Supplementary Information

Neuronal reference frames for social decisions in primate frontal cortex

Steve W. C. Chang, Jean-Francois Gariépy, and Michael L. Platt

Supplementary Information Table: 1
Supplementary Information Figures: 9

Supplementary Table 1

Supplementary Table 1. Classification of the reward type, trial type, and reward size selectivities at the level of individual neurons from OFC, ACCs, and ACCg, based on analysis of variance.

Area	Proportion of significant neurons by factors (reward epoch)		Proportion of significant neurons between different rewards (reward epoch)		Proportion of reference frame types (reward or choice/cue epoch)	
ACCg	*Reward Outcome*	57% (*n*=81)	*Self vs. Neither*	31% (*n*=81)	*SELF frame of reference*	15% (*n*=12/81) [**38%**] (*n*=12/32) (MY: 48%; MO: 18%)
	Trial Type	36% (*n*=81)				
	Reward Volume	12% (*n*=81)	*Self vs. Other*	36% (*n*=81)		
	Reward Outcome x Trial Type	30% (*n*=81)	*Other vs. Neither*	25% (*n*=81)	*OTHER frame of refernece*	12% (*n*=10/81) [**31%**] (*n*=10/32) (MY: 19%; MO: 55%)
	Reward Volume x Reward Outcome	7% (*n*=81)	*Self (Self:Other) vs. Self (Self:Neither)*	12% (*n*=81)		
	Reward Volume x Trial Type	4% (*n*=81)			*BOTH frame of reference*	12% (*n*=10/81) [**31%**] (*n*=10/32) (MY: 33%; MO: 27%)
	Reward Outcome x Trial Type x Reward Volume	7% (*n*=81)				
ACCs	*Reward Outcome*	72% (*n*=101)	*Self vs. Neither*	57% (*n*=101)	*SELF frame of reference*	51% (*n*=51/101) [**72%**] (*n*=51/71) (MY: 82%; MO: 61%)
	Trial Type	52% (*n*=101)				
	Reward Volume	25% (*n*=76)	*Self vs. Other*	53% (*n*=101)		
	Reward Outcome x Trial Type	36% (*n*=101)	*Other vs. Neither*	20% (*n*=101)	*OTHER frame of refernece*	10% (*n*=10/101) [**14%**] (*n*=10/71) (MY: 9%; MO: 20%)
	Reward Volume x Reward Outcome	16% (*n*=76)	*Self (Self:Other) vs. Self (Self:Neither)*	5% (*n*=101)		
	Reward Volume x Trial Type	4% (*n*=76)			*BOTH frame of reference*	10% (*n*=10/101) [**14%**] (*n*=10/71) (MY: 9%; MO: 19%)
	Reward Outcome x Trial Type x Reward Volume	8% (*n*=76)				
OFC	*Reward Outcome*	57% (*n*=85)	*Self vs. Neither*	37% (*n*=85)	*SELF frame of reference*	42% (*n*=36/85) [**80%**] (*n*=36/45) (MY: 86%; MO: 70%)
	Trial Type	45% (*n*=85)				
	Reward Volume	24% (*n*=62)	*Self vs. Other*	42% (*n*=85)		
	Reward Outcome x Trial Type	37% (*n*=85)	*Other vs. Neither*	14% (*n*=85)	*OTHER frame of refernece*	5% (*n*=4/85) [**9%**] (*n*=4/45) (MY: 7%; MO: 12%)
	Reward Volume x Reward Outcome	10% (*n*=62)	*Self (Self:Other) vs. Self (Self:Neither)*	13% (*n*=85)		
	Reward Volume x Trial Type	10% (*n*=62)			*BOTH frame of reference*	6% (*n*=5/85) [**11%**] (*n*=5/45) (MY: 7%; MO: 18%)
	Reward Outcome x Trial Type x Reward Volume	11% (*n*=62)				

The bold percentages shown inside the brackets on the 4th column show the proportions out of *classfied* neurons. Shown below in parentheses are these proportions for each monkey. Significance in all panels was based on $P < 0.05$ (analysis of variance and tukey HSD tests).

Supplementary Figure & Legends

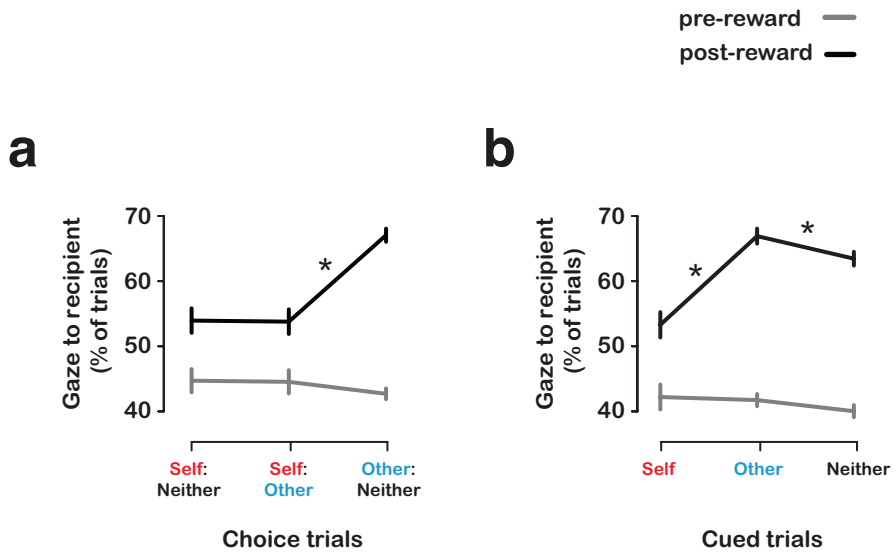

Supplementary Figure 1

Supplementary Figure 1 Percentage of gaze shifts (mean ± s.e.m.) directed toward the recipient prior to reward delivery (pre-reward epoch; grey) and following the onset of reward delivery (post-reward epoch; black) on choice trials (**a**) and cued trials (**b**). Asterisks indicate significant differences ($P < 0.05$, Welch two sample *t*-test). *Choice trials*. During the delay period between choice and reward delivery (pre-reward epoch, 0–0.9 sec; see **Fig. 1d**), gaze frequencies were comparable across trials involving actors' received and allocated rewards to other (44.75 ± 1.78% [mean ± s.e.m.], 44.57 ± 1.78%, 42.74 ± 0.83% on *Self:Neither*, *Self:Other*, *Other:Neither* trials, respectively; all comparisons $P > 0.17$, paired *t*-test). Following the onset of reward delivery (post-reward epoch, 1 sec), however, these frequencies were significantly higher on *Other:Neither* trials (67.01 ± 1.01%) compared to *Self:Other* (53.78 ± 1.88%) and *Self:Neither* trials (53.95 ± 1.88%) (both, $P < 0.0001$, paired *t*-test). *Cued trials*. During the delay period between cue offset and reward delivery (pre-reward epoch), gaze frequencies were comparable across cued *self*, cued *other* and cued *neither* trials (42.19 ± 1.90, 41.72 ± 0.92, and 40.03 ± 0.92, respectively; all comparisons $P > 0.20$). Following the onset of reward delivery (post-reward epoch), however, gaze frequencies were the highest for cued *other* trials (66.70 ± 1.15), compared to cued *self* (53.19 ± 1.93) or cued *neither* trials (63.26 ± 1.05) (both $P < 0.05$). Therefore, actors looked at the recipient at different rates depending on reward outcomes, as reported previously[33,34].

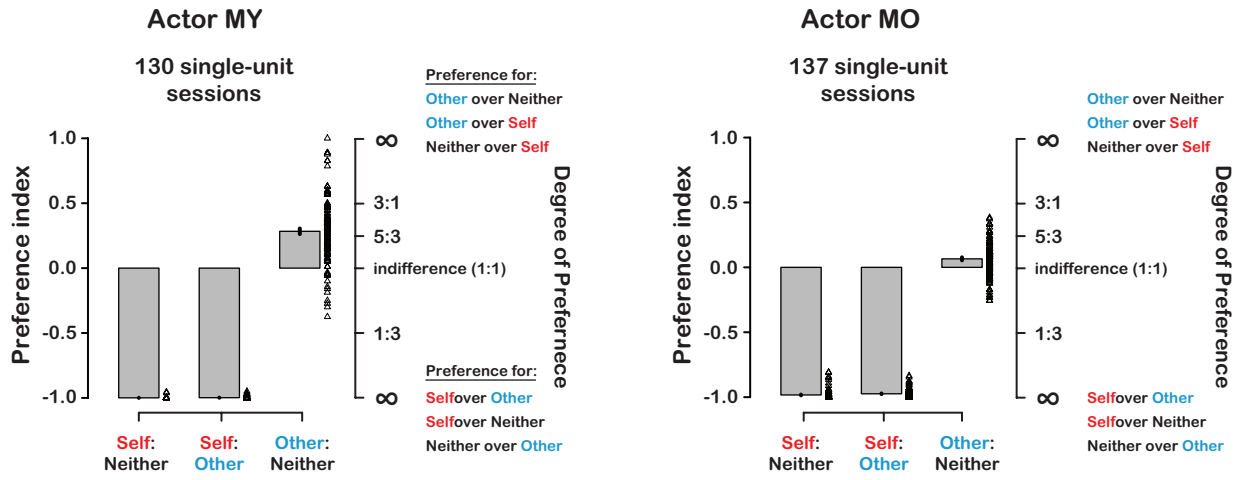

Supplementary Figure 2

Supplementary Figure 2 Choice preferences for each actor monkey (MY and MO). Shown are the preference indices as a function of reward outcome contrasts (i.e., choice contexts) for actor MY (left panel) (130 single-unit sessions) and for actor MO (right panel) (137 single-unit sessions). Data points next to each bar show individual sessions. The preference index for actor MY was –1.00 ± 0.00 (mean ± s.e.m.) for *Self:Neither*, –1.00 ± 0.00 for *Self:Other*, and 0.28 ± 0.02 for *Other:Neither* trials (significantly different from zero: all $P < 0.0001$, one-sample *t*-test). For actor MO, the preference index for *Self:Neither* was –0.98 ± 0.00, *Self:Other* was –0.97 ± 0.00, and *Other:Neither* was 0.07 ± 0.01 (significantly different from zero: all $P < 0.0001$, one-sample *t*-test). These choice behaviors are consistent with our previous studies using a similar behavioral paradigm, which also demonstrated differential reward allocation preferences depending on the familiarity and social status between the two animals[33] and the causal role of neuropeptide oxytocin in modulating these preferences[34].

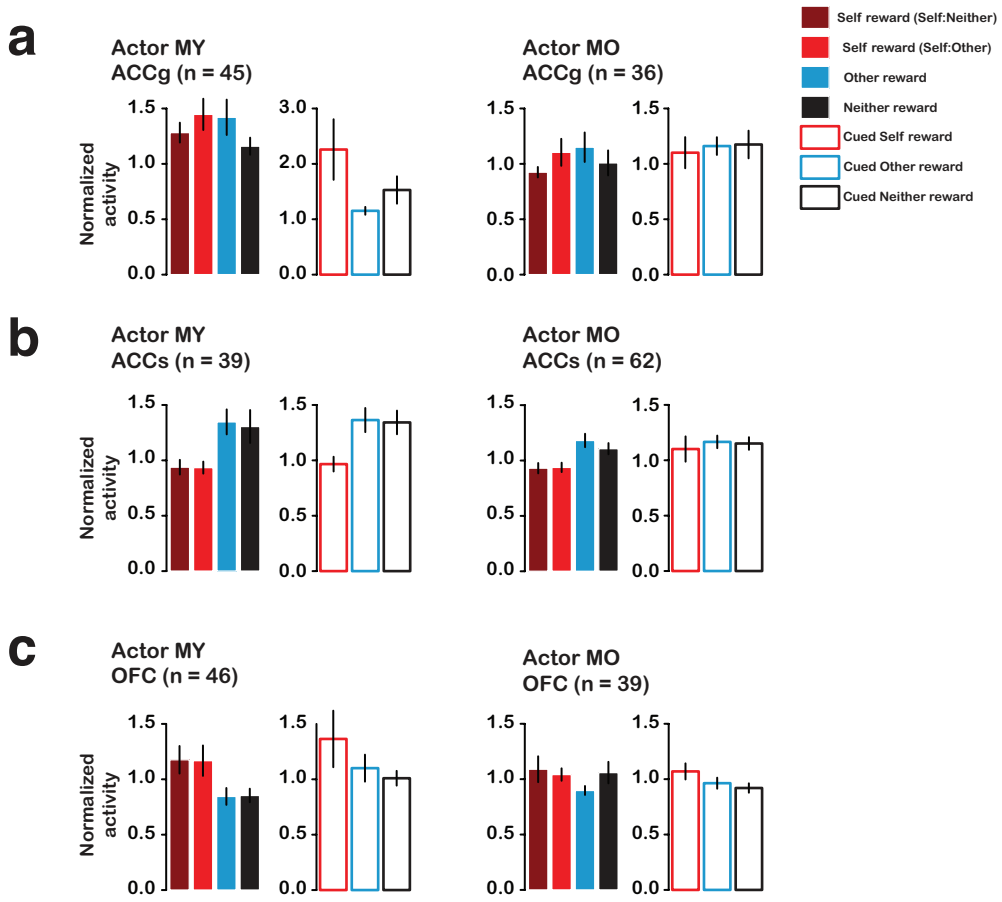

Supplementary Figure 3

Supplementary Figure 3 Preferential reward encoding biases in each actor monkey (MY and MO). Shown are the normalized responses (mean ± s.e.m.) to different reward outcomes during the reward epoch from ACCg (**a**), ACCs (**b**), and OFC (**c**). The inset shows the color coding scheme.

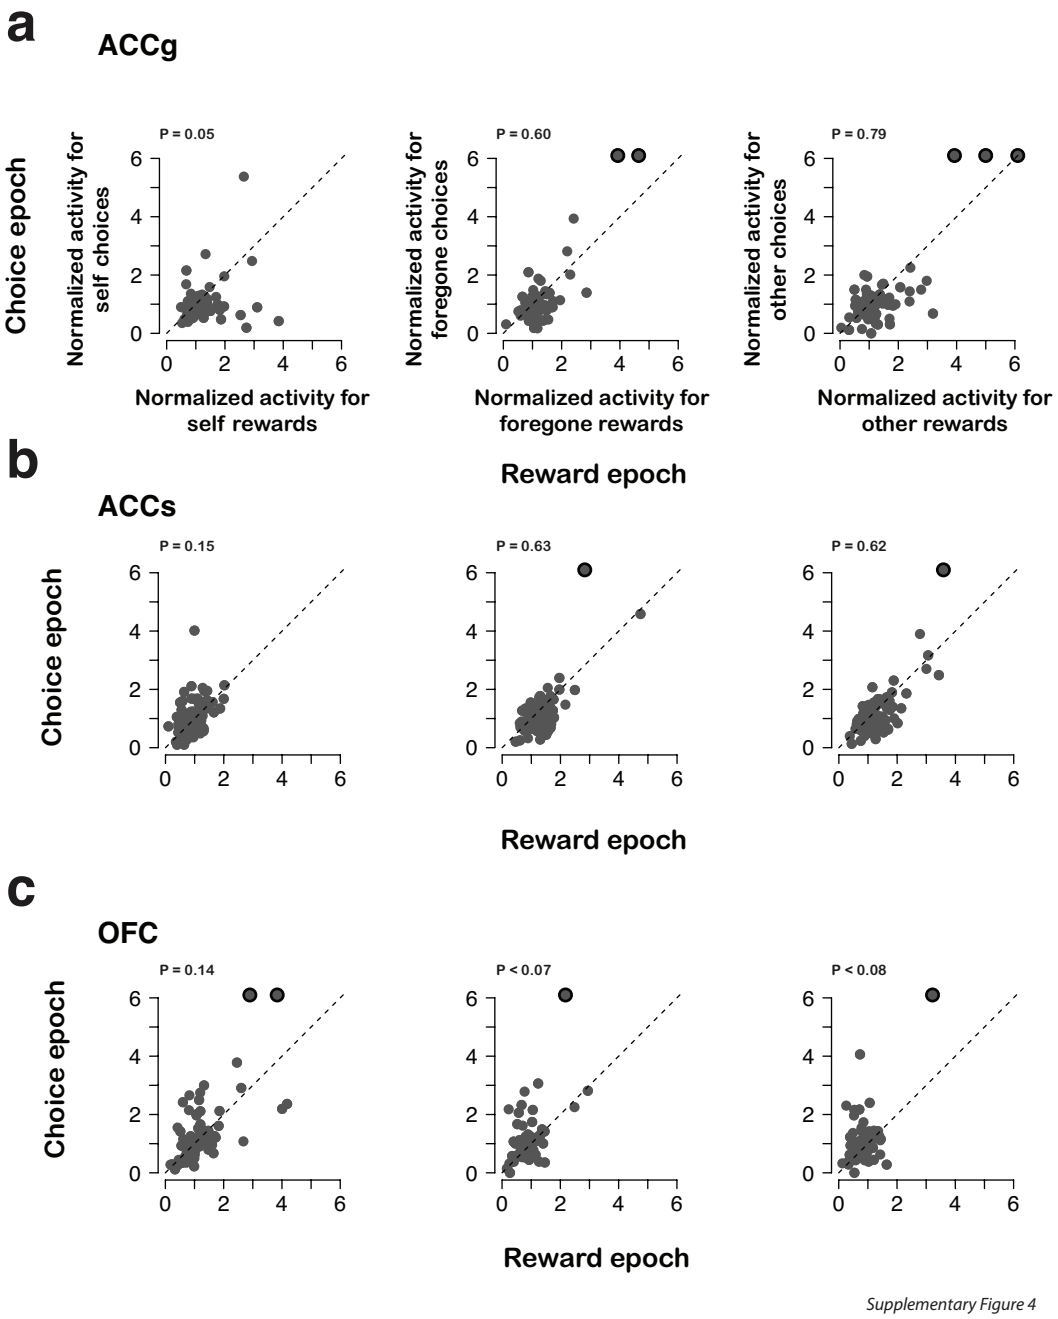

Supplementary Figure 4 Comparisons of neuronal responses across the choice/cue epoch and the reward epoch for (**a**) ACCg (n = 81), (**b**) ACCs (n = 101), and (**c**) OFC (n = 85). Plotted are the normalized responses from the two epochs for the following comparisons: *self* choices versus *self* rewards, foregone choices versus foregone rewards, and *other* choices versus *other* rewards. Data points with black outlines indicate that these values were truncated for the purpose of these displays. Significance values are shown at the top of each panel comparing the ordinate and abscissa at the population level using a paired t-test. At the population level, we observed a remarkable resemblance in neuronal activity across the choice/cue epoch and reward epoch in ACCg, ACCs, and OFC (**Fig. 3 & 4**). To quantify this similarity, we directly compared

normalized activity from the two epochs corresponding to the following pairs: *self* choices (choices leading to *self* rewards) versus *self* rewards (delivery of *self* rewards), foregone choices (leading to *other* or *neither* rewards) versus foregone rewards (delivery of *other* or *neither* rewards), and, finally, *other* choices versus *other* rewards. Although the majority of comparisons resulted in similar responses across the two epochs, there were some differences. The following summarizes the population level effects that we have observed here. ACCg as a population showed greater activity for *self* rewards during the time of reward delivery compared to the time around making choices leading to *self* rewards ($P = 0.05$, paired t-test). On the other hand, responses of ACCs neurons were similar in magnitude across all three comparisons (all $P > 0.15$, paired t-test). In contrast, OFC neurons showed a trend toward greater responses to foregone choices compared to foregone rewards ($P < 0.07$, paired t-test), as well as greater responses to *other* choices compared to *other* rewards ($P < 0.08$). At the individual cell level, however, a substantial number of neurons from each area showed significantly modulated activity across the two epochs. In ACCg, 38% ($n = 31$), 24% ($n = 19$), and 21% ($n = 17$) of neurons showed significantly modulated activity across the two epochs for *self*, *foregone*, and *other* choices versus rewards, respectively. In ACCs, these proportions were 63% ($n = 64$), 39% ($n = 39$), and 30% ($n = 30$), whereas in OFC, these proportions were 60% ($n = 51$), 40% ($n = 34$), and 28% ($n = 24$).

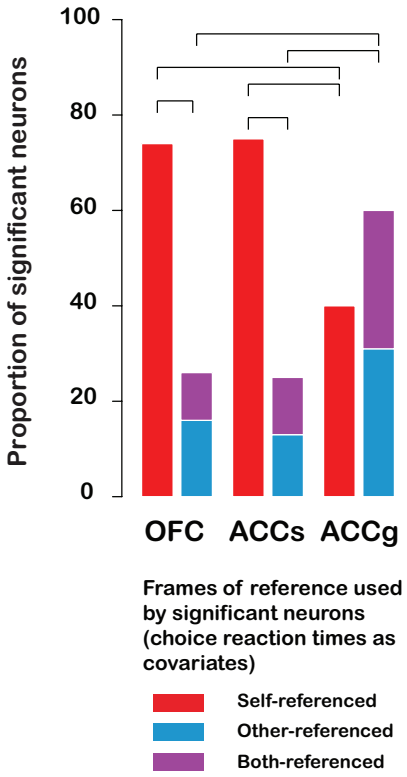

Supplementary Figure 5

Supplementary Figure 5 Proportion of neurons (out of significantly classified neurons) after correcting for eye movement choice reaction times (i.e., by using reaction times as covariates in the general linear model) from OFC, ACCs, and ACCg using self-referenced, other-referenced, and both-referenced frames for representing reward outcomes. Inset shows color codes used in the bar graph. Bars indicate significant differences in proportions ($P < 0.05$, χ^2 test). The proportions of reference frame types across the three areas remain similar even after correcting for trial-by-trial reaction times (compare to **Fig. 5d**). **Figure 2b** clearly indicates that different choices are associated with different choice reaction times. Therefore, it is possible that differential encoding schemes reported here might be simply driven by the subjective value of different choices (as inferred from reaction times). For example, if neurons were merely computing the subjective value associated with different choices, one might expect choice reaction times to explain a large amount of variance in neuronal response. We directly tested this hypothesis by including trial-by-trial reaction times as covariates in the ANOVA and re-calculated the proportion of neurons classified within different functional categories. This figure shows the distribution of different reference frame types across each area after taking into account choice reaction times. The results are virtually identical to those shown in **Fig. 5d**, suggesting that self-referenced, other-referenced, and both-referenced neurons are not the products of encoding the subjective value (as revealed by reaction times) associated with different choices. Therefore, the neurons in the current study appear to signal specific decision outcomes during social decision-making, rather than directly encoding their subjective value.

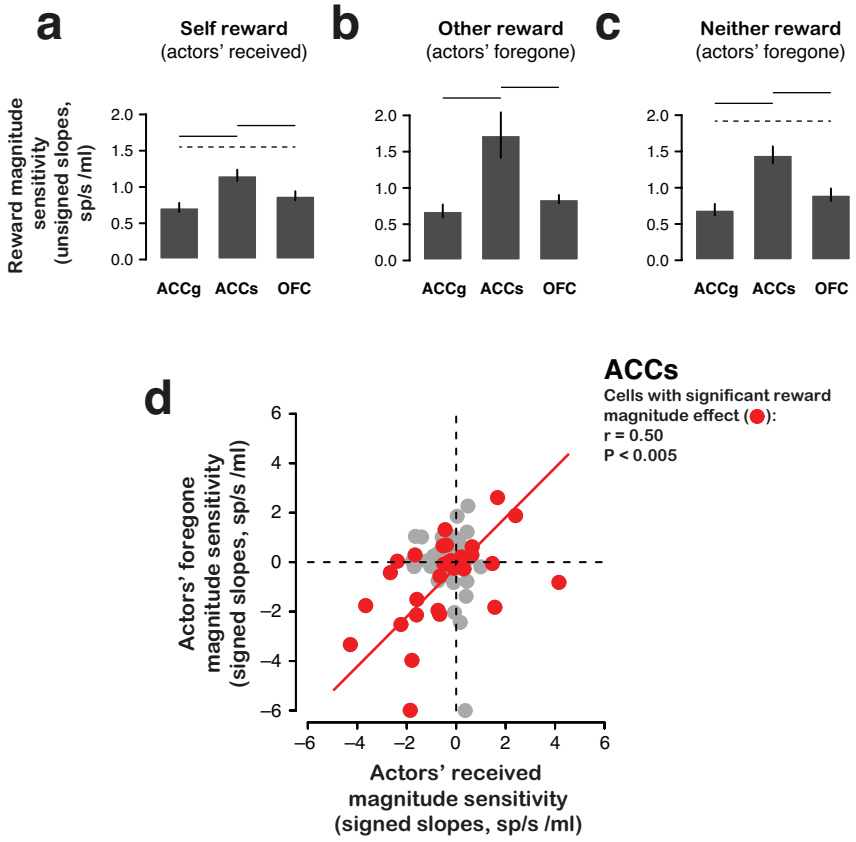

Supplementary Figure 6

Supplementary Figure 6 Effects of value on *self*, *other*, and *neither* reward responses. (**a**–**c**) Reward magnitude sensitivity for ACCg, ACCs, and OFC (unsigned slopes, sp/s/ml, mean ± s.e.m.) computed from linear regression of reward epoch activity as a function of increasing value for *self* (**a**), *other* (**b**), or *neither* (**c**) outcomes. Horizontal bars above the histograms indicate significance (solid: $P < 0.05$, dashed: $P < 0.10$, Welch two sample *t*-test). (**d**) Value sensitivity for actors' received and foregone rewards are positively correlated in ACCs. Slopes of linear regressions of reward epoch activity as a function of reward volume for actors' received rewards (*Self:Neither*, *Self:Other*, and cued *self*) and foregone rewards (*other*, *neither* rewards, *other* cued, and *neither* cued). Red points indicate significant effects of reward value or interactions (ANOVA). Line indicates type II regression for neurons with significant reward value or interaction effects. *r* and *p* reflect Pearson's correlations for significant cells. Here we examined response modulations by the magnitude of reward delivered to *self*, *other*, and *neither*. Based on the ANOVA on reward epoch responses, 40% of OFC (out of 62 cells collected in a task with a reward magnitude cue), 40% of ACCs (of 76), and 21% of ACCg (of 81) showed

either a significant effect of reward magnitude or a significant interaction involving reward magnitude (**Supplementary Table 1**). ACCg contained a significantly smaller proportion of neurons modulated by reward magnitude compared to either OFC or ACCs (both $P < 0.05$, χ^2 test), whereas OFC and ACCs did not differ ($P = 1$). Out of these regions, ACCs neurons showed the greatest sensitivity to reward magnitude based on the slopes of the regression line for each neuron across all outcomes, consistent with a prominent role for ACCs in behavioral adjustment[35–37] in an environment with constantly changing reward types and contexts. We next explored in detail how the magnitude of foregone rewards and *self* rewards was encoded in each area. We found a significant positive relationship between actors' received and forgone rewards in the sample of ACCs neurons showing significant effects of reward magnitude (significant cells: $r = 0.50$, $P < 0.005$; all cells: $r = 0.33$, $P < 0.005$, Pearson's correlation) (**d**). By contrast, we did not observe this relationship in the ACCg or OFC neurons with significant reward magnitude effects: both regions $|r| < 0.28$, $P > 0.18$; all cells: both $r < 0.21$, $P > 0.11$; Pearson's correlation) (see **Supplementary Figure 7**). Thus, the ACCs, but not the ACCg or OFC, processes actors' direct and forgone rewards in a similar manner (i.e., scale in the same direction), consistent with its hypothesized role in learning from both experience and observation[8,11].

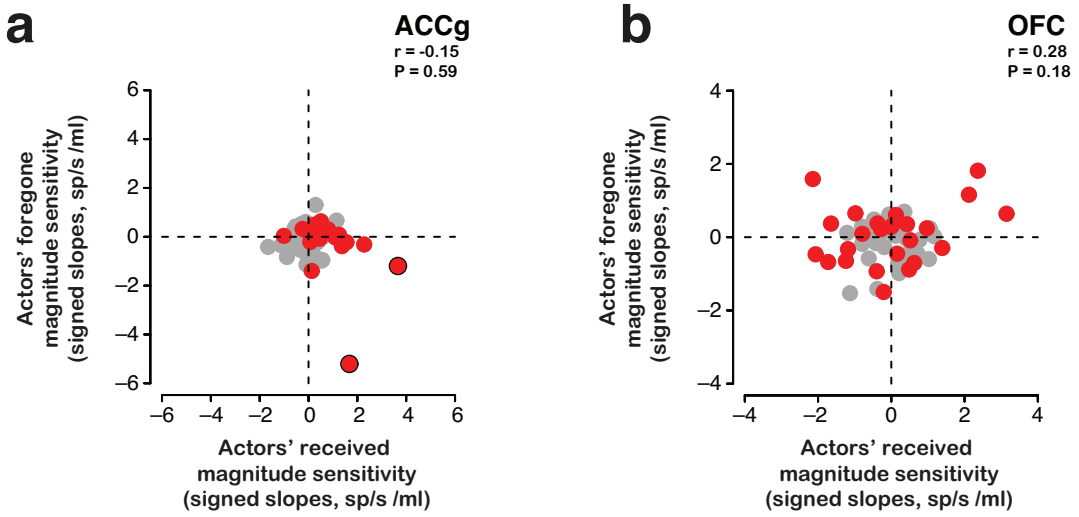

Supplementary Figure 7

Supplementary Figure 7 Reward magnitude sensitivity between actors' received and foregone rewards in ACCg (**a**) and OFC (**b**). Plotted are slopes from a linear regression of reward epoch activity as a function of different reward volumes between actors' received rewards (*Self:Neither*, *Self:Other*, and cued *self*) and foregone rewards (*other*, *neither* rewards, *other* cued, and *neither* cued). Red data points indicate significant reward magnitude main or interactions effects (ANOVA). Shown as texts in the inset are *r* and significance from Pearson's correlation for the significant neurons. Data points with black outline show the outlier cells excluded from the correlation analysis.

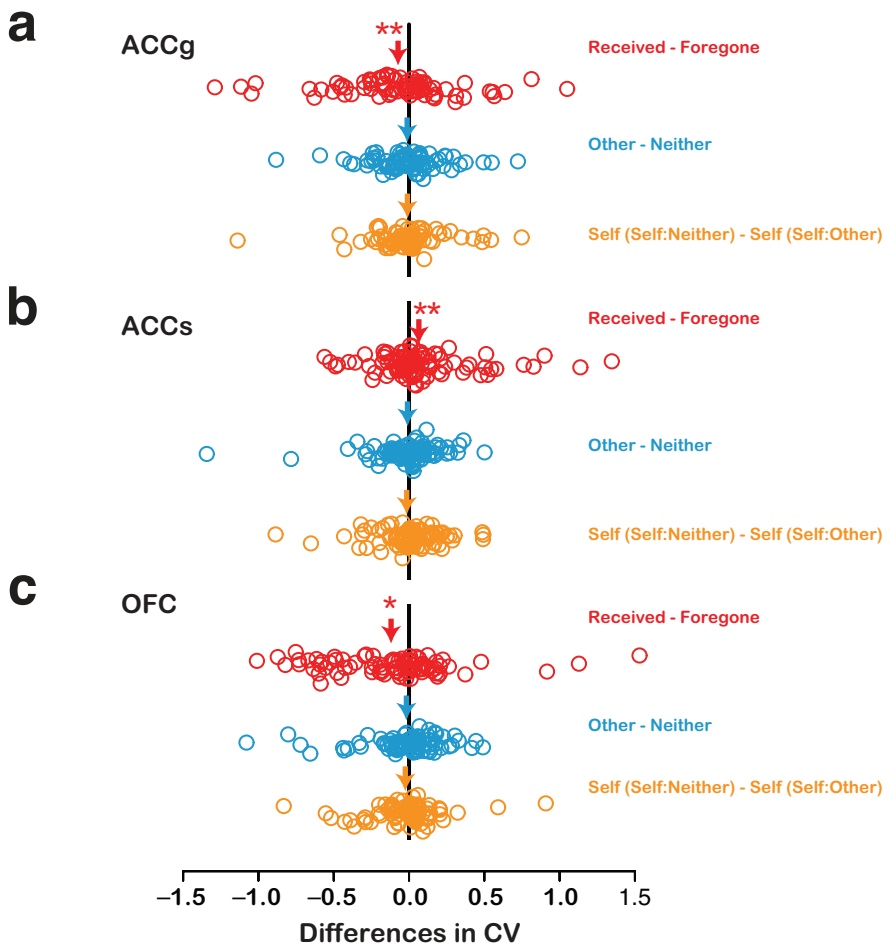

Supplementary Figure 8

Supplementary Figure 8 Differences in coefficient of variation (CV) across different reward outcomes reflect reward bias in ACCg (**a**), ACCs (**b**) and OFC (**c**). Plotted are differences in CV between a pair of reward categories (as indicated on the right of each distribution). We compared individual neuron averages of all trials in which the actors received rewards against all trials in which the actors did not receive rewards (*Received – Foregone*) (top of each panel). We also compared individual neuron averages of trials in which the recipient received the rewards against trials in which no one received rewards (*Other – Neither*) (middle of each panel), and between trials in which the actors received rewards in *Self:Neither* against *Self:Other* contexts (*Self (Self:Neither) – Self (Self:Other)*) (bottom of each panel). If applicable, the data were collapsed across choice and cued trials for this analysis. Data points are jittered in the vertical dimensions for visibility. Asterisks above the data points indicate significance (**:$P < 0.05$, *: $P < 0.10$, one sample *t*-test) in the distribution. An alternative way to examine neuronal information encoding is to assess whether lower trial-to-trial variability is associated with preferred outcomes. We tested whether the coefficient of variation in firing rates (CV; **Online Methods** Eq. 2) was systematically lower for preferred reward outcomes (based on response magnitude) compared to

non-preferred reward outcomes. We found this to be the case. The OFC population showed a lower CV for *self* rewards vs. rewards delivered to *other* or *neither* (*Received – Foregone*, –0.12 ± 0.04 [mean ± s.e.m.], $P < 0.01$, one-sample *t*-test), whereas the ACCs population showed a lower CV for rewards delivered to *other* or *neither* (*Foregone*) (0.07 ± 0.03, $P < 0.05$, one sample *t*-test). In ACCg, where some neurons preferred *self* and some preferred *other* rewards, we found a lower CV only for actors' received rewards compared to foregone rewards (–0.07 ± 0.04, $P < 0.09$, one sample *t*-test, $P < 0.05$, bootstrap test), but no difference between *other* and *neither* rewards, or between the two contexts of receiving *self* rewards (all $P > 0.34$). Thus, the most robust responses of neurons in all three areas were also the most reliable.

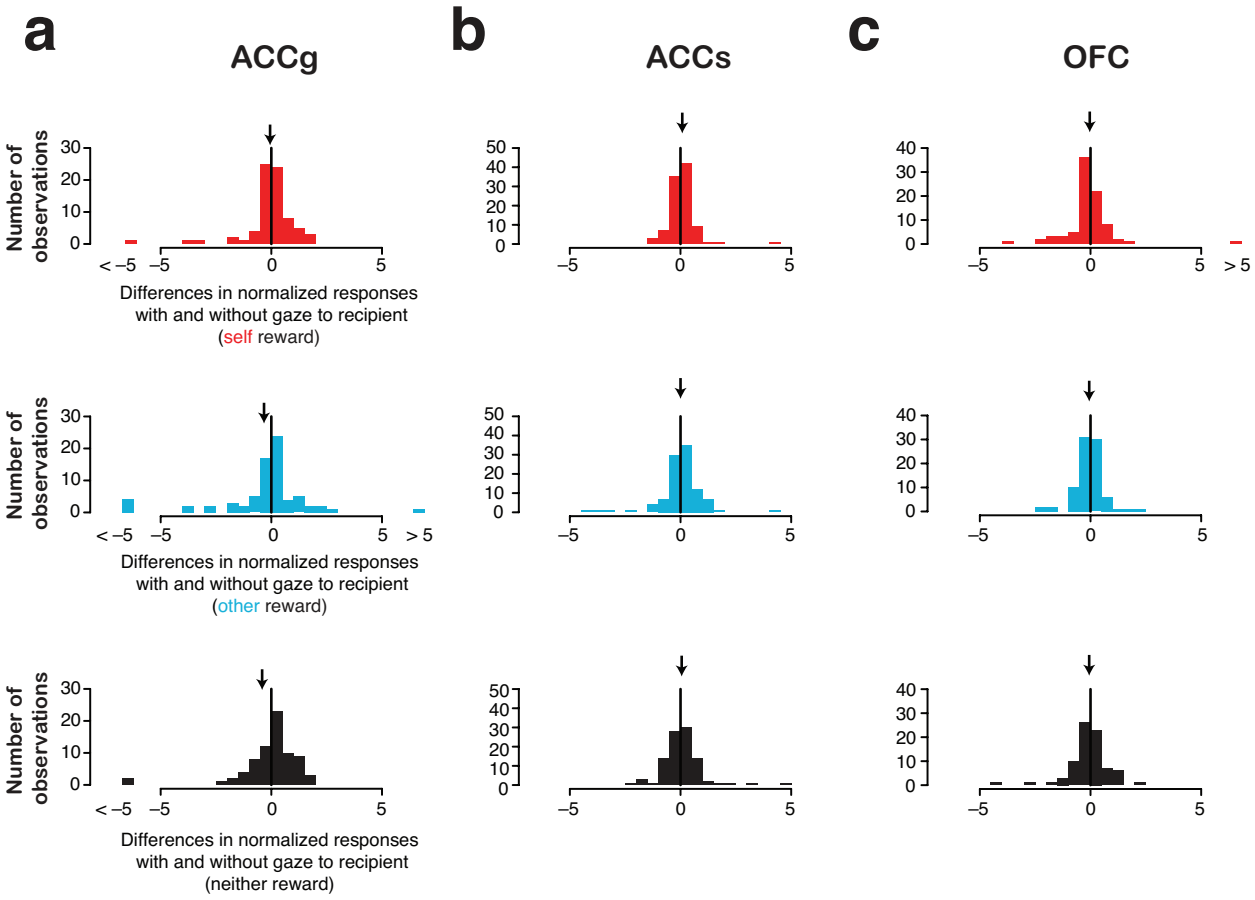

Supplementary Figure 9

Supplementary Figure 9 Reward coding is not driven by gaze shifts directed at the recipient. Shown are histograms of the differences in normalized reward epoch responses between trials with gaze shifts and without gaze shifts (responses 'with' – responses 'without' gaze shifts), for trials in which rewards were delivered to *self* (top), *other* (middle), or *neither* (bottom), for ACCg (**a**), ACCs (**b**), and OFC (**c**) populations. Arrows indicate distribution means. In the reward-allocation task, actors were allowed to look at the recipient (**Fig. 1d** & **Supplementary Figure 1**). To rule out the possibility that preferential reward responses of neurons in these areas were simply driven by where the actors looked on a given trial, we compared reward epoch responses between trials *with* and *without* gaze shifts to the recipient. We found no systematic differences in these reward responses at the population level (each areas and each reward outcome: all $P > 0.20$, Wilcoxon signed rank test). The reward-related responses in the three regions are thus neither simply driven by preparation to look at the recipient nor elicited as a consequence of inspecting the recipient.