

Published in final edited form as:

Nat Neurosci. ; 15(7): 960–961. doi:10.1038/nn.3140.

A mechanism for value-guided choice based on the excitation-inhibition balance in prefrontal cortex

Gerhard Jocham¹, Laurence T Hunt^{1,2}, Jamie Near¹, and Timothy EJ Behrens^{1,3}

¹FMRIB Centre, University of Oxford, John Radcliffe Hospital, Oxford OX3 9DU, United Kingdom

²Department of Experimental Psychology, University of Oxford, South Parks Road, Oxford OX1 3UD, United Kingdom

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

Abstract

Although the ventromedial prefrontal cortex (vmPFC) has long been implicated in reward-guided decision making, its exact role in this process has remained an unresolved issue. Here, we show that vmPFC levels of GABA and glutamate in human volunteers are predictive of both behavioural performance and the dynamics of a neural value comparison signal in a manner as predicted by models of decision-making. These data provide evidence for a neural competition mechanism in vmPFC supporting value-guided choice.

Organisms are constantly required to choose between options that differ in their expected reward value. Whilst neural signals reflecting these values are widespread throughout the brain¹⁻⁵, the ventromedial prefrontal cortex (vmPFC) has attracted particular interest. Neural signals recorded in this region appear to reflect choice^{1, 5} and damage to these areas in humans leads to deficits in decision-making⁶. However, the nature of the computations underlying these decisions has remained elusive.

One popular mechanism for decision-making is competition by mutual inhibition – a class of models in which representations of each available option inhibit one another until activity remains in only one. Such models can be implemented in abstract⁷ or biophysical⁸ forms, and dynamic neural signals consistent with these models can be found in the vmPFC⁹. A crucial prediction is that performance will depend heavily on the levels of inhibition relative to excitation in the network. If the vmPFC implements a decision-making process based on such an inhibitory competition mechanism, then both behavioural performance and the neural dynamics of the vmPFC value comparison signal should depend on the levels of the major excitatory and inhibitory neurotransmitters, glutamate and GABA, in the vmPFC. Such a finding would tie a neurochemical underpinning to our computational understanding

To whom correspondence should be addressed: Gerhard Jocham, FMRIB Centre, University of Oxford, John Radcliffe Hospital, Headington, Oxford OX3 9DU, United Kingdom, gjocham@fmrib.ox.ac.uk, Phone: +44(0)1865 222 738.

Author contributions: GJ designed research, acquired and analysed the data. LTH modified the biophysical model and analysed the biophysical model predictions. JN acquired and analysed MRS data. TEJB designed research and analysed the data. All authors were involved in writing the manuscript.

of value-guided choice, and provide a mechanistic explanation for interindividual variability in choice behaviour.

We tested these predictions by using magnetic resonance spectroscopy (MRS) to obtain measures of each individual's basal GABA and glutamate levels from the vmPFC and a right parietal region in the intraparietal sulcus (IPS) of 25 healthy male volunteers (see supplementary materials). Note that these neurochemical data are not time-resolved or choice-related. They reflect the baseline neurotransmitter levels in each subject at rest collected each from a single voxel in the vmPFC and a single voxel in the IPS. The IPS was chosen as control region because it has also been shown to encode value- and decision-related parameters in a number of studies^{2, 10}. Time constraints and methodological considerations precluded us from using additional control regions such as lateral orbitofrontal cortex (see supplementary materials). Following MRS acquisition, subjects underwent a short version of a reward-guided decision-making task (Fig. 1A and supplementary Fig. 1A) during scanning with fMRI. After scanning, participants completed a longer version of the decision-making task. During the task, subjects repeatedly made choices between two options of differing reward magnitude and reward probability.

Subject performance in this task can be characterised by using a standard Prospect theory model (see online methods). Whereas optimal behaviour on the task would be to multiply magnitude and probability and to choose the option with the highest Pascalian value, the model has two parameters that warp probability and reward space to match subject behaviour. A third, crucial parameter (the softmax inverse temperature, τ), reflects the accuracy of subject decisions. Subjects with a low τ require a substantial value difference in order to select the higher-value option reliably (see supplementary fig. 1D). If decisions are made by mutual inhibition, τ should increase with higher GABA and lower Glutamate levels (see supplementary fig. S4A for this relation in one instantiation of one mutual-inhibition model⁹).

In the decision-making task, GABA and glutamate concentrations indeed predicted average choice accuracy, τ . Because GABA is synthesized from glutamate, the levels of the two neurotransmitters are highly correlated within each brain region. It is therefore crucial to orthogonalize all effects with respect to the other neurotransmitter¹¹ – i.e. to compute partial correlations. Doing so, we found that vmPFC GABA and glutamate had opposing effects: τ was highest in subjects with high levels of GABA ($r=0.76$, $p < 0.00001$, Fig. 1D) and low levels of glutamate ($r = -0.598$, $p=0.001$, Fig. 1). This pattern was specific to the vmPFC, as no such relation was found with IPS GABA ($r=0.25$, $p=0.13$) and glutamate levels ($r = -0.29$, $p=0.1$). Furthermore, when the two pairs of neurotransmitter concentrations were formally compared in a single linear model (see supplement) vmPFC GABA had a significantly bigger positive effect than IPS GABA ($t=2.05$, $p=0.027$), and vmPFC glutamate had a significantly bigger negative effect than IPS glutamate ($t=2.29$, $p=0.017$).

In agreement with previous studies⁵, BOLD activity in the vmPFC recorded in the fMRI session correlated positively with the value difference between chosen and unchosen options on each trial (Fig. 2A, B). If the evolution of this value difference signal is indeed dependent on a balance between mutual inhibition and recurrent excitation, then it should ramp up

faster and ramp down earlier in subjects with relatively high ratios of excitation to inhibition. We therefore computed the temporal derivative of the value difference signal throughout the trial (Fig. 2C), and examined its correlation with vmPFC GABA and glutamate levels. Early in the trial, individuals with high glutamate ($t=4.96$, $p<0.00005$, Fig. 2D) and low GABA ($t=-3.05$, $p=0.0029$) levels exhibited higher derivatives, or faster ramping up. Late in the trial, the same individuals exhibited the most negative derivatives, indicating faster ramping-down (GABA: $t=3.42$, $p=0.0012$; glutamate: $t=-4.25$, $p<0.0002$). Importantly, this effect was specific to the value difference correlate, no such pattern was found on the raw BOLD signal, making it unlikely that the current finding is due to a general effect on the BOLD signal (supplementary fig. 3).

We have shown that interindividual variability in vmPFC GABA and glutamate concentrations explains variability in choice behaviour and fMRI signals recorded during value-guided choice. Taken together, the present findings suggest that value-guided choice is governed by a competition by mutual inhibition that is mediated by a balance between GABAergic inhibition and glutamatergic excitation in the vmPFC. We have recently shown using magnetoencephalography that vmPFC exhibits dynamics predicted by neural competition⁹. The current findings further support the idea that vmPFC has a central role to play not only in valuation¹² but also in choice¹³. However, in demonstrating that this competitive process is predictably dependent on GABA and glutamate concentrations, the current findings provide a clear link from neurochemical to computational mechanisms and thence to economic choice. Such an understanding of neurochemical mechanisms has potential clinical relevance. For example, It is noteworthy that altered prefrontal levels of GABA and glutamate have been reported in individuals with major depressive disorder¹⁴, a condition that has impairments in decision-making as one of its diagnostic criteria.

Online Methods

Participants

25 healthy male participants (18 to 35 years) participated in the experiment after written informed consent was obtained. The local ethics committee approved all procedures.

Magnetic resonance (MR) imaging

MR data were acquired at 3T on a Siemens Trio using a 32-channel coil. First, a high-resolution T1-weighted scan was acquired using an MPRAGE sequence¹⁵. Spectroscopy voxel placement was based on this scan. MR spectroscopy (MRS) data were acquired as described previously¹⁵, with TR = 3200 ms. The VMPFC voxel (AP $1.5 \times$ ML $3.0 \times$ DV 1.0 cm) was mediolaterally centred on the midline and dorsoventrally on the genu of the corpus callosum, with the posterior voxel boundary just rostral to the genu. The parietal voxel ($2.0 \times 2.0 \times 2.0$ cm) was centred on the right intraparietal sulcus (IPS). We acquired 180 averages for the parietal and 360 averages for the VMPFC voxel to compensate for the reduced SNR. A rather thin (dorsoventrally) voxel was used for the VMPFC to reduce the impact of field inhomogeneities in this brain area. The presence of field inhomogeneities was also the reason for choosing IPS as control site, rather than lateral orbitofrontal cortex. Temporal constraints precluded acquisition of more than two voxels. MRS was followed by

functional magnetic resonance imaging (fMRI) during which subjects performed the task. 45 slices with voxel resolution of 3 mm isotropic were obtained using a sequence optimized for the orbitofrontal cortex¹⁶. Field maps were acquired using a dual echo 2D gradient echo sequence with TR = 488 ms and TE of 7.65 and 5.19 ms on a 64×64×40 grid.

Behavioural task

The task involved repeatedly choosing between two options to obtain monetary reward (supplementary figure 1A). Each option consisted of one horizontal bar (reward magnitude) and a percentage written underneath it (reward probability). On a subset of trials ('no brainer' trials), both the magnitude and probability of one option was higher than on the alternative option. The reward schedule was designed to minimize the correlation between chosen and unchosen value (mean $r = 0.21$ and 0.33 across subjects for the fMRI and post-scanning task). 100 trials were presented during fMRI. After scanning, participants undertook another 400 trials of this task (without the initial viewing period and with quicker timing) on a laptop outside the scanner.

Data analysis

Behavioural data—Subjective reward magnitudes and probabilities were derived by fitting utility functions according to prospect theory¹⁷:

$$r_s = r_o^\alpha$$

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

Where the objective reward magnitude and probability r_o and p_o are transformed into subjective magnitude and probability, r_s and p_s , respectively. From these values, subjective expected values can be calculated as

$$sEV = r_s * p_s$$

The modelled probability to choose either of the two options was given by a softmax rule:

$$P_{(C=K)} = \frac{e^{sEV_K * \tau}}{\sum_{i=1}^n e^{sEV_i * \tau}}$$

Where K = choice made by subject, n = number of options, τ = softmax inverse temperature. We also fitted two models with only two free parameters, where either α or γ were fixed at 1. Calculation of Bayes information criterion (BIC) showed that model fits under the two alternative models were significantly worse, ($p < 0.008$, supplementary table 2).

We custom-implemented a Bayesian estimation procedure in MATLAB (The Mathworks, USA) to obtain the best-fitting parameters α , γ and τ (supplementary fig. 1 B, C). Choice probabilities as a function of τ are shown in supplementary figure 1D. To test whether subjects integrated probabilities and magnitudes, we performed a logistic regression analysis. Reward probability, magnitude, choice on the previous trial and outcome on the previous trial were entered as independent variables X to predict the binary outcome Y (choices, 0 or 1 for left and right choices, respectively, supplementary figure 1E). A further linear regression tested the impact of value difference, value sum and no-brainer trials on subjects' log reaction times (supplementary fig. 1F).

Processing and analysis of MRS data—A semi-automated MATLAB-based preprocessing routine was applied to all spectra prior to analysis. Motion corrupted signal averages were identified and removed, frequency and phase-drift corrections were performed to ensure exact alignment of the remaining averages, followed by signal averaging. Fully processed spectra were then analysed as in¹⁵. GABA and glutamate values are reported as ratio to creatine. GABA and glutamate were successfully detected in 24 (VMPFC) and 22 (Parietal cortex) volunteers. Only two spectra from VMPFC had a water line-width slightly > 10 Hz (11.5 and 11.66). As reported in the supplements, excluding them did not change the pattern of results, which is why the data reported in the main text include those two VMPFC data sets. The T1-weighted anatomical scans were segmented into grey and white matter using FAST (FMRIB's Automated Segmentation Tool)¹⁸ to calculate relative volumes of grey matter, white matter and cerebrospinal fluid in the MRS voxels. The reported concentrations of GABA and glutamate are corrected for the relative grey matter volumes (supplementary table 1).

Analysis of fMRI data—Analysis of fMRI data was performed using tools from the FMRIB Software Library (FSL¹⁹), using the same routine as in²⁰, with the spatial filter set to 6 mm full-width at half maximum. To investigate activity related to the value difference between the chosen and unchosen options, we set up a GLM that contained the following nine regressors: value difference, value sum, outcome value (reward vs. no reward obtained), reward magnitude obtained, main effect from stimulus presentation to response onset, main effect from response onset to outcome delivery, main effect of outcome phase, and two stick functions modelling left and right button presses, respectively. In addition, the six motion parameters from motion correction were included. Contrast images from the first level were then taken to group level using a random effects analysis. Results are reported at an activation level threshold of $p < 0.005$ ($Z > 2.58$) combined with a cluster-forming threshold of $p < 0.05$.

Region of interest (ROI) analyses—The above whole-brain analysis yielded a positive effect of value difference in the VMPFC (supplementary figure S2A). BOLD timecourses were extracted from the resulting activation. Each volunteer's timeseries was cut into trials with a duration of 16.2 s (Symbol presentation at 0 s, response onset at 4.58 s, outcome presentation at 8.12 s, corresponding to the mean onsets across subjects and trials). Timeseries were resampled to a resolution of 300 ms. A GLM containing the parameters of interest was then fitted at each time point for each volunteer. This resulted in a timecourse of

effect size for each regressor and for each volunteer. Timecourses were then averaged across participants (Fig. 02, supplementary figure 2B/C).

Correlation analyses—As mentioned in the main text, GABA and glutamate are correlated. Therefore, we performed partial correlation analyses. To test whether GABA and glutamate levels from vmPFC were better at predicting τ than IPS levels of those transmitters, GABA and glutamate from both regions were entered as regressors in the design matrix (along with a constant term) to predict the data, τ . The effects of vmPFC GABA and glutamate were directly contrasted with those of IPS GABA and glutamate.

For the time-resolved analysis of GABA and glutamate on the slope of the value difference signal, both GABA and glutamate were entered into one single GLM such that again, the reported effects are exclusively on the orthogonal, non-shared variance between the two neurotransmitters. To test whether the pattern of results could be due to a general relation between GABA and glutamate levels and the BOLD response, we conducted the same analysis, however, this time looking at the effects of the two neurotransmitters on the main effect, rather than on the value difference slope (supplementary fig. S3).

Modelling—We implemented a mean-field reduction of the spiking neuronal network model described in²¹. Full details are given in²². It is important to note that a number of models exist that implement a neural competition based on mutual inhibition, e.g.²¹⁻²⁴. We do not claim that our results are specific to this particular model; instead we choose this model as an example of the class. We have used this model successfully in a recent study²² to predict local field potential data, which are more closely related to the fMRI BOLD signal than neuronal spiking activity, which is why we decided to base our predictions on this model.

Model analysis—For model predictions of cross-subject behavioural variation, we fit softmax functions to model choice behaviour in exactly the same way as for individual subjects, choosing the softmax parameter that maximised the log-likelihood of each model instantiation's choices. The regression of these softmax parameters against $w+$ is plotted in supplementary figure 4A.

For model predictions of cross-subject neural activity variation, we first estimated a linear regression model for each instantiation of the model, with value difference and overall value as independent variables, and the network's response as the dependent variable. As for the fMRI data, we calculated the first temporal derivative of the parameter estimate for value difference. We then ran a second level linear regression analysis (*across* model instantiations) in which this temporal derivative was the dependent variable, and $w+$ was an independent explanatory variable. The T-statistic from this second level analysis is plotted in supplementary fig. 4B.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

This work was supported by a Wellcome Trust Research Career Development fellowship to TB (WT088312AIA). LH was supported by a 4-year DPhil studentship from the Wellcome Trust (WT080540MA). JN is supported by the Medical Research Council. We are very grateful to Alireza Soltani for providing the MATLAB code for the biophysical model. We thank Steve Knight for his valuable help in data acquisition.

References for main text

1. Wunderlich K, Rangel A, O'Doherty JP. *Proc Natl Acad Sci U S A*. 2009; 106:17199–17204. [PubMed: 19805082]
2. Gershman SJ, Pesaran B, Daw ND. *J Neurosci*. 2009; 29:13524–13531. [PubMed: 19864565]
3. Padoa-Schioppa C, Assad JA. *Nature*. 2006; 441:223–226. [PubMed: 16633341]
4. Cai X, Kim S, Lee D. *Neuron*. 2011; 69:170–182. [PubMed: 21220107]
5. Boorman ED, Behrens TE, Woolrich MW, Rushworth MF. *Neuron*. 2009; 62:733–743. [PubMed: 19524531]
6. Bechara A, Tranel D, Damasio H. *Brain*. 2000; 123(Pt 11):2189–2202. [PubMed: 11050020]
7. Ratcliff R, Rouder JN. *Psychological Science*. 1998; 9:347–356.
8. Wong KF, Wang XJ. *J Neurosci*. 2006; 26:1314–1328. [PubMed: 16436619]
9. Hunt LT, et al. *Nat Neurosci*. 2012
10. Hare TA, Schultz W, Camerer CF, O'Doherty JP, Rangel A. *Proc Natl Acad Sci U S A*. 2011; 108:18120–18125. [PubMed: 22006321]
11. Stagg CJ, Bachtiar V, Johansen-Berg H. *Curr Biol*. 2011; 21:480–484. [PubMed: 21376596]
12. Schoenbaum G, Roesch MR, Stalnaker TA, Takahashi YK. *Nat Rev Neurosci*. 2009; 10:885–892. [PubMed: 19904278]
13. Rushworth MF, Noonan MP, Boorman ED, Walton ME, Behrens TE. *Neuron*. 2011; 70:1054–1069. [PubMed: 21689594]
14. Hasler G, et al. *Arch Gen Psychiatry*. 2007; 64:193–200. [PubMed: 17283286]

References for online methods

15. Stagg CJ, et al. *J Physiol*. 2011; 589:5845–5855. [PubMed: 22005678]
16. Deichmann R, Gottfried JA, Hutton C, Turner R. *Neuroimage*. 2003; 19:430–441. [PubMed: 12814592]
17. Tversky A, Kahneman D. *Journal of Risk and Uncertainty*. 1992; 5:297–323.
18. Zhang Y, Brady M, Smith S. *IEEE Trans Med Imaging*. 2001; 20:45–57. [PubMed: 11293691]
19. Smith SM, et al. *Neuroimage*. 2004; 23(Suppl 1):S208–219. [PubMed: 15501092]
20. Boorman ED, Behrens TE, Woolrich MW, Rushworth MF. *Neuron*. 2009; 62:733–743. [PubMed: 19524531]
21. Wang XJ. *Neuron*. 2002; 36:955–968. [PubMed: 12467598]
22. Hunt LT, et al. *Nat Neurosci*. 2012
23. Usher M, McClelland JL. *Psychol Rev*. 2001; 108:550–592. [PubMed: 11488378]
24. Ditterich J, Mazurek ME, Shadlen MN. *Nat Neurosci*. 2003; 6:891–898. [PubMed: 12858179]

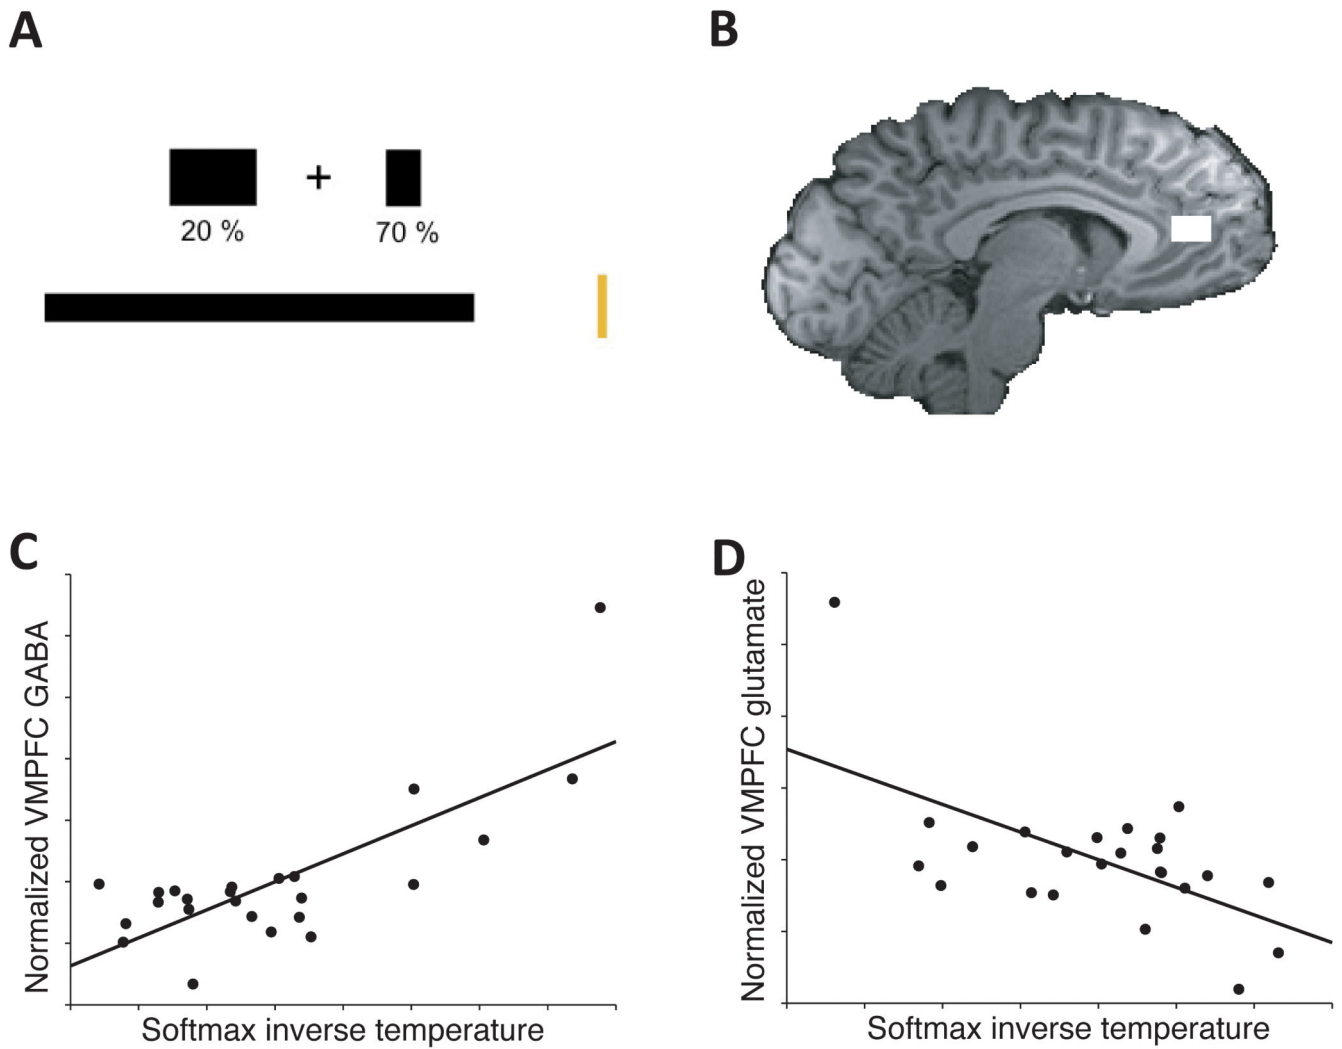


Figure 1.

Experimental task and correlation of spectroscopy data with behaviour. A: Task schematic. B: Location of the MRS voxel in the vmPFC. C and D: Performance was best in subjects with high levels of GABA (C) and low levels of glutamate (D). Values represent arbitrary units, as glutamate (C) and GABA (D) respectively had been regressed out of both variables to demonstrate effects that are orthogonal with respect to the other neurotransmitter.

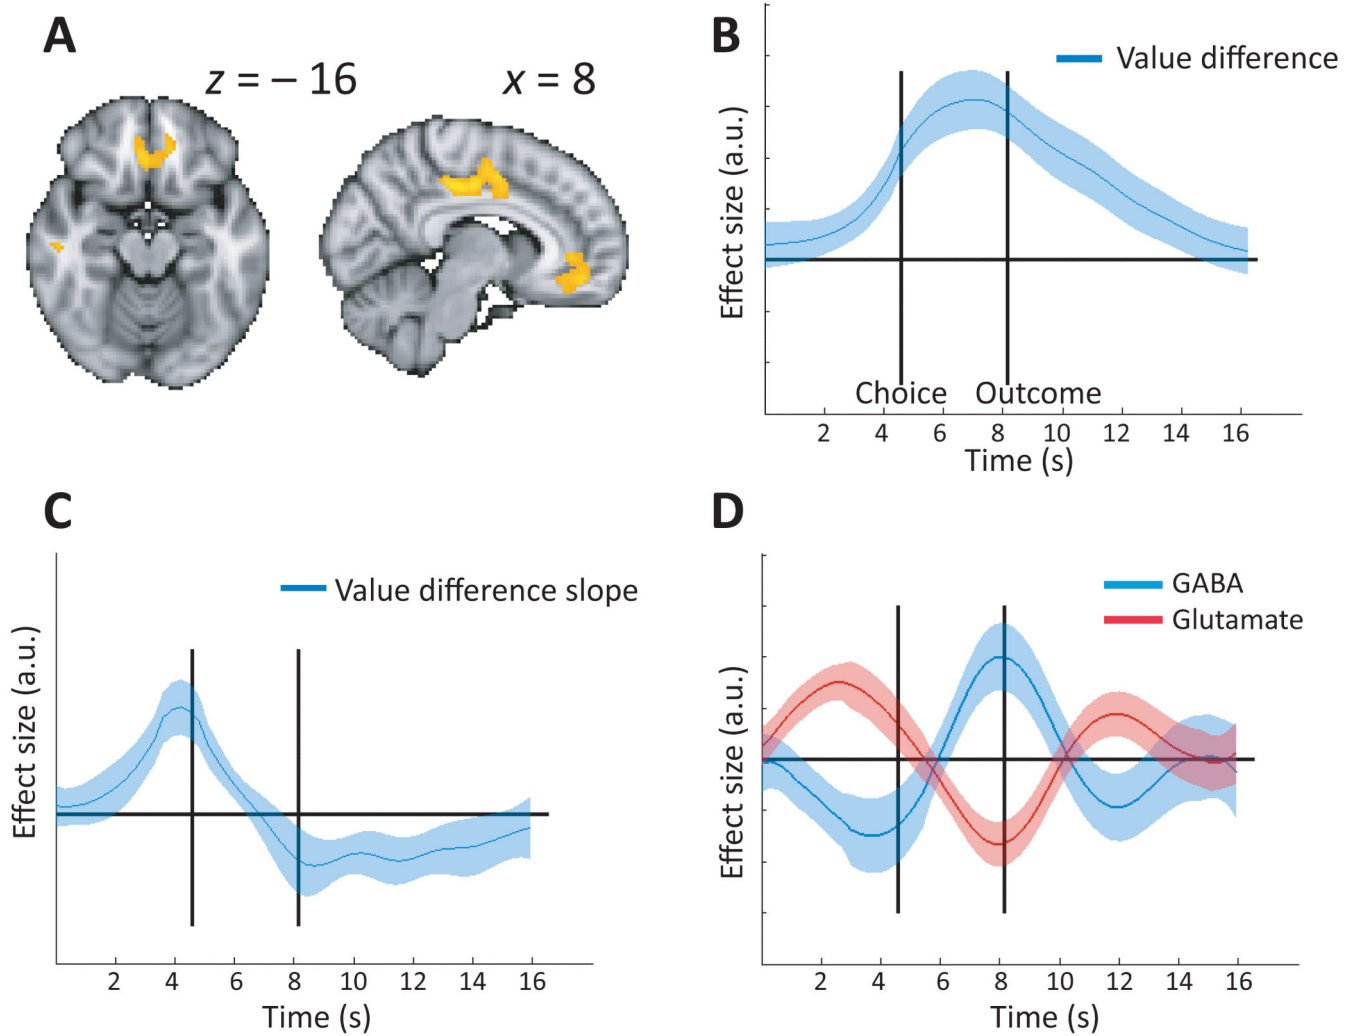


Figure 2. Relationship between GABA and Glutamate and value difference coding. **A:** Whole-brain analysis showing the effects of value difference. **B:** Time-resolved regression of value difference against the fMRI signal from the vmPFC. **C:** Slope (temporal derivative) of the value difference signal in (B). In B and C, the solid line represents the mean across subjects, the shaded area the standard error of the mean. **D:** The slope of the vmPFC value difference signal (C) was driven in opposite directions by local concentrations of GABA and glutamate.