TITLE

Neurocomputational mechanisms of confidence in self and others

AUTHORS

Dan Bang[1,2,*], Rani Moran[1,3], Nathaniel D. Daw[4] & Stephen M. Fleming[1,3,5,*]

AFFILIATIONS

[1]Wellcome Centre for Human Neuroimaging, University College London, London WC1N 3BG, UK

[2]Department of Experimental Psychology, University of Oxford, Oxford OX1 3UD, UK

[3]Max Planck UCL Centre for Computational Psychiatry and Ageing Research, University College London, London WC1B 5EH, UK

[4]Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, New Jersey 08544, USA

[5]Department of Experimental Psychology, University College London, London WC1H 0AP, UK

*Correspondence: danbang.db@gmail.com and stephen.fleming@ucl.ac.uk

ABSTRACT

Computing confidence in one's own and others' decisions is crucial for success in many social situations. While there has been substantial progress in our understanding of confidence in oneself, little is known about how we form confidence about others. Here, we address this question by asking subjects undergoing fMRI to place bets on perceptual decisions made by themselves or one of three other players of varying ability. We show that subjects compute confidence in another player's decisions by combining distinct estimates of player ability and decision difficulty – allowing them to predict that a good player may get a difficult decision wrong and that a bad player may get an easy decision right. We find that this computation is supported by an interaction between brain systems traditionally implicated in decision-making (LIP) and theory of mind (TPJ, dmPFC). These results reveal a neurocomputational interplay between self- and other-related processes during social cognition.

INTRODUCTION

A key feature of adaptive behaviour is an ability to estimate confidence in our decisions in the absence of immediate feedback. For example, we may recognize that a choice was based on weak evidence and change our mind as conflicting but stronger evidence comes to light[1,2]. Recent years have seen substantial progress in our understanding of decision confidence. At a computational level, an emerging consensus is that our confidence in a choice reflects the probability that the choice is correct as computed using Bayesian inference[3–7]. At a neural level, converging evidence points to a central role for prefrontal cortex in the computation of decision confidence [8–10], the mapping of this internal estimate onto explicit reports[11,12], and the integration of single-trial confidence estimates into a long-run estimate of task performance[13].

Humans are, however, highly social and many situations require that we compute confidence not only in our own decisions but also in those of others. In competitive tasks, having a sense of who is more likely to succeed is useful for deciding whether to compete or opt out[14]. In cooperative tasks, having a sense of who is more likely to be correct is critical for assigning appropriate weight to others' advice[15,16], or deciding whose opinion to follow[17]. In addition, the confidence that others report is often biased[6], and can serve to manage social influence[12,18–20], and it is therefore useful to independently verify such reports. However, in contrast to the case of confidence in one's own decisions, little is known about the neurocomputational basis of confidence in decisions made by others.

Intuitively, the starting point for a social confidence computation is the same as in the individual case[7]. We should first assess the difficulty of the decision: the easier the decision is, the higher the probability that others will get it right. However, it is typically not sufficient to simply simulate what choice we would have made and how confident we would have felt about this choice. Instead, we should recognise that others' ability may be different from our own: what is hard for us may be easy for others and what is easy for us may be hard for others[7]. Previous studies have investigated how we track others' ability[15,21,22], but it remains unknown whether, or how, we integrate knowledge about ability with our own assessment of decision difficulty when computing confidence in others' choices. Critically, it is only by taking both of these factors into account that we can predict that a high-ability individual may get a difficult decision wrong and that a low-ability individual may get an easy decision right.

This requirement for integrating estimates of both decision difficulty and others' ability predicts that multiple brain systems support confidence in others' decisions. The first component, the estimation of decision difficulty, is computationally similar in both the individual and the social case and may therefore be supported by similar neural substrates. For example, in the context of perceptual decision-making, where decision difficulty is varied by changing the uncertainty associated with a sensory stimulus, we would expect areas that support sensory processing – such as visual motion area MT+ and lateral intraparietal sulcus (LIP)[8,23] – to be recruited when computing confidence in others' decisions. However, sensory representations on their own are unlikely to be sufficient because they do not reflect the characteristics of others' sensory systems. Instead, the second component to a social confidence computation, the estimation of others' ability, may be furnished by the so-called theory of mind (ToM) network – including temporoparietal junction (TPJ) and dorsomedial prefrontal cortex (dmPFC)[24–26] – believed to support the representation of others' attributes as distinct from one's own. In

support of this prediction, human fMRI has indicated that both TPJ and dmPFC are involved in the maintenance of running estimates of others' task performance[22,27,28].

Here, we develop a novel paradigm to disentangle the components of a social confidence computation, and we use behavioural analysis together with computational modelling to show that people take into account both decision difficulty and others' ability. More specifically, we asked subjects to place bets (post-decision wagers) on decisions on a random dot motion task[29] made by either themselves (self-trials) or one of three other players (other-trials) who differed in terms of their ability. To maximise reward on other-trials, subjects should combine an estimate of decision difficulty with distinct estimates of the other players' abilities. In this way, subjects can predict that a good player may get a difficult stimulus wrong and that a bad player may get an easy stimulus right. We refer to this account as the ToM-model and we show that it provides a better fit to subjects' post-decision wagers than a set of alternative models (e.g., relying solely on self-projection or long-run estimates of others' task performance). In addition, analysis of fMRI data acquired during task behaviour indicated that the combination of distinct estimates of decision difficulty and others' ability was supported by an interaction between brain systems classically involved in decision-making (LIP) and ToM (TPJ and dmPFC). Overall, our results reveal a neurocomputational interplay between self- and other-related processes when forming confidence about others' decisions.

RESULTS

Experimental paradigm

Subjects ($N = 21$) placed post-decision wagers (PDWs) on either their own or another player's choices on a random dot motion task while undergoing fMRI (**Figure 1**). At the start of a trial, a screen indicated whether they had to perform a self-trial or an other-trial. On both trial types, subjects viewed a field of moving dots inside a circular aperture. On each update of the motion stimulus, a fraction of dots moved coherently to the left or the right (coherence range: .005-.5), whereas the remainder moved randomly. On self-trials, subjects had to actively decide whether the average direction of dot motion was to the left or the right. On other-trials, subjects watched the motion stimulus while the other player decided; subjects were not told which decision the other player made. On both trial types, a PDW screen then appeared requiring subjects to choose between a safe and a risky option. The safe option delivered a small but certain reward (safe range: 5-20 points). In contrast, the risky option delivered a greater reward if the self- or other-choice was correct and a corresponding loss if incorrect (risky range: 25-50 points). Each PDW thus provided an incentive-compatible metric of subjects' confidence in a choice being correct. Finally, subjects received feedback about the PDW outcome – subjects could infer the accuracy of the self- or other-choice from this feedback. Subjects were paired with three players of varying average choice accuracy (low: .55; medium: .75; high: .95) in a block-wise manner. The other players were created using pilot data from subjects who had performed the random dot motion task without a social component. Subjects were told that the other players were not currently present and had made their decisions at another time. We highlight that our three main factors – motion coherence, others' ability and the reward difference between the risky and the safe option – were, by design, uncorrelated.

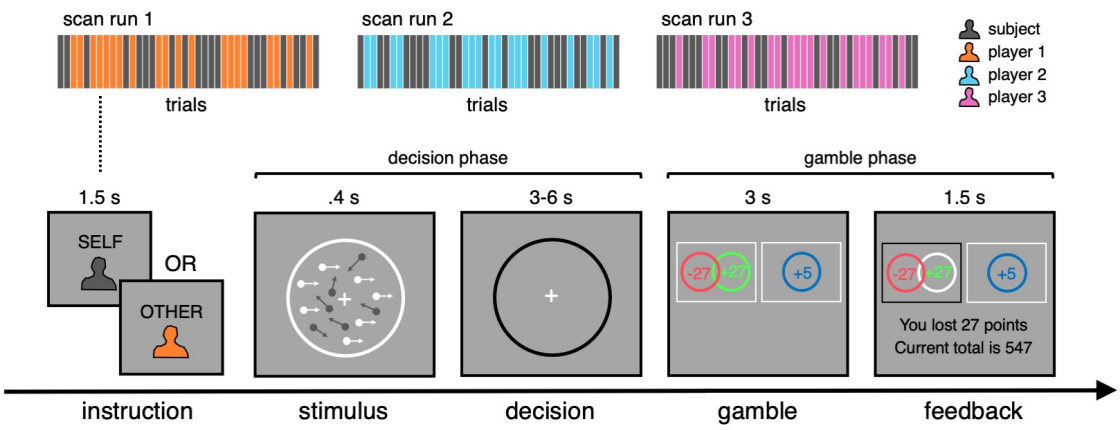

Figure 1. Social random dot motion task. On each trial, subjects placed a bet on a motion discrimination judgement (left or right) made by either themselves or one of three other players who varied in terms of their ability. The bet was made by choosing between a safe option, which yielded a small but guaranteed reward, and a risky option, which delivered a greater reward if the self- or other-choice was correct and a corresponding loss if incorrect. In this example, the subject chose the risky option (black box), but the self- or other-choice was incorrect. If subjects chose the safe option, then they were informed what the outcome would have been had they selected the risky option. Subjects were instructed how to calculate choice accuracy from the feedback screen. Subjects were paired with the three other players in a block-wise manner, and the trial type was signalled at the start of a trial. For fMRI analysis, we refer to the time window from the onset of the motion stimulus to 3 s after the onset of the decision screen as the decision phase and the time window from the onset of the gamble screen to the offset of the feedback screen as the gamble phase.

Post-decision wagering on self- and other-trials

We first sought to characterise subjects' PDWs behaviourally. Intuitively, the probability that subjects select the risky option should depend on the confidence that a self- or other-choice is correct and the difference between the rewards available under the risky and the safe option. The higher the confidence, the higher the probability that the risky option returns a reward, and the larger the reward difference, the higher the potential earnings from selecting the risky option. Importantly, self- and other-trials should differ in terms of the factors that enter into a confidence computation: on both trial types, confidence should take into account the difficulty of the decision, but on other-trials, confidence should also take into account the ability of the other player. To disentangle these contributions, we used a multiple logistic regression model to predict subjects' trial-by-trial PDWs, with predictors including trial type (other = -.5, self = .5), motion coherence, reward difference, others' ability (low = -.5, medium = 0, high = .5) and the interactions between trial type and the three other terms (see **Figure S1** for regression coefficients under full model and separate models for self- and other-trials).

This analysis showed that the probability that subjects selected the risky option increased with both motion coherence (**Figure 2A**; coherence: $t(20) = 7.78$, $p < .001$) and reward difference (**Figure 2B**; reward: $t(20) = 5.02$, $p < .001$), and that the impact of these factors did not depend on trial type (coherence x type: $t(20) = 0.18$, $p = .858$; reward x type: $t(20) = -0.44$, $p = .664$). The probability that subjects selected the risky option also increased with others' ability (ability: $t(20) = 2.85$, $p = .010$), but this effect depended on trial type (type x ability: $t(20) = -4.60$, $p < .001$). Importantly, there was no effect of others' ability on self-trials (blue bars in **Figure 2C**; ability under separate model for self-trials: $t(20) = -1.46$, $p = .161$), whereas on other-trials, the higher the ability of the other player, the higher the probability that subjects selected the risky option (yellow bars in **Figure 2C**; ability under separate model for other-trials: $t(20) = 4.52$, p

< .001). Overall, subjects were more likely to select the risky option on self- than other-trials (type: $t(20) = 4.47$, $p < .001$) – with subjects on average selecting the risky option on 72% of self-trials and 61% of other-trials.

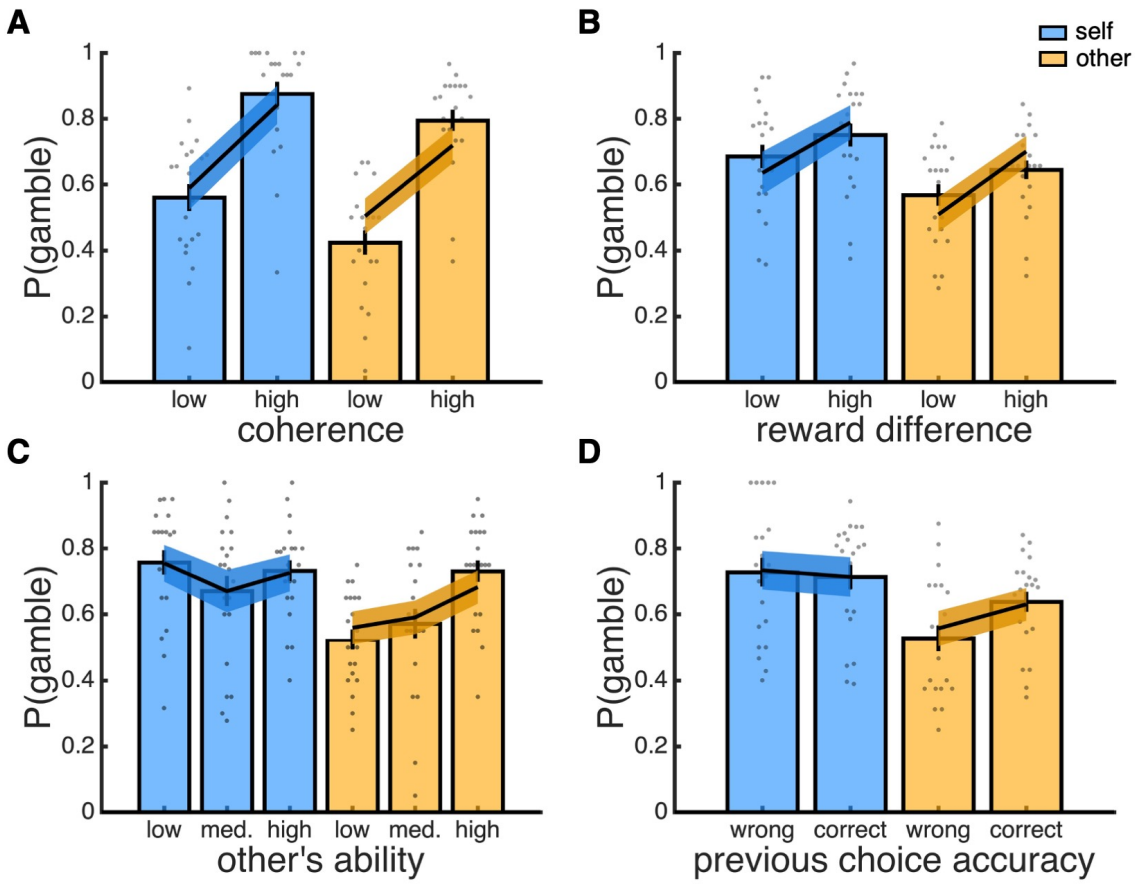

Figure 2. Behavioural results. In each panel, data are split into self-trials (blue) and other-trials (yellow). **(A)** Probability gamble as a function of coherence (median split). **(B)** Probability gamble as a function of the reward difference between the risky and the safe option (median split). **(C)** Probability gamble as a function of the ability of the other player. **(D)** Probability gamble as a function of choice accuracy on the previous trial of the same type. **(A-D)** Bar charts are empirical data, with individual subjects overlaid as dots. Lines are data simulated under the best-fitting model (ToM-model). Empirical and simulated data are represented as group mean ± 95% CI.

Critically, subjects were not informed about the ability of each player but had to learn it from experience. An obvious learning signal in our task is the feedback about choice accuracy available at the end of each trial. In line with this hypothesis, subjects were more likely to gamble on another player's choice when the player's previous choice was correct compared to when it was incorrect (yellow bars in **Figure 2D**; paired *t*-test: $t(20) = 2.97$, $p = .008$) – a history effect which was not present for self-trials (blue bars in **Figure 2D**: paired *t*-test, $t(21) = -0.35$, $p = .733$). Taken together, these results show (1) that subjects were equally sensitive to decision difficulty and reward characteristics on self- and other-trials and (2) that subjects tracked the abilities of the other players and used this knowledge to guide their PDWs in a player-dependent manner. We note that there was no effect of others' ability on subjects' task performance on self-trials, which we assessed using choice accuracy (logistic regression; coherence: $t(20) = 7.90$, $p < .001$; ability: $t(20) = -0.78$, $p = .444$) and reaction time (linear regression; coherence: $t(20) = -11.98$, $p < .001$; ability: $t(20) = 0.43$, $p = .674$).

Computational model of confidence in self and others

Qualitatively, the behavioural results fit with a ToM account in which the confidence in another player's decisions is based on both a self-related component, which reflects decision difficulty, and an other-related component, which reflects the ability of the other player. In order to test this interpretation formally, and to derive markers of the latent variables that underpin subjects' behaviour for fMRI analysis, we compared candidate models of subjects' PDWs.

In our task, subjects should decide whether to gamble on a choice by comparing the expected value of the risky and the safe option, $\Delta\mathrm{EV} = \mathrm{EV(risky)} - \mathrm{EV(safe)}$. We modelled the mapping between $\Delta\mathrm{EV}$ and gamble choices using a softmax function, $\mathrm{P(gamble)} = \mathrm{softmax}(\Delta\mathrm{EV})$. We allowed the parameters of this function to vary between trial types to accommodate individual variation in the tendency to gamble on choices made by oneself versus others. The expected value of the safe option is simply the amount offered, $\mathrm{EV(safe)} = \mathrm{V(safe)}$. By contrast, the expected value of the risky option is computed by weighting the risky offer according to the confidence in a choice being correct, $\mathrm{EV(risky)} = \mathrm{P(correct)} \times \mathrm{V(risky)} - [1 - \mathrm{P(correct)}] \times \mathrm{V(risky)}$. Here, we first consider the computation of confidence on self-trials, before turning to other-trials (models are described in mathematical detail in the **Methods**).

Building on earlier work on decision confidence[6,7], we modelled subjects' confidence on self-trials using Bayesian decision theory (self-trials in **Figure 3A**). In particular, we assumed that subjects represent the stimulus space as comprised of a set of distinct motion stimuli, each defined by a direction and a coherence. In the first version of the model, this representation consisted of linearly spaced motion stimuli. In a modified version, we allowed for an over-representation of low-coherence motion stimuli, consistent with a Weber-Fechner law in which the resolution of sensory perception diminishes for stimuli of greater magnitude. On each trial, subjects receive noisy sensory evidence, modelled as a sample from a Gaussian distribution centred on the true motion stimulus. Subjects then compute a belief state over stimulus space – that is, the probability that each motion stimulus is shown on the current trial – given the sensory evidence and their own level of sensory noise. Finally, subjects use this belief state to generate both a decision and a confidence in this decision being correct.

There are a number of ways in which a ToM account of a social confidence computation could be implemented. For consistency with earlier work on decision confidence[6,7], here we directly extended the Bayesian framework to the social case (see Discussion for an alternative heuristic ToM account). This ToM-model assumes that subjects compute confidence in the other player's decision by integrating their belief state over stimulus space with a distinct representation of the other player's expected accuracy for each motion stimulus (other-trials in **Figure 3A**). The ToM-model integrates these representations across the full stimulus space because subjects do not know whether the other player decided left or right. The ability representation – equivalent to a belief about the other player's psychometric function – is directly derived from an estimate of the other player's sensory noise. This estimate is updated at the end of each trial according to the difference between the observed and the predicted success of the other player's choice – a social prediction error.

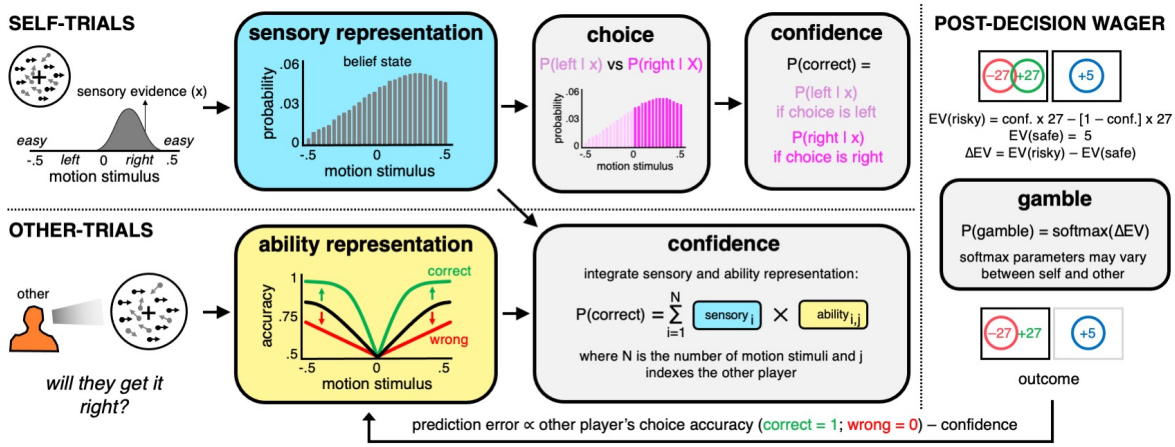

Figure 3. Illustration of a confidence computation for self and others under the ToM-model. All models assumed that subjects' choices and confidence were computed on self-trials according to Bayesian decision theory. Subjects represent the stimulus space as comprised of a set of distinct motion stimuli, each defined by a direction (sign) and a coherence (absolute value). On each trial, subjects receive sensory evidence sampled from a Gaussian distribution whose mean is given by the true motion stimulus and whose standard deviation is given by the subject's sensory noise. Subjects compute a belief state over the stimulus space given the sensory evidence and their sensory noise (blue box) and use this belief state to generate both a decision about the motion direction and their confidence in this decision being correct. The ToM-model assumes that subjects compute confidence in others' choices by combining their belief state over stimulus space with a representation of the other player's expected accuracy for each motion stimulus (yellow box). The latter representation – in effect, a psychometric function – is derived from an estimate of the other player's sensory noise. This estimate is updated at the end of each trial using a Rescorla-Wagner learning rule that takes into account the accuracy of the other player's choice and the confidence in this choice being correct. As a result of this update, the psychometric function becomes steeper after a correct choice (green) and shallower after an incorrect choice (red). On both trial types, confidence is used to calculate the difference in expected value between the risky and the safe options. Finally, the difference in expected value is entered into a softmax function to obtain the probability of selecting the risky option. The softmax parameters were allowed to vary between self- and other-trials.

Under a ToM account, confidence on other-trials is influenced by both decision difficulty and others' ability. For completeness, and to rule out extraneous explanations of the behavioural results, we also considered models that only included one of these two components. First, we considered a self-projection model (S-model), which assumes that subjects first ask themselves which choice they would have made and how confident they would have felt about this choice and then project this confidence estimate on to the other player. Under the S-model, confidence on other-trials is influenced by decision difficulty but not by others' ability. Second, building on previous work[21], we considered a performance-tracking model (Q-model), which assumes that subjects use a Rescorla-Wagner rule to learn about the value of each player based on their historical choice accuracy and use this value as a proxy for confidence in their decisions. Under the Q-model, confidence on other-trials is influenced by others' ability but not by decision difficulty.

In keeping with the behavioural results, and in support of a ToM account of a social confidence computation, the ToM-model provided the best fit to subjects' PDWs – with the winning model including both a Weber-Fechner law and separate softmax parameters for self and other (see **Figure S2** for model comparison results). As shown in **Figure 2**, the ToM-model captured all qualitative features of the behavioural data (see **Figure S3** for behavioural predictions under

the best-fitting version of each model class). In addition, the ToM-model captured the trial-by-trial evolution of subjects' PDWs for each of the three other players (**Figure S4**).

We performed two control analyses to assess the robustness of our modelling approach. First, we performed a model identifiability analysis using the best-fitting versions of each model class[30]. Specifically, we simulated data under each model and then fitted each model to these data. This analysis showed that the models were discriminable within the constraints of our experimental paradigm (**Figure S5**). Second, the models infer subjects' true confidence from their PDWs – raising the possibility that risk and/or loss aversion may have biased the model comparison results and/or distort the model-based estimates of confidence for subsequent fMRI analysis[31]. To rule out this possibility, we re-fitted the best-fitting version of each model class after adding a utility function, which transforms the expected value of the gamble options into subjective utilities and allows for individual differences in risk and/or loss aversion[31]. This analysis confirmed the ToM-model as providing the best account of the behavioural data and showed that the model-based estimates of confidence remained the same when inferred with or without a utility function (**Figure S6**).

Neural basis of a social confidence computation

We next turned to the fMRI data acquired during task behaviour to identify neural substrates that may contribute to a social confidence computation. The behavioural and modelling results supported a ToM account in which subjects compute confidence in another player's decisions by integrating distinct estimates of decision difficulty and others' ability. The first component, a representation of decision difficulty, is directly linked to the motion stimulus and in particular motion coherence: the higher the coherence, the easier the decision. The second component, a representation of others' ability, must be maintained separately from one's own ability and learned from feedback about the other player's choice accuracy. We would therefore expect brain systems traditionally involved in sensory and social processing to interact during a social confidence computation and in turn that the putative social areas carry learning signals relating to others' task performance.

To identify brain regions that were differentially activated by self- and other-related processing, we first estimated whole-brain contrasts between self- and other-trials during the decision and the gamble phase (**Figure 1**). During the stimulus phase, the self > other contrast identified classic perceptual decision-making regions, including extrastriate cortex, posterior parietal cortex and cingulate cortex, and – in line with only self-trials requiring active choice – motor regions (**Figure 4**). By contrast, the other > self contrast identified a classic ToM network, including middle temporal gyrus, TPJ and dmPFC (**Figure 4**). During the gamble phase, the self > other contrast did not identify differential activity, whereas the other > self contrast again identified the ToM network (results not shown for gamble phase – see **Table S2** and **Table S3** for all clusters surviving whole-brain correction for decision and gamble phases).

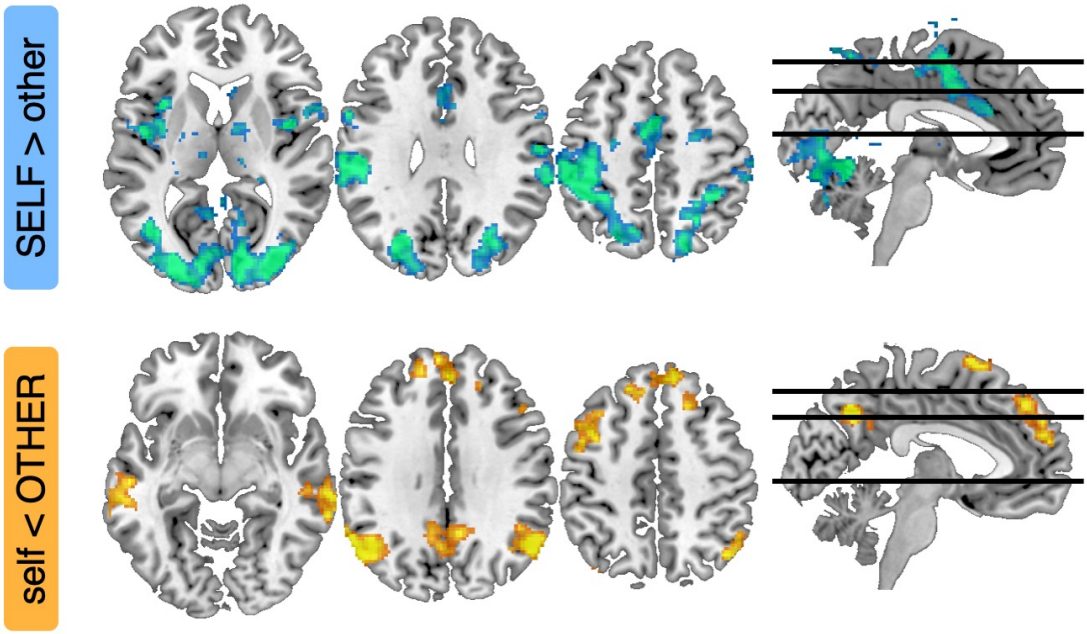

Figure 4. Whole-brain activations for self- and other-related processing during decision phase. Images display clusters that survive whole-brain correction ($p < .05$, FWE-corrected for multiple comparisons at a cluster-defining threshold of $p < .001$, uncorrected) for the contrast between self- and other-trials during the decision phase (cold: SELF > other; warm: self < OTHER). Images are shown at $p < .001$, uncorrected. All clusters surviving whole-brain correction during the stimulus and gamble phases are detailed in **Table S2** and **Table S3**.

We next focused on a subset of these brain regions, which were defined using independent ROI masks (**Figure 5**), to assess the neural evidence for a ToM account of a social confidence computation. First, we hypothesised that visual motion area MT+, identified in a separate localiser scan, and a human homologue of monkey LIP, encompassing posterior parts of the superior parietal lobule and the intraparietal sulcus[32], may support sensory representations of the motion stimulus which informs confidence in both one's own and others' choices. On this account, MT+ and LIP – implicated in sensory integration on the random dot motion task[8,33–36] – should encode motion coherence on both trial types. We surmise that the higher baseline activity in MT+ and LIP on self-trials (**Figure 4**) was driven by decision processes specific to self-trials such as the requirement for making an active choice. Second, we hypothesised that TPJ and dmPFC – implicated in social inference[21,22,37] – may combine a sensory representation of the motion stimulus with a distinct representation of others' ability as required by a social confidence computation. On this account, TPJ and dmPFC should track confidence in others' choices after controlling for the sensory representation of motion coherence. We tested these predictions using a complementary analysis approach: we (1) visualised temporally resolved neural encoding profiles by applying sliding-window regressions to up-sampled single-trial ROI activity time courses and then (2) estimated single-trial canonical hemodynamic response functions (c-HRF) for significance testing.

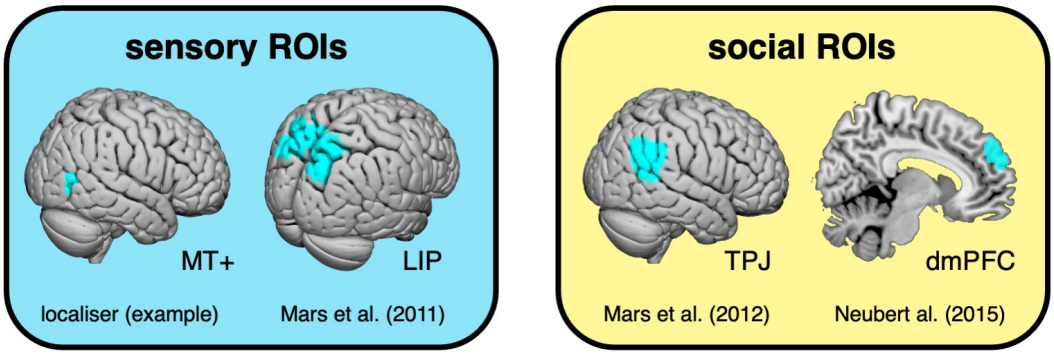

Figure 5. Anatomical masks for ROI analyses. The MT+ mask was specified for each subject using a localiser scan. The other masks were specified using published connectivity-based parcellation atlases.

A ToM account predicts that regions supporting the formation of a sensory representation of the motion stimulus should also contribute to a social confidence computation. In other words, we would expect our sensory ROIs to carry information about motion coherence on both self- and other-trials. To test this hypothesis, we quantified the neural impact of trial type (social; now coded as self = -.5 and other = .5), motion coherence and their interaction. As expected under a ToM account, MT+ and LIP tracked motion coherence on both trial types (**Figure 6**). In MT+, the response profile was the same on self- and other-trials: the higher the coherence, the higher the activity (c-HRF regression; social: $t(20) = -7.72$, $P < .001$; coherence, $t(20) = 5.42$, $P < .001$; interaction, $t(20) = -0.85$, $P = .408$; coherence on other-trials, $t(20) = 3.94$, $P < .001$; coherence on self-trials, $t(20) = 3.65$, $P = .002$). In contrast, the LIP response profile differed between trial types: the higher the coherence, the higher the activity on other-trials, but the lower the activity on self-trials (c-HRF regression; social: $t(20) = -7.86$, $P < .001$; coherence, $t(20) = 0.20$, $P = .844$; interaction, $t(20) = 3.07$, $P = .006$; coherence on other-trials, $t(20) = 2.94$, $P = .008$; coherence on self-trials, $t(20) = -2.14$, $P = .045$).

The LIP response pattern is consistent with previous observations of posterior parietal fMRI activity during active choice (characteristic of self-trials) versus passive viewing (characteristic of other-trials). When subjects are required to make an active choice, posterior parietal activity has been reported to decrease with motion coherence[35]. In contrast, during passive viewing, posterior parietal activity has been found to increase with motion coherence[38]. Such a flip is what we would expect if LIP neurons integrate sensory information into a belief state over stimulus space – equivalent to the sensory representation in our computational models. During active choice, the sensory integration process terminates earlier for high-coherence than low-coherence motion stimuli – a common result[23] – and bulk neural activity is therefore likely to be lower for high-coherence than low-coherence motion stimuli. In contrast, during passive viewing, such early termination of a choice process does not occur and motion coherence itself is likely to drive bulk neural activity. In support of this explanation, previous work has shown that, when motion coherence and reaction time are dissociated during active choice, posterior parietal activity increases with motion coherence, just as observed on other-trials[8].

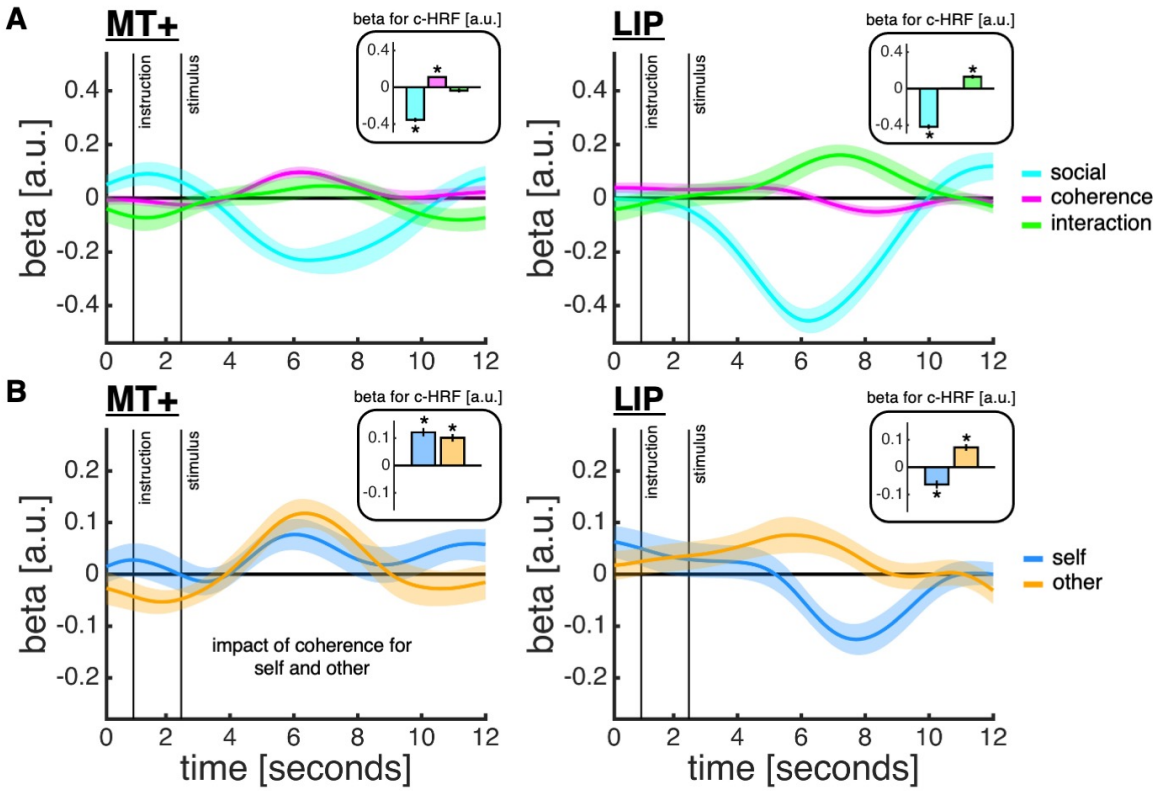

Figure 6. Encoding of motion coherence in MT+ and LIP. **(A)** The time courses are coefficients from a regression in which we predicted z-scored MT+ (*left*) and LIP (*right*) activity time courses using trial type (cyan; other = .5; self = -.5), z-scored motion coherence (pink) and their interaction (green). The insets show coefficients from an analysis of canonical HRFs (c-HRFs; asterisk indicates statistical significance, $P < .05$, one-sample t-test against zero). **(B)** Same approach as in panel A, but now quantifying the impact of motion coherence separately for each trial type. **(A-B)** Data are represented as group mean ± SEM.

Having found that MT+ and LIP track motion coherence on both trial types, we next turned to the social ROIs in order to assess their roles in a social confidence computation. The whole-brain analysis showed that the ToM network – including TPJ and dmPFC – was more active on other- than self-trials during both trial phases. However, a ToM account predicts that this network should not only discriminate between self and other but also differentially encode confidence for self and other. While previous work suggests that TPJ should be selective for a social confidence computation[39,40], the potential role of dmPFC is less clear: recent research indicates that dmPFC does not selectively encode social information but instead supports a separation between self- and other-related information when required by the task[41–43]. To test these hypotheses, we quantified the neural impact of trial type (social; self = -.5 and other = .5), our model-based confidence estimates (as computed under the ToM-model for both self and other) and the interaction between these terms. To control for low-level sensory effects, we orthogonalised the model-based confidence estimates with respect to motion coherence.

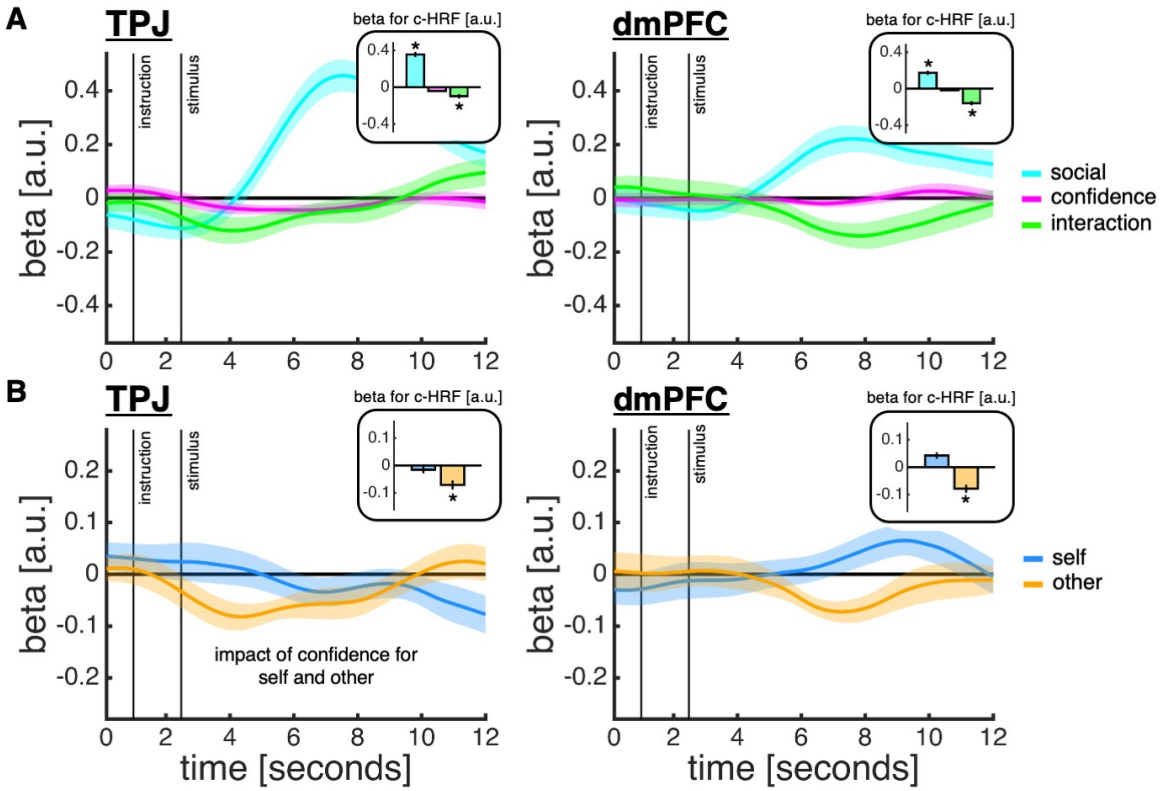

Figure 7. TPJ and dmPFC support a social confidence computation. **(A)** The time courses are coefficients from a regression in which we predicted z-scored TPJ (*left*) and dmPFC (*right*) activity time courses using trial type (cyan; other = .5; self = -.5), z-scored model-based confidence estimates as computed under the ToM-model (pink; first orthogonalised with respect to motion coherence) and their interaction (green). The insets show coefficients from an analysis of canonical HRFs (c-HRFs; asterisk indicates statistical significance, $P < .05$, one-sample t-test against zero). **(B)** Same approach as in panel A, but now quantifying the impact of model-based confidence estimates separately for each trial type. **(A-B)** Data are represented as group mean ± SEM.

This analysis showed that both TPJ and dmPFC encoded an interaction between trial type and the model-based confidence estimates (**Figure 7A**; c-HRF regression; TPJ social: $t(20) = 6.33$, $P < .001$; TPJ confidence, $t(20) = -1.79$, $P = .089$; TPJ interaction, $t(20) = -2.10$, $P = .049$; dmPFC social: $t(20) = 3.70$, $P = .001$; dmPFC confidence, $t(20) = -1.15$, $P = .264$; dmPFC interaction, $t(20) = -3.44$, $P = .003$). However, the nature of this interaction effect differed between TPJ and dmPFC (**Figure 7B**). In support of a selectivity for social inference, TPJ encoded the model-based confidence estimates on other-trials only, with higher activity when confidence was low (c-HRF regression; confidence on other-trials, $t(20) = -2.17$, $P = .043$; confidence on self-trials, $t(20) = -0.60$, $P = .555$). By contrast, dmPFC encoded model-based confidence estimates on both trial types: activity decreased with confidence on other-trials, but increased with confidence on self-trials, although the effect on self-trials did not reach statistical significance (c-HRF regression; confidence on other-trials, $t(20) = -2.67$, $P = .015$; confidence on self-trials, $t(20) = 1.79$, $P = .089$). The net effect of this inverse response profile in dmPFC is a larger neural distinction between other- and self-trials when confidence is low compared to when confidence is high.

Having established that TPJ and dmPFC track confidence estimates about others' choices, we next turned to the within- and across-trial dynamics of a social confidence computation. We reasoned that, if the ToM network receives a sensory representation of the motion stimulus

from classic perceptual decision-making regions and in turn furnishes it with a representation of another player's ability during a social confidence computation, then functional connectivity between our sensory and social ROIs should be higher on other- than self-trials. To test this hypothesis, we performed a psychophysiological interaction (PPI) analysis using MT+ or LIP as seed region and trial type as psychological variable (**Figure 8A**). In line with a ToM account, the PPI analysis revealed that connectivity between LIP and TPJ/dmPFC was higher on other- compared to self-trials (**Figure 8B-C**; one-sample *t*-test: TPJ, $t(20) = 2.50$, $p = .021$; dmPFC, $t(20) = 2.56$, $p = .019$). These relationships were not found for MT+ (one-sample *t*-test: TPJ, $t(20) = 0.51$, $p = .615$; dmPFC, $t(20) = -0.80$, $p = .432$).

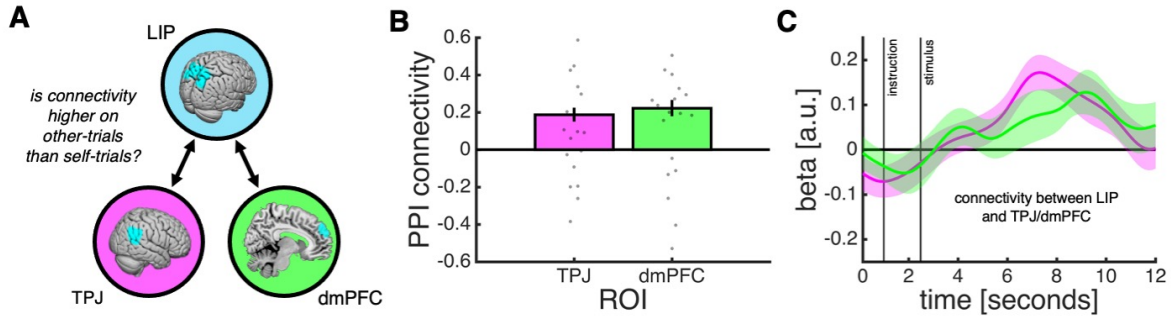

Figure 8. Functional coupling between sensory and social ROIs during social confidence computation. **(A)** Schematic of PPI analysis testing whether the correlation between LIP activity (seed region) and TPJ/dmPFC activity is higher on other- than self-trials. **(B)** Contrast estimates from PPI analysis (other > self) as implemented by the Generalised PPI toolbox. **(C)** Visualisation of activity time courses driving effects documented in panel B. The time courses are coefficients from a regression in which we predicted z-scored TPJ/dmPFC activity time courses using an interaction between z-scored LIP activity time courses and trial type (self = -.5; other = .5), while controlling for the main effect of each term. **(B-C)** Data are represented as group mean ± SEM.

Finally, having examined how sensory and social regions interact to compute confidence in others' choices, we analysed how the representation of another player's ability is updated with task experience. The behavioural and modelling results indicated that subjects revised their estimate of the other player's sensory noise based on the difference between the accuracy of the other player's choice and their confidence in this choice being correct (see schematic of learning mechanism in **Figure 3**). Consistent with a role in supporting this estimate, both TPJ and dmPFC tracked social prediction errors as computed under the ToM-model (**Figure 9**; (c-HRF regression; TPJ, $t(20) = -3.18$, $P = .005$; dmPFC, $t(20) = -3.12$, $P = .005$). We note that these response profiles remained after orthogonalising the social prediction error with respect to either the outcome term (accuracy) or the prediction term (confidence).

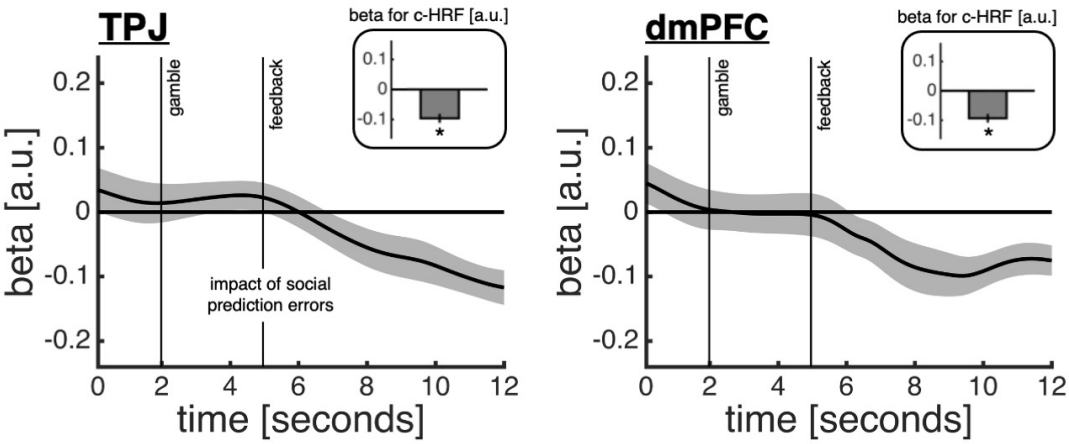

Figure 9. Encoding of social prediction errors in TPJ and dmPFC. The time courses are coefficients from a regression in which we predicted z-scored TPJ (*left*) and dmPFC (*right*) activity time courses using z-scored social prediction errors as computed under the best-fitting ToM-model. The insets show coefficients from an analysis of canonical HRFs (c-HRFs; asterisk indicates statistical significance, $P < .05$, one-sample *t*-test against zero). Data are represented as group mean ± SEM.

DISCUSSION

Computing confidence in others' decisions is critical for success on a range of competitive and cooperative tasks. While previous research has examined how we track others' ability, little is known about whether, or how, we combine this knowledge with an assessment of the decision problem at hand to form a confidence estimate. In many situations, simply knowing someone else's ability is insufficient for an accurate confidence computation as the difficulty of a decision may vary substantially. For example, an excellent medical doctor may not be able to diagnose a complex case, whereas someone who is generally bad at mathematics may still be able to solve a simple arithmetical problem. In situations like these, distinct estimates of others' ability and decision difficulty are needed in order to predict whether someone else will be successful at a given moment in time.

Here we addressed how people solve this problem using a perceptual decision task in which subjects placed bets on choices about stimuli of varying difficulty made by either themselves (self-trials) or one of three other players (other-trials) who differed in terms of their ability. By combining behavioural analysis and computational modelling, we found that subjects tracked the ability of each player and combined this knowledge with an estimate of decision difficulty in order to decide whether to gamble on the other players' choices being correct. Using fMRI, we found that a social confidence computation was supported by an interaction between brain systems involved in perceptual decision-making (MT+ and LIP) and social cognition (TPJ and dmPFC): MT+ and LIP tracked decision difficulty on both self- and other-trials; coupling between LIP and TPJ/dmPFC increased on other-trials; and TPJ/dmPFC tracked confidence in the other players' choices after controlling for decision difficulty and prediction error learning about the ability of the other players. Taken together, these results indicate that TPJ/dmPFC augment a representation of decision difficulty provided by LIP with a representation of the other players' ability. More broadly, our results reveal a neurocomputational interplay between self- and other-related processes during a social confidence computation.

Our winning model, the ToM-model, is a natural extension of Bayesian models of confidence from the individual to the social case[7]. In this model, decision difficulty and others' ability are

taken into account by integrating a belief state over stimulus space with a representation of the other players' expected accuracy for each stimulus. Our goal was to test whether a social confidence computation takes into account both decision difficulty and others' ability, rather than seeking to arbitrate between alternative algorithms that may support this computation. For example, people may solve the problem in a more model-free manner, by first calculating the confidence in the decision that they themselves would have made and then adjusting this estimate by a scalar value that reflects the performance of the other player (e.g., a combination of the S- and Q-models). Both this heuristic account and the ToM-model embody a hypothesis that a social confidence computation involves self- and other-related components, but they make subtly different predictions (e.g., the difference in confidence for a low-ability and high-ability player depends on the level of motion coherence under the ToM-model but is constant across levels of motion coherence under this more model-free alternative). Future studies are needed to disambiguate these related accounts (e.g., extensive psychophysical mapping of the relationship between motion coherence and confidence in others' choices).

Our confidence models were all grounded in Bayesian decision theory on self-trials, but they differed in their solution to a social confidence computation on other-trials. The Q-model solves this problem in a model-free manner, using running estimates of the other players' choice accuracy as a proxy for confidence in their choices. In contrast, the best-fitting model, the ToM-model, solves the problem in a Bayesian manner, by integrating distinct representations of decision difficulty and others' ability. We acknowledge that the computational basis of (non-social) confidence is still debated – with different studies supporting either Bayesian[7,44–47] or non-Bayesian[48–50] solutions. However, the grounding of our confidence models in Bayesian decision theory is unlikely to have biased our results. In keeping with standard signal detection theory, we modelled the sensory evidence as a single sample from a Gaussian distribution. Under these constraints, Bayesian and non-Bayesian confidence estimates are monotonically related and will therefore make similar predictions[6].

The question of how people infer what others think, feel or intend, and the relationship between inference about one's own and others' internal states, has attracted widespread attention in philosophy, psychology and neuroscience[26,51,52]. Our findings are in line with a view that self- and other-inference are distinct processes that involve distinct representations of oneself and others[24,53] – and go against a view that other-inference relies on self-projection[54,55] or that self-inference co-opts processes for inference about others[52]. At a behavioural level, we found that a social confidence computation involves distinct self- and other-related components – people did not simply simulate which choice they would have made and how confident they would have felt about this choice. At a neural level, we found that self- and other-related processes activated distinct brain systems. In support of this neural dissociation, a repetition suppression analysis – indexing the information-selectivity of neuronal populations[56] – revealed that activity in LIP (but not MT+) was highest on self-trials preceded by other-trials and lowest on other-trials preceded by other-trials, whereas activity in both TPJ and dmPFC was highest on other-trials preceded by self-trials and lowest on self-trials preceded by self-trials (**Figure S7**).

Our study characterises the neurocomputational basis of a common social problem – how we can make predictions about the choice accuracy of another system other than oneself. As is true for most social neuroscience studies, it is possible that one could devise an analogous task where the other system is not human but artificial (e.g., a motion detection algorithm) and find that this task recruits similar neurocomputational mechanisms. However, this result would

not alter the importance of these neurocomputational mechanisms for understanding human social behaviour. While our study was not designed to address the issue of social uniquenes[57], the involvement of a ToM network in the computation of confidence in other players' choices supports that our task taps into fundamental social processes. In addition, the recruitment of a ToM network supports our interpretation that subjects computed social confidence by estimating the other players' sensory noise – a process akin to inference about latent mental states which similarly can explain variability in others' behaviour – and not simply by averaging trial-by-trial feedback about the other players' choices.

More broadly, our study extends our understanding of the neural basis of social cognition. TPJ and dmPFC have previously been shown to support learning of summary statistics that reflect others' ability on a variety of tasks, including the reliability of advice on a bandit task[37], recency-weighted performance on a reaction-time task[21] or prediction accuracy on a stock-market task[22]. Our results show that TPJ and dmPFC also support the integration of such knowledge about others' ability with information about the current task state – a process that facilitates accurate predictions about others under conditions of varying task difficulty. There were, however, subtle differences between the activity profiles of TPJ and dmPFC. In line with a hypothesis that TPJ is selective for social inference, TPJ tracked the model-based confidence estimates only on other-trials. By contrast, dmPFC tracked these estimates on both trial types but in an inverse manner – resulting in a larger neural distinction between self and other when confidence was low. This result fits with evidence that dmPFC supports a separation between self- and other-related information in social contexts[41–43].

While TPJ involvement was restricted to other-evaluation in the current study, it remains plausible that TPJ may be involved in self-evaluation in other situations. On other-trials, but not self-trials, subjects had to learn a new mapping between the stimulus space and choice accuracy on each block to infer the other's ability. In addition, unlike on self-trials, subjects did not observe the other player's choice and had to consider that either choice could have been made. It remains to be seen whether self-evaluation co-opts putative social brain areas in situations where counterfactual thinking is required[58] or where one's own ability needs to be learned across time[52].

In summary, our study shows (1) that people form confidence in others' decisions by flexibly combining knowledge about others' ability with an assessment of the decision problem at hand and (2) that this process is supported by an interaction between brain systems traditionally involved in decision-making and social cognition. We highlight, however, that our study only addresses one of the two sides of a social confidence computation. In many cases, adaptive control of social behaviour requires not only that we can predict the success of others' choices – as examined here – but also that we can infer how confident others themselves feel about these choices. For example, in a strategic game like chess, our response to an opponent's mistake may depend on whether we think that they themselves realise that they have erred. How our confidence in others interacts – in a potentially reciprocal manner – with inference about others' confidence awaits further study.

AUTHOR CONTRIBUTIONS

S.M.F. and N.D.D. conceived the study. S.M.F. collected the data. D.B., R.M., N.D.D and S.M.F developed the models. D.B. analysed the data. D.B., R.M., N.D.D. and S.M.F. interpreted the results. D.B. wrote the manuscript. R.M., N.D.D. and S.M.F. provided critical comments. All authors approved the final manuscript.

ACKNOWLEDGEMENTS

We thank Sara Ershadmanesh and Matthew Rushworth for helpful comments on an earlier version of the manuscript. The Wellcome Centre for Human Neuroimaging is supported by core funding from the Wellcome Trust (203147/Z/16/Z). D.B. was supported by a Sir Henry Wellcome Postdoctoral Fellowship funded by the Wellcome Trust (213630/Z/18/Z). N.D.D. was supported by the John Templeton Foundation (61454). S.M.F. was supported by a Sir Henry Dale Fellowship jointly funded by the Wellcome Trust and the Royal Society (206648/Z/17/Z). For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript version arising from this submission.

SUBJECT DETAILS

Subjects performed a random dot motion task in a pre-scan and a scan session conducted on the same day. Twenty-two adults with no history of psychiatric or neurological disorders took part in the study. One subject was excluded due to poor performance in the pre-scan session, leaving twenty-one subjects for analysis (12 female, mean ± SD age = 22.6 ± 4.4 years). Subjects provided informed consent, and the study was approved by New York University's University Committee on Activities Involving Human Subjects. Subjects received a show-up fee of US$15, an additional US$25 for completing the MRI scan and were informed they could earn an additional performance-based bonus (in reality, all subjects received a US$9 bonus).

METHOD DETAILS

Experimental paradigm

Apparatus

The experiment was programmed in MATLAB 2014b using Psychtoolbox 3.0.12[59,60]. In the pre-scan session, stimuli were presented on an iMac desktop monitor viewed at a distance of approximately 45 cm. In the scan session, stimuli were presented via a projector at an approximate viewing distance of 58 cm.

Random dot kinematograms

Subjects viewed so-called random dot kinematograms (RDKs) contained in a circular aperture (7 degrees in diameter). Each RDK was made up of three independent sets of dots (each dot was 0.12 degrees in diameter) shown in consecutive frames. Each set of dots were shown for one frame (about 16 ms) and then replotted again three frames later (about 50 ms). Each time a set of dots was replotted, a subset of the dots, determined by the motion coherence, θ, was displaced in the direction of motion at a speed of 6 degrees s^{-1}, whereas the rest of the dots were displaced at random locations within the aperture. The motion direction, k, was to the left or to the right along the horizontal meridian. The use of three sets of dots means, for example, that dot positions in frame one are correlated with dot positions in frame four, seven and so forth. The dot density was fixed at 30 dots $degrees^{-2}$ s^{-1}.

Pre-scan session: familiarisation

Subjects were familiarised with the basic random dot motion task prior to the scan session. Each trial began with the presentation of a fixation cross at the centre of a circular aperture. After a fixed delay (.5 s), subjects viewed an RDK (.4 s). Once the RDK had terminated, subjects were asked to indicate whether the net direction of dot motion was to the left or to the right along the horizontal meridian. In particular, subjects were presented with a white box on the left and a white box on the right, corresponding to the two choice options. Once a choice had been made, the outline of the chosen option turned black (.5 s), and subjects received feedback about choice accuracy (low-pitched tone: error; high-pitched tone: correct). If subjects did not respond within 4 s, then the choice was scored as an error. The coherence, θ, was taken from the set, $\Theta \in \{.03, .06, .12, .24, .48, 1\}$, counterbalanced across trials. The motion direction, k, was sampled randomly from left or right. Subjects completed 240 trials.

Scan session: self-other

Subjects performed the same random dot motion task again, except that they were now also required to place post-decision wagers (PDWs) on the accuracy of either their own or another player's choices.

Each trial began with the presentation of a fixation cross at the centre of a circular aperture (1.5 s). Subjects then viewed a message instructing either a 'self-trial' or an 'other-trial' (1.5 s). On both trial types, after a brief delay (.5 s), subjects viewed an RDK (.4 s). After the RDK had terminated, subjects entered the choice screen – with the fixation cross and the circular aperture remaining visible. On self-trials, subjects actively decided about the net direction of dot motion. On other-trials, subjects passively viewed the screen while the other player made a decision about the motion stimulus – subjects were not informed whether these decisions were to the left or right. On both trial types, the circular aperture changed its colour to black once a choice had been made. The other player's choice reaction time was sampled uniformly from the range .5-1 s. On self-trials, if subjects did not respond within 3 s, then the trial was aborted – subjects had to wait for the duration that a full trial would have taken before proceeding to the next trial. The delay between the offset of the RDK and the offset of the choice screen was sampled uniformly from the range 3-6 s.

Subjects were then presented with a PDW screen (3 s) containing a safe and a risky option, indicated by a white box on each side of the screen (random assignment). The safe option, indicated by a blue circle, delivered a small but certain reward regardless of choice accuracy. The risky option delivered a larger reward, indicated by a green circle, if the self- or other-choice was correct, and a corresponding loss, indicated by a red circle, if the self- or other-choice was wrong. The safe value was drawn uniformly from the range 5-20 points, and the risky value was drawn uniformly from the range 25-50 points. Once subjects had made a PDW, they were presented with a feedback screen (1.5 s). The screen indicated the chosen PDW option (relevant box highlighted in black) and the accuracy of the choice (if correct, the green circle maintained its colour, whereas the red circle turned whited, and if incorrect, the red circle maintained its colour, whereas the green circle turned white). In this way, subjects were informed about the outcome on the current trial as well as the outcome that would have been obtained had they chosen the other PDW option. In addition, the screen also indicated the total earnings accumulated up to the current trial. Subjects started the scan session with 500 points and were told that their performance-based bonus depended on the total amount of points accumulated across the experiment. If subjects did not make a PDW within 3 s of the PDW onset, then they lost the points associated with the safe option.

Subjects performed 3 blocks (scan runs) consisting of 40 trials, with 20 self-trials and 20 other-trials randomly interleaved (120 trials in total). On each block, subjects were paired with a new player. All subjects faced the same three players. We created the three players using a pilot dataset of random dot motion decisions obtained from previous subjects who had visited the lab. We identified three pilot subjects who on average achieved about 60% (low ability), 75% (medium ability) and 90% (high ability) choice accuracy and then sub-selected 20 trials which covered a range of coherences (.001-.5), had an even split of leftwards and right motion and satisfied the above requirements for choice accuracy. On each block, the same set of motion stimuli were shown on self-trials and other-trials but in a randomised order. The order of the three players across blocks was randomised across subjects.

FMRI

Procedure

The scan session consisted of 5 scan runs. First, we acquired structural images. Second, subjects performed the self-other task over 3 runs (3 x 40 = 120 trials; 252 volumes per run). Lastly, we ran a localiser scan in which subjects viewed alternating displays (12 s) of static and coherently moving dots (2 x 10 = 20 displays; 215 volumes).

Acquisition

MRI data were acquired on a 3T Siemens Allegra scanner at New York University's Center for Brain Imaging. T1-weighted structural images were acquired using a 3D MPRAGE sequence: 1 x 1 x 1 mm resolution voxels; 176 sagittal slices. BOLD T2*-weighted functional images were acquired using a Siemens epi2d BOLD sequence: 3 x 3 x 3 mm resolution voxels; 42 transverse slices, 64 x 64 matrix; TR = 2.24s; TE = 30ms; slice tilt = -30 degrees T > C; slice thickness = 3 mm; interleaved slice acquisition). Local field maps were recorded for distortion correction of the acquired EPI data.

Pre-processing

MRI data were pre-processed using SPM12 (Wellcome Trust, London). The first 5 volumes of each functional run were discarded to allow for T1 equilibration. Functional images were slice-time corrected, re-aligned and un-warped using the collected field maps[61]. Structural T1-weighted images were co-registered to the mean functional image of each subject using the iterative mutual-information algorithm. Each subject's structural image was segmented into grey matter, white matter and cerebral spinal fluid using a nonlinear deformation field to map it onto a template tissue probability map[62]. These deformations were applied to both structural and functional images to create new images spatially normalised to the Montreal Neurological Institute (MNI) space and interpolated to 2 x 2 x 2 mm voxels. Normalized images were spatially smoothed using a Gaussian kernel with full-width half-maximum of 8 mm. Motion correction parameters estimated from the re-alignment procedure and their first temporal derivatives (12 regressors in total) were included as confounds in the first-level analysis for each subject.

Quantification and statistical analysis

Trial exclusion

We excluded trials when a decision was not made during the choice window or when a gamble was not made during the gamble window.

Behavioural analysis

We used multiple logistic regression to predict subjects' PDWs (**Figure S1**). We z-transformed all continuous predictors. We performed separate regressions for each subject. We tested group-level significance by comparing coefficients pooled across subjects to zero ($P < .05$, one-sample *t*-test).

Computational models of confidence

On each trial, a subject is asked to make a post-decision wager (PDW) about either their own choice or the choice of another player. On self-trials, the subject actively makes a choice about the motion stimulus. On other-trials, the subject is presented with the motion stimulus but not the other's choice. The subject has to decide between a risky and a safe option, by combining

the option values with an estimate of the probability that the choice is correct. We first consider models of PDWs on self-trials, before turning to other-trials. The models are all grounded in Bayesian decision theory[7,63].

Sensory sample

On each trial, a subject, s, receives a sensory sample, x, randomly sampled from a Gaussian distribution, $x \in N(k\theta_m, \sigma_s)$, where $m \in M = \{1,2,\dots,n\}$, $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$ indicates the motion coherence, $k \in K = \{-1,1\}$ indicates the motion direction (-1: left; 1: right), n is the number of levels of coherence and σ_s indicates the subject's level of sensory noise.

Self-trials

The subject represents the task as comprised of a set of states (k, m), corresponding to the set of possible motion stimuli defined by direction and coherence, $k\theta_m$. The subject computes a belief distribution over states given their sensory sample, x. The probability of each state conditional on x is given by:

$$P(k,m|x) = \frac{P(x|m,k)P(m,k)}{\sum_{k'}\sum_{i=1}^{n} P(x|m_i,k')\,P(m_i,k')} = \frac{P(x|m,k)}{\sum_{k'}\sum_{i=1}^{n} P(x|m_i,k')} \qquad \text{Eq. 1}$$

Note that we relied on the fact that all states are equally probable by design (with probability $1/2n$). The likelihood of x conditional on a state is given by:

$$P(x|k,m) = \varphi(x; k\theta_m, \sigma_s) \qquad \text{Eq. 2}$$

where φ is the normal probability density function evaluated at x:

$$\varphi(x;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2\sigma^2}} \qquad \text{Eq. 3}$$

The subject uses the belief distribution over states to compute the probability that the motion direction is leftwards ($k = -1$) and rightwards ($k = 1$):

$$P(k=-1|x) = \frac{\sum_{m\in M} P(x|m,k=-1)}{\sum_{k'}\sum_{m\in M} P(x|m,k')}$$

$$P(k=1|x) = \frac{\sum_{m\in M} P(x|m,k=1)}{\sum_{k'}\sum_{m\in M} P(x|m,k')} \qquad \textit{Eq. 4}$$

The subject makes their decision, r, by comparing the posteriors over motion direction:

$$r(x) = \begin{cases} -1 & \text{if } P(k=-1|x) > P(k=1|x) \\ 1 & \text{if } P(k=-1|x) < P(k=1|x) \end{cases} \qquad \textit{Eq. 5}$$

Finally, the probability that the choice is correct can be computed directly using the posteriors over motion direction relative to the chosen action:

$$P(\text{correct}|x) = P(k=r|x) = \begin{cases} P(k=-1|x) & \text{if r} = -1 \\ P(k=1|x) & \text{if r} = 1 \end{cases} \qquad \textit{Eq. 6}$$

Other-trials

Subjects were paired with three players of varying ability, $j \in J = \{1,2,3\}$. We consider below different models of how subjects compute confidence in player j's choices and how subjects learn about player j's ability.

Self-projection (S-model)

A simple way to compute confidence in another player's choice is to apply the same model to the other player as one applies to oneself. Under this model, the subject simulates the other's choices and confidence as if they were the other player.

The other player's posteriors over motion direction are calculated as:

$$P(k=-1|x,j) = P(k=-1|x,s)$$
$$P(k=1|x,j) = P(k=1|x,s) \qquad \text{Eq. 7}$$

The other player's choice is calculated as:

$$r(x) = \begin{cases} -1 & \text{if } P(k=-1|x,j) > P(k=1|x,j) \\ 1 & \text{if } P(k=-1|x,j) < P(k=1|x,j) \end{cases} \qquad \text{Eq. 8}$$

The subject's confidence in the other player's choice is calculated as:

$$P(\text{correct}|x,j) = P(k=r|x,j) = \begin{cases} P(k=-1|x,j) & \text{if r} = -1 \\ P(k=1|x,j) & \text{if r} = -1 \end{cases} \qquad \text{Eq. 9}$$

Note that the S-model does not learn about the other player's ability.

Simple value learner (Q-model)

An alternative strategy for estimating confidence in another player's choice is to keep a running estimate of the value (correctness) of their choices, $Q(j)$. This can be done compactly using a Rescorla-Wagner update rule:

$$Q(j)_{t+1} \coloneqq Q(j)_t + \alpha\delta_t \qquad \text{Eq. 10}$$

where t indicates the trial, α is a learning rate and δ_t is the difference between the observed outcome o_t (correct = 1; incorrect = 0) and the current estimate, $Q(j)_t$:

$$\delta_t = o_t - Q(j)_t \qquad \text{Eq. 11}$$

The subject's confidence in the other player's choice is then calculated as:

$$P_j(\text{correct}|j) = Q(j) \qquad \text{Eq. 12}$$

We set $Q(j)_{t=0} = .75$ for each player j. Note that, unlike the S-model, the Q-model is blind to the effect of coherence on the other player's chance of success.

Theory of mind (ToM-model)

A more sophisticated strategy is to maintain a running estimate of another player's sensory noise – that is, a model of their sensory system – and use this estimate to compute confidence in their choices. In the absence of shared error variance between the subject's and the other player's sensory sample (i.e., no noise correlation), the subject's prediction about the other player's success should take into account (1) the posterior probability that the world is in a given state, (k,m), given their own sensory sample, x, and (2) the other player's state-dependent performance (i.e., their expected accuracy in each state) given their estimated sensory noise, σ_j. The prediction about the other player's success is thus given by:

$$P_j(\text{correct}|x) = \sum_k \sum_{m \in M} P(m,k|x) P_j(\text{correct}|m,k) \qquad \text{Eq. 13}$$

where P_j denotes the probability that the other agent j makes a correct response. The first

term, the belief distribution over states, is calculated using σ_s, and the second term, the state-dependent performance of the other player, is calculated using σ_j:

$$P_j(\text{correct}|m,k) = \begin{cases} \phi\big(0; k\theta_m, \sigma_j\big) & \text{if } k\theta_m < 0 \\ 1 - \phi\big(0; k\theta_m, \sigma_j\big) & \text{if } k\theta_m > 0 \end{cases} \qquad \textit{Eq. 14}$$

where ϕ is the cumulative normal distribution:

$$\phi(y;\mu,\sigma) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^{y} e^{\frac{-(y-\mu)^2}{2\sigma^2}}\, dy \qquad \textit{Eq. 15}$$

The estimate of the other player's sensory noise is updated using an approximation to the Rescorla-Wagner rule:

$$\sigma_{j,t+1} = \sigma_{j,t} + \alpha\delta_t / D(x, \sigma_{j,t}) \qquad \textit{Eq. 16}$$

Where δ is a prediction error as in the Q*-model and $D(x,\sigma_j)$ is the following derivative:

$$D\big(x,\sigma_j\big) = \frac{dP_j(correct|x)}{d\sigma_j} \qquad \textit{Eq. 17}$$

In short, we want to define an update to the estimate of the other player's sensory noise, η_t, such that the prediction about the other player's success, $P_j(\text{correct}|x, \sigma_{j,t} + \eta_t\ )$, will be equal to $P_j\big(\text{correct}\big|x, \sigma_{j,t}\big) + \alpha\delta_t$:

$$P_j(\text{correct}|x, \sigma_{j,t} + \eta_t) = P_j\big(\text{correct}\big|x, \sigma_{j,t}\big) + \alpha\delta_t \qquad \textit{Eq. 18}$$

We can approximate the left-hand term around $\sigma_{j,t}$ linearly. So, we obtain:

$$P_j\big(\text{correct}\big|x, \sigma_{j,t}\big) + D\big(x, \sigma_{j,t}\big)\eta_t = P_j\big(\text{correct}\big|x, \sigma_{j,t}\big) + \alpha\delta_t \qquad \textit{Eq. 19}$$

Or

$$\sigma_{j,t+1} = \sigma_{j,t} + \alpha\delta_t / D\big(x, \sigma_{j,t}\big) \qquad \textit{Eq. 20}$$

The derivative of the prediction with respect to the parameter is calculated as:

$$\begin{aligned} D(x,\sigma_j) &= \sum\nolimits_{m\in M, k=\pm 1} \Big(\frac{dP(m,k|x)}{d\sigma_j} P_j[\text{correct}|m,k] + P(m,k|x)\frac{dP_j[\text{correct}|m,k]}{d\sigma_j}\Big) \\ &= \sum\nolimits_{m\in M, k=\pm 1} \Big(P(m,k|x)\frac{dP_j[\text{correct}|m,k]}{d\sigma_j}\Big) \end{aligned} \qquad \textit{Eq. 21}$$

Because $P(m,k|x)$ is not a function of σ_j, its derivative is 0.

The term $\frac{dP_j[\text{correct}|m,k]}{d\sigma_j}$ is calculated as:

$$\frac{dP_j[\text{correct}|m,k]}{d\sigma_j} = -\frac{\theta_m}{\sigma_j^2\sqrt{2\pi}}\, e^{\frac{-(\theta_m)^2}{2\sigma_j^2}} = -\frac{\theta_m}{\sigma_j}\, \varphi\big(0; \theta_m, \sigma_j\big) \qquad \textit{Eq. 22}$$

We thus have:

$$D(x,\sigma_j) = -\frac{1}{\sigma_j} \sum\nolimits_{m\in M, k=\pm 1} \big[P(m,k|x)\theta_m\, \varphi\big(0; \theta_m, \sigma_j\big)\,\big] \qquad \textit{Eq. 23}$$

The linear approximation to the Rescorla-Wagner rule may result in numerical overflow in a small number of cases (e.g., when σ_j is extremely low and x is extremely high). To avoid this issue, we implemented a proposal mechanism under which we iteratively adjusted the update term, $\alpha\delta_t/D(x, \sigma_{j,t})$, until $|P_j(\text{correct}|x, \sigma_{j,t} + \eta_t) - (P_j(\text{correct}|x, \sigma_{j,t}) + \alpha\delta_t)| < .001$.

Post-decision wagering

In all models, PDWs are based on the expected values of the risky and the safe options:

$$EV_{risky} = P(\text{correct}|x, a)V_{risky} - [1 - P(\text{correct}|x, a)]V_{risky}$$
$$EV_{safe} = V_{safe} \qquad \text{Eq. 24}$$

where a indicates the agent (subject or other) and V indicates the option value.

To account for individual differences in the propensity to gamble on a decision and sensitivity to changes in expected value, we modelled the probability of selecting the risky option using a softmax function:

$$P(risky) = \frac{1}{1 + e^{-[\beta_0 + \beta_1(EV_{risky} - EV_{safe})]}} \qquad \text{Eq. 25}$$

Model variations

We considered versions of the above models where the softmax bias, β_0, and/or temperature, β_1, could vary between self- and other-trials and where the representation of the stimulus space, M, was rescaled according to a power law (using a factor p). In brief, we rescaled each state as $m' = m^p$ and normalised the resulting state space, M', using the minimum and maximum values of the original stimulus space, M. The rescaling allowed for a higher density of motion stimuli at low ($p > 1$) or high ($p < 1$) stimulus intensities (where $p = 1$ preserves the original stimulus space). There were thus 24 models: 3 (S, Q, ToM) x 2 (same versus different bias, β_0, for self and other) × 2 (same versus different temperature, β_1, for self and other) × 2 (linearly-spaced versus rescaled stimulus space).

Model fitting and comparison

We used variational Bayesian inference as implemented in STAN[64] to fit the models to the behavioural data. We used a hierarchical fitting procedure where group-level parameters constrain subject-level parameters (see **Table S1** for full parameter space). We fitted each model using the following specifications: maximum number of iterations = 4,000; number of samples for Monte Carlo estimation of objective function = 200; number of iterations between evaluation of objective function = 100; convergence tolerance on the relative norm of the objective = .0001. To generate trial-by-trial predictions about each subjects' behaviour, and obtain estimates of the latent model variables, we drew 500 samples from the posterior distributions over fitted parameters (the 'generated quantities' block in STAN), calculated the trial-by-trial values for each set of samples and then averaged across these 500 iterations. To assess the ability of each model to account for a subject's data, we computed the log-likelihood of their trial-by-trial PDWs under each draw of parameters and applied leave-one-out cross-validation using Pareto-smoothed importance sampling (PSIS-LOO) and the widely applicable information criterion (WAIC) – both methods for estimating pointwise out-of-sample prediction accuracy under a Bayesian model[65]. We performed these steps 4 times for each model using

different random seeds and averaged all outputs across the 4 runs. The results of the model comparison are shown in **Figure S2**.

Model identifiability

As a quality control of our modelling approach, we conducted a model identifiability analysis[30]. We selected the best-fitting version of each model class (**Figure S2**), simulated 200 datasets for each of the three models – with each dataset mirroring our study in terms of subjects, task variables and trials – and then fitted the three models to the simulated datasets. Based on the model comparison results for all 200 iterations, we computed a confusion and an inversion matrix. A confusion matrix shows the probability that model Y provided the best fit to data generated by model X. By contrast, an inversion matrix shows the probability that model X generated the data given that model Y provided the best fit – in other words, if model Y wins the model comparison, which model is most likely to be the true model? The inversion matrix is closely related to the goal of computational modelling in cognitive neuroscience – where we want to know which model among a set of candidate models best describes the underlying cognitive process. As shown in **Figure S5**, the model identifiability analysis shows that all three models were discriminable within the constraints of our experimental paradigm.

Risk and/or loss aversion

Our models inferred subjects' confidence based on PDWs. Previous research has shown that PDWs are subject to risk and loss aversion[31] – raising the possibility that these factors may have biased the model comparison and/or distorted the model-based estimates of confidence for fMRI analysis. To rule out this possibility, we selected the best-fitting version of each model class (**Figure S2**) and re-fitted these models to the behavioural data after including a utility function. This function transforms the expected values of the gamble options into subjective utilities and allows for individual variation in risk and/or loss aversion. Formally, we computed the expected values of the risky and the safe options as follows:

$$EV_{risky} = P(\text{correct}|x,a)V_{risky}^{R} - [1 - P(\text{correct}|x,a)]SV_{risky}^{T}$$
$$EV_{safe} = V_{risky}^{R} \qquad \textit{Eq. 26}$$

where R, T and S parameterise the utility function – with S controlling loss aversion and R and T controlling risk-seeking/risk-aversion in the gain and loss domains, respectively. In keeping with previous work[31], we restricted R and T to the range 0-1 (see **Table S1** for all parameters). Importantly, the ToM-model continued to provide the best fit to the behavioural data, and the model-based confidence estimates were highly correlated (Pearson's $r > 0.9$ for all subjects) between the original ToM-model and the ToM-model with a utility function (**Figure S6**).

For completeness, we visualised the fitted group-level utility function under the ToM-model. In keeping with classic results in behavioural economics[66] and earlier results on the drivers of PDWs[31], this visualisation showed that losses loomed larger than gains (**Figure S6**).

Although not shown here, we highlight that all model-based fMRI results remained the same when performed using the ToM-model with a utility function.

FMRI

Whole-brain general linear models

The whole-brain analysis shown in **Figure 4** was based on a single event-related general linear model (GLM1). We separated trials into self-trials and other-trials and specified separate 'condition' regressors for the decision phase (i.e., boxcar from RDK onset to 3 s after decision onset) and the gamble phase (i.e., boxcar from gamble onset until feedback offset) for each trial type (yielding 4 condition regressors in total). We included motion parameters as 'nuisance' regressors. Regressors were convolved with a canonical hemodynamic response function. Regressors were modelled separately for each scan run and constants were included to account for between-run differences in mean activation and scanner drifts. A high-pass filter (128s cut-off) was applied to remove low-frequency drifts. Group-level significance was assessed by applying one-sample *t*-tests against zero to the first-level contrast images. We report clusters significant at $p < .05$, FWE-corrected for multiple comparisons, with a cluster extent of 10 voxels or more and a cluster-defining threshold of $p < .001$, uncorrected. Numerical simulations and tests of empirical data collected under the null hypothesis show that this combination of cluster-defining threshold and random field theory produces appropriate control of false positives[67].

Regions of interest

ROI masks for MT+ were created using the localiser scan: we created a group mask using the second-level contrast between dynamic and static motion, and then, for each subject, created a MT+ mask (8-mm sphere) around their peak activity inside the group mask. ROI masks for LIP, TPJ and dmPFC were created using published connectivity-based parcellation atlases: LIP was defined as the union of areas SPLD and SPLE in the atlas developed by Mars et al.[32]; TPJ was defined as area TPJp in the atlas developed by Mars et al.[68]; and dmPFC was defined as area 9 in the atlas developed by Neubert et al.[69]. All ROI masks were bilateral.

Psychophysiological interaction analysis

To assess changes in connectivity between visual and social ROIs as a function of trial type, we carried out a psychophysiological interaction (PPI) analysis. PPI is a measure of context-dependent connectivity, explaining the regional activity of other brain regions in terms of the interaction between responses in a seed region and a cognitive or sensory process. We used GLM1 as basis for PPI analysis. We carried out the PPI analysis using the Generalised PPI toolbox for SPM (gPPI; http://www.nitrc.org/projects/gppi/). gPPI creates a new GLM in which the deconvolved activity of the seed region is assigned to separate regressors dependent on the status of the original psychological variable (self versus other during the decision phase), and re-convolved with the hemodynamic response function. We extracted average contrast estimates within an ROI for each subject and then used these estimates for group-level testing.

Single-trial activity time courses

We transformed each ROI mask from MNI to native space and extracted preprocessed BOLD time courses as the average of voxels within the mask. For each scan run, we regressed out variation due to head motion as specified above, applied a high-pass filter (128 s cut-off) to remove low-frequency drifts, and oversampled the BOLD time course by a factor of ~23 (time resolution of .144 s; spline interpolation). For each trial, we extracted activity estimates in a 14 s window around an event of interest (i.e., RDK or feedback onset).

Canonical hemodynamic response functions

In addition to the whole-brain GLM analyses described above, we also estimated whole-brain GLMs which included a separate regressor for each trial[70]. For the stimulus-related analyses (**Figure 6**, **Figure 7** and **Figure S7**), the regressors were boxcars spanning RDK presentation. For the feedback-related analysis (**Figure 9**), the regressors were boxcars spanning feedback delivery. Each of these regressors was convolved with a canonical hemodynamic response function. We included motion parameters as 'nuisance' regressors. Regressors were modelled separately for each scan run and constants were included to account for between-run differences in mean activation and scanner drifts. A high-pass filter (128 s cut-off) was applied to remove low-frequency drifts. One consideration when obtaining single-trial activity estimates as a beta time series is that a beta for a given trial can be affected by acquisition artefacts that occur together with that trial (e.g., scanner pulse artefacts). Therefore, for each subject, we computed the grand-mean beta estimate across both voxels and trials and excluded any trial whose mean beta estimate across voxels was 3 SDs below or above this grand mean[70]. Finally, we used the ROI masks to extract single-trial ROI activity estimates under the canonical hemodynamic response function.

DATA AND SOFTWARE AVAILABILITY

Anonymized behavioural data and code supporting our main analyses are available on GitHub (https://github.com/danbang/article-self-other). Unthresholded group-level statistical maps are available on NeuroVault (https://neurovault.org/collections/9553/).

REFERENCES

1. Fleming, S. M., Van Der Putten, E. J. & Daw, N. D. Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience* **21**, 617–624 (2018).
2. Van Den Berg, R. *et al.* A common mechanism underlies changes of mind about decisions and confidence. *eLife* **5**, 1–21 (2016).
3. Ma, W. J. & Jazayeri, M. Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience* **37**, 205–220 (2014).
4. Pouget, A., Drugowitsch, J. & Kepecs, A. Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience* **19**, 366–374 (2016).
5. Meyniel, F., Sigman, M. & Mainen, Z. F. Confidence as Bayesian probability: from neural origins to behavior. *Neuron* **88**, 78–92 (2015).
6. Aitchison, L., Bang, D., Bahrami, B. & Latham, P. E. Doubly Bayesian analysis of confidence in perceptual decision-making. *PLoS Computational Biology* **11**, (2015).
7. Fleming, S. M. & Daw, N. D. Self-evaluation of decision-making: A general bayesian framework for metacognitive computation. *Psychological Review* **124**, 91–114 (2017).
8. Bang, D. & Fleming, S. M. Distinct encoding of decision confidence in human medial prefrontal cortex. *Proceedings of the National Academy of Sciences* **115**, 6082–6087 (2018).
9. Gherman, S. & Philiastides, M. G. Human VMPFC encodes early signatures of confidence in perceptual decisions. *bioRxiv* 224337 (2017) doi:10.1101/224337.
10. Pereira, M. *et al.* Disentangling the origins of confidence in speeded perceptual judgments through multimodal imaging. *Proceedings of the National Academy of Sciences of the United States of America* **117**, 8382–8390 (2020).
11. Fleming, S. M., Huijgen, J. & Dolan, R. J. Prefrontal contributions to metacognition in perceptual decision making. *Journal of Neuroscience* **32**, 6117–6125 (2012).
12. Bang, D., Ershadmanesh, S., Nili, H. & Fleming, S. M. Private–public mappings in human prefrontal cortex. *eLife* **9**, e56477 (2020).
13. Rouault, M. & Fleming, S. M. Formation of global self-beliefs in the human brain. *Pnas* **in press**, 1–9 (2020).
14. Johnson, D. D. P. & Fowler, J. H. The evolution of overconfidence. *Nature* **477**, 317–320 (2011).
15. Mahmoodi, A. *et al.* Equality bias impairs collective decision-making across cultures. *Proceedings of the National Academy of Sciences* **112**, 201421692 (2015).
16. Yaniv, I. Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes* **93**, 1–13 (2004).
17. Bang, D. *et al.* Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and Cognition* **26**, 13–23 (2014).
18. Hertz, U. *et al.* Neural computations underpinning the strategic management of influence in advice giving. *Nature Communications* **8**, 1–12 (2017).
19. Mahmoodi, A., Bang, D., Ahmadabadi, M. N. & Bahrami, B. Learning to make collective decisions: the impact of confidence escalation. *PLoS ONE* **8**, (2013).
20. Bang, D. *et al.* Confidence matching in group decision-making. *Nature Human Behaviour* **1**, (2017).
21. Wittmann, M. K. *et al.* Self-other mergence in the frontal cortex during cooperation and competition. *Neuron* **91**, 482–493 (2016).
22. Boorman, E. D., O'Doherty, J. P., Adolphs, R. & Rangel, A. The behavioral and neural mechanisms underlying the tracking of expertise. *Neuron* **80**, 1558–1571 (2013).
23. Shadlen, M. N. & Kiani, R. Decision making as a window on cognition. *Neuron* **80**, 791–806 (2013).
24. Saxe, R. Uniquely human social cognition. *Current Opinion in Neurobiology* **16**, 235–239 (2006).

25. Amodio, D. M. & Frith, C. D. Meeting of minds: the medial frontal cortex and social cognition. *Nature Reviews Neuroscience* **7**, 268–277 (2006).
26. Frith, C. D. & Frith, U. Mechanisms of social cognition. *Annual Review of Psychology* **63**, 287–313 (2012).
27. Wittmann, M. K. *et al.* Self-other mergence in the frontal cortex during cooperation and competition. *Neuron* **91**, 482–493 (2016).
28. Behrens, T. E. J., Hunt, L. T., Woolrich, M. W. & Rushworth, M. F. S. Associative learning of social value. *Nature* **456**, 245–249 (2008).
29. Newsome, W. T., Britten, K. H. & Movshon, J. A. Neuronal correlates of a perceptual decision. *Nature* **341**, 52–54 (1989).
30. Wilson, R. C. & Collins, A. G. E. Ten simple rules for the computational modeling of behavioral data. *eLife* **8**, 1–33 (2019).
31. Fleming, S. M. & Dolan, R. J. Effects of loss aversion on post-decision wagering: Implications for measures of awareness. *Consciousness and Cognition* **19**, 352–363 (2010).
32. Mars, R. B. *et al.* Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and Macaque resting-state functional connectivity. *Journal of Neuroscience* **31**, 4087–4100 (2011).
33. Britten, K. H., Shadlen, M. N., Newsome, W. T. & Movshon, J. A. Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience* **10**, 1157–1169 (1993).
34. Britten, K. H., Newsome, W. T., Shadlen, M. N., Celebrini, S. & Movshon, J. A. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience* **13**, 87–100 (1996).
35. Hebart, M. N., Donner, T. H. & Haynes, J. D. Human visual and parietal cortex encode visual choices independent of motor plans. *NeuroImage* **63**, 1393–1403 (2012).
36. Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* **86**, 1916–1936 (2001).
37. Behrens, T. E. J., Hunt, L. T., Woolrich, M. W. & Rushworth, M. F. S. Associative learning of social value. *Nature* **456**, 245–9 (2008).
38. Braddick, O. J. *et al.* Brain areas sensitive to coherent visual motion. *Perception* **30**, 61–72 (2001).
39. Saxe, R. Uniquely human social cognition. *Current Opinion in Neurobiology* **16**, 235–239 (2006).
40. Frith, C. D. & Frith, U. Mechanisms of social cognition. *Annual Review of Psychology* **63**, 287–313 (2012).
41. Wittmann, M. K. *et al.* Self-other mergence in the frontal cortex during cooperation and competition. *Neuron* **91**, 482–493 (2016).
42. Wittmann, M. K. *et al.* Causal manipulation of self-other mergence in the dorsomedial prefrontal cortex. *Neuron* 1–9 (2021) doi:10.1016/j.neuron.2021.05.027.
43. Nicolle, A. *et al.* An agent independent axis for executed and modeled choice in medial prefrontal cortex. *Neuron* **75**, 1114–21 (2012).
44. Bang, D. & Fleming, S. M. Distinct encoding of decision confidence in human medial prefrontal cortex. *Proceedings of the National Academy of Sciences* **115**, 6082–6087 (2018).
45. Kiani, R. & Shadlen, M. N. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* **324**, 759–764 (2009).
46. Khalvati, K., Kiani, R. & Rao, R. P. N. explains perceptual con fi dence and its deviations. *Nature Communications* 1–16 doi:10.1038/s41467-021-25419-4.
47. Fleming, S. M., Van Der Putten, E. J. & Daw, N. D. Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience* **21**, 617–624 (2018).

48. Samaha, J., Barrett, J. J., Sheldon, A. D., LaRocque, J. J. & Postle, B. R. Dissociating Perceptual Confidence from Discrimination Accuracy Reveals No Influence of Metacognitive Awareness on Working Memory. *Frontiers in Psychology* **7**, 851 (2016).
49. Zylberberg, A., Barttfeld, P. & Sigman, M. The construction of confidence in a perceptual decision. *Frontiers in Integrative Neuroscience* **6**, 1–10 (2012).
50. Adler, W. T. & Ma, W. J. Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLoS Comput Biol* **14**, e1006572 (2018).
51. Adolphs, R. The social brain: neural basis of social knowledge. *Annual review of psychology* **60**, 693–716 (2009).
52. Carruthers, P. How we know our own minds: The relationship between mindreading and metacognition. *Behavioral and Brain Sciences* **32**, 121–182 (2009).
53. Saxe, R. Against simulation: The argument from error. *Trends in Cognitive Sciences* **9**, 174–179 (2005).
54. Mitchell, J. P. Inferences about mental states. *Philosophical Transactions of the Royal Society B: Biological Sciences* **364**, 1309–1316 (2009).
55. Buckner, R. L. & Carroll, D. C. Self-projection and the brain. *Trends in Cognitive Sciences* **11**, 49–57 (2007).
56. Barron, H. C., Garvert, M. M. & Behrens, T. E. J. Repetition suppression: a means to index neural representations using BOLD? *Philosophical Transactions of the Royal Society B: Biological Sciences* **371**, 20150355 (2016).
57. Lockwood, P. L., Apps, M. A. J. & Chang, S. W. C. Is There a 'Social' Brain? Implementations and Algorithms. *Trends in Cognitive Sciences* **xx**, 1–12 (2020).
58. Mazor, M., Friston, K. J. & Fleming, S. M. Distinct neural contributions to metacognition for detecting, but not discriminating visual stimuli. *eLife* **9**, 1–34 (2020).
59. Brainard, D. H. The Psychophysics Toolbox. *Spatial Vision* **10**, 433–436 (1997).
60. Pelli, D. G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision* vol. 10 437–442 (1997).
61. Andersson, J. L. R., Hutton, C., Ashburner, J., Turner, R. & Friston, K. Modeling geometric deformations in EPI time series. *NeuroImage* **13**, 903–919 (2001).
62. Ashburner, J. & Friston, K. J. Unified segmentation. *NeuroImage* **26**, 839–851 (2005).
63. Green, D. & Swets, J. *Signal Detection Theory and Psychophysics*. (Wiley, 1966).
64. Kucukelbir, A., Ranganath, R., Gelman, A. & Blei, D. M. Automatic variational inference in Stan. in *Advances in Neural Information Processing Systems 28* 568–576 (2015).
65. Vehtari, A., Gelman, A. & Gabry, J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing* **27**, 1413–1432 (2017).
66. Kahneman, D. & Tversky, A. Prospect theory: an analysis of decision under risk. *Econometrica* **47**, 263–291 (1979).
67. Eklund, A., Nichols, T. E. & Knutsson, H. Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proceedings of the National Academy of Sciences* **113**, 7900–7905 (2016).
68. Mars, R. B. *et al.* Connectivity-based subdivisions of the human right 'temporoparietal junction area': evidence for different areas participating in different cortical networks. *Cerebral Cortex* **22**, 1894–1903 (2012).
69. Neubert, F.-X., Mars, R. B., Sallet, J. & Rushworth, M. F. S. Connectivity reveals relationship of brain areas for reward-guided learning and decision making in human and monkey frontal cortex. *Proceedings of the National Academy of Sciences* **112**, E2695–E2704 (2015).
70. Atlas, L. Y., Bolger, N., Lindquist, M. A. & Wager, T. D. Brain mediators of predictive cue effects on perceived pain. *Journal of Neuroscience* **30**, 12964–12977 (2010).

Supplementary Materials

Neurocomputational mechanisms of confidence in self and others

Dan Bang*, Rani Moran, Nathaniel Daw and Stephen M. Fleming*

*Correspondence: danbang.db@gmail.com and stephen.fleming@ucl.ac.uk

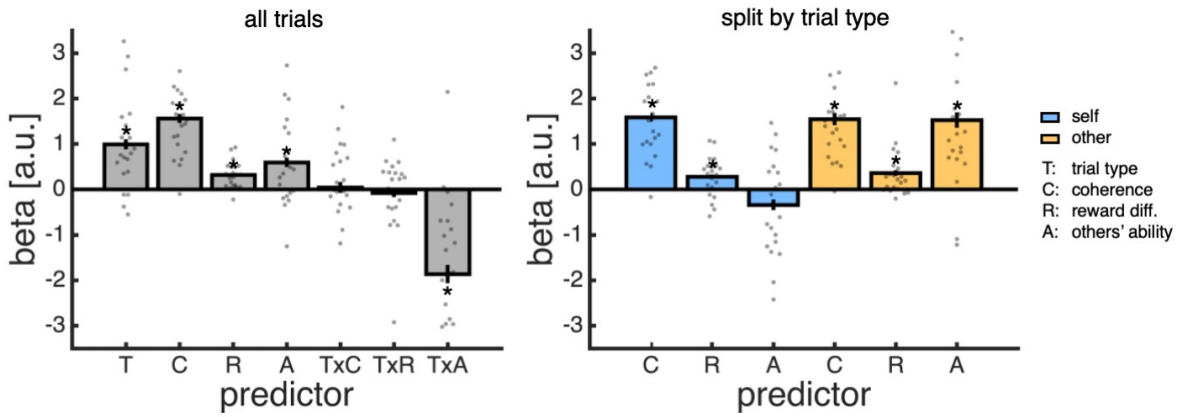

Figure S1. Multiple logistic regression analysis of post-decision wagers. **(A)** We first predicted gamble choices (i.e., selecting the risky option over the safe option) on all trials using trial type (T; other = -.5; self = .5), coherence (C; z-score), reward difference (R; z-score), other's ability (A; low = -.5, medium = 0, high = .5) and the interactions between trial type and the three other terms (TxC, TxR and TxA). **(B)** We then predicted gamble choices separately for self- (yellow) and other-trials (blue) using coherence, reward difference and other's ability. **(A-B)** We performed separate regressions for each subject. We tested group-level significance (asterisk) by comparing coefficients pooled across subjects to zero (one-sample *t*-test). Data are represented as group mean ± SEM, with individual subjects overlaid as dots.

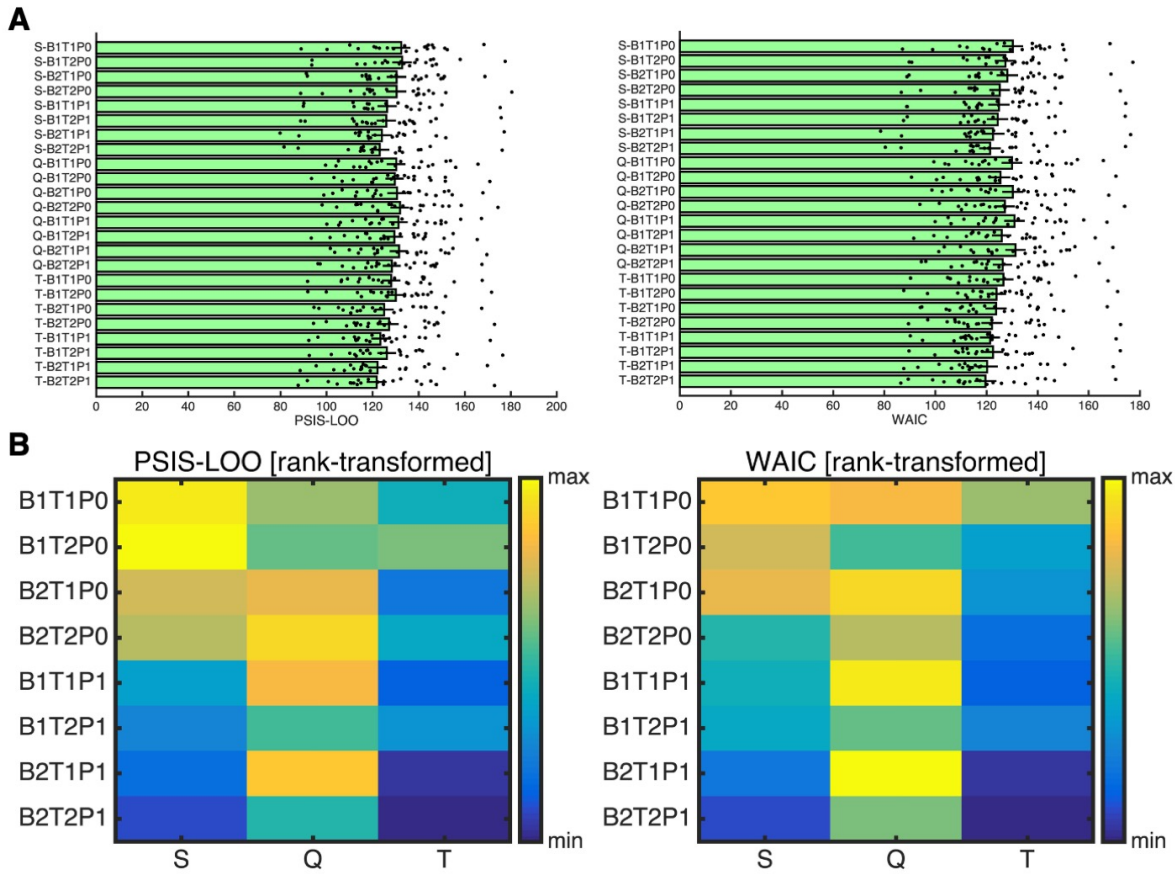

Figure S2. Model comparison based on leave-one-out cross-validation using (*left*) Pareto-smoothed importance sampling (PSIS-LOO) or (*right*) the widely applicable information criterion (WAIC). Here smaller values indicate better fit. **(A)** Metrics for each model represented as group mean ± SEM, with individual subjects overlaid as dots. **(B)** Rank-transformed average model metrics. **(A-B)** Model classes are abbreviated as follows: S, S-model; Q, Q-model; T, ToM-model. Model versions are abbreviated as follows: B, softmax bias (1: same bias for self and other; 2: different bias for self and other); T, softmax temperature (1: same temperature for self and other; 2: different temperature for self and other); P, rescaling of stimulus space (0: linear spacing; 1: re-scaled according to fitted factor). See **Table S1** for overview of model parameters.

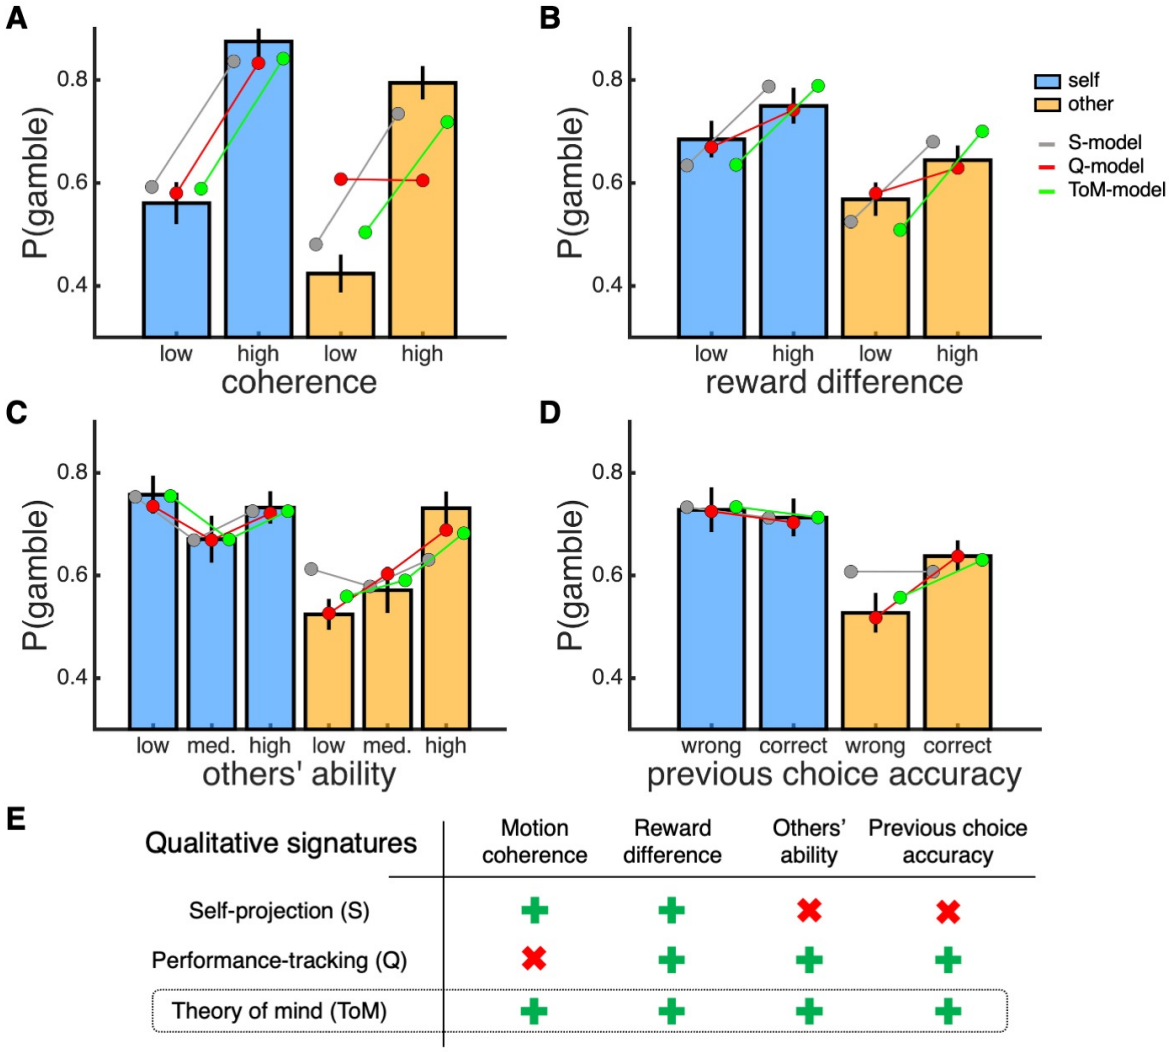

Figure S3. Model predictions. In each panel, data are split into self-trials (blue) and other-trials (yellow). **(A**) Probability gamble as a function of coherence (median split). **(B)** Probability gamble as a function of the reward difference between the risky and the safe option (median split). **(C)** Probability gamble as a function of the ability of the other player. **(D)** Probability gamble as a function of choice accuracy on the previous trial of the same type. **(E)** While all models predict the observed pattern of results on self-trials, only the ToM-model predicts the observed pattern of results on other-trials. **(A-D)** Bar charts are empirical data. Dots show data simulated under the best-fitting version of each model class (S-B2T2P1, Q-B2T2P1 and ToM-B2T2P1). Models are denoted by colour and displayed in the following order: S-model, Q-model and ToM-model. Empirical data are represented as group mean ± 95% CI and model predictions are represented as group mean.

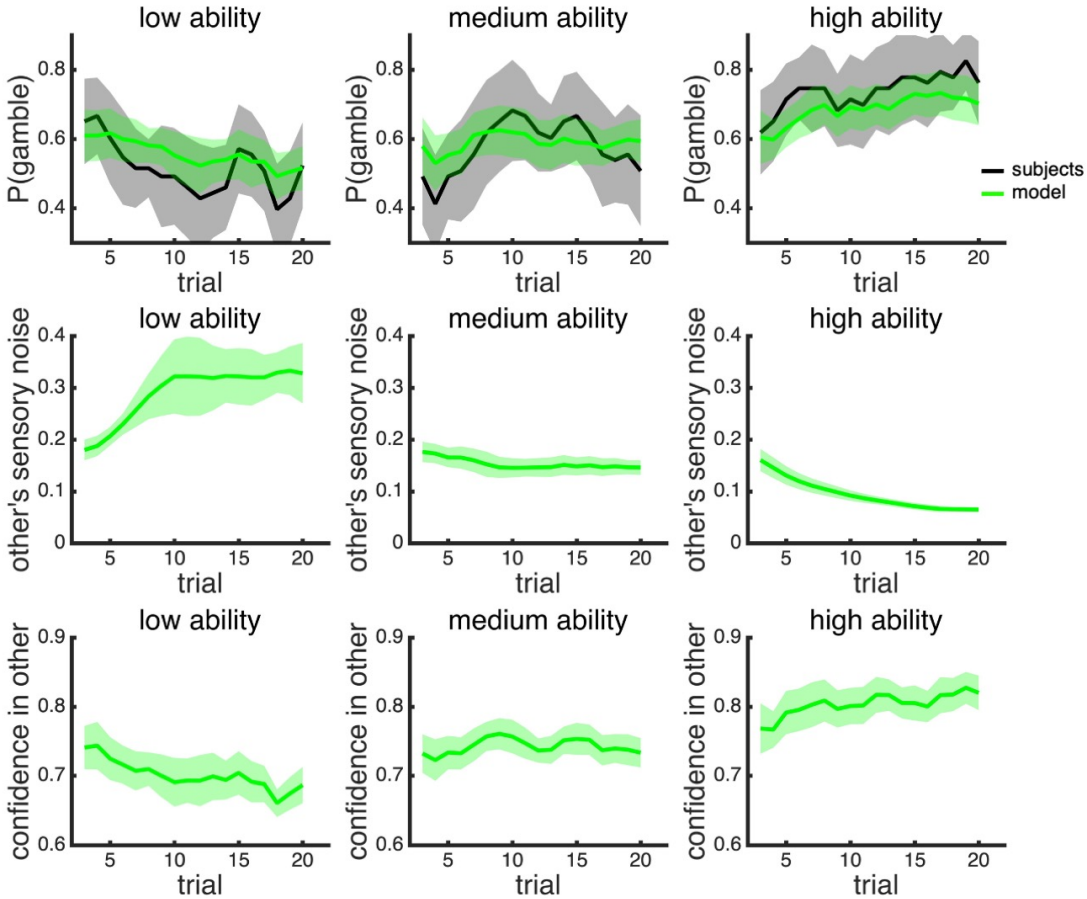

Figure S4. Temporal dynamics of model variables. **(A)** Evolution of probability gamble within each block as observed in the behavioural data and as predicted under the ToM-model. **(B)** Evolution of the estimate of the other players' sensory noise under the ToM-model. **(C)** Evolution of confidence in the other players' choice under the ToM-model. For each variable, we applied a running average where the value on trial *t* is the average over trials *t*-2 to *t* – the data are therefore shown from trial 3 onwards. Empirical and simulated data are represented as group mean ± 95% CI.

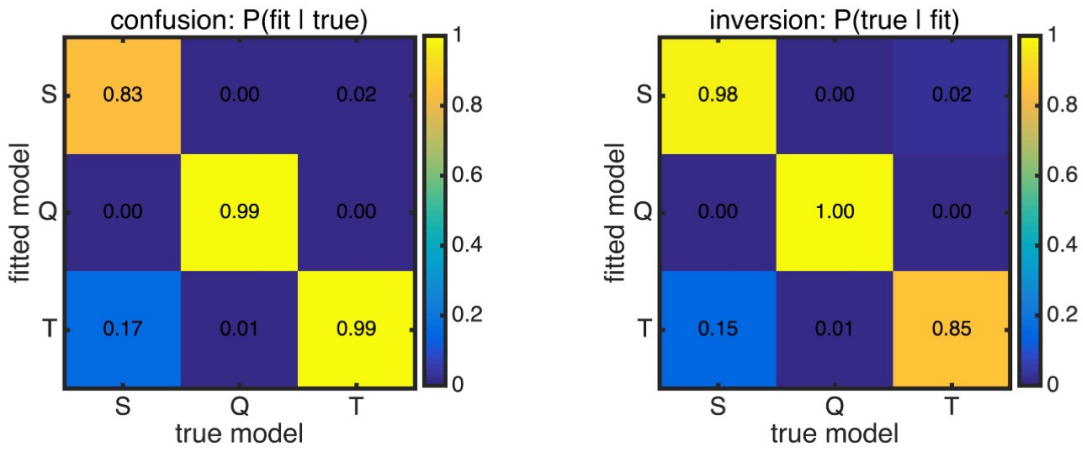

Figure S5. Model recoverability analysis. We selected the best-fitting version of each model class (S-B2T2P1, Q-B2T2P1 and ToM-B2T2P1), simulated 200 datasets for each model – with each dataset mirroring our study in terms of subjects, task variables and trials – and then fitted the models to the simulated datasets. Based on the model comparison results (WAIC) for all 200 iterations, we computed a confusion and an inversion matrix. A confusion matrix shows the probability that model Y provided the best fit to data generated by model X. By contrast, an inversion matrix shows the probability that model X generated the data given that model Y provided the best fit. We calculated the inversion matrix directly from the confusion matrix and assumed that the three models were a priori equally likely. Model classes are abbreviated as follows: S, S-model; Q, Q-model; T, ToM-model.

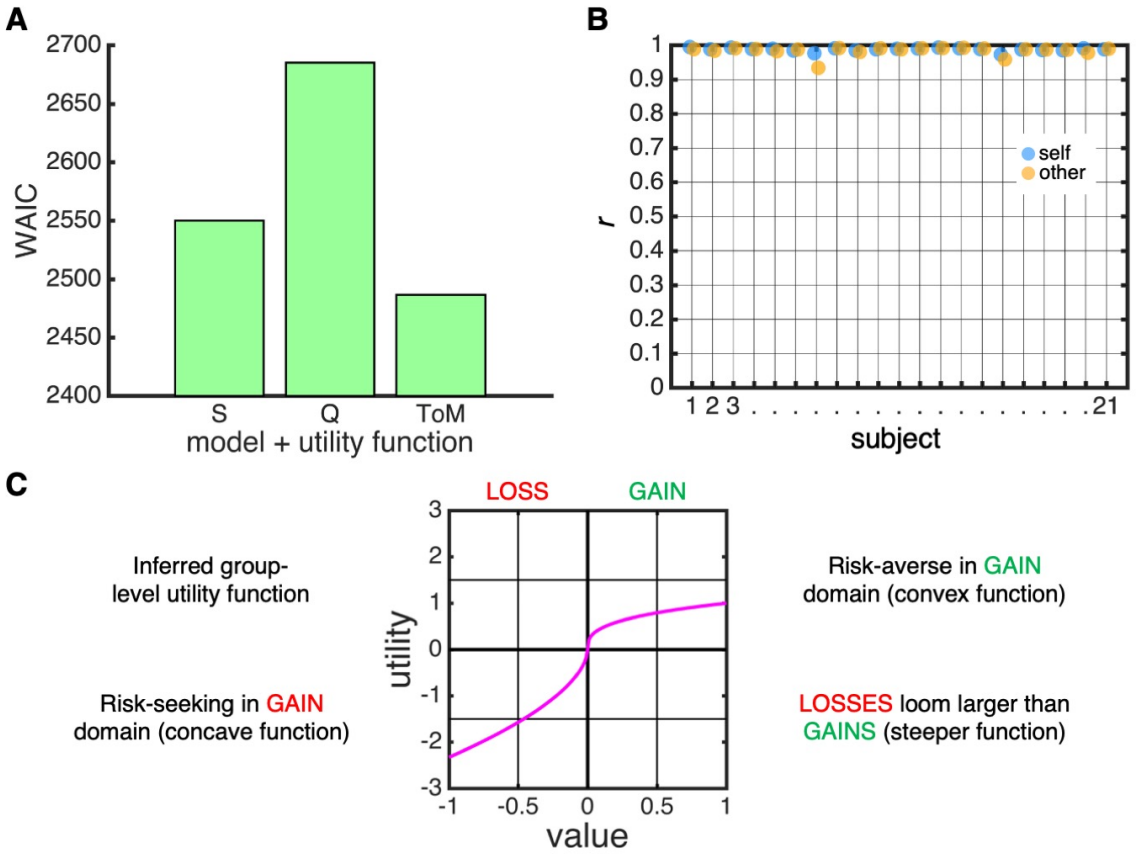

Figure S6. Risk and/or loss aversion. To assess the potential impact of risk and/or loss aversion on the model-based confidence estimates, we ran a control analysis in which we re-fitted the best-fitting version of each model class (S-B2T2P1, Q-B2T2P1 and ToM-B2T2P1) after including a utility function. **(A)** The ToM-model still provided the best account of the behavioural data (sum of WAIC across subjects) – indicating that value-based processes did not bias our model comparison results. **(B)** The model-based confidence estimates under both the original ToM-model and the ToM-model with a utility function were highly correlated (Pearson's *r*) within every subject – supporting that the original ToM-model was sufficient to infer subjects' confidence from their PDWs. **(C)** Visualisation of the inferred group-level utility function obtained from the fit of the augmented ToM-model. We normalised the value range for ease of visualisation.

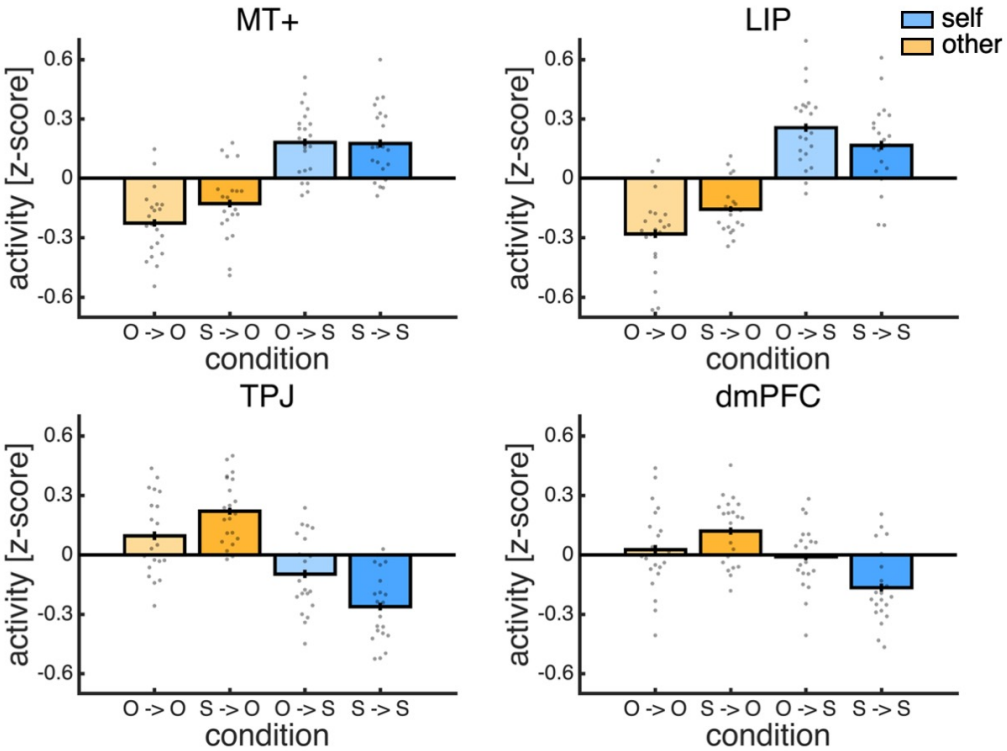

Figure S7. Repetition suppression analysis of agent-specific effects in sensory and social ROIs. In each panel, the stimulus-related activity estimates under the canonical hemodynamic response function are split according to the current trial type (yellow: other, O; blue: self, S) and whether the previous trial type was the same (light: same; dark: different). In MT+, there was only an effect of the current trial type (linear regression: current, $t(20) = 7.93$, $p < .001$; previous, $t(20) = 1.18$, $p = .252$; interaction, $t(20) = -1.10$, $p = .284$). However, in addition to an effect of the current trial type, there was an interaction between the previous and the current trial type (i.e. repetition suppression) in LIP (linear regression: current, $t(20) = 7.73$, $p < .001$; previous, $t(20) = 0.36$, $p = .719$; interaction, $t(20) = -3.42$, $p = .003$), TPJ (linear regression: current, $t(20) = -6.16$, $p < .001$; previous, $t(20) = -0.52$, $p = .609$; interaction, $t(20) = -3.40$, $p = .003$) and dmPFC (linear regression: current, $t(20) = -3.61$, $p = .002$; previous, $t(20) = -0.66$, $p = .516$; interaction, $t(20) = -2.80$, $p = .011$). Under the repetition suppression framework, if a region is primarily involved in self-inference, then activity should be highest on self-trials preceded by other-trials, and lowest on other-trials preceded by other-trials – a pattern observed for LIP. By contrast, if a region is primarily involved in other-inference, then activity should be highest on other-trials preceded by self-trials, and lowest on self-trials preceded by self-trials – a pattern observed for TPJ and dmPFC. Data are represented as group mean ± SEM. Each dot is a subject.

parameter	range	sampling	note
sensation (all)			
noise, σ	[.01,2]	$\sigma \sim normal(a,b)$ where a and b are obtained from a posterior distribution over sensory noise as fitted in the prescan session	
sensory sample, x	[-1,1]	$x \sim normal(a,b)$ where a is the motion stimulus, $k\theta$, and b is noise, σ	
softmax (all)			
bias, β_0	[-10,10]	$\beta_0 \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(0,10)$ and a: $[-10,10]$, $b \sim uniform(.01,10)$	same (B1) or different (B2) bias for self and other
temperature, β_1	[.01,2]	$\beta_1 \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(.2,2)$ and a: $[.01,2]$, $b \sim uniform(.01,2)$	same (B1) or different (B2) temperature for self and
stimulus space			
factor, p	[.01,10]	$p \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(2,4)$ and a: $[.01,10]$, $b \sim uniform(.01,4)$	linearly-spaced (P0) or re-scaled (P1) stimulus space
learning (Q, T)			
alpha, α	[0,1]	$\alpha \sim beta(a,b)$, where a and b are group-level parameters, $a = x * y^{-2}$, $x \sim normal(.2,1)$ and x: $[0,1]$, $y \sim uniform(0,2)$, $b = y^{-2} - a$	
Utility (S, Q, T)			
risk loss, R		$R \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(.6,1)$ and a: $[.01,1]$, $b \sim uniform(.01,1)$	
risk gain, T		$T \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(.4,1)$ and a: $[.01,1]$, $b \sim uniform(.01,1)$	
loss aversion, S		$S \sim normal(a,b)$, where a and b are group-level parameters, $a \sim normal(1,4)$ and a: $[.01,10]$, $b \sim uniform(.01,10)$	

Table S1. Overview of model parameters and priors as implemented in STAN. Model classes are abbreviated as follows: S, S-model; Q, Q-model; T, ToM-model; T. See online repository for code implementation. Here '∼' indicates 'is distributed as' and ':' indicates 'range'.

contrast	label	voxels at $p < .001$	peak z-score	p (cluster FWE corrected)	peak voxel MNI coordinates			
self > other	extrastriate, parietal, motor, cingulate	23632	6.25	<.001	-34	-16	68	LR
	thalamus	252	5.95	<.001	-22	-26	-6	L
	thalamus	145	5.49	<.001	22	-26	4	R
	inferior frontal gyrus	116	5.22	.002	-58	6	26	L
	insula	656	5.04	<.001	42	0	8	R
	striatum	222	4.40	<.001	12	18	-8	R
	striatum	192	4.36	<.001	18	-10	22	R
other > self	dorsomedial frontal	1941	5.30	<.001	14	30	52	LR
	temporoparietal	892	5.30	<.001	-48	-56	26	L
	temporoparietal	741	5.23	<.001	50	-62	36	R
	temporal	848	4.91	<.001	-58	-38	0	L
	precuneus	607	4.85	<.001	2	-60	38	LR
	middle frontal gyrus	177	4.75	<.001	48	24	36	R
	temporal	525	4.42	<.001	66	-28	-8	R
	middle frontal gyrus	467	4.41	<.001	-36	6	46	L
	temporal	82	3.85	.014	58	-8	-30	R

Table S2. Summary of significant activations for GLM1: decision phase. All reported activations are significant at $p < .05$, FWE-corrected for multiple comparisons. Cluster-defining threshold: $p < .001$, uncorrected. FWE: familywise error. MNI: Montreal Neurological Institute. L: left. R: right.

contrast	label	voxels at $p < .001$	peak z-score	p (cluster FWE corrected)	peak voxel MNI coordinates			
self > other	posterior cingulate	87	4.16	.009	-12	-52	10	LR
	ventromedial prefrontal	66	3.95	.037	2	38	-4	LR
other > self	middle frontal gyrus	258	5.44	<.001	-42	10	48	L
	temporoparietal	719	5.30	<.001	-48	-58	24	L
	temporal	503	5.30	<.001	64	-24	-12	R
	temporoparietal	550	5.23	<.001	54	-60	30	R
	dorsomedial frontal	1751	5.04	<.001	0	32	52	LR
	temporal	467	5.01	<.001	-64	-46	-2	L
	precuneus	176	4.96	<.001	6	-66	42	LR
	middle frontal gyrus	316	4.74	<.001	44	22	38	R
	temporal	414	4.68	<.001	-50	12	-28	L
	orbitofrontal	140	4.63	<.001	42	52	-12	R
	orbitofrontal	529	4.48	<.001	-34	54	-10	L
	orbitofrontal	97	4.44	.004	52	22	-6	R
	cerebellum	107	4.43	.002	-8	-80	-28	LR
	orbitofrontal	177	4.36	<.001	-28	22	-8	L
	precuneus	83	4.26	.011	6	-52	34	LR
	orbitofrontal	89	4.21	.007	-10	48	-24	LR
	temporal	84	4.17	.010	58	-38	2	R
	cerebellum	116	4.16	.001	-30	-84	-30	L
	orbitofrontal	75	3.98	.019	30	18	-6	R

Table S3. Summary of significant activations for GLM1: gamble phase. All reported activations are significant at $p < .05$, FWE-corrected for multiple comparisons. Cluster-defining threshold: $p < .001$, uncorrected. FWE: familywise error. MNI: Montreal Neurological Institute. L: left. R: right.