



Predeliberation activity in prefrontal cortex and striatum and the prediction of subsequent value judgment

Uri Maoz^{1*}, Ueli Rutishauser^{1,2}, Soyoun Kim³, Xinying Cai⁴, Daeyeol Lee³ and Christof Koch^{1,5}

¹ Division of Biology, California Institute of Technology, Pasadena, CA, USA

² Department of Neurosurgery and Neurology, Cedars Sinai Medical Center, Los Angeles, CA, USA

³ Department of Neurobiology, Yale University School of Medicine, New Haven, CT, USA

⁴ Department of Anatomy & Neurobiology, Washington University, St. Louis, MO, USA

⁵ Allen Institute for Brain Science, Seattle, WA, USA

Edited by:

Hauke R. Heekeren, Freie
Universität Berlin, Germany

Reviewed by:

Geoffrey Schoenbaum, University of
Maryland School of Medicine, USA
Shunsuke Kobayashi, Fukushima
Medical University, Japan

*Correspondence:

Uri Maoz, Division of Biology,
California Institute of Technology,
1200 E California Blvd, MC 216-76,
Pasadena, CA 90024, USA
e-mail: urim@caltech.edu

Rational, value-based decision-making mandates selecting the option with highest subjective expected value after appropriate deliberation. We examined activity in the dorsolateral prefrontal cortex (DLPFC) and striatum of monkeys deciding between smaller, immediate rewards and larger, delayed ones. We previously found neurons that modulated their activity in this task according to the animal's choice, while it deliberated (choice neurons). Here we found neurons whose spiking activities were predictive of the spatial location of the selected target (spatial-bias neurons) or the size of the chosen reward (reward-bias neurons) before the onset of the cue presenting the decision-alternatives, and thus before rational deliberation could begin. Their predictive power increased as the values the animals associated with the two decision alternatives became more similar. The ventral striatum (VS) preferentially contained spatial-bias neurons; the caudate nucleus (CD) preferentially contained choice neurons. In contrast, the DLPFC contained significant numbers of all three neuron types, but choice neurons were not preferentially also bias neurons of either kind there, nor were spatial-bias neurons preferentially also choice neurons, and vice versa. We suggest a simple winner-take-all (WTA) circuit model to account for the dissociation of choice and bias neurons. The model reproduced our results and made additional predictions that were borne out empirically. Our data are compatible with the hypothesis that the DLPFC and striatum harbor dissociated neural populations that represent choices and predeliberation biases that are combined after cue onset; the bias neurons have a weaker effect on the ultimate decision than the choice neurons, so their influence is progressively apparent for trials where the values associated with the decision alternatives are increasingly similar.

Keywords: pre-deliberation decision bias, value-based decision-making, decision circuit-modeling, free-choice decision making, dorsolateral prefrontal cortex, ventral striatum, caudate nucleus, monkey single-neuron recording

INTRODUCTION

There is a long-standing debate on the proper interpretations of brain signals that predict upcoming free decisions before subjects report having decided. In such free decisions the choice is up to the subject and there is no right or wrong answer (Libet et al., 1983; Libet, 1985). Recent results suggest that motor intentions can be decoded in these cases from neural activity up to about a second before action onset using EEG (Haggard and Eimer, 1999; Bai et al., 2011) or single-neuron activity in humans (Fried et al., 2011) and up to several seconds before onset, at about 60% accuracy, using functional magnetic-resonance imaging (fMRI; Soon et al., 2008; Bode et al., 2011). It was also demonstrated that more-abstract decisions (such as adding or subtracting two small numbers) could be decoded before the decision alternatives became available (Haynes et al., 2007; Soon et al., 2013). Interestingly, models of these experimental results focus on the dorsolateral prefrontal cortex (DLPFC) as the locus of intentional action selection (Pockett, 2006; Pacherie and Haggard, 2011).

However, the choices in the above experiments were generally between randomly raising the left or right hand, with no consequences and for no reason or purpose (although early, predictive local-field potentials may also exist for deliberate decision making that relies on rational deliberation; Maoz et al., 2012). So, importantly, it is not clear to what extent these early predictive signals are part of the neural process involved in deliberate decisions. An alternative hypothesis is that in the absence of reasons to prefer one option over the other, the decision network may also rely on activity that occurs before the decision options are known and rational deliberation can begin. Such slow background fluctuations, for example, could be used to break the symmetry and bias the decision toward one of the equally valued alternative (Haynes, 2011). In contrast, for divergently valued alternatives, the predeliberation fluctuations that occur before rational deliberation can begin are generally not strong enough to influence the outcome, so they would have little predictive power over the decisions. An empirical prediction of this predeliberation-biases hypothesis is,

therefore, that outcomes of decisions between similarly valued options could be, to some extent, predicted from brain activity occurring before the subject knows what the decision alternatives will be.

Neuronal activity before the presentation of the decision alternatives that is predictive of the eventual choice has been shown before, including in prefrontal cortex and the caudate nucleus (CD). But this was in situations where the decision alternatives associated with maximal reward were stable across multiple trials, and the optimal choice could thus be computed before the decision alternatives were revealed (Coe et al., 2002; Lauwereyns et al., 2002; Ding and Hikosaka, 2006). Perceptual judgments were also shown to partially depend on ongoing neural activity prior to the onset of the sensory stimuli (Sadaghiani et al., 2010), including in CD—but were not found in the DLPFC (Kim and Shadlen, 1999; Shadlen and Newsome, 2001; Williams et al., 2003; Ding and Gold, 2010). Moreover, an integrate-and-fire attractor model of decisions based on random spontaneous firing suggested that decisions can be predicted before cue onset from ongoing noise (Rolls and Deco, 2011). However, little is known about the neural correlates of biases during value-based decision-making, where computation of the optimal decision cannot begin before the decision alternatives are presented. In this type of decisions subjects' choices are freely based on the subjective values they associate with the decision alternatives (Rangel et al., 2008), and it is difficult to define correct and incorrect responses.

A common example of conflict among value-based decision options is intertemporal choice between a small, immediate reward and a larger, delayed one (Frederick et al., 2002). Previous studies found that the activity of individual neurons in the DLPFC, parietal cortex, and basal ganglia (BG) encode the subjective value of a delayed reward, namely, the magnitude of reward discounted by its delay (Kim et al., 2008; Louie and Glimcher, 2010; Cai et al., 2011). But how such temporally discounted values are used for decision making in the brain is not well understood. The prefrontal cortex, including DLPFC, has been implicated in multiple aspects of decision making in both humans and non-human primates (Rangel et al., 2008; Wallis and Kennerley, 2010), such as coding subjective values for alternative options or coding the option chosen by the animal (Barraclough et al., 2004; Padoa-Schioppa and Assad, 2006; Kennerley et al., 2009). Previous neuroimaging, lesion, and stimulation studies also highlighted the role of BG and cortico-BG communication in regulating decisions involving temporal delays (Kable and Glimcher, 2007, 2010; Luhmann et al., 2008; Pine et al., 2009; Cavanagh et al., 2011).

Following the above, we therefore focused here on an intertemporal choice in value-based decision-making to empirically test a prediction of the predeliberation-biases hypothesis above. Namely, we wanted to investigate the extent to which neural information that was represented in the DLPFC, ventral striatum (VS) and CD before cue onset—hereafter *bias activity*—influenced later decisions. We also wanted to know whether there was variability in bias-activity representation across brain regions. We were further interested in how bias activity differs, if at all, from value-related information that appeared after cue onset during deliberation—hereafter *choice activity*—and how the two are integrated. Our working hypothesis was that the bias activity was

dissociated from the choice activity, and had a weaker influence on the ultimate decision than that choice activity, so that its effect was mainly apparent when the values of the decision alternatives were similar.

MATERIALS AND METHODS

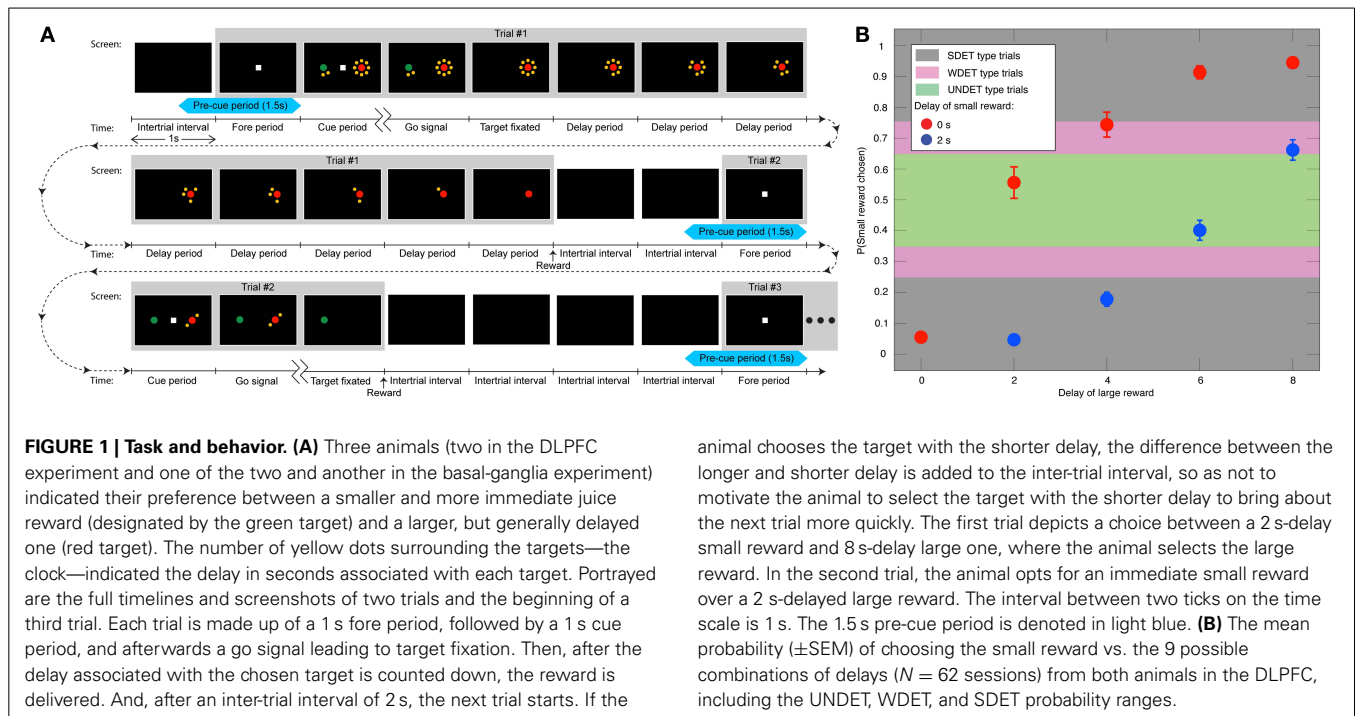
ANIMAL PREPARATION

All data analyzed in this study were from three male rhesus monkeys (D and J for DLPFC, H and again J for BG). Behavioral and neurophysiological data were collected as previously described (Kim et al., 2008; Cai et al., 2011). All procedures used in the present study were approved by the Institutional Animal Care and Use Committee at Yale University and The University of Rochester Committee on Animal Research, and conformed to the *PHS Policy on Humane Care and Use of Laboratory Animals* and the *Guide for the Care and Use of Laboratory Animals*.

BEHAVIORAL TASK

The animals performed an intertemporal choice and a control task. During each trial of the intertemporal choice task, the animal fixated a small white square that appeared at the center of the screen during a 1 s fore period (**Figure 1A**). Two peripheral targets were next presented along the horizontal meridian during the 1 s cue period. The central square was then extinguished—the go signal—and the animal was required to shift its gaze toward one of the two targets. Choosing the red target resulted in a large reward (0.4 ml of apple juice), while the green target delivered a smaller reward (0.26 ml). Target positions were randomized and counterbalanced across trials. Each target appeared either by itself or surrounded by a clock consisting of 2, 4, 6, or 8 yellow dots. This corresponded to no delay or to 2, 4, 6, or 8 s delay before reward delivery, respectively. Once the animal fixated its chosen target, the other target disappeared, and the clock surrounding the chosen target began counting down the dots at a rate of 1/s (**Figure 1A**). The animal was rewarded when the last yellow dot disappeared. The screen then went blank for a 2 s inter-trial interval. When the animal chose the target with the shorter delay, the inter-trial interval was increased by the difference between the longer and shorter delay, so that the onset of the next trial was not influenced by the animal's choice (**Figure 1A**). The delay associated with the small-reward target was 0 or 2 s long, and that associated with the large-reward was 0, 2, 4, 6, or 8 s. For the striatum data all 10 possible delay pairs were used, while for the DLPFC the combination of 2 s delay for the small reward and 0 s delay for the large reward was not used. Each combination was presented twice (four times) in a block of DLPFC (striatum) intertemporal choice trials with the target positions counterbalanced; 18 (40) trials per block in experiments in the DLPFC (striatum).

The control task was identical to the intertemporal choice task, except for two changes. First, the central fixation square was either green or red, which informed the animal of the target it was required to select. Incorrect trials were not rewarded. Second, the reward amount was fixed (at 0.13 ml) and awarded with no delay. Blocks of intertemporal choice and control trials were presented alternately. The experiment was set up so that each session would consist of the same number of trials in the intertemporal choice



and control tasks. Of the 164 DLPFC neurons, 135 were tested in 10 blocks for each condition (360 trials) and the rest were tested for at least six blocks for each condition (216 trials). Of the 183 striatal neurons (93 CD, 90 VS), 181 were tested over at least three blocks (120 trials) and on average more than four blocks were tested per neuron (167.4 ± 3.7 and 162.4 ± 4.1 trials for CD and VS, respectively).

When the animal broke fixation before the go signal, the trial was restarted with the same rewards and delays. If fixation was broken after cue onset (in 4855 of overall 36738, 13%, and in 369 trials of overall 34052 trials, 1%, for DLPFC and striatal recordings, respectively), the animal could have gained information about the decision option it was about to face in the upcoming restarted trial. Trials that immediately followed those where fixation breaking occurred after cue onset were therefore not included in the analysis. This, together with the randomization of the small and large reward positions and their associated delays for the intertemporal choice and control blocks, meant that during the pre-cue period analyzed in the present study, the animal would know neither the positions nor the delay of the targets about to appear. Therefore, in our task, cue onset is also when rational deliberation can begin. Our data is taken from Experiment I of Kim et al. (2008) and from Cai et al. (2011), where more details can be found.

ANALYSIS OF BEHAVIORAL DATA—PARTITIONING THE TRIALS INTO FOUR TYPES

Following our working hypothesis, we expected the bias activity to increasingly influence the animals' decisions when the subjective values that the animals associated with the decision alternatives became more similar—i.e., when the temporally discounted reward magnitudes associated with the two targets were

comparable for the animal. In other words, the less the decision was determined by the cue information, the more we expected it to be predictable from the bias activity. So we divided the experimental trials into four types according to the apparent similarity of the values of the decision alternatives. The frequency with which the animals chooses a particular target over its alternative within every pair of reward/delay combinations is a well-established estimate of the subjective value of that option (Sugrue et al., 2004). We therefore used it here too. More specifically, the animals increasingly discounted the value of the reward as a function of the reward's delay (Figure 1B) (Kim et al., 2008; Hwang et al., 2009). For each recording session, there were 9 or 10 possible short/long delay pairs (see above), each associated with the ratio of trials in which the animal chose the small reward to the overall number of trials (Figures 1B, Figure S1). The four trial types were: The *undetermined* trials [where the probability of selecting the small reward was in the range (0.35, 0.65)]—denoted *UNDET*—that had the most similar discounted values for the two targets, making the decision largely undetermined by the visual cues. The *weakly determined* trials [choice ratio range (0.25,0.35) or (0.65,0.75)]—*WDET*—that had the discounted values for the two targets farther apart. The *strongly determined* trials [ratio in (0,0.25) or (0.75,1)]—denoted *SDET*—that had the most divergent discounted values. Lastly, *no choice* trials—*NC*—included all trials in the control experiment, where the animal was instructed which target to fixate.

ANALYSIS OF NEURAL DATA

Overall, 164 DLPFC neurons were recorded (77 from animal D, 87 from J), 90 VS neurons (33 from H, 57 from J), and 93 CD neurons (32 from H, 61 from J). Great care was taken to insure

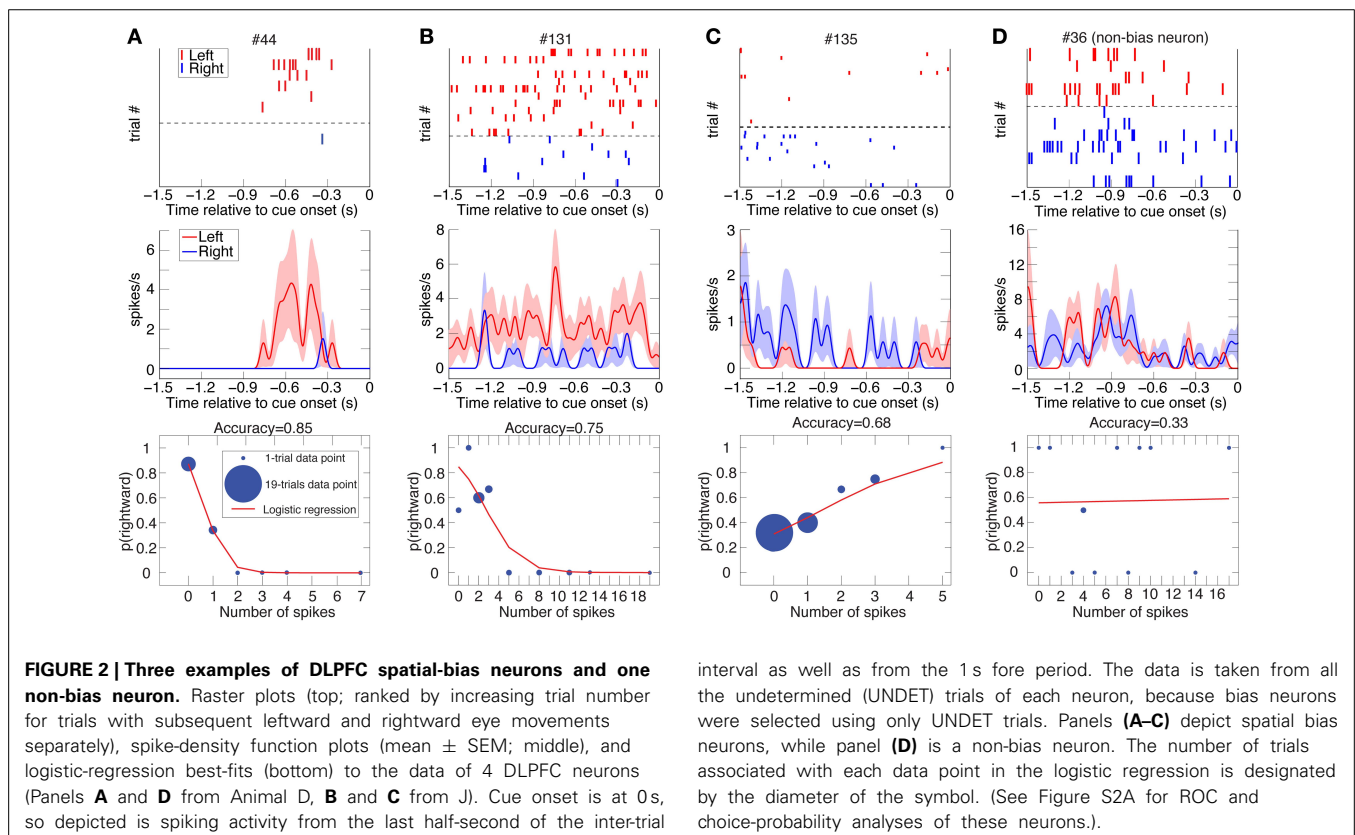
that all spikes originated from single neurons. To determine bias neurons we looked at spiking activity 1.5 s before cue onset in all analyses—henceforth the *pre-cue period*. This time frame was comprised of the last 0.5 s of the inter-trial interval and the 1 s fore period (Figure 1A). We included only the last half-second of the inter-trial interval because we did not want activity from the previous trial, which might linger in the beginning of the inter-trial interval, to confound our results.

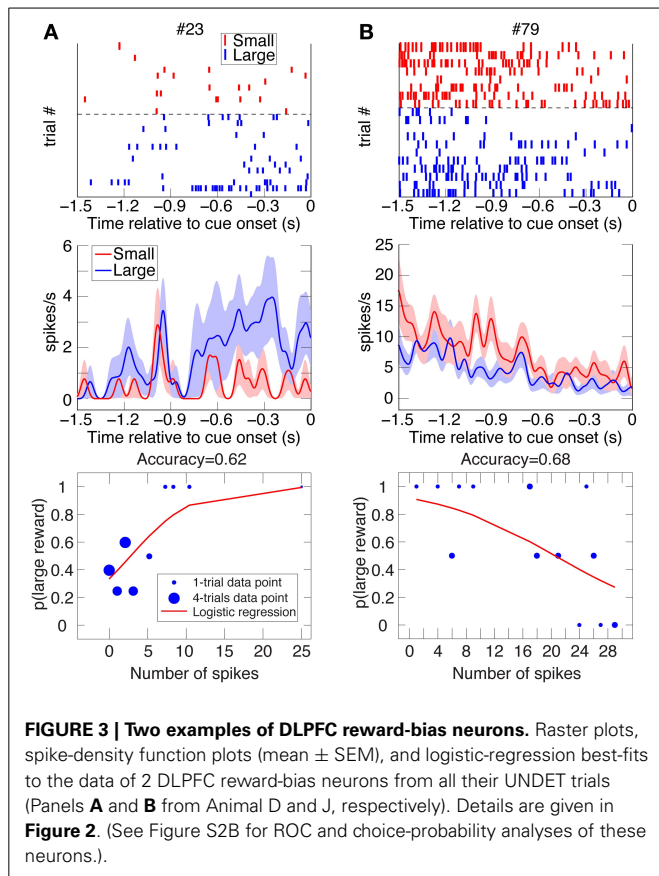
In addition, only neurons that satisfied three criteria were included in this study. First, they were recorded during sessions containing at least one UNDET-type reward-delay combination (which on average \pm SEM consisted of 32 ± 3 , 23 ± 3 , and 30 ± 6 trials for the DLPFC spatial-bias, DLPFC reward-bias, and VS spatial-bias, respectively). Second, the mean firing rate of the neurons 1.5 s before cue onset was at least 0.5 spikes/s. Third, the animal did not display a very strong behavioral spatial preference toward the left or the right—i.e., it had to choose the left and right target at least 25% of the trials. Importantly, any deviations from an even left/right split are specifically taken into account by our method for identifying bias neurons and by the proportion explained measure (see below). When analyzing the DLPFC data for spatial biases, 105 of the 164 neurons (64%) met all three criteria and were subjected to further analyses (44 from animal D). In the analysis of reward biases, the third criterion was not imposed because it is dissociated from a choice based on reward size. Hence, for the reward-size analysis, 118 of the 164 neurons (72%) were subjected to further

analysis (53 from animal D). For VS, 52 and 56 of the 90 neurons (58 and 62%) were included in the spatial-bias and reward-bias analyses (38 and 40 from J), respectively. For CD, 54 and 60 of the 93 neurons (58 and 65%) were included in the spatial-bias analysis and reward-bias analyses (43 and 44 from J), respectively.

The 105 DLPFC neurons used in the spatial-bias analysis were recorded over 62 sessions. Given our neuron inclusion criteria, there was at least one UNDET-type reward/delay combination in each session, but only 30 of the 62 sessions (48%) also contained WDET-type trials. The 118 DLPFC neurons used in the reward-bias analysis were recorded over 69 sessions. WDET-type trials appeared in 35 of those 69 sessions (51%). The 52 VS neurons were recorded over 51 sessions, 34 of which (65%) also contained WDET-type trials. The VS did not contain more reward-bias neurons than expected by chance and the CD did not contain more spatial- or reward-bias neurons than expected by chance (Table S1). So those breakdowns are not included here.

For each of the participating neurons, we counted the number of spikes in the 1.5 s before cue onset—during the last 0.5 s of the inter-trial interval and the entire 1 s fore period (Figure 1A)—for every participating trial. Spike bursts—action potentials fired less than 10 ms apart—were consolidated into a single spike (Bair et al., 1994) (omitting this step did not affect any of our main result). We also computed a spike density function (SDF) for each spike train by convolving it with a 100 ms Gaussian window (Figures 2–4), which was used for plotting purposes only.

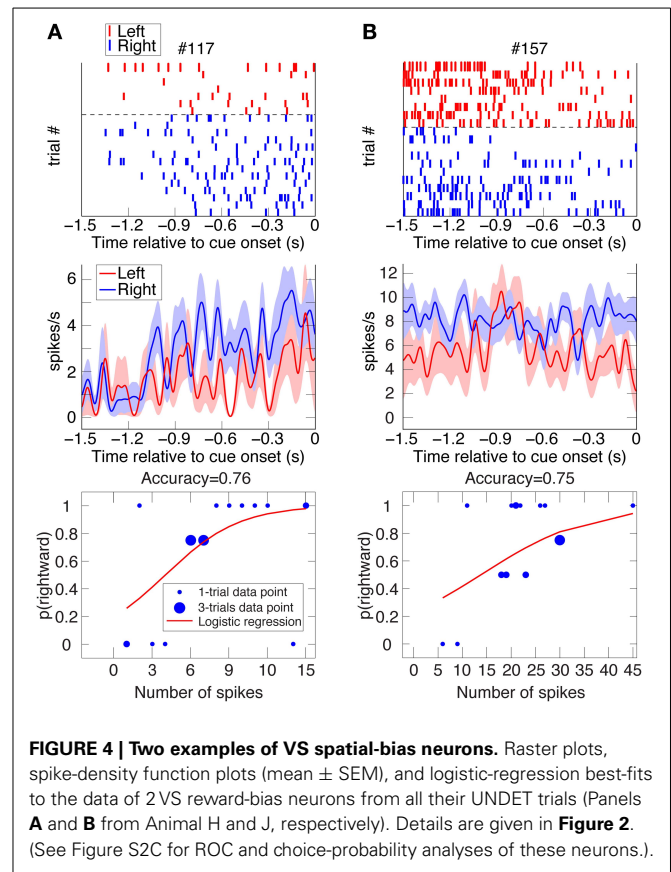




Computing prediction accuracy, proportion explained, and bias neurons

Counting the number of spikes during the 1.5 s prior to cue onset in each trial, we calculated the ratio of trials culminating in rightward saccades or high-reward choices to overall trials for each unique number of spikes (**Figures 2–4**). Analyzing each neuron separately, we then used binomial logistic-regression in a leave-one-out cross-validation procedure—training the model on all but one trial and testing it on the remaining, statistically independent trial—to predict rightwards saccades or large-reward choice for the remaining test-set trial. We labeled left (low-reward) as 0 and right (high-reward) as 1, and predicted a leftward (low-reward) choice for a trial if the predicted probability from the logistic regression was <0.5 and rightward (high-reward) otherwise (a log-likelihood-ratios analysis of the predicted probabilities suggested that 0.5 was the optimal threshold to use for our data). The prediction accuracy for each neuron was defined as the ratio of correctly predicted trials to overall trials.

For the spatial-bias analysis, the numbers of left and right choices for each neuron were not generally similar (**Figures 2–4**), while the number of high- and low-reward choices for each neuron tended to be more similar. Therefore, chance-level prediction varied among the neurons according to the behavioral preference of the animal, and we computed it using bootstrap resampling of the original data. For each neuron we randomly shuffled the left and right labels or high- and low-reward labels among the trials and computed the prediction accuracy with the same logistic



regression method. We repeated this procedure 10,000 times to create an expected distribution of chance prediction-accuracy (Efron and Tibshirani, 1993). The mean of this chance distribution was designated the chance-level prediction-accuracy for this neuron.

We computed the proportion of the variance that remained unexplained above chance level and was explained by the activity of the neuron—termed *proportion explained*—for all trial types. More specifically, designating the prediction accuracy by a_i and chance level by c_i for neuron i , we calculated $(a_i - c_i)/(1 - c_i)$, where the denominator, $(1 - c_i)$, is the variance left unexplained by the chance level. For example, in a session where the monkey saccaded to the left 65% of the time, a neuron predicting 75% of the trials correctly would have a proportion explained value of $(0.75 - 0.65)/(1 - 0.65) = 0.29$. Neurons for which the prediction accuracy was smaller than chance ($a_i < c_i$) therefore had negative proportion explained. We used this proportion explained as the basic measure of accuracy in our data for plotting purposes (**Figures 5A–C, 8**), though not for finding bias neurons (see below). We also then calculated in which percentile of the chance-level distribution lay the prediction accuracy, and for how many neurons the prediction accuracy lay in the top 50-percentile of their chance-level distribution (**Figure 6**).

We sought neurons whose pre-cue activity—i.e., activity before the decision alternatives were revealed—was significantly predictive of the animal's choice after cue onset. To accomplish this, we focused on UNDET trials and tested for each

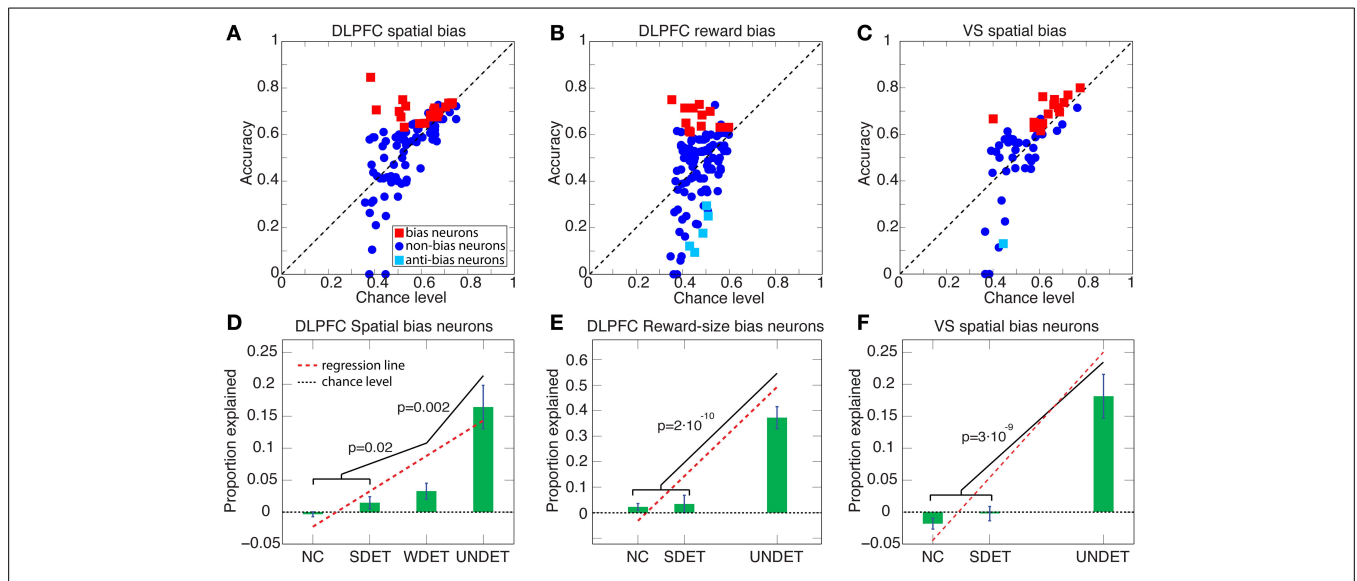


FIGURE 5 | Accuracy vs. chance level and proportion explained above chance level for DLPFC spatial- and reward-bias neurons and VS spatial-bias neurons. (A–C) Scatter plots of accuracy vs. chance level are depicted for UNDET trials. Bias neurons and anti-bias neurons (see **Figure 6**) are designated in red and cyan squares, respectively. **(D)** Average (\pm SEM) proportion explained above chance level over DLPFC spatial-bias neurons during UNDET, weakly determined (WDET), strongly determined (SDET), and no-choice control (NC) trials. The proportions explained for UNDET (0.16 ± 0.03 ; \pm SEM) and WDET (0.03 ± 0.01) are significantly greater than expected by chance, while those at SDET (0.01 ± 0.01) and NC (-0.005 ± 0.004) are at chance (paired one-tailed t -test, $p = 2 \cdot 10^{-5}$, $p = 0.01$, $p = 0.09$, and $p = 0.9$, respectively). In addition, the slope of the regression line through the four trial types was significantly positive ($p < 2 \cdot 10^{-8}$). **(E)** Average (\pm SEM) proportion explained above chance level over DLPFC reward-bias neurons during UNDET, SDET, NC trials. The proportion explained for UNDET (0.37 ± 0.04) was

significantly greater than expected by chance while the SDET (0.03 ± 0.03) and NC (0.02 ± 0.01) ones were not (paired one-tailed t -test, $p = 2 \cdot 10^{-6}$, $p = 0.17$, and $p = 0.06$, respectively). The slope of the regression line through the data was also significantly positive ($p < 4 \cdot 10^{-7}$). WDET trials were not included because there were too few trials—only 5 of 12—for meaningful statistical analysis. **(F)** Average (\pm SEM) proportion explained above chance level over VS spatial-bias neurons during UNDET, SDET, and NC trials. The proportion explained for UNDET (0.17 ± 0.03) was significantly greater than expected by chance, while the SDET (-0.004 ± 0.01) and NC (-0.02 ± 0.01) values were not (paired one-tailed t -test, $p = 8 \cdot 10^{-5}$, $p = 0.65$, and $p = 0.9$, respectively). The slope of the regression line through the data was also significantly positive ($p < 1 \cdot 10^{-6}$). WDET trials were not included because there were too few trials—only 9 of 14—for meaningful statistical analysis. P -values in the figure are for one-tailed t -tests. (See **Figure 6** for histograms of the percentiles of the prediction accuracies in the chance-level distributions.).

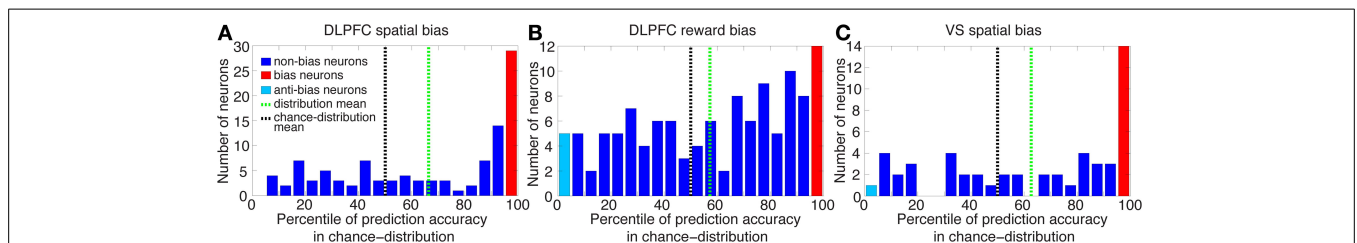


FIGURE 6 | Histograms of percentiles of prediction accuracies in the distributions of chance level over the entire neuron populations. For each neuron in the population we calculated its prediction accuracy, and then by repeatedly and randomly shuffling its left/right or large-/small-reward labels also its chance-level distribution. We then computed in which percentile of its chance-level distribution the neuron’s prediction accuracy lay. A histogram of these percentiles over all neurons is plotted here. Neurons at the top/bottom 5-percentile are designated bias/anti-bias neurons, respectively. **(A)** The mean percentile of all neurons involved in DLPFC spatial-bias analysis ($n = 105$) is 66% (green dotted line), and is significantly greater than chance-distribution mean (at 50%, black dotted

line; Wilcoxon test $p = 2 \cdot 10^{-7}$). Also, 69 of the 105 DLPFC neurons are in the top 50-percentile of the chance-distribution mean, which is significantly more than expected by chance (binomial test $p = 8 \cdot 10^{-4}$). **(B)** The mean percentile of all neurons involved in DLPFC reward-bias analysis ($n = 118$) is 57% and significantly greater than the chance-distribution mean (Wilcoxon test $p = 0.01$); 70 of 118 neurons are in the top 50-percentile, significantly more than expected by chance (binomial test $p = 0.026$). **(C)** The mean percentile of all neurons involved in VS spatial-bias analysis ($n = 52$) is 63% and significantly greater than chance-distribution mean (Wilcoxon test $p = 0.004$); 33 of 52 neurons are in the top 50-percentile, significantly more than would be expected by chance (binomial test $p = 0.035$).

neuron whether its pre-cue prediction accuracy was significantly above chance at $p = 0.05$. Therefore, for each neuron we tested whether its UNDET prediction accuracy was among the top 5% of the 10,000 UNDET chance-level values calculated in the

bootstrapping procedure above, for the spatial-bias and reward-bias analyses (**Figure 6**). Neurons that passed this test were designated *bias neurons* and their spiking activity was termed *bias activity*. We tested this bias activity over the pre-cue, cue, and go

periods. The go period is defined as the first 500 ms after the onset of the go signal (**Figure 1**). Similarly, neurons whose UNDET prediction accuracy was among the bottom 5% of chance values (0 of 105, 0%, 5 of 118, 4%, and 1 of 52, 2%, for DLPFC spatial-bias, DLPFC reward-bias, and VS spatial-bias analyses, respectively; all within the proportions expected by chance; **Figures 5A–C, 6**) were designated *anti-bias neurons*.

Error-trials and standardized firing-rate analysis

Trials where both rewards were offered with no delay yet the monkey selected the small immediate reward over the large immediate reward occurred rarely (3% of trials with immediate rewards), and were designated *error trials*. We wanted to understand whether the firing rate of bias neurons for trials in which their preferred direction aligned with the target associated with the smaller reward was higher than over the rest of the trials, which might suggest that the bias activity contributes to committing the error. We therefore searched for such trials recorded for bias neurons and found 24 trials distributed over 7 of the 29 DLPFC spatial-bias neurons, but no trials involving striatal spatial-bias neurons. In order to combine the firing rate across trials from different neurons with different baseline firing rates, we standardized the firing rates by transforming them into z -values (i.e., for each neuron separately, we subtracted the mean from each firing rate and divided by the standard deviation).

CIRCUIT MODEL AND SIMULATIONS

Two neural network circuits were constructed to capture the essential aspects of the neural dynamics of decision-making. The first is a simpler circuit encompassing only spatial bias activity (**Figure 9A**). The second is full model that includes both spatial and reward-bias activities (**Figure S3A**), and is described below. The simpler circuit is composed of 5 mean-rate units, modeled as non-saturating half-wave rectified units of the form $f(x) = \max(x - T, 0)$, where $T = 1$ is a constant activation threshold. Each unit represents the average firing rate of a network of biological spiking neurons (Wong and Wang, 2006). Three of these units (x_L , x_R , and h) form a simple winner-take-all (WTA) network. The other two units bias the activity of the WTA dynamics. The network reaches a decision through competition between excitatory units x_L and x_R ,—which stand for left and right choice, respectively, and model the population of choice neurons. The choice units are recurrently excitatory (with weight $\alpha = 1.5$) and mutually inhibitory through a shared inhibitory unit h (the connections to/from h have the weight $\beta_1 = 3/\beta_2 = 0.4$, respectively). This WTA circuit thus suppresses/amplifies the weaker/stronger external input I , respectively (Douglas et al., 1995). After convergence only x_L or x_R remain active. The two inputs, I_L and I_R , represent the temporally discounted values of the sensory evidence for choosing either the left or right target, respectively. Without any bias activity, the unit receiving the highest discounted value will win (**Figure 9B**) and determine the direction of the animal's saccade (e.g., through projections of x_L and x_R downstream to motor areas like the frontal eye field or superior colliculus) (Lo and Wang, 2006).

To model the population of spatial bias neurons, we added excitatory bias units, p_L and p_R (implemented as pointer

neurons; Hahnloser et al., 1999), which were bi-directionally connected to x_L and x_R , respectively (with weight $\delta = 0.1$). The dynamics of the choice units are modeled as $\tau \dot{x}_i + Gx_i = f(I_i + \alpha x_i + \delta p_i - \beta_1 h - T_i)$, and similarly for the other units (see below for detailed equations). Here $\tau = 0.01$ s is the membrane time constant and $G = 1$ is the load. Activation of a bias unit introduces an additional recurrent feedback loop to its choice unit, and thus the effective gain of that choice unit is dynamically increased relative to the other choice unit. Hence, a choice unit with external bias input can win the competition despite having a smaller input value relative to the other choice unit (**Figure 9C**). The stronger the bias input b_i , the more the competition is biased in favor of its corresponding state (**Figure 9D**). We assume that their activity fluctuates slowly between trials and stems from the positive part of a normal distribution with mean 0 and a small standard deviation. Numerical simulations were performed using Euler integration with $\Delta = 0.01$ s in Matlab (Mathworks, Inc.), with α , β , and δ fixed at the above values. This circuit has a high gain but is nevertheless stable through the balance of excitation and inhibition for a wide range of parameters (Rutishauser et al., 2011).

Steady-state analysis of the spatial-bias circuit model

Given the neural network schematized in **Figure 7**, the dynamic state equations for $i = L, R$ are:

$$\tau \dot{x}_i + Gx_i = f(I_i + \alpha x_i + \delta p_i - \beta_1 h - T_i), \quad (\text{A.1})$$

$$\tau \dot{h} + Gh = f\left(\beta_2 \sum_{j=L}^R x_j - T_h\right), \quad (\text{A.2})$$

and

$$\tau \dot{p}_i + Gp_i = f(\delta x_i + b_i - T_{p_i}) \quad (\text{A.3})$$

Here b_i is the constant bias input into the bias unit p_i , which has a two-way connection with its corresponding choice unit x_i , for $i = L, R$. I_i is the external input variable over time into x_i for $i = L, R$. The firing rate activation function of the units is the non-saturating rectification nonlinearity:

$$f(z) = \max(0, z) \quad (\text{A.4})$$

If we take $G = 1$ and $T_i = T_h = T_{p_i} = T$ for $i = L, R$, the steady-state value of the winning unit x_w (where w is either L or R) for constant input $I_w > 0$ and a constant bias $b_w > 0$ (that is sufficient to activate the bias unit p_w) is:

$$x_w = \frac{I_w + \delta(b_w - T) + T(\beta_1 - 1)}{1 - \alpha + \beta_1 \beta_2 - \delta^2}. \quad (\text{A.5})$$

Compare this to:

$$x_w = \frac{I_w + T(\beta_1 - 1)}{1 - \alpha + \beta_1 \beta_2} \quad (\text{A.6})$$

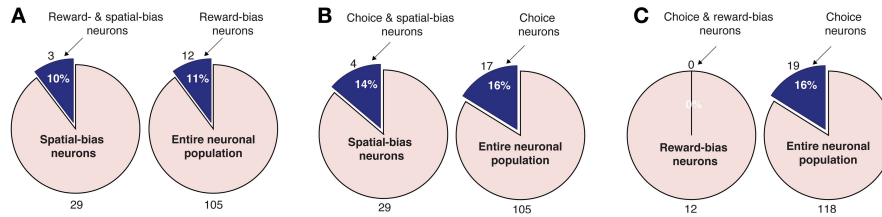


FIGURE 7 | Pairwise intersections of spatial-bias neurons, reward-bias neurons, and choice neurons in the DLPFC. (A) The proportion of reward-bias neurons among spatial bias neurons is not significantly different than the proportion of reward-bias neurons among the entire neuronal population ($p = 0.98$, χ^2 -test). **(B)** The proportion of choice

neurons among spatial-bias neurons is not significantly different than their proportion among the entire neuronal population ($p = 0.73$, χ^2 -test). **(C)** The proportion of choice neurons among reward-bias neurons is not significantly greater than their proportion among the entire neuronal population ($p = 0.87$, χ^2 -test).

[see also Equation 2.3 in Rutishauser and Douglas (2009)]. Equation A.6. is the steady state activity of the winner choice unit in a network without a bias (i.e., when $\delta = 0$).

Equation A.5 therefore shows that inserting the bias neurons adds two terms to Equation A.6. One is a constant linear term $\delta(b - T)$ multiplied by the gain, $1 / (1 - \alpha\beta_1\beta_2 - \delta^2)$, (the gain is the amount by which the choice unit, x_i , amplifies its input, I_i). The other is a modification (increase) of the gain by the subtraction of δ^2 from its denominator. Functionally, for the choice unit whose corresponding bias unit is active (that is for p_i , $i = L$, or R , where $\delta x_i + b - T > 0$), the gain is higher and thus the unit is more likely to win. A choice unit whose bias unit is not active effectively has $\delta = 0$ and is described by Equation A.5, thus ending up with a reduced gain. Hence, for a fixed δ , the only non-linearity is when the bias unit switches from inactive to active; thereafter the network is fully linear in b (though non-linear as a function of δ).

The full circuit-model

The full circuit model (see Figures S5E,F), encompassing both the spatial and reward-size biases, is an extension of the reduced—spatial-bias-only—circuit-model described in the Materials and Methods (Figure 7). It includes six additional units on top of the five making up the reduced circuit model. Two are the reward-size bias units q_{lo} and q_{hi} , representing a bias toward the low and high rewards, respectively. These units receive inputs c_{lo} and c_{ho} , respectively. The other four units are devoted to the conversion of low- and high-reward biases into motion-applicable left or right biases after cue onset. Thus, q_{lo} is connected d_{loL} and d_{loR} and q_{hi} is connected d_{hiR} and d_{hiL} , all with connection strength $\delta_3 = 5$. Then, d_{loL} and d_{hiL} are bi-directionally connected to x_L while d_{hiR} and d_{loR} are bi-directionally connected to x_R , all with connection strength $\delta_3 = 0.1$. It is important to keep the reward-size bias activity from affecting the choice neurons until the cue signal, which pairs the high and low rewards with left and right targets, reaches the circuit. This pairing is achieved by turning on one of two input signals—either $v_{hi=L}$ or $v_{hi=R}$. Here $v_{hi=L}$ inputs into d_{hiL} and d_{loR} while $v_{hi=R}$ inputs into d_{loL} and d_{hiR} . This is also achieved by keeping the activation threshold of the 4 d units a higher than the rest of the units, at $T_d = 8$. This structure enables the reward-size bias to counter and even later overturn the effects

of the spatial bias and the inputs on the choice units (see Figure S5F).

RESULTS

PREFRONTAL AND STRIATAL ACTIVITY RELATED TO BIASES IN DECISION MAKING

The choices of animals that were free to select between a delayed, larger reward and an immediate, smaller reward (Figure 1A) indicated that they increasingly discounted the value of a reward with longer delay (Kim et al., 2008; Hwang et al., 2009) (Figure 1B, Figure S1). Our goal was to study the extent to which the spiking activity of DLPFC, VS, and CD neurons before cue onset predicts the subsequent decision of the animal to select one or the other alternative and how that activity was integrated with the post-cue choice activity. Hence, we tested for correlation and not causation. We hypothesized that such pre-cue bias activity was dissociated from the post-cue choice activity, but that the two would be combined to make a decision at the go signal. Yet, this bias activity would have a weaker influence on the ultimate decision than the choice activity. So the bias activity would mainly have an effect when the values of the decision alternatives were similar.

We therefore divided all trials into four types according to the similarity of the discounted values of the targets, which is to say how much the animal's decision was determined by the stimulus. Trials with most similar target values, and thus relatively undetermined by the reward/delay stimulus, were designated UNDET (see Materials and Methods and Figures 1B, S1). Trials with somewhat similar target values were designated WDET. And trials with more divergent target values were designated SDET. Last, control trials, where the animals were instructed what target to select, were designated *no-choice trials* (NC; see Materials and Methods for details). For neurons coding a bias, the pre-cue activity level (spiking rate) should thus maximally influence decisions between eventual leftward- vs. rightward-movement (or smaller vs. larger reward) in UNDET trials, and that influence should progressively decrease for WDET and then SDET and NC trials.

We investigated two types of bias. The first is a spatial bias, favoring either the target on the left or on the right. The second is a reward-bias, favoring the target with the smaller or larger reward. For each neuron, we calculated the accuracy (proportion of trials with correctly predicted decisions from overall trials) over the four trial types using logistic-regression in a leave-one-out

cross-validation scheme (see Materials and Methods; **Figures 2–4** show raster plots, spike-density functions and the best-fit logistic regression for DLPFC spatial bias, DLPFC reward bias, and VS spatial bias, respectively). As we were interested in bias activity, we focused on the 1.5 s before cue onset. Using bootstrapping, we then computed the chance level for each neuron, from which we could compute the proportion of the variance that the neuron's activity explained above chance level for every trial type (hereafter, *proportion explained*; see Materials and Methods). We describe the results from 105 DLPFC neurons, 52 VS neurons, and 54 CD neurons for spatial-bias calculations. For the reward-bias analysis we included 13, 4, and 6 additional DLPFC, VS, and CD neurons that met the analysis criteria, respectively (see Materials and Methods).

We expected the pre-cue firing rate of bias neurons (henceforth *bias activity*) to be most predictive of the decision during UNDET trials. The accuracy in UNDET trials for spatial bias was significantly greater than chance for 29 DLPFC neurons (28%) and 14 VS neurons (27%), while 12 DLPFC neurons (10%) showed significant reward bias (at $p = 0.05$, see Materials and Methods; **Figures 5, 6**, Table S1). These were designated *DLPFC spatial-bias neurons*, *VS spatial-bias neurons*, and *DLPFC reward-bias neurons*, respectively. These numbers of bias neurons are significantly larger than expected by chance ($p = 3 \cdot 10^{-14}$, $p = 2 \cdot 10^{-7}$, and $p = 0.015$, binomial test, respectively; **Figures 5A–C**). As a control, we ran the same bias-neurons search procedure for SDET trials, where the stimulus should be the overwhelming influence on the decision. We therefore did not expect to be able to extract a significant number of bias neurons there. Indeed, we found only 8 DLPFC spatial-bias neurons and 3 VS spatial-bias neurons, which is within what would be expected by chance ($p = 0.16$ and $p = 0.49$, binomial test, respectively).

To make sure that our results are not due to our selection criteria for bias neurons, we also examined the entire neuronal populations for spatial and reward biases in the DLPFC and for spatial bias in the VS. We calculated the percentile of the prediction accuracy of each neuron in the population in its chance-level distribution (see Materials and Methods). And we found that the mean of the percentiles for the 105 neurons that took part in DLPFC spatial-bias calculations, the 118 neurons that participated in DLPFC reward-bias calculations, and the 52 neurons that were involved in VS spatial-bias calculations was 66, 57, and 63%, respectively; all significantly greater than the corresponding chance-distribution means ($p = 2 \cdot 10^{-7}$, $p = 0.01$, and $p = 0.004$, respectively; Wilcoxon tests; **Figure 6**). The numbers of neurons in each population in the top 50-percentile—69 of 105 for DLPFC spatial-bias analysis, 70 of 118 for DLPFC reward-bias analysis, and 33 of 52 for VS spatial-bias analysis—were also significantly above chance for all three populations ($p = 8 \cdot 10^{-4}$, $p = 0.026$, and $p = 0.035$, respectively; binomial tests; **Figure 6**).

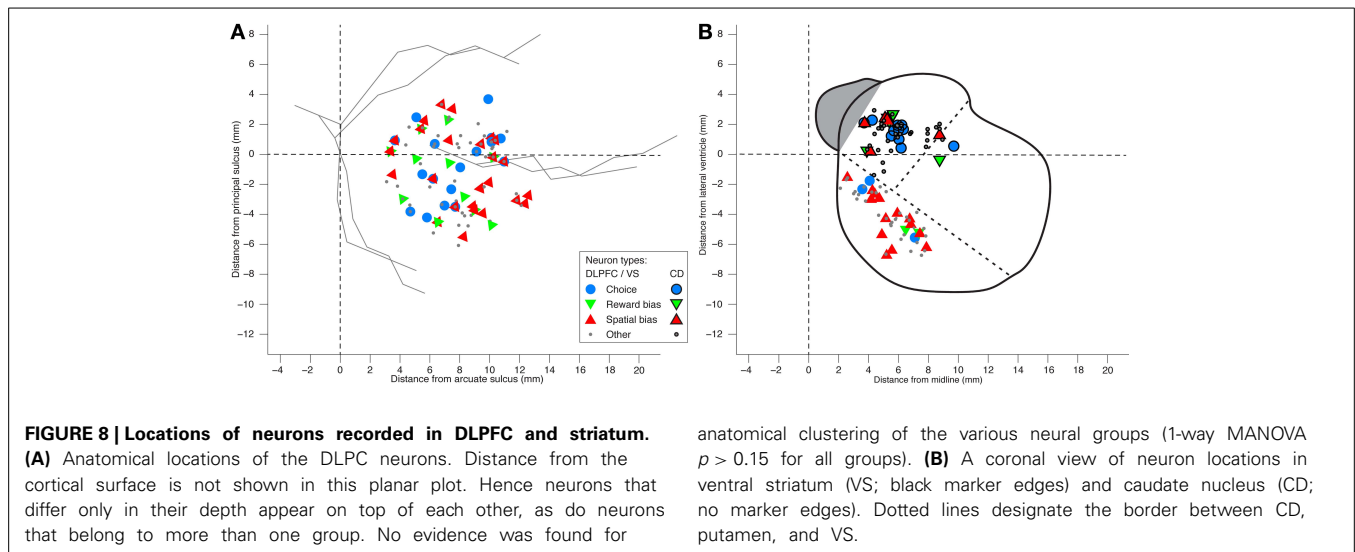
In addition, we expected a progressive decrease in the predictive power of the bias activity as the difference in the temporally discounted values of the two targets increased. We therefore calculated the predictive power of the bias neurons also over WDET, SDET, and NC trials. And we found the expected decrease in our data from UNDET trials, through WDET, to SDET and NC trials (**Figures 5D–F**). We conclude that the neural activity of the

29 DLPFC and 14 VS spatial-bias neurons as well as 12 DLPFC reward-bias neurons predicted the upcoming choice of movement direction or reward size before any information about the reward-delay combination of the alternatives became available to the animal.

We also analyzed the data in search of choice neurons. Using a regression model that included the temporally discounted values of both targets, we discovered that 17 of the 105 DLPFC neurons (16%; $p = 2 \cdot 10^{-5}$, binomial test) and 12 of the 54 CD neurons (22%; $p = 1 \cdot 10^{-5}$, binomial test) had significantly altered activity during the cue period (**Figure 1A**) according to the animal's choice of target (Table S1) (Kim et al., 2008; Cai et al., 2011). We refer to them as *choice neurons*. It was of interest to investigate to what extent spatial-bias neurons were also reward-bias neurons and choice neurons, and vice versa. In the DLPFC, the intersection of spatial- and reward-bias neurons with the choice neurons was not significantly larger than would be expected by chance when these effects were combined independently ($p = 0.73$ and $p = 0.87$, respectively, χ^2 -test; **Figures 7B,C**), indicating that choice neurons were not preferentially also bias neurons of either kind. Moreover, we found no evidence for anatomical clustering of spatial or reward-bias neurons or of choice neurons in the DLPFC of the two animals (**Figure 8A**). Unlike the DLPFC, in the striatum we found that only VS contained a significant number of spatial-bias neurons, and the proportion of spatial bias neurons was significantly higher in VS than in CD (**Figure 8B**, Table S1). In addition, only the CD contained a significant number of choice neurons, and the proportion of such neurons was again significantly greater than that of these neurons in VS. Neither VS nor CD contained a significant number of reward-bias neurons (**Figure 8B**, Table S1).

False-negatives analysis

We used logistic-regression in a leave-one-out cross-validation procedure to keep the false-positives rate for discovering bias neurons low. This comes at the expense of an increased probability of mistakenly rejecting neurons that are actually bias neurons, thus increasing the false-negative rate. We therefore developed a method to estimate a lower bound on this false-negatives rate. Given our recording procedure, we generally have access to only a single neuron at a time, and always to only one bias neuron at a time. And in our limited sample of trials, activity patterns of a true bias neuron might, due to noise for example, not coincide with that of the overall bias population that—according to our model—determines the decision of the monkey. To assess this false-negatives rate, we calculated the preferred and non-preferred reward size and the number of trials associated with each DLPFC reward-bias neuron. Taking the mean firing rate and number of trials for the preferred and non-preferred reward sizes, we ran a simulation, sampling from a Poisson process with these firing rates and number of trials over a 1.5 s duration (the duration of our pre-cue period). We ran our bias-neurons identification procedure on the simulated results to test whether such a simulated reward-bias neuron would be identified by our procedure. We ran this bootstrapping procedure 1000 times, and for the DLPFC reward-bias neurons found that in only 70% of these runs the simulated neuron was identified as a bias neuron. Therefore,



our population of 118 neurons involved in reward-bias analysis may in fact contain $12/0.7 \approx 17$ reward-bias neurons instead of only 12 bias neurons.

Importantly, other factors that are harder to simulate may increase the false-negatives rate: The recorded neuron might belong to, say, the high-reward bias population and we may happen to record a majority of trials where the monkey opted for the low reward in our limited sample. Alternatively, interactions between reward-bias neurons and spatial-bias neurons may cancel each other out, especially in a limited sample (Figure S3B). Therefore, the 70% of trials where a simulated reward-bias neuron was correctly identified should be viewed as an upper bound on the false-negatives rate inherent in our procedure, and it is possible that there are more than 17 reward-bias neurons in the recorded neuronal population.

POTENTIAL SOURCES OF THE BIAS ACTIVITY

Our analysis suggests that a significant number of DLPFC and VS neurons show single-trial pre-cue biases. We tested and excluded several possible sources of this bias activity. If the animal's current choice, especially for UNDET trials, is significantly correlated with its choice in the preceding trials, the pre-cue biases might be a reflection of persistent activity associated with the choice in the previous trial or of behavioral choice patterns across trials. But we found no evidence for correlations between choices in UNDET and preceding trials, nor between choices when examining only UNDET trials (as far as 10 trials back; Supplemental Text, Figure S4). In addition, we often found left/right and small-/large-reward differences in the bias activity throughout the 1.5 s pre-cue period (see examples in Figures 2–4). Had the bias activity reflected lingering neuronal activity from the previous trial, we would have expected to find these differences mainly in the beginning of the pre-cue period. This also suggests that it is not only bias activity close to cue onset that affects the animal's eventual decision.

For each neuron we also calculated whether its firing rate tended to be higher for all trials culminating in leftward (or

rightward) saccades for spatial-bias neurons or in saccades toward the smaller (or larger) reward for reward-bias analysis, during the pre-cue epoch. We found that 17 of the 29 DLPFC spatial-bias neurons and 9 of the 14 VS spatial-bias neurons were more active toward the left, and that 7 of the 12 reward-bias neurons were more active toward the small reward; none of these are significantly different from an even split ($p = 0.23$, $p = 0.21$, and $p = 0.39$, binomial test, respectively). We therefore found no evidence that bias neurons had a preferential preferred direction.

Lastly, in the DLPFC, 8 (9) of the 11 (18) bias neurons that were recorded from the left (right) hemisphere—i.e., recorded from animal D (J) (see Figure 6A)—had higher firing rate toward the left (right), across all trials in the pre-cue epoch. There was therefore no significant correlation between the hemispheric laterality and the target toward which the neuron was more active ($p = 0.34$ and $p = 0.29$, χ^2 -test of 8/11 vs. 17/29 and of 9/18 vs. 11/29, respectively). So we could not find evidence for correlation between the brain hemisphere from which the DLPFC neurons were recorded and their preferred direction. All VS neurons were recorded from the right hemisphere (Figure 6B). Following this and the even split for left/right preferred directions reported above, our data shows no spatial preference in terms of the recorded brain hemisphere and the behaviorally chosen side.

PREDICTIONS FROM A CIRCUIT MODEL

To better understand the nature of this pre-cue bias activity, we developed a plausible neural circuit model of a biased WTA network (Figure 9A, Materials and Methods). Unlike a simple WTA circuit where the identity of the winner depends only on the inputs (Figure 9B), in our model the activity of a divergent population of bias units also influences the winner selection (Figures 9C,D). Hereafter we refer to the elements of the circuit model as *units* and to actual DLPFC or BG cells as *neurons*; each unit represents a population of neurons. The model had two versions—one that focused on the spatial bias activity alone (Figure 9) and another, expanded version encompassing both spatial and reward biases (Figures S3A,B). The model does not

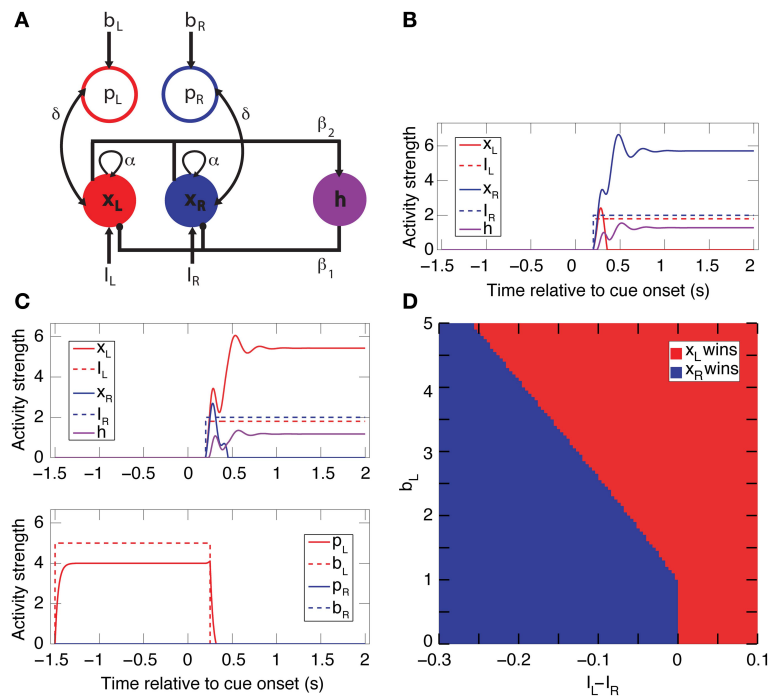


FIGURE 9 | A simple biased WTA circuit model accounts for the spatial-bias data and makes additional predictions. (A) A circuit for spatial-bias activity. Units x_L and x_R designate populations of left and right choice neurons, which are self-excitatory and mutually inhibitory (through inhibitory neuron h). Each choice unit is bi-directionally connected to a bias unit (p_L and p_R). Excitatory (inhibitory) connections are designated by arrows (circles). **(B)** Without bias input ($b_L = b_R = 0$), this is a simple WTA network, so for inputs $I_R > I_L$ ($I_R = 2$ and $I_L = 1.8$, in this case) x_L 's activity diminishes to

zero and x_R 's increases and then stabilizes. **(C)** When the left bias is turned on ($b_L = 5$, $b_R = 0$), x_L now wins for the inputs in B. Note that the bias activity was turned off (at $t = 250$ ms) shortly after the onset of the competition (at $t = 200$ ms). **(D)** A phase-space diagram simulating various values for b_L and $I_L - I_R$ when $b_R = 0$ and $I_R = 2$. The separatrix between the basin of attraction of x_L winning and x_R winning is constant until approximately 1, when b_L passes its activation threshold, and then linear. (See Figure S3 for the full model, including both spatial- and reward-bias activity).

incorporate noise for the sake of simplicity. Nevertheless, simulations we performed indicate that, even in the presence of noise, the results remain the same on average.

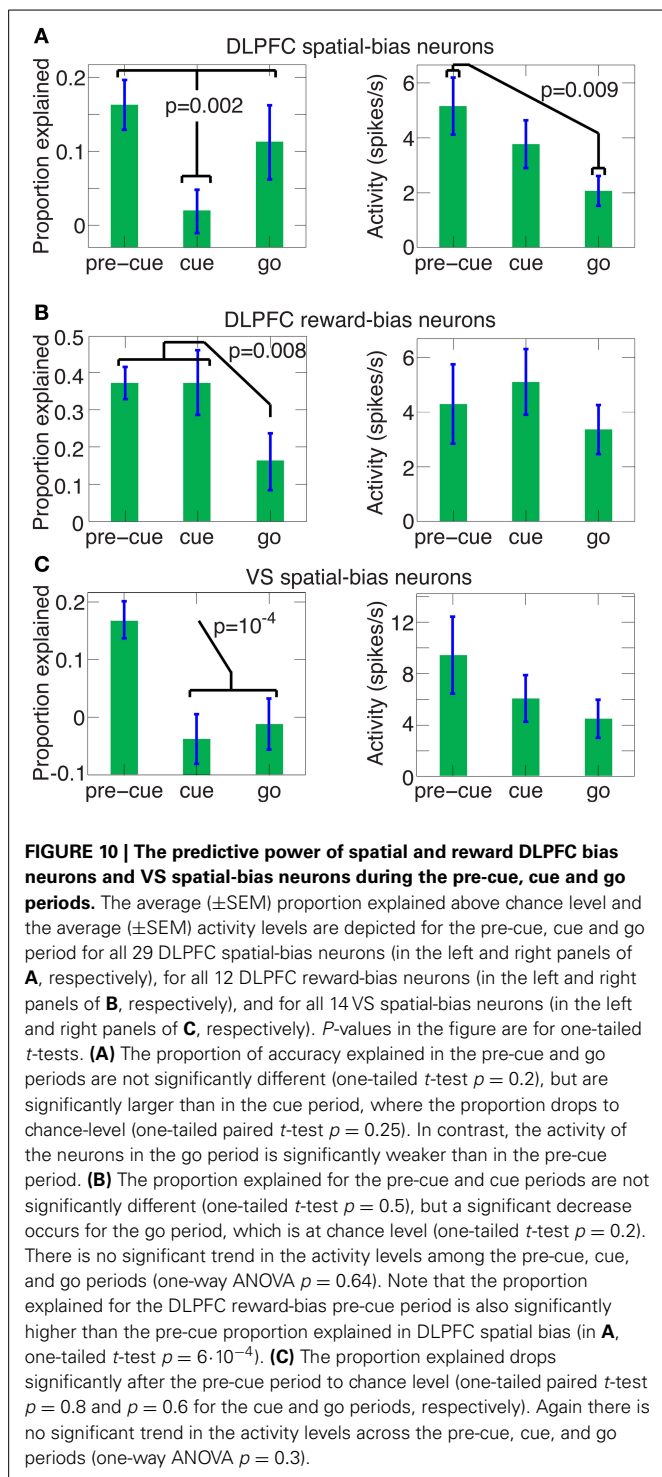
A main thrust of our model is that the bias and choice units are dissociated, as indicated by our data: We discovered that bias neurons are not preferentially also choice neurons in DLPFC (Figures 7B,C, 8A). And we found that, of the three types of neurons, there are only spatial-bias neurons in significant numbers in VS and only choice neurons in significant numbers in CD (Figure 8B). Our simulations also demonstrate that it is sufficient for the spatial-bias units to be active at the very beginning of the competition to influence the identity of the final winner: After giving an initial boost to the network dynamics, the WTA network amplifies any initial small differences. We therefore stop the input to the model's bias units shortly after the onset of the inputs to the choice units (Figure 9C). Our circuit model thus suggests that the prediction accuracy of the spatial-bias neurons may decrease with time during the cue period. This prediction is of interest because it relates to the manner by which the bias and choice activities are combined after cue onset by postulating how bias neurons would behave after cue onset, when rational deliberation can begin. We tested this model prediction by further analysis of our experimental data during the cue period and the

first second after the go signal, and found it to be true. For DLPFC spatial-bias neurons, there was no significant decrease in the proportion explained between the pre-cue period and go period, while the proportion explained during the cue-period decreased to chance level (Figure 10A). For VS spatial-bias neurons, the proportion explained decreased to chance level during both the cue and go periods (Figure 10C). We think that this might be because DLPFC spatial-bias neurons may be more involved in the motor aspects of the task than VS spatial-bias neurons.

For reward biases, in contrast, the mapping between the high and low rewards and the left and right targets is only revealed by the stimulus. So, reward-biases can be translated into spatial coordinates only after stimulus onset (see full model, Figure S3A). To influence the decision effectively, the reward-bias thus needs to remain active longer than the spatial bias (Figure S3B). The neural data we gathered also reflects this; the DLPFC reward-bias neurons are the only ones remaining predictive during the cue-period (compare Figure 10B and Figures 10A,C).

"Error-trials" analysis

In value-based decisions, it is generally not possible to know with certainty whether the chosen target was associated with the subject's larger discounted value or not. Except in some special



circumstances (see below), it is thus not possible to declare a choice correct or incorrect. However, when both the small and large rewards are offered with no delay (an SDET trial), animals maximizing the temporally discounted value of reward (Figure 1B) should always select the larger immediate reward (which they did for 97% of 1536 and 99% of 1698 of such DLPFC and striatal trials, respectively). Following the model prediction

above, immediate reward trials where the animal nevertheless chose the small reward—hereafter *error trials*—should be associated with especially strong bias activity, and hence with higher firing rate of the bias neurons. Consistent with this prediction, the mean (\pm SEM) standardized firing rate (see Materials and Methods) over all the DLPFC spatial-bias neurons whose preferred direction aligned with the target associated with the smaller discounted value during error trials was 0.68 ± 0.44 , which was significantly higher than their standardized firing rate during the rest of the trials, at 0 ± 0.03 (one-tailed *t*-test $p = 0.008$; the value of 0 for the mean firing rate results from the standardization procedure, see Materials and Methods). The error trials above were distributed throughout the experimental sessions, suggesting that the animal's motivation level and alertness did not factor into the firing rate difference. The striatum data contained too few error trials for similar statistical analysis.

DISCUSSION

We examined single-neuron activity from the DLPFC, VS, and CD of monkeys that freely choose between a delayed, larger reward and a more-immediate, smaller one. Our working hypothesis was that such activity before the cue revealed the decision alternatives—and thus before rational deliberation could begin—had some influence on the ultimate choice of the animals. We speculated that such bias activity was dissociated from the post-cue choice activity, and had a weaker influence on the decision than that choice activity. So, when it was later integrated with the choice activity, the effect of the bias activity would be mainly apparent where the reward/delay combinations would have similar values for the animal (i.e., in UNDET trials, in our case).

We found that when the animals were confronted with two targets with similar temporally discounted values for the rewards associated with the targets, more than a quarter of the analyzed neurons in DLPFC and in VS signaled the direction of the eventual eye movement before cue onset. In addition, 1 in 10 DLPFC neurons coded the eventual choice of the animal between a larger or smaller reward—or equivalently, between the red and green targets, respectively—before cue onset. In such trials, a choice between the large and small reward usually also translated into a choice between the short and long delay, respectively. Moreover, we found significant numbers of both types of bias neurons as well as of choice neurons in the DLPFC; both types of bias neurons were not preferentially also choice neurons, and vice versa. Yet, we found only spatial-bias neurons in significant numbers in the VS and only choice neurons in significant numbers in the CD. The spatial- and reward-bias activities were found, on average, over the entire DLPFC neuronal population, as was the spatial-bias activity found on average over the entire VS population. This suggests that the presence of bias neurons was not merely an artifact of our selection criteria. Our results are therefore compatible with our hypothesis.

We considered the implications of the presence of bias neurons for decision-making using a computational approach. In our task, we viewed decision making as a competition between the two choices represented by the targets. And this choice activity exerted a stronger influence on the results than the bias activity. The

function of the bias neurons was to more-weakly favor one of the alternatives, thus making that choice more likely, especially when the choice activity only marginally swayed the competition. While there are different approaches and levels of abstraction to modeling such decision-making, almost all regard making a decision as a competition. Drift-diffusion and race-to-threshold models (Smith and Ratcliff, 2004) have been used extensively and account for data acquired in perceptual discrimination tasks (Shadlen and Newsome, 2001; Mazurek et al., 2003; Palmer et al., 2005; Heekeren et al., 2008) as well as value-based decisions (Basten et al., 2010; Krajbich et al., 2010). However, their abstract mathematical nature does not easily lend itself to exploration of circuit mechanisms. In particular, it is not clear how to represent choice and bias neurons as dissociated populations (Figure 9A, Figure S3A) in these more abstract models. In contrast, biophysically based cortical microcircuit models for decision-making based on attractor dynamics explain decision-making data well (Wang, 2002; Wong and Wang, 2006), while suggesting mechanisms for the instantiation of these decisions in neural substrate (Rolls and Deco, 2011). We thus focused our modeling efforts on circuits implementing competition using shared inhibition between two pools of competing excitatory units. In this framework, bias neurons can be included naturally as additional units that bias the competition between different pools.

We derived several empirically testable predictions from our model, which we in turn tested in the data. First, the model suggested that spatial-bias, reward-bias, and choice neurons might be pairwise-divergent populations. We found that this was the case in the striatum, with the VS and CD containing only spatial-bias and choice neurons in significant numbers, respectively (Figure 8B). In the DLPFC we found that choice neurons are not preferentially also either type of bias neurons, nor are spatial-bias neurons preferentially reward-bias neurons, and vice versa (Figure 7). Moreover, we found different activity patterns for DLPFC spatial-bias, reward-bias and choice neurons across the pre-cue and cue periods (Figure 10). In that vein, the model also suggested that while spatial-biasing activity is required only in the beginning of the competition, the reward-biases would remain on longer after cue onset (Figure S3B). Also inline with our model's predictions, in "error" trials (where the animals chose a small immediate reward over a large immediate one) the firing rate of the bias neurons, whose preferred direction aligned with the target associated with the smaller temporally discounted choice, was higher than in other trials. However, as only a small number of neurons and trials were involved in this analysis, this result can only be taken as preliminary evidence that bias activity may occasionally also affect the monkey's decision when the values of the two decision options are substantially different.

In other tasks, when the most rewarding choice could be calculated before cue onset, predictive pre-cue neural activity was found in CD choice neurons (Coe et al., 2002; Lauwereyns et al., 2002; Ding and Hikosaka, 2006). In contrast, in our task where the most rewarding choice could only be deliberated upon after cue onset (due to the random shuffling of the reward size and delay among trials), we found only post-cue CD choice activity. Combined, these data are therefore consistent with the hypothesis that deliberation about an upcoming choice, be it before or

after cue onset, involves choice neurons (in DLPFC and CD, in our case), while pre-deliberation bias activity involves bias neurons (spatial ones in DLPFC and VS and reward ones in DLPFC alone).

Other studies utilized a perceptual judgment task (random-dot motion discrimination) to study choice activity. In that task, pre-deliberation spatial-bias activity was found in monkey LIP (Shadlen and Newsome, 2001; Williams et al., 2003) and CD (Ding and Gold, 2010), but not in DLPFC (Kim and Shadlen, 1999). There are three important differences between those findings and ours. First, while in perceptual judgments the reward is usually contingent upon the correct response to the perceptual cue, for value-based decisions there are seldom correct and incorrect answers and the animal is free to choose its reward. Second, previous studies reported spatial biases whereas we also found non-spatial biases that reflect preferences for reward size and are dissociated from spatial bias activity. Third, these studies focused on pre-cue bias activity in choice neurons, while we searched for bias activity in the entire neural population, and found bias activity that was dissociated from choice activity in DLPFC, and with anatomical specificity for bias and choice neurons in the striatum. We need not commit to a strict distinction between perceptual and value-based decisions, but the above may explain why we found DLPFC bias activity when previous studies did not and why we found only choice activity in CD and others found pre-cue spatial-bias activity in a small fraction of choice neurons there (Ding and Gold, 2010). Also, whereas a fluctuating bias may reduce the reward in a perceptual decision task, it may well be part of the subjective aspect of value-based decisions (Gold and Shadlen, 2007). Last, our animals might have used deliberative goal-directed judgments to accomplish their task, while subjects making perceptual decisions relied on intuitive-habitual judgments, which may involve different brain systems (Daw et al., 2005; Rangel et al., 2008).

The bias activity we recorded does not seem to be related to past behavioral patterns, nor does it have preferential preferred direction or hemispheric specificity. This is in contrast to neural activity in forced-choice memory tasks, where striatal neurons, but not DLPFC neurons, exhibit strong contralateral spatial preference (Kawagoe et al., 1998; Kobayashi et al., 2007). Moreover, the structure of our task, together with the random shuffling of the reward size and delay across trials, left no beneficial role for anticipation (Lauwereyns, 2010). Nor is it likely that visual attention played a role in bias formation, because during the pre-cue period the screen was empty but for a small central fixation target. But, as we found that DLPFC spatial-bias neurons were similarly predictive during the pre-cue and go period, yet not predictive during the cue period (Figure 10A), this DLPFC bias activity may include a motor component. Going back to the question with which we started, it seems that in our experimental setting, the bias activity is unlikely to be decision related, as it is recorded before the decision alternatives are presented and rational deliberation can begin. It may be related to top-down attention (Hopfinger et al., 2000; Changizi and Shimojo, 2008), or to other factors not immediately related to the neural mechanisms of decision making. In particular, the single-trial pre-cue bias activity we found may be related to slowly fluctuating spontaneous

brain activity (Fox et al., 2007) thought to be involved in responding to stimuli in conjunction with stimulus-evoked activity (Arieli et al., 1996; Wang, 2002; Leopold et al., 2003; Ress and Heeger, 2003; Fox et al., 2006; Vincent et al., 2006). If the bias activity is dependent on such fluctuations, it could explain why the bias and choice neurons may be distinct populations: the bias neurons may belong to a network of spontaneous activity, distinct from the decision network. A recent model of the neural substrates underlying unmotivated free-choices suggests that the onset of such choices depends on spontaneous fluctuations in neuronal activity (Schurger et al., 2012). So, assuming that the bias-neuronal activity we found is based on spontaneous fluctuations, might suggest that another aspect of free actions—action selection—also relies on these fluctuations. Neural fluctuations as the basis of our results are also congruent with the finding that inhibiting DLPFC activity in humans using transcranial magnetic stimulation reduces their ability to generate random motor sequences (Jahanshahi and Dirnberger, 1999). Moreover, it is inline with the proposed role of the BG and frontal cortex in biasing competitions among neural activations, which originate from sensory information and result in action (Cisek and Kalaska, 2010). Perhaps while the VS holds predeliberation biases, the CD informs the decisions, as they are made, through choice neurons.

As for testing the predeliberation-biases hypothesis, do decisions—especially indeliberate ones, between similarly valued options—also rely on predeliberation activity to bias the choice toward one of the alternatives? Our results appear consistent with that hypothesis. And, as the model demonstrates, even weak bias activity can end up greatly influencing the identity of the selected decision alternative. What is more, our results suggest that these bias brain-signals may not be part of a decision process that is normally involved in rational deliberation, because they happen before the decision alternatives are revealed and rational deliberation can begin. Instead, the bias signals we measured may depend on a dissociated predeliberation network. Raising the left or right hand for no reason or purpose and with little consequence (Haggard and Eimer, 1999; Haynes et al., 2007; Soon et al., 2008, 2013; Bode et al., 2011; Fried et al., 2011) is the epitome of such random, indeliberate decisions. So the weak-to-moderate prediction ability well before movement onset that is exhibited in these experiments (60–70% accuracy) might rely on such fluctuating bias activity rather than be part of the decision network normally involved in rational deliberation. It is therefore of interest that some recent attempts to model indeliberate decisions focus on the DLPFC as the locus of intentional action selection (Pockett, 2006; Pacherie and Haggard, 2011). Further research is required to shed more light on the origins of this early prediction ability.

But similarly valued decision options are not restricted only to such random choices. Arguably, some of the biggest and toughest decisions in life—such as selecting a partner or career path—are difficult also because the decision alternatives are associated with similar values (e.g., is a job that is better paying but less interesting and more demanding than the current one preferable?). Our results—to the extent they can be generalized to such situations—suggest that these decisions may well end up being considerably influenced by neural biases that are not part of the

rational decision process and start before rational deliberation can even begin.

ACKNOWLEDGMENTS

This work was funded by the Ralph Schlaeger Charitable Foundation, FSU Big Questions in Free Will Initiative, Mathers Foundation, ONR, DARPA, and NSF, as well as the Max Planck Society and NIDA (RL1 DA024855 and P20 DA027844). We are grateful to R. Rivlin, E. Portugaly, and S. Ullman for many insightful discussions.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <http://www.frontiersin.org/journal/10.3389/fnins.2013.00225/abstract>

REFERENCES

- Arieli, A., Sterkin, A., Grinvald, A., and Aertsen, A. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* 273, 1868. doi: 10.1126/science.273.5283.1868
- Bai, O., Rathi, V., Lin, P., Huang, D., Battapady, H., Fei, D. Y., et al. (2011). Prediction of human voluntary movement before it occurs. *Clin. Neurophysiol.* 122, 364–372. doi: 10.1016/j.clinph.2010.07.010
- Bair, W., Koch, C., Newsome, W., and Britten, K. (1994). Power spectrum analysis of bursting cells in area MT in the behaving monkey. *J. Neurosci.* 14, 2870.
- Barraclough, D., Conroy, M., and Lee, D. (2004). Prefrontal cortex and decision making in a mixed-strategy game. *Nat. Neurosci.* 7, 404–410. doi: 10.1038/nn1209
- Basten, U., Biele, G., Heekeren, H. R., and Fiebach, C. J. (2010). How the brain integrates costs and benefits during decision making. *Proc. Natl. Acad. Sci. U.S.A.* 107, 21767. doi: 10.1073/pnas.0908104107
- Bode, S., He, A., Soon, C., Trampel, R., Turner, R., and Haynes, J. (2011). Tracking the unconscious generation of free decisions using ultra-high field fMRI. *PLoS ONE* 6:e21612. doi: 10.1371/journal.pone.0021612
- Cai, X., Kim, S., and Lee, D. (2011). Heterogeneous coding of temporally discounted values in the dorsal and ventral striatum during intertemporal choice. *Neuron* 69, 170–182. doi: 10.1016/j.neuron.2010.11.041
- Cavanagh, J. E., Wiecki, T. V., Cohen, M. X., Figueroa, C. M., Samanta, J., Sherman, S. J., et al. (2011). Subthalamic nucleus stimulation reverses mediofrontal influence over decision threshold. *Nat. Neurosci.* 14, 1462–1467. doi: 10.1038/nn.2925
- Changizi, M. A., and Shimojo, S. (2008). A functional explanation for the effects of visual exposure on preference. *Perception* 37, 1510–1519. doi: 10.1068/p6012
- Cisek, P., and Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annu. Rev. Neurosci.* 33, 269–298. doi: 10.1146/annurev.neuro.051508.135409
- Coe, B., Tomihara, K., Matsuzawa, M., and Hikosaka, O. (2002). Visual and anticipatory bias in three cortical eye fields of the monkey during an adaptive decision-making task. *J. Neurosci.* 22, 5081–5090.
- Daw, N., Niv, Y., and Dayan, P. (2005). Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* 8, 1704–1711. doi: 10.1038/nn1560
- Ding, L., and Gold, J. (2010). Caudate encodes multiple computations for perceptual decisions. *J. Neurosci.* 30, 15747–15759. doi: 10.1523/JNEUROSCI.2894-10.2010
- Ding, L., and Hikosaka, O. (2006). Comparison of reward modulation in the frontal eye field and caudate of the macaque. *J. Neurosci.* 26, 6695–6703. doi: 10.1523/JNEUROSCI.0836-06.2006
- Douglas, R., Koch, C., Mahowald, M., Martin, K., and Suarez, H. (1995). Recurrent excitation in neocortical circuits. *Science* 269, 981. doi: 10.1126/science.7638624
- Efron, B., and Tibshirani, R. (1993). *An Introduction to the Bootstrap*, Vol. 57. Boca Raton, FL: Chapman and Hall/CRC.
- Fox, M., Snyder, A., Vincent, J., and Raichle, M. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron* 56, 171–184. doi: 10.1016/j.neuron.2007.08.023

- Fox, M., Snyder, A., Zacks, J., and Raichle, M. (2006). Coherent spontaneous activity accounts for trial-to-trial variability in human evoked brain responses. *Nat. Neurosci.* 9, 23–25. doi: 10.1038/nn1616
- Frederick, S., Loewenstein, G., and O'Donoghue, T. (2002). Time discounting and time preference: a critical review. *J. Econ. Lit.* 40, 351–401. doi: 10.1257/002205102320161311
- Fried, I., Mukamel, R., and Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. *Neuron* 69, 548–562. doi: 10.1016/j.neuron.2010.11.045
- Gold, J. I., and Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574. doi: 10.1146/annurev.neuro.29.051605.113038
- Haggard, P., and Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Exp. Brain Res.* 126, 128–133. doi: 10.1007/s002210050722
- Hahnloser, R., Douglas, R., Mahowald, M., and Hepp, K. (1999). Feedback interactions between neuronal pointers and maps for attentional processing. *Nat. Neurosci.* 2, 746–752. doi: 10.1038/11219
- Haynes, J. D. (2011). Decoding and predicting intentions. *Ann. N. Y. Acad. Sci.* 1224, 9–21. doi: 10.1111/j.1749-6632.2011.05994.x
- Haynes, J. D., Sakai, K., Rees, G., Gilbert, S., Frith, C., and Passingham, R. E. (2007). Reading hidden intentions in the human brain. *Curr. Biol.* 17, 323–328. doi: 10.1016/j.cub.2006.11.072
- Heekeren, H., Marrett, S., and Ungerleider, L. (2008). The neural systems that mediate human perceptual decision making. *Nat. Rev. Neurosci.* 9, 467–479. doi: 10.1038/nrn2374
- Hopfinger, J., Buonocore, M., and Mangun, G. (2000). The neural mechanisms of top-down attentional control. *Nat. Neurosci.* 3, 284–291. doi: 10.1038/72999
- Hwang, J., Kim, S., and Lee, D. (2009). Temporal discounting and intertemporal choice in rhesus monkeys. *Front. Behav. Neurosci.* 3:9. doi: 10.3389/neuro.08.009.2009
- Jahanshahi, M., and Dirnberger, G. (1999). The left dorsolateral prefrontal cortex and random generation of responses: studies with transcranial magnetic stimulation. *Neuropsychologia* 37, 181–190. doi: 10.1016/S0028-3932(98)00092-X
- Kable, J. W., and Glimcher, P. W. (2007). The neural correlates of subjective value during intertemporal choice. *Nat. Neurosci.* 10, 1625–1633. doi: 10.1038/nn2007
- Kable, J. W., and Glimcher, P. W. (2010). An “as soon as possible” effect in human intertemporal decision making: behavioral evidence and neural mechanisms. *J. Neurophysiol.* 103, 2513–2531. doi: 10.1152/jn.00177.2009
- Kawagoe, R., Takikawa, Y., and Hikosaka, O. (1998). Expectation of reward modulates cognitive signals in the basal ganglia. *Nat. Neurosci.* 1, 411–416. doi: 10.1038/1625
- Kennerly, S., Dahmubed, A., Lara, A., and Wallis, J. (2009). Neurons in the frontal lobe encode the value of multiple decision variables. *J. Cogn. Neurosci.* 21, 1162–1178. doi: 10.1162/jocn.2009.21100
- Kim, J., and Shadlen, M. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* 2, 176–185. doi: 10.1038/5739
- Kim, S., Hwang, J., and Lee, D. (2008). Prefrontal coding of temporally discounted values during intertemporal choice. *Neuron* 59, 161–172. doi: 10.1016/j.neuron.2008.05.010
- Kobayashi, S., Kawagoe, R., Takikawa, Y., Koizumi, M., Sakagami, M., and Hikosaka, O. (2007). Functional differences between macaque prefrontal cortex and caudate nucleus during eye movements with and without reward. *Exp. Brain Res.* 176, 341–355. doi: 10.1007/s00221-006-0622-4
- Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nat. Neurosci.* 13, 1292–1298. doi: 10.1038/nn.2635
- Lauwereyns, J. (2010). *The Anatomy of Bias: How Neural Circuits Weigh the Options*, (Cambridge, MA: MIT Press). doi: 10.7551/mitpress/9780262123105.001.0001
- Lauwereyns, J., Watanabe, K., Coe, B., and Hikosaka, O. (2002). A neural correlate of response bias in monkey caudate nucleus. *Nature* 418, 413–417. doi: 10.1038/nature00892
- Leopold, D. A., Murayama, Y., and Logothetis, N. K. (2003). Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging. *Cereb. Cortex* 13, 422–433. doi: 10.1093/cercor/13.4.422
- Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. *Behav. Brain Sci.* 8, 529–539. doi: 10.1017/S0140525X00044903
- Libet, B., Gleason, C., Wright, E., and Pearl, D. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential): the unconscious initiation of a freely voluntary act. *Brain* 106, 623. doi: 10.1093/brain/106.3.623
- Lo, C., and Wang, X. (2006). Corticobasal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* 9, 956–963. doi: 10.1038/nn1722
- Louie, K., and Glimcher, P. (2010). Separating value from choice: delay discounting activity in the lateral intraparietal area. *J. Neurosci.* 30, 5498. doi: 10.1523/JNEUROSCI.5742-09.2010
- Luhmann, C. C., Chun, M. M., Yi, D. J., Lee, D., and Wang, X. J. (2008). Neural dissociation of delay and uncertainty in intertemporal choice. *J. Neurosci.* 28, 14459. doi: 10.1523/JNEUROSCI.5058-08.2008
- Maoz, U., Ye, S., Ross, I., Mamelak, A., and Koch, C. (2012). “Predicting action content on-line and in real time before action onset – an intracranial human study,” in *Advances in Neural Information Processing Systems 25*, eds P. Bartlett, F. C. N. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, 881–889. Available online at: http://books.nips.cc/papers/files/nips25/NIPS2012_0403.pdf
- Mazurek, M., Roitman, J., Ditterich, J., and Shadlen, M. (2003). A role for neural integrators in perceptual decision making. *Cereb. Cortex* 13, 1257. doi: 10.1093/cercor/bhg097
- Pacherie, E., and Haggard, P. (2011). “What are intentions,” in *Conscious Will and Responsibility*, eds W. Sinnott-Armstrong and L. Nadel (Oxford University Press), 70–84.
- Padoa-Schioppa, C., and Assad, J. (2006). Neurons in the orbitofrontal cortex encode economic value. *Nature* 441, 223–226. doi: 10.1038/nature04676
- Palmer, J., Huk, A., and Shadlen, M. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *J. Vis.* 5, 376–404. doi: 10.1167/5.5.1
- Pine, A., Seymour, B., Roiser, J., Bossaerts, P., Friston, K., Curran, H., et al. (2009). Encoding of marginal utility across time in the human brain. *J. Neurosci.* 29, 9575. doi: 10.1523/JNEUROSCI.1126-09.2009
- Pockett, S. (2006). “The neuroscience of movement,” in *Does Consciousness Cause Behavior?* eds S. Pockett, W. P. Banks, and S. Gallagher (Cambridge, MA: MIT Press), 9–24. doi: 10.7551/mitpress/9780262162371.003.0002
- Rangel, A., Camerer, C., and Montague, P. (2008). A framework for studying the neurobiology of value-based decision making. *Nat. Rev. Neurosci.* 9, 545–556. doi: 10.1038/nrn2357
- Ress, D., and Heeger, D. (2003). Neuronal correlates of perception in early visual cortex. *Nat. Neurosci.* 6, 414–420. doi: 10.1038/nn1024
- Rolls, E. T., and Deco, G. (2011). Prediction of decisions from noise in the brain before the evidence is provided. *Front. Neurosci.* 5:33. doi: 10.3389/fnins.2011.00033
- Rutishauser, U., and Douglas, R. (2009). State-dependent computation using coupled recurrent networks. *Neural Comput.* 21, 478–509. doi: 10.1162/neco.2008.03-08-734
- Rutishauser, U., Douglas, R. J., and Slotine, J. J. (2011). Collective stability of networks of winner-take-all circuits. *Neural Comput.* 23, 735–773. doi: 10.1162/NECO_a_00091
- Sadaghiani, S., Hesselmann, G., Friston, K. J., and Kleinschmidt, A. (2010). The relation of ongoing brain activity, evoked neural responses, and cognition. *Front. Syst. Neurosci.* 4:20. doi: 10.3389/fnsys.2010.00020
- Schurger, A., Sitt, J. D., and Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc. Natl. Acad. Sci. U.S.A.* 109, E2904–E2913. doi: 10.1073/pnas.1210467109
- Shadlen, M., and Newsome, W. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* 86, 1916.
- Smith, P., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168. doi: 10.1016/j.tins.2004.01.006
- Soon, C., Brass, M., Heinze, H., and Haynes, J. (2008). Unconscious determinants of free decisions in the human brain. *Nat. Neurosci.* 11, 543–545. doi: 10.1038/nn.2112
- Soon, C., He, A., Bode, S., and Haynes, J. (2013). Predicting free choices for abstract intentions. *Proc. Natl. Acad. Sci. U.S.A.* 110, 6217–6222. doi: 10.1073/pnas.1212218110
- Sugrue, L., Corrado, G., and Newsome, W. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782. doi: 10.1126/science.1094765
- Vincent, J., Snyder, A., Fox, M., Shannon, B., Andrews, J., Raichle, M., et al. (2006). Coherent spontaneous activity identifies a hippocampal-parietal memory network. *J. Neurophysiol.* 96, 3517. doi: 10.1152/jn.00048.2006

- Wallis, J., and Kennerley, S. (2010). Heterogeneous reward signals in prefrontal cortex. *Curr. Opin. Neurobiol.* 20.2, 191–198. doi: 10.1016/j.conb.2010.02.009
- Wang, X. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* 36, 955–968. doi: 10.1016/S0896-6273(02)01092-9
- Williams, Z. M., Elfar, J. C., Eskandar, E. N., Toth, L. J., and Assad, J. A. (2003). Parietal activity and the perceived direction of ambiguous apparent motion. *Nat. Neurosci.* 6, 616–623. doi: 10.1038/nn1055
- Wong, K., and Wang, X. (2006). A recurrent network mechanism of time integration in perceptual decisions. *J. Neurosci.* 26, 1314. doi: 10.1523/JNEUROSCI.3733-05.2006

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Received: 16 August 2013; accepted: 05 November 2013; published online: 26 November 2013.

Citation: Maoz U, Rutishauser U, Kim S, Cai X, Lee D and Koch C (2013) Predeliberation activity in prefrontal cortex and striatum and the prediction of subsequent value judgment. Front. Neurosci. 7:225. doi: 10.3389/fnins.2013.00225
This article was submitted to Decision Neuroscience, a section of the journal Frontiers in Neuroscience.

Copyright © 2013 Maoz, Rutishauser, Kim, Cai, Lee and Koch. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.