

Mechanisms underlying cortical activity during value-guided choice

Laurence T Hunt^{1,2}, Nils Kolling¹, Alireza Soltani³, Mark W Woolrich^{2,4}, Matthew F S Rushworth^{1,2} & Timothy E J Behrens^{2,5}

When choosing between two options, correlates of their value are represented in neural activity throughout the brain. Whether these representations reflect activity that is fundamental to the computational process of value comparison, as opposed to other computations covarying with value, is unknown. We investigated activity in a biophysically plausible network model that transforms inputs relating to value into categorical choices. A set of characteristic time-varying signals emerged that reflect value comparison. We tested these model predictions using magnetoencephalography data recorded from human subjects performing value-guided decisions. Parietal and prefrontal signals matched closely with model predictions. These results provide a mechanistic explanation of neural signals recorded during value-guided choice and a means of distinguishing computational roles of different cortical regions whose activity covaries with value.

Deciding on the best course of action amongst a range of competing alternatives has been a fundamental problem that has been addressed in the fields of economics¹, psychology², behavioral ecology³, machine learning⁴ and, more recently, cognitive neuroscience^{5–8}. For neural circuits to select the choice yielding the greatest long-term reward, it has been proposed that these circuits should take inputs reflecting the subjective value of alternatives and compare these inputs to form a categorical decision⁸. Representations of value have been found in many cortical and subcortical brain regions^{9–18}, but whether and how activity changes in these representations might constitute the decision process itself are unknown. This uncertainty is partly a consequence of not knowing how the signature of a decision would manifest itself at the level of the activity that can be recorded in a population of neurons.

One potential neuronal mechanism for value comparison is competition by mutual inhibition^{19,20}. In this class of models, separate pools of neurons representing different options are excited by the value of their respective options, but inhibit each other such that activity only survives in the eventual winning pool. This mechanism is particularly attractive, as it can be implemented in networks of neurons that respect known neurobiology²⁰. Indeed, such models accurately predict single-cell activity in the parietal cortex during perceptual decisions²¹.

It has been proposed that similar mechanisms might also underlie value-guided choice, but this proposal has rarely been tested empirically^{10,22}. An important problem is that the model predictions are of single-unit activity, but it is impossible to simultaneously measure this across the many brain regions that exhibit value-related activity. However, if such inhibitory mechanisms were to exhibit

a characteristic signature that could be measured in the summed activity of the local network rather than in single-cell activity, then we could use imaging techniques to search for this signature across the entire brain and isolate those regions that are fundamental to value comparison.

We adopted such an approach. We analyzed a biophysically realistic network model of decision-making to generate predictions of the temporal dynamics of value correlates in local field potentials. We then applied the exact same analysis to source-reconstructed magnetoencephalography (MEG) data, a whole-brain human imaging technique that affords the requisite temporal resolution to test model predictions. Notably, MEG allows coverage of signal from the entirety of neocortex, allowing for predictions to be tested in multiple brain regions simultaneously with high temporal resolution. Regions of ventromedial prefrontal and superior parietal cortex matched well with the biophysical model, implicating them in value comparison. Value correlates in other cortical regions matched poorly, suggesting that they are involved in separate computational processes that covary with value.

RESULTS

Biophysical model predictions

We used a mean-field version²³ of a biophysical cortical attractor network model²⁰ to derive predictions of the temporal dynamics of activity in a cortical region that selects between inputs reflecting the value of two options. The model comprises two populations of excitatory pyramidal cells selective for each option, with strong recurrent excitation between cells of similar selectivity and effective inhibition between the two pools mediated by inhibitory interneurons²⁰

¹Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford, UK. ²FMRIB Centre, University of Oxford, John Radcliffe Hospital, Oxford, UK. ³Howard Hughes Medical Institute and Department of Neurobiology, Stanford University School of Medicine, Stanford, California, USA. ⁴Oxford Centre for Human Brain Activity (OHBA), University Department of Psychiatry, Warneford Hospital, Oxford, UK. ⁵Wellcome Trust Centre for Neuroimaging, University College London, London, UK. Correspondence should be addressed to L.T.H. (lhunt@fmrib.ox.ac.uk) or T.E.J.B. (behrens@fmrib.ox.ac.uk).

Received 13 October 2011; accepted 16 November 2011; published online 8 January 2012; doi:10.1038/nn.3017

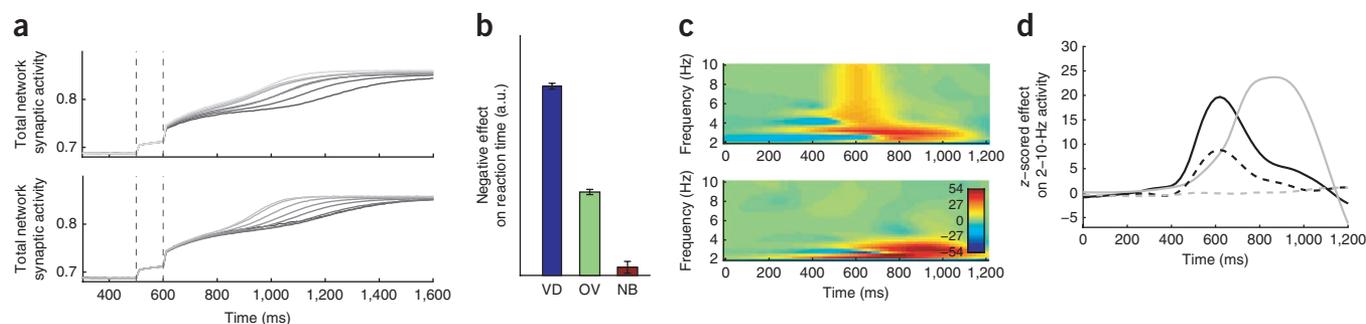


Figure 1 Predictions of neural activity from cortical attractor network model. **(a)** Top, summed network postsynaptic currents as a function of time through trial, sorted and binned into trials with high overall value (lighter shades of gray) through trials with low overall value (dark gray and black). Bottom, as top panel, resorted and binned by value difference between chosen and unchosen options. **(b)** Effect of value difference (VD), overall value (OV) and no brainer (NB) trials on reaction time, estimated using multiple regression (mean \pm s.e. of effect size; y axis is flipped, so positive values equate to a negative effect on reaction times). a.u., arbitrary units. **(c)** Time-frequency spectra of effects of overall value (top) and value difference (bottom) on network model activity, estimated with multiple regression. Color indicates z statistic. **(d)** z-scored effect of overall value (on frequency range 3–9 Hz, black lines) and value difference (on frequency range 2–4.5 Hz, gray lines); solid lines indicate correct trials and dashed lines indicate incorrect trials.

(see Online Methods). This effective inhibition mediates a competition between the two excitatory pools, with one pool ending up in a high-firing attractor state (chosen option) and the other pool staying in a low-firing attractor state (unchosen option). Neurons selective for option o receive inputs r_o at firing rates proportional to the subjective value of that option, sEV_o . The neurons also receive background noise inputs and currents from other cells in the network. Notably, the network has very few free parameters that are not otherwise constrained by their biophysical plausibility. The behavior of single units in the network has been described elsewhere^{20,24}; here we focus on predictions suited to investigation with MEG, namely behavior of the summed input currents to all pyramidal cells²⁵.

We simulated network behavior using a set of trials with varying sEV_o (as used in the human experiment, below). We sorted trials by overall value ($sEV_1 + sEV_2$; **Fig. 1a**) and value difference ($sEV_{\text{chosen}} - sEV_{\text{unchosen}}$; **Fig. 1a**). In both cases, the network attracted faster to a decision when overall value or value difference were higher, yielding the prediction of decreased reaction times under these conditions. We tested this prediction more formally using a multiple regression in which model reaction times were predicted as a function of both overall value and value difference; both variables were found to have a negative effect on reaction time (**Fig. 1b**). The model reaches an asymmetric attractor state more quickly when the basin of attraction for this option is larger as a result of larger value difference (which determines the difference between the two inputs). An increase in overall value causes the network activity to rise faster and diverge faster, which also results in faster reaction time.

We then performed a time-frequency analysis of network responses, which aided our subsequent comparison of model predictions with MEG data. We used Morlet wavelets to decompose network activity on each trial²⁶ and regressed the decomposed data onto overall value and value difference. Network transitions typically took several hundred milliseconds to occur, and most of the important model predictions were therefore limited to frequencies ranging from approximately 2–10 Hz (**Fig. 1c**). Overall value had a broadband effect on model activity in the 3–9-Hz frequency range soon after selective inputs were delivered to the network (**Fig. 1c**), whereas value difference had a later and slightly lower frequency effect, predominantly in the 2–4.5-Hz range (**Fig. 1c**). The different frequencies reflect the fact that overall value affects the population synaptic input earlier and over a shorter time period than value difference. The effect of the two regressors on network responses is a reflection of the fact

that network transitions occur at different speeds depending on the input presented; thus, the network does not explicitly ‘represent’ such quantities, but these effects are a manifestation of trial-to-trial variability in the speed of the different network transitions. If we collapsed across the relevant frequencies, the temporal progression from an overall value signal to a value difference signal could be clearly seen (**Fig. 1d**). It was also found that, on trials in which the network model made an error (that is, $sEV_{\text{chosen}} < sEV_{\text{unchosen}}$), there was an effect of overall value on the model’s activity, but no clear effect of value difference (**Fig. 1d**). However, it should be noted that error trials inherently covered a smaller range of value differences than correct trials, which may have caused the absence of any effect. Thus, the key predictions that we derived from the model were the temporal evolution from an overall value signal to a difference value signal, the difference in the frequency of the response, with

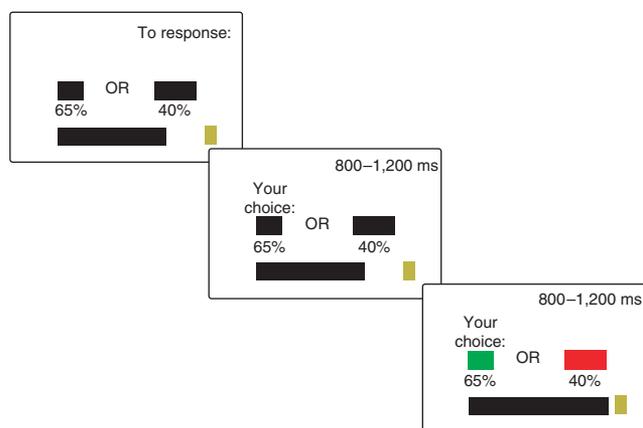


Figure 2 Value-based decision task. Task schematic. Subjects repeatedly chose between two risky prospects to obtain monetary reward. Stimuli consisted of a rectangular bar, whose width determined the amount of reward available, and a number presented underneath the bar, whose value determined the probability of receiving reward on that option. Stimuli were drawn such that reward magnitude and probability were never identical across the two options; subjects needed to integrate across stimulus dimensions to make optimal choices. On some trials, however, both probability and magnitude were larger on one side than the other (no brainer trial). Subjects had unlimited time to respond and received feedback on both chosen and unchosen options, green for the rewarded option(s) and red for the non-rewarded option(s).

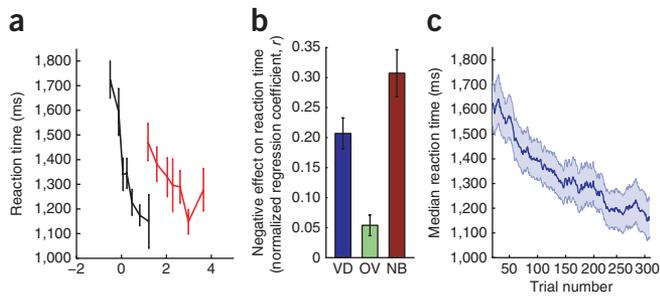


Figure 3 Subject behavior. (a) Reaction time (mean \pm s.e.) for an example subject, as a function of subjective value difference (black) and subjective overall value (red). (b) Effects of value difference (VD), overall value (OV) and no brainer trials (NB) on subject reaction times (mean \pm s.e. across subjects), estimated using linear regression. y axis is flipped; positive values equate to a negative effect on reaction times. (c) Running group mean \pm s.e. of reaction time (smoothed across 40 trials) as a function of trial number (also see **Supplementary Figs. 4 and 5**).

value difference dominating responses at lower frequencies than overall value, and the presence of an overall, but little or no difference, signal on error trials. These predictions were found to be robust to variation in model parameters determining the degree of recurrent excitation in the model (**Supplementary Discussion, Supplementary Table 1 and Supplementary Figs. 1–3**).

A distributed network of task-sensitive areas

We designed a simple value-guided choice task to test these predictions. Subjects ($n = 30$) repeatedly selected between two options of differing value (**Fig. 2**) while undergoing MEG. Each option had a certain number of points available, represented by the width of an onscreen bar, and a probability of obtaining those points, represented by a percentage underneath the bar. The aim was to accumulate points (displayed on a progress bar) to reach a gold target, at which point monetary reward was delivered and the progress bar was reset to its initial position. To accumulate maximal returns, subjects need to compute the objective Pascalian value (bar width multiplied by probability of winning, denoted EV_o for option o) and select the

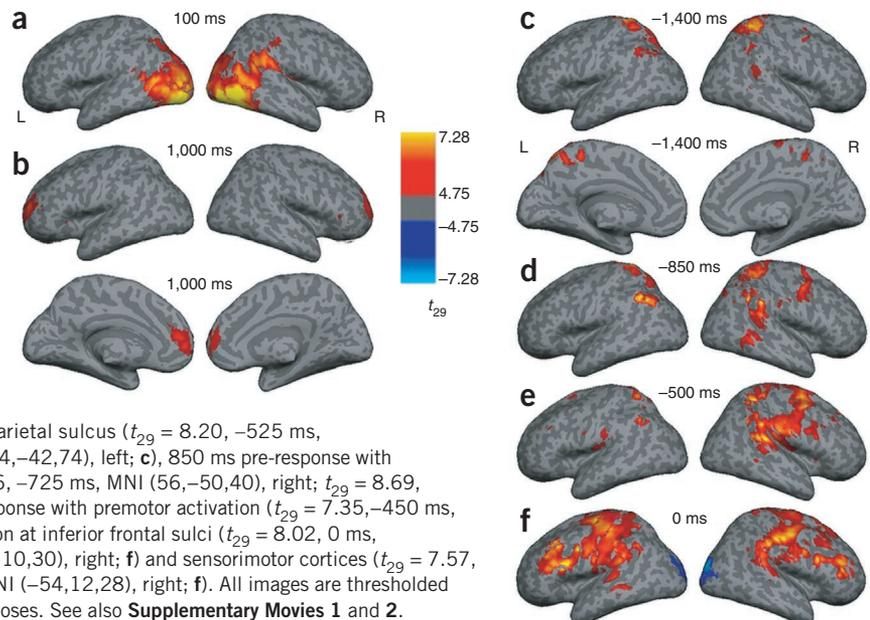
option with the higher value on each trial. In fact, most subjects tended to overweight low probabilities of winning and underweight high probabilities, and exhibited a concave utility function, consistent with predictions from prospect theory (**Supplementary Fig. 1 and Supplementary Table 2**)²⁷. Subject reaction times correlated negatively with both the difference in subjective option values ($sEV_{\text{chosen}} - sEV_{\text{unchosen}}$) and with the overall value of the decision ($sEV_1 + sEV_2$), consistent with model predictions (**Fig. 1b**) (**Fig. 3a,b**). We carried out a multiple linear regression of value difference ($t_{29} = -7.98$, $P < 0.0005$) and overall value ($t_{29} = -2.36$, $P < 0.05$) on reaction times across all subjects (**Fig. 3b and Supplementary Fig. 5**). We also included some trials in which both reward magnitude and probability were higher on one option than the other. There was an additional bonus in speed beyond that related to value for these ‘no brainer’ trials ($t_{29} = -8.32$, $P < 0.0005$; **Fig. 3b**). Subjects were therefore faster on average on these trials than on those in which probability and magnitude advocated opposing choices, and therefore needed to be translated into a ‘common currency’ in which the two stimulus features could be equated. There was a steady decrease in reaction time as subjects progressed through the task (**Fig. 3c**), without any coincident change in parameters describing choice behavior (**Supplementary Table 3**), suggesting that subjects became less deliberative and more automated in their choices as they became familiar with the task.

We used linearly constrained minimum variance beamforming²⁸ to spatially filter MEG data to locations in source space. We epoched data with respect to both stimulus onset and subject response and focused our analyses on responses in the 2–10-Hz frequency range, in accordance with model predictions. We first used a whole-brain statistical parametric mapping analysis to look for areas showing a main effect of performing the task relative to a pre-stimulus (–300 ms to –100 ms) or post-response (+100 ms to +300 ms) baseline. We hypothesized that, in addition to areas that are important to stimulus valuation such as ventromedial prefrontal cortex, the stimulus-locked analysis would reveal early visual areas that are involved in basic stimulus processing and the response-locked analysis would reveal areas that are involved in visually guided manual movements in parietal and premotor cortices²⁹, in addition to primary motor areas.

Figure 4 Main effect of task performance on activity in the 2–10-Hz frequency range.

(a,b) Stimulus-locked activity. Group t maps of effect of task performance relative to a –300 to –100 ms (pre-stimulus) baseline are shown at 100 ms post-stimulus with early visual activation (peak $t_{29} = 10.00$, 100 ms, MNI (40,–74,6)), 1,000 ms post-stimulus with activation at frontal pole ($t_{29} = 7.23$, 1,125 ms, MNI (22,58,26); b) and ventromedial prefrontal cortex ($t_{29} = 5.20$, 1,000 ms, MNI (43,60,35); b).

(c–f) Response-locked activity. The effects of task performance relative to a +100 ms to +300 ms (post-response) baseline are shown at 1,400 ms pre-response with activation at pSPL and posterior cingulate ($t_{29} = 7.05$, –1,625 ms (pre-response), MNI (18,–44,62); c) and mid-intraparietal sulcus ($t_{29} = 8.20$, –525 ms, MNI (30,–46,56), right; $t_{29} = 7.55$, –700 ms, MNI (–24,–42,74), left; c), 850 ms pre-response with activation at angular and supramarginal gyri ($t_{29} = 8.46$, –725 ms, MNI (56,–50,40), right; $t_{29} = 8.69$, –725 ms, MNI (–50,–60,42), left; d), 500 ms pre-response with premotor activation ($t_{29} = 7.35$, –450 ms, MNI (38,–2,64); e), the time of response with activation at inferior frontal sulci ($t_{29} = 8.02$, 0 ms, MNI (–54, 12, 28), left; $t_{29} = 7.55$, –75 ms, MNI (48,10,30), right; f) and sensorimotor cortices ($t_{29} = 7.57$, –75 ms, MNI (–50,–28,58), left; $t_{29} = 8.02$, 0 ms, MNI (–54,12,28), right; f). All images are thresholded at $t > 4.75$ ($P < 5 \times 10^{-5}$ uncorrected) for display purposes. See also **Supplementary Movies 1 and 2**.



A distributed network of areas was found to be task sensitive at these frequencies (Fig. 4a–f, Supplementary Movies 1 and 2). Stimulus-locked, early visual cortex activation (Fig. 4a) was followed by slowly ramping bilateral activation at the frontal pole and ventromedial prefrontal cortex (Fig. 4b). Although 2–10-Hz activity in these frontal regions peaked relatively late in the trial (1,000 ms after stimulus onset), it ramped from a much earlier point in the trial. Response-locked, prolonged activation spread from a mid-posterior portion of the superior parietal lobule, which extended medially into the marginal ramus of the posterior cingulate sulcus (Fig. 4c), to a bilateral medial portion of the mid-intraparietal sulcus (Fig. 4d). This was followed by bilateral activation of the angular/supramarginal gyri (Fig. 4d) and right premotor cortex (Fig. 4e), and, finally, bilateral inferior frontal sulci and primary sensorimotor cortices (Fig. 4f) were activated at the time of the response.

Predicted model activity in parietal and prefrontal cortex

Having isolated areas that showed changes in activity relative to baseline, we then examined whether activity in these regions covaried with decision values and where this activity matched with predictions derived from the biophysical decision model. Notably, by selecting regions on the basis of the main effect of task versus baseline, we ensured that we would not be subject to a selection bias when examining these regions for value-related activity. We also investigated activity in several a priori defined areas that are commonly found to be important in functional magnetic resonance imaging (fMRI) studies of decision making, bearing in mind that value correlates might not be restricted to regions showing a main effect of task versus baseline. We applied the exact same analysis to the time series from the source-reconstructed MEG data as we had applied to the biophysical model (Fig. 1).

We found that activity in the right posterior superior parietal lobule (pSPL) bore several hallmarks of the biophysical model (Fig. 5a). On trials in which subjects chose the option

with higher subjective value (correct trials), activity in pSPL showed a broad correlate of overall value between 3 and 10 Hz ($P < 0.0005$, permutation test, cluster corrected for multiple comparisons across time), followed by a lower frequency (2–4 Hz) correlate of value difference ($P < 0.01$, corrected; Fig. 5a), as predicted by the model (Fig. 1c). When we collapsed across the relevant frequencies (Fig. 5a), activity in these correct trials differed from that in error trials; error trials showed a positive correlate of overall value ($P < 0.05$, corrected; Fig. 5a), but no such positive correlate of value difference ($P > 0.5$; Figs. 1d and 5a). Finally, we tested the model prediction that across subjects there would be a behavioral speed-accuracy tradeoff, elicited by varying the degree of recurrent excitation in the network model, and that this would predict cross-subject variance in neural data. This prediction was also found to hold in pSPL (see Supplementary Discussion, Supplementary Figs. 6 and 7).

We also investigated whether the main effects of task performance in this region were affected by factors that have been shown behaviorally to modulate reaction time independently of value. We looked for changes in activity in early trials relative to late trials (in which reaction time was speeded; Fig. 3c) and compared activity in trials in which reward magnitude and probability advocated opposing choices with activity on no brainer trials (in which an additional bonus to reaction time was present beyond that explained by overall value or value difference; Fig. 3b). There was some difference between the patterns of activity in pSPL on these trials; an increase in 2–5-Hz power relative to baseline that was present on the first half of trials was largely absent on the second half of trials (Fig. 5a). A similar distinction could be seen

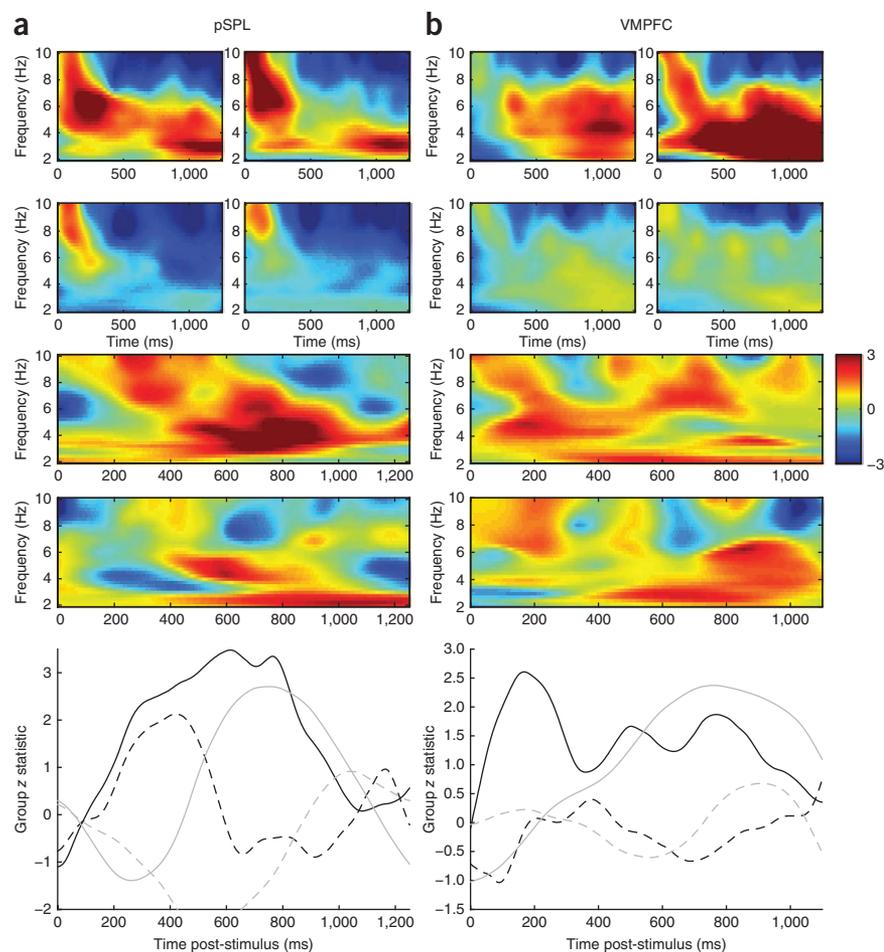


Figure 5 pSPL (MNI 18, –44, 62 mm) and VMPFC (MNI 6, 28, –8 mm) show several value-related hallmarks of the biophysical network model. **(a)** pSPL results. Top panels, main effect of task performance in pSPL relative to pre-stimulus baseline on first half of trials (top left) and second half of trials (bottom left); main effect of task performance on trials where reward magnitude and probability advocate opposing choices (top right), and no brainer trials (bottom right). Color indicates group z statistic. Middle panels, time-frequency spectra of effects of overall value (top) and value difference (bottom) on activity in pSPL, estimated using multiple regression. Analysis is equivalent to that performed in Figure 1c on biophysical model. Color indicates group z statistic. Bottom, effect of overall value (3–9 Hz, black) and value difference (2–4.5 Hz, gray) on correct (solid) and error trials (dashed lines, respectively). **(b)** VMPFC results. Data are presented as in **a**. Top panels, main effect of task in VMPFC. Middle panels, VMPFC effects of overall value and value difference, as for pSPL, but restricted to first half of experiment. Bottom, VMPFC collapsed value effects, as for pSPL, but restricted to first half of experiment.

between activity on trials in which reward magnitude and probability advocated opposing choices and a common currency representation might need to be formed, and no brainer trials (Fig. 5a).

We also investigated value-related activity in ventromedial prefrontal cortex (VMPFC), focusing our analyses on a subregion that has often been shown to signal value-related metrics during decision tasks^{11–16,30,31}. Notably, there has been debate over the precise role of this region in value-guided choice^{6,32}, perhaps triggered by the heterogeneity of responses that have been observed; in some fMRI studies, VMPFC has been found to signal a difference between chosen and unchosen values^{14,31}, whereas in others it has appeared to signal the overall value of available reward¹⁵ or the value of just the chosen option¹⁶. In VMPFC, there was an even more notable distinction between those situations in which subjects would be more deliberative and exhibit slower reaction times versus later or no brainer trials (Fig. 5b). VMPFC recruitment steadily decreased through the task, as could be seen more clearly when trials were further subdivided into separate quartiles of the experiment (Supplementary Fig. 8). We found that this region transitioned from signaling overall value ($P < 0.05$, corrected) to value difference ($P < 0.05$, corrected) specifically if we restricted our analysis to the first half of trials in which it was task active (Fig. 5b). When we directly contrasted the effect of overall value and value difference on early and late trials, we found that only the value difference signal was significantly stronger on earlier trials in this region ($P < 0.05$, corrected; Supplementary Fig. 9). There was not a significant effect of either overall value or value difference on error trials ($P > 0.5$; Fig. 5b), although the somewhat weaker signals in this region relative to pSPL may result from the relative insensitivity of MEG to deep, anterior sources, as opposed to posterior, superficial ones^{33,34}, and from the analysis including only half the number of trials.

One possible concern with the differences between the first and second halves of the experiment is that it might reflect more trivial cognitive differences, such as subject fatigue, rather than a change in the cortical networks underlying choice behavior. To address these concerns, we performed an additional whole-brain analysis in which we searched for regions coding more strongly for value difference in the second than in the first half of the experiment, that is, the opposite pattern of activity as that witnessed in VMPFC. A bilateral portion of the anterolateral intraparietal sulcus, more lateral than the main effect pSPL activation described above, selectively reflected value difference in the second half of trials (Supplementary Fig. 10). In this region, there were also no clear differences between the main effect of task performance on early versus late trials or on harder trials versus no brainer trials (Supplementary Fig. 11).

Lastly, we also searched for effects of value in other regions identified in the main effect contrast of task versus baseline (Fig. 4) and in several regions defined a priori from previous fMRI studies of value-based choice. In these analyses, we found that several areas exhibited value-dependent activity, but none of these regions matched well with predictions from the biophysical decision model (Supplementary Fig. 12). We hypothesize that the value correlates in these regions might be better described by appealing to their role in other computational processes that are likely to covary with value, such as attention or response preparation. Alternatively, it may be the case that these other regions are involved in value comparison, but in a manner that is different from that proposed using the biophysical modeling approach.

DISCUSSION

The cortical correlates of value during decision under risk are typically spread over a distributed network of areas, but the unique

contribution of each of these areas to choice is unclear. A region involved in value comparison should receive inputs relating to the value of available options and then transform these inputs into a categorical choice. We used a biophysically plausible model that exhibits this property to derive predictions of the temporal dynamics of cortical activity. We applied linear regression to investigate the time points at which, and in which frequency bands, value correlates could be found in network activity. These responses typically occurred at low frequencies (<10 Hz), consistent with a slow integrative process. We then applied the same analysis to source-reconstructed MEG data, to identify regions involved in value comparison. A distributed network of areas were task sensitive at the relevant frequencies, but only pSPL and VMPFC closely matched predictions of the biophysical model, with the latter doing so selectively in trials early in the experiment. Other regions were found to show value correlates, but did not match closely with predictions from the biophysical model; this suggests that extensions to the model are necessary to fully capture the role of different brain regions in the task. Furthermore, MEG is limited in its ability to resolve sources from deep brain structures that do not possess an open field layout, such as in striatum¹⁷; thus, we could not address the role of alternative mechanisms for selection (dependent on cortico-basal ganglia loops).

An important feature of the biophysical model is the ability to slowly integrate value-related inputs that is afforded by its recurrent excitatory structure and long synaptic time constants mediated by NMDA receptors. It is not immediately obvious that value comparison should be subject to a process of integration in the same manner as a noisy sensory stimulus. However, the observed distribution of reaction times fits well with a process of integration, as has been investigated more closely in previous studies that used a drift diffusion model to predict reaction times^{35,36}. The drift diffusion model was originally designed to make predictions of behavioral data and has often been used to make predictions of single-unit activity during perceptual choice. However, because it essentially describes differences in activity between different populations of selective cells and ignores any nonselective activity, it is unclear how the model's output should be translated into a prediction of imaging measures such as MEG or fMRI. We elected to use a biophysical implementation of a competition model, which makes clear and explicit predictions of the measurable data. When we used the pseudo-variable in the drift diffusion model as a marker for integrated brain activity, we found differences between the predictions (Supplementary Discussion and Supplementary Fig. 13).

The predictions from the model also form a marked example of the distinction between two types of representation, content and functional representations, in cortical circuits³⁷. To the external observer, recording with an imaging technique (or an electrode), the content of the network appears to represent first the overall value and then the value difference between the two options. In contrast, the functional representations in the network, those used by the brain, are quite different. There is a representation of option values on the input to the network and a representation of choice on the outputs of the network, as should be decoded by a suitable downstream observer. The reason that the network shows value-related activity is simply that the same network transitions occur faster on high value and high value difference trials. Thus, although neural activity in the network may covary with the overall value and value difference, this content need never be decoded by another brain region. Thus, the extent to which the network can be said to functionally represent these two quantities in a meaningful way is questionable³⁷.

The region in pSPL that we isolated as matching with model predictions is close to the cytoarchitectonic region hIP3 (ref. 38), which may

be the human homolog of the medial intraparietal area (MIP). It is also referred to as IPS4 and DIPSA, which resembles macaque MIP³⁹. In the macaque, this region has often been implicated in visually guided movements of the forelimbs²⁹. Thus, it may have a role in integrating information to guide limb movements that is analogous to the role of LIP in generating saccades. This process of saccade generation is closely linked to the tracking of value associated with generating a saccade in a particular direction^{9,10}.

The region in VMPFC that we examined has often been found, using fMRI, to be responsive to the value of stimuli during decision tasks^{11–16,30,31}, but its precise role has been debated^{6,32}, perhaps as a result of the relative absence of published single-unit recording data in comparison with the nearby lateral orbitofrontal cortex^{5,7}. In early trials, this region was found to transition from signaling overall value to signaling value difference. Notably, this same transition was also recently found in single-unit recordings from the most ventral portion of the striatum¹⁷, which receives a particularly dense projection from VMPFC⁴⁰, and in prefrontal cortex⁴¹. In that study, similar to ours, monkeys combined two stimulus properties to form their decision, namely the reward magnitude and the delay to reward delivery. In our task, VMPFC was selectively activated in trials in which subjects had to combine probability and magnitude information to choose accurately. This is also consistent with the finding that lesions to this area, but not nearby lateral orbitofrontal cortex, produce impairments in value comparison³² and, more specifically, produce changes in tasks in which multiple dimensions have to be considered in forming a choice⁴².

Previous studies have attempted to apply a modeling approach to capture signals from distributed cortical regions during choice, measured using fMRI. These studies have made predictions on the basis of either drift diffusion models⁴³ or biophysically plausible networks⁴⁴, but the predictions of these models are heavily dependent on whether fMRI signal is assumed to reflect activity from all time points, including the point after a decision has been formed⁴⁴, or whether it only reflects activity until the decision threshold is reached⁴³. Moreover, several key predictions of these models also relate to how their activity evolves over time as a decision is made, and the slow hemodynamic response indicates that fMRI is limited in how well it can tease apart these predictions of temporal dynamics. We argue that it is important to use a time-resolved technique, such as MEG, to test these predictions.

Biophysically inspired models have also been used to infer the structure of connections between or within different cortical areas from MEG and electroencephalography data⁴⁵. However, these studies have not inferred the specific neuronal mechanism underlying a particular cognitive process, as we have proposed here. Our model performs the critical computation of transforming value-related inputs into a choice and does so in a way that captures single-unit activity during perceptual decision tasks. The application of this computational biophysical modeling approach may not be limited to decision-making protocols. Predictions might, for instance, be derived from biophysical models that have already been designed to capture single-unit data in inhibitory control or working memory processes⁴⁶. In models of working memory, for instance, gamma-band (30–70 Hz) responses can be elicited⁴⁶, and parametric modulation of input to these models may explain variation in gamma-band frequencies that have been observed during working memory tasks in frontal cortex⁴⁷. Alternatively, by varying internal parameters of a biophysical model, new predictions might be derived of the effects of cross-subject variation on cortical responses measurable with MEG and electroencephalography (see also **Supplementary Discussion**).

Because these parameters relate to specific biophysical properties, such as the density of network connectivity or the concentration of a specific neurotransmitter, it may be possible to directly relate these parameters to cross-subject variation in these properties, for instance, via local measurements of neurotransmitter concentrations⁴⁸ or perhaps genetic polymorphism or pharmacological challenge.

METHODS

Methods and any associated references are available in the online version of the paper at <http://www.nature.com/natureneuroscience/>.

Note: Supplementary information is available on the Nature Neuroscience website.

ACKNOWLEDGMENTS

We thank V. Litvak and G. Barnes for many helpful discussions, C. Stagg and K. Friston for comments on the manuscript, T. Nichols for assistance implementing four-dimensional cluster-based permutation testing, and S. Braeutigam and A. Rao for help with data collection. This work was supported by the Wellcome Trust (L.T.H., N.K., M.W.W. and T.E.J.B., grant reference numbers WT088312 and WT080540), the Consortium Of Neuroimagers for Non-invasive Exploration of Brain Connectivity and Tracts (CONNECT; L.T.H. and T.E.J.B.), the UK Medical Research Council (M.F.S.R.) and the UK Engineering and Physical Sciences Research Council (M.W.W.). The project CONNECT acknowledges the financial support of the Future and Emerging Technologies (FET) program in the Seventh Framework Program for Research of the European Commission, under FET-Open grant number 238292.

AUTHOR CONTRIBUTIONS

L.T.H., T.E.J.B. and M.F.S.R. designed experiment. L.T.H. and N.K. collected data. A.S. and L.T.H. built models and analyzed model predictions. M.W.W. wrote code for source reconstruction. L.T.H., T.E.J.B., N.K. and M.W.W. analyzed data. L.T.H., M.F.S.R. and T.E.J.B. wrote the paper. All of the authors discussed the results and commented on the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Published online at <http://www.nature.com/natureneuroscience/>.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

1. von Neumann, J. & Morgenstern, O. *Theory of Games and Economic Behavior* (Princeton University Press, 1944).
2. Dickinson, A. & Balleine, B. Motivational control of goal-directed action. *Learn. Behav.* **22**, 1–18 (1994).
3. Stephens, D.W. & Krebs, J.R. *Foraging Theory* (Princeton University Press, 1986).
4. Sutton, R. & Barto, A. *Reinforcement Learning: an Introduction* (MIT Press, 1998).
5. Rushworth, M.F. & Behrens, T.E.J. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nat. Neurosci.* **11**, 389–397 (2008).
6. Kable, J.W. & Glimcher, P.W. The neurobiology of decision: consensus and controversy. *Neuron* **63**, 733–745 (2009).
7. Padoa-Schioppa, C. Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**, 333–359 (2011).
8. Rangel, A., Camerer, C.F. & Montague, P.R. A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* **9**, 545–556 (2008).
9. Platt, M.L. & Glimcher, P.W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
10. Sugrue, L.P., Corrado, G.S. & Newsome, W.T. Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787 (2004).
11. Tom, S.M., Fox, C.R., Trepel, C. & Poldrack, R.A. The neural basis of loss aversion in decision-making under risk. *Science* **315**, 515–518 (2007).
12. Plassmann, H., O'Doherty, J.P. & Rangel, A. Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J. Neurosci.* **27**, 9984–9988 (2007).
13. Knutson, B., Taylor, J., Kaufman, M., Peterson, R. & Glover, G. Distributed neural representation of expected value. *J. Neurosci.* **25**, 4806–4812 (2005).
14. Boorman, E.D., Behrens, T.E.J., Woolrich, M.W. & Rushworth, M.S.F. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* **62**, 733–743 (2009).
15. Blair, K. *et al.* Choosing the lesser of two evils, the better of two goods: specifying the roles of ventromedial prefrontal cortex and dorsal anterior cingulate in object choice. *J. Neurosci.* **26**, 11379–11386 (2006).
16. Kable, J.W. & Glimcher, P.W. The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* **10**, 1625–1633 (2007).

17. Cai, X., Kim, S. & Lee, D. Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. *Neuron* **69**, 170–182 (2011).
18. Kennerley, S.W., Dahmubed, A.F., Lara, A.H. & Wallis, J.D. Neurons in the frontal lobe encode the value of multiple decision variables. *J. Cogn. Neurosci.* **21**, 1162–1178 (2009).
19. Usher, M. & McClelland, J.L. The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* **108**, 550–592 (2001).
20. Wang, X.J. Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* **36**, 955–968 (2002).
21. 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).
22. Soltani, A. & Wang, X.J. A biophysically based neural model of matching law behavior: melioration by stochastic synapses. *J. Neurosci.* **26**, 3731–3744 (2006).
23. Wong, K.F., Huk, A.C., Shadlen, M.N. & Wang, X.J. Neural circuit dynamics underlying accumulation of time-varying evidence during perceptual decision making. *Front. Comput. Neurosci.* **1**, 6 (2007).
24. Wong, K.F. & Wang, X.J. A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* **26**, 1314–1328 (2006).
25. Hämäläinen, M.S., Hari, R., Ilmoniemi, R.J. & Knuutila, J. Magnetoencephalography: theory, instrumentation and applications to noninvasive studies of the human brain. *Rev. Mod. Phys.* **65**, 413–497 (1993).
26. Tallon-Baudry, C., Bertrand, O., Delpuech, C. & Pernier, J. Oscillatory gamma-band (30–70 Hz) activity induced by a visual search task in humans. *J. Neurosci.* **17**, 722–734 (1997).
27. Tversky, A. & Kahneman, D. Advances in prospect theory: cumulative representation of uncertainty. *J. Risk Uncertain.* **5**, 297–323 (1992).
28. Van Veen, B.D., van Drongelen, W., Yuchtman, M. & Suzuki, A. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Trans. Biomed. Eng.* **44**, 867–880 (1997).
29. Johnson, P.B., Ferraina, S., Bianchi, L. & Caminiti, R. Cortical networks for visual reaching: physiological and anatomical organization of frontal and parietal lobe arm regions. *Cereb. Cortex* **6**, 102–119 (1996).
30. Gershman, S.J., Pesaran, B. & Daw, N.D. Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *J. Neurosci.* **29**, 13524–13531 (2009).
31. Serences, J.T. Value-based modulations in human visual cortex. *Neuron* **60**, 1169–1181 (2008).
32. Noonan, M.P. *et al.* Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc. Natl. Acad. Sci. U S A* **107**, 20547–20552 (2010).
33. Hillebrand, A. & Barnes, G.R. A quantitative assessment of the sensitivity of whole-head MEG to activity in the adult human cortex. *Neuroimage* **16**, 638–650 (2002).
34. Marinkovic, K., Cox, B., Reid, K. & Halgren, E. Head position in the MEG helmet affects the sensitivity to anterior sources. *Neurol. Clin. Neurophysiol.* **2004**, 30 (2004).
35. Sigman, M. & Dehaene, S. Parsing a cognitive task: a characterization of the mind's bottleneck. *PLoS Biol.* **3**, e37 (2005).
36. Milosavljevic, M., Malmaud, J., Huth, A., Koch, C. & Rangel, A. The drift diffusion model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgm. Decis. Mak.* **5**, 437–449 (2010).
37. deCharms, R.C. & Zador, A. Neural representation and the cortical code. *Annu. Rev. Neurosci.* **23**, 613–647 (2000).
38. Scheperjans, F. *et al.* Observer-independent cytoarchitectonic mapping of the human superior parietal cortex. *Cereb. Cortex* **18**, 846–867 (2008).
39. 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. *J. Neurosci.* **31**, 4087–4100 (2011).
40. Haber, S.N., Kim, K.S., Maily, P. & Calzavara, R. Reward-related cortical inputs define a large striatal region in primates that interface with associative cortical connections, providing a substrate for incentive-based learning. *J. Neurosci.* **26**, 8368–8376 (2006).
41. Kim, S. & Lee, D. Prefrontal cortex and impulsive decision making. *Biol. Psychiatry* **69**, 1140–1146 (2011).
42. Fellows, L.K. Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain* **129**, 944–952 (2006).
43. Basten, U., Biele, G., Heekeren, H.R. & Fiebach, C.J. How the brain integrates costs and benefits during decision making. *Proc. Natl. Acad. Sci. USA* **107**, 21767–21772 (2010).
44. Rolls, E.T., Grabenhorst, F. & Deco, G. Choice, difficulty, and confidence in the brain. *Neuroimage* **53**, 694–706 (2010).
45. Kiebel, S.J., Garrido, M.I., Moran, R., Chen, C.C. & Friston, K.J. Dynamic causal modeling for EEG and MEG. *Hum. Brain Mapp.* **30**, 1866–1876 (2009).
46. Compte, A., Brunel, N., Goldman-Rakic, P.S. & Wang, X.J. Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. *Cereb. Cortex* **10**, 910–923 (2000).
47. Meltzer, J.A. *et al.* Effects of working memory load on oscillatory power in human intracranial EEG. *Cereb. Cortex* **18**, 1843–1855 (2008).
48. Muthukumaraswamy, S.D., Edden, R.A., Jones, D.K., Swettenham, J.B. & Singh, K.D. Resting GABA concentration predicts peak gamma frequency and fMRI amplitude in response to visual stimulation in humans. *Proc. Natl. Acad. Sci. USA* **106**, 8356–8361 (2009).

ONLINE METHODS

Experimental task. Subjects repeatedly chose between two risky prospects to obtain monetary reward. Stimuli consisted of a rectangular bar, whose width determined the amount of reward available, and a number presented underneath the bar, whose value determined the probability of receiving reward on that option. The probabilities of winning on each option were independent; thus, on any given trial, both, neither or either option(s) might yield reward. Stimuli were drawn such that reward magnitude and probability were never identical across the two options; thus, subjects needed to integrate across stimulus dimensions to make optimal choices (see below). On some trials, however, both probability and magnitude were larger on one side than the other, a decision we classify as a no brainer. By design, the mean correlation between overall value and value difference (chosen–unchosen value) was kept at 0.31 ± 0.08 (mean \pm s.d.), allowing them to explain largely separate portions of variance in behavioral and neural data.

Decisions were presented onscreen until a response was made. After selection, the chosen option was highlighted for 800–1,200 ms jittered and outcomes were presented for 800–1,200 ms jittered. Feedback was presented on both chosen and unchosen options by turning a rewarded option green and an unrewarded option red. Stimuli were then removed and an intertrial interval of 500–800 ms was presented.

On choosing a rewarded option, a ‘winnings bar’ displayed at the bottom of the screen increased in magnitude in proportion to the width of the chosen option. When this winnings bar reached a gold target on the far right of the screen, £2 was added to subjects’ earnings and the winnings bar reset itself to its original size. Total typical earnings for the task ranged from £26 to £34. We collected a secondary dataset with high-resolution eyetracking to exclude ocular artifacts as a possible confound to signals recorded during the task. No major differences in task-related activity were seen, so data was collapsed across the two experiments.

Stimuli were presented on a screen situated 1.5 m away from the subject, inside the magnetically shielded room; stimuli were displayed via projector (refresh rate of 60 Hz) situated outside the room. Stimulus presentation and timing was controlled using Presentation software (Neurobehavioral Systems). All subjects provided informed written consent in accordance with local ethical guidelines, and the experiment was reviewed by Oxfordshire Research Ethics Committee C.

Behavioral analysis. Subjective utility functions were derived from prospect theory and were of the form

$$v(r_o) = r_o^\alpha$$

$$w(p_o) = \frac{p_o^\gamma}{(p_o^\gamma + (1 - p_o)^\gamma)^\gamma}$$

where r_o and p_o are the reward magnitude and probability of gaining reward, respectively, on outcome o . The subjective expected value of outcome o was calculated as:

$$sEV_o = v(r_o) \times w(p_o)$$

The probability of choosing each option was then calculated using a softmax choice rule

$$P(C = o) = \frac{e^{\frac{sEV_o}{\tau}}}{\sum_{i=1}^n e^{\frac{sEV_i}{\tau}}}$$

where n is the number of options (two for this study) and τ is a temperature parameter that determines the stochasticity of action selection. Values of α , γ and τ were fit by maximizing the likelihood of each subject’s choices in the experiment, using nonlinear fitting routines in Matlab (Mathworks); parameter values and comparison to a reduced (objective value) model are given in

Supplementary Table 2. The fitted values were used to calculate subjective expected values, which have been found to provide a better fit to neural data in value-guided decision tasks^{16,49}, which were then used as trial-wise regressors in analysis of MEG data.

We examined the effects of value, trial number and no brainer trials on each subject’s reaction time data using multiple regression. We entered log(reaction time) as the dependent variable, as it has a distribution that is closer to normal than that of reaction time. In one analysis (Fig. 2c), we entered the following regressors as independent variables: a constant (to model mean reaction time), a term to capture speeding of reaction times for left versus right choices (data not shown and non-significant; $t_{29} = -0.99$, $P = 0.33$), the difference in subjective values between chosen and unchosen options ($sEV_{\text{chosen}} - sEV_{\text{unchosen}}$), the summed overall subjective value of both options ($sEV_1 + sEV_2$), a term to capture any linear change in reaction times as a function of performing the task (Fig. 2d), a term to capture autocorrelation (containing the reaction time from the previous trial), and a term to capture any additional bonus for the trial being a no brainer (containing a 1 wherever a no brainer trial occurred and a 0 otherwise). We normalized third and fourth regressors before entry into the design matrix to account for any differences in the relative scales of subjective value functions across subjects. We plotted the mean \pm s.e. across subjects of parameter estimates for the third, fourth and seventh regressors from this regression, and tested for statistical significance using a two-tailed one-sample t test across subjects.

In a subsidiary analysis we included two additional regressors: the objective value difference and the objective overall value of each trial, which were both normalized. In this analysis, we orthogonalized the third regressor with respect to the objective value difference, and the fourth regressor with respect to the objective overall value. These orthogonalized regressors provide a further test of the nonlinearity of the subjective value functions used in the experiment (Supplementary Fig. 5).

MEG and MRI data acquisition. MEG data were sampled at 1,000 Hz on a 306-channel VectorView system (Elekta Neuromag), with one magnetometer and two orthogonal planar gradiometers at each of 102 locations distributed in a hemispherical helmet across the scalp, in a magnetically shielded room. A band-pass filter of 0.03–330 Hz was applied during acquisition. Head position was monitored at the beginning of each run and at 20-min intervals during each run using four head position indicator (HPI) coils attached to the scalp. HPI coil locations, head points from across the scalp and three anatomical fiducial locations (nasion, left and right pre-auricular points) were digitized using a Polhemus Isotrak II before data acquisition. Vertical electro-oculogram and electro-cardiogram were also measured to detect eye blinks and heartbeat, respectively. MRI data for forward model generation were acquired using an MP-RAGE sequence on a Siemens 3T TRIO scanner, with voxel resolution of $1 \times 1 \times 1 \text{ mm}^3$ on a $176 \times 192 \times 192$ grid (echo time = 4.53 ms, inversion time = 900 ms, repetition time = 2,200 ms).

MEG data preprocessing. External noise was removed from MEG data using the signal-space separation method, and adjustments in head position across runs (detected using HPI) were compensated for using MaxMove software, both implemented in MaxFilter version 2.1 (Elekta Neuromag). Continuous data were down-sampled to 200 Hz and low-pass filtered at 40 Hz, before conversion to SPM8 format (<http://www.fil.ion.ucl.ac.uk/spm/>). Eye blinks were detected from the electro-oculogram channel. Detected eye blinks were used to generate an average eye blink time course, on which principle components analysis was run to obtain spatial topographies describing the average eye blink; these were regressed out of the continuous data (as described in ref. 50, without inclusion of brain source vectors as co-regressors; see http://www.fmrib.ox.ac.uk/~lhunt/artifact_session.zip for an SPM-based tutorial). Data were epoched with respect to stimulus onset (–1,000 to 2,000 ms around stimulus, with –200 to 0 ms pre-stimulus baseline) and button press (–2,000 to 1,000 ms around response, again with –200 to 0 ms pre-stimulus baseline). Artifactual epochs and bad channels were detected and rejected via visual inspection, using FieldTrip visual artifact rejection routines.

MEG source reconstruction. All source reconstruction was performed in SPM8. Magnetic resonance images were segmented and spatially normalized to an MNI template brain in Talairach space; the inverse of this normalization was used to

warp a cortical mesh derived from the MNI template to each subject's MRI space. Digitized scalp locations were registered to head model meshes using an iterative closest point algorithm, to affine register sensor locations to model meshes. Forward models were generated on the basis of a single shell using superposition of basis functions that approximately corresponded to the plane tangential to the MEG sensor array. The forward models are implemented in FieldTrip's *forwinv* toolbox.

Linearly constrained minimum variance beamforming was used to reconstruct data to a grid across MNI space, sampled with a grid step of 7 mm. Beamforming constructs a spatial filter at each grid location to spatially filter the sensor space data, y , to the grid location of interest, r_i , with the aim of achieving unit bandpass response at the location of interest while minimizing the variance passed from all other locations. The data at the source location of interest, $d(r_i)$, is given by multiplying the beamformer weights vector, $w(r_i)$ by the original sensor data

$$d(r_i) = w(r_i) \times y$$

This can be repeated across all grid locations to give a whole-brain image.

The sensor covariance matrix for beamforming was estimated using data that was bandpass-filtered to the frequency band of interest, 2–10 Hz, using 0% regularization. For stimulus-locked analyses, we included all non-artifactual trials from stimulus onset to 1 s after stimulus onset. For response-locked analysis, we included all non-artifactual trials from 1.5 s before response onset to the time of the response.

Frequency domain analysis of MEG data and linear regression. At each trial, the source data $d(r_i)$ was decomposed into ten time-frequency bins linearly spaced between 2 and 10 Hz, by convolving the data with Morlet wavelets (Morlet factor 5). This yielded, at each trial, tr , frequency f , and time point, t , an instantaneous estimate of the power at that frequency. For contrasts of main effect versus baseline (Fig. 4 and Supplementary Movies 1 and 2), we subtracted the power of the data from -300 ms to -100 ms pre-stimulus (stimulus locked) or +100 ms to +300 ms post-response (response locked). Linear regression was then used to estimate the contribution of experimental variables that varied across trials to this value

$$d(r_i)^{tr,f,t} = \beta_0^{f,t} + \beta_1^{f,t} \times OV^{tr} + \beta_2^{f,t} \times VD^{tr}$$

where OV is the subjective overall value ($sEV_{\text{chosen}} + sEV_{\text{unchosen}}$) and VD is the subjective value difference ($sEV_{\text{chosen}} - sEV_{\text{unchosen}}$). Overall value and value difference were normalized before regression so that they occupied a similar range of values across subjects. $\beta_0^{f,t}$, $\beta_1^{f,t}$, $\beta_2^{f,t}$ and their associated variances, $\text{var}(\beta_0^{f,t})$, $\text{var}(\beta_1^{f,t})$ and $\text{var}(\beta_2^{f,t})$, were estimated using ordinary least-squares regression.

The parameter estimates for $\beta_1^{f,t}$ and $\beta_2^{f,t}$, normalized by $\text{var}(\beta_1^{f,t})$ and $\text{var}(\beta_2^{f,t})$, were then averaged across frequencies to gain a single estimate of the contribution of overall value and value difference to 2–10-Hz activity at each time point through the experiment. These data were then subsampled at 20 Hz and submitted to second-level analysis.

For the whole-brain analyses of each time point and each location, a one-sample t test of first-level statistics was performed across subjects. A nonparametric permutation test was used to correct for multiple comparisons across voxels and time. The principle behind this statistical test is identical to that of the cluster-size permutation test that is frequently used during inference on functional MRI statistical images, with the slight complication that clusters are now formed in four dimensions (X, Y, Z and time), rather than three. We ran a Matlab-based clustering algorithm (available at http://www.fmrrib.ox.ac.uk/~lhunt/randomise_4D_clusters.zip) in combination with FMRIB Software Library's randomize version 2.6 to identify the size of each four-dimensional cluster. From each permutation, we took the maximum cluster size and used this to build a null distribution of cluster sizes. We then compared the size of clusters from the true t statistic image with that of the null distribution, and reported those clusters with a significance of $P < 0.01$, corrected for multiple comparisons across both space and time.

We performed additional tests on data beamformed to the clusters identified in the whole-brain analysis, to show the time course of low-frequency activity in this region and crucially to identify regions whose activity evolved through time from an initial representation of overall value to a later representation of value difference. Notably, we only performed statistical inference on tests orthogonal

to those originally used to identify the region of interest, namely, the main effect of task versus baseline. Frequency-domain analysis was performed as in the whole-brain analysis without subsampling the data to 20 Hz; time courses show the group z statistic of the averaged (normalized) β values from 2–4.5 Hz (for value difference) and 3–9 Hz (for overall value), based on predictions from the biophysical model.

For inference on the effects of overall value and value difference on region of interest data, we performed a cluster-based permutation test at the group level after collapsing across the relevant frequencies. We generated 5,000 randomly permuted t statistics for each time point (Fig. 5a,b and Supplementary Fig. 10). We then thresholded each permutation's t statistic time series at a threshold of $t_{29} > 2.1$ (equivalent to $P < 0.05$ uncorrected) and measured the maximum size of any cluster passing this threshold in the time series to build a null distribution of cluster sizes. We then compared the size of clusters from the true t statistic time series to those from the null distribution. We report clusters at a significance level of $P < 0.05$, corrected for multiple comparisons across time.

Computational model. We implemented a mean-field reduction of a previously described spiking neuronal network model²⁰ (full details of the reduction are given in refs. 23,24).

The reduced model consists of two units ($i = 1,2$), each selective for one option, with an excitatory recurrent coupling ($J_{A,ii}$) onto each unit and an effective inhibitory coupling to the other unit ($J_{A,ij}$). Each unit receives external input currents that are proportional to the value of its favored option, as well as noisy background inputs that resemble endogenous noise in the cortex. The firing rate in each population of selective neurons is a function of the total synaptic input to this pool

$$r_i = f(I_i) = \frac{aI_i - b}{1 - \exp(-d(aI_i - b))}$$

where a , b and d determined the input-output relationship for a neuronal population and were set to 270 Hz nA⁻¹, 108 Hz and 0.154 s, respectively.

The total synaptic currents to each pool of neurons represented by (in nA) $I_i = J_{A,ii}S_i - J_{A,ij}S_j + I_0 + J_{A,\text{ext}}(r_i + r_{\text{vis}}) + I_{\text{noise},i}$ where S_i is the NMDA synaptic gating variable related to neural pool i . I_0 represents the synaptic input current from external inputs to both pools and was fixed at 0.3297 nA, $I_{\text{noise},i}$ is white noise filtered by a synaptic time constant of 2 ms and an amplitude of 0.009 nA, $J_{A,\text{ext}}$ represents the strength of synaptic coupling constant from external sources and was set to 0.0011215 (nA Hz⁻¹) and r_{vis} represents input firing rates of neurons that respond to the presentation of the visual stimulus, fixed at 7.5 Hz. r_i represents the input firing rate proportional to the value of each option presented, given by the equation $r_i = r_{\text{dec}}(1 + k_{\text{dec}}sEV_i)$, where r_{dec} and k_{dec} are constants, and sEV_i is the subjective expected value on option i derived above from prospect theory. We set r_{dec} to be 10 Hz and k_{dec} to be 0.1125 for the simulations shown in Figure 1. For a typical subject in experiment 1 (prospect theory parameters $\alpha = 0.63$, $\gamma = 0.64$), r_i would therefore range between 10.63 Hz for the lowest value option on offer in the experiment to 14.03 Hz for the highest value option. Note that the values of r_{vis} and r_{dec} can be scaled, as their product with $J_{A,\text{ext}}$ determines the selective inputs to neuronal pools in the network. Finally, $J_{A,ii}$ was set to 0.3539 for the simulations shown in Figure 1, and varied between 0.3166 and 0.3725 for the cross-subject variation simulations (Supplementary Fig. 6). $J_{A,ij}$ was set at 0.0966.

S_i for populations $i = 1,2$ are dynamical variables representing the slow synaptic currents attributable to NMDA receptor activation, given by the equation

$$\frac{dS_i}{dt} = -\frac{S_i}{\tau_S} + (1 - S_i)\xi f(I_i)$$

where τ_S is the NMDA receptor decay time constant, set at 60 ms, and ξ is a parameter that relates the presynaptic input firing rate to the synaptic gating variable, set at 0.641. We used a total simulation period of 2,500 ms, with time step dt of 0.2 ms. Stimulus onset ($I_{\text{vis}} = 7.5$ nA) was from 500 ms, with reward-dependent inputs delivered from 600 ms ($I_{\text{opt}} = 10$ nA); both inputs were offset at 2,000 ms. The decision was made when the firing rate of one of the populations reached a threshold of 30 Hz.

When analyzing the model's behavior, we no longer investigated the firing of individual selective neuronal populations (as described in ref. 20), but instead the summed synaptic inputs to both populations within the network, $I_1 + I_2$. We chose this measure as MEG is more sensitive to the dipolar currents produced by postsynaptic potentials than the quadrupolar currents produced by action potentials²⁵, and because the lack of separation between the neuronal pools means their activity is likely to be mixed when viewed at the macroscopic spatial scale resolved by MEG. However, neuronal firing rates and synaptic input currents are highly correlated in the model, and similar results could be obtained using firing rates as the dependent variable.

For predictions relating to a single subject (Fig. 1), we simulated 6,480 trials generated from the same stimulus set as used in experiment 1, with $\alpha = 0.63$ and $\gamma = 0.64$. We plotted the activity of the model as a function of the overall value ($sEV_1 + sEV_2$; Fig. 1a) of the decision and as a function of the value difference ($sEV_{\text{chosen}} - sEV_{\text{unchosen}}$; Fig. 1b). We then treat the model outputs, m , in the same

way as the beamformed data at each location of interest (d_i). First, we performed a time-frequency decomposition of the data on each trial from 2–10 Hz using Morlet wavelets (Morlet factor 5). The decomposed data was then treated as the dependent variable as a function of overall value and value difference

$$m^{tr, f, t} = \beta_0^{f, t} + \beta_1^{f, t} \times OV^{tr} + \beta_2^{f, t} \times VD^{tr}$$

where $\beta_0^{f, t}$, $\beta_1^{f, t}$, $\beta_2^{f, t}$ and their associated variances, $\text{var}(\beta_0^{f, t})$, $\text{var}(\beta_1^{f, t})$ and $\text{var}(\beta_2^{f, t})$ are estimated using ordinary least-squares regression.

49. Hsu, M., Krajbich, I., Zhao, C. & Camerer, C.F. Neural response to reward anticipation under risk is nonlinear in probabilities. *J. Neurosci.* **29**, 2231–2237 (2009).
50. Berg, P. & Scherg, M. A multiple source approach to the correction of eye artifacts. *Electroencephalogr. Clin. Neurophysiol.* **90**, 229–241 (1994).