

Decision value computation in DLPFC and VMPFC adjusts to the available decision time

Peter Sokol-Hessner,¹ Cendri Hutcherson,¹ Todd Hare^{1,2} and Antonio Rangel^{1,3}

¹Division of the Humanities and Social Sciences, California Institute of Technology, 1200 E. California Blvd, Pasadena, CA 91125, USA

²Laboratory for Social and Neural Systems Research, Department of Economics, University of Zurich, Zurich, Switzerland

³Computational and Neural Systems, California Institute of Technology, Pasadena, CA, USA

Keywords: choice, finite impulse response, fMRI, hemodynamic deconvolution, information accumulation

Abstract

It is increasingly clear that simple decisions are made by computing decision values for the options under consideration, and then comparing these values to make a choice. Computational models of this process suggest that it involves the accumulation of information over time, but little is known about the temporal course of valuation in the brain. To examine this, we manipulated the available decision time and observed the consequences in the brain and behavioral correlates of choice. Participants were scanned with functional magnetic resonance imaging while they chose to eat or not eat basic food items, in two conditions differing in the amount of time provided for choice. After identifying valuation-related regions with unbiased whole-brain general linear models, we analyzed two regions of interest: ventromedial prefrontal cortex (VMPFC) and dorsolateral prefrontal cortex (DLPFC). Finite impulse response models of the upsampled estimated neural activity from those regions allowed us to examine the onset, duration and termination of decision value signals, and to compare across regions. We found evidence for the immediate onset of value computation in both regions, but an extended duration with longer decision time. However, this was not accompanied by behavioral changes in either the accuracy or determinants of choice. Finally, there was modest evidence that DLPFC computation correlated with, but lagged behind, VMPFC computation, suggesting the sharing of information across these regions. These findings have important implications for models of decision value computation and choice.

Introduction

There is a growing consensus that the brain makes simple decisions by assigning a decision value (DV) to the options under consideration, and then comparing these values to make a choice (Rangel & Hare, 2010; Padoa-Schioppa, 2011; Rushworth *et al.*, 2011). In this framework, understanding the DVs' properties is critical because they affect the choices that are eventually made. For example, if DVs are noisy when decisions are made under time pressure then so are choices (Milosavljevic *et al.*, 2010, 2011).

A considerable amount has been learned about the computational and neurobiological properties of the DV signals. A large number of studies, using different species and techniques, consistently link activity in the ventromedial prefrontal cortex (VMPFC) with the computation of DVs at the time of choice (Rangel & Hare, 2010; Wallis & Kennerley, 2010; Padoa-Schioppa, 2011). Activity in dorsolateral prefrontal cortex (DLPFC) also frequently correlates with DVs, though less consistently (Kable & Glimcher, 2007; Plassmann *et al.*, 2007, 2010; Litt *et al.*, 2010). In several decision tasks involving self-control, the DV signals in VMPFC seem to be

constructed using inputs from DLPFC areas that do not encode for DVs themselves (Hare *et al.*, 2009, 2011a; Figner *et al.*, 2010; Baumgartner *et al.*, 2011). Finally, recent studies have shown that similar areas of VMPFC and DLPFC are part of the network involved in comparing DVs to make a choice (Basten *et al.*, 2010; Hare *et al.*, 2011b).

Although these findings have established the critical role of VMPFC and DLPFC in the computation of DVs at the time of choice, the nature of the DV signals encoded, as well as how they relate to choices, have not been systematically investigated. Four questions are of particular interest. First, does the available decision time affect the duration and quality of decision value signals in VMPFC and DLPFC? Second, do the changes in the DV signals have an effect on the quality of choices? Third, are there differences between regions in how decision value signals evolve over time within decision episodes? And fourth, are there differences in the stimulus information that is used to compute the decision value signals in VMPFC and DLPFC (e.g., taste vs. health attributes for foods)?

The answers to these questions are important because they have direct bearing on the properties and quality of the decision-making circuitry, and how it is affected by environmental variables such as time pressure or distraction. For example, it is not known whether the brain keeps refining its estimates of DV if given extra decision time or

Correspondence: Peter Sokol-Hessner, as above.
E-mail: psh@caltech.edu

Received 25 November 2011, revised 9 February 2012, accepted 12 February 2012

whether it stops computing them as soon as the estimate is satisfactory, even if extra time remains (Kiani *et al.*, 2008). In the former case extra decision time could have a positive effect on the quality of choices, but in the latter it is just wasted time. The relationship between the two DV signals is also poorly understood. Some have argued that VMPFC signals serve as precursors to DLPFC signals, which eventually drive choices (Wallis & Miller, 2003). In contrast, in other decision paradigms the influence appears to go in the other direction (Hare *et al.*, 2009; Baumgartner *et al.*, 2011).

Here we describe the results of a functional magnetic resonance imaging (fMRI) experiment in which participants were asked to make consumption decisions about snack foods under two externally controlled decision speeds: fast (1 s) and slow (4 s) choices. The results of the experiment allow us to systematically test specific hypotheses associated with each of these three questions (see Results section for details).

Materials and methods

Participants

Twenty-eight participants completed both sessions of the study (eight females; mean \pm SD age, 22.4 ± 0.6 years). Of these, six participants were excluded from data analysis: one due to scanner malfunction during the study, two for excess head motion during scanning (total translation > 3 mm or total rotation $> 2.45^\circ$ in any single volume), and three due to highly atypical behavioral responses that suggested that they were not taking the task seriously (i.e. responding without considering the stimuli, as demonstrated by having identical consecutive responses in more than 50% of trials). Twenty-two participants were used in the analyses (six females, mean \pm SD age, 22.6 ± 0.6 years). All participants gave written informed consent, and all procedures were approved by Caltech's Institutional Review Board.

Task

The experiment took place over two sessions, separated by 1–4 days. On day 1, we collected basic subject-specific information about the stimuli used in the experiment. On day 2, participants performed an in-scanner choice task.

Day 1

For the first behavior-only session, participants were asked to fast for at least 4 h prior to the experiment. Participation in the task was contingent on signing a statement affirming that they had not eaten or

drunk anything except water for the past 4 h, and reporting no food allergies or intolerances. During this session, participants rated 249 color images of different appetitive and aversive snack food items (e.g. Kit-Kat bar, apple slices, Spam; see Supporting Information for a complete list) on three different five-point scales: Liking (Strongly Dislike to Strongly Like), Tastiness (Very Bad to Very Good) and Healthiness (Very Unhealthy to Very Healthy). Ratings were blocked, so that all ratings for a given type were completed before moving on to ratings for a different type. Liking ratings for all foods were assessed first, randomly followed by ratings of Healthiness and Tastiness (order counterbalanced across participants).

On every trial, participants were shown one of the food items for 1 s, with a white box around it, during which they could not respond. Immediately after the box disappeared, participants entered ratings at their own pace. Their response was followed by a 1-s feedback screen displaying the rating they had selected. Trials were separated by a uniform random 1- to 3-s screen with a centered fixation cross. Ratings scales were displayed and explained before each block, and the left–right order of each rating type was separately randomized across participants.

Day 2

The same 4-h fasting and screening criteria were implemented in the second session. After detailed instruction, participants performed a simple choice task in the scanner. Every trial they were shown one food item and had to decide whether they wanted to eat it at the end of the experiment. They indicated their choices using a four-point scale [Strong Yes (SY), Yes (Y), No (N), Strong No (SN); left–right ordering randomized across participants], which allowed us to measure simultaneously their decisions and the strength of their preferences. Participants cared about their decisions because they knew that they would have to stay in the lab for 20 min at the end of the experiment, and that the only thing that they would be allowed to eat would be based on their decisions. Specifically, one trial was randomly selected and participants were required to eat that food item if their response was SY or Y, and were not allowed to eat it otherwise.

The structure of the decision task was similar to the ratings task (Fig. 1). The only differences were as follows: (i) the initial 'viewing period' of the foods now lasted either 1 s (Short condition) or 4 s (Long condition); (ii) participants had to indicate their decision within 1 s of the beginning of the response period or the trial aborted; and (iii) the intertrial intervals were drawn uniformly from 1 to 6 s. To maintain the spacing across trials, the unused fraction of the response period was added to the intertrial interval. Stimuli were divided between the Short and Long conditions based on each participant's liking ratings in the behavioral session, to equate the distribution of

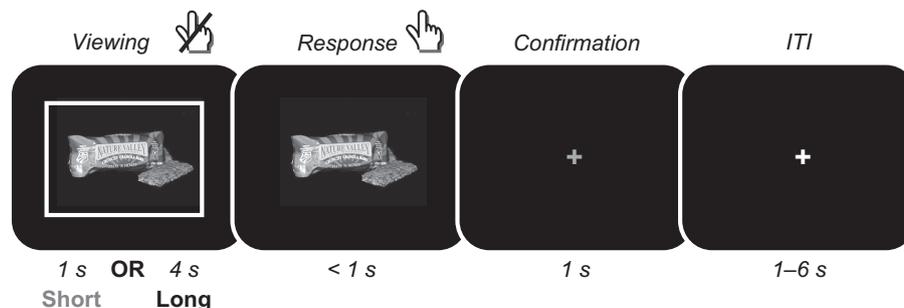


FIG. 1. Scanner choice task. On each trial, a color image of a food item appeared surrounded by a white box for either 1 s (Short condition) or 4 s (Long condition). During this time participants were not allowed to enter a response. Afterwards the white box disappeared and participants had up to 1 s to enter their choice on a 4-point scale [Strong Yes (SY), Yes (Y), No (N) or Strong No (SN)]. Immediately after the response, a 1-s confirmation screen appeared with a green fixation cross. Trials were separated by an intertrial interval (ITI) of 1–6 s random duration.

Liking ratings across conditions. Conditions were blocked (25 trials per block), and block order was randomized subject to the constraint that there were two blocks of each type (Short; Long) during each functional scan. Blocks were preceded with an instruction cue indicating the duration of the viewing period in the upcoming block of trials (shown for 5 s, followed by a 1-s fixation cross) informing participants about the upcoming condition. There were 200 trials in each condition (400 trials total; eight blocks of each trial type, 16 blocks total). See the Supporting Information for details on how foods were distributed between conditions.

fMRI – data collection

Scanning was performed at the Caltech Brain Imaging Center with a Siemens 3T TIM-Trio full body scanner and a Siemens 32-channel phased array head coil. High-resolution anatomical images were acquired using a T1-weighted protocol (FOV, 256; 176 slices, $1 \times 1 \times 1$ mm). Functional imaging used a gradient echo EPI sequence (TR, 2530 ms; TE, 30 ms; FOV, 192; anterior-to-posterior phase encoding, ascending slice acquisition), acquiring four functional runs of 334 volumes, each with 40 oblique axial slices aligned 30° off the AC–PC plane (to improve signal in the orbitofrontal cortex; Deichmann *et al.*, 2003), 3-mm isometric voxels and a 0.3-mm between-slice gap.

fMRI – preprocessing

Two volumes were discarded before the beginning of data collection in each run to allow for equilibration of the magnetic field. Data were preprocessed with SPM8 software (Statistical Parametric Mapping 8, Wellcome Trust Center for Neuroimaging; <http://www.fil.ion.ucl.ac.uk/spm/>), including slice-time correction, realignment (motion correction) and spatial smoothing (isotropic 6-mm FWHM Gaussian kernel). Functional runs were coregistered and normalized to the standard Montreal Neurological Institute EPI template. Data were high-pass filtered prior to analysis (cutoff, 128 s).

fMRI – whole-brain analyses

Analysis was performed with SPM8 and custom MATLAB scripts (Mathworks, Natick, MA, USA). Individual level whole-brain general linear models (GLMs) with first order autoregression (AR(1)) and SPM8's standard hemodynamic response function were estimated in three steps. First, we estimated the model separately for each individual. Second, we calculated contrast statistics at the individual level. Third, we computed second-level statistics by carrying out one-sample *t*-tests on the single-subject contrast coefficients.

GLM 1

This model contained the following three regressors of interest: R1, regressor for initial stimulus presentation, in the form of a boxcar function that ran from image onset to response, pooling across both the Short and Long conditions; R2, Decision Value regressor, created by modulating R1 by the subject's decision value on each trial (SN, -1.5 ; N, -0.5 ; Y, $+0.5$, SY, $+1.5$); and R3, Response Time regressor, in the form of a stick function at the time of the response.

GLM 2

The second GLM consisted of the following regressors of interest: R1, regressor for Short trials only, in the form of a boxcar from image

onset to response; R2, Tastiness rating regressor for Short trials, created by modulating R1 by the Tastiness rating (Very Bad, -2 to Very Good, $+2$); R3, Healthiness rating regressor for Short trials, created by modulating R1 by the Healthiness rating (Very Unhealthy, -2 to Very Healthy, $+2$); R4, regressor for Long trials only, in the form of a boxcar from image onset to response; R5, Tastiness rating modulated regressor for Long trials; R6, Healthiness rating modulated regressor for Long trials; R7, Indicator for response in Short trials; and R8, Indicator for response in Long trials.

In addition, both GLMs also included the following nuisance regressors – estimates of motion from preprocessing and, separately for each condition, boxcars for the pre-block instruction cues (5 s), transient predictors for item presentation when participants failed to respond, transient predictors for item presentation when participants responded too early, and transient predictors for the response itself when participants responded too early.

In our whole-brain analysis we localize activity for the contrasts of interest using the liberal but common statistical threshold $P < 0.001$ uncorrected with a five contiguous voxel extent threshold. We chose this criterion to give the best chance of identifying the desired regions of interest (ROIs) for further investigation.

fMRI – ROI definition

Most analyses in the paper examined the patterns of activation in two pre-specified ROIs: the VMPFC and left DLPFC regions exhibiting blood oxygenation level-dependent (BOLD) responses modulated by DVs (R2) in GLM 1. The resulting functionally specified ROIs are depicted in Fig. 2. The VMPFC ROI consisted of all voxels significant in this contrast at $P < 0.00001$, uncorrected (78 voxels). This stringent threshold was necessary to prevent the inclusion of voxels from other distinct areas of the brain (e.g. ventral striatum, cingulate). The DLPFC ROI consisted of all voxels significant at $P < 0.001$, uncorrected (seven voxels). After identifying ROIs in whole-brain analyses, the responses within these regions were investigated in two ways: (i) raw BOLD data was extracted from these ROIs for deconvolved (neural signal) finite impulse response (FIR) analyses, and (ii) mean parameters within the ROIs from the whole-brain GLM 2 (Taste and Health) were extracted for use in conventional BOLD analyses.

fMRI – neural estimate analyses

Although standard approaches often present analyses based on BOLD responses, hemodynamic responses are substantially delayed and prolonged relative to the neural signals that generate them. This means that analyses sensitive to small differences in timing (~ 1 – 3 s) can often be difficult to visualize or interpret using these approaches. To better address questions specific to the timing of value signals, we therefore turned to deconvolution methods designed to extract the 'neural estimate', i.e. an estimated neural response (Gitelman *et al.*, 2003). This deconvolution procedure offers the additional benefit of using a constrained set of hemodynamic response function (HRF) parameters rather than a single fixed HRF. This accounts for some of the differences that may arise between regions, lowering the risk of systematic bias and increasing our ability to compare the neural estimates across regions (see Supporting Information for similar analyses performed on the non-deconvolved BOLD responses).

This analysis proceeded in the following steps.

First, for every voxel within an ROI, the raw BOLD signal at each TR was extracted, mean-corrected and adjusted for motion using standard SPM tools.

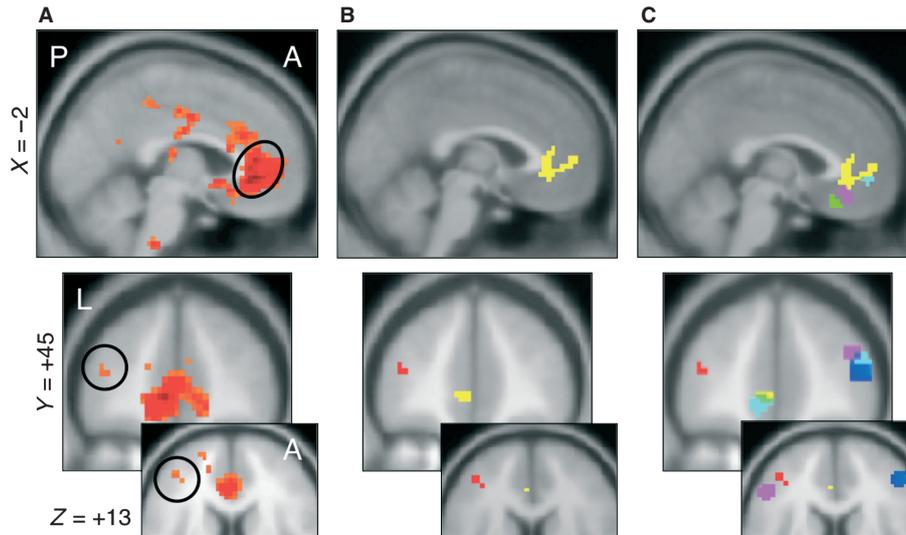


FIG. 2. Whole-brain localization of DV signals and ROIs. (A) Areas in which BOLD responses during the initial period of stimulus presentation were positively modulated by decision values, regardless of condition ($P < 0.001$ uncorrected, cluster threshold = 5; see GLM 1). Hot colors indicate positive contrast values, cold colors indicate negative contrast values. The circles identify the VMPFC (top) and DLPFC (bottom) regions correlated with decision value. (B) ROIs used in further analyses – VMPFC, yellow; DLPFC, red. (C) Relationships between the ROIs and areas in VMPFC and DLPFC that have been shown elsewhere to correlate with DVs at the time of choice – Plassmann *et al.* (2010), left DLPFC, purple; Plassmann *et al.* (2007), right DLPFC, blue; McClure *et al.* (2004), VMPFC, cyan; and Hare *et al.* (2009), VMPFC, green.

Second, the resulting timecourses for each individual's extracted data were then subjected to singular value decomposition, and the first eigenvariate (i.e. the first principal component of the raw time series from all voxels within the ROI) was used as the index of the regional response (Friston *et al.*, 1996; O'Reilly *et al.*, 2010; Mars *et al.*, 2011a,b). This technique captures the signal representing the greatest proportion of variance within the ROI (i.e. that which is most coherent across the voxels), making it less susceptible to noise from any given voxel compared to a simple average (Friston *et al.*, 1996; Gitelman *et al.*, 2003).

Third, a neural estimate was constructed by deconvolving the BOLD timecourse with the formula for SPM8's canonical HRF using standard SPM scripts employing a parametric empirical Bayes formulation (Gitelman *et al.*, 2003). This deconvolution entailed assumptions common to all analyses using HRFs (including classic whole-brain GLMs), such as assumptions of HRF linearity, shape and consistency across regions. As in any analyses requiring these or similar assumptions, any inferences are predicated on the accuracy of those assumptions.

Fourth, the neural estimate resulting from this deconvolution was linearly interpolated (up-sampled) to a resolution of 1/16th of a TR (0.158 s). This fine temporal resolution is preferred in order to reduce the temporal blurring associated with rounding event onset times to the nearest data timepoint in the finite impulse response analyses described below.

Finally, we estimated FIR models on these data for each individual, ROI and scan session. In an FIR model, events of interest are modeled with multiple regressors, each estimating the signal at a single given timepoint relative to event onset. Note that because this technique estimates the signal at discrete timepoints, regressors must be similarly aligned, incurring temporal blurring proportional to the timecourse resolution. Events of interest were modeled with 19 sequential non-overlapping transient predictors in the Short (1 s) condition (covering the 3 s from the onset of the image to the end of the trial), and 38 transient predictors in the Long (4 s) condition (to cover the 6 s from image onset to the end of the trial). The resulting models included

separate regressors for each response type (SY, Y, N and SN) and a session constant. After estimation, parameter estimates were averaged across the four scan sessions for each participant. Because the hemodynamic response was deconvolved from the ROI data, no hemodynamic lag was expected in the neural estimates. One participant was excluded from the FIR analyses because they never provided a SN response.

We report three different types of FIR analyses. First, we examine the instantaneous differences in the FIR estimates between SY and SN trials. This provides a measure of the instantaneous DV signal (i.e., the difference score increases as the neural responses associated with DV signals become more discriminating between SY and SN trials). Second, because computational models of choice behavior suggest that decisions result from the accumulation of information over time (Kiani *et al.*, 2008; Ratcliff & McKoon, 2008; Kiani & Shadlen, 2009), we also looked at differences in the cumulative value signal. In particular, for any time t , let $v(t)$ be the SY – SN difference score at that time. The cumulative value difference at time T is then given by the sum of all $v(t)$ from $t = 0$ to $t = T$. This analysis allows us to test whether there are differences across conditions or regions in how the DV signals accumulate over time. Third, we examined differences in the timing of valuation by (i) computing the time of the peak value computation (t at which $SY - SN_t = \max [SY - SN]$), and (ii) computing the number of timepoints we had to shift the Short SY – SN curve to best match the Long SY – SN curve, by minimizing the sum of the squared differences at the overlapping points (intuitively, the point at which the two curves were the most similar).

fMRI – BOLD ROI analyses

Using the VMPFC and DLPFC ROIs, we also conducted a standard analysis based on the BOLD signal. Using GLM 2, we extracted average beta coefficients for the weighting of Tastiness and Healthiness in the Short and Long conditions from each ROI for each subject, and subjected them to a detailed analysis examining the

influence of decision time on the neural representation of each attribute.

Results

This section is organized as follows. First, we describe the results of a whole-brain analysis of the BOLD data designed to identify areas of VMPFC and DLPFC in which BOLD responses were correlated with DVs, for both experimental conditions. The results of this analysis were used to define two functional ROIs, one for VMPFC and one for DLPFC, which were found to be similar to those identified in previous studies. Second, we carried out various ROI-based analyses of the properties of the DV signals encoded in these two areas; these analyses were used to systematically address the four questions posed in the introduction.

Whole-brain analysis – correlation with DVs during the evaluation period

We estimated a general linear model of brain responses in which activity during the entire evaluation period (i.e., from stimulus onset to response prompt, 1 s in Short condition and 4 s in Long condition) was modulated by DVs (see GLM 1 in Materials and Methods for details). DVs were measured using the four-point response scale (SN, N, Y, SY). Figure 2A depicts regions in which the BOLD responses were positively correlated with DVs which, as expected, included areas of VMPFC and left DLPFC (see Table S1 for full list of activations).

These results were used to functionally define the ROIs in VMPFC and DLPFC, depicted in Fig. 2B, that were used in all subsequent analyses. Note that their location aligned well with previous reports (Fig. 2C; McClure *et al.*, 2004; Plassmann *et al.*, 2007; Hare *et al.*, 2008; Plassmann *et al.*, 2010) and that their definition was statistically independent of tests below that compare the two conditions, as GLM 1 imposes a common estimate for the DV regressor for both conditions.

Question 1 – does available decision time influence the timing of DV signals?

Our primary research question concerned whether the amount of available decision time influenced the timing of decision value computations in the brain (e.g., changing the onset, duration or trajectory, as compared to the null hypothesis that time has no influence on these computations). Evidence for such changes would help to explain behavioral findings that time pressure can lead to less accurate choices (Milosavljevic *et al.*, 2010, 2011) and might complement evidence from the perceptual literature that evidence accumulation begins immediately but continues only until a decision threshold is crossed, even if more time and information are available (Kiani *et al.*, 2008; Kiani & Shadlen, 2009). This evidence suggests that we should observe changes in the timing of subjective decision value computation as a function of available decision time, although the precise nature of these changes could take several forms, three of which we test below.

We tested this hypothesis using extracted neural estimates of value-related activity in VMPFC and DLPFC and FIR analyses that estimate neural responses in each ROI, time bin and condition separately for each response type (SY, Y, N, and SN). This combination of methods allows us to compare signals, without lag, across regions because it does not assume the same shape for the hemodynamic response in

each area (see Materials and Methods for full details and justification). For analyses reported in the paper, we defined an index of the instantaneous DV signal in each time bin by the difference between the SY and SN responses (Fig. 3A and B; see Fig. S1 for separate SY, Y, N and SN responses). The larger the SY – SN score (i.e. the more strongly items were differentiated based on value), the more value computation was taking place. Intuitively, an SY – SN difference score of zero at a given moment indicates no value computation at that time. As the SY – SN difference score becomes significantly different from zero, it indicates a stronger and more consistent representation of stimulus values.

We first tested for differences in onset of DV computations in Short vs. Long trials. We found that in both ROIs the onset of value computation appeared to be immediate and of identical magnitude. In other words, the magnitude of the response during the first second was significantly above zero at all timepoints, and did not differ between Short and Long trials in either region (all timepoints n.s.).

We next compared differences in the duration of DV computation. This comparison revealed significantly longer DV computation in the Long condition, in both regions. In other words, whereas in the Short condition the DV signal became statistically indistinguishable from zero ~2 s after stimulus onset (DLPFC timepoint 15, 2.37 s; VMPFC timepoint 12, 1.89 s), in the Long condition the DV signal was present for nearly twice as long (DLPFC timepoint 26, 4.11 s; VMPFC timepoint 25, 3.95 s). Note that although the termination of computation aligned to the end of the viewing window in the Long condition, it extended a full second past this point in the Short condition. To test these observations of duration differences more

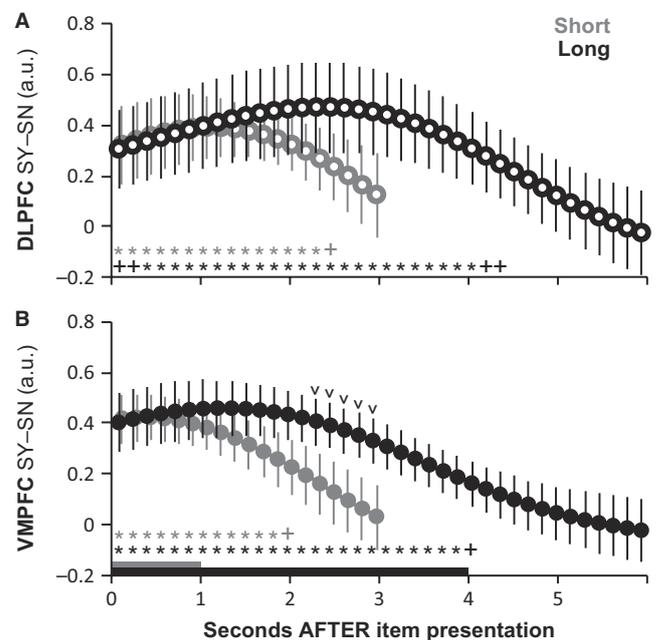


FIG. 3. Instantaneous neural estimate analyses. (A) Estimate of the difference in activity levels in the DLPFC ROI between trials with SY and SN responses, for each time bin and condition. Gray and black bars at the bottom of the figure indicate the length of time the foods were initially shown in each condition. Gray and black symbols indicate that SY – SN was significantly different from zero at $*P < 0.05$ and marginally different at $^+P < 0.1$, based on a one-sample t -test computed against zero. (B) Analogous estimates for the VMPFC ROI. Black carets indicate a marginal difference between the estimates for the two conditions ($^{\vee}P < 0.1$, paired t -test). Error bars are SEM calculated across participants. See Materials and Methods and Results sections for details on how the measures are computed; y-axes are in arbitrary units (a.u.).

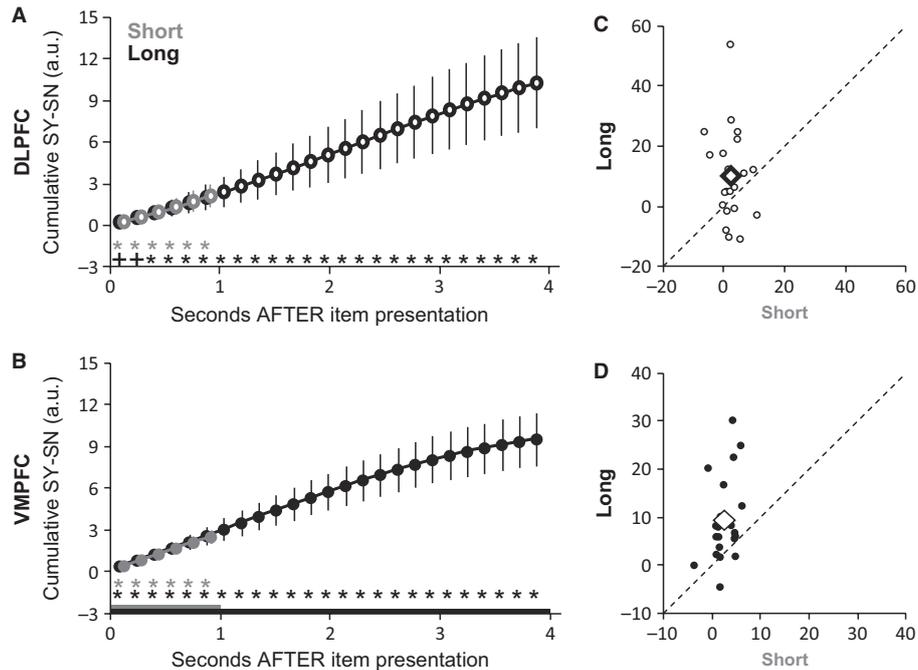


FIG. 4. Cumulative neural estimate analyses. (A) Estimate of the cumulative difference in activity levels in the DLPFC ROI between trials with SY and SN responses, for each time bin and condition. The cumulative level of activity for any time bin t , subject and condition was computed by adding the instantaneous activity levels depicted in Fig. 3, from bin 0 to bin t . Gray and black bars at the bottom of the figure indicate the length of time the foods were initially shown in each condition. Gray and black symbols indicate that the cumulative SY-SN difference was significantly different from zero at $*P < 0.05$ (or marginally different at $^*P < 0.1$), based on a one-sample t -test computed against zero. Note that the two curves lie almost on top of each other. (B) Analogous estimates for the VMPFC ROI. Error bars are SEM calculated across participants. (C) Individual estimates of the cumulative difference in DLPFC activity between trials with SY and SN responses, measured at the end of the initial non-response viewing period for each condition. Each circle denotes one subject. Black diamonds are the group mean. The dashed line indicates the line of identity. (D) Analogous plot for the VMPFC ROI. A paired t -test across subjects indicates that in both cases the cumulative value signal at the end of initial item presentation was greater in the Long condition (DLPFC, $P = 0.04$; VMPFC, $P = 0.0009$; $N = 21$); y-axes are in arbitrary units (a.u.).

formally, we defined a measure of the cumulative value computation during the full viewing window. For any time bin T , the measure was constructed for each subject by adding up the instantaneous estimates reported in Fig. 3 from all bins between 0 and T , which is basically a measure of the area under the instantaneous curve (Fig. 4A; see Materials and Methods for full details). Larger cumulative value scores indicate a longer duration of value computation. To compare conditions within each ROI, we computed the value of the cumulative signal at the end of the stimulus presentation period (1 s for Short, 4 s for Long). In both regions, the mean cumulative signal at the end of the period was larger in the Long condition (VMPFC, $t_{20} = 3.9$, $P = 0.0009$; DLPFC, $t_{20} = 2.2$, $P = 0.04$), demonstrating that longer viewing windows resulted in a significantly longer duration of value integration.

Taken together, this set of results indicates the following robust properties of the DV signals in VMPFC and DLPFC. First, the DV signals were indistinguishable across conditions for the first second of computation, when the stimulus was present in both cases. Second, computation began to differ after that – the signal ramped down in the Short condition, but remained above zero for significantly longer in the long condition.

Question 2 – does extra decision time result in more accurate choices?

The computational models of choice that state that decision values must be computed and compared to make a choice suggest that the additional computation time we observed in VMPFC and DLPFC should result in more accurate choices. To determine whether this was the case, we tested for differences either in average DVs or in the

TABLE 1. Mixed effect estimates for a linear regression of choices on liking ratings.

	Mean estimated coefficients		R^2
	Constant	Liking rating	
Short	-0.111 ($P = 0.048$)	0.481 ($P = 1.6 \times 10^{-15}$)	0.57
Long	-0.107 ($P = 0.046$)	0.494 ($P = 5.4 \times 10^{-15}$)	0.56
Short vs. Long*	$P = 0.81$	$P = 0.36$	$P = 0.80$

For each individual and condition, the decisions made in the scanner (represented by their DVs; see Methods) were regressed against the liking ratings provided in the first experimental session. The table reports the mean estimated coefficients of the regression, as well as P -values for two-tailed groupwise one-sample t -tests against the null hypothesis that the mean estimated coefficient is zero ($N = 22$).

*Two-tailed paired sample t -test.

relationship between DVs and the temporally unconstrained Liking ratings (mean \pm SEM, 0.12 ± 0.08) made during Session 1. Differences in the averages, or lower explained variance in the Short vs. Long conditions, would indicate reduced consistency or accuracy in choices. However, we observed no evidence of changes in either measure. The mean decision values and the variance explained by Liking ratings did not differ across conditions (mean DV Short, -0.04 ± 0.06 ; mean DV Long, -0.03 ± 0.06 ; $R^2_{\text{Short}} = 0.57$, $R^2_{\text{Long}} = 0.56$; see Table 1). To ensure that lack of differences did not result from extra time taken by participants during the response

period we also examined reaction times. If decisions are not computed fully within 1 s, participants may use more of the response period to achieve greater accuracy. However, participants were slightly faster in the Short condition than the Long condition (average participant median \pm SE reaction times – in Short, 345 ms \pm 12 ms; in Long, 376 \pm 8 ms; $t_{21} = 3.56$, $P = 0.002$). Explanations for this small but counterintuitive difference could be the release of response inhibition in the Long condition or the greater predictability of response onset in the Short condition. Future research might seek to examine the source of this discrepancy and its consequences for the control of valuation. [We also observed that extreme (SY and SN) responses were faster than weak (Yes and No) responses – average participant median \pm SE reaction times for extreme responses, 351 \pm 9 ms; for weak responses, 370 \pm 9 ms; $t_{21} = 3.9$, $P = 0.0008$), but this effect did not differ by condition.] Regardless, these results suggest that choice accuracy was identical in the Short and Long conditions, and at ceiling. Despite evidence of neural differences, a computation time of 1 s was enough to saturate the value estimates driving behavior.

Question 3 – do decision value signals evolve differently over time in VMPFC vs. DLPFC?

Although we observed some differences in the temporal course of DV signals in both VMPFC and DLPFC, it is also possible that there are differences between these two regions specifically in their evolution over time. This hypothesis was motivated by the informal observation that presentation times of < 2 s were sufficient to reliably engage VMPFC, but such short times led to inconsistent findings of DLPFC (Kim *et al.*, 2008; Litt *et al.*, 2010; Harris *et al.*, 2011), compared to a presentation time of 4 s which tended to engage both (McClure *et al.*, 2004; Plassmann *et al.*, 2007; Hare *et al.*, 2008, 2009).

We tested the hypothesis that these regions display differences in the timing of their value computation signals (which could explain these observations) against the null hypothesis that they proceed identically and in parallel. To do so, we performed two closely related analyses: (i) identification of the peak timepoint of individuals' SY – SN curves in each region of interest and condition; and (ii) computing the best fit shift for each region of the Short SY – SN curve to match the Long curve (intuitively, the point at which the two curves were the most similar).

The identification of peaks in the SY – SN score serve as a crude measure of the relative timing of a response. This analysis indicated no significant difference in the timing of the peaks in the Short condition (DLPFC peak, 9.8; VMPFC peak, 8.8 timepoints). However, in the Long condition, the DLPFC (peak = 21.7) peaked significantly later than the VMPFC (peak = 13.4 timepoints; $t(20) = 2.7$, $P = 0.01$).

To identify the shift, we minimized the sum of squared differences between the overlapping points of the Short and Long curves for each region. This yielded a shift of anywhere from zero (no shift) to 19 timepoints (3 s), which averaged 11.8 timepoints (1.87 s) in the DLPFC and 8.5 timepoints (1.34 s) in the VMPFC. This corroborates the peak analysis, suggesting that in the Long condition the DLPFC response was shifted later compared to VMPFC, relative to the respective original responses in the 1-s condition. However, estimates of this shift were highly non-normal, bounded, and quantized, violating common fundamental assumptions of parametric tests, for which reason we limit these results to descriptive statistics.

In concert, these findings suggest subtle but significant differences in the timing of computations performed in VMPFC and DLPFC. Given this, we turned next to an investigation of differences in the content of those value signals.

Question 4 – are there differences in the information that is integrated to compute decision value signals in VMPFC and DLPFC?

Computational models posit that decision values are the result of the accumulation and integration of multiple sources of information (e.g. different stimulus attributes; Ratcliff & McKoon, 2008). Recent work suggests that value computations in VMPFC may be biased toward low-level primary attributes of a stimulus, such as Tastiness, and that DLPFC may be required to incorporate more abstract attributes, such as Healthiness, into decisions (McClure *et al.*, 2004; Hare *et al.*, 2009). A natural hypothesis based on the above observations is that the differences we observed in timing may result from difference in content.

Before examining neural differences, we asked whether computation time affects the determinants of value signals driving decisions as observed behaviorally. We regressed each participant's choices (DVs; see Materials and Methods) on their ratings of Tastiness and Healthiness of those same items, separately for the Short condition and the Long condition. If these attributes differentially predict choice as a function of time we should see differences in the parameters across the two conditions. However, we observed little evidence of such changes. Across participants, mean \pm SEM for ratings of the food items were 0.24 \pm 0.08 for Taste and -0.44 ± 0.06 for Health. In both conditions, Tastiness was the primary driver of choice ($\beta_{\text{Short}} = 0.474$, $\beta_{\text{Long}} = 0.491$; both $P < 0.01$), while Healthiness played a much smaller role ($\beta_{\text{Short}} = 0.033$, $\beta_{\text{Long}} = 0.056$). There were no significant differences between the Short and Long conditions for any of the parameters of interest (Table 2), although Healthiness did play a marginally significant role in choices made in the Long condition ($P = 0.063$) but not in the Short condition ($P = 0.18$). Individual-level parameter significance tests also revealed no consistent patterns across conditions (not shown). In summary, there was no strong behavioral evidence that different attributes of the stimuli were driving decision values in the Short vs. Long conditions.

Although we did not observe differences in the behavioral weighting of these two properties, it is still possible that different regions may differentially represent these components of value, or that such a representation interacts with the available decision time. We thus tested for neural differences in the representation of Tastiness and Healthiness in VMPFC and DLPFC by condition. To do this, we ran a second GLM regressing Tastiness and Healthiness against the BOLD response separately in the Short and Long conditions (see Materials

TABLE 2. Mixed effect estimates for a linear regression of choices on health and taste ratings.

	Mean estimated coefficients			R^2
	Constant	Health rating	Taste rating	
Short	-0.158 ($P = 0.0016$)	0.033 ($P = 0.18$)	0.474 ($P = 2.7 \times 10^{-12}$)	0.54
Long	-0.155 ($P = 0.0012$)	0.056 ($P = 0.063$)	0.491 ($P = 6.8 \times 10^{-12}$)	0.56
Short vs. Long*	$P = 0.86$	$P = 0.11$	$P = 0.23$	$P = 0.42$

For each individual and condition, the decisions made in the scanner (represented by DVs; see Methods) were regressed against the Taste and Health ratings provided in the first experimental session. The table reports the mean estimated coefficients of the regression, as well as P -values for two-tailed groupwise one-sample t -tests against the null hypothesis that the mean estimated coefficient is zero ($N = 22$).

*Two-tailed paired sample t -test.

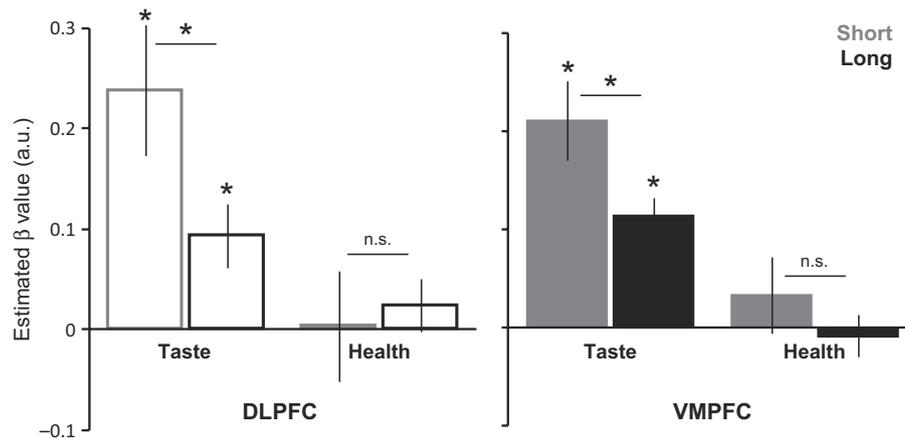


FIG. 5. ROI analysis of BOLD responses to health and taste ratings. Average estimated beta values for each type of rating, ROI and condition (see GLM 2). Asterisks indicate a significant difference at $*P < 0.05$ (one-sample or paired two-tailed *t*-tests, as applicable). In all ROIs and conditions, Taste betas were significantly greater than Health betas ($P < 0.05$).

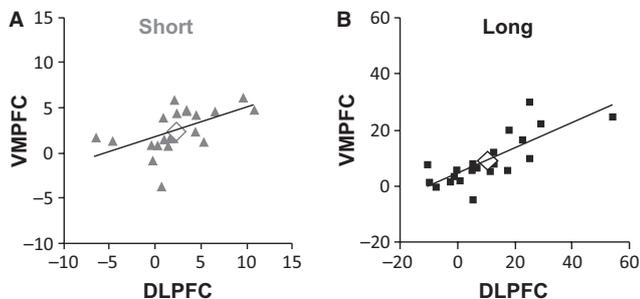


FIG. 6. Correlations of cumulative neural estimate analyses across regions. (A) For Short trials only – estimates of the cumulative difference in activity between trials with SY and SN responses, measured at the end of the initial 1-s viewing period, in the DLPFC and VMPFC ROIs. Each triangle denotes one subject. Large diamonds are the group mean. The line indicates the best fit regression. (B) Analogous plot for the Long condition. Each square denotes one subject. Plots show that in both conditions there is a positive correlation between the ROIs' estimates (Short, $r_{19} = 0.53$, $P = 0.01$; Long, $r_{19} = 0.78$, $P = 0.00004$; Pearson's correlation; $N = 21$).

and Methods, GLM 2). We then extracted parameter estimates for the VMPFC and DLPFC for each attribute and condition and subjected them to a 2 (Region: DLPFC & VMPFC) \times 2 (Attribute: Health & Taste) \times 2 (Condition: Short & Long) repeated-measures ANOVA. This analysis indicated a strong main effect of Attribute ($F_{1,21} = 22.4$, $P = 0.0001$; Taste > Health), a weak main effect of Condition ($F_{1,21} = 3.5$, $P = 0.08$; Short > Long), and a weak Condition \times Attribute interaction ($F_{1,21} = 4.1$, $P = 0.06$), driven by the decrease in responses to Taste in the Long relative to the Short condition in both regions (DLPFC $t_{21} = 2.1$, $P = 0.05$; VMPFC $t_{21} = 2.2$, $P = 0.04$), with no similar decrease in responses to Health (all $P > 0.35$; Fig. 5). Note that conclusions based on these differences are subject to the caveat that the canonical hemodynamic response, as implemented in the GLM, fits the brain activity equally well in both conditions.

Taken together, these results suggest that there are few if any behavioral or neural differences in the attributes driving the computation of decision value.

The combination of differences in timing but not content in VMPFC and DLPFC value signals could result from a sharing of information between the two. This suggests two complementary hypotheses: (i) value signals in the two regions should be correlated; and (ii) greater computation times should allow for greater convergence in

their signals, leading to stronger correlations in the Long vs. Short condition. To test these hypotheses, we calculated the correlation between the total accumulated DV signal in these regions across participants. In both Short and Long conditions, these measures were highly correlated (Short, $r_{19} = 0.53$, $P = 0.01$; Long, $r_{19} = 0.78$, $P = 0.00004$). Moreover, the correlation in the long condition was marginally greater (one-tailed Fisher's *r*-to-*z*, $z = 1.37$, $P = 0.09$; Fig. 6). These results provide modest support for the sharing of value-related information between these regions.

Discussion

In this study, we address four questions related to whether and how the available computation time affects neural decision value signals and behavioral responses. We found that, while the available computation time did not change choices or the attributes driving them, decision value signals in the DLPFC and VMPFC showed marked intra- and interregional differences, providing novel constraints on our understanding of decision values.

Previous studies involving different species and techniques have found that activity in VMPFC and DLPFC reflects the computation of DV signals (Rangel & Hare, 2010; Wallis & Kennerley, 2010; Padoa-Schioppa, 2011), which are then compared to make a choice (Basten *et al.*, 2010; Milosavljevic *et al.*, 2010; Hare *et al.*, 2011b). Activity in the same regions also correlated with DVs in the current study, in both the Short and the Long conditions. This provides additional support in favor of the hypothesis that these two areas are central to the computation of DVs and hence choices. However, the main goal of the study was to extend this literature by characterizing several unknown properties of the DV signals encoded in these two regions, including their relationship to choices, and examine how signals in these areas differ and interact. Our findings support the following conclusions.

Finding 1 – the temporal profile of DV signals depended on the available decision time

In particular, during the first second of the trial, the signals were indistinguishable across the Short and Long conditions but began to diverge thereafter. This suggests that these value computations were sensitive to the presence of the stimulus and/or the onset of the response window. In particular, in the Short condition, DV signals in

both regions began ramping down after the viewing window ended and a response was required, but remained high in the Long condition through its entire 4-s length, only ramping down at the end.

This finding provides useful insights about the nature of the DV computations. First, it shows that in our task the computation of DVs expands to fit the available computation time (and in fact may extend beyond this point in the Short condition). This stands in contrast to neurophysiological studies which suggest that evidence accumulation in parietal neurons during perceptual tasks continues only until a decision threshold is crossed and no further, even if there is more computation time available (Kiani *et al.*, 2008; Kiani & Shadlen, 2009). As there are many differences between the two tasks, an important question for future work will be to characterize what determines the duration of basic perceptual and valuation computations, across a wide variety of tasks. Second, although the start and end of DV computation aligns well with the presence of the stimulus, suggesting that valuation is strictly stimulus-locked, we think our data suggest a more complex relationship. Specifically, there is evidence that these computations may not be strictly stimulus-locked as DV computation in the Short condition continues after, and in the Long condition stops before, the response has been entered and the stimulus is removed from the screen.

It is important to acknowledge that there may be alternative explanations for these response profiles. For example, neural adaptation could account for some of the differences across the Short and Long conditions; alternatively, the observed value signals may have represented the maintenance of a stored value signal and not its ongoing computation. The overwhelming evidence that these regions engage in active and consequential DV computation argues against these interpretations, but future work will be necessary to fully address these alternatives. For example, to test whether adaptation explains the patterns of the observed neural estimates, future studies could vary the length of the viewing window, repeat items a variable number of times, prime item presentations with other objects or vary the stimuli to be simpler or more complex. Observing how computation might change in these cases could shed light on these possible explanations.

Finding 2 – the Long condition resulted in greater cumulative DV signals in both ROIs

This larger cumulative signal could be interpreted to indicate a stronger representation of value, which could have implications for the quality and reliability of downstream processes. Although the increase in available decision time did not lead to improvements in choice quality in this study, our results are consistent with previous studies of speed–accuracy tradeoffs in valuation (Ratcliff & McKoon, 2008; Milosavljevic *et al.*, 2010, 2011).

This finding also provides additional insights into the nature of the valuation process. First, it shows that the DVs are not computed using a computationally efficient process. Specifically, we found that they continued after the response in the Short condition, and that choices were no different across conditions despite an almost five-fold difference in the total amount of computation. This seems highly inefficient as computation is metabolically costly (Kandel *et al.*, 2000), and in the above cases either cannot or did not affect choices. It is possible the additional computation would have effects on other aspects of behavior we did not observe (e.g. learning, changes over time etc.), but this remains a question for future research.

This inefficiency also suggests that DV signals need only be of sufficient quality for consistent accurate choice (for example, to average out the noise). This is consistent with a growing number of studies finding that values are compared to make a choice using algorithms that resemble a Drift-Diffusion Model (Gold & Shadlen, 2002; Bogacz *et al.*, 2006;

Basten *et al.*, 2010; Krajbich *et al.*, 2010; Milosavljevic *et al.*, 2010; Hare *et al.*, 2011b). In these models the comparison process terminates when the cumulative DV signal crosses a pre-specified barrier. As a consequence, further improvements in the signal beyond that point in time do not affect choices. Participants' choices in our study were highly accurate in both conditions, which suggests that even in the Short condition the amount of signal computed was sufficient to cross the associated barriers. This could explain why the additional calculation of the DVs in the Long condition did not improve choice quality.

How can we reconcile our findings of consistently accurate choices with evidence that neural estimates of value signals extend beyond the effective or necessary amount of time? Recent research suggests that value signals may be used after choice for other reasons, including online response adjustment or inter-trial learning, as proposed elsewhere (Resulaj *et al.*, 2009; Ding & Gold, 2011).

Finding 3 – in the Long condition, DV signals of equivalent magnitude appeared later in DLPFC than in VMPFC

A similar difference could not be found for the Short condition, though a careful look at Fig. 3 suggests that the null results for the Short condition might be the product of limited statistical power. Nevertheless, this finding is consistent with previous studies of the DV signals in monkeys (Wallis & Miller, 2003), which also found that the DLPFC signal lagged the one in VMPFC. The authors' interpretation of the delay was that the VMPFC value signals were subsequently passed to DLPFC to influence the selection of motor responses. This interpretation is inconsistent with recent fMRI studies of the network involved in making choices (Basten *et al.*, 2010; Hare *et al.*, 2011b), which find that the DLPFC plays a role in controlling the comparison process, but that its activity does not reflect the computation of DVs. Regardless of the precise computational relationship between the two regions, we demonstrate that they do not have the same temporal profile and, specifically, that computation in the DLPFC is shifted later relative to the VMPFC. Such a temporal difference could arise, for example, if the two areas used different inputs to compute the DVs, which were themselves computed with different latencies, or if value information was shared between these regions, as mentioned above.

Finding 4 – VMPFC and DLPFC were equally sensitive to Tastiness and Healthiness of foods

Contrary to hypotheses that these regions would represent different attributes, VMPFC and DLPFC were equally sensitive to Tastiness and Healthiness of foods, two attributes that can be used to compute DVs (Basten *et al.*, 2010; Rangel & Hare, 2010). Others have argued that VMPFC signals may reflect the integration of 'basic attributes' (e.g., taste) whereas DLPFC may reflect the integration of 'abstract attributes' (e.g., health) (McClure *et al.*, 2004; Hare *et al.*, 2011a). However, an important caveat to conclusions in the current study was the weak behavioral relationship we observed between Healthiness and DVs. It is possible that if participants had been more motivated to consider Healthiness in their choices (e.g. were dieting) we would have observed differences within and between the DLPFC and VMPFC signals. Regardless, our findings are not fully consistent with this view, as the relative weight of Tastiness, which was tightly linked to decision value, was represented similarly in the two areas. The results also show that the ability of DLPFC to consider more abstract attributes, such as health, does not necessarily increase with additional computational time. Future research will be needed to examine how healthiness, and other attributes not considered here, may play a role

in the computation of decision value and may interact with longer amounts of decision time.

By controlling the amount of time in which a very simple choice must be reached, this study provides a lower bound on the consequences of altering the temporal constraints of valuation. Value circuitry shifts computation to accommodate the time given. That it did so was not a foregone conclusion. This suggests both a high degree of neural flexibility at the cost of efficiency and also limited exploitation of the available time. Value computation may be basic, or even primitive, but it is also a remarkably flexible process, sensitive to both external constraints and internal control.

Supporting Information

Additional supporting information can be found in the online version of this article:

Data S1. Additional information on experimental stimulus sorting and stimulus randomization.

Fig. S1. ‘Neural Estimate’ FIR model estimates of responses by condition.

Fig. S2. BOLD FIR model estimates of responses by condition.

Table S1. Contrast table, GLM 1, Decision Value Regressor (R2).

Table S2. Full list of food items used in experiment.

Please note: As a service to our authors and readers, this journal provides supporting information supplied by the authors. Such materials are peer-reviewed and may be re-organized for online delivery, but are not copy-edited or typeset by Wiley-Blackwell. Technical support issues arising from supporting information (other than missing files) should be addressed to the authors.

Acknowledgements

This research was supported by the NSF (SES-0851408, SES-0926544, SES-0850840), NIH (R01 AA018736), the Betty and Gordon Moore Foundation, and the Lipper Foundation to A.R.

Abbreviations

BOLD, blood oxygenation level-dependent (signal); DLPFC, dorsolateral prefrontal cortex; DV, decision value; FIR, finite impulse response; fMRI, functional magnetic resonance imaging; GLM, general linear model; HRF, hemodynamic response function; N, No; ROI, region of interest; SN, Strong No; SY, Strong Yes; VMPFC, ventromedial prefrontal cortex; Y, Yes.

References

Basten, U., Biele, G., Heekeren, H.R. & Fiebach, C.J. (2010) How the brain integrates costs and benefits during decision making. *PNAS*, **107**, 21767–21772.

Baumgartner, T., Knoch, D., Hotz, P., Eisenegger, C. & Fehr, E. (2011) Dorsolateral and ventromedial prefrontal cortex orchestrate normative choice. *Nat. Neurosci.*, **14**, 1468–1474.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P. & Cohen, J.D. (2006) The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol. Rev.*, **113**, 700–765.

Deichmann, R., Gottfried, J.A., Hutton, C. & Turner, R. (2003) Optimized EPI for fMRI studies of the orbitofrontal cortex. *NeuroImage*, **19**, 430–431.

Ding, L. & Gold, J.I. (2011) Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb. Cortex*.

Figner, B., Knoch, D., Johnson, E.J., Krosch, A.R., Lisanby, S.H., Fehr, E. & Weber, E.U. (2010) Lateral prefrontal cortex and self-control in intertemporal choice. *Nat. Neurosci.*, **13**, 538–539.

Friston, K.J., Poline, J.-B., Holmes, A.P., Frith, C.D. & Frackowiak, R.S.J. (1996) A multivariate analysis of PET activation studies. *Hum. Brain Mapp.*, **4**, 140–151.

Gitelman, D.R., Penny, W.D., Ashburner, J. & Friston, K.J. (2003) Modeling regional and psychophysiological interactions in fMRI: the importance of hemodynamic deconvolution. *Neuroimage*, **19**, 200–207.

Gold, J.I. & Shadlen, M.N. (2002) Banburismus and the brain: decoding the relationship between sensory stimuli, decisions, and reward. *Neuron*, **36**, 299–308.

Hare, T.A., O’Doherty, J.P., Camerer, C.F., Schultz, W. & Rangel, A. (2008) Dissociating the role of the orbitofrontal cortex and the striatum in the computation of goal values and prediction errors. *J. Neurosci.*, **28**, 5623–5630.

Hare, T.A., Camerer, C.F. & Rangel, A. (2009) Self-control in decision-making involves modulation of the vmPFC valuation system. *Science*, **324**, 646–648.

Hare, T.A., Malmaud, J. & Rangel, A. (2011a) Focusing attention on the health aspects of foods changes value signals in vmPFC and improves dietary choice. *J. Neurosci.*, **31**, 11077–11087.

Hare, T.A., Schultz, W., Camerer, C.F., O’Doherty, J.P. & Rangel, A. (2011b) Transformation of stimulus value signals into motor commands during simple choice. *PNAS*, **108**, 18120–18125.

Harris, A., Adolphs, R., Camerer, C.F. & Rangel, A. (2011) Dynamic construction of stimulus values in the ventromedial prefrontal cortex. *PLoS ONE*, **6**, e21074.

Kable, J.W. & Glimcher, P.W. (2007) The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.*, **10**, 1625–1633.

Kandel, E.R., Schwartz, J.H. & Jessell, T.M., (Eds) (2000) *Principles of Neural Science*. McGraw-Hill, New York.

Kiani, R. & Shadlen, M.N. (2009) Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, **324**, 759–764.

Kiani, R., Hanks, T.D. & Shadlen, M.N. (2008) Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci.*, **28**, 3017–3029.

Kim, S., Hwang, J. & Lee, D. (2008) Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron*, **59**, 161–172.

Krajchich, I., Armel, C. & Rangel, A. (2010) Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.*, **13**, 1292–1298.

Litt, A., Plassmann, H., Shiv, B. & Rangel, A. (2010) Dissociating valuation and saliency signals during decision-making. *Cereb. Cortex*, **21**, 95–102.

Mars, R.B., Jbabdi, S., Sallet, J., O’Reilly, J.X., Croxson, P.L., Olivier, E. *et al.* (2011a) 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.

Mars, R.B., Sallet, J., Schüffegen, U., Jbabdi, S., Toni, I. & Rushworth, M.F.S. (2011b) Connectivity-based subdivisions of the human right ‘temporoparietal junction area’: evidence for different areas participating in different cortical networks. *Cereb. Cortex*.

McClure, S.M., Laibson, D.I., Loewenstein, G. & Cohen, J.D. (2004) Separate neural systems value immediate and delayed monetary rewards. *Science*, **306**, 503–507.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C. & Rangel, A. (2010) The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgm. Decis. Mak.*, **5**, 437–449.

Milosavljevic, M., Koch, C. & Rangel, A. (2011) Consumers can make decisions in as little as a third of a second. *Judgm. Decis. Mak.*, **6**, 520–530.

O’Reilly, J.X., Beckmann, C.F., Tomassini, V., Ramnani, N. & Johansen-Berg, H. (2010) Distinct and overlapping functional zones in the cerebellum defined by resting state functional connectivity. *Cereb. Cortex*, **20**, 953–965.

Padoa-Schioppa, C. (2011) Neurobiology of economic choice: a good-based model. *Annu. Rev. Neurosci.*, **34**, 333–359.

Plassmann, H., O’Doherty, J.P. & Rangel, A. (2007) Orbitofrontal cortex encodes willingness to pay in everyday economic transactions. *J. Neurosci.*, **27**, 9984–9988.

Plassmann, H., O’Doherty, J.P. & Rangel, A. (2010) Appetitive and aversive goal values are encoded in the medial orbitofrontal cortex at the time of decision making. *J. Neurosci.*, **30**, 10799–10808.

Rangel, A. & Hare, T. (2010) Neural computations associated with goal-directed choice. *Curr. Opin. Neurobiol.*, **20**, 262–270.

Ratcliff, R. & McKoon, G. (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.*, **20**, 873–922.

Resulaj, A., Kiani, R., Wolpert, D.M. & Shadlen, M.N. (2009) Changes of mind in decision-making. *Nature*, **461**, 263–266.

Rushworth, M.F., Noonan, M.P., Boorman, E.D., Walton, M.E. & Behrens, T.E. (2011) Frontal cortex and reward-guided learning and decision-making. *Neuron*, **70**, 1054–1069.

Wallis, J.D. & Kennerley, S.W. (2010) Heterogeneous reward signals in prefrontal cortex. *Curr. Opin. Neurobiol.*, **20**, 191–198.

Wallis, J.D. & Miller, E.K. (2003) Neuronal activity in primate dorsolateral and orbital prefrontal cortex during performance of a reward preference task. *Eur. J. Neurosci.*, **18**, 2069–2081.