Reward expectation differentially modulates attentional behavior and activity in visual area V4

Jalal K Baruni[1,6], Brian Lau[1,5,6] & C Daniel Salzman[1–4]

Neural activity in visual area V4 is enhanced when attention is directed into neuronal receptive fields. However, the source of this enhancement is unclear, as most physiological studies have manipulated attention by changing the absolute reward associated with a particular location as well as its value relative to other locations. We trained monkeys to discriminate the orientation of two stimuli presented simultaneously in different hemifields while we independently varied the reward magnitude associated with correct discrimination at each location. Behavioral measures of attention were controlled by the relative value of each location. By contrast, neurons in V4 were consistently modulated by absolute reward value, exhibiting increased activity, increased gamma-band power and decreased trial-to-trial variability whenever receptive field locations were associated with large rewards. These data challenge the notion that the perceptual benefits of spatial attention rely on increased signal-to-noise in V4. Instead, these benefits likely derive from downstream selection mechanisms.

Spatial attention is associated with enhanced perception at specific locations in the visual field. These perceptual benefits are typically measured as improved task performance and faster reaction times[1], which we refer to as attentional behavior. Top-down factors, such as goals, motor planning and reinforcement history, and the intrinsic salience of visual stimuli can influence spatial attention[2–4]. Neurons with receptive fields (RFs) at attended locations display modulated responses throughout striate and extrastriate visual cortex[5]. Visual area V4 has often been the focus of physiological studies investigating the neural correlates of spatial attention, and directing attention into the RF of V4 neurons enhances firing rate[6] as well as gamma-band power and synchronization[7–9] while diminishing trial-to-trial variability and inter-neuronal noise correlations[10–12].

Most physiological studies control the locus of top-down spatial attention by manipulating subjects' expectations about which locations are more likely to be associated with reward[13]. Correct task performance at one location is associated with higher reward magnitude or probability (labeled variously in previous studies as the 'target', 'attend-in' or 'relevant' location), whereas other locations are associated with lower or zero reward magnitude or probability ('distracter', 'attend-out' or 'irrelevant' locations). Thus, changes in neural activity accompanying task manipulations can be equivalently described in terms of attentional allocation or reward configuration. Notably, in terms of reward, these characteristic manipulations alter in tandem both the absolute expectation of reward at one location and the expectation of reward relative to other locations in the visual field.

Relative value, in addition to its important role in choice and motor behavior[14–17], is also likely to be relevant for spatial attention. The relative values of locations are useful for directing preferential processing to one location at the expense of others[13,18,19]. The absolute expectation of reward at one location contributes to its relative value, but it also contributes to overall reward expectation. For example, doubling the rewards associated with all locations increases both the absolute reward at each location as well as the average value of a trial. This can change animals' vigilance, motivational state and arousal level in manners not linked to spatial location. Measures of these processes include increased anticipatory licking[20], higher trial completion rates[21,22], faster reaction times[21,22] and increased pupil dilation[22].

Prior studies of value-based decision-making have manipulated reward expectations to determine whether absolute or relative reward better accounts for neural activity. For example, in the lateral intraparietal cortex (LIP), neural activity predominantly tracks relative value when making value-based decisions between two locations[16,17]. This relative value encoding may underlie perceptual enhancements during spatial attention. However, in these studies of relative value in LIP, attentional behavior was not measured, and the focus was not on how neural modulations relate to perceptual performance. In visual cortex, one study found that V1 neurons are modulated by relative value when stimuli appear at locations that are close together in space[19]. However, it is not known whether relative value across hemifields modulates visual cortex or whether the relative or absolute values of RF stimuli better account for neural modulation in visual area V4.

The characterization of how absolute and relative reward values modulate neural activity in V4 may help to elucidate mechanisms underlying the perceptual benefits of attention. Three mechanisms have been proposed to link neural modulations of visual representations with attentional behavior[23]: increased signal, decreased noise and efficient selection of sensory responses. Signal-to-noise mechanisms yield perceptual benefits through higher fidelity sensory representations, whereas efficient selection yields perceptual benefits

[1]Department of Neuroscience, Columbia University, New York, New York, USA. [2]Kavli Institute for Brain Sciences, Columbia University, New York, New York, USA. [3]Department of Psychiatry, Columbia University, New York, New York, USA. [4]New York State Psychiatric Institute, New York, New York, USA. [5]Present address: Brain and Spine Institute, Sorbonne Universités, UPMC University of Paris 06, UMR S 1127, CNRS UMR 7225, Paris, France. [6]These authors contributed equally to this work. Correspondence should be addressed to B.L. (brian.lau@upmc.fr) or C.D.S. (cds2005@columbia.edu).

Received 28 April; accepted 17 September; published online 19 October 2015; doi:10.1038/nn.4141

through selective pooling of sensory signals[23–25]. Each of these mechanisms potentially contributes to the benefits of attention, but their relative contributions are unclear[10,24,26–28].

We reasoned that independently changing absolute and relative reward expectations could provide insight into the mechanisms of the perceptual benefits of attention. We trained monkeys to perform a perceptual task at two locations simultaneously while we independently varied the relative and absolute rewards associated with correct performance at the two locations. We found that increasing the absolute reward value associated with the RF location enhanced firing rate, increased gamma-band power in local field potentials (LFPs) and decreased trial-to-trial variability in V4 neural responses. Changes in behavioral measures of spatial attention, on the other hand, tracked relative value manipulations. Indeed, attentional behavior could vary without large changes in V4 activity, and V4 activity could vary without large changes in attentional behavior. Modulation of V4 activity representing a particular spatial location therefore does not necessarily predict improved perceptual performance at that location. This finding places a strong and unexpected constraint on models of how V4 activity relates to perception and behavior.

RESULTS

Reward expectation controls attentional behavior and motivational state

We trained monkeys to perform a dual two-alternative forced choice (dual 2-AFC) orientation discrimination task in which we independently varied rewards associated with correct performance at two different locations (**Fig. 1a**). On each trial, monkeys acquired central fixation, after which two colored discs briefly appeared to cue the potential reward associated with correct performance at each location. Then, at each location, a rapid stream of oriented Gabor patches cycled randomly and independently through different orientations. At a random time, the final oriented stimuli were presented for a longer duration before being abruptly extinguished and masked. Next, we queried monkeys about the final orientation of one of the Gabor patches by presenting two choice targets around one of the two locations. Monkeys indicated with an eye movement whether the final orientation at the queried location was more horizontal or more vertical. Each location was queried on half of trials selected randomly, allowing us to measure performance at each location. By employing a dual discrimination task, as opposed to a detection task, we ensured that performance differences across conditions did not depend on adaptive shifts in response bias[29].

We characterized how the distribution of reward contingencies across the visual field modulated neural activity and behavior by independently varying the reward associated with correct discrimination at each location. Correct performance at each location was associated with either a large or small reward, yielding four trial conditions that were pseudorandomly interleaved (**Fig. 1b**). We refer to the final oriented stimulus at the queried location as the 'target' and the opposite stimulus as the 'distracter'. In the two unbalanced conditions, the target and distracter differ in potential reward (LS: large reward target, small reward distracter; SL: small reward target, large reward distracter) and the locus of spatial attention presumably shifts to the large reward location. These reward manipulations are similar to those yielding 'attend-in' and 'attend-out' conditions in other attention studies[13]. In the two balanced conditions, both target and distracter have the same potential reward (LL: large reward target, large reward distracter; SS: small reward target, small reward distracter) and there is no advantage to shifting spatial attention. Thus in the balanced conditions (LL and SS), the absolute target value changes similarly to the unbalanced conditions (LS and SL), but the relative target value remains fixed and equal. By contrast, in the balanced conditions, the average trial value (mean of the potential rewards at both locations) is different, and monkeys' motivation or arousal may be modulated accordingly. For example, in the LL condition, because correct performance at either location is associated with a large reward, the overall expected value of the trial is high and monkeys may therefore be more motivated or aroused.

The allocation of spatial attention, measured using percent correct and reaction time, was governed largely by relative target value. The percent correct was highest when relative target value was highest (77.1% in LS condition, Cochran-Mantel-Haenszel test (CMH), $P < 10^{-10}$ for all comparisons, uncorrected; **Fig. 1c**), intermediate in balanced conditions (72.7% and 70.7% in LL and SS conditions) and

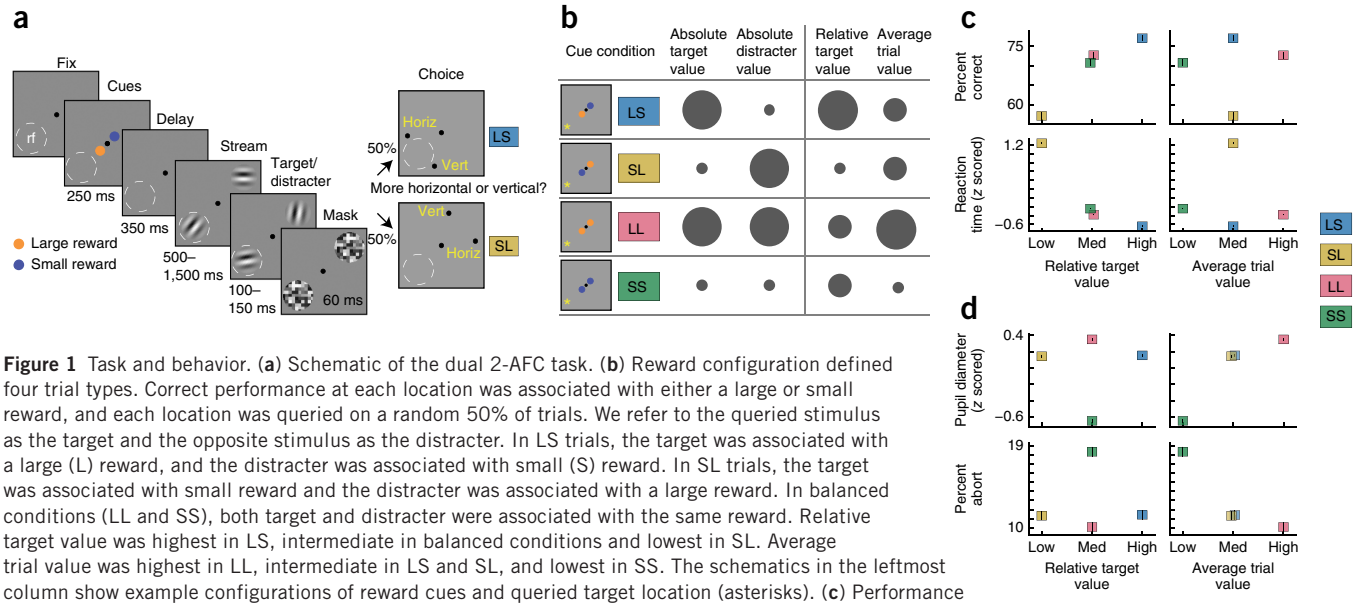

Figure 1 Task and behavior. (**a**) Schematic of the dual 2-AFC task. (**b**) Reward configuration defined four trial types. Correct performance at each location was associated with either a large or small reward, and each location was queried on a random 50% of trials. We refer to the queried stimulus as the target and the opposite stimulus as the distracter. In LS trials, the target was associated with a large (L) reward, and the distracter was associated with small (S) reward. In SL trials, the target was associated with small reward and the distracter was associated with a large reward. In balanced conditions (LL and SS), both target and distracter were associated with the same reward. Relative target value was highest in LS, intermediate in balanced conditions and lowest in SL. Average trial value was highest in LL, intermediate in LS and SL, and lowest in SS. The schematics in the leftmost column show example configurations of reward cues and queried target location (asterisks). (**c**) Performance and reaction time tracked relative target value. Behavioral data are presented as function of relative target value (left), as well as average trial value (right). Error bars indicate ±s.e.m. Overlapping points are slightly offset for visualization. (**d**) Pupil diameter and abort rate track average trial value. Differences in pupil diameter were sustained for the duration of the Gabor stream (**Supplementary Fig. 6**). Data are presented as in **c**.

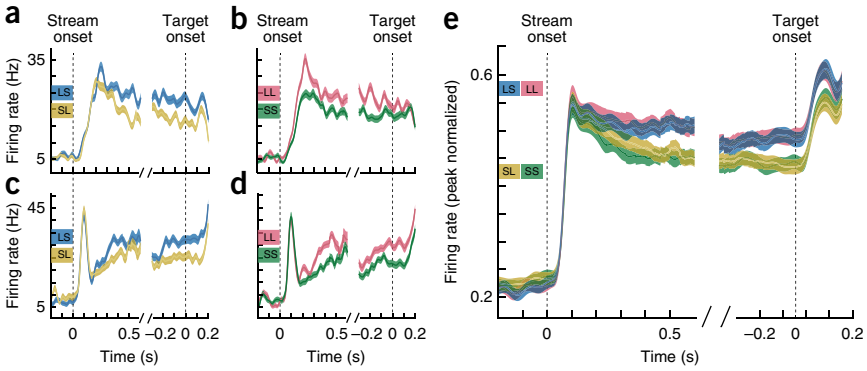

Figure 2 Firing rate modulation in V4 reflects absolute value of RF stimuli. (**a**–**d**) Spike density functions (σ = 15 ms; shading, ±s.e.m.) for two single units in V4. Responses are aligned to two events in a trial. On the left, responses are aligned to the onset of the streaming Gabor stimulus. On the right, responses are aligned to the target onset. Distracter appearance also occurs at this time. (**a**,**c**) Responses in LS (large reward in RF, small reward opposite RF) and SL (small reward in RF, large reward opposite RF) trials. (**b**,**d**) Responses in LL (large rewards both in and opposite RF) and SS (small rewards both in and opposite RF) trials. (**e**) Peak-normalized population average firing rate (n = 190 units; 92 single, 98 multi; 106 from M1, 84 from M2).

lowest when relative target value was lowest (57.2% in SL condition, $P < 10^{-10}$ for all comparisons). Reaction times mirrored the percent correct and were fastest in the LS condition, intermediate in the balanced conditions and slowest in the SL condition (Wilcoxon rank-sum test (WRS), $P < 10^{-10}$ for all comparisons, uncorrected; **Fig. 1c**). There was a small, but significant, benefit to performance (CMH, $P < 10^{-3}$) and reaction time (WRS, $P < 10^{-10}$) on LL compared with SS trials, suggesting that, although relative value accounts for the bulk of the performance differences, average trial value also has a minor role. These behavioral effects were observed in both monkeys, at both locations, and across task difficulty (**Supplementary Figs. 1–3**). Overall, we observed that the major behavioral benefits of attention occurred when reward expectation differed across the visual field, with preferential selection of stimuli with high relative value.

We found that arousal and motivational state, measured using pupil diameter[22] and trial abort rates[21,22], respectively, were correlated with the average trial value. Pupils were most dilated in the LL condition, intermediate in unbalanced conditions and smallest in the SS condition (WRS, $P < 10^{-10}$ for all comparisons that differed in average trial value; **Fig. 1d**). Average trial value also influenced trial abort rates (**Fig. 1d**), with monkeys more frequently aborting trials of the lowest expected value (SS; CMH, $P < 10^{-10}$ for all comparisons). One monkey also aborted LL trials significantly less frequently than LS and SL trials (CMH, $P < 10^{-6}$ for both comparisons; **Supplementary Fig. 2**), whereas the other showed no significant difference in abort rate between LL and the unbalanced conditions (CMH, $P > 0.1$ for both comparisons; **Supplementary Fig. 2**). All other behavioral measures were similar in both monkeys (**Supplementary Fig. 2**). The pattern of results for pupil diameter and abort rate differed markedly from that for percent correct and reaction time, indicating that the dual 2-AFC task effectively isolated changes in arousal and motivational state from changes in attentional allocation.

Visual cortical activity reflects the reward associated with receptive field stimuli

We predicted that V4 neural activity would track the relative value of RF stimuli if this activity conferred the performance and reaction time benefits associated with attention (for example, see refs. 19,29,30). In this case, firing rates would be highest in the LS condition (large reward in RF, small reward opposite RF), intermediate in balanced conditions and lowest in the SL condition (small reward in RF, large reward opposite RF). **Figure 2** shows the responses of two neurons. Both neurons increased firing rate when the relative RF value was high (LS compared with SL trials; WRS, $P < 0.01$ for both comparisons; **Fig. 2a**,**c**), reproducing prior findings of higher activity for attend-in compared with attend-out trials. To our surprise, these neurons also showed similar firing rate enhancement in LL trials compared with SS trials (WRS, $P < 0.01$ for both comparisons; **Fig. 2b**,**d**), despite minimal differences in behavioral performance measures between these conditions. Finally, firing rates were not different for both neurons

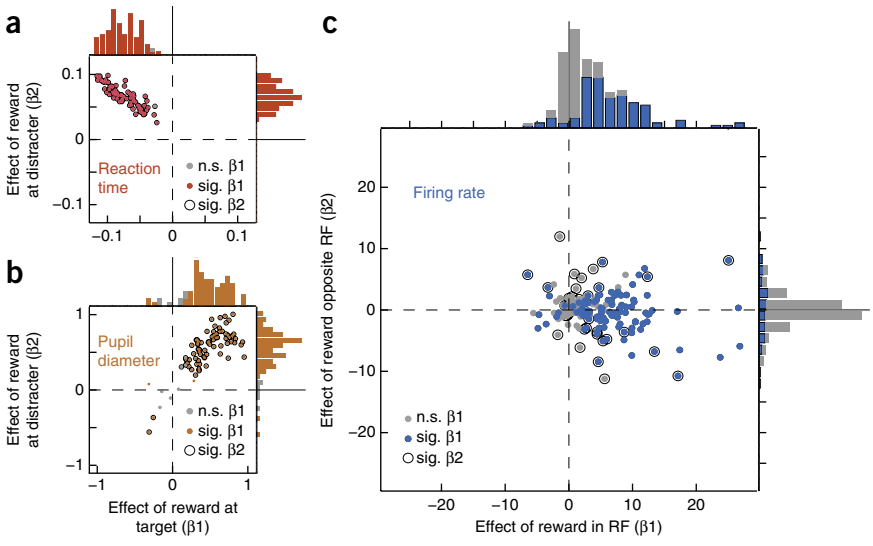

Figure 3 Regression summary of independently changing reward at two spatial locations. Joint and marginal distributions of regression coefficients across all behavioral sessions (n = 79) and recorded units (n = 190). (**a**) Reaction times reflected relative target value in all sessions. Associating the larger reward with the target location decreased reaction times (WSR, $P < 10^{-10}$), whereas large rewards at distracter locations increased reaction times (WSR, $P < 10^{-10}$). Symbol style denotes significance for individual data points ($P < 0.05$). Colored shading in marginal histograms indicates significance for individual data points ($P < 0.05$). (**b**) Pupil diameter reflected average trial value. Associating large rewards with either the target (WSR, $P < 10^{-10}$) or distracter (WSR, $P < 10^{-10}$) increased pupil diameters. (**c**) Neuronal modulation reflected RF value. Large rewards at the RF location were associated with increased firing rates (WSR, $P < 10^{-10}$). On average, large rewards opposite the RF were not associated with a significant change in firing rate (WSR, $P = 0.964$).

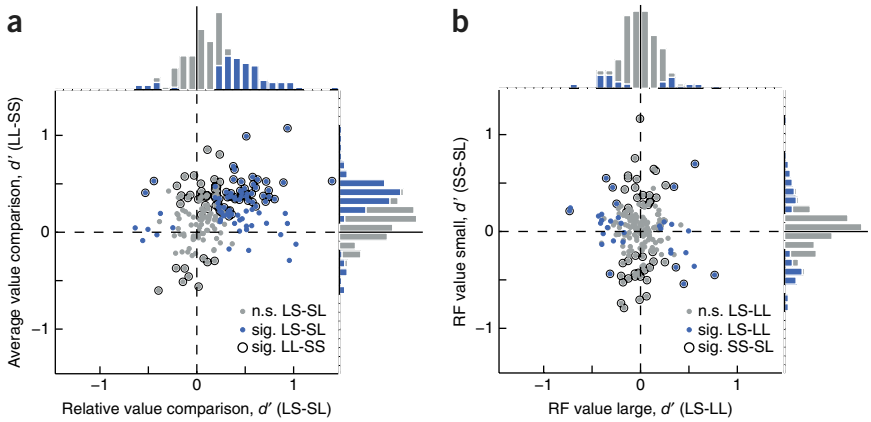

Figure 4 Correlations between neural effects. Distributions of effect sizes (d') for all recorded units ($n = 190$). (**a**) Relative RF value (LS-SL) and average trial value (LL-SS) similarly modulated neuronal activity. Both manipulations increased firing rate (WSR, $P < 10^{-10}$), with effect sizes positively correlated across units ($r = 0.31$, $P < 10^{-4}$). Symbol style indicates significance of selectivity for each unit (randomization test, $P < 0.05$). (**b**) Absolute RF value was fixed in two condition comparisons (LS-LL and SS-SL). For both comparisons, effect size distributions did not significantly differ from zero (WSR, $P > 0.3$) and effect sizes were not positively correlated across units ($r = -0.092$, $P = 0.207$).

(WRS, $P > 0.1$) when comparing conditions in which the value of the stimulus in the RF was the same (SS versus SL and LS versus LL), despite the large differences in performance and reaction time observed between these conditions. This pattern of results was true across the population, with neurons showing higher firing rates in LS and LL trials than in SL and SS trials (**Fig. 2e**).

We used multiple linear regression to quantify how the reward contingency at each location affected neuronal activity and behavior across the population (Online Methods). For behavioral data, the regression coefficients separately characterize the influence of associating large rewards with target (β_1) and distracter (β_2) locations. For neural data, regression coefficients separately characterize the influence of associating large rewards with RF (β_1) and opposite-RF (β_2) locations. Consistent with average reaction times tracking relative value (**Fig. 1c**), we found that associating a large reward with the target uniformly decreased reaction times, whereas associating large rewards with distracter locations uniformly increased reaction times (Wilcoxon signed-rank test (WSR), $P < 10^{-10}$; **Fig. 3a**). Similarly, consistent with average pupil diameters tracking average trial value (**Fig. 1d**), large rewards at either target or distracter locations increased pupil diameters (WSR, $P < 10^{-10}$; **Fig. 3b**).

In contrast with behavior, neurons were modulated primarily by the value of the stimulus in the RF. Associating the larger reward with RF stimuli consistently and robustly enhanced neural activity, with a mean firing rate increase of 4.05 ± 0.40 spikes per s across all units (WSR, $P < 10^{-10}$; **Fig. 3c**). 97 units displayed a significantly positive (nine significantly negative, $P < 0.05$) effect of large reward in the RF. By contrast, associating larger rewards with locations opposite the RF did not significantly affect firing rate (mean = -0.046 ± 0.21 spikes per s, WSR, $P = 0.964$; **Fig. 3c**). These effects were similar in each monkey (**Supplementary Fig. 4**). Thus, V4 neural modulation primarily reflects the absolute rather than the relative value of rewards associated with correctly discriminating stimuli in the receptive field.

The regression coefficients compactly describe reward modulation in our task by summarizing firing rate changes across the set of six possible condition comparisons. To further explore how reward modulation varied across the population of recorded units, we estimated effect sizes (d') for four key condition comparisons (**Fig. 4**). We first compared LS and SL trials, which is the comparison analogous to the attend-in versus attend-out difference explored in most attention studies. Across these conditions, the relative value of RF stimuli differed, but the average trial value did not. Firing rates were increased in LS compared with SL trials (mean $d' = 0.196$, WSR, $P < 10^{-10}$; **Fig. 4a**), corresponding to a mean increase of 10.2% and mean neural modulation index (MI, Online Methods) of 0.038, effect sizes that are similar to those reported previously for V4 (refs. 10,11,29).

When the average trial value differed, but the relative target value did not (LL and SS), we observed minimal effects on performance and reaction time (**Fig. 1c**). Nevertheless, increasing average trial value enhanced firing rates (mean $d' = 0.192$, WSR, $P < 10^{-10}$; **Fig. 4a**) to the

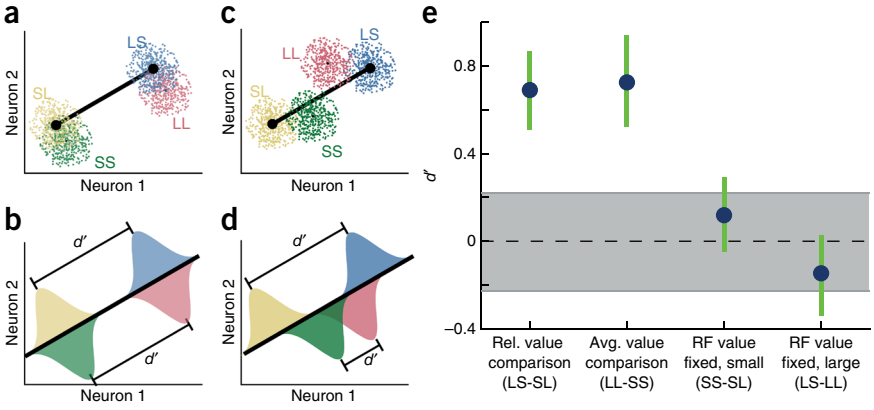

Figure 5 Differences in population activity when changing average trial value are similar to those when changing relative value. (**a**–**d**) Characterizing population activity along the attention axis[31]. (**a**) Population activity is represented by a point in the space defined by the activity of each neuron on each trial (two neurons for illustration). The attention axis (bold black line) is the vector connecting the mean responses in LS and SL trials, which represents the axis along which changes in population activity are accompanied by changes in behavioral measures of spatial attention (**Fig. 1**). (**b**) Projecting the population activity for each trial onto the attention axis to measure the discriminability (d') between population activity in different conditions. In the scenario depicted in **a** and **b**, changes in population activity during LL and SS trials resembled those during LS and SL trials, yielding similar d'. In the scenario depicted in **c** and **d**, changes in population activity during LL and SS trials were different from those during LS and SL trials, yielding smaller d'. (**e**) Discriminability (d') for different condition comparisons. Average trial value modulation (LL-SS) was similarly discriminable to relative value modulation (LS-SL) along the attention axis. Changes of population activity in conditions in which the absolute RF value was constant, but relative value changed (SS-SL and LS-LL), were not well discriminated along the attention axis (d' values not significantly different from 0). Circles denote mean d' values, green bars indicate 95% confidence intervals, and the shaded area indicates 95% confidence intervals associated with condition-shuffled data.

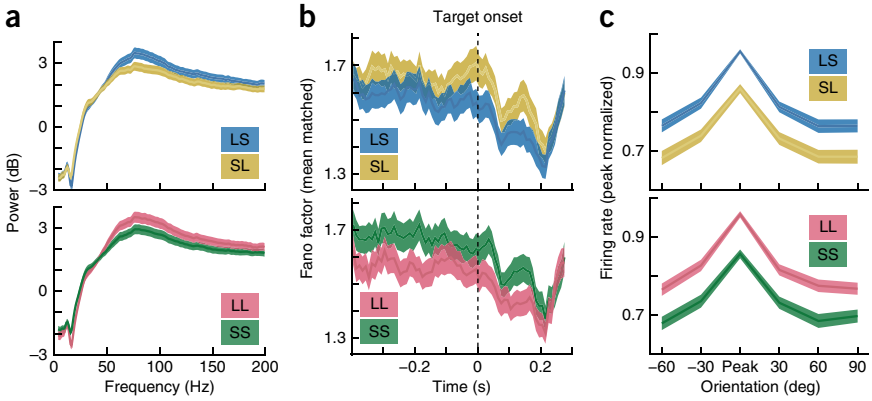

Figure 6 Relative RF value and average trial value similarly modulate power spectra, trial-to-trial reliability and orientation tuning (**a**) Increases in relative RF value and average trial value both increased power in the gamma band (40–80 Hz, $P < 10^{-6}$). Normalized power spectra are shown for each condition ($n = 69$ sites). Shading indicates ±s.e.m. (**b**) Increases in relative RF value and average trial value both increased trial-to-trial reliability. Mean-matched Fano factors aligned on target/distracter onset (bin width of 50 ms, step size of 10 ms). Shading indicates ±s.e.m. (**c**) Relative RF value and average trial value similarly modulated orientation tuning. Population orientation tuning functions ($n = 105$) from subspace reverse correlation. Orientation tuning functions for each unit were peak normalized and aligned on the preferred orientation before averaging. Shading indicates ±s.e.m.

same degree as increasing relative RF value (WSR, $P = 0.885$). Thus, despite having markedly different effects on attentional behavior and arousal, changing the relative target and average trial values similarly modulated V4 activity.

Relative value (LS-SL) and average trial value (LL-SS) manipulations similarly modulated neural activity (**Fig. 4a**), but had separable effects on attentional behavior and arousal (**Fig. 1c,d**). This dissociation between neural modulation and behavior could be reconciled if relative value and average trial value signals were encoded by distinct neuronal populations. We did not find support for this (**Fig. 4a**). First, similar proportions of neurons exhibited activity modulated by relative and average value. 143 of 190 units (75.2%) were positively modulated by relative RF value (LS-SL), achieving significance ($P < 0.05$) in 73 units (38.4%). 149 of 190 (78.4%) units were positively modulated by average trial value (LL versus SS), achieving significance ($P < 0.05$) in 80 units (42.1%). 49 units exhibited significantly positive effects ($P < 0.05$) for both relative RF value and average trial value (49 of 73, 67.1%). The proportions of units significantly positively modulated in these two comparisons were not significantly different (χ^2 test, $P = 0.633$). Moreover, the proportion of units exhibiting a significantly positive effect of both relative RF value and average trial value was greater than that expected by chance (χ^2 test, $P < 10^{-7}$). Finally, the degree of modulation observed in each comparison was correlated across units ($r = 0.31$, $P < 10^{-4}$; **Fig. 4a**). We conclude that manipulating the relative RF and average trial values lead to similar modulations in the same neurons.

For two trial comparisons (LS-LL and SS-SL; **Fig. 4b**), both relative RF value and average trial value differed, but the absolute value of the RF stimuli was constant. For example, the absolute reward associated with RF stimuli was the same on LS and LL trials, but the relative value was higher and the average trial value was lower on LS trials. Firing rates did not differ significantly between LS and LL trials (mean $d' = -0.008$, WSR, $P = 0.533$) or between SS and SL trials (mean $d' = 0.011$, WSR, $P = 0.378$) (**Fig. 4b**). We next considered whether effect sizes across these two comparisons were related. If neurons were modulated by relative value, they should be positively modulated in both LS-LL and SS-SL condition comparisons (**Fig. 4b**). We found, however, that the proportions of units exhibiting positive (χ^2 test, $P = 0.329$) or significantly positive (χ^2 test, $P = 0.946$) modulation in both comparisons were not significantly different than chance. Moreover, the two effects were not significantly correlated ($r = -0.092$, $P = 0.207$; **Fig. 4b**). Thus, modulation of neuronal responses in V4 was well described by absolute reward, and not by relative reward.

Across conditions, attentional behavior is dissociated from firing rate modulation when examining neurons individually. We considered whether a population-based analysis of neuronal data would reveal differences between conditions accounting for the behavioral effects in our task. We addressed this by asking how changes in population activity associated with attentional behavior (LS and SL) are related

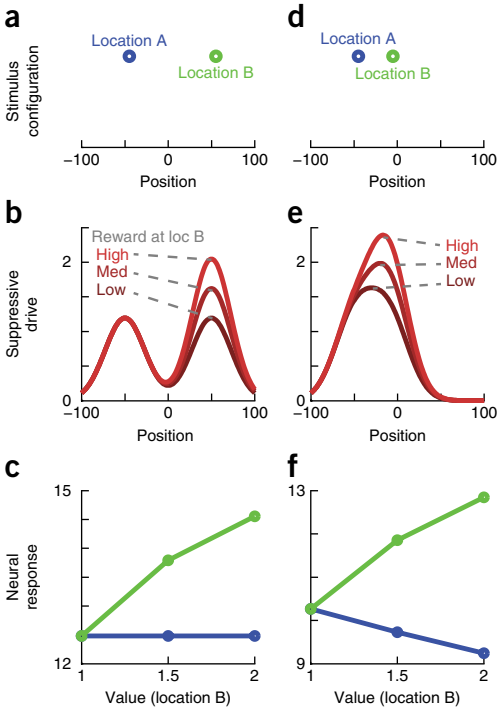

Figure 7 Normalization predicts interaction between reward and spatial scale. (**a**–**c**) When stimuli are far apart, top-down modulation by absolute reward at location B does not affect responses at location A. (**a**) Locations of the two Gabor stimuli, used as inputs to the normalization model. (**b**) Each stimulus drives divisive suppression of neurons with nearby RFs. Suppressive drive is here plotted as a function of position, for three different reward levels at location B. (**c**) Neuronal response from a unit with RF at location A (blue), and another unit with RF at location B (green). When the reward associated with location B increases, firing rates at location B increase, but responses at location A are unchanged. (**d**–**f**) When stimuli are close together, top-down modulation by absolute reward at location B decreases responses at location A. When stimuli are close together, suppressive fields overlap (**e**), and increases in reward at location B are associated with decreased responses at location A (**f**).

to changes in population activity in other condition comparisons. To do this, we treated the activity of all neurons in the LS and SL trials as points in a space in which the activity of each neuron defines a dimension (**Fig. 5a–d** and Online Methods). In this space, we defined an 'attention axis'[31] as the vector linking the means of LS and SL trials. This allowed us to ask how population activity in different conditions varies along the axis most associated with the perceptual benefits of spatial attention.

We characterized modulation across conditions by projecting onto the attention axis and quantifying discriminability (d'). Consistent with the individual neuron analysis (**Fig. 4a**), we found that average trial value (LL-SS, mean $d' = 0.726$) and relative RF value (LS-SL, mean $d' = 0.691$) similarly modulated population representations (**Fig. 5e**). Discriminability of the two pairwise comparisons did not differ significantly (randomization test, $P > 0.05$) and both differed significantly from chance (randomization test, $P < 0.05$). That is, the pattern of population activity varied similarly during average trial value (LL-SS) and relative value (LS-SL) manipulations. By contrast, in both pairwise comparisons in which reward in the RF was fixed (SS-SL and LS-LL), discriminability along the attention axis was markedly decreased. For SS compared with SL trials, discriminability was slightly, but not significantly, positive (mean $d' = 0.119$, randomization test, $P > 0.05$). For the LS-LL trial comparison, discriminability was slightly, but not significantly, negative (mean $d' = -0.147$, randomization test, $P > 0.05$). Thus, changes in population activity did not account for modulations in behavioral measures of spatial attention.

In addition to firing rate changes, prior studies have also reported that attention lowers trial-to-trial variability in neural responses to attended stimuli[11], changes gamma-band LFP power[7–9] and modifies stimulus tuning functions[32–34]. We found that relative RF value and average trial value manipulations had similar effects for all these measures. First, both had similar effects on LFPs (**Fig. 6a** and **Supplementary Fig. 5**), producing increased gamma band power (40–80 Hz, WSR, $P < 10^{-6}$). Power modulation by relative RF value and average trial value did not differ significantly in the gamma band (WSR, $P = 0.776$). Second, both similarly decreased trial-to-trial variability measured using the Fano factor (LS-SL: MI = −0.022, WSR, $P = 0.014$; LL-SS: MI = −0.029, WSR, $P = 0.001$; **Fig. 6b**), and the extent of modulation by relative and average trial value did not differ significantly (WSR, $P = 0.920$). Finally, we characterized orientation tuning functions using reverse correlation and found that both relative RF value and average trial value manipulations predominantly affected tuning functions by additive shifts (WSR, $P < 10^{-4}$) rather than multiplicative scaling (WSR, $P > 0.05$). Although this contrasts with findings of multiplicative effects in visual cortex[32], modulations inconsistent with pure multiplicative gain effects are also frequently observed[34–36]. Furthermore, recent studies have found weak evidence for clearly distinguishing between additive and multiplicative models of response enhancement by attention[34,36]. Notably, the extent of additive and multiplicative changes to orientation tuning functions did not differ between relative RF value and average trial value manipulations (WSR, $P > 0.2$; **Fig. 6c**). Thus, modulation resulting from average trial value was largely indistinguishable from modulation resulting from relative RF value, suggesting that absolute reward is a key determinant of neuronal modulation in V4.

DISCUSSION

We discovered that neural modulation in V4 can be dissociated from the perceptual benefits of spatial attention; performance could be enhanced without neural modulation and neural activity could be modulated without substantial perceptual improvement. These findings are inconsistent with signal-to-noise based mechanisms of attention that depend on the quality of V4's sensory representation at a given location. Our data instead suggest that modulation of neural activity in V4 reflects the absolute value of stimuli in the RF, and behavioral benefits accrue only when there is an imbalance in response modulation across V4. These observations imply that, at least in our task, efficient selection is the critical mechanism linking neural modulation in V4 to the behavioral benefits of attention. According to this hypothesis, the characteristic response enhancement observed in attentive states may primarily reflect a weighting of sensory signals by their associated absolute reward, which serves to bias competitive selection by downstream brain areas[4,24,26,28]. Together, these data argue that the link between neural modulation in V4 and its perceptual consequence may depend critically on downstream readout.

Although our results may appear to be at odds with prior studies of attention in V4, our results were actually consistent with previous findings when considering comparable reward configurations. Consistent with prior reports, when we manipulated relative value (LS-SL), we observed large changes in attentional behavior that were associated with enhanced firing rate[6], gamma-band power[7–9] and trial-to-trial reliability[11]. In several prior studies, investigators have described robust trial-to-trial correlations between neural correlates of attention in V4 and behavior[9,31], suggesting that neural modulation in V4 may be causally involved in conferring the associated behavioral changes. However, the observation of trial-to-trial correlations does not indicate the mechanism by which V4 neural modulation may affect behavior. Fluctuations in V4 activity at relevant locations could be correlated with perceptual performance either by modulating the signal-to-noise ratio in the relevant sensory representation or by biasing selection of relevant locations by downstream processes. These mechanisms can be distinguished by considering two reward conditions (LL and SS), which have rarely been studied in the context of visual attention. Varying average trial value while holding relative value fixed (LL-SS) had minimal effects on behavioral measures of attention, but the effects on V4 activity were comparable to relative value manipulations (LS-SL). This marked dissociation of neural activity and perceptual performance suggests that, at least in our task, V4 neural modulation contributes to performance by promoting selection rather than by enhancing signal-to-noise.

Enhanced performance in the detection tasks used in many attention studies can result from both increasing sensitivity as well as shifting criterion to respond more readily to cued than to un-cued locations. A recent study[29] manipulated the relative and absolute reward of two locations to induce changes in sensitivity without modifying criterion; other reward manipulations changed criterion without affecting sensitivity. The authors found that modulation of neural activity in V4 was associated with sensitivity changes, but not criterion shifts. Our data are also compatible with these results. We used a dual discrimination task in which discriminanda are equally likely to fall on either side of an orientation category boundary; biasing choices toward horizontal or vertical categories was therefore maladaptive. Consequently, our task conditions are most analogous to conditions in which the previous study[29] changed sensitivity, not response bias. When considering comparable conditions as the previous study, we observed that V4 activity was enhanced along with behavioral sensitivity (for example, LS versus SL). However, we independently manipulated relative and absolute reward; data from these conditions demonstrate that behavioral sensitivity can increase without V4 modulation (SS versus SL) and that V4 modulation is not always

accompanied by commensurate changes in behavioral sensitivity (LL versus SS). These data reveal that the sensitivity changes quantified by signal detection theory are likely implemented by selection, rather than by signal-to-noise, mechanisms.

The manipulation of average trial value in our task may recruit similar neural mechanisms to those involved with arousal[37–39] and the heightened 'cognitive effort' associated with more difficult tasks[30,40,41]. These non-spatial vigilance factors are associated with enhanced neural activity[37–39], but the link between enhanced neural activity in these settings and perception remains unclear. The neural correlates of average trial value have rarely been reported in primate visual cortex. Beyond sensory cortex, neural activity in many brain areas is modulated by reward-related parameters that can be used to compute average trial value, including orbitofrontal cortex[21,42], dorsolateral prefrontal cortex[43], cingulate cortex[44], LIP[16,17], striatum[45], basal forebrain[46] and the amygdala[18,20]. Some of these brain areas may provide direct or indirect inputs to V4 that underlie the modulation that we observed.

Given that manipulating relative reward value across the visual field shifts spatial attention, brain areas that reflect the selection underlying attentional behavior should exhibit signals modulated by relative, rather than absolute, reward value. However, the presence of modulation by relative reward may also reflect the spatial scale of the task used to characterize neural response properties. For example, modulation of neural activity in area V1 is correlated with relative value when two stimuli appear in the visual field[19]. One critical difference between the task used to make this observation and ours is the spatial separation of visual stimuli. Although our data demonstrate that rewards associated with stimuli in the opposite hemi-field do not modulate activity in V4, the prior study employed competing saccade targets placed near each other in the same quadrant of the visual field. Notably, the pattern of results observed in both studies is predicted by the normalization model of attention[47], where the response of a neuron to a stimulus in its RF is suppressed (or normalized) by a pool of neurons responsive to a broader range of features and spatial positions. We demonstrated this by using the normalization model of attention to simulate the response of a population of neurons to two visual stimuli that vary in location and associated reward (**Fig. 7**). We found that top-down response enhancement by absolute reward leads to suppression of responses to stimuli presented at locations nearby, but not far away, with the spread of suppression determined by the spatial extent of neurons in the normalization pool. If these considerations are correct, then V1 activity would not reflect relative value when tested with two stimuli located far enough from each other such that neurons in the normalization pool do not respond to both stimuli. In the model, the details of this change are dictated by the spatial scale of the top-down enhancement (attention field) and the spatial scale of the normalization pool. The spatial scale of top-down enhancement is likely task dependent[47]; this emphasizes the critical importance of studying different brain areas with identical tasks to determine the contribution of relative and absolute reward on neural firing rates.

One brain area that exhibits relative value modulation at a larger spatial scale is the lateral intraparietal area (LIP). Our data demonstrate that the classic attention-related modulations of V4 activity can be decoupled from behavioral measures of spatial attention. LIP receives inputs from V4, but displays relative value modulation characteristic of competitive selection both in[48] and across hemi-fields[16,17,49], suggesting that LIP may transform V4 activity into relative value through response normalization[49]. LIP is unlikely to be solely responsible for selecting behaviorally relevant sensory signals[50]. Similar transformations may occur at multiple points along sensorimotor pathways, which is supported by recent evidence in visual search[25]. Characterization of these transformations is likely to be critical for understanding the sequential and interactive processing carried out in different brain areas to confer the perceptual benefits of attention. Our data indicate that the manipulation of absolute and relative reward values while measuring attentional behavior is a powerful tool for providing this characterization.

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

We thank S. Bouret, E. Bromberg-Martin, V. Ciaramitaro, K. Louie and F. Pestilli for helpful comments on an earlier version of the manuscript, S. Dashnaw for MRI support, G. Asfaw for veterinary support, and K. Marmon for technical support. This research was supported by grants to C.D.S. from the US National Institute of Mental Health (NIMH) (R01 MH082017) and by a core grant from the US National Eye Institute (NEI) (P30-EY19007) to Columbia University. B.L. received support from the NIMH (T32-MH015144) and the Helen Hay Whitney Foundation. J.K.B. received support from the NEI (T32-EY013933) and the Columbia University Medical Scientist Training Program.

AUTHOR CONTRIBUTIONS

B.L. designed the experiment. J.K.B. and B.L. collected the data and wrote the manuscript. J.K.B. analyzed the data with assistance from B.L. C.D.S. supervised and provided input about all aspects of the project and edited the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Carrasco, M. Visual attention: the past 25 years. *Vision Res.* **51**, 1484–1525 (2011).
2. Awh, E., Belopolsky, A.V. & Theeuwes, J. Top-down versus bottom-up attentional control: a failed theoretical dichotomy. *Trends Cogn. Sci.* **16**, 437–443 (2012).
3. Desimone, R. & Duncan, J. Neural mechanisms of selective visual attention. *Annu. Rev. Neurosci.* **18**, 193–222 (1995).
4. Krauzlis, R.J., Bollimunta, A., Arcizet, F. & Wang, L. Attention as an effect not a cause. *Trends Cogn. Sci.* **18**, 457–464 (2014).
5. Maunsell, J.H.R. & Cook, E.P. The role of attention in visual processing. *Phil. Trans. R. Soc. Lond. B* **357**, 1063–1072 (2002).
6. Moran, J. & Desimone, R. Selective attention gates visual processing in the extrastriate cortex. *Science* **229**, 782–784 (1985).
7. Fries, P., Reynolds, J.H., Rorie, A.E. & Desimone, R. Modulation of oscillatory neuronal synchronization by selective visual attention. *Science* **291**, 1560–1563 (2001).
8. Gregoriou, G.G., Gotts, S.J., Zhou, H. & Desimone, R. High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science* **324**, 1207–1210 (2009).
9. Womelsdorf, T., Fries, P., Mitra, P.P. & Desimone, R. Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).
10. Cohen, M.R. & Maunsell, J.H.R. Attention improves performance primarily by reducing interneuronal correlations. *Nat. Neurosci.* **12**, 1594–1600 (2009).
11. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Differential attention-dependent response modulation across cell classes in macaque visual area V4. *Neuron* **55**, 131–141 (2007).
12. Mitchell, J.F., Sundberg, K.A. & Reynolds, J.H. Spatial attention decorrelates intrinsic activity fluctuations in macaque area V4. *Neuron* **63**, 879–888 (2009).
13. Maunsell, J.H.R. Neuronal representations of cognitive state: reward or attention? *Trends Cogn. Sci.* **8**, 261–265 (2004).
14. Lauwereyns, J., Watanabe, K., Coe, B. & Hikosaka, O. A neural correlate of response bias in monkey caudate nucleus. *Nature* **418**, 413–417 (2002).
15. Pastor-Bernier, A. & Cisek, P. Neural correlates of biased competition in premotor cortex. *J. Neurosci.* **31**, 7083–7088 (2011).
16. Platt, M.L. & Glimcher, P.W. Neural correlates of decision variables in parietal cortex. *Nature* **400**, 233–238 (1999).
17. Sugrue, L.P., Corrado, G.S. & Newsome, W.T. Matching behavior and the representation of value in the parietal cortex. *Science* **304**, 1782–1787 (2004).

18. Peck, C.J., Lau, B. & Salzman, C.D. The primate amygdala combines information about space and value. *Nat. Neurosci.* **16**, 340–348 (2013).
19. Stănişor, L., van der Togt, C., Pennartz, C.M. & Roelfsema, P.R. A unified selection signal for attention and reward in primary visual cortex. *Proc. Natl. Acad. Sci. USA* **110**, 9136–9141 (2013).
20. Paton, J.J., Belova, M.A., Morrison, S.E. & Salzman, C.D. The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature* **439**, 865–870 (2006).
21. Roesch, M.R. & Olson, C.R. Neuronal activity related to reward value and motivation in primate frontal cortex. *Science* **304**, 307–310 (2004).
22. Varazzani, C., San-Galli, A., Gilardeau, S. & Bouret, S. Noradrenaline and dopamine neurons in the reward/effort trade-off: a direct electrophysiological comparison in behaving monkeys. *J. Neurosci.* **35**, 7866–7877 (2015).
23. Serences, J.T. & Kastner, S. A multi-level account of selective attention. In *The Oxford Handbook of Attention* (Nobre, A.C. & Kastner, S.) 76–104 (Oxford University Press, 2014).
24. Pestilli, F., Carrasco, M., Heeger, D.J. & Gardner, J.L. Attentional enhancement via selection and pooling of early sensory responses in human visual cortex. *Neuron* **72**, 832–846 (2011).
25. Mirpour, K. & Bisley, J.W. Dissociating activity in the lateral intraparietal area from value using a visual foraging task. *Proc. Natl. Acad. Sci. USA* **109**, 10083–10088 (2012).
26. Chen, Y. & Seidemann, E. Attentional modulations related to spatial gating but not to allocation of limited resources in primate V1. *Neuron* **74**, 557–566 (2012).
27. Itthipuripat, S., Ester, E.F., Deering, S. & Serences, J.T. Sensory gain outperforms efficient readout mechanisms in predicting attention-related improvements in behavior. *J. Neurosci.* **34**, 13384–13398 (2014).
28. Zénon, A. & Krauzlis, R.J. Attention deficits without cortical neuronal deficits. *Nature* **489**, 434–437 (2012).
29. Luo, T.Z. & Maunsell, J.H.R. Neuronal modulations in visual cortex are associated with only one of multiple components of attention. *Neuron* **86**, 1182–1188 (2015).
30. Spitzer, H., Desimone, R. & Moran, J. Increased attention enhances both behavioral and neuronal performance. *Science* **240**, 338–340 (1988).
31. Cohen, M.R. & Maunsell, J.H.R. A neuronal population measure of attention predicts behavioral performance on individual trials. *J. Neurosci.* **30**, 15241–15253 (2010).
32. McAdams, C.J. & Maunsell, J.H. Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
33. Reynolds, J.H., Pasternak, T. & Desimone, R. Attention increases sensitivity of V4 neurons. *Neuron* **26**, 703–714 (2000).
34. Williford, T. & Maunsell, J. Effects of spatial attention on contrast response functions in macaque area V4. *J. Neurophysiol.* **96**, 40–54 (2006).
35. Thiele, A., Pooresmaeili, A., Delicato, L.S., Herrero, J.L. & Roelfsema, P.R. Additive effects of attention and stimulus contrast in primary visual cortex. *Cereb. Cortex* **19**, 2970–2981 (2009).
36. Sanayei, M., Herrero, J.L., Distler, C. & Thiele, A. Attention and normalization circuits in macaque V1. *Eur. J. Neurosci.* **41**, 949–964 (2015).
37. McGinley, M.J. *et al.* Waking state: rapid variations modulate neural and behavioral responses. *Neuron* **87**, 1143–1161 (2015).
38. Harris, K.D. & Thiele, A. Cortical state and attention. *Nat. Rev. Neurosci.* **12**, 509–523 (2011).
39. Reimer, J. *et al.* Pupil fluctuations track fast switching of cortical states during quiet wakefulness. *Neuron* **84**, 355–362 (2014).
40. Boudreau, C.E., Williford, T.H. & Maunsell, J.H.R. Effects of task difficulty and target likelihood in area V4 of macaque monkeys. *J. Neurophysiol.* **96**, 2377–2387 (2006).
41. Ruff, D.A. & Cohen, M.R. Global cognitive factors modulate correlated response variability between V4 neurons. *J. Neurosci.* **34**, 16408–16416 (2014).
42. Padoa-Schioppa, C. & Assad, J.A. Neurons in the orbitofrontal cortex encode economic value. *Nature* **441**, 223–226 (2006).
43. Leon, M.I. & Shadlen, M.N. Effect of expected reward magnitude on the response of neurons in the dorsolateral prefrontal cortex of the macaque. *Neuron* **24**, 415–425 (1999).
44. McCoy, A.N., Crowley, J.C., Haghighian, G., Dean, H.L. & Platt, M.L. Saccade reward signals in posterior cingulate cortex. *Neuron* **40**, 1031–1040 (2003).
45. Kim, H.F. & Hikosaka, O. Distinct basal ganglia circuits controlling behaviors guided by flexible and stable values. *Neuron* **79**, 1001–1010 (2013).
46. Peck, C.J. & Salzman, C.D. The amygdala and basal forebrain as a pathway for motivationally guided attention. *J. Neurosci.* **34**, 13757–13767 (2014).
47. Reynolds, J.H. & Heeger, D.J. The normalization model of attention. *Neuron* **61**, 168–185 (2009).
48. Falkner, A.L., Krishna, B.S. & Goldberg, M.E. Surround suppression sharpens the priority map in the lateral intraparietal area. *J. Neurosci.* **30**, 12787–12797 (2010).
49. Louie, K., Grattan, L.E. & Glimcher, P.W. Reward value-based gain control: divisive normalization in parietal cortex. *J. Neurosci.* **31**, 10627–10639 (2011).
50. Moore, T., Armstrong, K.M. & Fallah, M. Visuomotor origins of covert spatial attention. *Neuron* **40**, 671–683 (2003).

ONLINE METHODS

Animals and implantation. Two male rhesus monkeys (*Macaca mulatta*, 8–13 kg, 7–8 years old) were used in these experiments. All experimental procedures complied with US National Institutes of Health guidelines and were approved by the Institutional Animal Care and Use Committees at the New York State Psychiatric Institute and Columbia University. Prior to training, we implanted a plastic head post secured to the skull using ceramic bone screws. Surgery was conducted using aseptic techniques under isoflurane anesthesia, and analgesics and antibiotics were administered postsurgically. After the monkeys were behaviorally trained, we acquired T1-weighted MRIs with fiducial markers attached to the head post. In a second surgery, we implanted a plastic recording chamber over dorsal visual area V4 guided by a neuronavigation system (Brainsight, Rogue Research) registered to the MRI for each monkey. Recordings targeted the lunate gyrus posterior to the junction of the superior temporal sulcus and the sylvian fissure.

Data acquisition. Monkeys were seated and head-restrained in a darkened sound-attenuating booth. Eye position and pupil diameter were monitored using an infrared camera sampled at 1,000 Hz (Eyelink, SR Research). Visual stimuli were generated using EXPO (Center for Neural Science, New York University) and were displayed on a CRT monitor positioned 57 cm away from the monkey.

We recorded from the right dorsal V4 of each monkey using one or two electrodes individually lowered with a multiple-electrode microdrive (NaN Instruments). Extracellular activity was recorded using epoxylite-insulated tungsten electrodes (8–10-MΩ impedance, FHC) or glass-coated tungsten electrodes (0.5–2.0-MΩ impedance, Alpha Omega). Analog signals were amplified, band-pass filtered (250–7,500 Hz) and sampled (30 kHz) for unit isolation (Blackrock Microsystems). Units were isolated using manual clustering on the basis of several waveform parameters including principal components, peak and trough amplitudes, as well as the presence of a refractory period (Plexon Offline Sorter, Plexon). LFPs were filtered between 0.3 and 500 Hz and sampled at 1 kHz.

Behavioral task and visual stimuli. Monkeys performed a dual 2-AFC orientation discrimination task for liquid reward. Trials were initiated when monkeys fixated a central spot. Monkeys were required to maintain fixation throughout the trial in a window of radius 1.0–1.5°. After 300 ms of fixation, two reward cues appeared for 250 ms. Reward cues were uniform chromatic discs of 0.5° diameter, 2.0–2.5° eccentricity, and an angular position matched to subsequently appearing Gabor patches. Reward cues indicated the amount of juice associated with correct discrimination of the associated Gabor patch. For each monkey, we employed two different, luminance-matched cue sets. Cue colors were defined in DKL color space (large reward azimuths: 220° and 340° (M1), 20° and 135° (M2); small reward azimuths: 90° and 280° (M1), 220° and 280° (M2)). Cue sets did not substantially affect any of the reported behavioral results. Considering separately the subset of trials from each cue set in each monkey, performance was always highest in LS (CMH, $P < 10^{-4}$) and lowest in SL (CMH, $P < 10^{-10}$), reaction times were fastest in LS (WRS, $P < 10^{-10}$) and slowest in SL (WRS, $P < 10^{-10}$), pupils were most dilated in LL (CMH, $P < 10^{-6}$) and least dilated in SS (CMH, $P < 10^{-10}$), and aborted trials were most frequent in SS (CMH, $P < 0.05$). Similarly, performing the regression analysis separately on the subset of trials from each cue set in each monkey yielded similar results. For each cue set in each monkey, associating a large reward with RF stimuli increased firing rates (WSR, $P < 10^{-6}$), and associating a large reward with the location opposite RF did not significantly alter firing rates (WSR, $P > 0.1$).

The offset of the reward cues was followed by a 350 ms delay, after which two streams of stimuli appeared, one in the receptive field of the neuron under study, and the other diametrically opposed in the opposite hemi-field. The streams were composed of presentations of Gabor patches lasting 20 ms, interleaved with blank stimuli of the same duration matched to the background luminance of the monitor (probability of blank = 0.15). The orientation of each Gabor presentation was independently and randomly drawn from a set of 6 equally spaced orientations (0°, 30°, 60°, 90°, 120° and 150° from horizontal). The spatial frequency and size of Gabor patches were tailored to the neuron under study. On each trial, the streams of Gabor presentations were stopped at a random time, which was determined by adding a fixed minimal duration to a random draw from an exponential distribution (truncated at a maximum of 2 s) to approach a flat hazard rate[51]. Across all experiments, the average stream duration was 685.7 ms.

Following the streams of Gabor presentations, a final pair of Gabors appeared at the two stimulus locations. We term these two final Gabor presentations the target and distracter (together, the discriminanda). Discriminanda were distinguished from prior stimuli in the stream primarily by a longer presentation duration. Mean discriminanda duration was 96.7 ± 2.0 ms, which was adjusted to maintain consistent performance. The discriminanda were followed by noise masks lasting 60 ms. Choice targets then appeared at one of the two locations, which was the monkeys' first indication of the identity of the target and distracter. In other words, the location about which monkeys would be asked to render a perceptual decision was not revealed to the monkey until after extinction and masking of the discriminanda. Thus, the behavioral relevance of discriminanda was exclusively determined by the associated reward signaled by prior appearing cues. Choice targets flanked the location of extinguished discriminanda, and monkeys reported whether the orientation of the target was more horizontal or more vertical by saccading to the choice target nearest the horizontal meridian or the choice target nearest the vertical meridian, respectively. Discriminanda orientations were randomly selected on each trial and varied between 3° (difficult) and 45° (easy) away from the category boundary (corresponding to 0–90° from horizontal).

Correct trials were rewarded as indicated by the cue associated with the target. Large rewards were 4–5-fold larger than small rewards. For M1, the small reward was 0.10 ml, and the large reward was 0.50 ml. For M2, the small reward was 0.07 ml, and the large reward was 0.31 ml.

Fixation breaks before the appearance of choice targets resulted in trials in which the payoff structure, queried location and discriminanda orientation (but not the duration of Gabor streams) were repeated. This ensured that monkeys could not increase their reward rate by aborting trials of low expected value. Following completed trials, both correct and incorrect, the subsequent trial type was determined pseudorandomly by sampling without replacement from a 16-element matrix (4 conditions × 2 cue sets × 2 repetitions).

Data analysis. All statistical tests were two-sided. In most cases, we employed non-parametric statistical tests. When parametric tests were employed, data distributions were assumed to be normal but this was not formally tested. We did not perform analyses blind to the identity of trial types. We did not run any statistical test to determine sample sizes a priori, but our sample sizes are similar to those generally employed in the field.

Conditions. For performance and reaction time data, conditions are defined with respect to the queried location (**Fig. 1b**). For example, for performance and reaction time data, LS refers to the unbalanced condition where the monkey is asked to report the orientation of the stimulus associated with large reward. By contrast, firing rates were taken from an epoch before the appearance of saccade targets, before the queried location is determined. Thus, for neural data, the unbalanced conditions are defined with respect to the location of the receptive field. When referring to neural data, LS refers to the unbalanced condition where the large reward stimulus is in the receptive field.

To characterize the influence of reward expectation on neural activity and behavior, we distinguish between three types of value. Absolute value refers to the reward associated with queried or RF locations, independent of the value of other stimuli in the visual field. Relative value refers to the fractional payoff associated with the queried or RF location. Average trial value describes the overall reward expectancy of the trial as the average reward for a trial before identification of target and distracter locations.

Behavior. For proportion data (percent correct and abort rate), we assessed statistical significance using the Cochran-Mantel-Haenszel test. Aborted trials were trials where monkeys broke fixation after onset of reward cues. For comparisons of reaction time and pupil diameter across conditions, we assessed statistical significance using two-tailed Wilcoxon rank-sum tests. Reaction times were defined as the beginning of a choice target-directed saccade. Pupil diameter values were taken from the same epoch as spike counts (300 ms before onset of discriminanda) and z-scored by session.

Regression. We used multiple linear regression to quantify how the reward contingency at each location affected neuronal activity and behavior across the population. The regression was performed on spike rates in the 300 ms epoch before

doi:10.1038/nn.4141

onset of discriminanda. We chose this time epoch because it is close to measured behavior, yet before the appearance of discriminanda. Therefore, firing rates in this epoch are not affected by ultimate discriminanda orientation, attendant differential reward expectation (for easy versus difficult discriminanda), and/or presaccadic activity related to choice. For each neuron, we regressed the firing rate in this epoch onto two predictors and a constant term.

$$FR = \beta_0 + \beta_1 * x_1 + \beta_2 * x_2$$

where x_1 is a categorical predictor indicating whether the receptive field location is associated with large reward ($x_1 = 1$ for LS and LL trials, $x_1 = 0$ for SL and SS trials), and x_2 is a categorical predictor indicating whether the opposite location is associated with large reward ($x_2 = 1$ for SL and LL trials, $x_2 = 0$ for SS and LS trials). Therefore, β_1 indicates the effect on firing rate of high value in the receptive field and β_2 indicates the effect of high value opposite the receptive field. If neuronal modulation reflected the relative value of the discriminanda in the receptive field, this would be captured in the regression coefficients by oppositely signed β_1 and β_2. To assess statistical significance of regression coefficients, we performed two-tailed Wilcoxon signed-rank tests. We also examined an expanded model, which included reaction time and pupil diameter as additional predictors, which yielded similar results (**Supplementary Fig. 7**).

To characterize reward modulation of behavior, we performed the same regression on our single-trial resolved measures of behavior. For each session, we regressed reaction times and pupil diameters from the 300 ms epoch before discriminanda onset onto two predictors and a constant term.

$$RT = \beta_0 + \beta_1 * x_1 + \beta_2 * x_2$$

where x_1 is a categorical predictor indicating whether the target is associated with a large reward, and x_2 is a categorical predictor indicating whether the distracter associated with a large reward. As above, β_1 indicates the effect on behavior of high value at the queried location and β_2 indicates the effect of high value opposite the queried location.

Effect size (*d′*) analysis. We assessed differences in firing rate across conditions using *d′*, defined as

$$d' = (\mu_1 - \mu_2)/\sqrt{((SS_1 + SS_2)/(df_1 + df_2))}$$

where μ_X is the mean, SS_X is the sum of squares and df_X is degrees of freedom (number of trials − 1) for each condition. We assessed the statistical significance of *d′* values across all neurons using two-tailed Wilcoxon signed-rank tests. We assessed the statistical significance of individual neuron *d′* values by a randomization test. To perform randomization tests, reference distributions were constructed by computing *d′* values for 10,000 random assignments of conditions to trials. Comparisons were deemed significant if >97.5% of the reference distribution fell on one side of the observed value (equivalent to a two-tailed test at $\alpha = 0.05$). *d′* values were computed from firing rates in the 300 ms epoch before onset of discriminanda.

We also quantified effect sizes using MI, defined as $MI = (a - b)/(a + b)$, where *a* is the mean in the modulated condition (for example, LS and LL), and *b* is the mean in the reference condition (for example, SL and SS). Quantifying effect sizes using modulation indices did not alter any of the reported results.

We used χ^2 tests to assess the statistical significance of differences in the proportion of units modulated by each pairwise condition comparison (for example, relative value (LS-SL) versus average value (LL-SS)), as well as the independence of this modulation across units.

Attention axis analysis. Population activity was analyzed using peak-normalized spike counts from the 300-ms epoch preceding discriminanda onset. We sought to characterize how activity across a population of V4 neurons changes in the different reward conditions. Since we did not record activity from many neurons simultaneously, we constructed a population response for each trial using the following procedure. First, to equalize trial numbers for each neuron-condition, we randomly resampled trials with replacement. Trial order in each neuron-condition was then shuffled to generate trial activity for a population of the same size as the number of neurons we recorded ($n = 190$). Note that while this procedure allows us to examine population activity in a multivariate manner, it is limited in that we cannot determine the role of inter-neuronal correlations or across-trial fluctuations since neurons were not recorded simultaneously.

A random half of trials were selected to define the attention axis[31]. The remaining half of trials were projected onto the attention axis, and used to calculate the discriminability (*d′*) of pairwise condition comparisons. We repeated this process 1,000 times. In **Figure 5**, we plotted the mean *d′* across runs, as well as the intervals containing the middle 95% of values across runs. These confidence intervals were used to determine whether *d′* values differed significantly from each other. We determined whether *d′* values differed significantly from chance by using a randomization test. We randomly assigned conditions to trials and then computed discriminability (*d′*) values along the attention axis as described above. The 95% confidence intervals for this reference distribution are plotted as the shaded area in **Figure 5**. Mean *d′* values for a given comparison were deemed significant if >97.5% of the reference distribution fell on one side of the observed value (equivalent to a two-tailed test at $\alpha = 0.05$).

Spectral analysis. We estimated power spectra using a multi-taper algorithm[52] implemented in the Chronux toolbox (http://www.chronux.org), using seven tapers and a time-bandwidth product of 5. As with analyses of spike rate, we used the 300-ms epoch directly preceding onset of the discriminanda to compute power spectra. Raw power spectra during this period were converted to decibels with respect to reference power spectra collected during the 300-ms epoch preceding reward cue onset. To assess statistical significance in defined frequency bands, we used two-tailed Wilcoxon signed-rank tests.

Fano factor analysis. Fano factors were computed from spike counts in the 300-ms epoch directly preceding discriminanda onset. We characterized Fano Factor modulation using a modulation index, calculated as described above. Statistical significance of Fano factor modulation was assessed using a two-tailed Wilcoxon signed-rank test on the distribution of MIs. To determine whether Fano factor differences could be attributed to differences in firing rate across conditions, we performed a mean-matching procedure (adapted from[53]). We computed Fano factors in sliding time bins (50-ms width, 10-ms steps) using subdistributions of units with matched spike count distributions across all four conditions. Because our interest was in differences across conditions (but not across time), we allowed the mean count distributions to fluctuate across time bins.

Orientation tuning analysis. We characterized orientation-tuning functions using reverse correlation[54]. To construct orientation-tuning functions, we first determined a unit-specific time window in which to analyze responses to individual 20-ms Gabor presentations (**Supplementary Fig. 8**). Unit-specific time windows were employed because units displayed diverse temporal responses to orientation. For example, units varied substantially in the timescale over which orientation influenced firing rates and often displayed time-varying (for example, biphasic) orientation tuning kernels. In defining spike-counting windows, we aimed to capture the epoch where orientation tuning was maximally expressed and temporally consistent. To define the counting window for each unit, we first determined the time with respect to Gabor presentation onset with the largest effect of orientation on firing rate. Effect size was computed in 5-ms-wide bins, stepped every 1 ms, using a standard eta squared measure, $\eta^2 = (SS_{ori} / SS_{total})$, where SS_{ori} is the sum squared error in firing rates explained by orientation, and SS_{total} is the total sum squared error. Spike counting windows were centered on the time of peak orientation effect and extended in both directions as far as two criteria were met. First, we mandated that η^2 values remain above 10 s.d. beyond that measured during a baseline epoch (150 ms prior – 40 ms after presentation onset). Second, we computed a sliding Spearman's rank correlation with the time of peak orientation effect size, and mandated that correlation coefficients in the window be greater than 0.3. Rank correlation coefficients were performed on the set of six spike density functions (1-ms step, $\sigma = 5$ ms) split by orientation. This second correlation criterion was necessary to restrict the counting window to an epoch where orientation tuning was temporally consistent. Using this method, the mean window size across all units was 39.6 ± 2.1 ms. For the minority of orientation-tuned units that did not meet the above criteria (53 of 171 units with a significant effect of orientation by ANOVA), we employed a conservative window of 20 ms (equal to flicker duration) centered on the time of peak effect size.

Orientation tuning functions were calculated as the average spike rate elicited by stimuli of each orientation-condition. Because most units displayed pronounced transient responses to the onset of the stream of flickering stimuli, we excluded responses to the first three presentations in the stream. Included in the population plot (**Fig. 6c**) are all neurons (n = 105) that showed (i) positive effects of relative value (LS-SL), (ii) positive effects of average trial value (LL-SS), and (iii) a significant effect of orientation by ANOVA (df = 5, $P < 0.05$). Altering selection criteria had no effect on the qualitative outcome of orientation tuning analysis. For included units, the mean number of Gabor presentations used to construct tuning functions was 22,968 ± 1,261.

We characterized tuning function modulation by fitting a separate Gaussian to orientation-tuning functions for each of the four conditions. Each Gaussian had four free parameters: μ, σ, amplitude and asymptote. Differences in μ and σ were interpreted as shifts in orientation tuning and bandwidth respectively. To summarize multiplicative scaling and additive shift of orientation tuning functions, we computed the ratio of fitted amplitudes and asymptotes respectively. We assessed statistical significance of differences in amplitude and asymptote ratios using Wilcoxon signed-rank tests on log-transformed ratios. For all statistical comparisons of Gaussian fit parameters, we included only the subset of units in **Figure 6c** with significant Gaussian fits to all four conditions, assessed using an F-test (α = 0.05) that compared the Gaussian fits to fits to the mean response across all orientations (n = 61 units with significant fits). Both average trial value and relative RF value had weakly positive, but non-significant, effects on the amplitude ratio of the fitted Gaussians (LS-SL, median ratio: 1.03, WSR, $P > 0.05$, LL-SS median ratio: 1.006, WSR, $P > 0.05$). By contrast, both average trial value and relative RF value modulations had significantly positive effects on the asymptote ratio of Gaussian fits (LS-SL: median ratio 1.10, WSR, $P < 10^{-4}$; LL-SS: median ratio 1.14, WSR, $P < 10^{-6}$). There were no significant changes to tuning bandwidth (σ).

Normalization model. We simulated reward modulation of population activity using the normalization model[47] (http://www.cns.nyu.edu/heegerlab). The normalization model describes population activity as being shaped by an excitatory stimulus drive, a divisive suppressive drive, and a multiplicative attention field. In the spatial domain, the key parameters that determine whether the representations of two stimuli are mutually suppressive are the location and sizes of the stimuli themselves and the spatial spread of the suppressive field (*IxWidth*).

Figure 7 depicts two sets of simulations. For the simulation shown in **Figure 7a–c**, we used a stimulus consisting of two Gabors, each 10 units wide, with centers separated by 100 units. For the simulation shown in **Figure 7d–f**, we used a stimulus consisting of two Gabors, each 10 units wide, with centers separated by 20 units. In both cases, we simulated three levels of top-down modulation at location B by varying the *Apeak* parameter (*Apeak* = 1, 1.5, and 2). For all other parameters, we used the default values. To characterize the spatial profile of suppressive fields (**Fig. 7c**,**e**), we plot the suppressive drive to only the segment of the population with orientation preference matched to the stimuli. To characterize response modulation of individual units (**Fig. 7c**,**f**), we plot the response of two units, each having a receptive field center and orientation preference matched to the stimuli.

A **Supplementary Methods Checklist** is available.

51. Ghose, G.M. & Maunsell, J. Attentional modulation in visual cortex depends on task timing. *Nature* **419**, 616–620 (2002).
52. Mitra, P. & Bokil, H. *Observed Brain Dynamics* (Oxford University Press, 2007).
53. Churchland, M.M. *et al.* Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat. Neurosci.* **13**, 369–378 (2010).
54. Ringach, D.L., Hawken, M.J. & Shapley, R. Dynamics of orientation tuning in macaque primary visual cortex. *Nature* **387**, 281–284 (1997).

doi:10.1038/nn.4141

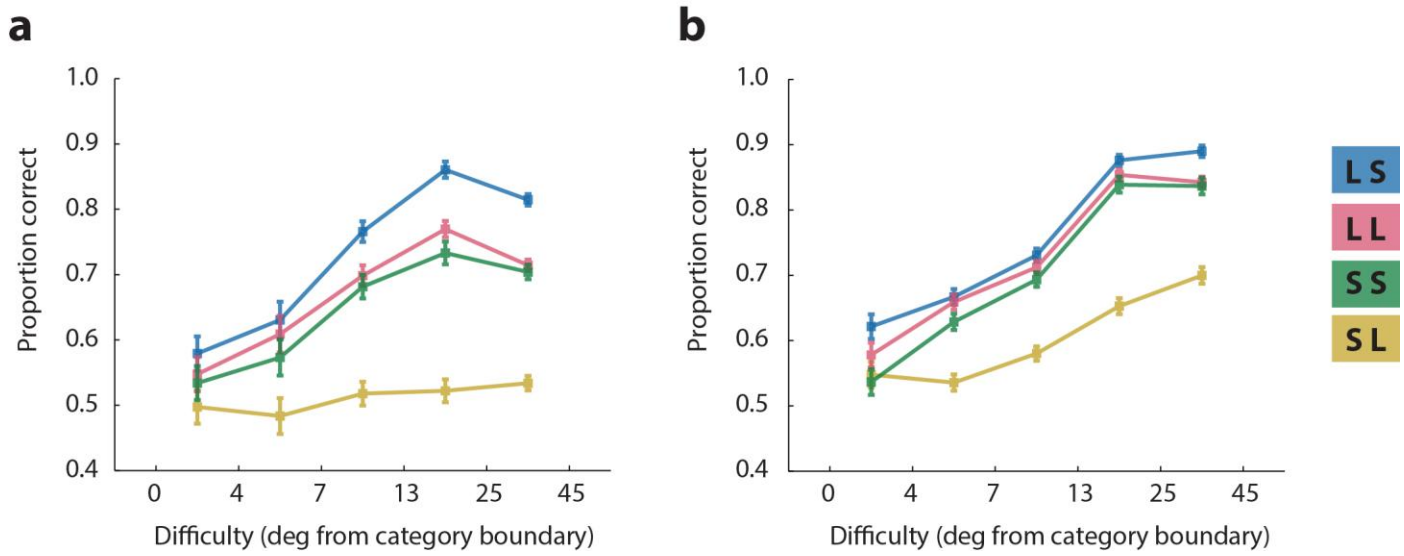

Supplementary Figure 1

Psychometric functions.

Proportion correct as a function of task difficulty, shown separately for Monkey 1 (**a**) and Monkey 2 (**b**). Labels on abscissa indicate orientation bin edges. Error bars indicate +/- 1 s.e.m.

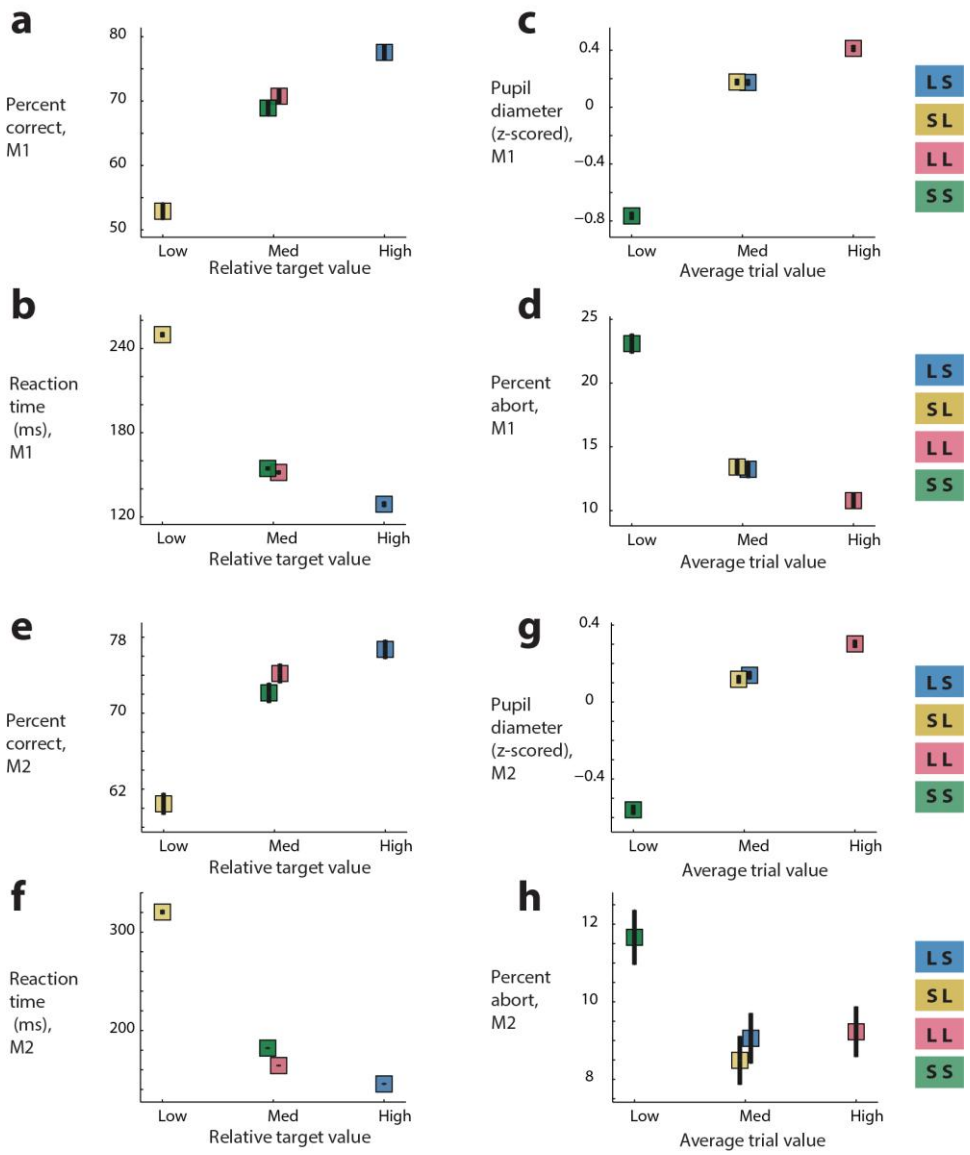

Supplementary Figure 2

Behavioral effects were consistent across monkeys.

Behavioral data are presented as function of relative target value or average trial value, from all trials completed during recording sessions in Monkey 1 (**a-d,** 5,174 LS; 5,230 LL; 5,312 SL; and 5,226 SS trials) and Monkey 2 (**e-h**, 6,881 LS; 6,965 LL; 7,018 SL; and 6,953 SS trials). Error bars indicate +/- 1 s.e.m. Overlapping points are slightly offset for visualization. **(a,b**) Performance and reaction time track relative target value in M1. Percent correct was highest in LS (CMH, $p<10^{-10}$ for all comparisons), and lowest in SL (CMH, $p<10^{-10}$ for all comparisons). Reaction times were fastest in LS (WRS, $p<10^{-10}$ for all comparisons), and slowest in SL (WRS, $p<10^{-10}$ for all comparisons). Performance was slightly higher (CMH, p=.0402) and reaction times slightly faster (WRS, $p<10^{-4}$**)** in LL compared to SS trials. **(c,d**) Pupil diameter and abort rate track average trial value in M1. Pupils were most dilated in LL (WRS, $p<10^{-10}$ for all comparisons) and least dilated in SS (WRS, $p<10^{-10}$ for all comparisons). Abort rate was highest in SS (CMH, $p<10^{-10}$ for all comparisons) and lowest in LL (CMH, $p<10^{-6}$ for all comparisons). **(e,f**) Performance and reaction time track relative target value in M2. Percent correct was highest in LS (CMH, $p<10^{-3}$ for all comparisons), and lowest in SL (CMH, $p<10^{-10}$ for all comparisons). Reaction times were fastest in LS (WRS, $p<10^{-10}$ for all comparisons), and slowest in SL (WRS, $p<10^{-10}$ for all comparisons). Performance was slightly higher (CMH, p=.0066) and reaction times faster (WRS, $p<10^{-10}$**)** in LL compared to SS trials. **(g,h**) Pupil diameter and abort rate track average trial value in M2. Pupils were most dilated in LL (WRS, $p<10^{-10}$ for all comparisons) and least dilated in SS (WRS, $p<10^{-10}$ for all comparisons). Abort rate was highest in SS (CMH, $p<10^{-5}$ for all comparisons), and did not significantly differ between the other three conditions (CMH, p>0.1 for all comparisons).

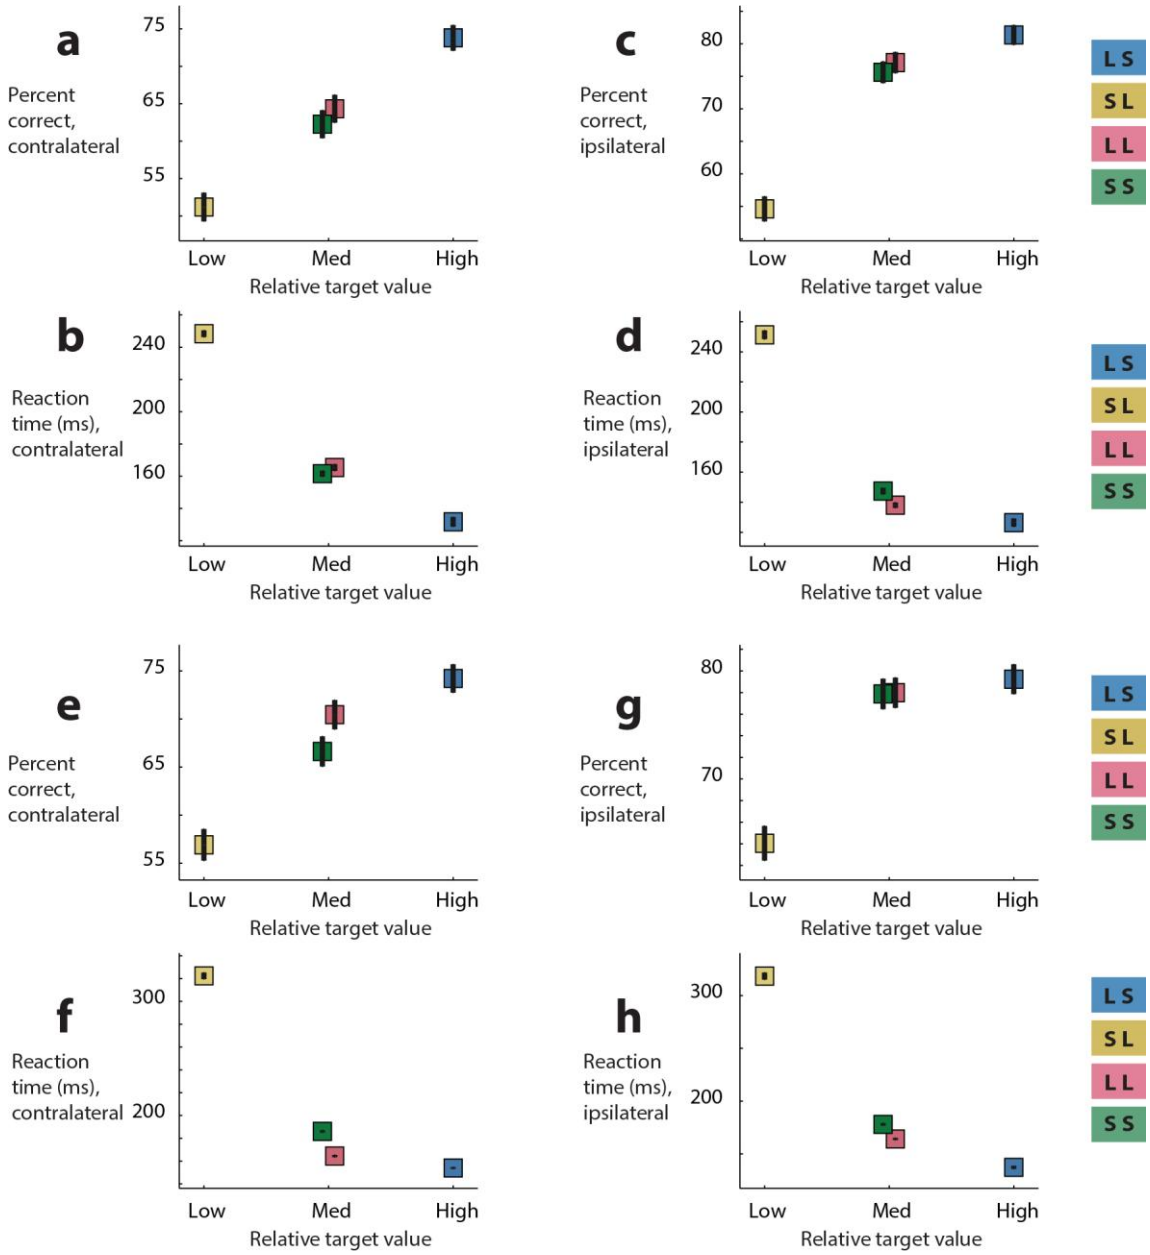

Supplementary Figure 3

Behavioral effects were consistent across locations.

Behavioral data are presented as function of relative target value, from all trials completed during recording sessions in Monkey 1 (**a-d**) and Monkey 2 (**e-h**). Behavioral data are shown separately for targets contralateral (**a,b,e,f**) and ipsilateral (**c,d,g,h**) to the recorded brain area. Error bars indicate +/- 1 s.e.m. Overlapping points are slightly offset for visualization. **(a,c)** Performance tracks relative target value at each location in M1. At each location, performance was highest in LS (CMH, $p<10^{-4}$) and lowest in SL (CMH, $p<10^{-10}$). **(b,d)** Reaction times track relative target value at each location in M1. At each location, reaction times were fastest in LS (WRS, $p<10^{-10}$), and slowest in SL (WRS, $p<10^{-10}$**)** **(e)** Performance tracks relative target value at the contralateral location in M2. Performance was highest in LS (CMH, $p<10^{-3}$) and lowest in SL (CMH, $p<10^{-10}$). (**g**) At the ipsilateral location, performance was highest in LS (79.24%), but not significantly higher than either LL (77.98%; CMH, $p>0.05$) or SS conditions (77.87%; CMH, $p>0.05$). Performance at the ipsilateral location was lowest in SL (64.05%, CMH, $p<10^{-10}$). **(f,h)** Reaction times track relative target value at each location in M2. At each location, reaction times were fastest in LS (WRS, $p<10^{-10}$) and slowest in SL (WRS, $p<10^{-10}$**)**.

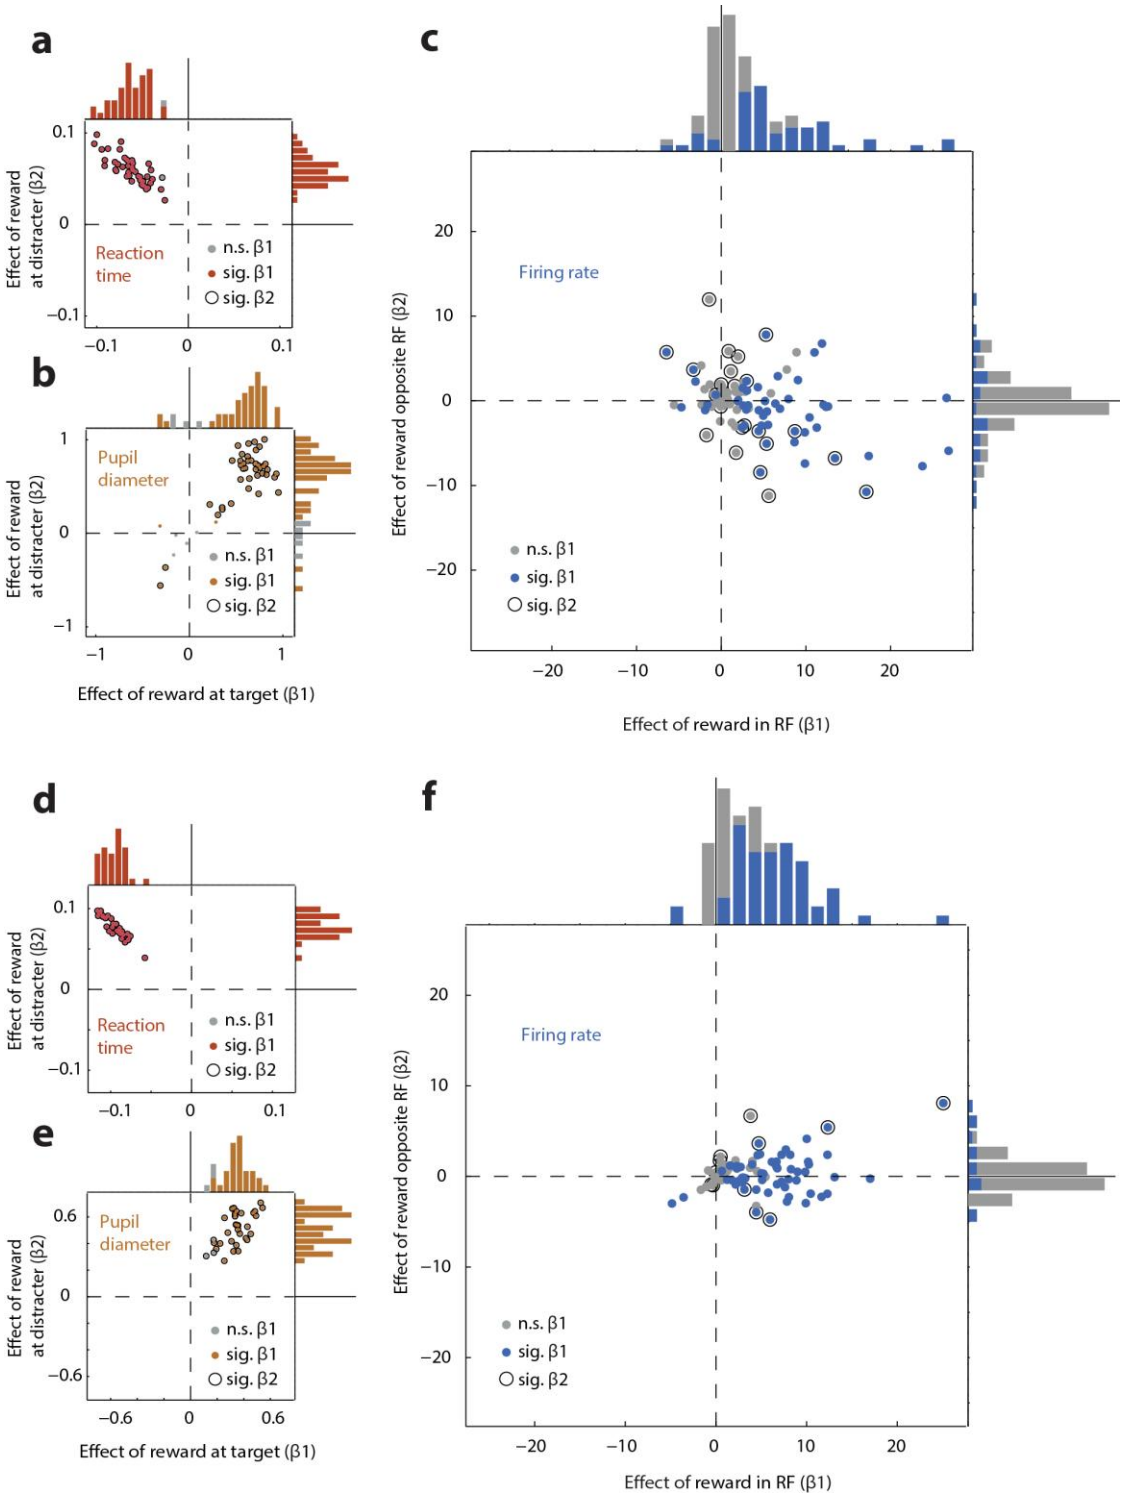

Supplementary Figure 4

Neural and behavioral effects across the population were similar in each monkey

Joint and marginal distributions of regression coefficients for Monkey 1 (**a-c**, n=46 behavioral sessions, n=106 recorded units) and Monkey 2 (**d-f**, n=33 behavioral sessions, n=84 recorded units). Symbol style denotes significance for individual sessions/units. (**a**) Reaction times reflect relative target value in all M1 sessions. Associating the larger reward with the target location decreased reaction

times (WSR, $p<10^{-8}$), whereas large rewards at distracter locations increased reaction times (WSR, $p<10^{-8}$). (**b**) Pupil diameter reflects average trial value in M1 sessions. Associating the large reward with either the target (WSR, $p<10^{-7}$) or distracter (WSR, $p<10^{-7}$) increased pupil diameters. (**c**) Neuronal modulation reflects RF value in M1. Large rewards at the RF location were associated with increased firing rates (mean= +3.55 spikes/s; WSR, $p<10^{-8}$). Large rewards opposite the RF were not associated with a significant change in firing rate (mean= -0.301 spikes/s; WSR, p=0.45). (**d**) Reaction times reflect relative target value in all M2 sessions. Associating the larger reward with the target location decreased reaction times (WSR, $p<10^{-6}$), whereas large rewards at distracter locations increased reaction times (WSR, $p<10^{-6}$). (**e**) Pupil diameter reflects average trial value in all M1 sessions. Associating the large reward with either the target (WSR, $p<10^{-6}$) or distracter (WSR, $p<10^{-6}$) increased pupil diameters. (**f**) Neuronal modulation reflects RF value in M1. Large rewards at the RF location were associated with increased firing rates (mean= +4.69 spikes/s; WSR, $p<10^{-10}$). Large rewards opposite the RF were not associated with a significant change in firing rate (mean= +0.276 spikes/s; WSR, p=0.32).

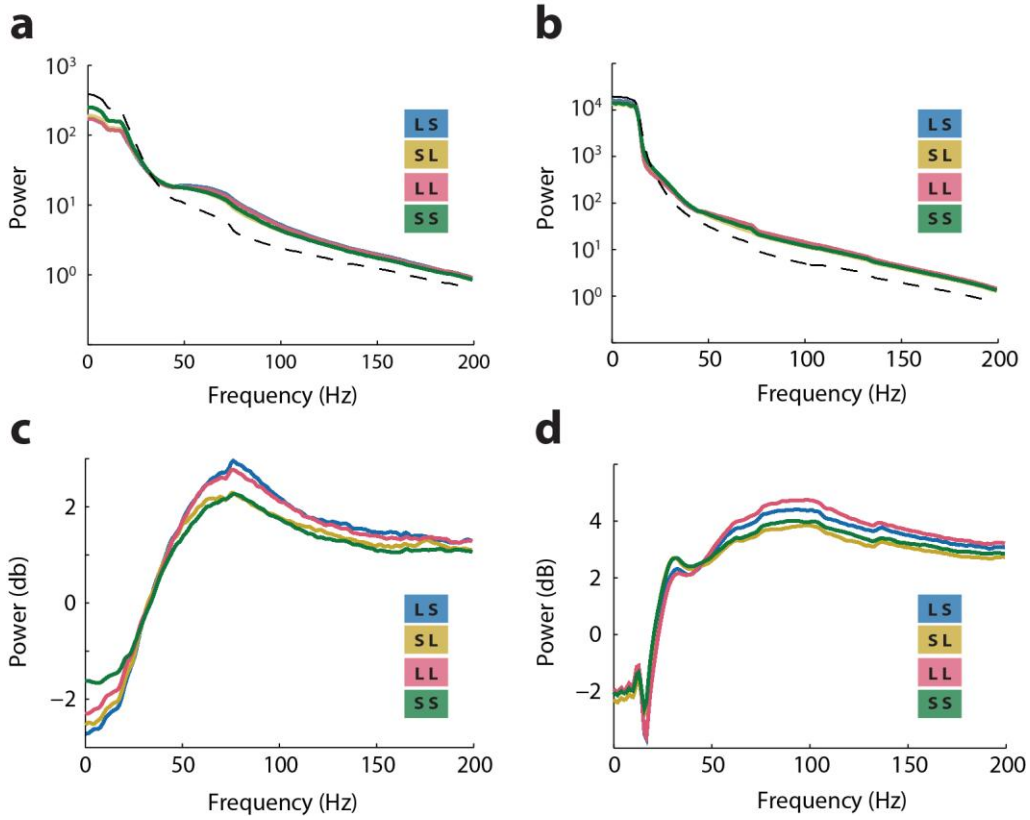

Supplementary Figure 5

Relative RF value and average trial value increase gamma band power in both monkeys.

Raw (**a,b**) and normalized (**c,d**) power spectra for each monkey. Power spectra from the reference epoch used for normalization are indicated by the dotted line (**a,b**). In Monkey 1 (**a,c,** n=40 sites), increases in relative RF value and average trial value both increased power in the gamma band (40-80hz; WSR, $p<10^{-3}$). In Monkey 2 (**b,d,** n=29 sites), increases in relative RF value and average trial value both increased power in the gamma band (40-80hz; WSR, $p<10^{-4}$).

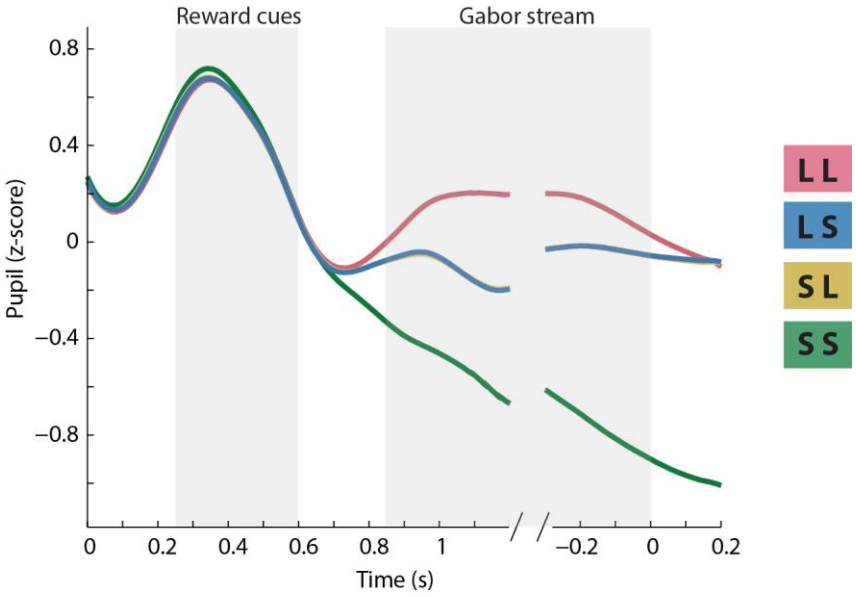

Supplementary Figure 6

Pupil diameters diverge after reward cue presentation.

Z-scored pupil diameters are aligned on two events in the trial. On the left, pupil values are aligned to acquisition of fixation. On the right, pupil values are aligned to onset of discriminanda. In all traces, line thickness exceeds +/- 1 s.e.m. After reward cue presentation (shaded box at left) and through streaming Gabor stimulus (shaded box at right), pupil diameters reflect average trial value. (n= 48,103)

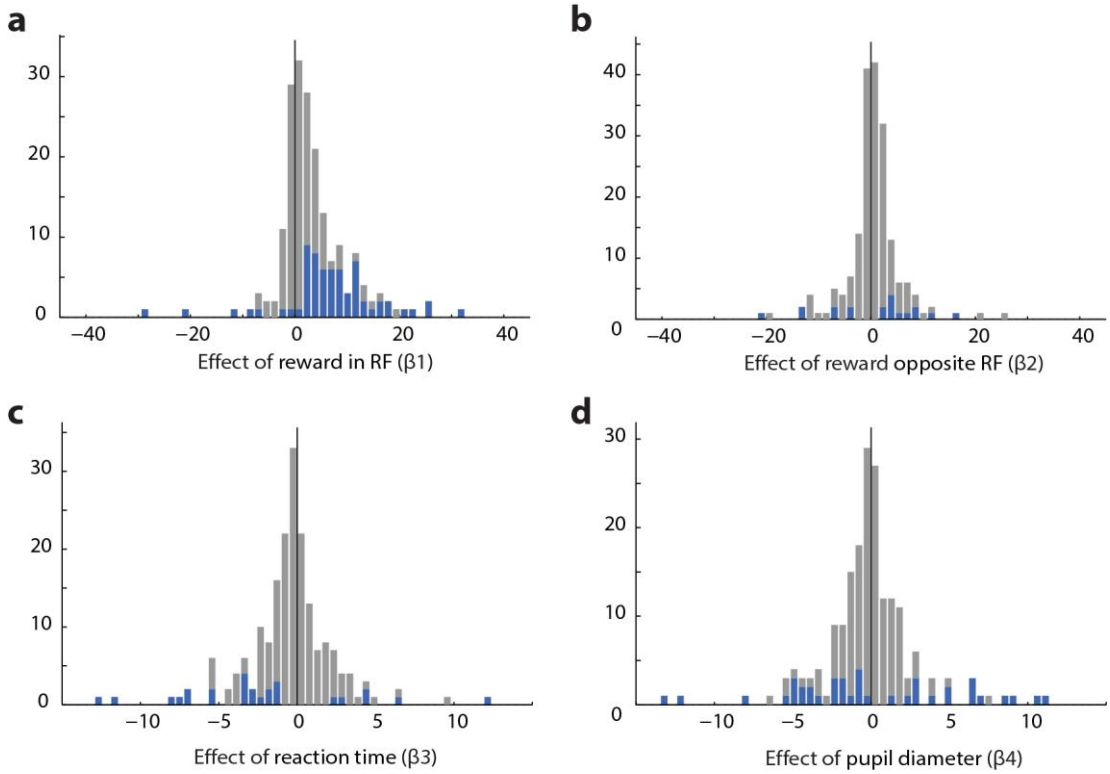

Supplementary Figure 7

Regression coefficients for expanded model.

Symbol style denotes significance for individual sessions/units (n=190). (**a**) Large rewards at the RF location were associated with increased firing rates (mean β_1= +3.678 spikes/sec). (**b**) Large rewards opposite the RF were associated with small increases in firing rate (mean β_2= +0.397 spikes/sec), inconsistent with relative value modulation (**c**) Increased reaction times were associated with small decreases in firing rate (mean β_3= -0.350 spikes/sec). (**d**) Increased pupil diameters were associated with small decreases in firing rate (mean β_4= -0.350 spikes/sec).

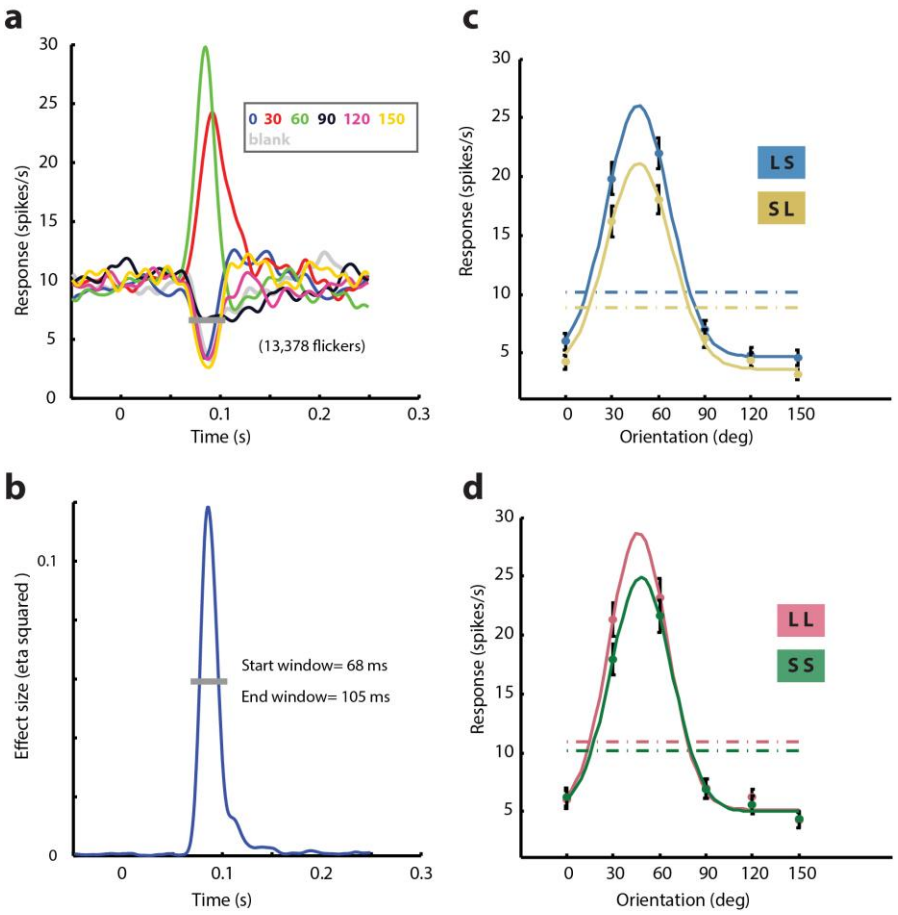

Supplementary Figure 8

Procedure for constructing orientation-tuning functions.

(**a**) Spike density functions (SDF) for an example unit. Responses to individual 20ms Gabor presentations within the stimulus stream were sorted by orientation, aligned on presentation onset, and averaged to generate spike density functions. (**b**) To determine a unit-specific spike count window, we computed the effect size (η^2) of Gabor orientation on firing rate in a sliding window (5 ms wide, 1 ms steps). Effect sizes, as well as correlation of SDFs across time, were used to define the window, denoted by the shaded line. (**c,d**) Spike counts were compiled into orientation tuning functions and fit with Gaussian functions. Error bars indicate +/- 1 s.e.m. Dashed lines indicate mean responses for each condition, averaged across orientation.