

# Hierarchical competitions subserving multi-attribute choice

Laurence T Hunt<sup>1,2</sup>, Raymond J Dolan<sup>1</sup> & Timothy E J Behrens<sup>1,3</sup>

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 that competitive inhibition may occur in early valuation stages, before option selection. We found that 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 within-attribute competition, competition between attributes and option selection. This multi-layered inhibition framework challenges the assumption that option values are computed before choice. Instead, our results suggest a canonical competition mechanism throughout all stages of a processing hierarchy, not simply at a final choice stage.

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 consumers<sup>1</sup>, foraging animals<sup>2</sup>, social interactions<sup>3</sup> and economic choice<sup>4</sup>. Although key aspects of multi-attribute choice behavior have been well characterized, their neural basis has not yet been systematically examined. This question is important 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.

When faced with such choices, human subjects dynamically construct their preferences online as opposed to merely revealing them<sup>1</sup>. A common assumption in neurobiological studies of reward-guided choice is that preference construction depends first on combining attributes in each option into an integrated value, followed by a process involving value comparison<sup>5–11</sup>. 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 one<sup>12</sup>. Some expressions of choice behavior indicate that subjects do indeed integrate different features of an option to form a unitary value<sup>13</sup>.

However, anomalies in human decision-making indicate that this normative explanation cannot fully account for subjects’ choice behavior by itself. For instance, behavioral economic studies have highlighted preference reversals between two options when a third option is introduced. As this is critically dependent on the degree of similarity between alternatives on specific attributes, this raises the likelihood of a within-attribute comparison process<sup>14</sup>. Studies of information gathering during multi-attribute choice, containing multiple options and multiple features, have also suggested that evidence is acquired within-attribute, at least initially<sup>15</sup>. Such behavior is explained by several alternative accounts, and these

depend on different combinations of within-attribute and within-option comparisons<sup>1,14,16–19</sup>.

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 before choice. One obstacle to exploring this conundrum has been the difficulty of designing tasks in which there is transparent behavioral evidence for within-attribute comparison<sup>20</sup>. In addition, there is a lack of a candidate neural mechanism by which this type of decision might be implemented.

Here we provide evidence that preference construction in multi-attribute choice may occur via a hierarchical competition process<sup>17–19</sup>. Such a model argues that because competition via mutual inhibition is a canonical feature of local neural circuits<sup>21</sup>, it should occur at all levels of representation. In this scheme, comparison would still occur at the level of option values<sup>5–10</sup>, but it also has 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 modeling in a novel multi-attribute choice task. 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. Notably, on each trial, one or another of the attributes was more salient (relevant) for guiding behavior, allowing us to investigate how competition for attribute salience is implemented neurally. Our key findings, based on functional magnetic resonance imaging (fMRI) data, were that intraparietal

<sup>1</sup>Wellcome Trust Centre for Neuroimaging, University College London, London, UK. <sup>2</sup>Sobell Department of Motor Neuroscience, University College London, London, UK. <sup>3</sup>Centre for Functional MRI of the Brain, University of Oxford, Nuffield Department of Clinical Neurosciences, John Radcliffe Hospital, Oxford, UK. Correspondence should be addressed to L.T.H. (laurence.hunt@ucl.ac.uk).

Received 23 May; accepted 12 September; published online 12 October 2014; doi:10.1038/nn.3836

sulcus (IPS) signaled a competition over which attribute was the most salient for the current decision, whereas portions of medial frontal cortex reflected an 'integrated' value signal. Consistent 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. 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

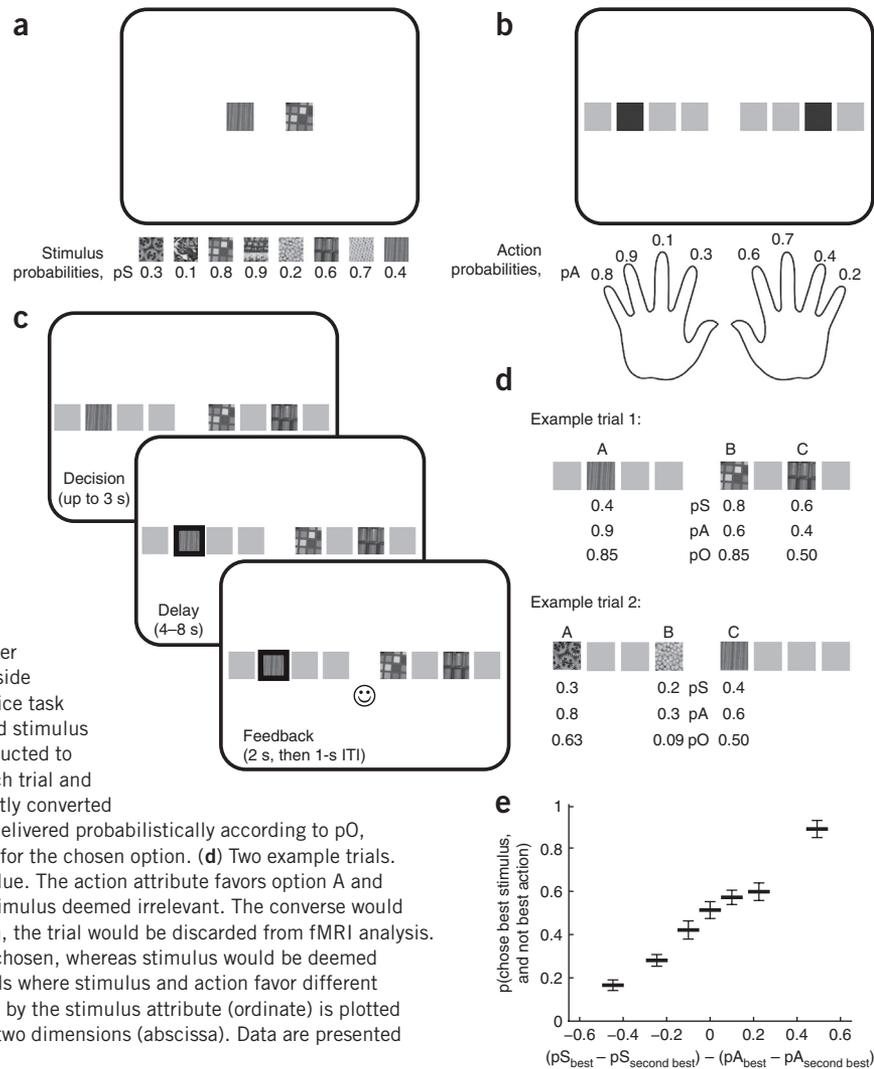
### A multi-attribute reward-guided decision task

We used a multi-attribute choice task, which forms the basis for our modeling and behavioral results. Subjects were trained on the relative likelihood of receiving a reward on a set of eight different images (not tied to any spatial location, hereafter referred to as stimuli; **Fig. 1a**) and a set of eight different button presses (tied to spatial locations onscreen, hereafter referred to as actions; **Fig. 1b**). Subjects learned action-reward probabilities (pA) and stimulus-reward probabilities (pS) 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 min, until they attained a preordained performance criterion.

Following training, subjects underwent fMRI scanning while 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 on pO, a Bayesian combination of probabilistic information from the stimulus (pS) and action (pA) (**Fig. 1d** and Online Methods). The logic behind this rule for combining the pre-learned 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 toward using either stimulus or action attribute (**Supplementary Fig. 1**).

Subjects successfully deployed information learned in the initial phase of the experiment to guide their decisions inside the scanner. Based on the Bayesian integrated values (pO), subjects chose the best option on  $79.0 \pm 3.8\%$  (mean  $\pm$  s.d.) of all trials. On more challenging 'conflict' trials, in which the best stimulus favored one option, but the best action favored a different option, this figure fell to  $70.9 \pm 4.6\%$ . On  $89.6 \pm 4.8\%$  of all of these more challenging trials, subjects selected



**Figure 1** Experimental task design. (a) Outside of the scanner, subjects learned reward probabilities (pS) associated with eight 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 eight actions (finger presses) via pairwise choices between actions (dark gray squares indicate currently available actions; each action was tied to an onscreen spatial location). Subjects were 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 before the fMRI experiment. (c) Inside of the scanner, subjects performed a three-option choice task 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 (Online Methods). Reward was delivered probabilistically according to pO, the optimal combination of pS and pA (equation (1)), 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). Data are presented as mean  $\pm$  s.e.m. across 19 subjects.

either the best pS or the best pA. However, there was a relatively high proportion of challenging trials ( $22.6 \pm 3.8\%$ ) in which subjects went with a high pS or pA when it was not the highest pO. 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 in which 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 ( $\sim 5$  per subject) in which the highest pO was neither the best pS nor the best pA.

A further simple measure of subjects' behavior is shown in **Figure 1e**. This measure focuses on the challenging conflict trials, showing the probability of using the stimulus attribute to guide behavior (that is, 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 on the basis of the largest stimulus, and vice versa. In addition, because the point of subjective equivalence ( $P = 0.5$ ) occurs when within-attribute differences are equal, this shows that subjects did not have any bias (on average) towards 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, in which the probability of choosing an option is a function of its value relative to other alternatives. To also capture subjects' reaction times (RTs), comparison may be modeled as a dynamic process in which evidence is accumulated through time. There are several closely interrelated dynamical models<sup>17–19,21,22</sup>. One class assumes leaky, competing accumulators for each option that receive value-related inputs and compete via inhibition<sup>18,21</sup>.

Notably, softmax choice models and evidence accumulator models make firm predictions as to which factors should influence choices and RTs in multi-alternative and multi-attribute decisions. First, in the softmax choice rule, the relative frequency of choosing option 1 ( $p(C = 1)$ ) over option 2 ( $p(C = 2)$ ) is independent of any remaining values in the decision; that is, the integrated value of option 3 (pO<sub>3</sub>) does not affect the ratio of choosing between options 1 and 2

$$\frac{p(C = 1)}{p(C = 2)} = \left( \frac{e^{pO_1}}{e^{pO_1} + e^{pO_2} + e^{pO_3}} \right) \div \left( \frac{e^{pO_2}}{e^{pO_1} + e^{pO_2} + e^{pO_3}} \right) = \frac{e^{pO_1}}{e^{pO_2}}$$

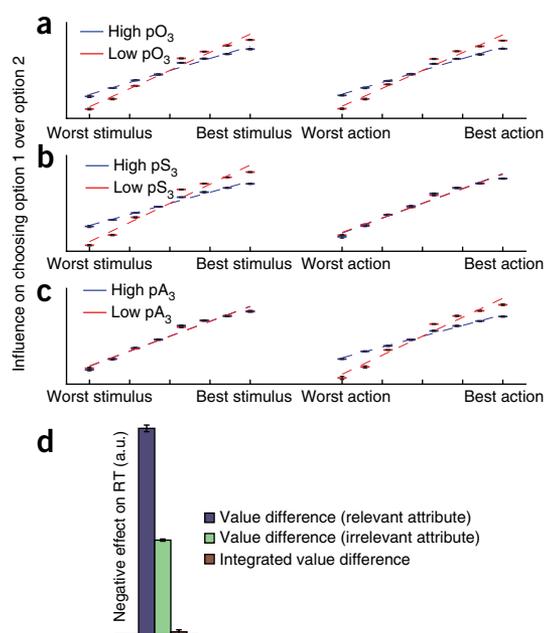
The softmax temperature parameter is omitted for clarity. Second, predictions of RTs and choices from accumulator models of decision-making suggest that they will principally be driven by variation in integrated values (pO) rather than the underlying constituents of these values (pA and pS).

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 1 and 2. One account of distractor effects appeals to a neural process of normalization by the set of choice alternatives on offer before comparison. This implies that high-valued third items prove to be more distracting than low-valued third items and so impair discrimination, matching empirical observations in humans and macaque monkeys<sup>23</sup>. Specific classes of distractor effects also occur in multi-attribute choice, such as 'compromise', 'similarity' and 'attraction' effects<sup>17</sup>. These necessitate comparison at the level of underlying attributes and not solely at the level of integrated values, and so have inspired models in which evidence accumulation and competition occur at multiple hierarchical levels<sup>17–19</sup>.

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 between attributes (in a competition for attribute relevance). Normalization is also included in the model, but occurs at the initial input level of individual attributes rather than on integrated values (see Online Methods for detailed mathematical description of the model).

The most novel feature of the model is the competition for attribute relevance. Inputs to this competition (**Supplementary Fig. 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 substantially better than the second best) dominate this competition and so become more relevant for guiding model choices. This is achieved via feedback connections that 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 RT, once activity in this node reaches a decision threshold.

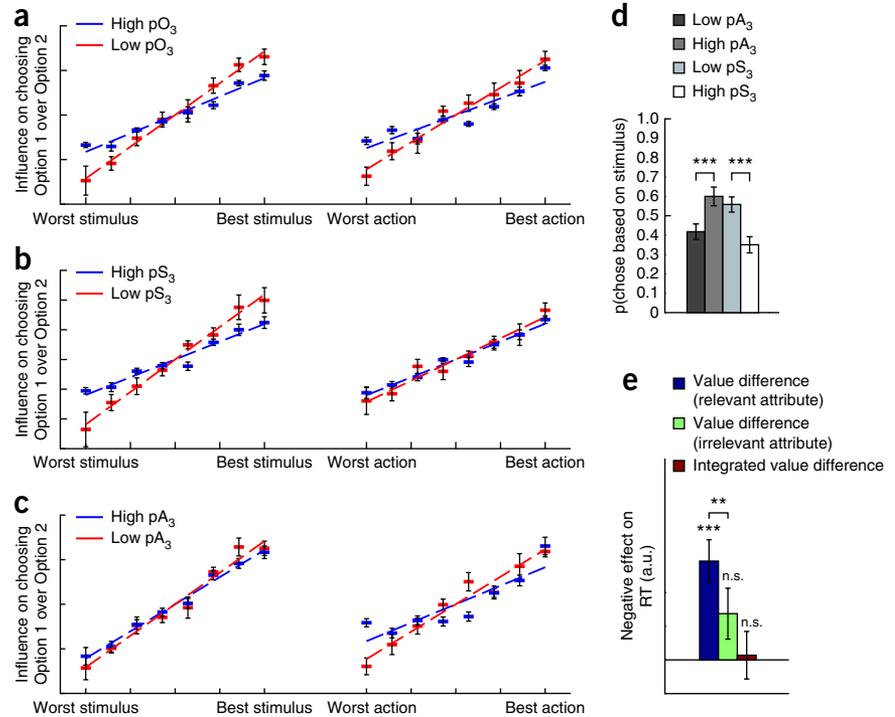


**Figure 2** Predictions of behavior from hierarchical model. (**a–c**) Distractor effects (see **Fig. 3a–c**). A distractor effect is when option 3 value affects choice probabilities between options 1 and 2, here assessed via logistic regression. The model shows a classic value-based distractor effect (**a**), but also a within-attribute distractor effect (**b,c**), as is found in subject behavior. (**d**) RT effects in model estimated via linear regression, comparable with those revealed in subject behavior in **Figure 3e**. The model is most heavily influenced by value difference on the relevant attribute, not the irrelevant attribute or integrated values. Data are presented as mean  $\pm$  s.e.

**Figure 3** Subject behavior ( $n = 19$  subjects).

(a) Value-based distractor effect. High values of option 3 make options 1 and 2 less discriminable. Data points show mean  $\pm$  s.e.m. (across subjects) of logistic regression parameters for each of the 16 stimuli and actions on the probability of choosing option 1 versus option 2 (Online Methods). Trials have been split into those where the third 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 indicate when option 3 pO is low (in bottom 33%). Lines show average best fit to data points; the slope of this line was significantly different between high  $pO_3$  versus low  $pO_3$  (stimulus influence,  $T(18) = 3.17$ ,  $P = 0.0053$ ; action influence,  $T(18) = 2.45$ ,  $P = 0.025$ ). (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

in which evidence given by stimulus and action are equal and opposite. Probability of choosing option 1, on trials where  $pS_1 > pS_2$ ,  $pA_2 > pA_1$ , and  $(pS_1 - pS_2) \approx (pA_2 - pA_1)$ . On such trials, choosing on the basis of stimulus (plotted on ordinate) is equivalent to choosing option 1. Option 3 is always one of the two unchosen options. Data are presented as mean  $\pm$  s.e. (across subjects).  $***P = 8.5 \times 10^{-4}$ , 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) RTs are more heavily influenced by the relevant attribute than the irrelevant attribute. Data are presented as mean  $\pm$  s.e. (across subjects) of effects of value difference on subject RTs, estimated via linear regression (y axis is flipped, that is, higher value differences typically elicit faster RTs).  $***P = 2.13 \times 10^{-4}$ , one-sample  $t$  test ( $T(18) = 4.62$ );  $**P = 0.0030$ , paired  $t$  test ( $T(18) = 3.43$ ); n.s. = non-significant one-sample  $t$  test ( $P = 0.08$  and  $P = 0.84$ ).



Notably, there are several behavioral observations from our experiment predicted by this framework that appear to be inconsistent with alternative or previous models (Supplementary Figs. 3 and 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' RTs. We examined these on the basis of observations made in subjects' actual behavior. The exact same analysis was applied to both subject behavior and that of the hierarchical model.

### A within-attribute distractor effect in choice

We first assessed the hierarchical model's predictions of distractor effects, where the value of option 3 affects discriminations between options 1 and 2. We analyzed 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. 2). This was also found to be true in subject behavior (Fig. 3a).

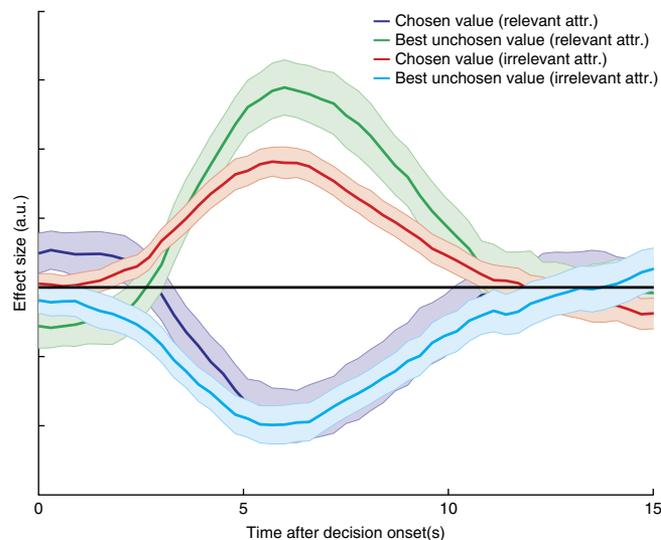
An important feature of this analysis was that it focused on just a pair of the options, including one chosen and one unchosen option at a time. By repeating this analysis on a subset of the data in which the third option, left out of the regression model, was high or low

in value, we were able to examine its effects on the discriminability of these two options<sup>23,24</sup> (third option here refers to any option that 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, discriminations between the values of the remaining two options appeared to be better than was the case when the third option had a high pO (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; stimulus:  $T(18) = 3.17$ ,  $P = 0.0053$ ; response:  $T(18) = 2.45$ ,  $P = 0.025$ ).

Notably, this effect replicates observations from a recent study of value-guided choice and can be explained in the context of divisive normalization<sup>23</sup>. However, it can also be seen in our hierarchical competition model, where divisive normalization is applied within-attribute rather than on integrated values. Moreover, the observed response pattern violates the axiom of independence of irrelevant alternatives and therefore 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 on what other alternatives are available to the subject.

We next split the third option not on the basis of its integrated value (pO), but instead on the basis of 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 normalization applied at the level of within-attribute competition rather than at the level of integrated option values.

**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 hemodynamic 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 RT included as a co-regressor of no interest). Data are presented as mean (solid lines)  $\pm$  s.e.m. (shaded areas) of parameter estimates from the regression across ten simulations of the model. The contrast (chosen value – best unchosen value)<sub>irrelevant</sub> – (chosen value – best unchosen value)<sub>relevant</sub> is encoded by the model.



By contrast, in alternative models in which integrated values ( $pO$ ) were divisively normalized before 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 3 affected subjects' ability to 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 3 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** and **Supplementary Fig. 5**). We focused on trials in which stimulus value difference on options 1 and 2 was approximately equal and opposite in sign to action value difference on options 1 and 2 (**Fig. 3d**). 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 that we observed 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 and 4** and **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 RTs. To enable this, we applied multiple regression with the logarithm of 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;  $T(18) = -9.22, P = 3.06 \times 10^{-8}$ ), whereas choices were slower when the best unchosen value increased (model:  $\beta = 1.10 \pm 0.01$ ; behavior:  $\beta = 0.31 \pm 0.07$ ;  $T(18) = 4.47, P = 2.96 \times 10^{-4}$ ). Thus, RTs decreased when ( $pO_{\text{chosen}} - pO_{\text{best unchosen}}$ ) increased (model:  $\beta = -1.91 \pm 0.02$ ; behavior:  $\beta = -1.46 \pm 0.18$ ;  $T(18) = -8.00, P = 2.44 \times 10^{-7}$ ). 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 before comparison. In our task, however, this appeared to be inconsistent with the presence of a significant within-attribute distractor effect. On the basis of the hierarchical model, we hypothesized that on trials in which attributes conflict—that is, on trials in which  $pS$  values favor one alternative, but  $pA$  values favor another—the two attributes might compete with one another when guiding the choice, with one attribute becoming more relevant for guiding behavior. Although the relevant attribute would vary from trial to trial, we could nevertheless use the subject's choice to disclose it. In trials in which  $pS_{\text{chosen}} > pS_{\text{best unchosen}}$  and  $pA_{\text{chosen}} < pA_{\text{best unchosen}}$ , we labeled 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 labeled action as relevant and stimulus as irrelevant. In the model, competition for relevance was realized by between-attribute competition (**Supplementary Fig. 2**).

The crucial question in this context is whether reward probabilities on the relevant attribute explain more variability in RT 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 (Online Methods). In the model, we found that the probability difference on the relevant attribute had a greater influence on RTs than the irrelevant attribute probability difference (**Fig. 2d**). Moreover, these two terms explained away the contribution of  $pO$  value difference to model RT (**Fig. 2d**). Notably, these predictions were not made of simpler models of evidence accumulation (**Supplementary Fig. 4**).

We applied the same analysis to subjects' RTs. Although the difference on reward probabilities on the relevant attribute influenced RTs ( $\beta = 1.46 \pm 0.31, T(18) = -4.62, P = 2.13 \times 10^{-4}$ ), 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 effect (**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

**Figure 5** IPS 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)<sub>irrelevant</sub> – (chosen value – best unchosen value)<sub>relevant</sub> 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) Time series analysis of this region, time-locked to decision phase, revealed negative correlates of (chosen value)<sub>irrelevant</sub> and positive correlates of (best unchosen value)<sub>relevant</sub>, but positive correlates of (chosen value)<sub>relevant</sub> and negative correlates of (best unchosen value)<sub>irrelevant</sub> (data are presented as mean (solid lines)  $\pm$  s.e. (shaded areas) across subjects). To avoid circular analysis, a leave-one-out cross-validation approach was used for time series extraction (Online Methods).

remained (paired  $T(18) = -2.45$ ,  $P = 0.02$ ) after normalizing all the regressors to account for any difference in their magnitude. In addition, these effects remained evident even when we considered alternative rules for constructing pO (mean, multiplication or sum; **Supplementary Table 1**).

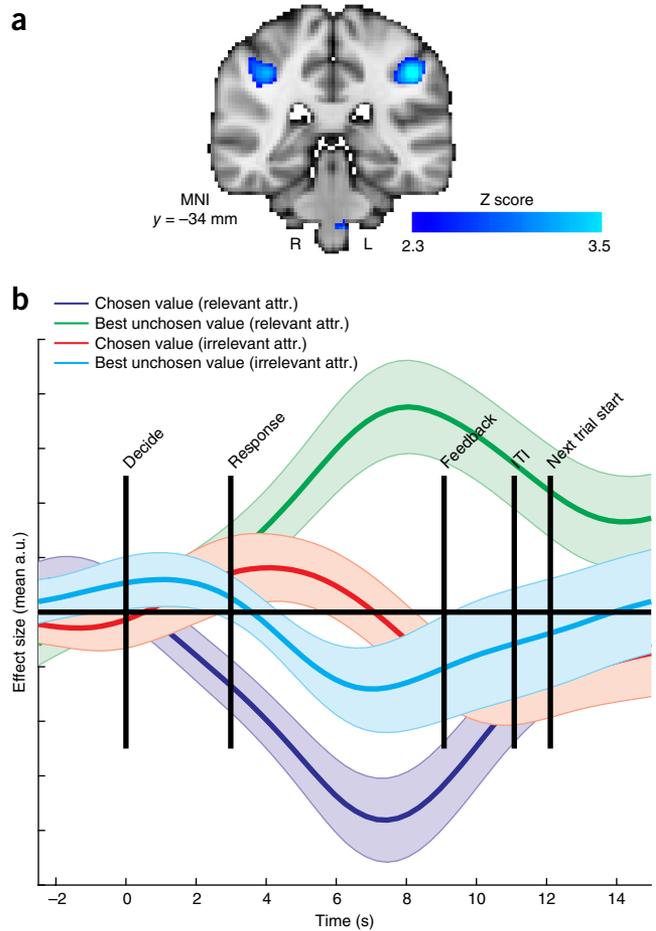
One potential final concern is that hierarchical competition could be replaced by a simpler model that retained within-attribute comparison, but randomly selected a given attribute to become relevant on each trial. We tested such a possibility by running a formal comparison on prediction of subjects'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

Our behavioral findings 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 term the relevant attribute. This is consistent with our model, in which competitions occur hierarchically.

### IPS and attribute relevance competition

We used fMRI data to investigate which brain regions subserve multi-attribute choice, tying our analysis to the modeling framework described 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 capitalized on recent reported dissociations between neural structures encoding action and stimulus values in cortical<sup>12</sup> and subcortical<sup>25,26</sup> 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 hypothesized that this region would be brought out using the contrast  $(pCh_{irrelevant} - pBUCh_{irrelevant}) - (pCh_{relevant} - pBUCh_{relevant})$ , where pCh is the chosen probability on a specific attribute and pBUCh is the best unchosen probability on a specific attribute. This contrast reflects



the difference in relative 'attribute goodness' between irrelevant and relevant attributes. Notably, it controls for areas simply encoding integrated value difference, as a result of the subtraction of relevant values from irrelevant values. Analogous to recent findings concerning value difference coding<sup>8,10,27,28</sup>, it assumes that a relevance competition signal reflects the difference of inputs that determine which attribute was more relevant (namely, within-attribute differences).

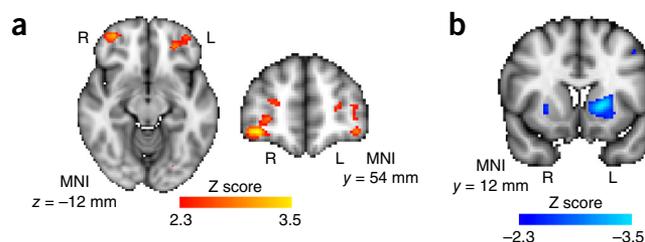
We first tested whether this contrast emerged from the attribute competition node in our hierarchical model by simulating fMRI data directly from activity in this node (**Supplementary Note**). We used similar simplifying assumptions to those used in other studies, namely that activity would be greatest in the node while competition was still being resolved and that this competition would continue until a decision was reached<sup>8,29</sup>. 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  $pCh_{relevant}$ ,  $pBUCh_{relevant}$ ,  $pCh_{irrelevant}$  and  $pBUCh_{irrelevant}$ . 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  $(pCh_{relevant} - pBUCh_{relevant})$  was predicted to have a negative regression coefficient on fMRI data, but  $(pCh_{irrelevant} - pBUCh_{irrelevant})$  was predicted to have a positive regression coefficient. This is a consequence of competition between attributes being greater on trials in which  $(pCh_{relevant} - pBUCh_{relevant})$  is closer in value to  $(pCh_{irrelevant} - pBUCh_{irrelevant})$ . 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<sup>8,29</sup>.

We therefore searched for brain regions that encoded the contrast  $(pCh_{irrelevant} - pBUnCh_{irrelevant}) - (pCh_{relevant} - pBUnCh_{relevant})$ . We found bilateral activation in the 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 and Supplementary Table 2; unthresholded  $z$  statistic maps for all contrasts are available 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 time course from IPS with a method that controls for circularity in statistical inference<sup>30</sup> by defining each subject's region of interest for analysis from the remaining  $(n - 1)$  subjects' data<sup>31</sup>. From this time course, it was apparent that all the necessary components of relevance competition were present in IPS signal. On the relevant attribute, IPS blood oxygen level-dependent (BOLD) fMRI signal correlated negatively with the value of the chosen option and positively with the value of the best unchosen option (Fig. 5b). This was shown formally by a negative effect of  $(pCh_{relevant} - pBUnCh_{relevant})$  ( $T(18) = -2.83$ ,  $P = 0.011$ ). This is similar to signals previously observed in IPS in value-guided choice tasks, in which only one attribute or integrated value was considered<sup>27,32</sup>. As predicted, however, this pattern was reversed on the irrelevant attribute: IPS signal correlated positively with the value of the chosen option and negatively with the value of the best unchosen option (Fig. 5b). This was again shown formally by a positive effect of  $(pCh_{irrelevant} - pBUnCh_{irrelevant})$  ( $T(18) = 2.50$ ,  $P = 0.022$ ). These effects are based on contrasts of parameter estimates from a general linear model (model 1, Online Methods), which partials out any covariance between irrelevant and relevant attributes. This signal profile would not be predicted for a region that compares 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  $\log(RT)$  in the same region, as is predicted by a dynamical model. Notably, however, the above effects (Fig. 5b) survived the inclusion of  $\log(RT)$  as a co-regressor (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, that is  $|(pCh_{irrelevant} - pBUnCh_{irrelevant}) - (pCh_{relevant} - pBUnCh_{relevant})|$  ( $r$  values across subjects ranged from 0.91 to 0.99). This mirrors an observation made previously in the context of value comparison<sup>8</sup>. Disambiguating this absolute contrast from our original contrast is therefore not possible with our 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 on which attribute was currently relevant. To test this hypothesis, we used a psychophysiological interaction analysis<sup>33</sup> 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 in which stimulus was relevant, IPS exhibited greater functional connectivity to a bilateral sector of anterior orbitofrontal/lateral frontopolar cortex (right peak  $Z = 3.52$ , MNI (36, 52, -10); left peak  $Z = 2.90$ , MNI (-38, 52, 10); Fig. 6a). This is a notable finding in that IPS has direct connections to OFC via the third branch of the superior longitudinal fasciculus<sup>34</sup>, and lesions to OFC are known to impair stimulus-based, but not response-based, value-guided choice<sup>12</sup>. By contrast, on trials in which action was relevant, IPS exhibited greater functional connectivity to bilateral putamen (left peak  $Z = -3.46$ , MNI



**Figure 6** Psychophysiological interaction with intraparietal cortex. (a) Functional connectivity with a bilateral portion of the anterior/lateral orbitofrontal cortex was greater on trials in which stimulus was the relevant attribute relative to trials in which action was the relevant attribute (peak  $Z = 3.52$ ; MNI = 36, 52, -10 mm (right OFC); peak  $Z = 2.91$ , MNI = -26, 42, -12 mm (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 in which action was the relevant attribute relative to trials in which 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, whereas the right putamen showed a similar, smaller activation).

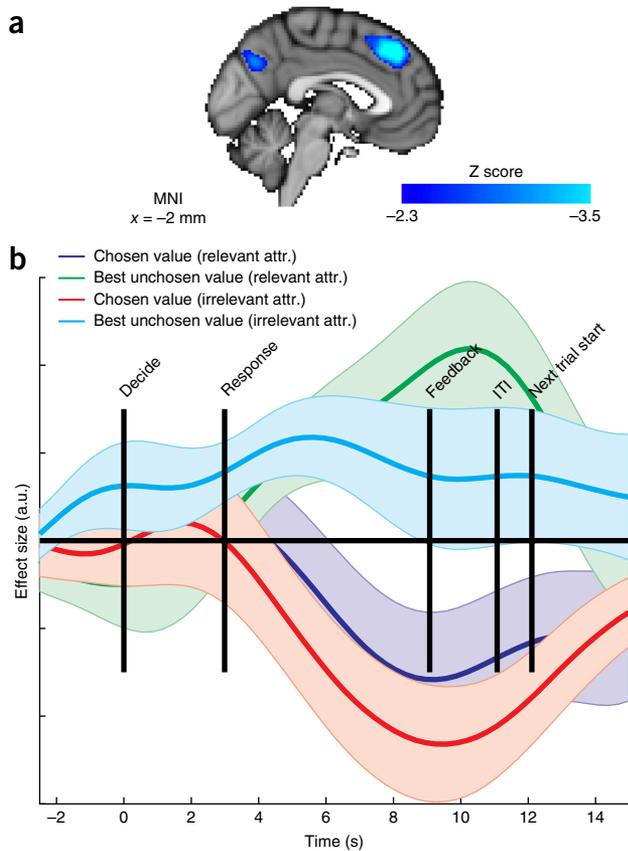
(-18, 12, 2); right peak  $Z = -2.69$ , MNI (26, 10, 0); Fig. 6b), a portion of striatum previously implicated in habitual action-guided choice<sup>25,26</sup>. The finding that both sets of activations occur bilaterally considerably reduces their probability of being false positives.

Finally, at the time of feedback, the ventral striatum encoded a classic reward prediction error, signaling 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 was encoded positively on the relevant attribute and the prediction of reward was encoded 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 a (inverted) prediction error signal suggestive of a role in relevance learning. These 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:  $(pCh_{relevant} - pBUnCh_{relevant}) + (pCh_{irrelevant} - pBUnCh_{irrelevant})$ . This contrast correlated negatively with BOLD fMRI signal in the dorsal medial frontal cortex (dmFC), in the vicinity of preSMA/paracingulate sulcus (peak  $Z = -3.85$ , MNI (0, 38, 42); Fig. 7a and Supplementary Table 3). This is in close proximity to a locus previously implicated in a comparison of integrated values during action selection<sup>8,35</sup>. Notably, this region is frequently co-activated with the IPS; both IPS and dmFC are recruited more strongly by



**Figure 7** dMFC 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 dMFC reflects this contrast (FWE-corrected,  $P = 0.0054$ , cluster-forming threshold  $Z < -2.3$ ; peak  $Z = -3.85$ , MNI =  $-2, 34, 46$  mm). Other regions surviving whole-brain correction are detailed in **Supplementary Table 3** ( $n = 19$  subjects). **(b)** Time series analysis of this region, time-locked to decision phase, revealed 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  $(\text{best unchosen value})_{\text{irrelevant}}$  (data are presented as mean  $\pm$  s.e. across subjects), slightly delayed in time relative to attribute comparison signal in IPS (compare with **Fig. 5b**). As before, cross-validated ROIs were used for time series extraction (Online Methods).

trials in which values are close together, and so correlate negatively with the value of the chosen option minus the unchosen option<sup>8,29</sup>. It is therefore particularly notable in our study that, although dMFC and IPS encode variables in the same direction on the relevant attribute, they do so in the opposite direction on the irrelevant attribute (**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<sup>36,37</sup>, but in which different studies have observed different task variables<sup>8,27</sup>. 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 items<sup>11</sup>, but it cannot explain several key behavioral observations. Among these are the effects of within-attribute similarity on choice, suggesting that some forms of comparison are enacted before integration, as well as the effects of relevance of a particular attribute modulating RTs. Subject behavior in our multi-attribute decision task showed both these effects, most notably via a previously undescribed within-attribute distractor effect.

To explain our data, we adopted a fundamentally different approach to value comparison, suggested previously in the behavioral literature<sup>16–19</sup>, but notably absent from many neuroscientific studies. In this scheme, competition (via mutual inhibition<sup>10,21</sup>) occurs not just at the higher level of integrated values, but also at lower levels (within attribute) and between levels (attention toward 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 occurs in different frames of reference. Such hierarchies mirror those proposed in the domain of rule-based action selection, or tasks requiring cognitive control<sup>38</sup>. 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 dMFC has not been described previously and it is notable that these areas have typically coactivated in studies of decision-making<sup>8,29</sup>, particularly where there is increased choice difficulty (decreased value difference). The portion of IPS that we identified is ideally placed, in terms of anatomical connectivity to prefrontal cortex<sup>34,39</sup> and its established role in attentional reorienting<sup>40</sup>, to compute which attribute to attend on a particular trial. Indeed, in number comparison tasks in which the relevant and irrelevant dimensions are explicitly cued rather than internally generated, IPS only reflects value differences in the relevant, and not the irrelevant, competition<sup>41</sup>. 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 information<sup>42</sup>.

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 as a result of 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 in which both attributes can be represented in stimulus or goal space, deficits in attribute comparison can be induced by lesions to vmPFC<sup>43</sup>. We therefore argue that attribute comparison is not a unique process with a single cortical focus, but an 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 specialization. For example, lateral OFC is likely to be involved in competition between stimuli given that it receives a highly processed sensory input<sup>44</sup> and is critical for tasks involving stimulus, but not action, comparison<sup>12</sup>. In keeping with this formulation, we observed increased functional connectivity of OFC with IPS on trials in which stimulus probabilities became relevant. By contrast, a dorsolateral portion of the striatum is likely to be involved in competition between actions, given evidence that it possesses anatomical connectivity with motoric structures<sup>45</sup> and is implicated in habitual action selection<sup>25,26</sup>. This region showed increased functional connectivity to IPS on trials in which action probabilities became relevant. Finally, both dmFC and ventral medial frontal cortex are implicated in competition between values<sup>8,10,11,27,28,35</sup>. Here, integrated value comparison signals were particularly prominent in dmFC at the time of choice and, similar to 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. One study<sup>23</sup> isolated distractor effects in subjects who chose between rewards of different value. As in our study, high value distractors impaired discrimination between two alternative options. The authors proposed a divisive normalization of values during choice. In our model, divisive normalization captured distractor effects (**Supplementary Figs. 3 and 4**), but, given the within-attribute distractor effect, we placed it at the level of attributes instead of the level of integrated values. Indeed, normalization may also reflect a canonical mechanism operating at all levels of a processing hierarchy<sup>46</sup>. Future models may unify the processes of competition and normalization in a single framework. Another study<sup>47</sup>, in a task in which decisions were 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 this task, distractors are made transiently available and then removed from the decision<sup>47</sup>. 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 found a similar (albeit weak) effect in our leaky, competing accumulator model when divisive normalization is switched off, and competition between options is high (**Supplementary Fig. 3b**). The conditions under which normalization occurs during competition remain debated<sup>48</sup>, but one hypothesis is that divisive normalization occurs over a different time course than competition via mutual inhibition, and may not occur when 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, in which introducing a new option reduces the probability of choosing options that are similar on multiple attributes to those that are dissimilar<sup>49</sup>. In our model, similar options (with small within-attribute differences) are down-weighted 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 effects<sup>17</sup>.

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 that suggest that comparison occurs within-attribute, as well as on integrated values. Our account also explains observations in which competition is observed in different frames of reference in

different brain structures, as well as how interactions between brain regions are modulated during a task<sup>50</sup>. More broadly, the findings imply that competition via mutual inhibition represents a canonical computation<sup>21,46</sup> that subserves distributed decision-making processes throughout the brain, rather than a process occurring solely at a single comparator node.

## METHODS

Methods and any associated references are available in the [online version of the paper](#).

*Note: Any Supplementary Information and Source Data files are available in the online version of the paper.*

## 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 (Sir Henry Wellcome Fellowship 098830/Z/12/Z to L.T.H., Senior Investigator Award 098362/Z/12/Z to R.J.D. and Career Development Fellowship 088312AIA to T.E.J.B.). The Wellcome Trust Centre for Neuroimaging is supported by core funding from the Wellcome Trust (091593/Z/10/Z).

## 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.

## COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

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

- Bettman, J.R. Constructive consumer choice processes. *J. Consum. Res.* **25**, 187–217 (1998).
- McFarland, D.J. & Sibly, R.M. The behavioral final common path. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* **270**, 265–293 (1975).
- Camerer, C.F. & Fehr, E. When does "economic man" dominate social behavior? *Science* **311**, 47–52 (2006).
- Keeney, R.L. & Raiffa, H. *Decisions With Multiple Objectives: Preferences And Value Functions* (Cambridge University Press, Cambridge, 1993).
- Rangel, A., Camerer, C. & Montague, P.R. A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* **9**, 545–556 (2008).
- Kable, J.W. & Glimcher, P.W. The neurobiology of decision: consensus and controversy. *Neuron* **63**, 733–745 (2009).
- Padoa-Schioppa, C. Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.* **34**, 333–359 (2011).
- Hare, T.A., Schultz, W., Camerer, C.F., O'Doherty, J.P. & Rangel, A. Transformation of stimulus value signals into motor commands during simple choice. *Proc. Natl. Acad. Sci. USA* **108**, 18120–18125 (2011).
- Lim, S.-L., O'Doherty, J.P. & 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.* **33**, 8729–8741 (2013).
- Hunt, L.T. *et al.* Mechanisms underlying cortical activity during value-guided choice. *Nat. Neurosci.* **15**, 470–476 (2012).
- FitzGerald, T.H., Seymour, B. & Dolan, R.J. The role of human orbitofrontal cortex in value comparison for incommensurable objects. *J. Neurosci.* **29**, 8388–8395 (2009).
- Rudebeck, P.H. *et al.* Frontal cortex subregions play distinct roles in choices between actions and stimuli. *J. Neurosci.* **28**, 13775–13785 (2008).
- Glöckner, A. & Betsch, T. Multiple-reason decision making based on automatic processing. *J. Exp. Psychol. Learn. Mem. Cogn.* **34**, 1055–1075 (2008).
- Tversky, A. & Simonson, I. Context-dependent preferences. *Manage. Sci.* **39**, 1179–1189 (1993).
- Payne, J.W. Task complexity and contingent processing in decision-making—information search and protocol analysis. *Organ. Behav. Hum. Perform.* **16**, 366–387 (1976).
- Vlaev, I., Chater, N., Stewart, N. & Brown, G.D. Does the brain calculate value? *Trends Cogn. Sci.* **15**, 546–554 (2011).
- Roe, R.M., Busemeyer, J.R. & Townsend, J.T. Multialternative decision field theory: a dynamic connectionist model of decision making. *Psychol. Rev.* **108**, 370–392 (2001).
- Usher, M. & McClelland, J.L. Loss aversion and inhibition in dynamical models of multialternative choice. *Psychol. Rev.* **111**, 757–769 (2004).
- Tsetsos, K., Usher, M. & Chater, N. Preference reversal in multiattribute choice. *Psychol. Rev.* **117**, 1275–1293 (2010).
- Park, S.Q., Kahnt, T., Rieskamp, J. & Heekeren, H.R. Neurobiology of value integration: when value impacts valuation. *J. Neurosci.* **31**, 9307–9314 (2011).

21. Wang, X.-J. Decision making in recurrent neuronal circuits. *Neuron* **60**, 215–234 (2008).
22. Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J.D. The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.* **113**, 700–765 (2006).
23. Louie, K., Khaw, M.W. & Glimcher, P.W. Normalization is a general neural mechanism for context-dependent decision making. *Proc. Natl. Acad. Sci. USA* **110**, 6139–6144 (2013).
24. Noonan, M.P. *et al.* Separate value comparison and learning mechanisms in macaque medial and lateral orbitofrontal cortex. *Proc. Natl. Acad. Sci. USA* **107**, 20547–20552 (2010).
25. Balleine, B.W. Neural bases of food-seeking: affect, arousal and reward in corticostriatal limbic circuits. *Physiol. Behav.* **86**, 717–730 (2005).
26. Tricomi, E., Balleine, B.W. & O'Doherty, J.P. A specific role for posterior dorsolateral striatum in human habit learning. *Eur. J. Neurosci.* **29**, 2225–2232 (2009).
27. Boorman, E.D., Behrens, T.E.J., Woolrich, M.W. & Rushworth, M.S.F. How green is the grass on the other side? Frontopolar cortex and the evidence in favor of alternative courses of action. *Neuron* **62**, 733–743 (2009).
28. Philastides, M.G., Biele, G. & Heekeren, H.R. A mechanistic account of value computation in the human brain. *Proc. Natl. Acad. Sci. USA* **107**, 9430–9435 (2010).
29. Heekeren, H.R., Marrett, S., Bandettini, P.A. & Ungerleider, L.G. A general mechanism for perceptual decision-making in the human brain. *Nature* **431**, 859–862 (2004).
30. Kriegeskorte, N., Simmons, W.K., Bellgowan, P.S. & Baker, C.I. Circular analysis in systems neuroscience: the dangers of double dipping. *Nat. Neurosci.* **12**, 535–540 (2009).
31. Boorman, E.D., Rushworth, M.F. & Behrens, T.E. Ventromedial prefrontal and anterior cingulate cortex adopt choice and default reference frames during sequential multi-alternative choice. *J. Neurosci.* **33**, 2242–2253 (2013).
32. Gershman, S.J., Pesaran, B. & Daw, N.D. Human reinforcement learning subdivides structured action spaces by learning effector-specific values. *J. Neurosci.* **29**, 13524–13531 (2009).
33. O'Reilly, J.X., Woolrich, M.W., Behrens, T.E., Smith, S.M. & Johansen-Berg, H. Tools of the trade: psychophysiological interactions and functional connectivity. *Soc. Cogn. Affect. Neurosci.* **7**, 604–609 (2012).
34. Jbabdi, S., Lehman, J.F., Haber, S.N. & Behrens, T.E. Human and monkey ventral prefrontal fibers use the same organizational principles to reach their targets: tracing versus tractography. *J. Neurosci.* **33**, 3190–3201 (2013).
35. Alexander, W.H. & Brown, J.W. Medial prefrontal cortex as an action-outcome predictor. *Nat. Neurosci.* **14**, 1338–1344 (2011).
36. Clithero, J.A. & Rangel, A. Informatic parcellation of the network involved in the computation of subjective value. *Soc. Cogn. Affect. Neurosci.* **9**, 1289–1302 (2014).
37. Levy, D.J. & Glimcher, P.W. The root of all value: a neural common currency for choice. *Curr. Opin. Neurobiol.* **22**, 1027–1038 (2012).
38. Collins, A.G. & Frank, M.J. Cognitive control over learning: creating, clustering, and generalizing task-set structure. *Psychol. Rev.* **120**, 190–229 (2013).
39. Mars, R.B. *et al.* Diffusion-weighted imaging tractography-based parcellation of the human parietal cortex and comparison with human and macaque resting-state functional connectivity. *J. Neurosci.* **31**, 4087–4100 (2011).
40. Thiel, C.M., Zilles, K. & Fink, G.R. Cerebral correlates of alerting, orienting and reorienting of visuospatial attention: an event-related fMRI study. *Neuroimage* **21**, 318–328 (2004).
41. Piazza, M., Izard, V., Pinel, P., Le Bihan, D. & Dehaene, S. Tuning curves for approximate numerosity in the human intraparietal sulcus. *Neuron* **44**, 547–555 (2004).
42. O'Reilly, J.X., Jbabdi, S., Rushworth, M.F. & Behrens, T.E. Brain systems for probabilistic and dynamic prediction: computational specificity and integration. *PLoS Biol.* **11**, e1001662 (2013).
43. Fellows, L.K. Deciding how to decide: ventromedial frontal lobe damage affects information acquisition in multi-attribute decision making. *Brain* **129**, 944–952 (2006).
44. Price, J.L. Definition of the orbital cortex in relation to specific connections with limbic and visceral structures and other cortical regions. *Ann. NY Acad. Sci.* **1121**, 54–71 (2007).
45. Haber, S.N., Fudge, J.L. & McFarland, N.R. Striatonigrostriatal pathways in primates form an ascending spiral from the shell to the dorsolateral striatum. *J. Neurosci.* **20**, 2369–2382 (2000).
46. Carandini, M. & Heeger, D.J. Normalization as a canonical neural computation. *Nat. Rev. Neurosci.* **13**, 51–62 (2012).
47. Chau, B.K., Kolling, N., Hunt, L.T., Walton, M.E. & Rushworth, M.F. A neural mechanism underlying failure of optimal choice with multiple alternatives. *Nat. Neurosci.* **17**, 463–470 (2014).
48. Rangel, A. & Clithero, J.A. Value normalization in decision making: theory and evidence. *Curr. Opin. Neurobiol.* **22**, 970–981 (2012).
49. Sjöberg, L. Choice frequency and similarity. *Scand. J. Psychol.* **18**, 103–115 (1977).
50. Hunt, L.T., Woolrich, M.W., Rushworth, M.F. & Behrens, T.E. Trial-type dependent frames of reference for value comparison. *PLoS Comput. Biol.* **9**, e1003225 (2013).

## ONLINE METHODS

**Experimental task.** 24 human volunteers (age range = 20–49 years, 13 female, 11 male, 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. Two subjects failed to reach the requisite learning criterion during training; three 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 University of College London local research ethics committee, and informed consent was obtained from all of the subjects included.

**Training.** Subjects learnt that certain stimuli and certain responses were more predictive of reward by performing pairwise choices within-attribute (that is, stimulus versus stimulus, action versus 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. In these trials, only infrequently (5% of trials) did the best  $pO$  occur for options with neither the best  $pS$  nor  $pA$ ; 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  $pA$  and  $pS$  being coincident with best  $pO$ , subjects showed no systematic bias toward 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 (Supplementary Data). Task timings are shown in Figure 1c. Subjects could view freely (that is, 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

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

**Behavioral data analysis.** Trials on which no response was made within the 3-s 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,  $\gamma$ , that was greater in subjects who used stimulus information more heavily, and a free choice parameter  $\beta$ .  $\gamma$  serves to transform the true stimulus and action probabilities of option  $i$  into subjective 'weights'

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

Each individual's  $\gamma$  and  $\beta$  parameter was fit via maximum likelihood estimation.

**Logistic regression of choices, revealing distractor effects (Figs. 2a–c and 3a–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<sup>23,24</sup>, 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 (that is, 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 data points: where the third (left-out) option was either high (>67th percentile) or low (<33rd percentile) in integrated value (Figs. 2a and 3a), high or low in stimulus probability (Figs. 2b and 3b), or high or low in action probability (Figs. 2c and 3c). The data points in these figures show the mean  $\pm$  s.e. (across subjects) of the contrasts of parameter estimates from the logistic regression.

**RT regression (Figs. 2d and 3e).** We assessed the influence of reward probabilities on RTs using multiple linear regression. Log(RT) 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 RT. Four regressors were included: the integrated value of the chosen option ( $pO_{\text{chosen}}$ ), the integrated value of the best unchosen option ( $pO_{\text{best unchosen}}$ ), the integrated value of the worst unchosen option ( $pO_{\text{worst unchosen}}$ ), and 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). Ten regressors were included in total: a constant term, the chosen, best unchosen and worst unchosen values on the relevant attribute, the chosen, best unchosen and worst

unchosen values on the irrelevant attribute, and 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 and best unchosen values on RT (that is, [1 -1] contrast on chosen and best unchosen values). Complete anonymous behavioral data sets and MATLAB analysis scripts are provided in the **Supplementary Data**.

**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 (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 modeled using a dynamical equation

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

where  $\mathbf{E}^l(t)$  is a  $3 \times 2$  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  $3 \times 2$  matrix, containing the within-attribute probabilities, divisively normalized within-attribute (that is, the columns of  $\mathbf{M}$  sum to 1);  $\mathbf{S}^l$  is a  $3 \times 3$  matrix, containing off-diagonal elements that determine competition ( $k_j$ ), and on-diagonal elements that determine the rate of self-decay ( $d_j$ );  $\mathbf{C}$  is a  $3 \times 3$  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  $3 \times 2$  matrix containing Gaussian noise, with each element drawn from the normal distribution with mean 0 and variance  $\sigma^2 dt$ ;  $\mathbf{0}$  is a  $3 \times 2$  matrix of zeros.

Accumulation of evidence integrated across attributes was then modeled using

$$\mathbf{E}^h(t + dt) = \mathbf{E}^h(t) + \mathbf{S}^h \mathbf{E}^h(t) dt + \mathbf{E}^l(t) \mathbf{w}^{-1}(t) dt; \mathbf{E}^h(0) = \mathbf{0} \quad (2)$$

where  $\mathbf{E}^h(t)$  is a  $3 \times 1$  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  $3 \times 3$  matrix, containing off-diagonal elements that determine competition ( $k_j$ ), and on-diagonal elements that determine the rate of self-decay ( $d_j$ );  $\mathbf{w}(t)$  is a  $2 \times 1$  vector, whose elements sum to 1 and are  $\geq 0$ , that determines the relative weight assigned to each lower level attribute;  $\mathbf{0}$  is a  $3 \times 1$  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 second best option on each attribute (as suggested by the signal observed in the intraparietal sulcus at the time of making the decision)

$$w_i(t) = \frac{1}{1 + e^{\beta(E_{diff_i} - E_{diff_{-i}})}} \quad (3)$$

$E_{diff_i}$  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.  $\beta$  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  $w_{-i}(t)$  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  $\theta$ . The RT 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 RTs 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 RTs—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,  $t_{nd} = 300$  ms [100:100:500]; threshold,  $\theta = 200$  [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 modeling was implemented in MATLAB (Mathworks); MATLAB code for models is available on request. For comparison with alternative models, and details of fMRI simulations, see **Supplementary Figures 3 and 4** and **Supplementary Note**.

**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 minimize dropout in both orbitofrontal cortex and amygdala<sup>51</sup>. Each volume contained 43 slices of 3-mm isotropic data; echo time = 30 ms, repetition time = 3.01 s per volume, echo spacing of 0.5 ms, slice tilt of  $-30^\circ$  ( $T > C$ ), 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 on participants' RTs). The first five volumes of each data set were discarded to account for T1 saturation effects, and so the experiment was not started until these five 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 =  $256 \times 240$ , matrix =  $256 \times 240$ , 1-mm isotropic resolution, TE = 2.48 ms, TR = 7.92 ms, flip angle  $16^\circ$ , inversion time = 910 ms. Total acquisition time was 12 min 51 s.

**Field maps.** Whole-brain field maps (3-mm isotropic) were acquired to allow for subsequent correction in geometric distortions in EPI data at high field strength. Acquisition parameters were 10-ms/12.46-ms echo times (short/long respectively), 37-ms 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, <http://www.fmrib.ox.ac.uk/fsl>)<sup>52</sup>. Region-of-interest time series analysis was performed using custom-written scripts in MATLAB (Mathworks).

**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 7.0 mm; 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 (Figs. 5 and 7, **Supplementary Figs. 6–9**, and **Supplementary Tables 2 and 3**), and model 2 to specify the psychophysiological interaction (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.

DECIDE\_ONSET: 1 during decision (from decision until response), 0 outside decision.

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.

DECIDE\_RELATT\_BEST\_UNCHPROB: As DECIDE\_RELATT\_CHPROB, containing probability of best unchosen option on the relevant attribute.

DECIDE\_RELATT\_WORST\_UNCHPROB: As DECIDE\_RELATT\_CHPROB, containing probability of worst unchosen option on the relevant attribute.

DECIDE\_IRRELATT\_CHOSEN\_PROB: As DECIDE\_RELATT\_CHPROB, containing probability of chosen option on irrelevant attribute.

DECIDE\_IRRELATT\_BEST\_UNCHPROB: As DECIDE\_RELATT\_CHPROB, containing probability of best unchosen option on irrelevant attribute.

DECIDE\_IRRELATT\_WORST\_UNCHPROB: As DECIDE\_RELATT\_CHPROB, containing probability of worst unchosen option on irrelevant attribute.

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.

RESPONSE\_ONSET: Value 1 during the 4–8-s period when the response was onscreen, 0 outside this period.

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 1 s, time-locked to response.

FEEDBACK\_ONSET: Value 1 during 3-s period when reward feedback was onscreen, 0 outside this period.

FEEDBACK\_REWARD: Same timings as FEEDBACK\_ONSET, value 1 if reward was delivered and 0 otherwise.

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.

FEEDBACK\_IRRELATT\_CHPROB: As FEEDBACK\_RELATT\_CHPROB, 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 time series. 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.

Regions encoding value difference differentially across relevant and irrelevant attributes (Fig. 5): (DECIDE\_IRRELATT\_CHPROB – DECIDE\_IRRELATT\_BEST\_UNCHPROB) – (DECIDE\_RELATT\_CHPROB – DECIDE\_RELATT\_BEST\_UNCHPROB).

Regions encoding value difference commonly across relevant and irrelevant attributes (Fig. 7): (DECIDE\_RELATT\_CHPROB – DECIDE\_RELATT\_BEST\_UNCHPROB) + (DECIDE\_IRRELATT\_CHPROB – DECIDE\_IRRELATT\_BEST\_UNCHPROB).

Regions encoding a traditional reward prediction error, tied to the relevant attribute (Supplementary Fig. 8a): (FEEDBACK\_REWARD – FEEDBACK\_RELATT\_CHPROB).

Regions encoding an 'attribute prediction error' (Supplementary Fig. 8c): (FEEDBACK\_RELATT\_CHPROB – 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 modeled other elements of the task (regressors of no interest).

DECIDE\_WENT\_W\_STIM: 1 during decision period (that is, lasting from decision onset until response), 0 outside decision period, on trials where stimulus was relevant attribute.

DECIDE\_WENT\_W\_RESP: 1 during decision period (that is, lasting from decision onset until response), 0 outside decision period, on trials where action was relevant attribute.

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.

IPS\_WENT\_W\_STIM\_INTERACTION: Interaction of DECIDE\_WENT\_W\_STIM and IPS\_TIMECOURSE; following the procedure outlined in<sup>33</sup>,

DECIDE\_WENT\_W\_STIM had zero center and IPS\_TIMECOURSE had zero mean.

IPS\_WENT\_W\_RESP\_INTERACTION: The interaction of DECIDE\_WENT\_W\_RESP and IPS\_TIMECOURSE; DECIDE\_WENT\_W\_RESP had zero center 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 DECIDE\_WENT\_W\_STIM, but on trials where neither the best stimulus nor best action was chosen), and six parametric regressors for the values of the best, second best and worst response, and best, second best and worst stimulus, all at the decision phase.

The key contrast of parameter estimates (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.<sup>7</sup>

**Region of interest analysis.** Time series for ROI plotting and PPI analyses were determined by: (i) thresholding Z statistic images from the second level analysis reported above; (ii) binarizing 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 3-mm<sup>3</sup> isotropic space in which EPI data were acquired; (iv) thresholding (at 0.3) and re-binarizing this back-projected mask; (v) extracting the mean time series within this region of interest from the pre-processed EPI data. We adopted a leave-one-out approach to ROI construction<sup>31</sup>, 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 time-locked to decision onset (Figs. 5b and 7b, and Supplementary Fig. 9) or feedback onset (Supplementary Fig. 8b,d) of each trial. This creates a data matrix with dimensions nTrials\*nTime points. Each time point 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 ref. 53.

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

IPS mask for **Figure 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 **Figure 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 8-mm sphere placed at a location consistent with findings from a recent meta-analysis of value-related activations<sup>37</sup>.

A **Supplementary Methods Checklist** is available.

51. 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* **33**, 493–504 (2006).
52. Woolrich, M.W. *et al.* Bayesian analysis of neuroimaging data in FSL. *Neuroimage* **45**, S173–S186 (2009).
53. Behrens, T.E., Hunt, L.T., Woolrich, M.W. & Rushworth, M.F. Associative learning of social value. *Nature* **456**, 245–249 (2008).