1 **Transient neuronal suppression for exploitation**

2 **of new sensory evidence**

3 Maxwell Shinn^{1,2}, Daeyeol Lee^{3,4,5,6}, John D. Murray^{1,2,7,8*}, Hyojung Seo^{1,2,8*}

4

5 ¹Interdepartmental Neuroscience Program, Yale University

6 ²Department of Psychiatry, Yale University

7 ³Zanvyl Krieger Mind/Brain Institute, Johns Hopkins University

8 ⁴Kavli Discovery Neuroscience Institute, Johns Hopkins University

9 ⁵Department of Psychological and Brain Sciences, Johns Hopkins University

10 ⁶Department of Neuroscience, Johns Hopkins University

11 ⁷Department of Physics, Yale University

12 ⁸Department of Neuroscience, Yale University

13

14 Correspondence to: john.murray@yale.edu, hyojung.seo@yale.edu

15

16 **Abstract**

17 In noisy but stationary environments, decisions should be based on the temporal
 18 integration of sequentially sampled evidence. This strategy has been supported by many
 19 behavioral studies and is qualitatively consistent with neural activity in multiple brain
 20 areas. By contrast, decision-making in the face of non-stationary sensory evidence
 21 remains poorly understood. Here, we trained monkeys to identify the dominant color of a
 22 dynamically refreshed bicolor patch that becomes informative after a variable delay.
 23 Animals' behavioral responses were briefly suppressed after evidence changes, and many
 24 neurons in the frontal eye field displayed a corresponding dip in activity at this time,
 25 similar to that frequently observed after stimulus onset. Generalized drift-diffusion
 26 models revealed consistency of behavior and neural activity with brief suppression of
 27 motor output, but not with pausing or resetting of evidence accumulation. These results
 28 suggest that momentary arrest of motor preparation is an important component of
 29 dynamic perceptual decision making.

30 **Introduction**

31 Momentary evidence from sensory stimuli or memory seldom provides sufficient
 32 information to choose the most appropriate action. Rather, speed and accuracy of choice
 33 behaviors in humans and animals are more consistent with models of integrators or
 34 accumulators, such as the drift-diffusion model (DDM), in which noisy evidence is
 35 temporally integrated as a decision variable that triggers an action upon crossing a
 36 threshold (Ratcliff 1978; Gold and Shadlen 2007). In addition, neural activity in multiple

brain areas might correspond to the trajectory of such decision variables (Hanes and Schall 1996; Zoltowski et al 2019). Yet, how this relatively simple strategy can be extended for time-varying sensory stimuli (Glaze et al., 2015) remains poorly understood. For example, in most previous neurophysiological studies, stimulus evidence had constant magnitude within a single trial and evidence onset was clearly demarcated. By contrast, in the present study, we investigated how the trajectory of neural activity related to evidence accumulation might be adjusted by the subtle onset of sensory signals decoupled from the onset of the stimulus itself.

We examined three different hypotheses regarding cognitive strategies for how the detection of informative stimuli might impact perceptual decision making. First, the neural activity related to the decision variable, such as that observed in lateral intraparietal cortex (LIP) and frontal eye field (FEF), shows a temporary dip after stimulus onset (Sato and Schall 2001; Roitman and Shadlen 2002; Kiani et al., 2008; Ding and Gold 2012; Hanks et al., 2014; Teichert et al. 2016), and this has been interpreted as the reset of an integrator for the decision variable. According to this “reset” model, the decision variable and its neural correlate might be fully or partially reset when the non-informative sensory stimulus is replaced by an informative stimulus. Second, rather than a reset, the decision variable might be temporally frozen so that the incoming stream of evidence during a volatile transition period can be ignored. Indeed, such a “pause” model has successfully accounted for the reaction time (RT) data during behavioral tasks, such as countermanding or double-step saccade tasks, in which the subjects must adapt to sudden and unpredictable changes in instructions (Reingold and Stampe 2002; Åkerfelt et al. 2005; Boucher et al. 2007; Logan et al. 2015; Salinas and Stanford 2018). Finally, decision makers might adapt to the unpredicted arrival of new

information simply by suppressing their motor outputs (Purcell et al. 2010; Bompas and Sumner 2011; Bompas et al. 2015; Wessel et al. 2016) without modifying the state of decision variables. Unlike the reset or pause models, this “motor suppression” model predicts that even motor outputs unrelated to the task, such as microsaccades, might be suppressed.

We tested these three alternative hypotheses by analyzing the behavioral data and neural activity recorded from the FEF in monkeys performing a perceptual decision-making task in which the stimulus onset was temporally decoupled from the onset of informative stimulus evidence. We observed temporary reductions in the RT distribution and in FEF activity after the onset of informative stimulus. Moreover, motor outputs and FEF activity were suppressed more by stronger sensory signals. Formal model testing showed that both behavioral and neural data were most consistent with the motor suppression model compared to the reset and pause models. These results suggest behavioral and neural signatures of motor suppression as a cognitive mechanism for the strategic use of changes in evidence during perceptual decision-making.

Results

Immediate effect of a change in evidence on the RT distribution

We trained two rhesus monkeys to perform a two-alternative forced-choice color-discrimination task (Figure 1). The stimulus for discrimination was a square patch consisting of blue and green pixels, and the relative number of pixels between the two colors, referred to as color coherence, determined the difficulty of discrimination. To

temporally dissociate the processes associated with the detection of a change in the evidence from those for the detection of the onset of the stimulus itself, stimulus presentation was divided into two consecutive periods containing an uninformative "presample" and informative "sample".

During the variable presample period (0, 400, or 800 ms), there were equal numbers of blue and green pixels displayed in the stimulus, corresponding to a color coherence of zero. During the sample period, color coherence of the stimulus changed to a non-zero value, which was fixed for a single trial but randomly selected from three values across trials. No explicit cue was presented to indicate the abrupt transition from presample to sample, and pixels were rearranged at 20 Hz during both periods to make this transition non-obvious. Reward cues surrounded the saccade targets and indicated whether a correct response to the designated target would result in a large or small reward (see Online Methods). Only correct choices made after the sample onset were rewarded. The animals' performance in this task changed with color coherence and presample duration. As reported in detail previously (Shinn et al., 2020a), choices were less accurate (Figure 1c, ED1a) and slower (Figure 1d, ED1b) during trials with a low coherence.

Figure 1: The color-discrimination task. (a) The temporal sequence of trial events in the color-discrimination task. Inset in the lower left are cues which indicate a large or small reward. (b) Schematics for the time course of color coherence throughout the trial for each presample duration. Gray indicates zero coherence, and colors indicate non-zero coherence. The thin gray horizontal lines denote the fixation period. (c) Psychometric function showing the mean probability of a correct response for Monkey 1. Error bars representing standard error are hidden beneath the markers. (d) Chronometric function showing the mean RT as the function of coherence for Monkey 1. Error bars representing standard error are hidden beneath the markers.

To gain a mechanistic understanding of how the monkeys might be performing the task, we fit a generalized drift-diffusion model (GDDM) (Shinn et al 2020b) to RT distribution data, a model previously shown to explain both the reward bias and timing in our dataset (Shinn et al 2020a). This model incorporated leaky integration, an urgency signal, and two forms of reward bias (Shinn et al. 2020a). As reported previously (Shinn et al 2020a), the GDDM predicts higher accuracy and shorter RT in trials with higher color coherence, and this was confirmed in the data (Figure 1c,d; ED1a,b).

116 Despite the general success of the GDDM in accounting for the complex
117 behavioral patterns observed in the RT and choice data, we found a behavioral feature
118 immediately after an abrupt change in evidence which cannot be explained by the model.
119 To examine the effect of changing evidence on the most rapid responses, we focused on
120 the RT in the longest presample condition, because in trials with shorter presample
121 durations, there were few responses within a 200-ms window immediately following the
122 change in evidence (Figure 2a, ED2a-c). We found that, despite an increase in evidence, a
123 change to a higher coherence resulted in a short-latency suppression of responses, visible
124 as a “dip” in the RT distribution (Figure 2b, ED2d). This RT dip was more pronounced
125 for larger changes in evidence, and its latency was similar to the results reported
126 previously (Reingold and Stampe 2000; Bompas and Sumner 2011; Salinas et al. 2019).
127 The coherence-dependence also implies it is not the result of anticipating the change.
128 This dip in the RT distribution is inconsistent with the standard models of evidence
129 accumulation such as the DDM, as well as the GDDM, which all predict that stronger
130 evidence will shorten RTs without any dip in the RT distribution (Figure 2c).

Figure 2: Transient effect of changes in evidence on RT. (a) The RT distribution for all trials, aligned to the presample onset, for Monkey 1. Highlighted in (b) is the RT distribution centered around the onset of the sample during trials with 800-ms presample. (c) Predictions from a generalized drift-diffusion model (GDDM) for the portion of the RT distribution shown in (b). The black bar indicates significance ($p < 0.05$, one-tailed test for difference of means, bootstrapping across trials). Shaded regions represent bootstrapped 95% confidence interval.

Coherence-dependent dip in FEF population activity

Activity in cortical areas such as FEF and LIP is often hypothesized to represent the accumulation of relative evidence favoring a particular behavioral response. According to this hypothesis, an increase in evidence favoring the target in the response field of a neuron should lead to an increase in the neuron's firing rate. As in the analysis of RT data, we focused on FEF activity during the period immediately following sample onset, and

examined mean-normalized activity averaged across all FEF neurons separately for each presample duration and coherence level.

During the trials with 400- or 800-ms presample duration, we found that FEF neurons showed a robust coherence-dependent reduction in their activity, such that higher coherence resulted in larger reduction in activity (Figure 3c,d, ED2g,h). The latency of this “evidence dip” in FEF activity was comparable to that of the RT dip. Consistent with previous findings, a dip in FEF activity was also seen immediately after the onset of the stimulus itself in trials without a presample period (0-ms presample duration). In comparison to the evidence dip, we refer to this as the “stimulus dip”, because unlike the evidence dip observed after the presample, the magnitude of the stimulus dip was largely independent of the color coherence of the stimulus (Figure 3b and Figure ED2f). Similar to the behavioral results, the presence of the evidence dip is not predicted by models of evidence accumulation. If FEF activity exclusively represents the total accumulated evidence, then it should not decrease when stronger evidence becomes available. Likewise, the mean trajectory of the decision variable simulated with the GDDM did not show a reduction in activity at any point in the first 300 ms after sample onset (Figure 3e-g).

Figure 3: Transient effect of changes in evidence on population FEF activity. (a) The normalized population activity for all trials, aligned to the presample onset, for Monkey 1. (b-d) Highlighted is activity centered around the onset of the sample during trials with a 0- (b), 400- (c), and 800-ms (d) presample duration. The black bar indicates significance ($p < 0.05$, one-tailed test for difference of means, bootstrapping across neurons). Shaded regions represent bootstrapped 95% confidence interval. (e-g) Predictions from a drift-diffusion model decision variable (DV) are shown below for the 0- (e), 400- (f) and 800-ms (g) presample durations. Light gray indicates no prediction within the model's non-decision time.

Coherence-dependent dip in individual FEF neurons

We next examined whether the evidence dip could be detected at the single-neuron level. We first used a linear regression model to determine how single-neuron activity was modulated by the experimental variables such as evidence (color coherence), reward, and choice, at the onsets of the presample, sample, and saccade (Eq. 1). We found that reward magnitude significantly modulated the activity of almost all neurons at all three time points, and that the animal's choice significantly modulated neural activity at the time of saccade onset (Figure ED3). In addition, at the sample onset, but not at the presample or saccade onset, neural activity was significantly modulated by coherence, such that the mean firing rate decreased with coherence, consistent with the dip we described above.

To explore whether this early modulation by coherence after the sample onset took the form of a dip, we fit a time-resolved regression model to the instantaneous firing rate of each neuron (Eq. 2). Our model included multiple kernels aligned to the sample onset, presample onset, and saccade onset. Consistent with the results from the analysis of the mean activity (Figure ED3), we made the presample-aligned kernel sensitive to reward magnitude, the sample-aligned kernel to coherence, and the saccade-aligned kernel to choice (Figure 4a). Here, we focus on the sample-aligned kernel scaled by coherence (referred to as evidence kernel) and the presample-aligned kernel (referred to as stimulus kernel). Thus, the evidence kernel represents the time course of coherence-dependent neural activity after the transition from presample to sample, and the stimulus kernel represents mean neural activity after stimulus onset, corresponding to the evidence dip and stimulus dip respectively, as described for the population activity. The evidence

196 kernel allowed us to examine the effect of coherence immediately after the onset of the
197 sample independently of saccadic activity.

198

199 *Figure 4: Evidence dip in individual FEF neurons for Monkey 1. (a) A schematic of the*

200 *regression model showing alignment of the kernels to the presample, sample, and*

saccade. (b) Smoothed firing rate for two example neurons across conditions. Colors are the same as in Figures 2 and 3. (c) Evidence kernel for the example neurons in (b). (d) For each point in time, the number of neurons in the population with significantly positive (top) or negative (bottom) evidence kernel is shown ($p < 0.05$, one-tailed test). (e) Mean evidence kernel across neurons. Error bars indicate 95% confidence interval of the mean. (f) The first singular vector of the evidence kernels is shown, along with (g) the corresponding factor scores of each neuron.

During the sample period, neural activity reflecting the value of the decision variable should increase with coherence, namely, the evidence kernel should always be non-negative. Contrary to this prediction, we found that for many FEF neurons, activity immediately after sample onset tended to decrease more for larger increases in the coherence of the sample stimulus favoring the action towards the neuron's response field, resulting in a negative evidence kernel. Some neurons show dip-like activity traces (Figure 4b, ED4b), and their evidence kernels were negative between 100 and 200 ms after the sample onset (Figure 4c, ED4c). Across the population, many more neurons show significant negative evidence kernel during the same period. For example, in the interval 150–175 ms after the sample onset, 39% vs. 0% of neurons for Monkey 1, and 52% vs. 9% of neurons for Monkey 2, show significant negative vs. positive evidence kernel, respectively ($p < 0.05$, one-tailed t-test; Figure 4d, ED4e). The fraction of neurons with significantly positive evidence kernel eventually began to increase about 200 ms prior to the saccade onset. Consistently, the average evidence kernel across all neurons also shows a significant dip in this time interval (Figure 4e, ED4d).

To confirm that the mean evidence kernel reflects the temporal pattern of a dip across single neurons, we computed the best rank-one approximation of the kernels, the first singular vector (Figure 4f, ED4f). We found that the temporal profile of weights for the first singular vector is highly correlated with the mean evidence kernel (Monkey 1 $r=0.970$, Monkey 2 $r=0.998$). This singular vector has a predominantly positive factor scores across neurons (Monkey 1, 89% neurons, Monkey 2, 100% neurons; Figure 4g, ED4g), indicating that the dip exhibited by the mean kernel is representative of the kernels of individual neurons. The presence of the dip in the evidence kernel suggests that the evidence dip is coherence-dependent. To confirm this is the case, we performed the above analysis on the coherence-independent sample-aligned kernel and found no evidence of a dip (Figure ED5).

We also tested whether the evidence dip in FEF activity might be confounded by saccade-related activity, using two different types of analyses. First, we extended the regression model to increase the lengths of the saccade and evidence kernels and make the saccade kernel sensitive to coherence (Eq. 3; Figure ED6a). We found that the evidence dip was still present in this extended regression analysis (Figure ED6b,h). This extended model revealed a broad dip of different form and timing than the evidence dip in its coherence-dependent saccade kernel (Figure ED6c,i), which is expected from a broader RT distribution and FEF activity that builds up more slowly during trials with low coherence.

Second, we also tested how the timing of the evidence dip was related to RT. We divided trials into RT quintiles, and determined the timing of the evidence dip for each quintile (see Online Methods). If the evidence dip were saccade-related, there would be a

positive correlation between evidence dip latency and RT quintile. We found that the timing of the evidence dip did not vary significantly with RT (Spearman correlation $r=0.0$, $p=1.0$ for both monkeys; Figure ED6d,f,j,m). We also tested whether the saccade-aligned kernel might show a dip, and examined how the timing of this saccade-related dip might be related to RT. For Monkey 2, the saccade-related kernel did not show a robust dip (Figure ED6k). For monkey 1, we reliably detected a dip in the saccade kernel, and the latency of this dip showed a negative correlation with RT, suggesting that its timing was largely locked at the sample onset. These results demonstrate that the evidence dip is unlikely to reflect pre-saccadic activity.

Coherence-dependent and -independent dips are correlated

Previous studies have reported a dip in neural activity in LIP and FEF when the stimulus, such as randomly-moving dots, is first presented. Indeed, we found a similar dip in FEF activity immediately after the stimulus was first presented regardless of whether it was presample or sample, e.g., in the 0-ms presample condition (Figure 3b, ED2f; Figure 5a, ED8a). We examined whether this stimulus dip might show the same properties, in the same neurons, as the coherence-dependent evidence dip. First, we found that approximately 100–200 ms after presample onset, many FEF neurons exhibited a significantly negative stimulus kernel but very few neurons exhibited positive stimulus kernel ($p<0.05$, one-tailed t-test; Figure 5b, ED8b). For example, during the interval 100–125 ms after stimulus onset, 49% vs 5% of neurons for Monkey 1, and 35% vs 17% of neurons for Monkey 2, showed a significant negative vs positive kernel. There was also a dip in the mean stimulus kernel (Figure 5c, brown). To determine whether individual neurons exhibited a similar dip, we computed the first two singular vectors of the

stimulus kernels. While the first singular vector showed a monotonic ramp (Figure 5d, ED8d, pink), the second singular vector resembled a dip (Figure 5d, ED8d, purple) which had high correlation with the mean stimulus kernel (Monkey 1 $r=0.856$, Monkey 2 $r=0.499$) and predominantly positive factor scores on individual neurons (Monkey 1, 82%, Monkey 2, 83%; Figure 5e, ED8e), demonstrating that individual neurons showed a stimulus dip.

Figure 5: Comparison of evidence and stimulus dips for Monkey 1. (a) Mean FEF activity is plotted for each presample and coherence condition at presample (400 or 800 ms presample) or sample (0 ms presample) onset. (b) The fraction of neurons at each point in time with significantly positive (top) or negative (bottom) stimulus kernel ($p < 0.05$, one-tailed test). (c) Mean stimulus kernel is shown with overlaid evidence kernel. Time is given as time since the sample (evidence kernel) or presample (stimulus kernel). Error bars indicate 95% confidence interval of the mean. (d) The first (SV1) and second (SV2) singular vector of the stimulus kernels are shown, along with (e) the

corresponding factor scores for each neuron on SV2. (f) Evidence dip index is plotted against stimulus dip index. Spearman correlation and p-value is inset.

Next, we compared the stimulus dip to the evidence dip. The timing of the stimulus dip was slightly earlier than the evidence dip (Figure 5c, ED8c), with a significantly earlier minimum of the mean stimulus kernel than the mean evidence kernel (50 ms, $p < 0.0001$ for Monkey 1, 100 ms, $p < 0.0001$ for Monkey 2, bootstrapped two-tailed confidence interval across neurons). We then compared the magnitudes of these two dips by establishing a standardized index to quantify the dip magnitude for individual neurons. The stimulus (evidence) dip index was computed by z-scoring each stimulus (evidence) kernel, and then finding the mean value at the time points in the interval 100 – 150 ms (125 – 175 ms) (see Online Methods). As expected, this index showed a strong correlation with the factor scores on the dip-like singular vectors (Figure ED7). We found that the evidence and stimulus dip indices were significantly correlated across FEF neurons ($p < 0.01$ in both monkeys) (Figure 5e, ED8e), demonstrating that FEF neurons with a strong evidence dip were likely to show a strong stimulus dip. However, we found no consistent association of evidence or stimulus dip index with other neuronal properties, including mean firing rate, directional selectivity index, visuomovement index, or spike width (Figure ED7).

Evidence dip and motor suppression

In order to understand the relationship between the evidence dip observed in FEF activity and the RT dip, we examined three potential mechanisms through which changes in evidence might influence behavior during perceptual decision-making. For all three

mechanisms, it is assumed that a change can be detected quickly, and that the detection event triggers a downstream change in the relatively slower evidence integration pathway. First, in the “pause model” (Figure 6a), when a change is detected, the stream of evidence is briefly interrupted, or “paused”, in order to avoid integrating irrelevant information. Second, in the “reset model” (Figure 6b), the change elicits a partial “reset” of the decision variable back towards its initial value, allowing the animal to ignore the irrelevant integrated signal. Finally, in the “motor suppression model” (Figure 6c), motor output is temporarily blocked, thereby suppressing responses during this period without impacting the decision variable. For each strategy, we extended the generalized drift-diffusion model (GDDM) described previously and fit parameters to RT distributions (see Online Methods). Of these models, the motor suppression model provided the best fit to the behavioral RT distribution for both monkeys as measured by Bayesian information criterion (BIC) (Figure 6d, ED9d).

Next, we examined the impact of each model on the speed-accuracy tradeoff. All three models increased accuracy at the expense of RT (Figure 6e-g, ED9e-g). The motor suppression model only resulted in slower RT in conditions where it also improved accuracy (Figure 6g, ED9g). By contrast, the pause and reset models increased RT of some conditions without improving the accuracy of those conditions (Figure 6e,f, ED9e,f). This shows how the motor suppression model can serve as an efficient mechanism for increasing accuracy without creating a generalized deficit in speed.

Figure 6: Comparison of computational models of the dip. (a-c) Schematic of the (a) pause model, (b) reset model, and (c) motor suppression model across stages of the decision-making process. (d) The fit of each GDDM to the RT distribution, as quantified by BIC, is shown for each of the three models. (e-g) For each coherence, presample, and reward condition, the difference in mean RT and accuracy, with and without the dip mechanism, are plotted for the (e) pause, (f) reset, and (g) motor suppression models. (h-j) Neural predictions of the (h) pause model, (i) reset model, and (j) motor suppression model, based on the mean decision variable or motor-decision variable (see Online Methods) for inside (in-RF, black) and outside (out-RF, red) the response field on correct high-coherence trials for models fit to the RT distribution of Monkey 1. (k) Population

activity from FEF neurons in Monkey 1 for correct responses inside (gray) and outside (red) the response field for high-coherence, 800-ms presample trials.

A comparison of the FEF population activity with the FEF activity predicted by each model also suggests that neural activity is best accounted for by the motor suppression model. For this analysis, we focused on the trials expected to show the strongest evidence dip, namely, those with 800-ms presample duration, high coherence, and the large reward target in the neuron's response field. The pause model predicts a flattening of the decision variable trace and hence a flattening of FEF activity, followed by an increase for trials with the chosen target inside the response field and a decrease for those with the chosen target outside the response field (Figure 6h, ED9h). This is inconsistent with the evidence dip in FEF. By contrast, the reset model predicted that the decision variable and population FEF activity should initially decrease in all trials but increase again only when the chosen choice is in the neuron's response field (Figure 6i, ED9i). However, population FEF activity decreased and then increased in all trials, before it eventually decreased in trials where the animal chose the target away from the neuron's response field, resulting in a small activity bump (Figure 6k).

The motor suppression model predicts that information about accumulated evidence should be maintained during the dip, but it also predicts that crossing the decision boundary should not trigger a saccade during the motor suppression interval. Thus, we hypothesized that FEF activity in this model might represent a "motor-decision variable", a version of the decision variable which is scaled-down by a constant factor during the suppression interval. Then, a saccade is triggered when the motor-decision

variable, not the decision variable, crosses the boundary. Unlike the pause or reset models, this can account for the rebound in the FEF activity observed regardless of whether the chosen target is inside or outside the response field (Figure 6j, ED9j). Therefore, the motor suppression model can parsimoniously account for the FEF population activity as well as the RT dip.

The motor suppression model makes another unique prediction about behavior, namely, that motor suppression triggered by the presample onset might affect other ongoing motor activity, such as microsaccades. To test this prediction, we examined the microsaccade rate over time for each coherence and presample condition. We indeed found that for the longest presample duration, there was a coherence-dependent reduction in microsaccade rate, such that high coherence changes elicited a greater reduction in microsaccade rate (Figure 7d, Figure ED9h). For one monkey, this difference in microsaccade rate was also significant for 0 ms and 400 ms presample durations (Figure 7b,c). This dip in microsaccade rate demonstrates that saccadic motor output is inhibited even when it does not involve a response to a target. Thus, the motor suppression model best fits the RT distribution, has a desirable speed-accuracy tradeoff, and is able to uniquely explain the rebound in FEF activity as well as the dip in microsaccade rate.

Figure 7: Dips in microsaccade rate. (a) Microsaccade rate is shown for all trials, aligned to the presample onset, for Monkey 1. Highlighted are the microsaccade rates centered around the onset of the sample during trials with (b) 0 ms, (c) 400 ms, and (d) 800 ms presample. The black bar indicates significance ($p < 0.05$, one-tailed test for the difference of means, bootstrapping across neurons). Shaded region represents bootstrapped 95% confidence interval.

Discussion

In the present study, we found that changes in evidence strength during a perceptual decision-making task led to a transient suppression of three separate measures, including RT distribution, population and single-neuron FEF activity, and microsaccade rate. Larger changes in evidence elicited a larger dip in all three measures, an observation that goes against classic models which predict that more evidence always shorten RTs and increase FEF activity for the corresponding behavioral response. By fitting models of three

potential cognitive strategies to the behavioral data, we found that transiently suppressing motor output after change detection explains our observations. Critically, it also explains the observations which motivated the other two strategies we examined (Engbert and Kliegl 2003; Kiani et al., 2008; Hafeed and Ignashchenkova 2013; Teichert et al., 2016; Buonocore et al. 2017; Salinas and Stanford 2018). While the motor suppression mechanism is sufficient to explain our observations, we cannot exclude the possibility of additional mechanisms acting in parallel. Our model also shares many similarities to previous models of saccade (Reingold and Stampe 2002; Bompas and Sumner 2011; Bompas and Sumner 2015) and motor (Purcell et al 2010; Wessel et al. 2016) inhibition.

Our results suggest that FEF does not directly encode the decision variable, but rather, it encodes a mixture of the decision variable and a motor signal. This means that FEF might not be directly responsible for integrating evidence, a conclusion supported by causal experiments (Peel et al. 2016) and by the fact that evidence presented during the dip period can still be integrated (Kiani et al., 2008; Huk and Meister 2012; Teichert et al. 2016). While this work could not address the neurobiological source of motor inhibition, one possibility is that ascending excitatory input synapses onto an inhibitory subpopulation within FEF, providing non-specific lateral inhibition across FEF. The presence of such an inhibitory population within FEF is supported by microstimulation studies, which show that sub-threshold stimulation of regions of FEF outside of the response field prevents task-related saccades during the stimulation period (Burman and Bruce 1997, Izawa et al. 2004), and by experiments showing that cooling FEF leads to increases in the duration of microsaccade suppression (Peel et al. 2016). The transient suppression of FEF may be mediated by the release of inhibition (Tehovnik et al., 1999; Opris et al., 2001; Izawa et al., 2009; Sommer and Wurtz, 2000).

An alternative view is that motor suppression arises from reduced ascending input, perhaps mediated by the subthalamic nucleus (Isoda and Hikosaka 2008; Pasquereau and Turner, 2017), and is common to many regions of the brain. In addition to FEF, a stimulus dip has been observed in LIP (Roitman and Shadlen 2002; Falkner et al. 2010; Bollimunta et al., 2011), superior colliculus (Ratcliff et al. 2003; Li et al. 2006), visual areas V1 and V2 (Nowak et al. 1995), and striatum (Aosaki et al. 1995; Ford and Everling 2009; Schulz and Reynolds 2013). Similar dips in neural activity have also been observed in FEF during a countermanding task (Brown et al. 2008) and striatum during an anti-saccade task (Ford and Everling 2009), two tasks which may require motor suppression. In behavioral experiments, the RT dip occurs even after task-irrelevant changes in the visual field (Reingold and Stampe 2000), after both high contrast changes and isoluminant changes (Reingold and Stampe 2000), and across the visual field (Reingold and Stampe 2002; Buonocore and McIntosh 2008; Edelman and Xu 2009). These studies collectively suggest that the dip, and motor inhibition more generally, are not localized exclusively to FEF or limited to our specific task, but may be a more general mechanism for dealing with a changing environment. Thus, dips caused by motor suppression may serve more broadly as a marker for attentional shifts or perception of stimulus changes, such as through the orienting reflex (Sokolov 1963; Courchesne et al. 1975; Bisley 2010).

436

Online Methods

Behavioral task

Two rhesus monkeys, 1 and 2, were trained to perform a two-alternative forced-choice color-discrimination task (Figure 1), as previously described in Shinn et al. (2020). In each trial, a central square stimulus was presented consisting of a 20×20 grid of green and blue pixels that rearranged randomly at 20Hz. The animal indicated its choice by shifting its gaze to one of two flanking choice targets, one green and one blue. The location of one target was chosen based on the response field of the neuron recorded in that session, determined through a memory saccade task, and the other target was opposite to the first. The trial was rewarded via juice delivery if the selected target color corresponded to the majority color of pixels in the sample. Reward cues were displayed surrounding the saccade targets which indicated whether a large or small reward would be delivered for a correct response to the corresponding target. Reward cues were randomly assigned to a target on each trial. 3 (2) drops of juice were given for the large (small) reward condition for Monkey 1, and 3-5 (1-2) drops for Monkey 2.

Stimulus presentation was divided into two consecutive periods containing an uninformative "presample" followed by an informative "sample". There was always an equal number of blue and green pixels displayed during the presample period. Task difficulty was manipulated by parametrically varying the difference in the fraction of pixels of each color in the sample, which we refer to as color coherence. A coherence of 0 indicated equal numbers of both colors, whereas 1 indicated a solid color. No explicit cue was presented to indicate the transition from presample to sample, and the change was

instantaneous. The presample duration was selected randomly from three possible time intervals—0, 400 or 800 ms—with equal probability. Animals were allowed to direct their gaze to a choice target any time after the onset of the sample. Eye movements were tracked at 225 hz with a high speed eye tracker (ET49; Thomas recording). A premature choice before the sample onset aborted the trial, and was punished by a 2-s timeout. In a small fraction of trials (5%), the sample was identical to the presample and maintained zero color coherence throughout the trial. On these trials, animals were allowed to respond at any point during stimulus presentation. The ratio of pixels in the high, medium, and low coherence trials for Monkey 1 were 70:30, 60:40, and 53:47, and for Monkey 2 were 63:37, 58:42, and 52:48, respectively.

Reaction time distribution and behavioral functions

The reaction time (RT) distribution and psychometric function were calculated using all the trials in which the monkey successfully completed a saccade to one of the two choice targets. This included the trials with incorrect choices, the trials with choices during the presample period, and the trials in which the animal failed to maintain fixation on the chosen target. The RT histogram was constructed using 20-ms bins and smoothed with an order 1 Savitzky-Golay filter of width 5. We estimated 95% confidence intervals in each coherence condition for visualization by bootstrapping across trials. We performed 10,000 resamplings with replacement from all trials with the given coherence across all sessions, computed the smoothed histogram of each resample, and then found the 2.5% and 97.5% quantiles at each time point.

Statistical significance for the difference in RT histogram at each point in time was determined using a separate bootstrapping procedure which compared the highest

and lowest coherence conditions.. This consisted of performing 10,000 resamplings with replacement from the highest and the lowest non-zero coherence trials, computing the confidence interval for mean RT difference between the two conditions at each point in time. The lowest non-zero coherence was used instead of zero-coherence trials due to the limited number of zero-coherence trials. The chronometric function was calculated using only correct choices made after the sample onset.

Analysis of microsaccades

We detected microsaccades using the method of Engbert and Kliegl (2003). Briefly, the time of microsaccade was determined as the center of an interval in which the eye velocity exceeded six times a robust estimator of the standard deviation. Large saccades, such as the saccades to choice targets and those resulting in fixation breaks, were excluded. The microsaccade rate was calculated in successive 20-ms bins aligned at the presample onset, and then smoothed using an order-1 Savitzky-Golay filter of width 3.

Visualization of population and single-neuron activity

We recorded from 57 neurons in the frontal eye field (FEF) of Monkey 1 and 23 neurons from FEF in Monkey 2. For each neuron, we constructed a mean spike density function for a given experimental condition by calculating mean spike counts in successive 20-ms bins aligned at the onset of the presample or sample, and smoothing with an order 1 Savitzky-Golay filter of width 3. This was then normalized by subtracting the mean firing rate of each neuron during the presample period of the trials with an 800 ms presample duration for all conditions. These spike density functions were averaged across all neurons together to obtain the population activity. Confidence intervals for visualization of the population activity were estimated by bootstrapping across neurons.

We performed 10,000 resamplings with replacement from all neurons across all sessions, computed the smoothed spike counts of each resample, and then found the 2.5% and 97.5% quantiles at each time point. Confidence intervals for statistical significance for the mean difference in spike count between highest and lowest coherence conditions were obtained from a separate bootstrapping procedure, where a sampling distribution of mean spike difference was formed from 10,000 resamplings with replacement of neurons for both conditions.

Activity shown for individual neurons was computed similarly to the population activity, but smoothed using an order 1 Savitzky-Golay filter of width 5 for visualization, and no normalization was applied.

Single-neuron regression analysis

In order to understand how task events and experimental conditions modulated the firing rate of each neuron over time, an ordinary least squares regression model was used to analyze the activity of individual neurons. Spike counts were predicted during three 100-ms intervals: the presample interval ($0 < t < 100$ ms), the sample interval ($P_i + 100 < t < P_i + 200$ ms), and the saccade interval ($S_i - 50 < t < S_i + 50$ ms), where t is the time since presample onset, P_i is the presample duration on trial i , and S_i is the time of the saccade on trial i . For each interval I , we predicted spike counts as

$$x_I^i = \beta_0 + \beta_1 C_i + \beta_2 R_i + \beta_3 F_i \quad (\text{Eq. 1})$$

where C_i is the color coherence on trial i , R_i is whether the large or small reward target was in the response field ($R_i = 1$ or -1 , respectively), and F_i is whether or not the choice was into or out of the response field ($F_i = 1$ or -1 , respectively).

To explore the change in spiking activity in response to task events, we implemented a time-resolved ordinary least squares regression analysis. Spikes were counted using 25-ms bins ($\Delta t=25$). Spike counts x_t^i at time t (measured from the presample onset) for trial i was predicted as

$$x_t^i = k_t^P + k_t^{PR} R_i + \left(k_{t-P_i}^E + k_{t-P_i}^{EC} C_i \right) \delta_{P_i \neq 0} + k_{t-S_i}^S + k_{t-S_i}^{SF} F_i \quad (\text{Eq. 2})$$

where, for trial i , P_i is the presample duration, S_i is the time of the saccade, C_i is the color coherence, R_i is whether the large or small reward target was in the response field ($R_i=1$ or -1 , respectively), F_i is whether or not the choice was into or out of the response field ($F_i=1$ or -1 , respectively), and δ is the indicator function. The k values correspond to kernels aligned to different events; k_0^j is the bin containing the given event for kernel j , and k_t^j is the bin t ms relative to the event, where $t \in [T_{start}^j, T_{start}^j + \Delta t, \dots, T_{end}^j]$ and $k_t^j=0$ if $t < T_{start}^j$ or $t > T_{end}^j$. Each kernel is identified by a short string, where the first letter indicates the alignment (“P” to presample, “E” to sample or evidence onset, and “S” to saccade) and subsequent letters indicate any conditions (“R” for reward magnitude, “C” for coherence, and “F” for the response field location). Presample-aligned kernels cover the entire trial duration, while sample-aligned kernels and saccade-aligned kernels span a short period surrounding or following the event for sample-aligned and saccade-aligned kernels, respectively. Since here we focus on the effect of evidence change (e.g. transition from presample to sample) on neural modulation, we estimated evidence kernel only for the trials where presample duration was not zero ($\delta_{P_i \neq 0}$). The six kernels in the equation above have the following constants:

Kernel (j)	Event aligned to	Scaling variable	Kernel start	Kernel end
			time (ms) (T_{start}^j)	time (ms) (T_{end}^j)
P	Presample	---	-500	2500
PR	Presample	Large reward in response field (R_i)	-500	2500
E	Sample	---	0	300
EC	Sample	Coherence (C_i)	0	300
S	Saccade	---	-200	200
SF	Saccade	Choice in response field (F_i)	-200	200

548 In the main text, we focus our analysis on the coherence-dependent evidence-aligned EC
549 kernel and the presample-aligned P kernel, which are expected to reflect transient
550 reductions, or dips, in FEF activity. Thus, we use the term “evidence kernel” to refer to
551 the EC kernel, and the term “stimulus kernel” to refer to the P kernel minus a baseline,
552 chosen as the mean value of the P kernel in the interval 0-100 ms.

553 To control for the dip as an artifact of saccadic activity, we fit a separate
554 regression model which differed only in the kernels used (Figure ED6). This model
555 includes extended saccade kernels scaled by the coherence. To distinguish the kernels in
556 this model from those of the previous model, we appended a "*" suffix. The full model is

$$x_t^i = k_t^{P*} + k_t^{PR*} R_i + \left(k_{t-P_i}^{E*} + k_{t-P_i}^{EC*} C_i \right) \delta_{P_i \neq 0} + k_{t-S_i}^{S*} + k_{t-S_i}^{SC*} C_i + \left(k_{t-S_i}^{SI*} + k_{t-S_i}^{SIC*} C_i \right) S_i. \quad (\text{Eq. 3})$$

The kernels used for this analysis were:

Kernel	Event aligned to	Scaling variable	Kernel start time (ms) (	Kernel end time (ms) (
			T_{start}^j)	T_{end}^j)
P*	Presample		-500	2500
PR*	Presample	Large reward in response field (R_i)	-500	2500
E*	Sample	---	0	400
EC*	Sample	coherence (C_i)	0	400
S*	Saccade		-800	200
SI*	Saccade	Choice in response field (S_i)	-800	200
SC*	Saccade	coherence	-1000	200
SIC*	Saccade	Choice in response field and coherence ($C_i \times S_i$)	-1000	200

Likewise, the “evidence* kernel” is defined as the EC* kernel, and the “stimulus* kernel” is defined as the P* kernel minus the mean of the P* kernel in the interval from 0-300 ms.

Analysis of transient activity

To examine whether the dips in the mean population kernels were representative of kernels at the single neuron level, singular value decomposition (SVD) was performed on the evidence and stimulus kernels across neurons. SVD is a standard technique in linear algebra which is similar in principle to principal component analysis (PCA), but does not normalize the single-neuron kernels by subtracting the mean kernel across the population at each time point. Briefly, for a matrix of kernels M , the singular vectors are defined to be the eigenvectors v_i of the matrix $M^T M$, sorted by decreasing eigenvalue. Likewise, the factor scores for the i 'th singular vector are defined as the projection of the data onto these singular vectors, namely Mv_i . Stimulus kernels were first truncated to the interval from 0 to 300 ms after stimulus onset before performing SVD.

To confirm that the dip in FEF activity was not confounded by pre-saccadic activity, we analyzed the timing of the dip separately for the trials of different RT quintiles. FEF population activity for each RT quintile was computed similarly to the FEF spike density function described previously by using a first-order Savitzky-Golay filter of width 5 (Figure ED6d-g,j-m). Fifteen resamplings of the mean spike density function were computed, and the first local minimum for each condition in each resampling was detected by a simple local minimum detection algorithm. The algorithm stepped through the timeseries until encountering a value exceeding the observed minimum value by a fixed tolerance determined empirically to be 0.12 spikes per second.

To understand the difference in latency between the dip in the evidence kernel and the dip in the stimulus kernel, we analyzed the mean of each of these kernels. We determined the minimum as the median of the three lowest values in the kernel within

300 ms from the sample or presample onsets for the evidence or stimulus kernels, respectively. We tested whether the minimum of the evidence kernel was later than the stimulus kernel by bootstrapping across neurons. We performed 10,000 resamplings with replacement from all evidence or stimulus kernels, computed the minimum of the mean using the procedure described above, and tested whether the minimum of the stimulus kernel was earlier than that of the evidence kernel.

The reduction in neural activity following sample or presample onset was quantified using the evidence and stimulus dip indices. These were used to provide a direct comparison between the evidence and stimulus dipoles, and to compare them to other physiological data. Dip indices were computed by z-scoring the kernel and then taking the mean z-scored evidence or stimulus kernel value from an interval I . We determined I separately for the evidence and stimulus dip indices by finding the two time points with the largest fraction of significant kernels (Figure 4d, 5b, ED4d, ED8b). For the evidence dip, we found $I=[125,175]$ ms for both monkeys. For the stimulus dip index, we found $I=[100,150]$ ms using the same procedure for Monkey 1. The two minimum time points were not consecutive for Monkey 2, so we used the same interval as in Monkey 1.

Generalized drift-diffusion model

The reaction times were modeled using the generalized drift-diffusion model (GDDM, Shinn et al, 2020b), which extends the standard drift-diffusion model (Ratcliff 1978) by allowing the model parameters to be arbitrary functions of time. The form of the GDDM used here is one of the models considered in Shinn et al. (2020a), and was previously found to have excellent performance during the task used in the present study. The model includes several extensions on the standard DDM to accommodate the temporal and

reward structure of the task. For the task’s temporal structure, the model includes leaky integration, as well as a delayed linear increase in gain during each trial. For the task’s reward structure, the model includes a baseline offset— impacting both starting position and the value to which leaky integration decays—and a “mapping error” in which the high reward choice is sometimes chosen after integrating sufficient evidence for a low reward choice.

The GDDM is given by the equation

$$dx = -l(x + mt)dt + I_{t > D} \Gamma(t) C s dt + \Gamma(t) dW$$

where s is the signal-to-noise ratio, l is the leak constant, m is the strength of the time-dependent reward bias, C is the coherence, D is the duration of the presample, W is a Wiener process, and

$$\Gamma(t) = \begin{cases} \gamma_0, & t < t_0 \\ \gamma_0 + m_\gamma(t - t_0), & t \geq t_0 \end{cases}$$

This GDDM model was fit to each monkey separately. It was fit using maximum likelihood on the full distribution through differential evolution (Shinn et al. 2020b). For robustness, fitting was performed using an exponential distribution mixture model with a rate and mixture strength fit to data (Shinn et al, 2020a). The model was simulated by solving the Fokker-Planck equation with a timestep of 5 ms and a space discretization of 0.005.

Modified GDDM to test potential mechanisms underlying the dip

In order to understand the relationship between the evidence, stimulus, and RT dips, we constructed modified GDDMs, implementing three potential cognitive mechanisms.. The

pause model was implemented by setting the drift rate and noise to 0 within an interval, $[t_{start}, t_{stop}]$ ms. The reset model was implemented by setting the leak l to a fixed value l_{dip} within an interval, $[t_{start}, t_{stop}]$ ms, where l_{dip} is fit to the data. The motor suppression model was implemented by introducing a motor-decision variable x' such that, for decision variable x ,

$$x' = \begin{cases} cx, & t_{start} \leq t \leq t_{stop} \\ x, & \text{otherwise} \end{cases}$$

where $0 < c < 1$. Here, we arbitrarily set $c=0.2$. We assume that, while the decision variable x continues to track integrated evidence, we only trigger a decision when the motor-decision variable x' crosses the boundary. For the purpose of efficient simulation, we use an equivalent formulation whereby the bound is increased by a factor of $1/c$ during the interval from $[t_{start}, t_{stop}]$ ms. In order to mitigate numerical artifacts caused by the abrupt change in bound, the change in bound height was implemented to be a smooth increase and decrease in the form of the probability density function of the Beta(3,3) distribution.

We assume that these dips were more likely to occur in trials with a higher color coherence, with a probability determined by a saturating sigmoidal curve with a fixed scale which was fit to the data. The probability of detecting the change on any given trial (and therefore invoking the given dip mechanism) was

$$p_{detect} = 2 / (1 + \exp(-\lambda C)) - 1$$

where C is the color coherence ranging from 0 for an equal pixel ratio to 1 for a solid color. A total of three additional parameters were fit for the pause and motor suppression models, and four additional parameters for the reset model. Four GDDMs were constructed, one for each of the three dip mechanisms described above, and one baseline

containing none of the dip mechanisms. Parameters were fit for each model by simulating with the Crank-Nicolson (pause and reset) or implicit (motor suppression) method, using differential evolution to optimize the likelihood over the full probability distribution (Shinn 2020b). Measurements of model BIC were computed using this full distribution likelihood. To compute ΔBIC , we subtracted each modified GDDM's BIC from the BIC of the unmodified GDDM.

To gain insights into the potential function of the dips in perceptual decision-making, we evaluated the impact of each dip mechanism on the RT and accuracy by comparing the RT and accuracy of each of the modified GDDMs to those from the same modified GDDM but with the modification disabled, i.e. keeping the same parameters for all other aspects of the model.

Predicted FEF activity was determined by transforming the decision variable trace. For each model, we simulated the monkey's highest coherence when the correct response was inside or outside the response field, with an 800-ms presample and the large reward target inside the response field. We recorded the time evolution of the decision variable distribution, and took the mean at each point in time to obtain a deterministic prediction of the decision variable value. For the motor suppression model, while we simulated by increasing the integration bounds, we interpret this equivalently as a temporary multiplicative suppression of the decision variable by the corresponding amount. As a result, FEF activity was predicted by dividing the decision variable value by the height of the bound. Finally, to simulate the spike density functions, decision variable traces were thresholded at 0 and filtered with a Gaussian kernel.

673 Simulations were performed using the PyDDM package (Shinn et al., 2020b). Upon
674 acceptance for publication, we will make all GDDMs available in the PyDDM
675 documentation. Other datasets generated during and/or analyzed during the current study
676 are available from the corresponding author on reasonable request.

677 **Acknowledgments**

678 We thank Matthew Gay, Joey Schnurr and Irina Bobeica for their technical support. This
679 study was supported by the National Institute of Health (R01 MH108629; R01
680 MH112746). MS was supported by the Gruber Foundation Science Fellowship.

681 **Author contributions**

682 DL and HS designed and performed the experiments. MS, JDM, and HS analyzed the
683 data. All authors interpreted the results. MS and DL wrote the paper. All authors
684 reviewed, edited, and approved the manuscript.

Extended Data Figures

Figure ED1

Figure ED1: Behavioral data and model characterization. (a-b) Psychometric (a) and chronometric (b) functions are shown for Monkey 1 (left) and Monkey 2 (right).

Figure ED2

691

692 *Figure ED2: Dips in RT distributions and FEF recordings for Monkey 2. (a-d) The RT*
 693 *distribution is shown for Monkey 2. Similar format as Figure 2. (e-h) The FEF*
 694 *population activity is shown for Monkey 2. Similar format as Figure 3.*

Figure ED3

696

697 *Figure ED3: Influence of task conditions on firing rate. For each neuron, a regression*
 698 *model was fit to a 100-ms interval at the onset of the presample, sample, and saccade.*
 699 *Each model included four terms—coherence, the large reward target in the response*
 700 *field, the choice into the response field, and an intercept. The number of neurons for*
 701 *which each coefficient showed a significantly positive ($p < 0.05$, two-tailed t -test) value is*
 702 *shown for each monkey. Dark colors indicate a significantly positive value, and light*
 703 *colors indicate a significantly negative value.*

704 **Figure ED4**

705

706 *Figure ED4: Evidence dip in individual FEF neurons for Monkey 2. Similar format as*
707 *Figure 4.*

Figure ED5

Figure ED5: The sample kernel (referred to as E kernel in the Methods) does not show a dip. Format is similar to Figure 4, except using the coherence-independent, sample-aligned E kernel instead of the evidence kernel. Monkey 1 is shown on the left, and 2 on the right.

715 **Figure ED6**

716

717 *Figure ED6: The evidence dip is not saccade-related. (a) A schematic of the alternative*
 718 *regression model, with kernels designated by a * suffix. (b,c,h,i) Format similar to Figure*
 719 *4d, except using the evidence* kernel instead of the evidence kernel, for Monkey 1 (b,c)*
 720 *and 2 (h,i). (d,j) Normalized FEF activity from high-coherence 800 ms presample trials is*

721 *aligned to the sample onset and plotted separately for each of 5 RT bins for Monkey 1 (d)*
722 *and 2 (j). (e,k) Normalized FEF activity from these trials is aligned to the saccade onset*
723 *and plotted separately for each of 5 RT bins for Monkey 1 (e) and 2 (k). (f,m) Time of the*
724 *first local minimum after the sample onset for the curves in (e) and (k) are plotted for*
725 *Monkey 1 (f) and 2 (m) against the mean RT of the trials in the RT bin. Points indicate*
726 *the median of 15 resamplings of the FEF activity, and error bars indicate the*
727 *interquartile range. (g) Time of the first local minimum before the saccade onset for the*
728 *curves in (f) are plotted for Monkey 1 against the mean RT of the trials in the RT bin.*
729 *Minima could not be reliably detected in for saccade-aligned activity in Monkey 2.*

Figure ED7

731

732 Figure ED7: Evidence and stimulus dip indices are not correlated with neuron
 733 properties. The evidence and stimulus dip indices for each neuron are plotted against the
 734 weight on SV2 of the evidence kernel, coefficient of partial determination (CPD) for the

735 *evidence kernel, firing rate, directional selectivity index (DSI), visuomovement (VM)*
736 *index, and the neuron's mean spike peak-to-peak width. Spearman correlation and p-*
737 *value is inset.*

Figure ED8

Figure ED8: Comparison of evidence and stimulus dipoles for Monkey 2. Same format as

Figure 5.

742 Figure ED9

743

744 Figure ED9: Comparison of computational models of the dip for Monkey 2. (a-h) Same

745 format as Figure 6. (i-m) Same format as Figure 7.

References

1. Annika Åkerfelt, Colonius, H. & Diederich, A. Visual-tactile saccadic inhibition. *Experimental Brain Research* **169**, 554–563 (2005).
2. Aosaki, T., Kimura, M. & Graybiel, A. M. Temporal and spatial characteristics of tonically active neurons of the primate's striatum. *Journal of Neurophysiology* (1995) doi:10.1152/jn.1995.73.3.1234.
3. Bisley, J. W. The neural basis of visual attention. *The Journal of Physiology* **589**, 49–57 (2010).
4. Bollimunta, A. & Ditterich, J. Local Computation of Decision-Relevant Net Sensory Evidence in Parietal Cortex. *Cerebral Cortex* **22**, 903–917 (2011).
5. Bompas, A. & Sumner, P. Saccadic Inhibition Reveals the Timing of Automatic and Voluntary Signals in the Human Brain. *Journal of Neuroscience* **31**, 12501–12512 (2011).
6. Bompas, A. & Sumner, P. Saccadic inhibition and the remote distractor effect: One mechanism or two? *Journal of Vision* **15**, 15 (2015).
7. Boucher, L., Palmeri, T. J., Logan, G. D. & Schall, J. D. Inhibitory control in mind and brain: An interactive race model of countermanding saccades. *Psychological Review* **114**, 376–397 (2007).
8. Brown, J. W., Hanes, D. P., Schall, J. D. & Stuphorn, V. Relation of frontal eye field activity to saccade initiation during a countermanding task. *Experimental Brain Research* **190**, 135–151 (2008).

- 767 9. Buonocore, A. *et al.* Alteration of the microsaccadic velocity-amplitude main
768 sequence relationship after visual transients: implications for models of saccade
769 control. *Journal of Neurophysiology* **117**, 1894–1910 (2017).
- 770 10. Buonocore, A. & McIntosh, R. D. Saccadic inhibition underlies the remote distractor
771 effect. *Experimental Brain Research* **191**, 117–122 (2008).
- 772 11. Burman, D. D. & Bruce, C. J. Suppression of Task-Related Saccades by Electrical
773 Stimulation in the Primate Frontal Eye Field. *Journal of*
774 *Neurophysiology* **77**, 2252–2267 (1997).
- 775 12. Courchesne, E., Hillyard, S. A. & Galambos, R. Stimulus novelty, task relevance and
776 the visual evoked potential in man. *Electroencephalography and Clinical*
777 *Neurophysiology* **39**, 131–143 (1975).
- 778 13. Ding, L. & Gold, J. I. Neural Correlates of Perceptual Decision Making before,
779 during, and after Decision Commitment in Monkey Frontal Eye Field. *Cerebral*
780 *Cortex* **22**, 1052–1067 (2012).
- 781 14. Edelman, J. A. & Xu, K. Z. Inhibition of Voluntary Saccadic Eye Movement
782 Commands by Abrupt Visual Onsets. *Journal of Neurophysiology* **101**, 1222–1234
783 (2009).
- 784 15. Engbert, R. & Kliegl, R. Microsaccades uncover the orientation of covert attention.
785 *Vision Research* **43**, 1035–1045 (2003).
- 786 16. Falkner, A. L., Krishna, B. S. & Goldberg, M. E. Surround Suppression Sharpens the
787 Priority Map in the Lateral Intraparietal Area. *Journal of Neuroscience* **30**, 12787–
788 12797 (2010).
- 789 17. Ford, K. A. & Everling, S. Neural Activity in Primate Caudate Nucleus Associated
790 With Pro- and Antisaccades. *Journal of Neurophysiology* **102**, 2334–2341 (2009).

- 791 18. Glaze, C. M., Kable, J. W. & Gold, J. I. Normative evidence accumulation in
792 unpredictable environments. *eLife* **4**, (2015).
- 793 19. Gold, J. I. & Shadlen, M. N. The Neural Basis of Decision Making. *Annual Review of*
794 *Neuroscience* **30**, 535–574 (2007).
- 795 20. Hafed, Z. M. & Ignashchenkova, A. On the Dissociation between Microsaccade Rate
796 and Direction after Peripheral Cues: Microsaccadic Inhibition Revisited. *Journal of*
797 *Neuroscience* **33**, 16220–16235 (2013).
- 798 21. Hanes, D. P. & Schall, J. D. Neural Control of Voluntary Movement Initiation.
799 *Science* **274**, 427–430 (1996).
- 800 22. Hanks, T., Kiani, R. & Shadlen, M. N. A neural mechanism of speed-accuracy
801 tradeoff in macaque area LIP. *eLife* **3**, (2014).
- 802 23. Huk, A. C. & Meister, M. L. R. Neural correlates and neural computations in
803 posterior parietal cortex during perceptual decision-making. *Frontiers in Integrative*
804 *Neuroscience* **6**, (2012).
- 805 24. Isoda, M. & Hikosaka, O. Role for Subthalamic Nucleus Neurons in Switching from
806 Automatic to Controlled Eye Movement. *Journal of Neuroscience* **28**, 7209–7218
807 (2008).
- 808 25. Izawa, Y., Suzuki, H. & Shinoda, Y. Response Properties of Fixation Neurons and
809 Their Location in the Frontal Eye Field in the Monkey. *Journal of Neurophysiology*
810 **102**, 2410–2422 (2009).
- 811 26. Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded Integration in Parietal Cortex
812 Underlies Decisions Even When Viewing Duration Is Dictated by the Environment.
813 *Journal of Neuroscience* **28**, 3017–3029 (2008).

- 814 27. Li, X., Kim, B. & Basso, M. A. Transient Pauses in Delay-Period Activity of Superior
815 Colliculus Neurons. *Journal of Neurophysiology* **95**, 2252–2264 (2006).
- 816 28. Logan, G. D., Yamaguchi, M., Schall, J. D. & Palmeri, T. J. Inhibitory control in
817 mind and brain 2.0: Blocked-input models of saccadic countermanding.
818 *Psychological Review* **122**, 115–147 (2015).
- 819 29. Nowak, L. G., Munk, M. H. J., Girard, P. & Bullier, J. Visual latencies in areas V1
820 and V2 of the macaque monkey. *Visual Neuroscience* **12**, 371–384 (1995).
- 821 30. Opris, I., Ferrera, V. P. & Barborica, A. On the gap effect for saccades evoked by
822 electrical microstimulation of frontal eye fields in monkeys. *Experimental Brain*
823 *Research* **138**, 1–7 (2001).
- 824 31. Pasquereau, B. & Turner, R. S. A selective role for ventromedial subthalamic nucleus
825 in inhibitory control. *eLife* **6**, (2017).
- 826 32. Peel, T. R., Hafed, Z. M., Dash, S., Lomber, S. G. & Corneil, B. D. A Causal Role for
827 the Cortical Frontal Eye Fields in Microsaccade Deployment. *PLOS Biology* **14**,
828 e1002531 (2016).
- 829 33. Purcell, B. A. *et al.* Neurally constrained modeling of perceptual decision making.
830 *Psychological Review* **117**, 1113–1143 (2010).
- 831 34. Ratcliff, R. A theory of memory retrieval. *Psychological Review* **85**, 59–108 (1978).
- 832 35. Ratcliff, R., Cherian, A. & Segraves, M. A Comparison of Macaque Behavior and
833 Superior Colliculus Neuronal Activity to Predictions From Models of Two-Choice
834 Decisions. *Journal of Neurophysiology* **90**, 1392–1407 (2003).
- 835 36. Reingold, E. M. & Stampe, D. M. Saccadic Inhibition and Gaze Contingent Research
836 Paradigms. in *Reading as a Perceptual Process* 119–145 (Elsevier, 2000).
837 doi:10.1016/b978-008043642-5/50008-5.

- 838 37. Reingold, E. M. & Stampe, D. M. Saccadic Inhibition in Voluntary and Reflexive
839 Saccades. *Journal of Cognitive Neuroscience* **14**, 371–388 (2002).
- 840 38. Roitman, J. D. & Shadlen, M. N. Response of Neurons in the Lateral Intraparietal
841 Area during a Combined Visual Discrimination Reaction Time Task. *The Journal of*
842 *Neuroscience* **22**, 9475–9489 (2002).
- 843 39. Salinas, E. & Stanford, T. R. Saccadic inhibition interrupts ongoing oculomotor
844 activity to enable the rapid deployment of alternate movement plans. *Scientific*
845 *Reports* **8**, (2018).
- 846 40. Salinas, E. *et al.* Voluntary and involuntary contributions to perceptually guided
847 saccadic choices resolved with millisecond precision. *eLife* **8**, (2019).
- 848 41. Sato, T. & Schall, J. Pre-excitatory pause in frontal eye field responses. *Experimental*
849 *Brain Research* **139**, 53–58 (2001).
- 850 42. Schulz, J. M. & Reynolds, J. N. J. Pause and rebound: sensory control of cholinergic
851 signaling in the striatum. *Trends in Neurosciences* **36**, 41–50 (2013).
- 852 43. Shinn, M., Ehrlich, D. B., Lee, D., Murray, J. D. & Seo, H. Confluence of Timing and
853 Reward Biases in Perceptual Decision-Making Dynamics. *The Journal of*
854 *Neuroscience* **40**, 7326–7342 (2020).
- 855 44. Shinn, M., Lam, N. H. & Murray, J. D. A flexible framework for simulating and
856 fitting generalized drift-diffusion models. *eLife* **9**, (2020).
- 857 45. Sokolov, E. N. Higher Nervous Functions: The Orienting Reflex. *Annual Review of*
858 *Physiology* **25**, 545–580 (1963).
- 859 46. Sommer, M. A. & Wurtz, R. H. Composition and Topographic Organization of
860 Signals Sent From the Frontal Eye Field to the Superior Colliculus. *Journal of*
861 *Neurophysiology* **83**, 1979–2001 (2000).

- 862 47. Tehovnik, E. J., Slocum, W. M. & Schiller, P. H. Behavioural conditions affecting
863 saccadic eye movements elicited electrically from the frontal lobes of primates.
864 *European Journal of Neuroscience* **11**, 2431–2443 (1999).
- 865 48. Teichert, T., Grinband, J. & Ferrera, V. The importance of decision onset. *Journal of*
866 *Neurophysiology* **115**, 643–661 (2016).
- 867 49. Wessel, J. R. & Aron, A. R. On the Globality of Motor Suppression: Unexpected
868 Events and Their Influence on Behavior and Cognition. *Neuron* **93**, 259–280 (2017).
- 869 50. Zoltowski, D. M., Latimer, K. W., Yates, J. L., Huk, A. C. & Pillow, J. W. Discrete
870 Stepping and Nonlinear Ramping Dynamics Underlie Spiking Responses of LIP
871 Neurons during Decision-Making. *Neuron* **102**, 1249–1258.e10 (2019).