Supporting Information

Foley et al. 10.1073/pnas.1613844114

SI Methods

Equations.

Validity and EIG. Explicit choices of information in humans have been shown to be sensitive to the MAP estimate—i.e., favor those cues that maximize the probability that one of the decision alternatives is correct (1)—and sampling policies based on MAP may also be implemented in oculomotor control (8). In mathematical notation,

$$MAP = \max(i)P, \quad \text{[S1]}$$

where i is the rewarded option and indexes the possible targets. In the current task, validity is defined in terms of the probability of signaling the correct target, and is precisely equal to MAP.

A second common measure of information gains is based on Shannon information, defined as

$$\text{SEIG} = H(0.50) - H(\text{Val}), \quad \text{[S2]}$$

where entropy (H) is given by

$$H(p) = -[p \cdot Log_2(p) + (1-p) \cdot Log_2(1-p)].$$

At the outset of each trial, the probability that a target location is correct is $P = 0.5$, and the decision entropy, $H(0.50)$, is 1 bit. After sampling motion direction information, decision entropy is reduced by $H(\text{Val})$, where Val is cue validity. The SEIG is the difference between the prior and posterior entropy, and equals 1.0, 0.278, and 0.007 bits for, respectively, Val = 1.0, 0.8, and 0.55. Thus, validity is monotonically related to SEIG.

Regressions. All regressors were standardized to a range of 0 to 1. For saccade velocity, latency, accuracy, and fixation duration, we converted the raw values into percentiles from 0 to 1. In two-choice conditions, validity has a lower bound at 0.5; therefore we stretched the full validity range (0.5 to 1.0) to fit into the range of 0 to 1 (or, equivalently, plotted regression coefficients as if the line intercepted the y axes at $x = 0.5$ rather than $x = 0$). Thus, the fitted coefficients capture the firing rate modulations across the full range of regressors that were used in the task.

All of the regressions were fit to the time resolved firing rates (FR), measured in a sliding window of 50 ms width with a 1-ms step. Thus, each beta coefficient is associated with a waveform across time. The β_0 coefficient captured the time course of the average response, including the baseline firing, visual onset, and presaccadic responses that were common across all conditions.

To estimate the effects of validity independent of saccade parameters for Fig. 2*B*, we used the equation

$$\begin{aligned}\text{FR} = {} & \beta_0 + \beta_1 \cdot \text{inRFValidity} + \beta_2 \cdot \text{oppRFValidity} + \beta_3 \cdot \text{SaccDir} \\ & + \beta_4 \cdot \text{SaccLatency} + \beta_5 \cdot \textit{SaccVelocity} \\ & + \beta_6 \cdot \text{SaccAccuracy} + \varepsilon,\end{aligned} \quad \text{[S3]}$$

where inRFValidity is the validity of the RF cue, oppRFValidity is the validity of the opposite RF cue, SaccLatency is saccade latency, SaccVelocity is peak saccade velocity, and SaccAccuracy is saccade angular accuracy.

To compare responses to informative cues and uninformative stimuli (Fig. 4 *B* and *C*), we fit firing rates in each condition to the equation

$$\text{FR} = \beta_0 + \beta_1 \cdot C + \varepsilon, \quad \text{[S4]}$$

where C codes the validity of the informative RF cue in informative trials and the reward probability of the yoked uninformative stimulus in uninformative trials. The fitted coefficients capture the visual response (β_0) and the validity/reward modulations (β_1).

To measure the impact of local reward history (Fig. 5*B*), we fit firing rates in each condition to the equation

$$\text{FR}_n = \beta_0 + \beta_1 \cdot \text{Val}_n + \beta_2 \cdot \text{Rew}_{n-1} + \varepsilon. \quad \text{[S5]}$$

In this equation, Rew_{n-1} is coded as 1 if the previous trial received a reward and 0 otherwise, regardless of the validity of the cue.

To test for possible cue-specific effects and for effects across multiple trials, we used the additional equation

$$\begin{aligned}\text{FR}_n = {} & \beta_0 + \beta_1 \cdot \text{Val}_n + \beta_2 \cdot \text{ValMatch} \cdot \text{Rew}_{n-1} + \beta_3 \cdot \text{ValMatch} \\ & \cdot \text{Rew}_{n-2} + \beta_4 \cdot \text{ValMatch} \cdot \text{Rew}_{n-3} + \varepsilon.\end{aligned} \quad \text{[S6]}$$

Here ValMatch was set to 1 if the validity of the cue on the prior trial matched that on the current trial, and 0 otherwise. Therefore, the equation models a scenario in which monkeys update reward values based on the reward history of individual cues. We computed the coefficients separately for each of the three validity levels. This resulted in nine reward coefficients, none of which were significantly above 0 during the delay epoch.

To determine the effects of EIG and fixation duration (Fig. 5*D*), we fit the time-resolved firing rates with

$$\text{FR} = \beta_0 + \beta_1 \cdot \text{Val} + \beta_2 \cdot \text{VT} + \varepsilon, \quad \text{[S7]}$$

where VT is the stamdardized duration (percentile) of the postsaccadic fixation (viewing time, Fig. 5 *C* and *D*).

To determine the effects of change in reward (Fig. 7 *B* and *C*), we fit firing with

$$\text{FR} = \beta_0 + \beta_1 \cdot \Delta\text{Val} + \varepsilon, \quad \text{[S8]}$$

where ΔVal codes the change in reward associated with a cue. In the informative condition, ΔVal had values of +0.22, +0.02, and −0.23 for, respectively, the 100%, 80%, and 55% informative cues. In the cue-change condition, the regressor took values of −0.20, 0.00, and +0.25 for, respectively, the 100%, 80%, and 55% first cues.

RL Simulations. RL simulations were performed to compute the values and RPEs that a reward maximizing process would attribute to each state. We modeled each setting using a Markov decision process (MDP) that includes the task states and transitions, and allowed the model to learn state values based on the final rewards using Q-learning, a standard trial and error learning algorithm. Q-learning is a well-known model-free RL technique that learns an action value function specifying the expected utility of taking a given action in a given state and following the optimal policy thereafter, and can be used to find an optimal action selection policy for any finite MDP (26).

To model the informative condition, the MDP had a cue state with two possible values corresponding to 55% and 80% validity cues, an action state with two values corresponding to the decision

alternatives, and a reward state with two values corresponding to reward or no reward. The MDP for the uninformative condition was similar except that an uninformative state was interposed between the cue and action states. To model the RPE for informative cues (Fig. 6*B*), we used a model that was equivalent to that for the informative cues except that it had an "initial" state preceding the cues. To model RPE, we used an initial cue state that had a 55%, 80%, or 100% valid cue and could transition to the same or a different validity at the next step. Note that we only modeled changes in validity and ignored mere changes in location. Thus, the probability of transitioning to the same validity was 87.5% for the 55% and 100% initial cues and 100% for the 80% cues [based on the probability of change trials (in which only one cue appeared, 75% of total) and "jump to self" trials in which the second cue signaled the same validity as the first].

At each time instant, the Q values are computed as follows: $Q(x,a) = Q(x,a) + a(r + \gamma\ max\ Q(x,b) - Q(x,a))$, where $r + \gamma\ max\ Q(x,b) - Q(x,a)$ is the standard temporal difference error or RPE. For each simulation, we computed the optimal values using an exact value iteration method. In each instance (Figs. 3*C* and 6*B*), we analyzed the z-scored Q values of the cue states computed over 1,000 simulations at the point when learning converged (with convergence evaluated by comparison with the optimal values).

Note that, because we did not aim to replicate the precise response magnitudes shown by the cells, we arbitrarily set the simulation parameters in a way that would approximate these magnitudes (learning rate $\alpha = 0.25$, discount factor $\gamma = 0.9$, and an ε-greedy policy of 50%) and focused on the model predictions about the relative values of the different cues and across the informative and uninformative conditions.

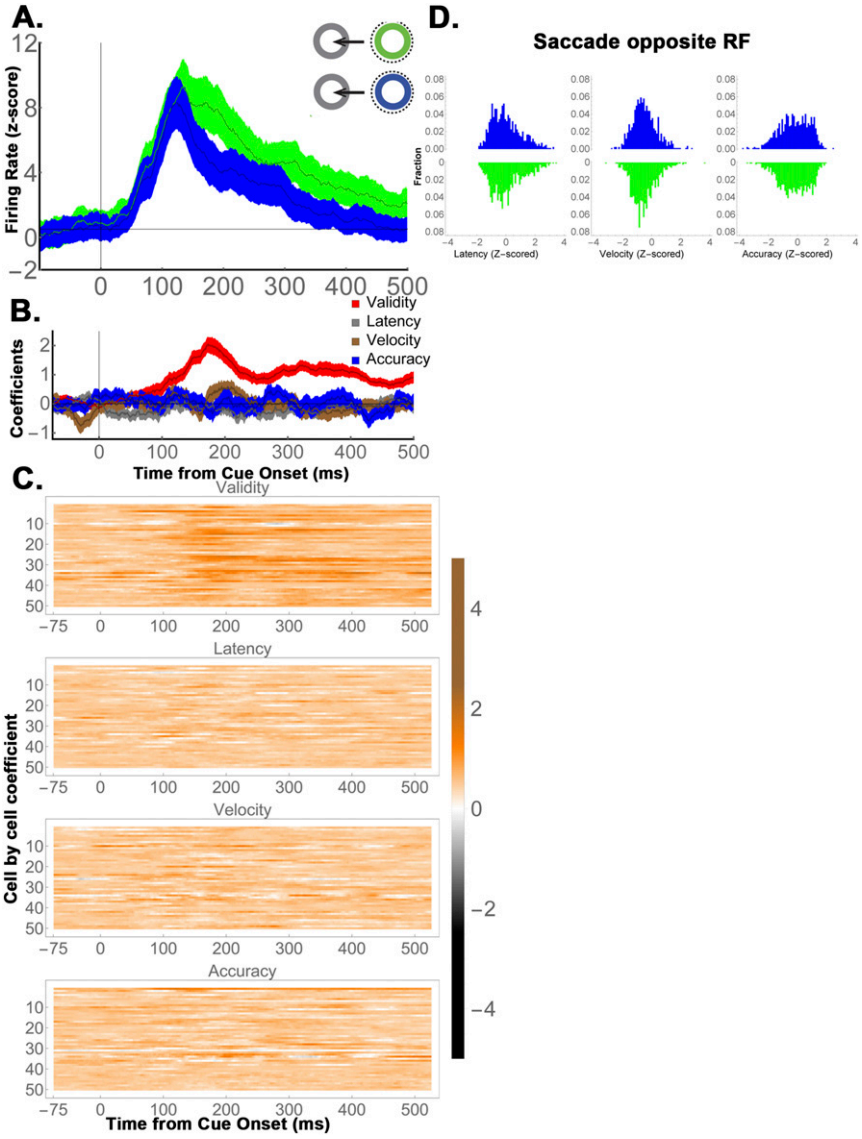

Fig. S1. Validity effects in subsets of trials equated for saccade direction and the validity of the opposite cue (refers to Fig. 2). To further verify the independence of EIG responses from saccade planning, we selected subsets of two-cue trials that differed only in the validity of the RF cue (55% vs. 80%) but were identical in terms of the saccade direction (away from the RF) and the validity of the opposite-RF cue (100%). (*A*) Average population responses for the two trial types (mean and SEM across 50 cells) show that the cells responded more for 80% relative to 55% RF cues. (*B*) Regression analysis (identical to Eq. **S3** except for the absence of a term for saccade direction and validity of the opposite cue, which did not vary in these trials) confirmed that the cells did not modulate as a function of saccade latency, velocity, or accuracy, but had a significant modulation with the validity of the RF cue ($P < 10^{-10}$ relative to 0 during the 125 to 250 ms after cue onset). (*C*) The results were consistent in individual cells. The time-resolved regression coefficients for each parameter are shown for individual cells. (*D*) Saccade latencies, velocities, and accuracies did not differ as a function of cue validity. Trials were pooled across all recording sessions after *z*-scoring within individual sessions. The EIG of the RF cue (blue vs. green) had negligible effects on saccade metrics. Paired comparisons using the Mann–Whitney *u* test yielded $P = 0.38$ for saccade latency, $P = 0.15$ for velocity, and $P = 0.86$ for accuracy. Note that, by *z*-scoring and pooling, we minimized intersession noise and maximized trial numbers ($n = 1{,}676$ trials), resulting in high statistical power and very low likelihood of a type II error.

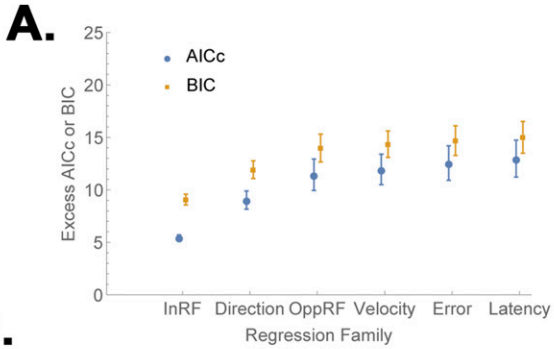

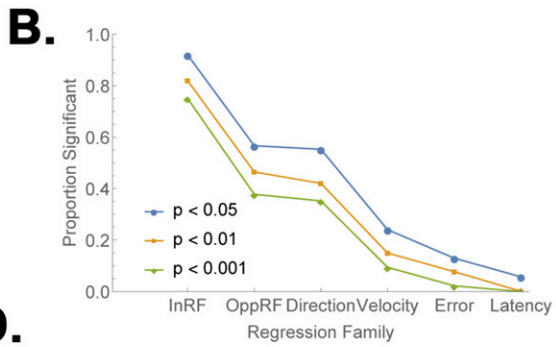

C.

AICc p value	InRF	Direction	OppRF	Velocity	Error	Latency
InRF		0.0783	0.0158	0.0112	0.0073	0.0055
Direction	0.0783		0.1590	0.1173	0.0801	0.0609
OppRF	0.0158	0.1590		0.4127	0.3154	0.2553
Velocity	0.0112	0.1173	0.4127		0.3959	0.3278
Error	0.0073	0.0801	0.3154	0.3959		0.4267
Latency	0.0055	0.0609	0.2553	0.3278	0.4267	

D.

BIC p value	InRF	Direction	OppRF	Velocity	Error	Latency
InRF		0.0548	0.0073	0.0051	0.0036	0.0026
Direction	0.0548		0.1124	0.0810	0.0590	0.0435
OppRF	0.0073	0.1124		0.4103	0.3311	0.2642
Velocity	0.0051	0.0810	0.4103		0.4157	0.3404
Error	0.0036	0.0590	0.3311	0.4157		0.4205
Latency	0.0026	0.0435	0.2642	0.3404	0.4205	

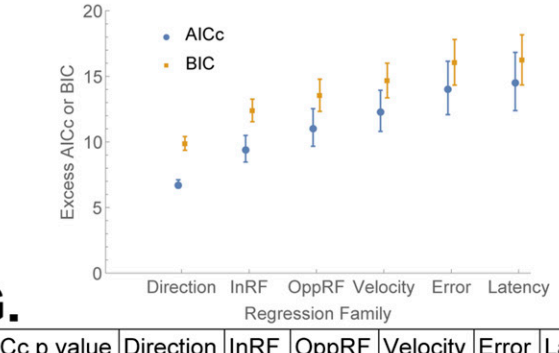

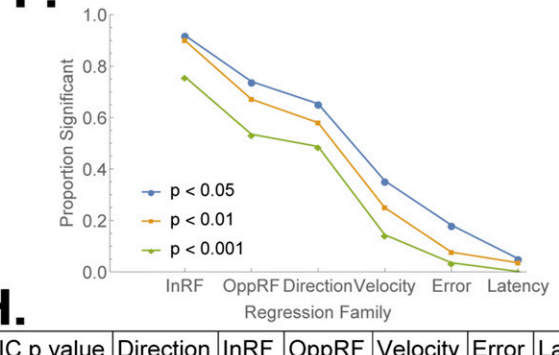

G.

AICc p value	Direction	InRF	OppRF	Velocity	Error	Latency
Direction		0.1345	0.0481	0.0206	0.0062	0.0044
InRF	0.1345		0.2452	0.1190	0.0388	0.0279
OppRF	0.0481	0.2452		0.2937	0.1104	0.0813
Velocity	0.0206	0.1190	0.2937		0.2299	0.1755
Error	0.0062	0.0388	0.1104	0.2299		0.4162
Latency	0.0044	0.0279	0.0813	0.1755	0.4162	

H.

BIC p value	Direction	InRF	OppRF	Velocity	Error	Latency
Direction		0.0752	0.0248	0.0082	0.0021	0.0017
InRF	0.0752		0.2385	0.0922	0.0248	0.0206
OppRF	0.0248	0.2385		0.2450	0.0750	0.0630
Velocity	0.0082	0.0922	0.2450		0.1998	0.1717
Error	0.0021	0.0248	0.0750	0.1998		0.4536
Latency	0.0017	0.0206	0.0630	0.1717	0.4536	

Fig. S2. Model comparison (refers to Fig. 2*B*). For each cell, we computed 63 regression models (corresponding to all possible combinations of the six regressors used in Fig. 2*B*) in the (*A*–*D*) early delay period (100 md to 300 ms) and (*E*–*H*) late delay period (300 ms to 500 ms) and evaluated each model using sample corrected Akaike Information Criterion (AICc) and the Bayesian information Criterion (BIC). AICc and BIC measure the goodness of fit of a model using different philosophies. AICc assumes that reality is much higher dimension than any model, and is thus biased toward overfitting the data (allowing too many free parameters), BIC, in contrast, assumes that one of the available models is a "true" best fit, and is relatively biased toward underfitting the data (rejecting an additional parameter that has explanatory power). Because it is not a priori clear which bias is more appropriate, we chose a standard strategy of computing both metrics and testing for disagreements between them. The two tests agreed in assigning the best performance (lowest scores) to models that included a term for the validity of the RF cue. (*A* and *E*) Mean and SEM of AICc and BIC values for subgroups of models that included each one of the six regressors (across with all possible combinations of the remaining regressors) arranged in order of model quality (lower values are better). AICc and BIC values were normalized for each cell by subtracting the cell's best model, and the points show the mean and SEM of the normalized values ("Excess") across cells and models in a given class. (*B* and *F*) The fraction of models (pooled across cells) for which the coefficient of each regressor was significant at $P < 0.05$, 0.01, and 0.001. (*C* and *D*) Comparison matrices of the AICc and BIC values in *A*, showing the likelihood that the differences between any pair of values arose by chance. Red indicates values that are significantly different at $P < 0.05$. (*G* and *H*) Same as *C* and *D*, for the data in *E*, analyzing 300 ms to 500 ms after cue onset. In the early delay period, models that included a term for the validity of the RF cue outperformed all other models, including those that included terms for saccade direction and the validity of the opposite cue, whether evaluated with the AICc or the BIC (*A*–*D*). In the late delay period, models that included the validity of the RF cue were only slightly worse than those including saccade direction (*E*–*H*). In both epochs, over 90% of the coefficients for the validity of the RF cue were significant at $P < 0.05$ (*B* and *F*). For all of the analyses, equivalent results were obtained when we reanalyzed the data using median rather than mean AICc and BIC values.

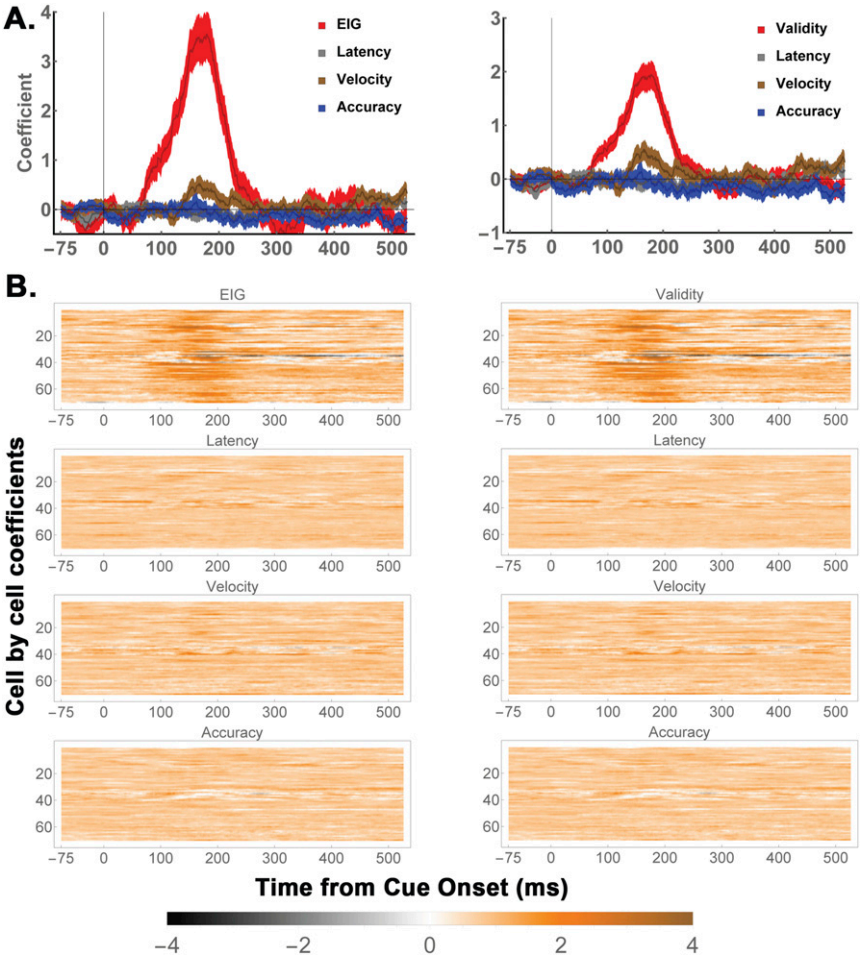

Fig. S3. Neurons encode EIG and validity but not saccade metrics in the one-cue task (refers to Fig. 4). The same analysis as in Fig. S1, applied to one-cue informative trials in which a single informative cue was in the RF (same cells that are analyzed in Fig. 4). The regressions were coded in terms of (*Left*) SEIG (Eq. **S2**) and (*Right*) validity. (*A*) The average coefficient value (mean and SEM across cells) and (*B*) the results for individual cells. EIG/validity modulations were robust in a majority of cells, but no neuron showed effects of saccade metrics.

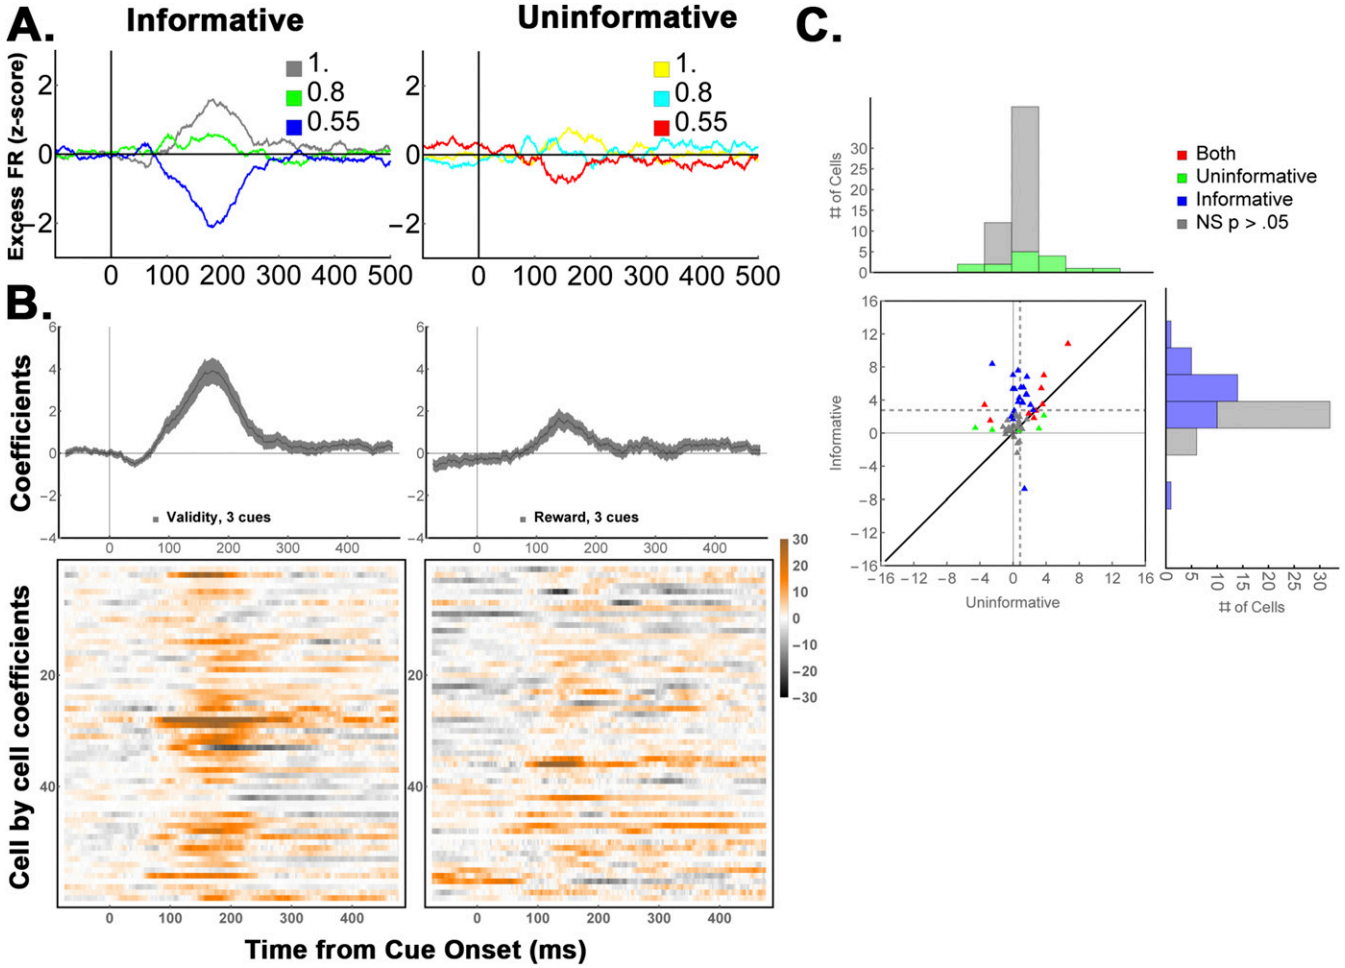

Fig. S4. Neurons modulate for informative but not uninformative cues when using three levels of EIG (refers to Fig. 4). The same analysis as in Fig. 4, applied to the data from M1, who was tested with all three informative cues and corresponding uninformative cues. The average coefficients for informative and uninformative cues were, respectively, 2.75 ± 0.48 vs. 0.85 ± 0.27 *z*-score units/validity, $P = 0.0001$, with 53% vs. 23% of cells showing significant modulations, $z = 3.37$, and $P = 0.0007$ for test of proportions. (*A*–*C*) Same conventions as Fig. 4. (*A*) Average firing rates. (*B*) Average regression coefficients. (*C*) Cell-by-cell regression coefficients in 125–250-ms window.

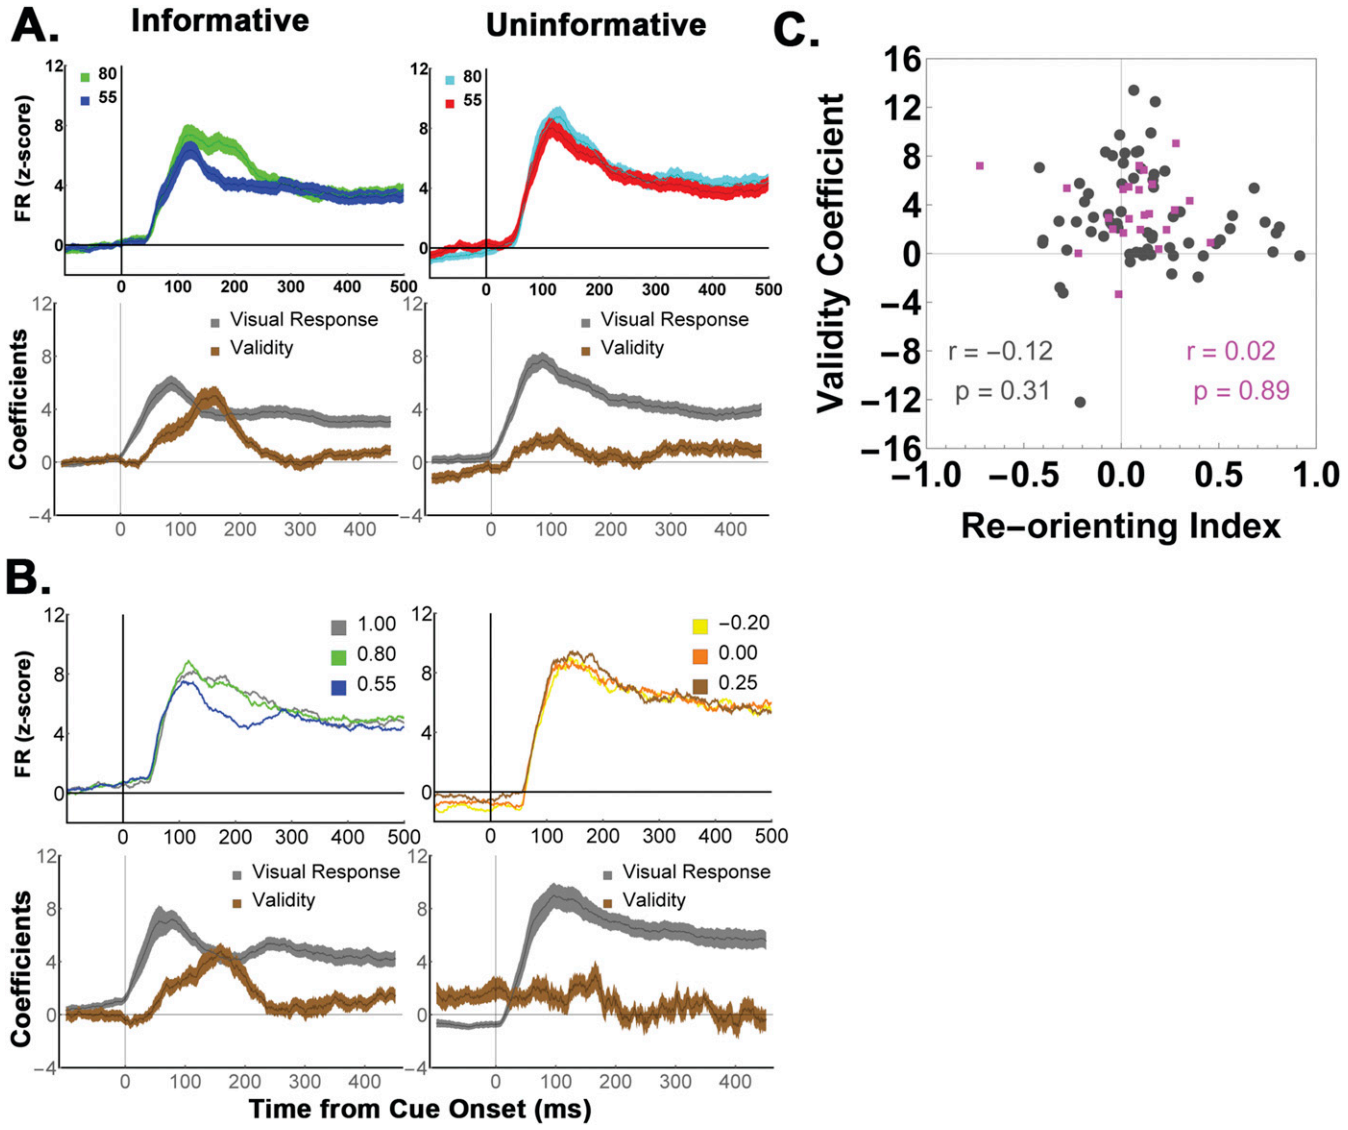

Fig. S5. Full responses to the uninformative and cue-change tasks (refers to Figs. 4 and 7). (*A*) (*Top*) The average population responses to the informative and uninformative cues that are compared in Fig. 4, before mean subtraction. The neurons had a visual response to the onset of both cue types, and the validity modulations (differences between the two cue types) rode on top of this response. This finding is the expected one in LIP neurons, which multiplex visual, saccade-related, and decision activity in their firing rates (14) and implies that validity is not encoded in terms of absolute firing rates but must be read out using more complex decoding mechanisms (34). (*Bottom*) Regression coefficients associated with the visual response (β_0 in Eq. **S4**, gray) and the validity/reward modulations (β_1 in Eq. **S4**, brown). (*B*) (*Top*) Same as *A*, but showing the responses to informative cues vs. cue-change trials that are analyzed in Fig. 7. (*Bottom*) Regression coefficients associated with the visual response (β_0 in Eq. **S8**, gray) and the validity/reward modulations (β_1 in Eq. **S8**, brown). (*C*) Reorienting effects cannot explain the validity encoding. Comparison of the visual onset responses shows that the overall visual responses were slightly higher for an uninformative stimulus relative to an informative cue in the RF (gray traces in *Top*, *Right* vs. *Left* in *A*), and also slightly higher for cue-change relative to informative trials (gray traces in *Top*, *Right* vs. *Left* in *B*), reflecting the reorienting of attention from the opposite-RF precue to the new stimulus in the RF (27). To determine whether this effect may have affected validity coding, we quantified the visual enhancement for each cell using a reorienting index (RI), $RI = (FR_{2nd} - FR_{1st})/(FR_{2nd} + FR_{1st})$, where FR_{first} is the *z*-scored visual response to an RF stimulus when that stimulus was the first one that appeared on a trial (i.e., the informative cue in the one-cue task). FR_{second} is the visual response to the RF stimulus when that was the second stimulus in the task (i.e., the uninformative stimulus in the uninformative task, or the second cue in the cue-change task). The RI was computed in the epoch 100 ms to 250 ms after cue onset to capture the peak enhancement. The scatterplot shows the cell-by-cell RI and validity modulation for the cells tested in the one-cue task (black points; $n = 69$; Fig. 4) and for those tested on the cue-change task (magenta, $n = 24$, described in Fig. 7). RIs were, on average, positive and showed a marginally significant trend for visual enhancement (average RI = 0.09, $P = 0.069$). Critically, there was no correlation with encoding of validity in either the uninformative or the cue-change test.

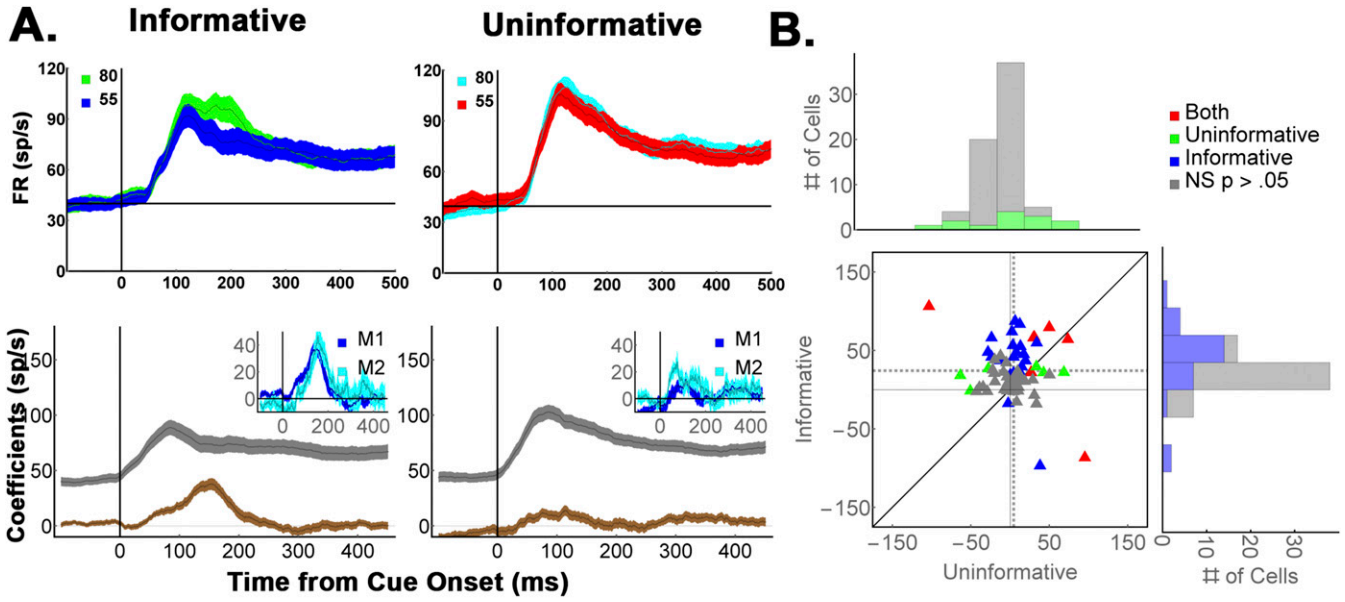

Fig. S6. Validity encoding is robust in nonnormalized firing rates (refers to Fig. 4). The same analysis as in Fig. 4, applied to raw firing rates before *z*-scoring or mean subtraction. Validity effects remain robust, with average regression coefficients of 24.15 ± 0.57 sp/s for informative cues vs. 4.68 ± 1.2 sp/s for uninformative stimuli ($P = 0.000008$) and 42% vs. 18% of cells showing significant modulations ($z = 2.96$, $P = 0.001$ for test of proportions). Note that, in this analysis, the β_0 coefficient (gray traces) tracks the baseline rate and raw visual response and hence is tonically elevated even before cue onset, whereas the validity β_1 coefficient tracks the relative modulation by cue type and is 0 before cue onset. (*A*) Average firing rates and average regression coefficients. (*B*) Cell-by-cell regression coefficients in 125–250-ms window.