

Published in final edited form as:

Nat Neurosci. 2014 November ; 17(11): 1613–1622. doi:10.1038/nn.3836.

Hierarchical competitions subserving multi-attribute choice

Laurence T Hunt^{1,2,*}, Raymond J Dolan¹, and Timothy EJ Behrens^{1,3}

¹ Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London WC1N 3BG, United Kingdom

² Sobell Department of Motor Neuroscience, University College London, Queen Square House, London WC1N 3BG, United Kingdom

³ Centre for Functional MRI of the Brain, University of Oxford, Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, Oxford OX3 9DU, United Kingdom

Abstract

Valuation is a key tenet of decision neuroscience, where it is generally assumed that different attributes of competing options are assimilated into unitary values. Such values are central to current neural models of choice. By contrast, psychological studies emphasize complex interactions between choice and valuation. Principles of neuronal selection also suggest competitive inhibition may occur in early valuation stages, before option selection. Here, we show behavior in multi-attribute choice is best explained by a model involving competition at multiple levels of representation. This hierarchical model also explains neural signals in human brain regions previously linked to valuation, including striatum, parietal and prefrontal cortex, where activity represents competition within-attribute, competition between attributes, and option selection. This multi-layered inhibition framework challenges the assumption that option values are computed before choice. Instead our results indicate a canonical competition mechanism throughout all stages of a processing hierarchy, not simply at a final choice stage.

Introduction

When choosing between different options, it is often the case that one alternative is preferable on one set of attributes, but another is preferred on others. Such trade-offs are ubiquitous in decisions affecting consumers¹, foraging animals², social interactions³, and economic choice⁴. Whilst key aspects of multi-attribute choice behavior are well characterised, their neural basis has not yet been systematically examined. This question is of significance not only in relation to the ecological validity of such decisions, but also because of the constraints they place on the implementation of choice in neural circuits.

Users may view, print, copy, and download text and data-mine the content in such documents, for the purposes of academic research, subject always to the full Conditions of use:http://www.nature.com/authors/editorial_policies/license.html#terms

*to whom correspondence should be addressed laurence.hunt@ucl.ac.uk Phone: +44 (0)203 448 4372; Fax: +44 (0)20 7813 1420 .
Author Contributions

All of the authors contributed to study design and preparation of the manuscript. L.T.H. acquired the data, and built the model. L.T.H. and T.E.J.B. analyzed data.

When faced with such choices, human subjects dynamically construct their preferences online as opposed to merely revealing them¹. A common assumption in neurobiological studies of reward-guided choice is that preference construction depends first upon combining attributes within each option into an integrated value, followed by a process involving value comparison⁵⁻¹¹. This ‘integrate-then-compare’ strategy, a classic solution to the decision problem in the behavioral literature, holds the appeal of a normative approach to choice, albeit a computationally expensive one¹². Some expressions of choice behavior indicate that subjects do indeed integrate different features of an option to form a unitary value¹³.

However, anomalies in human decision making indicate this normative explanation cannot, by itself, fully account for subjects’ choice behavior. For instance, behavioral economic studies highlight preference reversals between two options when a third option is introduced. As this is critically dependent upon the degree of similarity between alternatives on specific attributes, this raises the likelihood of a within-attribute comparison process¹⁴. Studies of information gathering during multi-attribute choice, containing multiple options and multiple features, also suggest evidence is acquired within-attribute, at least initially¹⁵. Such behavior is explained by several alternative accounts, and these depend upon different combinations of within-attribute and within-option comparisons^{1,14,16-19}.

The neural substrate of these alternative decision strategies remains unexplored, but is of great interest given that it violates the most common assumption in neural studies of decision-making – that values are integrated and then compared, or even that values are computed prior to choice. One obstacle to exploring this conundrum has been a difficulty in designing tasks where there is transparent behavioral evidence for within-attribute comparison²⁰. In addition, there is a lack of a candidate *neural* mechanism by which this type of decision might be implemented.

In this paper, we provide evidence that preference construction in multi-attribute choice may occur via a *hierarchical competition* process¹⁷⁻¹⁹. Such a model argues that because competition via mutual inhibition is a canonical feature of local neural circuits²¹, it should take place at *all* levels of representation. Within this scheme comparison would still occur at the level of option values⁵⁻¹⁰, but it also has the additional features of comparison occurring at the level of the component attributes, as well as at the level of which attribute is most salient for guiding choice. To test this mechanism we combined behavioral analysis, functional imaging data and computational modelling in a novel multi-attribute choice paradigm. Subjects were explicitly instructed to equally weight two different attributes in guiding choice. Crucially, the task was designed such that within-attribute comparisons might emerge naturally, even in a relatively simple (and experimentally tractable) three-option, two-attribute decision. Using this task we found clear behavioral evidence for a within-attribute comparison strategy. Critically, on each trial one or other of the attributes was more salient (‘relevant’) for guiding behavior, allowing us to investigate how competition for attribute salience is implemented neurally. The key findings here, based on functional magnetic resonance imaging (fMRI) data, were that intraparietal sulcus (IPS) signalled a competition over which attribute was the most salient for the current decision, whereas portions of medial frontal cortex reflected an ‘integrated’ value signal. In keeping

with this functional architecture, IPS altered its functional connectivity with regions subserving lower level (within attribute) comparisons as a function of which attribute was currently most relevant for guiding behavior. We argue our results provide evidence for a canonical inhibitory competition mechanism that is general throughout all layers of a stimulus processing hierarchy, and not simply present at a final choice stage.

Results

A multi-attribute reward-guided decision task

We first introduce the multi-attribute choice task, which forms the basis for the modelling and behavioral results set out below. Subjects were trained on the relative likelihood of receiving a reward on a set of eight different images (not tied to any spatial location (Fig. 1a), hereafter referred to as ‘stimuli’) and a set of eight different buttonpresses (tied to spatial locations onscreen (Fig. 1b), hereafter referred to as ‘actions’). Subjects learned action-reward probabilities (p_A) and stimulus-reward probabilities (p_S) separately from one another, by performing pairwise choices between two randomly selected alternatives from each set. If they chose the better of the two options they received positive feedback (a smiley face), and if they chose the worse they received negative feedback (a sad face). Subjects received interleaved blocks of stimulus and action trials across a training session that lasted approximately 45 minutes, until they attained a preordained performance criterion.

Following training, subjects underwent fMRI scanning whilst performing a monetary reward task in which three stimuli were pseudorandomly paired with three actions (Fig. 1c). Subjects were instructed to apportion equal weight to each of the two sources of information so as to select the best option. Rewards were determined probabilistically, based upon p_O – a Bayesian combination of probabilistic information from the stimulus (p_S) and action (p_A) (Fig. 1d and methods). The logic behind this rule for combining the prelearnt probabilities is straightforward – for example, two cues that both predict reward with probability of 0.5 combine to produce a net reward probability of 0.5, as would one cue with 0.8 probability combined with another cue with 0.2 probability. It also meant that the information provided by each of the two attributes was *balanced*; both stimulus and action provided equally relevant probabilistic information about reward likelihood. We adopted such an approach to avoid problems arising from subjects being biased to use each attribute differentially. The effectiveness of our approach was revealed in behavioral evidence that subjects were not (on average) biased towards using either stimulus or action attribute (Supplementary Fig. 1).

Subjects successfully deployed information learnt in the initial phase of the experiment to guide their decisions inside the scanner. Based on the Bayesian integrated values (p_O), subjects chose the best option on 79.0 \pm 3.8 % (mean \pm s.d) of all trials. On more challenging ‘conflict’ trials, where the best stimulus favoured one option but the best action favoured a different option, this figure fell to 70.9 \pm 4.6 %. On 89.6 \pm 4.8 % of all these more challenging trials, subjects selected either the best p_S or the best p_A . However, there was a relatively high proportion of challenging trials (22.6 \pm 3.8 %) where subjects went with a high p_S or p_A when it was not the highest p_O . Such choices accounted for the majority (78.2 \pm 10.3%) of errors on challenging trials. Moreover, a substantial portion of these errors (42.2 \pm 15.6 %) were those where the best individual attribute on the chosen

option exceeded or was equal to the best attribute available on the best option. There were few trials (~5 per subject) where the highest pO was neither the best pS or the best pA.

A further simple measure of subjects' behavior is shown in Figure 1e. This focusses on the challenging 'conflict' trials, and plots the probability of using the stimulus attribute to guide behavior (i.e. choosing the best stimulus) as a function of the within-attribute differences on stimulus and action attributes. When the within-attribute difference is much greater on the stimulus attribute than the action attribute, subjects became far more likely to select based upon the largest stimulus, and vice versa. Additionally, because the point of subjective equivalence (p=0.5) occurs when within-attribute differences are equal, this further shows that subjects did not have any bias (on average) to using one attribute over the other.

Hierarchical competition for multi-attribute choice

How might subjects solve such a task? The classical idea is that once the two attributes are integrated to form a unitary value (for example, by calculating pO for each option), the option values will then be compared with one another. One common model for comparison assumes a 'softmax' choice rule, where the probability of choosing an option is a function of its value relative to other alternatives. To additionally capture subjects' reaction times, comparison may be modelled as a dynamic process in which evidence is accumulated through time. There are several closely interrelated dynamical models^{17-19,21,22}. One class assumes accumulators for each option that receive value-related inputs and compete via inhibition^{18,21}.

Importantly, softmax choice models and evidence accumulator models make firm predictions as to which factors should influence choices and reaction times in multi-alternative and multi-attribute decisions. Firstly, in the softmax choice rule, the relative frequency of choosing option 1 over option 2 is *independent* of any remaining values in the decision - that is, pO₃ does not affect the ratio of choosing between options 1 and 2:

$$\frac{p(C=1)}{p(C=2)} = \left(\frac{e^{p_{O1}}}{e^{p_{O1}} + e^{p_{O2}} + e^{p_{O3}}} \right) \div \left(\frac{e^{p_{O2}}}{e^{p_{O1}} + e^{p_{O2}} + e^{p_{O3}}} \right) = \frac{e^{p_{O1}}}{e^{p_{O2}}}$$

(The softmax temperature parameter is omitted for clarity). Secondly, predictions of reaction times and choices from accumulator models of decision making suggest they will principally be driven by variation in integrated values (pO) rather than underlying constituents of these values (pA and pS).

However, several known effects in multi-attribute and multi-alternative choice violate these predictions. Many of these violations show that the value of a third option *can* influence the relative frequency of choosing between two options, serving as a 'distractor' from discriminating between options one and two. One account of distractor effects appeals to a neural process of *normalisation* by the set of choice alternatives on offer, prior to comparison. This implies high-valued third items prove more distracting than low-valued third items and so impair discrimination, matching empirical observations in humans and macaque monkeys²³. Specific classes of distractor effects also occur in multi-attribute

choice, such as ‘compromise’, ‘similarity’ and ‘attraction’ effects¹⁷. These necessitate comparison at the level of underlying attributes and not solely at the level of integrated values and so have inspired models where evidence accumulation and competition occur at multiple hierarchical levels¹⁷⁻¹⁹.

We propose a hierarchical competition model for our task (Supplementary Fig. 2). This is motivated by previous models, together with the idea that specific attributes are more influential or *relevant* on certain trials for guiding behavior, as evidenced by behavioral observations made below. Evidence accumulation and comparison occurs at multiple levels in the model: within-attribute, on integrated values, and also *between* attributes (in a competition for ‘attribute relevance’). Normalisation is also included in the model, but occurs at the initial input level of individual attributes rather than on integrated values. We provide detailed mathematical description of the model in the methods section.

The most novel feature of the model is the competition for attribute relevance. Inputs to this competition, denoted by ** in Supplementary Figure 2, reflect the absolute difference between the best and second best option on each attribute, calculated instantaneously from currently accumulated evidence in each node. Attributes with a large within-attribute difference (where the best option is significantly better than the second best) dominate this competition and so become more relevant for guiding model choices. This is achieved via feedback connections (dashed lines) which reweight the importance of each attribute as it projects forward to an ‘integrated’ value comparison. The output of the integrated value comparator is used to determine the model’s choice and reaction time, once activity within this node reaches a decision threshold.

Importantly, there are several behavioral observations from our experiment predicted by this framework that appear inconsistent with alternative or previous models (see Supplementary Figs. 3-4). These are a specific *within-attribute* distractor effect on subjects’ choices, and a selective effect of value difference on the *relevant attribute* on subjects’ reaction times (Fig. 2). In the next section, we elucidate these, based on subjects’ actual behavior. The exact same analysis is applied to both subject behavior and that of the hierarchical model.

A within-attribute distractor effect in choice

We first assess the hierarchical model’s predictions of *distractor effects*, where the value of option three affects discriminations between options one and two. We analysed model choices using binomial logistic regression. In this analysis, the presence or absence of a stimulus/action was denoted by an indicator variable. This allows the regression model to independently estimate the influence of each stimulus and each action on the probability of choosing the associated option, and so is agnostic to the probabilities assigned to each stimulus/action by the experimenter. It controls for the other attribute presented on the same option, as well as the attributes of the other option. As expected, better stimuli and better actions tied to a particular option led to a clear increase in the logistic regression weight for that option in the model’s choices (Fig. 2a). This was also found to be true in subject behavior (Fig. 3a).

An important feature of this analysis was that it focussed on just a *pair* of the options, including one chosen and one unchosen option at a time. By repeating this analysis on a subset of the data where the *third* option, left out of the regression model, was high or low in value, we could examine its effects on the discriminability of these two options^{23,24}. ('Third option' here refers to any option which is unchosen; each trial is included twice in the analysis, with each unchosen option becoming the 'third' option once). Crucially, when the left-out third option had a *low* pO (red datapoints in Fig. 2a), discriminations between the values of the remaining two options appeared *better* than was the case when the third option had a *high* pO (blue datapoints in Fig. 2a). This was again found to be true in subject behavior (Fig. 3a), reflected by a significant difference in the slopes of the regression lines on both stimulus and action attributes (T-test on difference of slopes between lines; $T(18)=3.17$, $p=0.0053$ (stimulus), $T(18)=2.45$, $p=0.025$ (response)).

Notably, this effect replicates observations from a recent study of value-guided choice, and can be explained in the context of divisive normalisation²³. However, it is also seen in our hierarchical competition model, where divisive normalisation is applied within-attribute rather than on integrated values. Moreover, the observed response pattern violates the axiom of independence of irrelevant alternatives, and so could not be predicted by softmax or simple value accumulator models (Supplementary Fig. 3a). These data indicate that a comparison of different options is fundamentally dependent upon what other alternatives are available to the subject.

We next split the third option based not upon its integrated value (pO) but instead based upon its stimulus value (pS; Fig. 2b) or its action value (pA; Fig. 2c). In doing so, we found that the distractor effect in the model's choices was restricted to the same attribute used to perform the split. This 'within-attribute distractor effect' is a consequence of a normalisation applied at the level of within-attribute competition rather than at the level of integrated option values. By contrast, in alternative models where integrated values (pO) were divisively normalised prior to comparison, distractor effects were distributed equally across both attributes, irrespective of which attribute was used to perform the split (Supplementary Fig. 3c).

We therefore asked whether the distractor effect occurred selectively within-attribute in subject behavior, or was distributed equally across both attributes. We found the stimulus value of option three affected subjects' ability of discriminate between stimulus values of options one and two ($T(18)=2.51$, $p=0.021$), but not their action values ($T(18)=0.82$, $p=0.82$) (Fig. 3b). Likewise, the action value of option three affected action ($T(18)=3.18$, $p=0.0052$) but not stimulus discriminations ($T(18)=1.78$, $p=0.092$) (Fig. 3c). Collapsing across both stimulus and action analyses, this resulted in a significant interaction between distractor effect magnitudes on different attributes, and the attribute used to perform the split ($T(18)=3.54$, $p=0.0023$). Similar behavioral observations were also made without use of a regression analysis, by examining raw choice probabilities (Fig. 3d; Supplementary Fig. 5). In Figure 3d, we focus on trials where stimulus value difference on options 1 and 2 was approximately equal and opposite in sign to action value difference on options 1 and 2. The probability of choosing the option with higher stimulus value (and hence lower action value) decreased when option 3 had a high stimulus value or low action value, but increased when

option 3 had a high action value or low stimulus value (interaction in two-way ANOVA, [attribute * high/low]: $F(1,72) = 16.62, p=0.00017$).

The within-attribute distractor effect we observe cannot be explained by models in which only integrated values are compared to solve the task. If such a model were true, then distractor effects would be spread equally across all attributes, irrespective of which attribute was used to parse the third option. Instead, the evidence highlights a dependence on comparisons being made within-attribute, shown formally by the comparison of the hierarchical model to alternative models (Supplementary Figs. 3-4; Supplementary Note). We conclude that there is a contribution of within-attribute comparison in our multi-attribute choice task.

RT is more influenced by one attribute than another

We next sought to isolate a signature of hierarchical competition in both model and subject reaction times. To enable this, we applied multiple regression with the logarithm of reaction times (RTs) as the dependent variable and the value of different options as independent variables. In a first analysis of both model and subjects' RTs, we entered integrated values (pO) as explanatory variables. Unsurprisingly, RTs scaled with trial difficulty. When the chosen value was higher, choices were made more rapidly (model: $\beta=-0.81\pm 0.01$, mean \pm s.e. across 10 simulations; behavior: $\beta=-1.13\pm 0.12$ across subjects; $T(18)=-9.22, p=3.06\cdot 10^{-8}$), whilst choices were slower when the best unchosen value increased (model: $\beta=1.10\pm 0.01$; behavior: $\beta=0.31\pm 0.07$; $T(18)=4.47, p=2.96\cdot 10^{-4}$). Thus, reaction times decreased when $(pO_{\text{chosen}} - pO_{\text{best unchosen}})$ increased (model: $\beta=-1.91\pm 0.02$; behavior: $\beta=-1.46\pm 0.18$; $T(18)=-8.00, p=2.44\cdot 10^{-7}$). The influence of the worst unchosen value on RTs was also significant (model: $\beta=1.42\pm 0.03$; behavior: $\beta=0.22\pm 0.09$; $T(18)=2.47, p=0.024$).

The above analysis is predicated on the idea that there is integration across attributes prior to comparison. In our task, however, this appears inconsistent with the presence of a significant within-attribute distractor effect. Based on the hierarchical model, we hypothesised that on trials where attributes *conflict* – that is, on trials where pS values favor one alternative but pA values favour another – the two attributes might compete with one another when guiding the choice, with one attribute becoming more *relevant* for guiding behavior. While the relevant attribute would vary from trial to trial, we could nevertheless use the subject's choice to disclose it. In trials where $(pS_{\text{chosen}} > pS_{\text{best unchosen}}$ and $pA_{\text{chosen}} < pA_{\text{best unchosen}}$), we labelled stimulus as the 'relevant' attribute and action as 'irrelevant' (or strictly, 'less relevant'). By contrast, if $(pS_{\text{chosen}} < pS_{\text{best unchosen}}$ and $pA_{\text{chosen}} > pA_{\text{best unchosen}}$), we labelled action as 'relevant' and stimulus as 'irrelevant'. In the model, competition for relevance is realised by the node labelled 'attribute competition' in Supplementary Figure 2.

The crucial question in this context is whether reward probabilities on the 'relevant' attribute explain more variability in reaction time than on the 'irrelevant' attribute, or instead whether RTs are better explained by the 'integrated' value difference used in our initial analysis. We again applied multiple regression, but now all three forms of value difference competed for variance in explaining RTs (see Methods for details). In the model, we found that the 'relevant' attribute probability difference had a greater influence on reaction times than the

‘irrelevant’ probability difference (Fig. 2d, blue/green bars). Moreover, these two terms explained away the contribution of pO value difference to model reaction time (Fig. 2d, red bar). Importantly, these predictions were not made of simpler models of evidence accumulation (Supplementary Fig. 4).

We applied the same analysis to subjects’ RTs. Whereas the difference on reward probabilities on the relevant attribute influenced RTs ($\beta=1.46\pm 0.31$; $T(18)=-4.62, p=2.13\cdot 10^{-4}$), neither the difference in probabilities on the irrelevant attribute ($\beta=0.68\pm 0.37$; $T(18)=-1.48, p=0.08$) nor the Bayesian integrated value difference ($\beta=0.06\pm 0.35$; $T(18)=-0.19, p=0.84$) had a significant impact (Fig. 3e). Consequently, the difference on the ‘relevant’ attribute showed a significantly greater effect than the ‘irrelevant’ difference (paired $T(18)=-3.43, p=0.003$). To assess the robustness of this result, we applied Levene’s test and confirmed that there was no significant difference in variances between the two sets of parameter estimates ($F_{1,36}=0.12, p=0.74$). We also confirmed the effect remained (paired $T(18)=-2.45, p=0.02$) after normalising all the regressors to account for any difference in their magnitude. Additionally, these effects remained evident even when we considered alternative rules for constructing pO (mean, multiplication or sum – Supplementary Table 1).

One potential final concern is that hierarchical competition could be replaced by a simpler model that retained within-attribute comparison but randomly selected a given attribute to become relevant on each trial. We tested such a possibility by running a formal model comparison on prediction of subjects’ choice behavior, comparing this simpler ‘random attribute’ model with a hierarchical competition model (Supplementary Note). Model evidence favoured the hierarchical model in 16 out of 19 individual subjects, reflected across the population by a significantly lower Bayesian Information Criterion (BIC) in the hierarchical model compared to a ‘random attribute model’ ($T(19)=-3.24, p=0.0046$; mean \pm s.e.: BIC = 19.56 \pm 6.05).

The behavioral findings we highlight are inconsistent with the hypothesis that subjects equally weight both sources of information on each trial. They instead indicate that subjects rely more heavily on one source of evidence than another on each trial, which we here term the ‘relevant attribute’. This accords with our model, in which competitions occur hierarchically. In the next section, we consider the neural basis by which these competitions are realised.

Intraparietal sulcus and attribute relevance competition

We used fMRI data to investigate which brain regions subserved multi-attribute choice, tying our analysis to the modelling framework above. A unique feature of our model (absent from ‘integrate-then-compare’ frameworks) is that the *outputs* of within-attribute competitions themselves form the *inputs* of a competition for which attribute becomes ‘relevant’ in the current trial. By tying one attribute to stimuli, and another attribute to actions, we capitalised upon recent reported dissociations between neural structures encoding action and stimulus values in cortical¹² and subcortical^{25,26} structures. As the competition for relevance was resolving, this would *upregulate* the relevant within-attribute competition and downregulate

the irrelevant one. This might be reflected by changes in functional connectivity between regions implicated in attribute comparison and those involved in lower-level competition.

A region that represents the ‘attribute competition’ process should compare the value difference on the ‘relevant’ attribute with the same value difference on the ‘irrelevant’ attribute. We therefore hypothesised this region would be brought out using the contrast $(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}}) - (pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$, where pCh is the chosen probability on a specific attribute and $pBUnCh$ is the best unchosen probability on a specific attribute. This contrast reflects the difference in relative ‘attribute goodness’ between irrelevant and relevant attributes. Importantly, it controls for areas simply encoding ‘integrated value difference’, due to the subtraction of relevant values from irrelevant values. Analogous to recent findings concerning value difference coding^{8,10,27,28}, it assumes that a ‘relevance competition’ signal reflects the *difference of inputs* that determine which attribute was more relevant (namely, within-attribute differences).

We first tested whether this contrast emerged from the ‘attribute competition’ node in our hierarchical model, by simulating fMRI data directly from activity within this node (Supplementary Note). We used similar simplifying assumptions to those used in other studies, namely that activity would be greatest in the node whilst competition was still being resolved, and this competition would continue until a decision was reached^{8,29}. We convolved the node’s total activity on each trial with a hemodynamic response function, added observation noise, and then regressed the simulated data against pCh_{relevant} , $pBUnCh_{\text{relevant}}$, $pCh_{\text{irrelevant}}$ and $pBUnCh_{\text{irrelevant}}$. All four elements of our contrast were found to affect the simulated fMRI signal in the direction specified by our ‘attribute comparison’ contrast (Fig. 4). Hence $(pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$ was predicted to have a negative regression coefficient on fMRI data, but $(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}})$ was predicted to have a positive regression coefficient. This is a consequence of competition between attributes being greater on trials where $(pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$ is closer in value to $(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}})$. It is analogous to the hypothesis that regions implementing a choice comparison process via evidence accumulation may exhibit greater activity when choice alternatives are closer in value^{8,29}.

We therefore searched for brain regions that encoded the contrast $(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}}) - (pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$. We found bilateral activation in the intraparietal sulcus (IPS) correlated with this contrast, with the left IPS (peak $Z = 3.65$, MNI (46, -34, 42)) surviving whole-brain correction ($p=0.0023$, cluster-corrected), and the right IPS showing a large but sub-threshold response in a similar location (Fig. 5a; Supplementary Table 2; unthresholded Z-statistic maps for all contrasts are available for download at <http://neurovault.org>). This was the only region to survive this contrast with whole-brain correction. A wider network of regions (notably including bilateral superior frontal sulcus) was present at an uncorrected threshold (Supplementary Fig. 6). Next, we extracted the timecourse from IPS with a method that controls for circularity in statistical inference³⁰ by defining each subject’s region of interest for analysis from the remaining (n-1) subjects’ data³¹. From this timecourse, it was apparent that all the necessary components of relevance competition were present in IPS signal. On the relevant attribute, IPS BOLD fMRI signal correlated negatively with the value of the chosen option (Fig. 5b, dark blue), and positively

with the value of the best unchosen option (Fig. 5b, green). This was shown formally by a negative effect of $(pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$ ($T(18)=-2.83, p=0.011$). This is similar to signals previously observed in IPS in value-guided choice tasks, in which only one attribute or integrated value was considered^{27,32}. As predicted, however, on the irrelevant attribute this pattern was reversed: IPS signal correlated positively with the value of the chosen option (Fig. 5b, red), and negatively with the value of the best unchosen option (Fig. 5b, cyan). This was again shown formally by a positive effect of $(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}})$ ($T(18)=2.50, p=0.022$). These effects are based upon contrasts of parameter estimates from a general linear model (model 1, see methods) which partials out any covariance between irrelevant and relevant attributes. This signal profile would not be predicted for a region that is comparing integrated values, but is instead predicted for a region that selects which feature of a decision to attend.

In a further analysis, we also found a positive correlate of $\log(\text{reaction time})$ in the same region, as is also predicted by a dynamical model. Importantly, however, the effects shown in Fig. 5b survived the inclusion of $\log(\text{reaction time})$ as a coregressor (supplementary Fig. 7), as did predictions from the hierarchical model.

It is also notable the signed 'relative attribute difference' is highly correlated with the absolute (unsigned) difference in within-attribute differences, i.e. $8(pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}}) - (pCh_{\text{relevant}} - pBUnCh_{\text{relevant}})$ (r-values across subjects ranged from 0.91 to 0.99). This mirrors an observation made previously in the context of value comparison⁸. Disambiguating signed and unsigned signals is therefore not possible with the present study design.

Next, we reasoned that if IPS forms part of a hierarchical comparison process, then this should be reflected in its functional connectivity with other brain regions in a manner dependent upon which attribute was currently relevant. To test this hypothesis, we used a psychophysiological interaction analysis³³ that compared IPS functional connectivity with the rest of the brain as a function of whether the stimulus attribute was currently relevant, or whether the action attribute was currently relevant. On trials where stimulus was relevant, IPS exhibited greater functional connectivity to a bilateral sector of anterior orbitofrontal/lateral frontopolar cortex (Fig. 6a; right peak $Z=3.52$, MNI (36, 52, -10); left peak $Z=2.90$, MNI (-38, 52, 10)). This is a notable finding in that IPS has direct connections to OFC via the third branch of the superior longitudinal fasciculus³⁴, and lesions to OFC are known to impair stimulus- but not response-based value-guided choice¹². By contrast, on trials where action was relevant, IPS exhibited greater functional connectivity to bilateral putamen (Fig. 6b, left panel; left peak $Z=-3.46$, MNI (-18, 12, 2); right peak $Z=-2.69$, MNI (26, 10, 0)), a portion of striatum previously implicated in habitual action-guided choice^{25,26}. The finding that both sets of activations occur bilaterally considerably reduces their probability of being false positives.

Finally, at the time of feedback the ventral striatum encoded a classic reward prediction error, signalling reward outcomes positively and predictions of reward negatively (on both relevant and irrelevant attributes; supplementary Fig. 8a/b). By contrast, the signal encoded in the IPS can also be interpreted as a prediction error, but acting on relevance rather than

value or reward (supplementary Fig. 8c/d). It encoded reward *negatively*, but the prediction of reward *positively* on the relevant attribute and the prediction of reward *negatively* on the irrelevant attribute. This signal can be construed as a prediction error on *attribute relevance*. If reward is greater than expected, more attention would be given to the relevant attribute on future trials, and less attention to the irrelevant attribute. It should be noted that this IPS relevance prediction error is inverted in sign, in the sense that a more positive attribute prediction error elicits a greater deactivation at the time of outcome.

In summary, at the time of the decision, IPS carried signals that implicate it in selecting which attribute is relevant, and changed its functional connectivity in accordance with an upregulation of lower-level competition. At the time of outcome, IPS showed an (inverted) prediction error signal suggestive of a role in relevance learning. These combined data provide evidence for a role for the IPS in selecting which attribute to attend to on any given trial.

Medial prefrontal cortex and integrated value comparison

If the IPS carries a signal that reflects a competition for attribute ‘relevance’ then do any other brain regions reflect a more traditional ‘integrated’ value competition? This can be tested by looking at a different contrast: $(pCh_{\text{relevant}} - pBUnCh_{\text{relevant}}) + (pCh_{\text{irrelevant}} - pBUnCh_{\text{irrelevant}})$. This contrast correlated negatively with BOLD fMRI signal in the dorsal medial frontal cortex (dmFC), in the vicinity of preSMA/paracingulate sulcus (Fig. 7a; Supplementary Table 3; peak $Z = -3.85$, MNI (0, 38, 42)). This is in close proximity to a locus previously implicated in a comparison of integrated values during action selection^{8,35}. Notably, this region is frequently co-activated with the IPS – both IPS and dorsal MFC are recruited more strongly by trials in which values are close together, and so correlate negatively with the value of the chosen option minus the unchosen option^{8,29}. It is therefore particularly striking in our study that whereas dmFC and IPS encode variables in the same direction on the relevant attribute, they do so in the *opposite* direction on the irrelevant attribute (compare Figs. 5b and 7b).

Finally, we asked what signals are encoded in ventromedial prefrontal cortex (VMPFC), a region commonly recruited in studies of value-guided choice^{36,37}, but in which different studies have observed different task variables^{8,27}. In the present study, we found that VMPFC encoded a chosen value signal, selectively for the relevant but not irrelevant attribute (Supplementary Fig. 9).

Discussion

It is commonly assumed that a key step in how the brain supports value-guided choice is through an integration of the different features of an option into a single value, so that values can be compared with one another in a common currency. Such a decision schema has obvious appeal, not least because it can support comparison of incommensurable items¹¹, but it cannot explain several key behavioral observations. Among these are the effects of within-attribute similarity on choice, suggesting some forms of comparison are enacted before integration, as well as the effects of relevance of a particular attribute modulating

reaction times. Subject behavior in our multi-attribute decision paradigm showed both these effects, most strikingly via a novel within-attribute distractor effect.

To explain our data we adopted a fundamentally different approach to value comparison, suggested previously in the behavioral literature¹⁶⁻¹⁹ but notably absent from many neuroscientific studies. In this scheme, competition (via mutual inhibition^{10,21}) occurs not just at the 'higher level' of integrated values, but also at lower levels (within attribute) and between levels (attention towards attributes). This was necessary to capture all key features of our behavioral data, which were not captured by more simple models of evidence accumulation. Apart from successfully explaining behavioral data, our model gained extra validity by explaining task-related neural activity as well as functional interactions between brain regions involved in supporting value comparison, indexed via fMRI.

Hierarchical competition invokes the notion of a canonical computation performed across multiple brain regions, and it is implicit that competition in different brain regions takes place in different *frames of reference*. Such hierarchies mirror those proposed in the domain of rule-based action selection, or tasks requiring cognitive control³⁸. In our task, the IPS carries a signal reflecting competition in an attribute frame of reference, whereas medial prefrontal structures carry signals reflecting competition in a frame of reference of options. Such a clear dissociation between computations in IPS and dmPFC has not been described previously, and it is notable that these areas have typically coactivated in studies of decision making^{8,29}, particularly where there is increased choice difficulty (decreased value difference). The portion of IPS identified in our study is ideally placed, in terms of anatomical connectivity to prefrontal cortex^{34,39} and its established role in attentional reorienting⁴⁰, to compute which attribute to attend on a particular trial. Indeed, in number comparison tasks where the relevant and irrelevant dimensions are explicitly cued rather than internally generated, IPS only reflects value differences within the relevant, not the irrelevant, competition⁴¹. When information is presented from two different sources whose distributions can be formally specified, fMRI signal in IPS reflects the Kullback-Leibler divergence (or degree of competition) between these two sources of information⁴².

However, we do not argue that IPS performs a general role of arbitrating between attributes at the top level of hierarchical competition. In our experiment subjects had to trade a competition that occurred in stimulus space in the OFC against a competition that occurred in action space in the motor system. IPS is ideally placed to resolve such a competition due to its monosynaptic projections with both structures. It is clear that unlike OFC or vmPFC, IPS has access to values selective for spatial locations, for specific actions and specific stimuli. Notably, in circumstances where both attributes can be represented in stimulus or goal space, deficits in attribute comparison can be induced by lesions to vmPFC⁴³. We therefore argue that attribute comparison is not a unique process with a single cortical focus, but a striking example of a general rule of competition in cortical processes. The critical neural substrate will depend on the relevant features of the decision at hand.

Similarly, the frames of reference and functional roles of other brain regions contributing to task performance reflect both their anatomical connectivity and functional specialisation. For example, lateral OFC is likely to play a role in competition between stimuli given that it

receives a highly processed sensory input⁴⁴ and is critical for tasks involving stimulus but not action comparison¹². In keeping with this formulation we observed increased functional connectivity of OFC with IPS on trials where stimulus probabilities became relevant. By contrast a dorsolateral portion of the striatum is likely to play a role in competition between actions, given evidence that it possesses anatomical connectivity with motoric structures⁴⁵ and is implicated in habitual action selection^{25,26}. This region showed increased functional connectivity to IPS on trials where action probabilities became relevant. Finally, both dorsal and ventral medial frontal cortex are implicated in competition between values^{8,10,11,27,28,35}. In the present study, integrated value comparison signals were particularly prominent in dMFC at the time of choice and like the IPS this region has monosynaptic connections to both OFC and motoric structures.

Recent studies have examined distractor effects in value-guided choice tasks with multiple options, but from a different perspective to that examined here. Louie *et al.* isolated distractor effects in subjects who chose between rewards of different value. As in our study, high value distractors impaired discrimination between two alternative options²³. They proposed a divisive normalisation of values during choice. In our model, divisive normalisation captured distractor effects (Supplementary Figs. 3-4), but because of the novel within-attribute distractor effect, we place it at the level of attributes instead of the level of integrated values. Indeed, normalisation may also reflect a canonical mechanism operating at all levels of a processing hierarchy⁴⁶. Future models may unify the processes of competition and normalisation in a single framework. Chau *et al.*, in a task where decisions are made under time pressure, reported a distractor effect that operates in the opposite direction, where high-valued third options are less distracting than low-valued third options. In their task distractors are made transiently available and then removed from the decision⁴⁷. Such an effect can be explained by the increased pulse of inhibition introduced into a competitive decision-making network by high-value distractors, resulting in slowing of a decision that renders it more accurate. We find a similar (albeit weak) effect in our LCA model when divisive normalisation is switched off, and competition between options is high (Supplementary Fig. 3b). The conditions under which normalisation occurs during competition remain debated⁴⁸, but one hypothesis is that divisive normalisation occurs over a different time-course than competition via mutual inhibition, and may not occur where one option is removed shortly after decision onset. Finally, the within-attribute distractor effect may mirror the well-known ‘similarity’ effect in multi-attribute choice, where introducing a new option reduces the probability of choosing options that are similar on multiple attributes to those that are dissimilar⁴⁹. In our model, similar options (with small within-attribute differences) will be downweighted by the attention competition relative to dissimilar options (with large within-attribute differences). Future work could directly test the ability of the model to reproduce this effect, alongside other known effects in multi-attribute choice such as the ‘attraction’ and ‘compromise’ effects¹⁷.

In summary, we propose that instead of competition being an isolated ‘component process’ in value-guided decision making, it occurs at multiple, distributed levels of representation. This hierarchical competition account of decision making reconciles conflicting behavioral observations which suggest that comparison occurs within-attribute, as well as on integrated

values. Our account also explains observations where competition is observed in different frames of reference within different brain structures, as well as how interactions between brain regions are modulated during a task⁵⁰. More broadly the findings imply that competition via mutual inhibition represents a canonical computation⁴⁶ that subserves distributed decision-making processes throughout the brain, rather than a process occurring solely at a single comparator node.

Online Methods

Experimental task

24 human volunteers (age range 20-49, 13F, 11M, recruited from a subject pool at University College London) participated in the experiment. Inclusion criteria were based on age (minimum 18 years, maximum 50 years), and screening for history of neurological or psychiatric illness. The sample size was based on similar sample sizes in recent fMRI studies of decision-making and methodological recommendations in the literature. 2 subjects failed to reach the requisite learning criterion during training ; 3 further subjects showed excessive head motion during fMRI (based on the motion correction report in FSL version 6.00). These subjects were excluded, leaving 19 subjects in all subsequent behavioral and neural analysis. The experimental protocol was approved by the UCL local research ethics committee, and informed consent was obtained from all subjects included.

Training—Subjects learnt that certain stimuli and certain responses were more predictive of reward by performing pairwise choices within-attribute (i.e. stimulus vs. stimulus, action vs. action). In each ‘action training’ trial, two locations from eight were highlighted onscreen and the subject aimed to select the better of the two actions to receive positive feedback (smiley face). Each spatial location was tied to a button press with a specific forefinger on a keyboard. In each ‘stimulus training’ trial, two stimuli from eight were presented in random spatial locations (either side of a fixation point), and subjects aimed to select the better of the two stimuli. Subjects performed four training blocks outside the scanner (two stimulus, two action, each consisting of a minimum of 70 trials and terminating once subjects had reached a performance criterion of 90% correct responses). Stimuli and actions were counterbalanced across subjects.

During training, feedback was delivered deterministically, according to whether the subjects chose the better of the two options. Deterministic feedback ensured subjects learnt equally about both the selected and unselected options throughout training (which would not be the case if reward was delivered probabilistically on the selected option). The rank of the best through to the worst options was translated into a probability of receiving reward (termed pS for stimuli, pA for responses) in the main experiment – scaled from worst to best as (0.1,0.2,0.3,0.4,0.6,0.7,0.8,0.9). These probabilities are almost perfectly collinear ($r=0.99$) with the experienced frequency of reward in the training phase (supplementary Fig. 10).

Main experimental task—Inside the scanner, subjects performed 200 trials of a three-option choice task, in which stimuli were pseudorandomly paired with different actions. Subjects were instructed to weight both stimulus and action attributes equally in making their choice. It was never the case that the same stimulus (or action) was available on more

than one option. Trials were selected such that key variables of interest for neural (and behavioral) analysis were decorrelated (as far as possible) in the design. In this schedule, as expected by chance, the best stimulus and best action suggested conflicting responses on 66% of trials. Within these trials, only infrequently (5% of trials) did the best pO occur for options with neither the best pS or pA; instead, the best pO typically aligned with the best pS or best pA (59% and 36% of trials, respectively). Across all trials, this meant that the best pO aligned with the best pS on 73% of trials, and with the best pA on 58% of trials. Despite this slight asymmetry between best pA/pS being coincident with best pO, subjects showed no systematic bias towards using one attribute over another (supplementary Fig. 1). The same schedule was used for each subject, but with high/low value stimulus identities and high/low value actions counterbalanced across subjects. Full details of schedules used can be obtained by examining the raw datafiles and analysis scripts (provided as supplementary material). Task timings are shown in Fig. 1c. Subjects could view freely (i.e. they were not required to hold fixation) during the decision phase.

True reward probabilities for each option were based on a posterior probability of reward (pO) that optimally weighted pS and pA for each alternative, using Bayes' Rule:

$$pO_i = \frac{pS_i * pA_i}{pS_i * pA_i + (1 - pS_i) * (1 - pA_i)}$$

Behavioral data analysis

Trials on which no response was made within the 3s time limit (on average <3% of total trials) were removed from subsequent behavioral analysis. The same analysis was applied to both predictions from the computational model of the task (Fig. 2) and subject behavior (Fig. 3).

Preference for using stimulus or action attribute (supplementary Fig. 1)—To estimate the average weight assigned to using the stimulus or action attribute, we fit a behavioral model to subjects' choices that contained a free parameter, Y , which was greater in subjects who used stimulus information more heavily, and a free choice parameter β . Y serves to transform the true stimulus and action probabilities of option i into subjective 'weights' via the following formulae:

$$\begin{aligned} \text{Stimuli: } p\hat{S}_i &= \frac{1}{1 + e^{-\gamma(pS_i - 0.5)}} \\ \text{Actions: } p\hat{A}_i &= \frac{1}{1 + e^{-\frac{1}{\gamma}(pA_i - 0.5)}} \\ \text{Integrated: } p\hat{O}_i &= \frac{p\hat{S}_i * p\hat{A}_i}{p\hat{S}_i * p\hat{A}_i + (1 - p\hat{S}_i) * (1 - p\hat{A}_i)} \\ \text{Choice probabilities: } p(Ch=i) &= \frac{e^{-\beta p \hat{O}_i}}{\sum_o e^{-\beta p \hat{O}_o}} \end{aligned}$$

Each individual's Y and β parameter was fit via maximum likelihood estimation.

Logistic regression of choices, revealing distractor effects (main Fig. 2(a)-(c)/ 3(a)-(c))—To estimate the influence of each stimulus and action on subjects' choices, we

performed a binomial logistic regression analysis. To perform this analysis in a manner amenable to revealing a third-option distractor effect, we performed the logistic regression on *pairs* of options taken from the three-option choice data. Within the logistic regression model, each trial featured twice: once with the chosen option and unchosen option A included, and again with the chosen option and unchosen option B included. As shown before^{23,24}, this allows interrogation of the regression coefficients as a function of when the option *left out* of the regression was of high or low value.

The subject's choice (i.e. chose option 2 coded as 1, chose option 1 coded as 0) was the dependent variable, and as independent variables we included separate indicator variables for each stimulus and each action for option 1 and 2 (valued 1 when that stimulus/action was present for option 1 or 2, and 0 otherwise). However, such an approach leads to rank deficiency in the design matrix, as the linear combination of each stimulus/action sums to produce a constant term. To finesse this, we removed the regressor corresponding to the eighth (best) stimulus and action for both options 1 and 2 (that is, we removed columns 8, 16, 24 and 32 from the design matrix), and added a single constant term. We then constructed contrasts of parameter estimates that recovered the mean parameter estimate of each stimulus/response, collapsed across options 1 and 2 (Supplementary Fig. 11).

This was repeated on subsets of datapoints: where the third (left-out) option was either high (>67th percentile) or low (<33rd percentile) in integrated value (main Fig. 2a/3a), high or low in stimulus probability (main Fig. 2b/3b), or high or low in action probability (main Fig. 2c/3c). The datapoints in these figures show the mean \pm s.e. (across subjects) of the contrasts of parameter estimates from the logistic regression.

Reaction time regression (main Fig. 2d/3e)—We assessed the influence of reward probabilities on reaction times using multiple linear regression. $\text{Log}(\text{reaction time})$ was used as the dependent variable as this approximates a normal distribution. In a first analysis, we only investigated the role of integrated values on reaction time. 4 regressors were included – (i) the integrated value of the chosen option (pO_{chosen}), (ii) the integrated value of the best unchosen option ($pO_{\text{best unchosen}}$), (iii) the integrated value of the worst unchosen option ($pO_{\text{worst unchosen}}$), and (iv) a constant term.

In a second analysis, we investigated the separable contributions of value difference on the 'relevant' attribute, the 'irrelevant' attribute, and the integrated value, on trials in which 'relevance' could be defined (where stimulus and action conflicted, and subjects either selected the best stimulus or the best response). 10 regressors were included in total; (i) a constant term, (ii)-(iv) the chosen, best unchosen and worst unchosen values on the relevant attribute; (v)-(vii) the chosen, best unchosen and worst unchosen values on the irrelevant attribute; (viii)-(x) the chosen, best unchosen and worst unchosen integrated values. Contrasts of parameter estimates were used to calculate the influence of the value difference on chosen ad best unchosen values on reaction time (i.e. [1 –1] contrast on chosen and best unchosen values).

Complete anonymised behavioral datasets and MATLAB analysis scripts are provided as supplementary material.

Computational model

We implemented a hierarchical leaky, competing accumulator model to capture behavioral and neural effects observed in our multi-attribute choice task. In the main model (results presented in main Figs. 2 and 4), we assumed that evidence accumulated and competed on each attribute (within-attribute competition), there was competition over which attribute to attend (between-attribute competition), and there was competition on integrated values to select the winning option.

Within-attribute evidence accumulation was modelled using the following dynamical equation:

$$\mathbf{E}^l(t+dt) = \mathbf{E}^l(t) + \mathbf{S}^l \mathbf{E}^l(t) dt + \mathbf{C} \mathbf{M} dt + \mathbf{C} \mathbf{N}(0, \sigma^2 dt); \mathbf{E}^l(0) = \mathbf{0} \quad (\text{eq. 1})$$

where: $\mathbf{E}^l(t)$ is a 3×2 matrix (superscripted with l to denote lower-level of hierarchical model), containing accumulated evidence on each option (row) and attribute (column) at time t ;

\mathbf{M} is a 3×2 matrix, containing the within-attribute probabilities, *divisively normalised* within-attribute (that is, the columns of \mathbf{M} sum to 1);

\mathbf{S}^l is a 3×3 matrix, containing off-diagonal elements that determine competition (k_l), and on-diagonal elements that determine the rate of self-decay (d_l);

\mathbf{C} is a 3×3 contrast matrix that subtracts the mean of all other alternatives from the firing rate of each element (on-diagonal elements set to 1, off-diagonal elements set to -0.5);

$\mathbf{N}(0, \sigma^2 dt)$ is a 3×2 matrix containing Gaussian noise, with each element drawn from the normal distribution with mean 0 and variance $\sigma^2 dt$;

$\mathbf{0}$ is a 3×2 matrix of zeros.

Accumulation of evidence *integrated across attributes* was then modelled using the following equation:

$$\mathbf{E}^h(t+dt) = \mathbf{E}^h(t) + \mathbf{S}^h \mathbf{E}^h(t) dt + \mathbf{E}^l(t) \mathbf{w}^{\rightarrow}(t) dt; \mathbf{E}^h(0) = \mathbf{0} \quad (\text{eq. 2})$$

where: $\mathbf{E}^h(t)$ is a 3×1 matrix (superscripted with h to denote higher-level of hierarchical model), containing accumulated evidence on each option at time t ;

\mathbf{S}^h is a 3×3 matrix, containing off-diagonal elements that determine competition (k_h), and on-diagonal elements that determine the rate of self-decay (d_h);

$\mathbf{w}(t)$ is a 2×1 vector, whose elements sum to 1 and are ≥ 0 , that determines the relative weight assigned to each lower level attribute;

$\mathbf{0}$ is a 3×1 matrix of zeros.

Finally, the weight assigned to each attribute is assumed to be calculated instantaneously by applying a softmax transformation to the within-attribute difference of the best and 2nd best option on each attribute (as suggested by the signal observed in the intraparietal sulcus at the time of making the decision):

$$w_i(t) = \frac{1}{1 + e^{\beta(\text{Ediff}_i - \text{Ediff}_{\sim i})}} \quad (\text{eq. 3})$$

Ediff_i refers to the difference between the highest value in $\mathbf{E}^l(t)$ and next highest value in $\mathbf{E}^l(t)$ for the i -th attribute. β is a free parameter that determines the influence of within-attribute differences on attention. When $w_i(t)$ is close to 0.5, both stimulus and action have an influence on guiding choice; when it is close to 1, then $w_{\sim i}(t)$ will be close to 0, and only the i -th attribute will have influence on choice.

A decision was made when any value of \mathbf{E}^h exceeded a decision threshold ϑ . The reaction time of the model was the time at which this occurred plus a non-decision time, t_{nd} .

The model has eight free parameters. We assumed that the dynamics of the leaky, competing accumulators were the same both within- and between- attributes (that is, $k_h = k_l$ and $d_h = d_l$), reducing the parameter space to six parameters. We performed a grid search across parameter space to fit basic properties of subject behavior – namely, the distribution of reaction times and the error rates of an average subject. The fit parameters are listed below. However, it is important to note that the key behaviors of the model – that is, the within-option distractor effect and the effect on reaction times – are *qualitative* rather than quantitative features of the model. These properties did not change fundamentally as a function of the parameter fit.

Model parameters were as follows (search grid resolution and limits specified in square brackets):

Decay rate, $d_h = d_l = -0.1$ [-0.5:0.1:0]

Inhibitory competition, $k_h = k_l = -0.05$ [-0.08:0.01:0]

Non-decision time, $t_{nd} = 300\text{ms}$ [100:100:500]

Threshold, $\vartheta = 200$ [100:100:1100]

Input noise, $\sigma = 1$ [0:1:10]

Between-attribute competition softmax temperature, $\beta = 0.1$. [0.1:10, logarithmically spaced in 10 bins]

All modelling was implemented in MATLAB (The Mathworks, Natick, MA); MATLAB code for models is available on request. For comparison with alternative models, and details of fMRI simulations, see Supplementary Figures 3-4 and Supplementary Note.

fMRI data acquisition

FMRI data—Whole-brain T2*-weighted echo-planar imaging (EPI) data were acquired using a Siemens Trio 3T scanner, using a 32-channel headcoil. The sequence chosen was selected to minimise dropout in both orbitofrontal cortex and amygdala⁵². Each volume contained 43 slices of 3mm isotropic data; echo time = 30ms, repetition time = 3.01s per volume, echo spacing of 0.5ms, slice tilt of -30° (T>C), Z-shim of -1.4 mT/m*ms, phase oversampling of 13%, ascending slice acquisition order. The mean number of volumes acquired per subject was 815 (the total number of volumes acquired varied depending upon participants' reaction times). The first 5 volumes of each dataset were discarded to account for T1 saturation effects, and so the experiment was not started until these 5 volumes had been acquired.

Structural data—Whole-brain T1-weighted structural data were acquired using a 3D MDEFT routine with sagittal partition direction, 176 partitions, field of view=256×240, matrix=256×240, 1mm isotropic resolution, TE=2.48ms, TR=7.92ms, flip angle 16° , inversion time=910ms. Total acquisition time was 12 minutes 51 seconds.

Field maps—Whole-brain field maps (3mm isotropic) were acquired to allow for subsequent correction in geometric distortions in EPI data at high field strength. Acquisition parameters were 10ms/12.46ms echo times (short/long respectively), 37ms total EPI readout time, with positive/up phase encode direction and phase-encode blip polarity -1 .

fMRI analysis

FMRI data processing was carried out using FEAT (FMRI Expert Analysis Tool) Version 6.00, part of FSL (FMRIB's Software Library, www.fmrib.ox.ac.uk/fsl)⁵³. Region-of-interest timeseries analysis was performed using custom-written scripts in MATLAB (The Mathworks, Natick, MA).

Preprocessing—The following pre-statistics processing was applied in FSL; motion correction using MCFLIRT; unwarping of B0 slice-timing correction using Fourier-space time-series phase-shifting; non-brain removal using BET; spatial smoothing using a Gaussian kernel of FWHM 7.0mm; highpass temporal filtering (Gaussian-weighted least-squares straight line fitting, with $\sigma=50.0$ s).

First-level general linear model—Time-series statistical analysis was carried out using FILM (FMRIB's Improved Linear Model) with local autocorrelation correction. Two event-related models were specified – model 1 with solely task-based regressors (results shown in main Figs. 5/7, Supplementary Figs. 6-9, Supplementary Tables 2 and 3), and model 2 to specify the psychophysiological interaction (shown in main Fig. 6).

Model 1 contained 14 regressors in total. The first eight were related to the decision phase of the trial; the next two were related to the response; the final four were related to feedback. The regressors were:

- a) DECIDE_ONSET: 1 during decision (from decision until response), 0 outside decision.

- b)** DECIDE_RELATT_CHPROB: Same timings as DECIDE_ONSET; parametric regressor for probability of the chosen option being rewarded on relevant attribute (chosen pS for ‘stimulus relevant’ trials, chosen pA for ‘action relevant’ trials). Only trials in which stimulus and action attributes had different ‘best’ options (and where one of these was selected) were included in this regressor.
- c)** DECIDE_RELATT_BEST_UNCHPROB: As (b), containing probability of best unchosen option on the relevant attribute.
- d)** DECIDE_RELATT_WORST_UNCHPROB: As (b), containing probability of worst unchosen option on the relevant attribute.
- e)** DECIDE_IRRELATT_CHOSEN_PROB: As (b), containing probability of chosen option on irrelevant attribute.
- f)** DECIDE_IRRELATT_BEST_UNCHPROB: As (b), containing probability of best unchosen option on irrelevant attribute.
- g)** DECIDE_IRRELATT_WORST_UNCHPROB: As (b), containing probability of worst unchosen option on irrelevant attribute.
- h)** DECIDE_EASY: Same timings as DECIDE_ONSET; value 1 on ‘easy’ decisions in which both stimulus and action indicate the same best option, 0 on other trials.
- i)** RESPONSE_ONSET: Value 1 during the 4-8s period when the response was onscreen, 0 outside this period.
- j)** RESPONSE_RIGHT>LEFT: Value 1 when a button was pressed with the right hand; -1 with the left hand; 0 outside this period. Stick function lasting 1s, time-locked to response.
- k)** FEEDBACK_ONSET: Value 1 during 3s period when reward feedback was onscreen, 0 outside this period.
- l)** FEEDBACK_REWARD: Same timings as FEEDBACK_ONSET, value 1 if reward was delivered and 0 otherwise.
- m)** FEEDBACK_RELATT_CHPROB: Same timings as FEEDBACK_ONSET; parametric regressor containing probability of chosen option on relevant attribute. Only included for same trials as DECIDED_RELATT_CHPROB.
- n)** FEEDBACK_IRRELATT_CHPROB: As (m), but containing probability of chosen option on irrelevant attribute.

All regressors were convolved with FSL’s canonical gamma hemodynamic response function, and temporally filtered with the same high-pass filter applied to the FMRI timeseries. Temporal derivatives of all regressors were included to account for variability in the hemodynamic response function.

The following contrasts of parameter estimates were constructed, and reported in the main text:

- (i) Regions encoding value difference *differentially* across relevant and irrelevant attributes (main Fig. 5): $(\text{DECIDE_IRRELATT_CHPROB} - \text{DECIDE_IRRELATT_BEST_UNCHPROB}) - (\text{DECIDE_RELATT_CHPROB} - \text{DECIDE_RELATT_BEST_UNCHPROB})$
- (ii) Regions encoding value difference *commonly* across relevant and irrelevant attributes (main Fig. 7): $(\text{DECIDE_RELATT_CHPROB} - \text{DECIDE_RELATT_BEST_UNCHPROB}) + (\text{DECIDE_IRRELATT_CHPROB} - \text{DECIDE_IRRELATT_BEST_UNCHPROB})$
- (iii) Regions encoding a traditional reward prediction error, tied to the relevant attribute (supplementary Fig. 8a): $(\text{FEEDBACK_REWARD} - \text{FEEDBACK_RELATT_CHPROB})$
- (iv) Regions encoding an ‘attribute prediction error’ (supplementary Fig. 8c): $(\text{FEEDBACK_RELATT_CHPROB} - \text{FEEDBACK_IRRELATT_CHPROB})$

Model 2 contained 17 regressors in total. The first two regressors comprised the ‘psychological’ component of the psychophysiological interaction; the next regressor comprised the ‘physiological’ component; the next two regressors made up the key ‘interaction’ contrast. The remaining 12 regressors modelled other elements of the task (regressors of no interest).

- a) DECIDE_WENT_W_STIM: 1 during decision period (i.e. lasting from decision onset until response), 0 outside decision period, on trials where stimulus was relevant attribute.
- b) DECIDE_WENT_W_RESP: 1 during decision period (i.e. lasting from decision onset until response), 0 outside decision period, on trials where action was relevant attribute.
- c) IPS_TIMECOURSE: The mean timeseries (after preprocessing) extracted from a region of interest based on analysis from model 1. The mask used is described below.
- d) IPS_WENT_W_STIM_INTERACTION: Interaction of DECIDE_WENT_W_STIM and IPS_TIMECOURSE; following the procedure outlined in ³³, DECIDE_WENT_W_STIM had zero centre and IPS_TIMECOURSE had zero mean.
- e) IPS_WENT_W_RESP_INTERACTION: The interaction of DECIDE_WENT_W_RESP and IPS_TIMECOURSE; DECIDE_WENT_W_RESP had zero centre and IPS_TIMECOURSE had zero mean.

The remaining twelve regressors (of no interest) were: RESPONSE_ONSET, RESPONSE_RIGHT>LEFT, FEEDBACK_ONSET, FEEDBACK_REWARD, DECIDE_NOBRAINER, DECIDE_WENT_W_NEITHER (as regressor (a), but on trials where neither the best stimulus nor best action was chosen), and six parametric regressors for the values of the best, 2nd best and worst response, and best, 2nd best and worst stimulus, all at the decision phase. The key contrast of parameter estimates (main Fig. 6) was

therefore: (IPS_WENT_W_STIM_INTERACTION - IPS_WENT_W_RESP_INTERACTION). A region significantly *greater* than zero for this contrast was showing greater functional connectivity with IPS on trials where *stimulus* was relevant, whereas a region significantly *less* than zero for this contrast was showing greater functional connectivity with IPS on trials where *action* was relevant.

Intersubject registration—Registration to high resolution structural images was carried out using FLIRT (FMRIB’s Linear Image Registration Tool). Registration from high resolution structural to standard space was then further refined using FNIRT nonlinear registration.

Second-level general linear model and statistical inference—For both models 1 and 2, higher-level analysis was carried out using FLAME (FMRIB’s Local Analysis of Mixed Effects) stage 1. A group mean was fit to contrasts of parameter estimates from the first-level analysis, and T-statistics were estimated at each voxel to test whether this mean was significantly different from zero. For model 1 (main task GLM), Z (Gaussianised T)-statistic images were thresholded using clusters determined by $Z > 2.3$ and a (whole-brain corrected, family wise error) cluster significance threshold of $P = 0.05$. (This was except for prediction error signals at feedback time, where the clusters extent determined by a threshold of $Z > 2.3$ were too large to make sensible inference. Hence, a more stringent cluster-forming threshold of $Z > 3.1$ was used). For model 2 (PPI GLM), Z (Gaussianised T)-statistic images were thresholded using clusters determined by $Z > 2.3$ and clusters exceeding 100 voxels in predefined regions of interest were reported. Clusters formed at this threshold were also tested for whole-brain significance using Monte Carlo simulation of a Gaussian random field, using the AlphaSim software released as part of AFNI version 1014, using a whole-brain significance threshold of $p < 0.05$; the degree of smoothing for this analysis was estimated using the FSL tool *smoothest* v2.1.’

Region of interest analysis—Timeseries for ROI plotting and PPI analyses were determined by: (i) thresholding Z-statistic images from the second level analysis reported above; (ii) binarising this thresholded image to form a mask; (iii) applying the inverse of each individual’s registration, calculated during intersubject registration, to project this mask from single-subject space back to the 3mm^3 isotropic space in which EPI data were acquired; (iv) thresholding (at 0.3) and re-binarising this back-projected mask; (v) extracting the mean timeseries within this region of interest from the pre-processed EPI data. We adopted a leave-one-out approach to ROI construction³¹, in which the mask used to extract each subject’s data was based upon a group analysis containing all the remaining ($n-1$) subjects, and then tested in the independent left-out subject.

To plot effects of individual regressors through time, the timeseries was upsampled, then timelocked to decision onset (Figs. 5b/7b, supplementary Fig. 9) or feedback onset (supplementary Fig. 8b, 8d) of each trial. This creates a data matrix with dimensions $n\text{Trials} \times n\text{Timepoints}$. Each timepoint was regressed against explanatory variables of interest for each subject. The mean \pm standard error (across subjects) of parameter estimates from this regression is plotted. A full description of this approach is given in ⁵⁴.

The masks for timeseries analysis and PPIs were determined as follows:

- IPS mask for main Fig. 5b, attribute-based prediction error (supplementary Fig. 8d), and PPI analysis: all voxels with $Z > 2.8$ for contrast (i) in GLM 1 that lay within an anatomical mask of the parietal cortex.
- Nucleus accumbens/VMPFC mask, feedback phase (supplementary Fig. 8b): all voxels with $Z > 3.1$ for contrast (iv) in model 1, described above, that lay within an anatomical mask encompassing the nucleus accumbens and VMPFC. A more stringent threshold was used for ROI generation than in IPS/dMFC to enable increased anatomical specificity; similar results could be obtained irrespective of exact threshold used.
- Cingulate/paracingulate cortex mask for main Fig. 7b: all voxels with $Z < -2.8$ for contrast (ii) in GLM 1 that lay within an anatomical mask of the cingulate and paracingulate cortices
- Ventromedial prefrontal cortex mask, decision phase (supplementary Fig. 9): an 8mm sphere placed at a location consistent with findings from a recent meta-analysis of value-related activations³⁷

A supplementary methods checklist is available.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank M. Rushworth, K. Tsetsos and M. Usher for helpful discussions, and H. Barron and P. Smittenaar for comments on the manuscript. This work was supported by the Wellcome Trust [L.T.H. Sir Henry Wellcome Fellowship 098830/Z/12/Z; R.J.D. Senior Investigator Award 098362/Z/12/Z; T.E.J.B. Career Development Fellowship 088312AIA]. The Wellcome Trust Centre for Neuroimaging is supported by core funding from the Wellcome Trust, 091593/Z/10/Z.

References

1. Bettman JR. Constructive consumer choice processes. *J Consum Res.* 1998; 25:187–217.
2. Mcfarland DJ, Sibly RM. Behavioral Final Common Path. *Philos T Roy Soc B.* 1975; 270:265–293.
3. Camerer CF, Fehr E. When does “economic man” dominate social behavior? *Science.* 2006; 311:47–52. [PubMed: 16400140]
4. Keeney, RL.; Raiffa, H. *Decisions with multiple objectives: preferences and value functions.* Cambridge University Press; Cambridge: 1993.
5. Rangel A, Camerer C, Montague PR. A framework for studying the neurobiology of value-based decision making. *Nat Rev Neurosci.* 2008; 9:545–556. [PubMed: 18545266]
6. Kable JW, Glimcher PW. The neurobiology of decision: consensus and controversy. *Neuron.* 2009; 63:733–745. [PubMed: 19778504]
7. Padoa-Schioppa C. Neurobiology of economic choice: A good-based model. *Annu Rev Neurosci.* 2011; 34:333–359. [PubMed: 21456961]
8. Hare TA, Schultz W, Camerer CF, O’Doherty JP, Rangel A. Transformation of stimulus value signals into motor commands during simple choice. *Proc Natl Acad Sci U S A.* 2011; 108:18120–18125. [PubMed: 22006321]

9. Lim S-L, O'Doherty JP, Rangel A. Stimulus value signals in ventromedial PFC reflect the integration of attribute value signals computed in fusiform gyrus and posterior superior temporal gyrus. *J Neurosci.* 2013; 33:8729–8741. [PubMed: 23678116]
10. Hunt LT, et al. Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci.* 2012; 15:470–476. [PubMed: 22231429]
11. FitzGerald TH, Seymour B, Dolan RJ. The role of human orbitofrontal cortex in value comparison for incommensurable objects. *J Neurosci.* 2009; 29:8388–8395. [PubMed: 19571129]
12. Rudebeck PH, et al. Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J Neurosci.* 2008; 28:13775–13785. [PubMed: 19091968]
13. Glockner A, Betsch T. Multiple-reason decision making based on automatic processing. *J Exp Psychol Learn.* 2008; 34:1055–1075.
14. Tversky A, Simonson I. Context-dependent preferences. *Manage Sci.* 1993; 39:1179–1189.
15. Payne JW. Task Complexity and Contingent Processing in Decision-Making - Information Search and Protocol Analysis. *Organ Behav Hum Perf.* 1976; 16:366–387.
16. Vlaev I, Chater N, Stewart N, Brown GD. Does the brain calculate value? *Trends Cogn Sci.* 2011; 15:546–554. [PubMed: 21983149]
17. Roe RM, Busemeyer JR, Townsend JT. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol Rev.* 2001; 108:370–392. [PubMed: 11381834]
18. Usher M, McClelland JL. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol Rev.* 2004; 111:757–769. [PubMed: 15250782]
19. Tsetos K, Usher M, Chater N. Preference reversal in multiattribute choice. *Psychol Rev.* 2010; 117:1275–1293. [PubMed: 21038979]
20. Park SQ, Kahnt T, Rieskamp J, Heekeren HR. Neurobiology of value integration: when value impacts valuation. *J Neurosci.* 2011; 31:9307–9314. [PubMed: 21697380]
21. Wang X-J. Decision making in recurrent neuronal circuits. *Neuron.* 2008; 60:215–234. [PubMed: 18957215]
22. Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev.* 2006; 113:700–765. [PubMed: 17014301]
23. Louie K, Khaw MW, Glimcher PW. Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci U S A.* 2013; 110:6139–6144. [PubMed: 23530203]
24. Noonan MP, et al. Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc Natl Acad Sci U S A.* 2010; 107:20547–20552. [PubMed: 21059901]
25. Balleine BW. Neural bases of food-seeking: affect, arousal and reward in corticostriatal limbic circuits. *Physiol Behav.* 2005; 86:717–730. [PubMed: 16257019]
26. Tricomi E, Balleine BW, O'Doherty JP. A specific role for posterior dorsolateral striatum in human habit learning. *Eur J Neurosci.* 2009; 29:2225–2232. [PubMed: 19490086]
27. Boorman ED, Behrens TEJ, Woolrich MW, Rushworth MSF. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron.* 2009; 62:733–743. [PubMed: 19524531]
28. Piliastides MG, Biele G, Heekeren HR. A mechanistic account of value computation in the human brain. *Proc Natl Acad Sci U S A.* 2010; 107:9430–9435. [PubMed: 20439711]
29. Heekeren HR, Marrett S, Bandettini PA, Ungerleider LG. A general mechanism for perceptual decision-making in the human brain. *Nature.* 2004; 431:859–862. [PubMed: 15483614]
30. Kriegeskorte N, Simmons WK, Bellgowan PS, Baker CI. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat Neurosci.* 2009; 12:535–540. [PubMed: 19396166]
31. Boorman ED, Rushworth MF, Behrens TE. Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J Neurosci.* 2013; 33:2242–2253. [PubMed: 23392656]

32. Gershman SJ, Pesaran B, Daw ND. Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *J Neurosci.* 2009; 29:13524–13531. [PubMed: 19864565]
33. O'Reilly JX, Woolrich MW, Behrens TE, Smith SM, Johansen-Berg H. Tools of the trade: psychophysiological interactions and functional connectivity. *Soc Cogn Affect Neurosci.* 2012; 7:604–609. [PubMed: 22569188]
34. Jbabdi S, Lehman JF, Haber SN, Behrens TE. Human and monkey ventral prefrontal fibers use the same organizational principles to reach their targets: tracing versus tractography. *J Neurosci.* 2013; 33:3190–3201. [PubMed: 23407972]
35. Alexander WH, Brown JW. Medial prefrontal cortex as an action-outcome predictor. *Nat Neurosci.* 2011; 14:1338–1344. [PubMed: 21926982]
36. Clithero JA, Rangel A. Informatic parcellation of the network involved in the computation of subjective value. *Soc Cogn Affect Neurosci.* 2013 doi: 10.1093/scan/nst106.
37. Levy DJ, Glimcher PW. The root of all value: a neural common currency for choice. *Curr Opin Neurobiol.* 2012; 22:1027–1038. [PubMed: 22766486]
38. Collins AG, Frank MJ. Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychol Rev.* 2013; 120:190–229. [PubMed: 23356780]
39. Mars RB, 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.* 2011; 31:4087–4100. [PubMed: 21411650]
40. Thiel CM, Zilles K, Fink GR. Cerebral correlates of alerting, orienting and reorienting of visuospatial attention: an event-related fMRI study. *NeuroImage.* 2004; 21:318–328. [PubMed: 14741670]
41. Piazza M, Izard V, Pinel P, Le Bihan D, Dehaene S. Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron.* 2004; 44:547–555. [PubMed: 15504333]
42. O'Reilly JX, Jbabdi S, Rushworth MF, Behrens TE. Brain systems for probabilistic and dynamic prediction: computational specificity and integration. *PLoS Biol.* 2013; 11:e1001662. [PubMed: 24086106]
43. Fellows LK. Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain.* 2006; 129:944–952. [PubMed: 16455794]
44. Price JL. Definition of the orbital cortex in relation to specific connections with limbic and visceral structures and other cortical regions. *Ann NY Acad Sci.* 2007; 1121:54–71. [PubMed: 17698999]
45. Haber SN, Fudge JL, McFarland NR. Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *J Neurosci.* 2000; 20:2369–2382. [PubMed: 10704511]
46. Carandini M, Heeger DJ. Normalization as a canonical neural computation. *Nat Rev Neurosci.* 2012; 13:51–62. [PubMed: 22108672]
47. Chau BK, Kolling N, Hunt LT, Walton ME, Rushworth MF. A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nat Neurosci.* 2014; 17:463–469. [PubMed: 24509428]
48. Rangel A, Clithero JA. Value normalization in decision making: theory and evidence. *Curr Opin Neurobiol.* 2012; 22:970–981. [PubMed: 22939568]
49. Sjöberg L. Choice Frequency and Similarity. *Scand J Psychol.* 1977; 18:103–115.
50. Hunt LT, Woolrich MW, Rushworth MF, Behrens TE. Trial-type dependent frames of reference for value comparison. *PLoS Comp Biol.* 2013; 9:e1003225.

Additional references for online methods

51. Sutton, RS.; Barto, AG. Reinforcement learning: an introduction. MIT Press; 1998.
52. Weiskopf N, Hutton C, Josephs O, Deichmann R. Optimal EPI parameters for reduction of susceptibility-induced BOLD sensitivity losses: a whole-brain analysis at 3 T and 1.5 T. *Neuroimage.* 2006; 33:493–504. [PubMed: 16959495]

53. Woolrich MW, et al. Bayesian analysis of neuroimaging data in FSL. *Neuroimage*. 2009; 45:S173–186. [PubMed: 19059349]
54. Behrens TE, Hunt LT, Woolrich MW, Rushworth MF. Associative learning of social value. *Nature*. 2008; 456:245–249. [PubMed: 19005555]

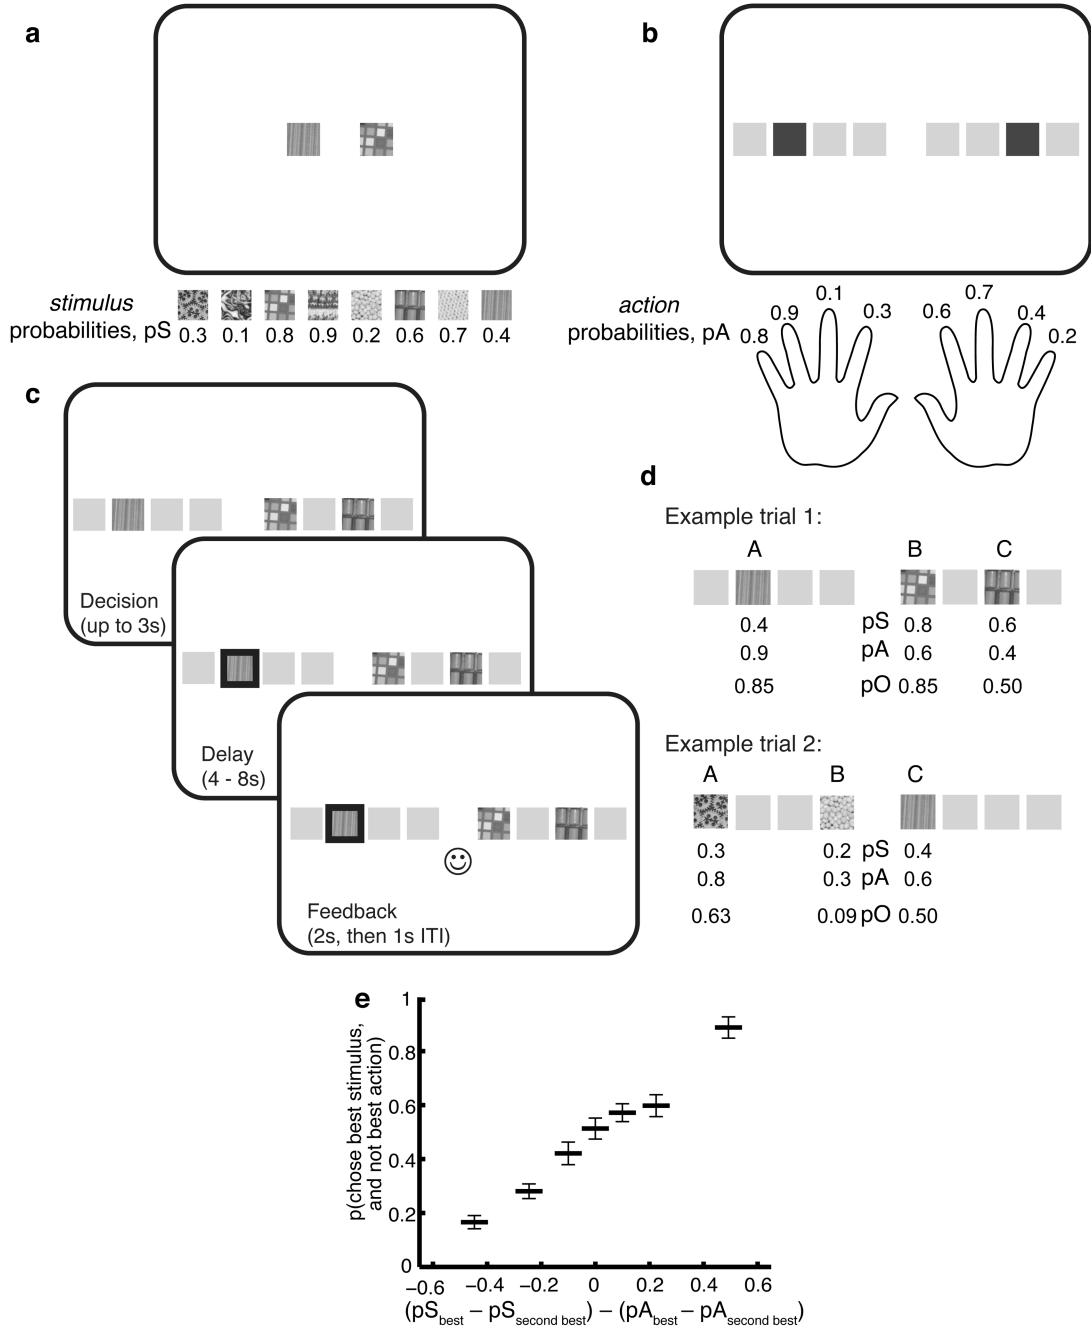


Figure 1. Experimental task design. (a) Outside scanner, subjects learned reward probabilities (pS) associated with 8 stimuli, via pairwise choices between stimuli. Subjects trained to reach a minimum performance criterion of 90% correct (choosing higher valued stimulus). (b) Subjects also learned reward probabilities (pA) associated with 8 actions (finger presses), via pairwise choices between actions (dark grey squares indicate currently available actions; each action tied to an onscreen spatial location). Subjects trained to same criteria as for stimuli. Stimulus and action training was alternated in blocks of at least 70 trials, with an

additional 40 trial ‘refresher’ block immediately prior to the fMRI experiment. (c) Inside the scanner, subjects performed a three-option choice experiment, in which each option comprised one previously learned stimulus and one previously learned action. Subjects were instructed to weight stimulus and action information equally on each trial, and select the best option to obtain points that subsequently converted into monetary reward (see methods). Reward was delivered probabilistically according to pO , the optimal combination of pS and pA (see equation 1 in methods), for the chosen option. (d) Two example trials. In trial 1, options A and B are of equal (integrated) value. The action attribute favors option A and so would be deemed ‘relevant’ if A were chosen, and stimulus deemed ‘irrelevant’. The converse would be true if option B were chosen. Were option C chosen, the trial would be discarded from fMRI analysis. In trial 2, action would be deemed relevant if A were chosen, whereas stimulus would be deemed relevant if C were chosen. (e) Choice behavior. On trials where stimulus and action favor different options, the probability of choosing the option favored by the stimulus attribute (ordinate) is plotted as a function of the difference in probabilities on the two dimensions (abscissa). Bars show mean \pm s.e.m. across 19 subjects.

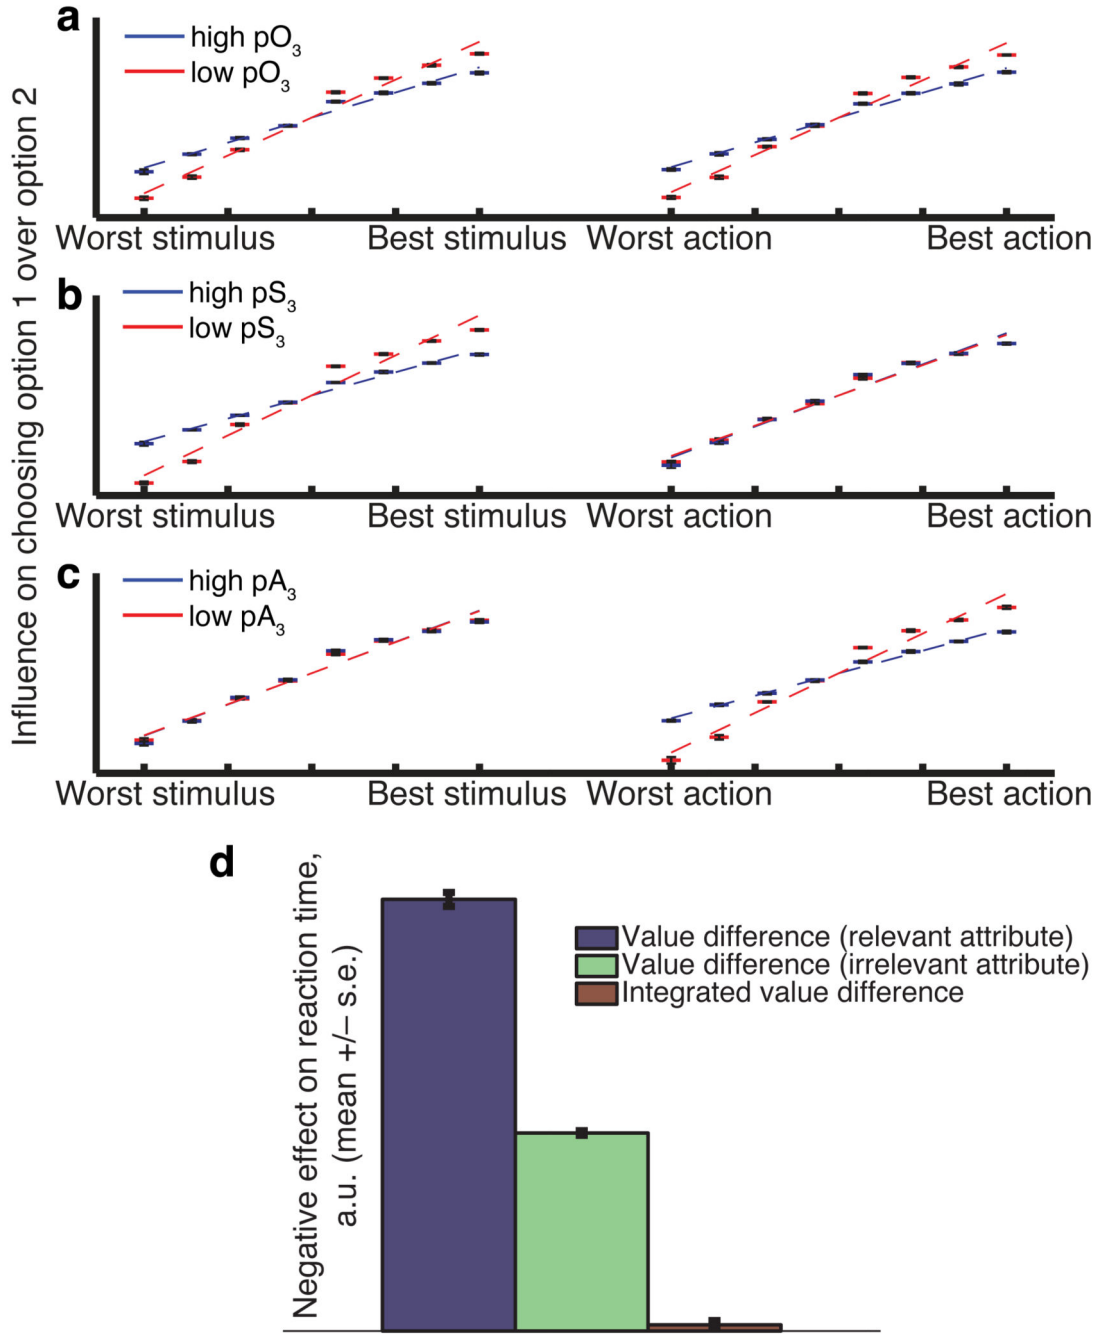


Figure 2. Predictions of behavior from hierarchical model. (a)-(c) Distractor effects, described in text and elucidated in subject behavior in figure 3(a)-(c). A distractor effect is where option 3 value affects choice probabilities between options 1 and 2, here assessed via logistic regression. The model shows a classic value-based distractor effect (a), but also a within-attribute distractor effect (b/c), as is found in subject behavior. (d) Reaction time effects in model estimated via linear regression, comparable to those revealed in subject behavior in

figure 3(e). The model is most heavily influenced by value difference on the relevant attribute, not the irrelevant attribute or integrated values.

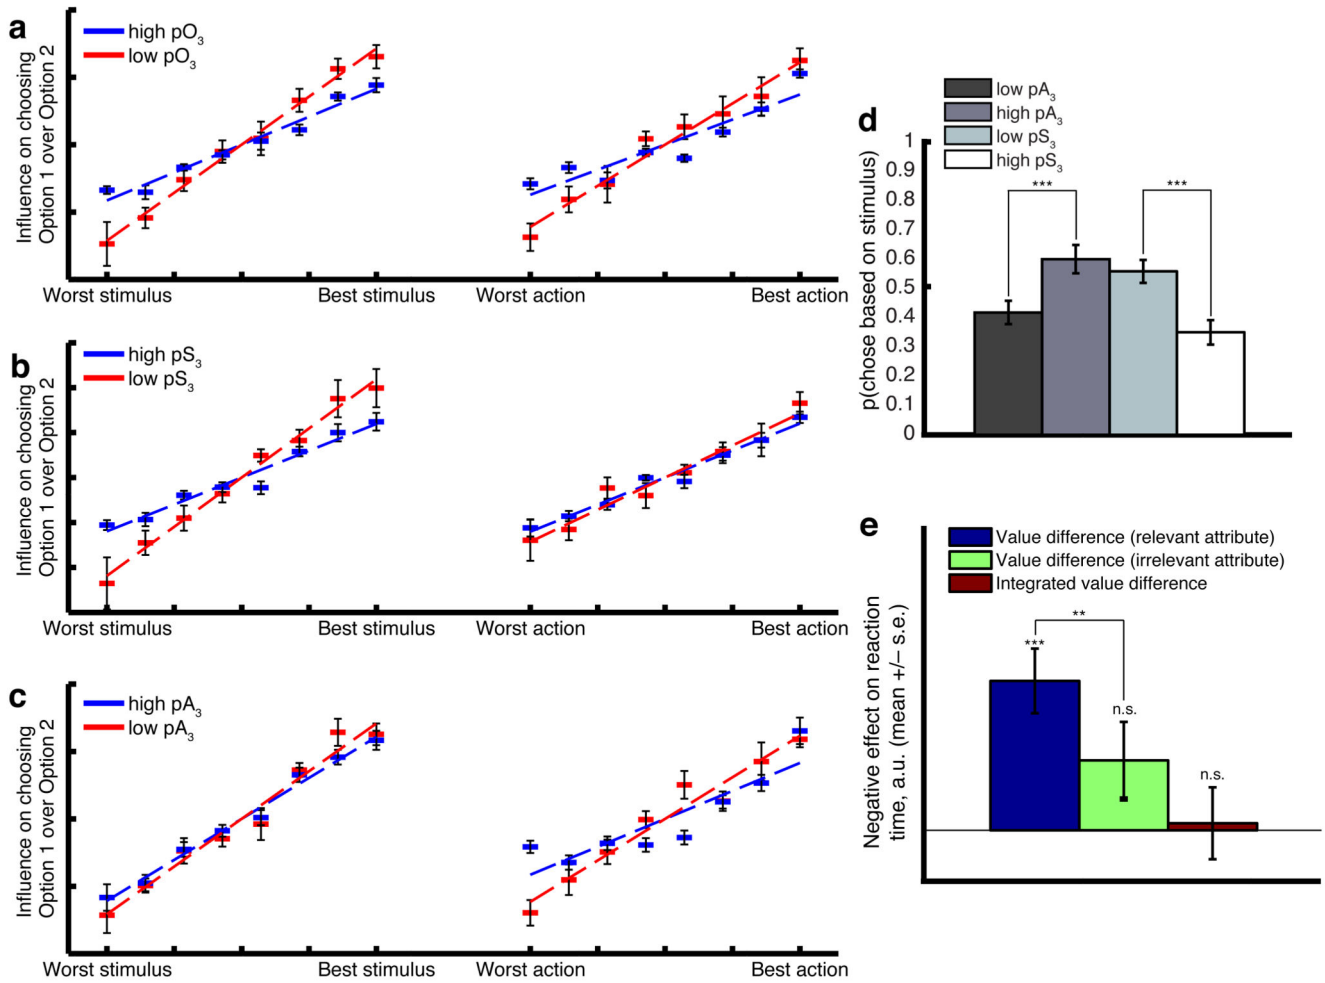


Figure 3.

Subject behavior (n=19 subjects). (a) Value-based distractor effect – high values of option 3 make options 1 and 2 less discriminable. Datapoints show mean \pm s.e.m. (across subjects) of logistic regression parameters for each of the 16 stimuli/actions on the probability of choosing option 1 vs option 2 (see methods for details). Trials have been split into those where the 3rd option has a high value, and those where it has a low value. Red points show effects when option 3 pO is high (in top 33% of values), blue points when option 3 pO is low (in bottom 33%). Lines show average best fit to datapoints; the slope of this line is significantly different between high pO₃ vs low pO₃ ($T(18)=3.17$, $p=0.0053$ (stimulus influence); $T(18)=2.45$, $p=0.025$ (action influence)). (b) When option 3 is split selectively on the stimulus attribute, the distractor effect remains on the stimulus discriminability of options 1 and 2 ($T(18)=2.51$, $p=0.021$), but not on the action discriminability ($T(18)=0.82$, $p=0.42$). (c) When option 3 is split selectively on the action attribute, the distractor effect remains on the action attribute ($T(18)=3.18$, $p=0.0052$), but not on the stimulus attribute ($T(18)=1.78$, $p=0.092$). (d) Analysis of trials where evidence given by stimulus and action are equal and opposite. Probability of choosing option 1, on trials where $pS1 > pS2$, $pA2 > pA1$, and $(pS1 - pS2) \approx (pA2 - pA1)$. On such trials, ‘choosing based on stimulus’ (plotted

on ordinate) is equivalent to 'choosing option 1'. Option 3 is always one of the two unchosen options. Bars show mean \pm s.e. (across subjects). *** denotes $p=8.5 \times 10^{-4}$, paired T-test between low and high pR_3 ($T(18)=-3.99$) and $p=9.0 \times 10^{-6}$, paired T-test between low and high pS_3 ($T(18)=6.11$). Interaction in two-way ANOVA, $F(1,72) = 16.62$, $p=1.7 \times 10^{-4}$. (e) Reaction times are more heavily influenced by the relevant attribute than the irrelevant attribute. Bars show mean \pm s.e.m. (across subjects) of effects of value difference on subject reaction times, estimated via linear regression (y-axis is flipped – i.e. higher value differences typically elicit faster reaction times). *** denotes $p=2.13 \times 10^{-4}$, one-sample T-test ($T(18)=4.62$); ** denotes $p=0.0030$, paired T-test ($T(18)=3.43$); n.s. = non-significant one-sample T-test.

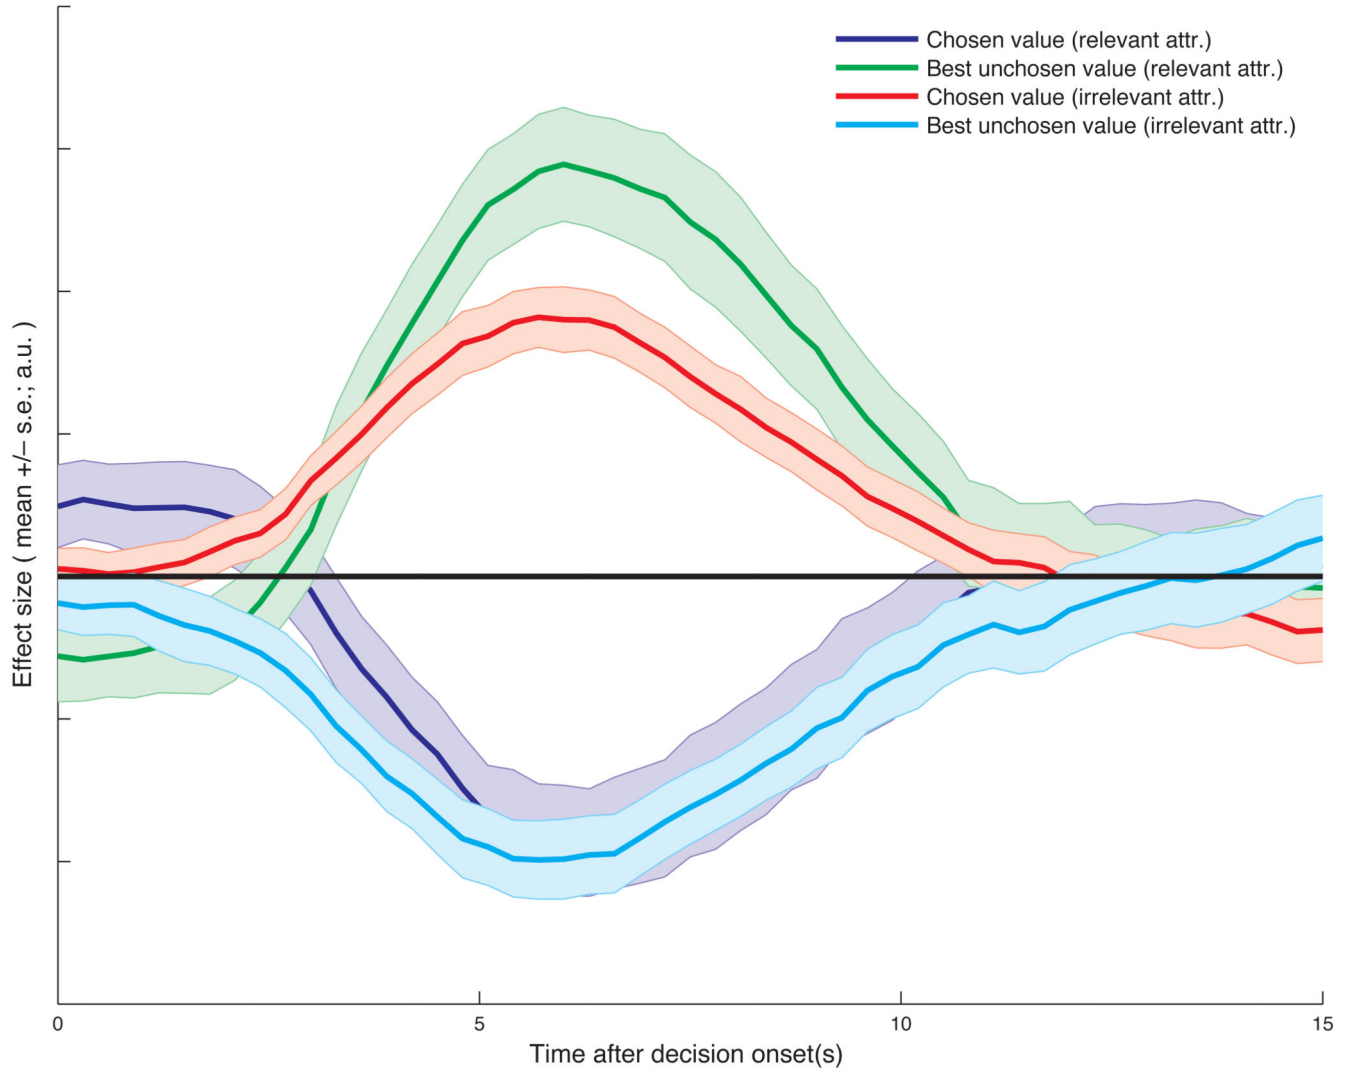


Figure 4. Model predictions of fMRI data, derived from ‘attribute comparison’ node of the hierarchical model. Model activity from each trial was convolved with a haemodynamic response function, and then regressed against chosen value and best unchosen value on both relevant and irrelevant attributes (together with a constant term, and model reaction time included as a coregressor of no interest). Bars show mean \pm s.e.m. of parameter estimates from the regression across 10 simulations of the model. The contrast $(\text{chosen value} - \text{best unchosen value})_{\text{irrelevant}} - (\text{chosen value} - \text{best unchosen value})_{\text{relevant}}$ is encoded by the model.

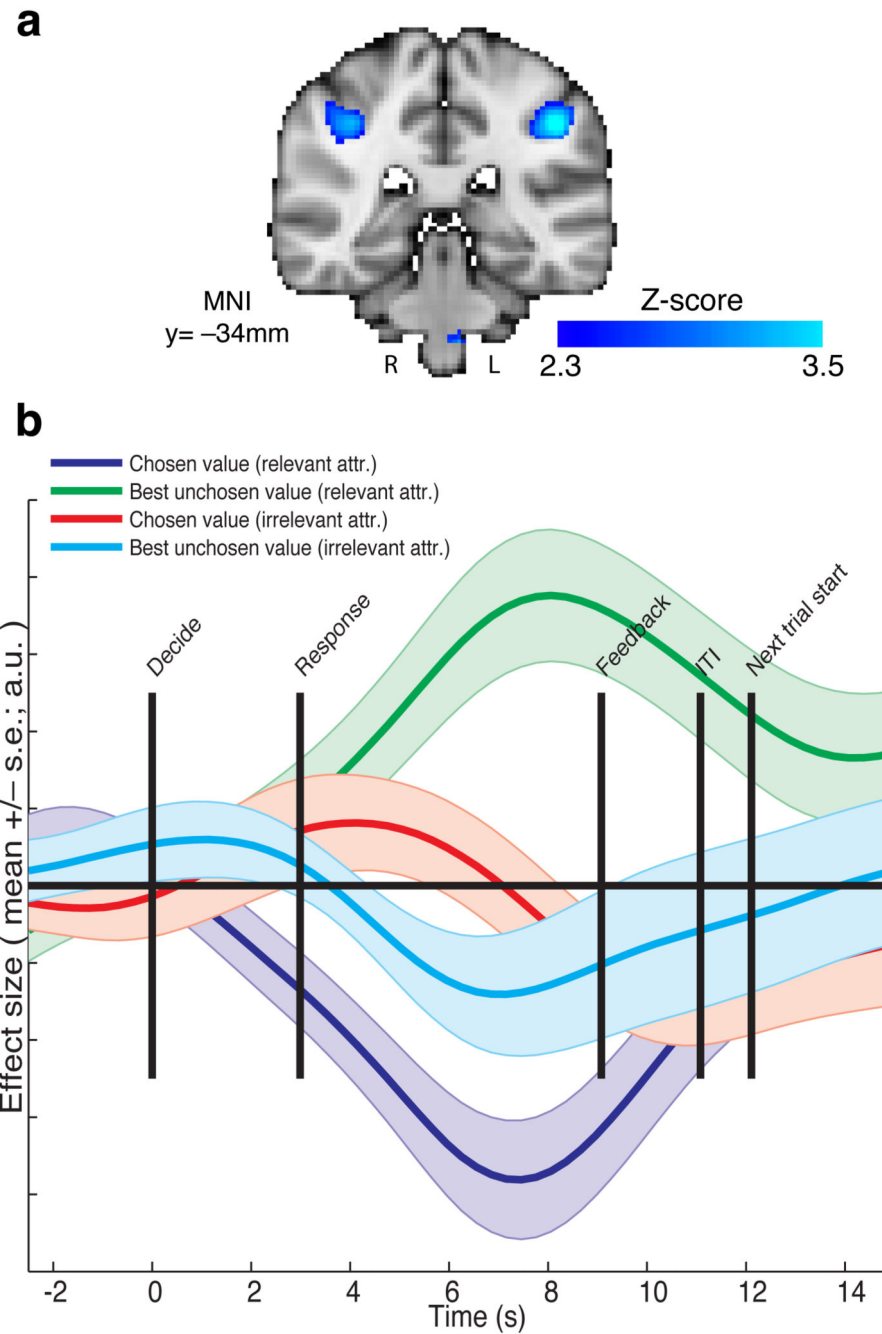


Figure 5. Intraparietal sulcus shows features of an attribute comparison signal, with opposing signs for relevant and irrelevant attributes. (a) Statistical parametric map of the contrast (chosen value – best unchosen value)_{irrelevant} – (chosen value – best unchosen value)_{relevant}, at the time of making the decision, thresholded at $Z > 2.3$ uncorrected for display purposes ($n = 19$ subjects). A bilateral portion of IPS reflects this contrast, with the left IPS surviving whole-brain correction (FWE-corrected $p = 0.0023$, cluster-forming threshold $Z > 2.3$; peak $Z = 3.65$, MNI = $-38, -42, 40$ mm). (b) Timeseries analysis of this region, time-locked to decision phase,

reveals negative correlates of (chosen value)_{irrelevant} and positive correlates of (best unchosen value)_{relevant}, but positive correlates of (chosen value)_{relevant} and negative correlates of (best unchosen value)_{irrelevant} (bars show mean \pm s.e. across subjects). To avoid circular analysis, a leave-one-out cross-validation approach was used for timeseries extraction (see methods).

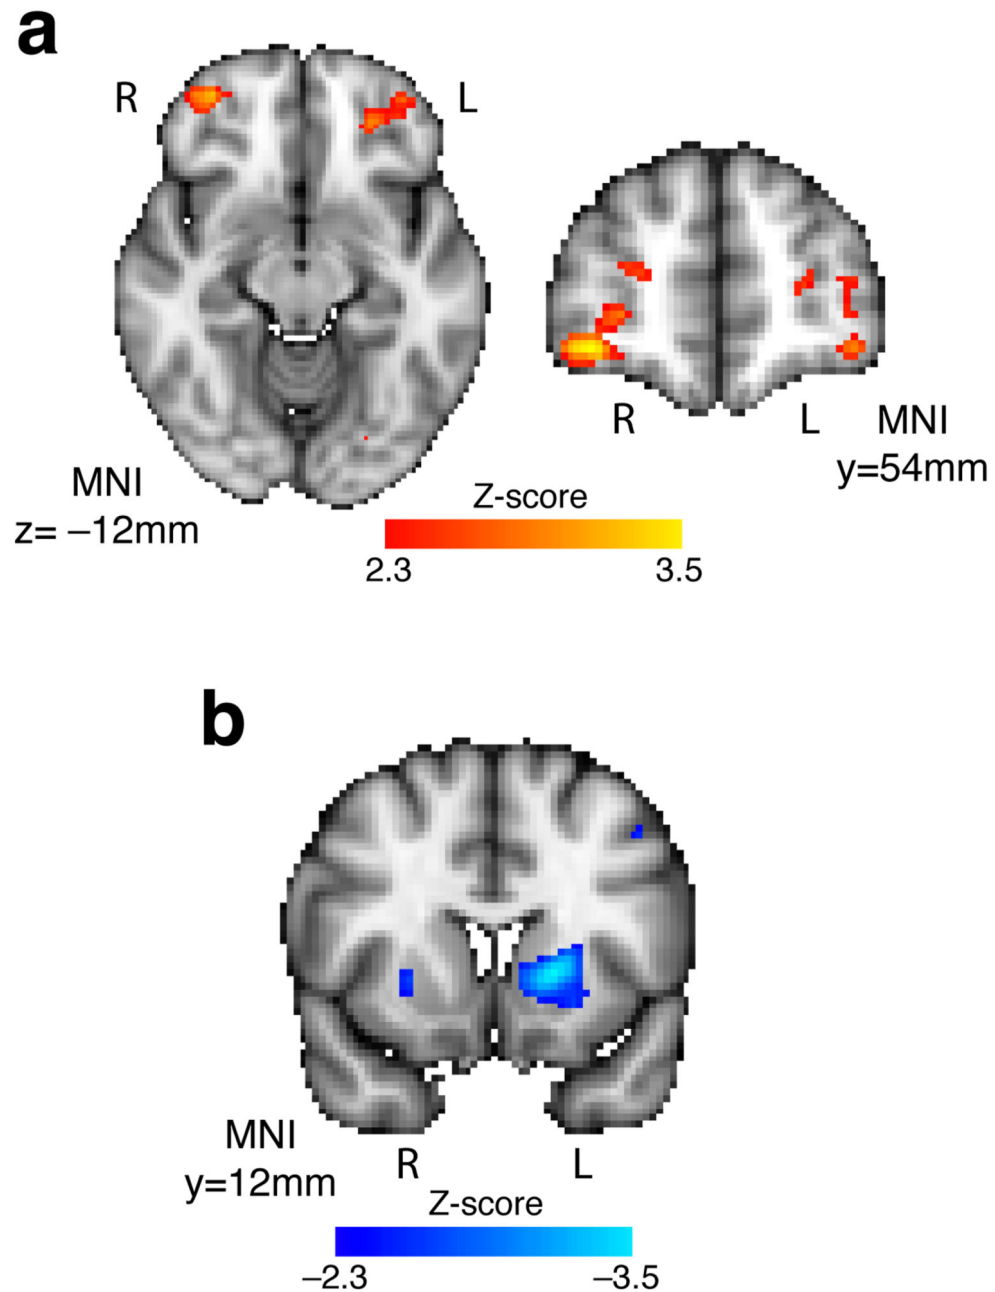


Figure 6. Psychophysiological interaction with intraparietal cortex. (a) Functional connectivity with a bilateral portion of the anterior/lateral orbitofrontal cortex was greater on trials where stimulus was the relevant attribute relative to trials where action was the relevant attribute (peak $Z = 3.52$; MNI = 36,52,-10 mm (right OFC); peak $Z = 2.91$, MNI = -26,42,-12mm (left OFC); all voxels with $Z > 2.3$ shown, both clusters contained > 100 voxels at this threshold) ($n = 19$ subjects). (b) Functional connectivity with a portion of the left putamen was greater on trials where action was the relevant attribute relative to trials where stimulus

was the relevant attribute (peak $Z = -3.46$; MNI = $-18,12,2$ mm; all voxels with $Z > 2.3$ shown, the left putamen contained >100 voxels at this threshold, whilst the right putamen showed a similar, smaller activation).

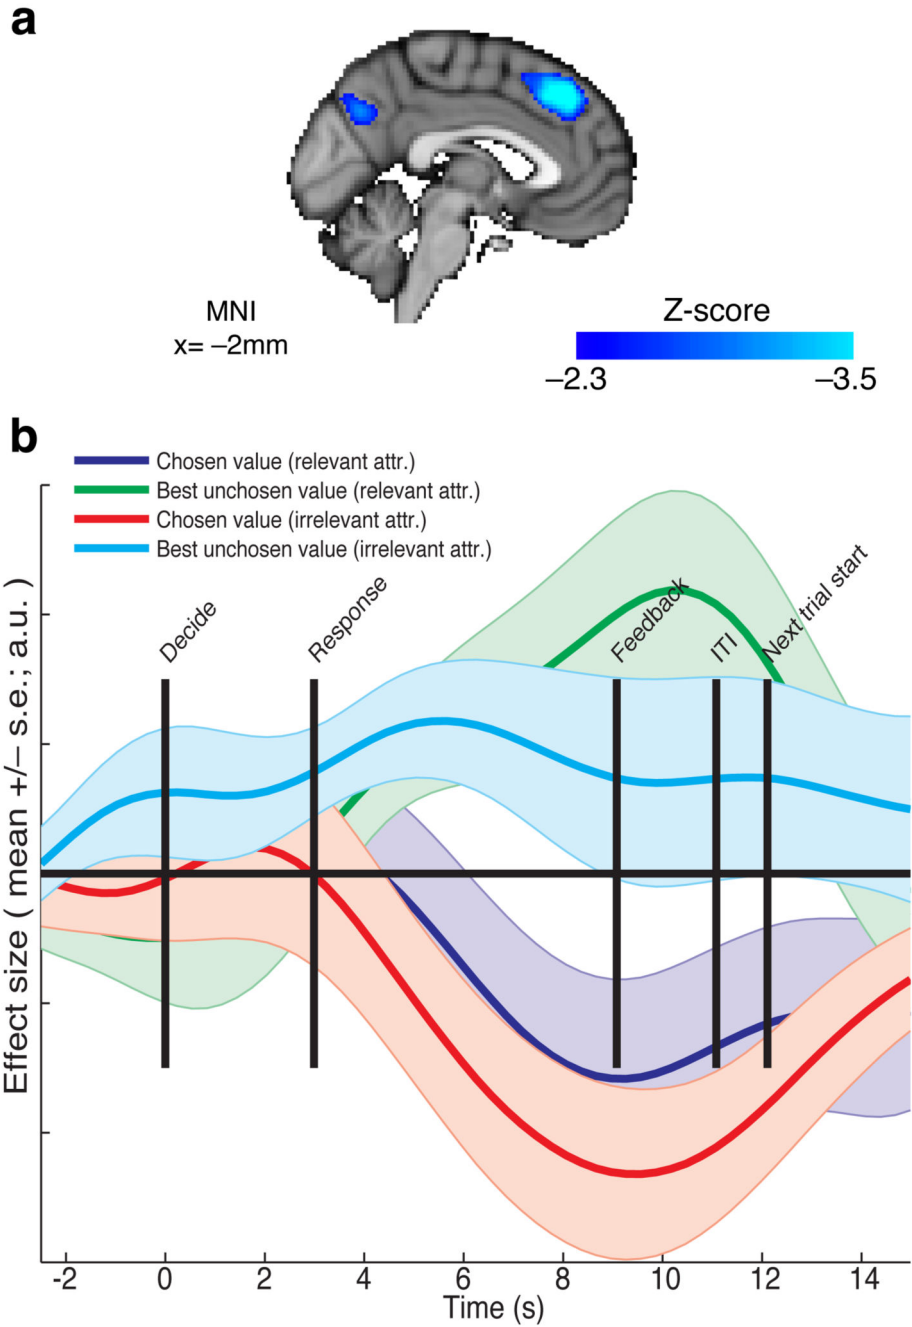


Figure 7. Dorsal medial frontal cortex shows an integrated value difference signal, with the same sign for both relevant and irrelevant attributes. (a) Statistical parametric map of the contrast $(\text{chosen value} - \text{best unchosen value})_{\text{relevant}} + (\text{chosen value} - \text{best unchosen value})_{\text{irrelevant}}$, at the time of making the decision, thresholded at $Z < -2.3$ uncorrected for display purposes. A portion of dorsal medial frontal cortex reflects this contrast (FWE-corrected $p = 0.0054$, cluster-forming threshold $Z < -2.3$; peak $Z = -3.85$, MNI = $-2,34,46\text{mm}$). Other regions surviving whole-brain correction are detailed in table S3. (n=19 subjects) (b) Timeseries

analysis of this region, time-locked to decision phase, reveals negative correlates of both $(\text{chosen value})_{\text{relevant}}$ and $(\text{chosen value})_{\text{irrelevant}}$, and positive correlates of $(\text{best unchosen value})_{\text{relevant}}$ and $(\text{best unchosen value})_{\text{irrelevant}}$ (bars show mean \pm s.e. across subjects), slightly delayed in time relative to attribute comparison signal in IPS (compare to figure 5b). As before, cross-validated ROIs were used for timeseries extraction (see methods).