward expectation error in the putamen and caudate nucleus during stimulus-action-reward association learning. *J Neurophysiol.* 2006; 95:948–59. [PubMed: 16192338]
19. D'Ardenne K, McClure SM, Nystrom LE, Cohen JD. BOLD responses reflecting dopaminergic signals in the human ventral tegmental area. *Science.* 2008; 319:1264–7. [PubMed: 18309087]
20. Kennerley SW, Walton ME, Behrens TE, Buckley MJ, Rushworth MF. Optimal decision making and the anterior cingulate cortex. *Nat Neurosci.* 2006; 9:940–7. [PubMed: 16783368]
21. Deaner RO, Khera AV, Platt ML. Monkeys pay per view: adaptive valuation of social images by rhesus macaques. *Curr Biol.* 2005; 15:543–8. [PubMed: 15797023]
22. Shepherd SV, Deaner RO, Platt ML. Social status gates social attention in monkeys. *Curr Biol.* 2006; 16:R119–20. [PubMed: 16488858]
23. Rudebeck PH, Buckley MJ, Walton ME, Rushworth MF. A role for the macaque anterior cingulate gyrus in social valuation. *Science.* 2006; 313:1310–2. [PubMed: 16946075]
24. O'Doherty JP. Reward representations and reward-related learning in the human brain: insights from neuroimaging. *Curr Opin Neurobiol.* 2004; 14:769–76. [PubMed: 15582382]
25. Kable JW, Glimcher PW. The neural correlates of subjective value during intertemporal choice. *Nat Neurosci.* 2007; 10:1625–33. [PubMed: 17982449]
26. Amodio DM, Frith CD. Meeting of minds: the medial frontal cortex and social cognition. *Nat Rev Neurosci.* 2006; 7:268–77. [PubMed: 16552413]
27. Allison T, Puce A, McCarthy G. Social perception from visual cues: role of the STS region. *Trends Cogn Sci.* 2000; 4:267–278. [PubMed: 10859571]
28. Castelli F, Frith C, Happe F, Frith U. Autism, Asperger syndrome and brain mechanisms for the attribution of mental states to animated shapes. *Brain.* 2002; 125:1839–49. [PubMed: 12135974]
29. Van Hoesen, GW.; Morecraft, RJ.; Vogt, BA. Neurobiology of cingulate cortex and limbic thalamus. Vogt, BA.; Gabriel, M., editors. Boston: Birkhauser; 1993.
30. Plassmann H, O'Doherty J, Rangel A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J Neurosci.* 2007; 27:9984–8. [PubMed: 17855612]
31. Smith SM, et al. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage.* 2004; 23(Suppl 1):S208–19. [PubMed: 15501092]
32. Woolrich MW, Ripley BD, Brady M, Smith SM. Temporal autocorrelation in univariate linear modeling of FMRI data. *Neuroimage.* 2001; 14:1370–86. [PubMed: 11707093]

**Figure 1.**

Experimental task and behavioural findings. (a) Experimental task (See methods and Supplementary information). Each trial consists of four phases. Subjects are presented with a decision (CUE), receive the advice (red square) of the confederate (SUGGEST) and respond using a button press (grey square). An INTERVAL period follows, before the correct outcome is revealed (MONITOR). If the subject chooses correctly the red bar is incremented by the number of points on the chosen option. (b,c) Reward schedules for reward (b) and social (c) information. Dashed lines show the true probability of blue being correct (b) and the true probability of correct confederate advice (c). Each schedule underwent periods of stability and volatility. Solid lines show the model's estimate of the

probabilities. (d) Optimal model estimates of the volatility of reward (green) and social (red) information. (e) Logistic regression on subject behaviour. Factors included were: The reward magnitude difference between options (**RMD**); the outcome probability derived from the model using reward outcomes (**RLO**); the outcome probability derived from the model using confederate advice (**RLC**); the possibility that the subjects would blindly follow the confederate without learning (**BFC**); and the possibility that subjects would assume the confederate would behave as in the previous trial (**CPT**). The logistic regression analysis revealed significant effects only on **RMD**, **RLO** and **RLC**.

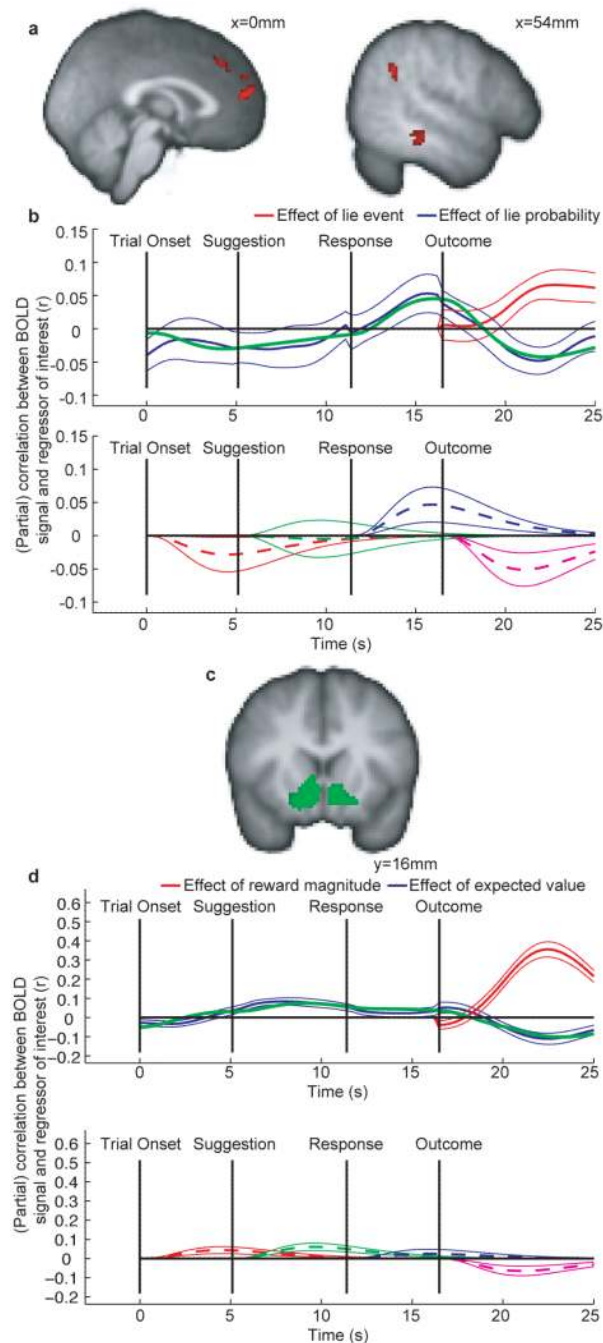


Figure 2.

Predictions and prediction errors in social and non-social domains. Timecourses show (partial) correlations \pm SEM. See figure S2. (a) Activation in the DMPFC, right TPJ/STS and MTG correlate with the social prediction error at the outcome (thresholded at $Z > 3.1$, cluster size > 50 voxels). (b) Deconstruction of signal change in the DMPFC. Similar results were found in the MTG and TPJ/STS. *Top panel:* Following the outcome, areas that encode prediction error correlate positively with the outcome and negatively with the predicted probability. Red: effect size of the confederate lie outcome (1 for lie, 0 for truth). Blue: effect size of the predicted confederate lie probability. To perform inference, we fit a hemodynamic model in each subject to the timecourse of this effect (i.e. to the blue line).

The green line in the top panel shows the mean overall fit of this hemodynamic model (for comparison with the blue line). *Bottom panel:* The effect of lie probability (blue line from top panel) is decomposed into an hrf at each trial event (fig S2). Dashed and solid lines show mean responses \pm s.e.m. Each region showed a significant positive effect of predicted confederate lie probability after the decision ($t(22)=1.96$ ($p<0.05$), 1.73 ($p<0.05$), 1.74 ($p<0.05$) for DMPFC, MTG and TPJ/STS respectively). Crucially, each brain region showed a significant negative effect of predicted confederate lie probability after the outcome ($t(22)=2.68$ ($p<0.005$), 2.35 ($p<0.05$), 3.27 ($p<0.005$)). (c) Ventral striatum is taken as an example of a number of regions revealed by the voxelwise analysis of reward prediction error (thresholded at $Z>3.1$, cluster size >100 voxels) (d) Panels are exactly as in (b), but coded in terms of reward and not in terms of confederate fidelity. Top panel shows the parameter estimate relating to the expected value of the trial (blue line) and, after the outcome, the parameter estimate relating to the magnitude of these rewards (grey line). To test for prediction error coding, we again fit a hemodynamic model to the expectation parameter estimate (shown by the green line, for comparison with blue line). Bottom panel: The timecourse showed a significant *positive* effect during the time of the decision ($t(22)=3.32$ ($p<0.002$)), and a *significant negative effect* after the trial outcome ($t(22)=2.50$, $p<0.05$) - see supplementary information for further discussion.

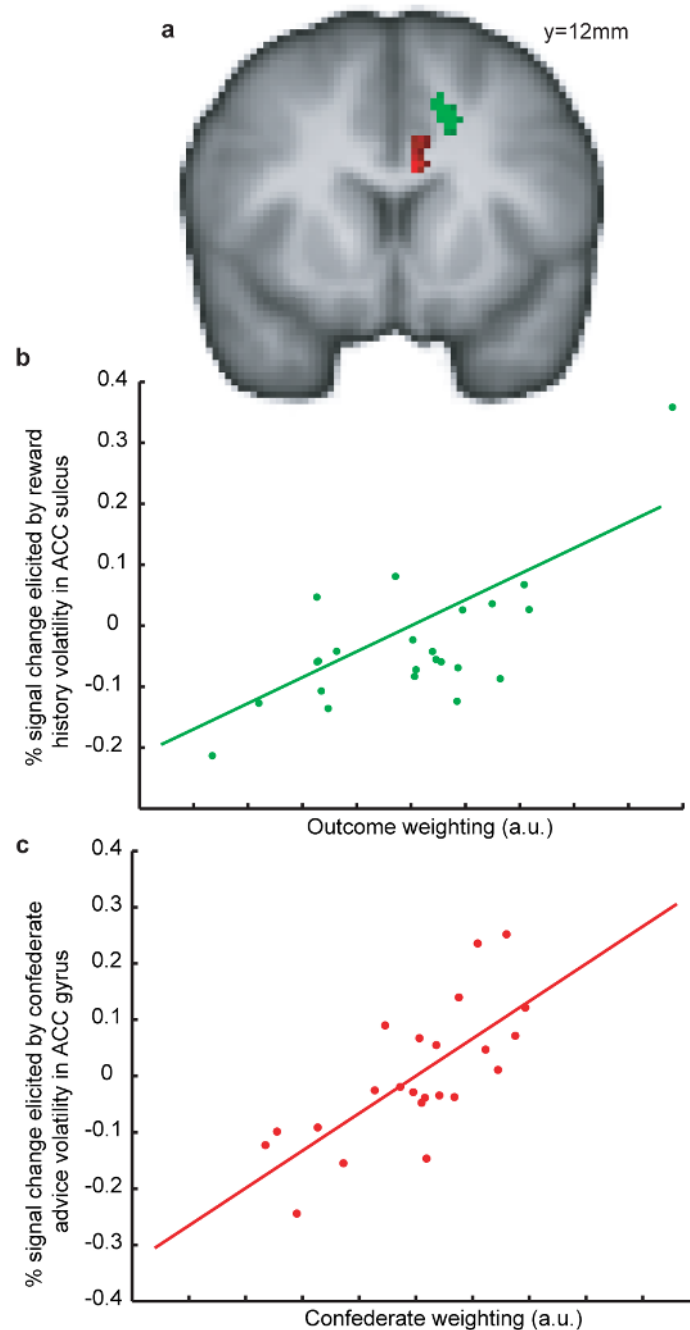


Figure 3.

Agency-specific learning rates dissociate in the ACC (a) Regions where the BOLD correlates of reward (green) and confederate (red) volatility predict the influence that each source of information has on subject behaviour ($Z > 3.1$, $p < 0.05$ cluster corrected for cingulate cortex). Subjects with high BOLD signal changes in response to reward volatility in the ACC *sulcus* are guided strongly by reward history information (max $Z = 3.7$, correlation ($R = 0.7163$, $p < 0.0001$) shown in (b)). Subjects with high BOLD signal changes in response to confederate advice volatility in the ACC *gyrus* are guided strongly by social information (max $Z = 4.1$, correlation ($R = 0.7252$, $p < 0.0001$) shown in (c)).

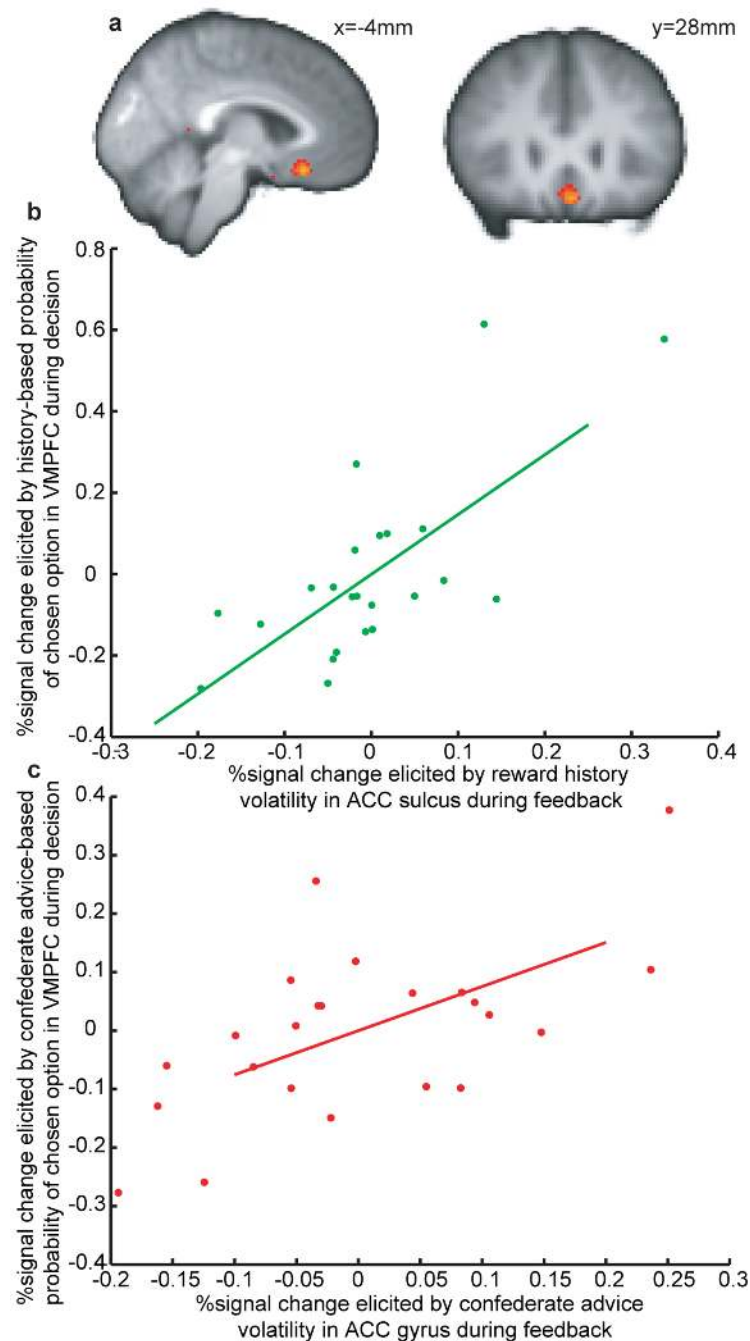


Figure 4.

Combination of expected value of chosen option in VMPFC. (a) Activation for the combination (mean contrast) of experience-based probability during CUE and SUGGEST phases, and advice-based probability during SUGGEST phase (thresholded at $Z > 3.1$, $p < 0.005$ cluster-corrected for VMPFC). These phases represent the times at which subjects had these probabilities available to them (see supplementary information and figure S4). (b) Correlation between effect of outcome-based probability in VMPFC during the decision and effect of outcome volatility in ACCs during MONITOR ($R = 0.6119$, $p < 0.0002$). (c) Correlation between effect of confederate-based probability in VMPFC during the decision and effect of confederate volatility in ACCs during MONITOR ($R = 0.6119$, $p < 0.0002$).