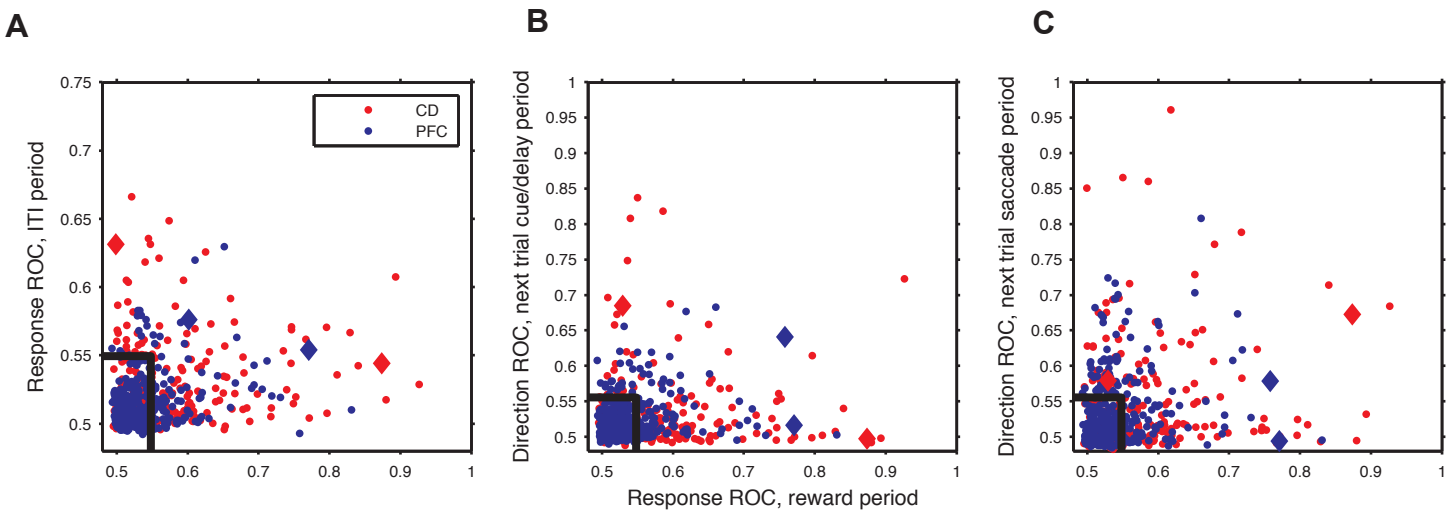

Supplementary figure 1: Comparison of direction and outcome ROC magnitudes.
A, scatterplot where each point is a cell; red points: Cd cells, blue: PFC cells. X-axis: correct/error ROC value for that cell during the reward period. Y-axis, correct/error ROC value in the inter-trial period. Black line: 3 std. dev. around the mean ROC expected by chance (calculated by shuffling trials). Highlighted cells are the single cells shown in Fig 2A-D.
B, scatterplot of correct/error ROC values (x-axis) and direction ROC on the next trial during the cue and delay periods. Highlighted cells are the single cells in Fig 2A-B and Fig 3.
C, scatterplot of correct/error ROC values (x-axis) and direction ROC on the next trial during the response period. Highlighted cells same as in B.

The strongest signals across the population are the transient reward signals (x-axis, B-C), and the direction signal at the time of saccade (y-axis, C). Furthermore, direction and reward signals appear to be largely distributed across the population. This is in contrast to the effects observed in the hippocampus by Wirth et al. (2009); they found that the same neurons that encoded outcome also tended to show sharpening of directional tuning. If that were the case, we would expect to see points in C distributed on the main diagonal.

While there is some evidence that these two variables might be represented by different neurons, our data suggests at minimum that these groups overlap to some degree.

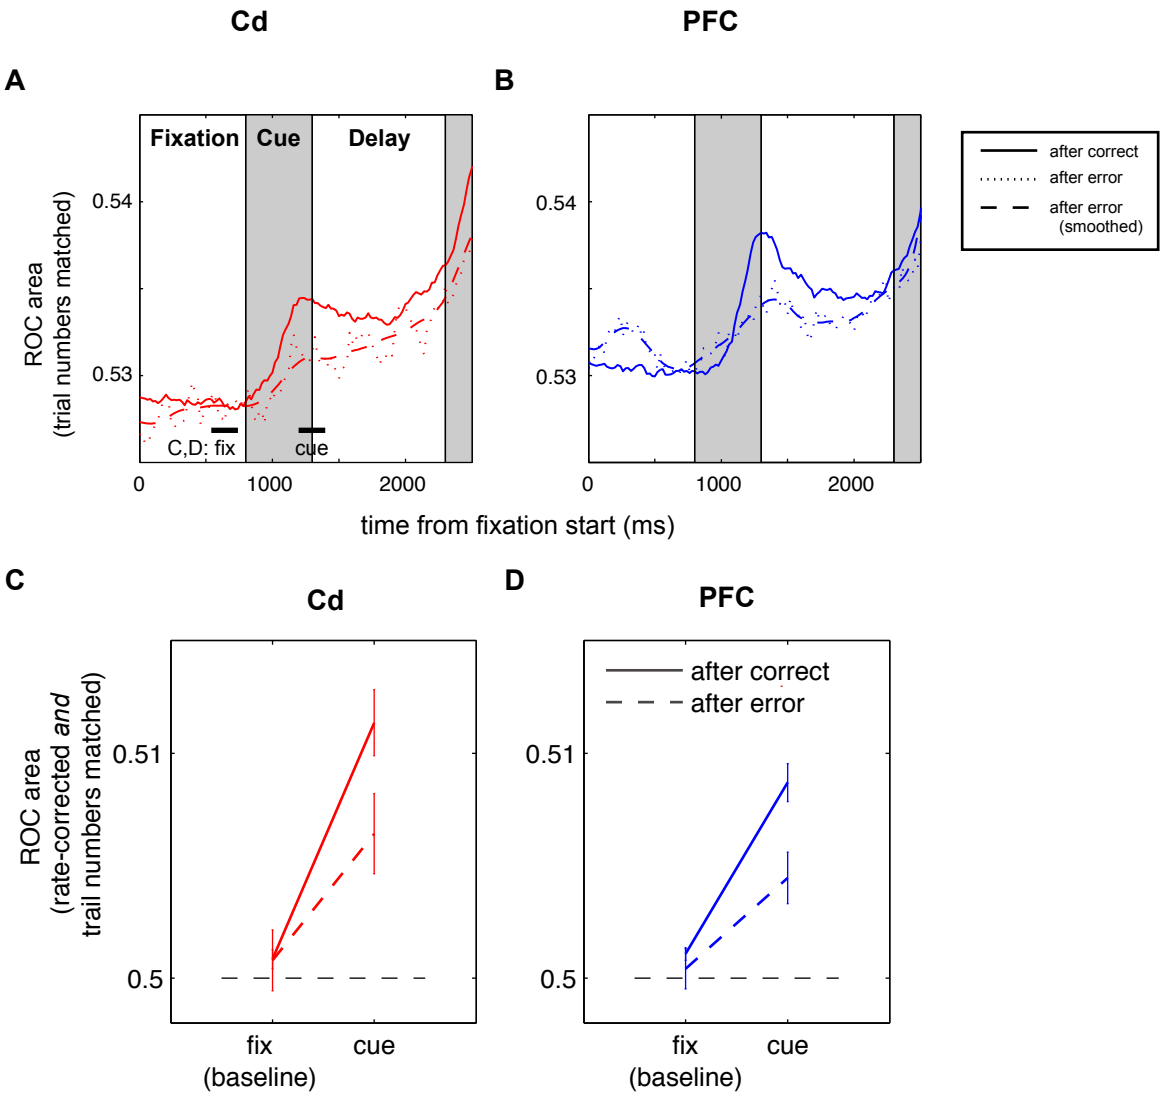

Supplementary Figure 2: ROC bias does not affect our results

Supplementary Figure 2: ROC bias does not affect our results

ROC area can have a different measured mean value than the true underlying mean (i.e. it is a biased statistic); the effect is analogous to that discussed by (Panzeri et al., 1999). This is in part because we cannot *a priori* choose a veridical preferred case for each neuron. For example for our outcome ROC (e.g. Fig. 2), we don't know whether neurons prefer correct or error (Supp. Fig. 3), and thus we rectify the ROC value around 0.5, so that ROC values range between 0.5-1.0. This causes a bias; the mean ROC value is slightly greater than 0.5 even when the null hypothesis of no tuning is true (c.f. Panzeri et al., 1999; Panzeri et al., 2007). This bias depends on both the number of trials in the dataset as well as the mean rate, likely because both cause departures from a continuous towards a discretized ROC curve which is necessarily not exactly the diagonal line which would give an area of exactly 0.5.

To control for these biases, in the main text and Figs. 2-5, we show the ROC value after shuffle correction. To compute the shuffle-corrected ROC, at each timepoint, we randomized the assignment of trials; for example, in the ROC area shown in Fig 2, we randomly assigned trials to correct and error outcome. Then, we subtracted this shuffled value from the measured ROC value.

To further confirm that different numbers of correct and error trials did not affect our results, we performed several other controls shown here. First, we computed the ROC area by matching the number of trials exactly (A-B). For each neuron, we first computed the direction ROC area for the smaller group, which was always the set of trials following error trials (mean over experimental days, N=148.2 +- std.dev. 70.2). Then, we repeatedly took a matching number of trials from the set of correct trials (N=559.9 ± 63.7; mean ± s.d), computed the ROC area, and averaged over repetitions of randomization. A, mean over Cd neurons; B, mean over PFC neurons, same data as in Fig. 4.

While this trial-matching procedure controlled for spurious effects that might result from different trial numbers, it could conceivably also be affected by differences in mean rates. Therefore, we additionally shuffle corrected the ROC values (C-D). We chose two representative timepoints because calculation was computationally expensive. *fix, cue*: 200 ms time windows indicated by black bars in A. Here, direction selectivity as shown by simultaneous trial-matched and shuffle-corrected ROC is still greater after correct than after error trials.

Finally, A-D show mean ROC values across our entire population of recorded neurons. To show that the most well-tuned neurons showed large ROC areas, we list ROC means across the top 10th percentile of neurons (data from C-D, trial-matched and shuffle-corrected):
Cd, cue, after correct: 0.602, after error: 0.523.
PFC, cue, after correct: 0.584, after error: 0.510.
Fixation values: 0.498-0.509.
S.E.M.s: all below 0.010.

Panzeri, S., Senatore, R., Montemurro, M. A., and Petersen, R. S. (2007). Correcting for the sampling bias problem in spike train information measures. Journal of Neurophysiology *98*, 1064-1072.
Panzeri, S., Treves, A., Schultz, S., and Rolls, E. T. (1999). On decoding the responses of a population of neurons from short time windows. Neural computation *11*, 1553-1577.

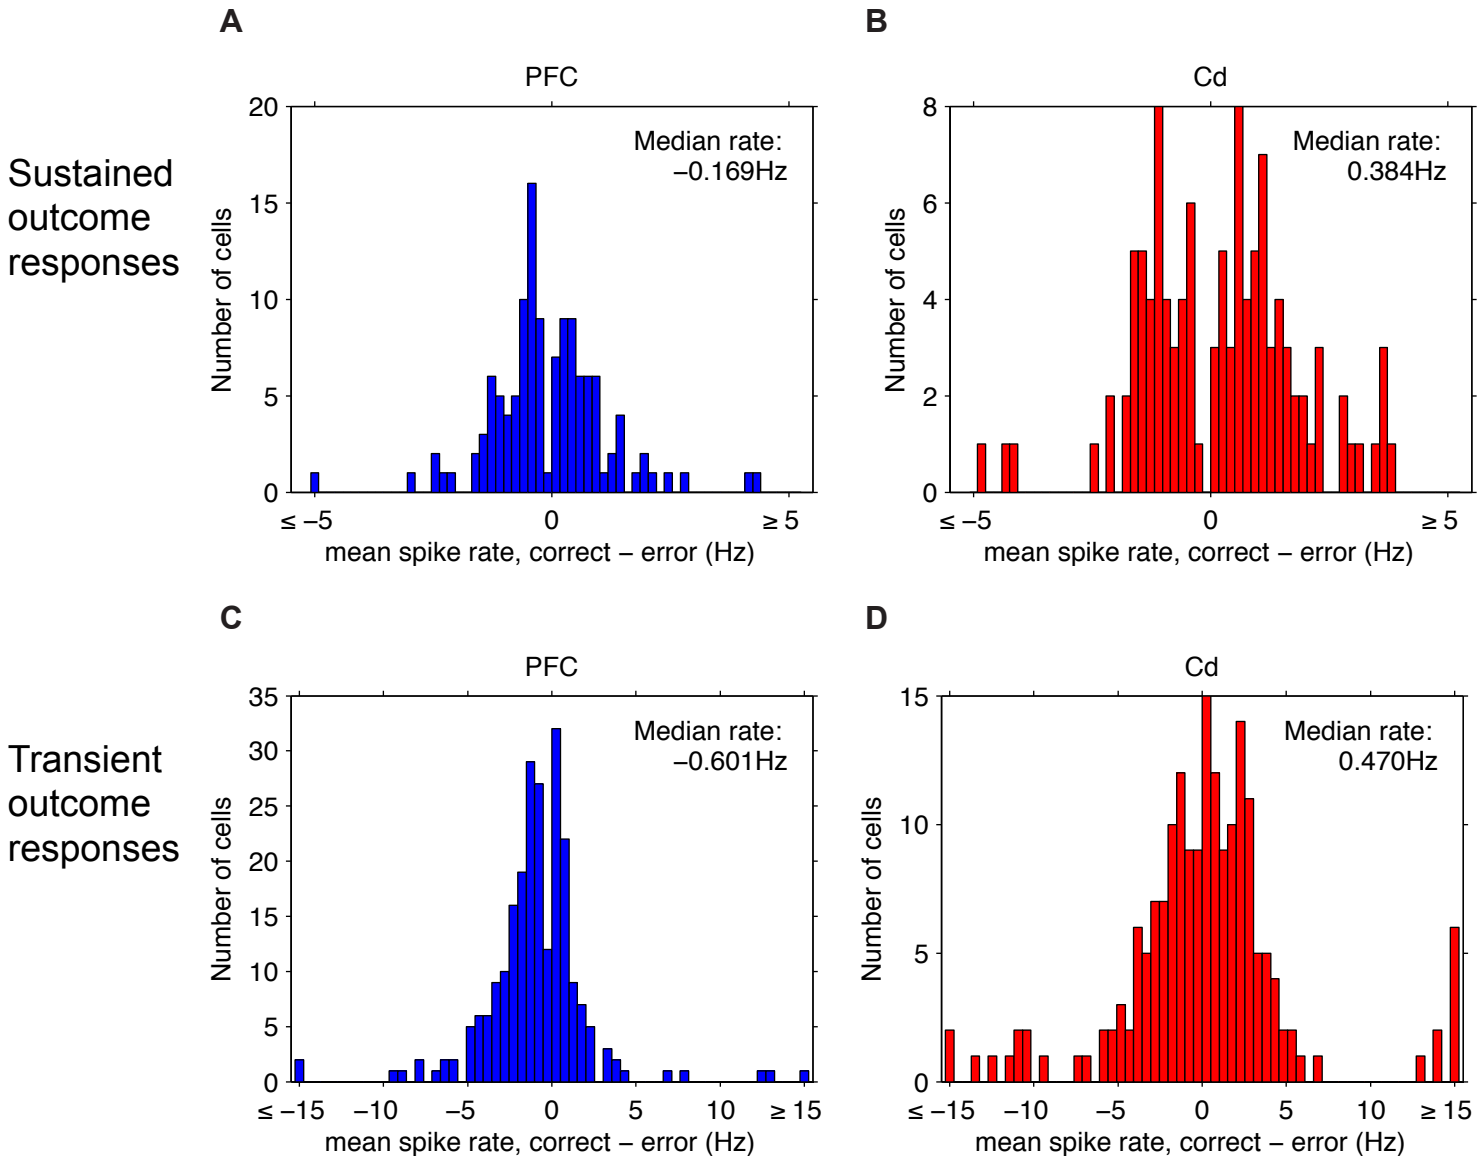

Supplementary figure 3: Similar numbers of cells increase firing rate for correct as for error

To examine whether cells are more often excited by correct responses or excited by error responses, we show here the difference in firing rate over the outcome periods. We found that nearly an equal number of cells prefer (increase their firing rate to) correct trials as to error trials.

A, Histogram of rate changes of PFC neurons during the sustained outcome period (the inter-trial interval). Time interval is a 1 s window lasting until the beginning of the next trial (defined as the onset of the fixation spot). B, Cd neurons, same time interval. C-D, PFC and Cd neurons' rate changes over the reward period, defined as a 500 ms window beginning 500 ms after outcome feedback; chosen to capture the peak of this response (see Fig. 2). In A-D, only neurons showing a statistically significant difference (Mann-Whitney/Wilcoxon test $p<0.05$) are shown, yielding a decrease in number of cells around 0 for each plot.

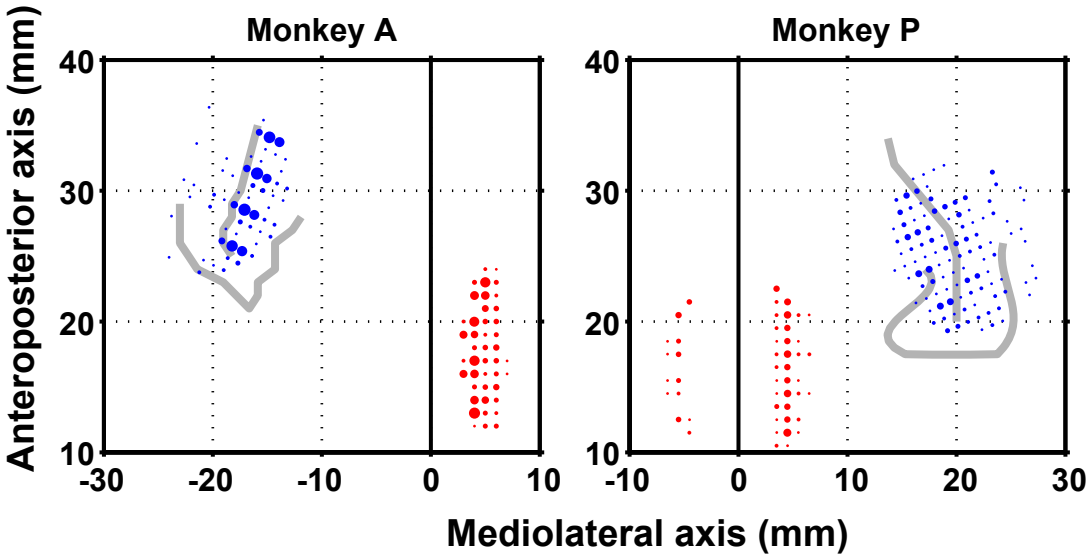

Supplementary figure 4: Locations of recording sites
PFC (red) and Cd (blue) recording sites for monkeys A (left) and P (right). Size of each dot reflects the number of electrode penetrations (range: 1 to 18 penetrations). The principal and arcuate sulci, reconstructed based on structural MRI, are indicated by gray lines. In monkey P, cells from the Cd of both hemispheres were recorded.

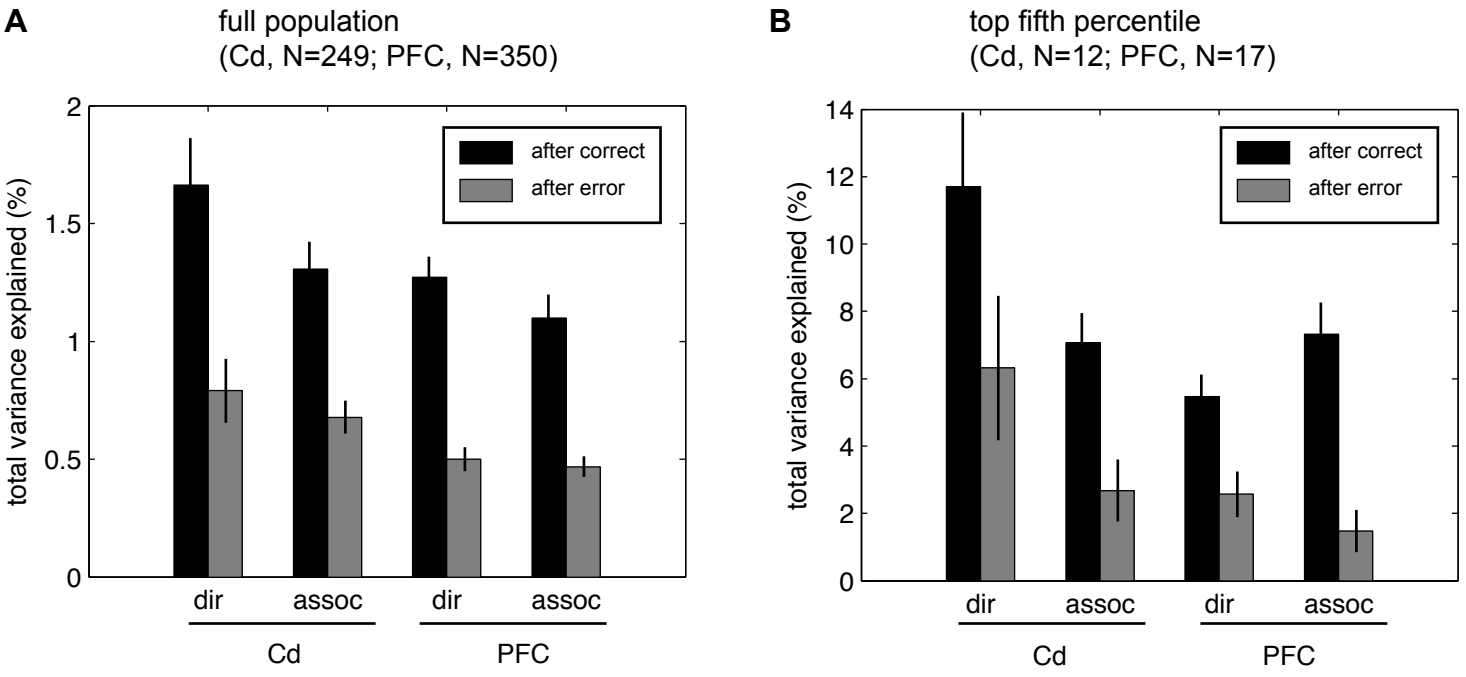

Supplementary figure 5: Stronger association as well as direction selectivity is seen after correct trials

In previous work (Pasupathy and Miller, 2005) we had computed neural selectivity for both direction and association. We partitioned the total variance into cue, direction, and association factors, using a linear model with one term for cue, one for direction, and an interaction (association) term.

In this work, we analyze learning-related changes using an ROC for direction, which will capture both direction tuning and a large proportion of association tuning (see Results), which are the main changes during learning (Pasupathy and Miller, 2005; Asaad et al. 2001).

However, to facilitate comparison to our prior work, here we quantify selectivity for both direction and association using the variance-partitioning approach. We find that both direction and association tuning are lower after an error trial than a correct trial.

A, average percent explained variance (PEV) across all neurons in Cd (left) and PFC (right), computed in the 500 ms cue period (Methods). B, average PEV over the top 5th percentile of recorded neurons, showing the larger proportion of variance explained by the most-tuned neurons. For each of the four groups of bars, to prevent spurious differences between correct and error, we found the top 5th percentile of the average PEV of correct and error. Error bars are S.E.M.