

# Bayesian inference by visuomotor neurons in prefrontal cortex

Thomas Langlois<sup>1\*</sup>, Julie A. Charlton<sup>2\*</sup>, Robbe L. T. Goris<sup>1+</sup>

\* equal contribution <sup>1</sup> Center for Perceptual Systems, The University of Texas at Austin, Austin, TX USA. <sup>2</sup> Princeton Neuroscience Institute, Princeton University, Princeton, NJ, USA. + email: robbe.goris@utexas.edu

**Perceptual judgements of the environment emerge from the concerted activity of neural populations in decision-making areas downstream of sensory cortex [1, 2, 3]. When the sensory input is ambiguous, perceptual judgements can be biased by prior expectations shaped by environmental regularities [4, 5, 6, 7, 8, 9, 10, 11]. These effects are examples of Bayesian inference, a reasoning method in which prior knowledge is leveraged to optimize uncertain decisions [12, 13]. However, it is not known how decision-making circuits combine sensory signals and prior expectations to form a perceptual decision. Here, we study neural population activity in the prefrontal cortex of macaque monkeys trained to report perceptual judgments of ambiguous visual stimuli under two different stimulus distributions. We analyze the component of the neural population response that represents the formation of the perceptual decision (the decision variable, DV), and find that its dynamical evolution reflects the integration of sensory signals and prior expectations. Prior expectations impact the DV's trajectory both before and during stimulus presentation such that DV trajectories with a smaller dynamic range result in more biased and less sensitive perceptual decisions. These results reveal a mechanism by which prefrontal circuits can execute Bayesian inference.**

Perceptual systems infer properties of the environment from sensory measurements that can be ambiguous. However, prior knowledge can be leveraged to disambiguate the interpretation [13]. This inference strategy typically manifests as perceptual interpretations that are biased towards the prior expectation. Such biases may reflect implicit knowledge of statistical regularities that are stable features of the environment, such as the tendencies of sunlight to come from above [4], image velocities to be slow [6, 7], and cardinal orientations to be over-represented in visual scenes [14]. However, biased perceptual interpretations can also reflect knowledge of statistical regularities that are context-specific and short-lived in nature [8, 9, 10, 11]. The diversity of experimental settings under which prior-induced perceptual biases occur suggests that a general neural mechanism may underlie these effects [15, 16, 17].

Prior expectations about sensory stimulation modulate neural activity in decision circuits in various ways. Context cues that signal specific environmental statistics can selectively modulate activity of single cells before stimulus onset [18], during stimulus presentation [19, 20, 21, 22], and while the perceptually-driven behavior is being produced [21, 22]. At the population level, these effects conspire to bias neural representations towards the prior expectation [22]. Because these neural effects co-occur with biases in perceptual reports, they are thought to reflect the neural computations that govern the perceptual interpretation of the environment [18, 20, 21, 22]. However, there is an alternative explanation that cannot be ruled out. In all of these previous studies, perceptual interpretations had a fixed relationship to the overt motor response. Biases in perceptual reports thus coincided with biases in motor responses. Decision-making circuits are typically involved in action planning [23, 24, 3]. It is thus not clear whether the reported neural correlates of perceptual expectation pertain to perceptual inference, motor planning, or a mixture of the two.

To obtain an unobstructed view on the neural correlates of perceptual expectation and inference, we used a task that requires flexible reporting of perceptual decisions [3]. Monkeys judged whether a visual stimulus was oriented clockwise or counterclockwise from vertical and communicated their decision with a saccadic eye movement towards one of two visual targets (Fig. 1a). The meaning of each response option was signaled by the target's orientation (clockwise vs counterclockwise) and was unrelated to its spatial position (one target was placed in the neurons' estimated motor response field, the other was placed on the opposite side of the fixation mark, see Methods). Because the spatial configuration of the choice targets varied randomly from trial to trial, changes in prior stimulus statistics biased the animals' perceptual reports, but not their overt motor responses. While the animals performed this task, we recorded extracellular responses from neural ensembles in the prearcuate gyrus (Supplemental Fig. 1), an area of prefrontal cortex (PFC) involved in the selection of saccadic eye movements [25, 26] that represents visuomotor deliberation [27, 28]. We previously found that the population activity initially represents the formation of a perceptual choice, before transitioning into the representation of the motor plan [3]. Here, we study how prior expectations and sensory signals shape the perceptual decision-making process at the single trial level. Our results reveal that prior expectations selectively change the dynamic evolution of the decision-related component of population activity in PFC. Expectations impact the starting point, slope, and dynamic range of the stimulus-driven trajectory in neural activity, culminating

51 in more biased and less sensitive perceptual decisions.

## 52 Results

### 53 Task and Behavior

54 Two rhesus macaques were trained to perform a visual orientation discrimination task in which they communicated perceptual  
55 decisions under two different stimulus-response mapping rules (Fig. 1a). The data were previously described in detail in ref.  
56 [3]. Task difficulty was controlled by manipulating stimulus orientation and contrast (Fig. 1b). Monkeys received a reward if  
57 they selected the appropriately oriented choice target. They performed the task similarly well under both mapping rules (median  
58 performance: monkey F = 79.7% correct; monkey J = 79.6% correct; median difference in performance across mapping rules:  
59 monkey F = 2.58%,  $P = 0.002$ ; monkey J = 0.5%,  $P = 0.57$ ; Wilcoxon signed-rank test). We manipulated the animals' prior  
60 expectations by sampling stimulus orientation from one of two skewed distributions (Fig. 1c). On each trial, the shape of the  
61 fixation mark signaled the current prior condition. The prior condition and stimulus contrast were switched across short blocks  
62 of trials (see Methods). The prior manipulation selectively biased the animals' perceptual judgments of ambiguous stimuli  
63 (median difference in choice bias for vertical gratings: monkey F = 1.22 log odds,  $P < 0.001$ ; monkey J = 1.49 log odds,  $P <$   
64  $0.001$ ), but not the motor responses used to communicate these decisions (median difference in motor bias: monkey F = 0.03  
65 log odds,  $P = 0.31$ ; monkey J = 0.03 log odds,  $P = 0.06$ ; Fig. 1d). Together, these results suggest that our paradigm engages  
66 brain mechanisms that specifically bias perceptual judgements of ambiguous sensory inputs.

67 What is the nature of the inference process that underlies these biased judgements of visual stimuli? A prominent hypothesis  
68 is that these biases naturally arise under an inference strategy that seeks to make the best possible decision given ambiguous  
69 sensory measurements and prior experience [12, 13]. This notion is formalized in the framework of Bayesian inference. An  
70 ideal Bayesian decision-maker computes the posterior evidence in support of each response option by multiplying the prior  
71 probability of each stimulus orientation with the likelihood that a specific stimulus orientation gave rise to the current sensory  
72 measurement (Fig. 1e, top). As a consequence, perceptual decisions are biased for all stimulus orientations, manifesting as a  
73 horizontal shift of the psychometric function (Fig. 1e, bottom, grey vs orange line). The impact of the prior on the decision  
74 depends on the ambiguity of the sensory response. The more ambiguous the sensory response, the broader the likelihood  
75 function, and the larger the impact of the prior on the posterior (Fig. 1e, bottom left vs right panel). A Bayesian inference  
76 strategy maximizes choice accuracy, even though it results in decision biases and systematic errors (Fig. 1f, grey vs orange  
77 lines). Does this strategy explain why the monkeys' judgments of ambiguous stimuli were biased? In keeping with this  
78 hypothesis, the monkeys' decisions were biased for all stimulus orientations, not just vertical gratings (Fig. 1g, red vs blue  
79 symbols). Moreover, they tended to make more biased decisions in task conditions associated with higher sensory uncertainty.  
80 We estimated decision bias by measuring the separation between the prior-conditioned psychometric functions, and sensory  
81 uncertainty by calculating the slope of the psychometric function (see Methods; Fig. 1g). For both monkeys, decision bias and  
82 sensory uncertainty were significantly correlated (Spearman rank correlation: Monkey J =  $-0.42$ ,  $P = 0.017$ ; Monkey F =  $-0.60$ ,  
83  $P = 0.0015$ ; Fig. 1h). This data pattern suggests that the animals used a perceptual decision-making strategy that resembles  
84 Bayesian inference.

### 85 Linking PFC population activity to perceptual inference at the single trial level

86 The animals' choice behavior was variable. For most experimental conditions, the overall proportion of clockwise choices was  
87 neither zero nor one (Fig. 1g). This choice variability may either arise from cross-trial variability in sensory measurements  
88 [29, 30, 31, 32] or from fluctuations in prior expectations [11]. Given that choice variability was minimal or absent for the most  
89 extreme stimulus orientations (Fig. 1g), less likely sources are noise in the choice-response mapping process or attentional  
90 "lapses". Because the choice variability has no obvious external origin, identifying the neural factors that determine the outcome  
91 of individual decisions ultimately requires a moment-to-moment analysis of neural population activity within single trials  
92 [33, 34].

93 We obtained a trial-by-trial measurement of the animals' evolving decision state by decoding a time-varying decision variable  
94 (DV) from jointly recorded neural responses using linear discriminant analysis (see Methods). The DV indicates how well the  
95 subject's upcoming perceptual choice can be predicted from a 50 ms bin of neural population activity. In previous work, we  
96 showed that trial-averaged DV trajectories exhibit key signatures of a decision-making process [3]. Most importantly, grouping  
97 trials by stimulus strength and choice accuracy revealed a graded representation of sensory evidence (the more the stimulus  
98 orientation deviates from vertical, the higher the sensory evidence). By contrast, the action-planning component of neural  
99 activity did not show these effects. We concluded that the decision-making process was implemented as a competition between  
100 possible perceptual interpretations, independent of the ensuing action. Here, we build on these findings to develop an analysis

101 aimed at uncovering how decision-making circuits integrate perceptual expectations and sensory signals.

102 Consider the DV trajectories of two example trials, measured from the same set of neurons in the same experimental condition  
103 (Fig. 2a, symbols). Both trials yielded identical overt choice behavior (a clockwise choice indicated with a rightward saccade).  
104 The raw DV trajectories are noisy, but share some clear commonalities. On both trials, the DV value hovers around zero before  
105 stimulus onset. Following stimulus onset, the DV ramps up to a peak value, after which it decays steadily. We speculate that the  
106 peak occurs around the time of choice commitment, after which the animal starts to prepare the corresponding motor response  
107 [3]. Closer inspection of the trajectories reveals that the build-up towards the peak differs in a number of ways. The excursions  
108 seem to start from a different baseline level and the ramping phase appears to differ in slope and amplitude. We hypothesize  
109 that these features capture an important aspect of the temporal evolution of the animal's decision state. We therefore sought to  
110 obtain quantitative estimates of these trajectory features for each trial.

111 We estimated key features of the DV trajectories by fitting a model in which trajectories evolve smoothly to the raw DV values  
112 (Fig. 2a, lines; see Methods). This way, we obtained a single "model DV-trajectory" for each trial. Due to the irregular  
113 nature of the raw DV trajectories, the correlation between model DV-trajectory and raw DV value was modest (median Pearson  
114 correlation: Monkey J = 0.63; Monkey F = 0.58; Fig. 2b). Nevertheless, the model captured the systematic structure of the trial-  
115 averaged data well. This was evident at the level of individual recording sessions (example shown in Fig. 2c). To test whether  
116 the model DV-trajectories afford additional insight into the animals' decision state, we conducted a logistic regression analysis  
117 of the choice behavior. For each recording session, we first measured the association between the experimentally controlled  
118 variables (stimulus prior, orientation, and contrast) and choice outcome. We then asked whether additionally including the  
119 peak value of the model DV-trajectory as a regressor helped to better predict choice outcome (see Methods). We computed  
120 a standard measure for prediction error (Akaike's Information Criterion, AIC) and found that including the peak value of the  
121 model DV-trajectories systematically improved prediction quality (manifesting as a difference in AIC that is larger than 0 in  
122 Fig. 2d). Together, these results suggests that for each trial, the model DV-trajectory provides a useful estimate of the DV's  
123 baseline level and the slope and magnitude of its ramping phase.

124 We have argued that the monkeys' behavior suggests that they combined prior expectations and sensory inputs in a manner  
125 that resembles Bayesian inference. We therefore hypothesized that the peak value of the neurally decoded DV should exhibit  
126 key signatures of a Bayesian posterior. To test this, we investigated how the model DV-trajectories' peak value depended  
127 on the stimulus prior, orientation, and contrast. In the vast majority of recording sessions, the peak DV value exhibited the  
128 hypothesized structure, as can be seen for an example session (Fig. 2e). Specifically, the peak DV grows with stimulus  
129 orientation, but is biased by the stimulus prior. The impact of the sensory input is larger when stimulus contrast is high, as is  
130 evident from the difference in the slope of a linear regression line (Fig. 2e, left vs right panel, slope = 0.18 for low contrast  
131 stimuli and 0.31 for high contrast stimuli). Conversely, the impact of the prior expectation is larger when stimulus contrast is  
132 low, as is evident from the difference in the offset of the regression lines (Fig. 2e, left vs right panel, offset = 0.86 for low  
133 contrast stimuli and 0.35 for high contrast stimuli). These effects are representative of both animals' data (median difference  
134 in slope between low and high contrast stimuli: monkey J = 0.02,  $P = 0.003$ ; monkey F = 0.01,  $P = 0.007$ ; one-sided Wilcoxon  
135 signed-rank test; Fig. 2f; median difference in offset: monkey J = -0.28,  $P = 0.016$ ; monkey F = -0.04,  $P = 0.05$ ; Fig. 2g). Thus,  
136 around the putative time of choice commitment, the neurally estimated decision-state exhibits key signatures of a Bayesian  
137 posterior. Having established this, we now turn to the question of how prior expectations impact the evolving decision-state  
138 over the course of a single trial.

## 139 **Prior expectations bias neural population trajectories**

140 The effects of prior expectation on the evolving decision-state can be appreciated by considering neural responses to low contrast  
141 stimuli whose orientation matches the categorization boundary (vertical gratings). These stimuli yield sensory responses that  
142 are, on average, completely ambiguous. It follows that any systematic choice biases, summarized in Fig. 1d, are purely driven  
143 by the subject's prior belief. As can be seen for an example recording session, the neural DV reflects this prior belief well  
144 before stimulus onset (Fig. 3a). At this point in time, the animal can only leverage knowledge about the skew of the stimulus  
145 distribution (i.e., whether a clockwise or counter-clockwise orientation is more likely) and the level of sensory uncertainty (i.e.,  
146 whether a low or high contrast stimulus is more likely). For both animals, this contextual information biases the DV's early  
147 value in the direction of the prior expectation, though the effect was more prominent for Monkey J than for Monkey F. This  
148 was evident from the DV traces averaged across all sessions (Fig. 3b) and from the model DV-trajectory value before stimulus  
149 onset (difference in median value for the -800 to -600 ms epoch: monkey J = 0.118,  $P < 0.001$ ; two-sided Wilcoxon rank  
150 sum test; monkey F = 0.008,  $P < 0.001$ ; Fig. 3c). As can also be seen from the average DV traces, the prior-induced bias  
151 is not static, but grows over time. To quantify this, we used the model DV-trajectories to compute the average DV value per  
152 recording session for an early and late temporal window (spanning the range from -800 to -600 ms and -500 to -300 ms).  
153 The first window precedes any stimulus-driven activity in PFC, while the second one overlaps with the putative time of choice

154 commitment. Plotting the late against the early DV value reveals a relationship that is steeper than the line of unity (Fig. 3d).  
155 For both monkeys, the gradient of the first principle component was approximately 72.5 degrees, significantly steeper than 45  
156 degrees (monkey J: 72.0 deg,  $P < 0.001$ , bootstrap analysis, see Methods; monkey F: 72.7 deg,  $P < 0.001$ ).

157 Why does the influence of the prior expectation grow during deliberation? We hypothesize that both animals integrated incom-  
158 ing sensory evidence over time and gave more weight to noisy orientation estimates that agreed with their prior expectation.  
159 Such a confirmation bias deviates from an ideal inference strategy, in which all sensory evidence is given the same weight,  
160 resulting in a flat average DV trace for ambiguous stimuli. However, a confirmation bias naturally arises when subjects use  
161 an inference strategy that approximates the posterior rather than computing it exactly [35]. Under approximate inference, as  
162 incoming sensory evidence is accumulated to update the posterior belief in the state of the world, the prior belief enters into the  
163 update multiple times, resulting in a positive feedback loop [35]. Critically, this explanation implies that a model comprised  
164 of an initial prior-induced offset and a bound that terminates the deliberation process does not suffice to capture the dynam-  
165 ically increasing influence of the prior on the DV. The extra ingredient that is needed is a prior-induced drift [35]. To test this  
166 explanation for our data, we simulated bounded accumulation of ambiguous sensory evidence and compared early and late DV  
167 values (see Methods; Fig. 3e). We varied the strength of the prior-induced effects and the height of the bound, resulting in a  
168 range of DV values (Fig. 3f). In the simulations that only included an early offset and lacked a prior-induced drift, early and  
169 late DV values were equal on average (Fig. (3f, left)). We found that inclusion of a prior-induced drift was necessary to produce  
170 late DV values that were greater than the early DV values, as seen in the physiological data Fig. (3f, right). Together, these  
171 results suggest that the subjects interpreted neutral sensory evidence in a biased fashion, thus further entrenching their prior  
172 expectations.

## 173 **Relating the dynamic range of neural trajectories to decision bias and sensory uncertainty**

174 Our analysis so far suggests that the decision-related component of neural population activity in PFC resembles a bounded  
175 accumulation process in which prior expectations influence the DV both before and during accumulation of sensory evidence.  
176 This hypothesis is in part motivated by the neural DV's temporal evolution during presentation of ambiguous sensory stimuli  
177 (Fig. 3). To test its explanatory power, we first asked whether such a process can capture the behavioral effects of prior  
178 expectation for ambiguous as well as non-ambiguous stimuli. To this end, we extended the drift-diffusion model simulation to  
179 non-ambiguous stimulus conditions (see Methods; Extended Data Fig. 2a). In our simulation, stimulus orientation determines  
180 the mean of the momentary sensory evidence distribution, while stimulus contrast governs its spread. We further assumed  
181 that prior expectations influence the DV's initial value and its stimulus-dependent drift in exactly the same manner across all  
182 stimulus conditions. A decision is reached when a bound is crossed, or when time is up. In the latter case, the sign of the DV  
183 determines choice outcome. Consider one specific instantiation of this process (Fig. 4a). As can be seen, prior expectations  
184 produce a horizontal shift of the psychometric functions, just like we observed in the animals' behavior (compare with Fig. 1g).  
185 Moreover, lowering stimulus contrast makes the psychometric function more shallow and increases the decision bias (Fig. 4a,  
186 light vs dark lines), again recapitulating the animals' behavior. Varying the model parameters can change the magnitude of the  
187 decision bias and the steepness of the psychometric functions, but it does not alter this basic pattern (Extended Data Fig. 2b).

188 The proposed decision-making process gives rise to a relationship between the dynamic range of the DV trajectory and decision  
189 bias. Specifically, simulated trials in which the eventual choice aligns with the prior expectation tend to have a smaller dynamic  
190 range (defined as the difference between the simulated DV's peak value and its initial value), while trials in which the choice  
191 deviates from the prior expectation tend to have a larger dynamic range (Fig. 4b). This relationship arises naturally if a prior  
192 expectation influences the DV's initial value while a fixed boundary terminates the deliberation process (or conversely, if a prior  
193 expectation influences the bounds without impacting the initial value). To examine the relationship between dynamic range and  
194 decision bias, we computed the distribution of the DV's dynamic range across all trials within a simulated experiment and  
195 separately analyzed the choice behavior for the trials whose dynamic range was below and above the median (Fig. 4c). We  
196 observed that the effect of the prior differs substantially across the low and high dynamic range trials in the simulated dataset.  
197 Specifically, when the DV's dynamic range is low, the decision bias is amplified and the psychometric function is more shallow  
198 (Fig. 4c, light vs dark lines). This pattern of results is specific to a decision-making process in which the prior changes the  
199 distance between the DV's starting point and the bounds. Alternative scenarios in which the prior only induces a drift during  
200 the accumulation of sensory evidence, or in which there is no terminating boundary, do not produce a relationship between  
201 dynamic range and decision bias (Extended Data Fig. 3).

202 To test whether these simulated effects are present in the real data, we conducted the same median-split analysis on the neural  
203 DVs, using the model DV-trajectories to estimate each trial's dynamic range (Methods). We discovered the same relationship  
204 between dynamic range and the animals' decision bias. To quantify this relationship, we computed the difference between the  
205 decision bias for low and high dynamic range trials,  $\Delta B = \Delta\beta_L - \Delta\beta_H$  (example shown in Fig. 4d). Across all experiments,  
206 the mean value of  $\Delta B$  was 0.45 ( $P < 0.001$ , Wilcoxon signed rank test; Fig. 4e, left). Considering the data at a more granular



207 level (per monkey and per contrast level) reveals the robustness of these results, though not every individual case reached  
208 statistical significance (Fig. 4e, left). We also measured the relationship between dynamic range and the steepness of the  
209 psychometric function, quantified as  $\Delta S = \Delta\sigma_L - \Delta\sigma_H$ , and again found a pattern of results consistent with the model  
210 simulation. Lower dynamic range trials tended to be associated with a shallower psychometric function (mean value of  $\Delta S$   
211 was 0.09,  $P < 0.01$ ; Fig. 4e, right). For some datasets, these effects were substantial, for others they were small. Interestingly,  
212 the effects tended to be larger for datasets where the model DV-trajectories' peak values better predicted the animal's choice  
213 behavior (Fig. 4f). To summarize these results, we computed the correlation between  $\Delta B$  and  $\Delta S$  on the one hand, and the  
214 previously introduced measure for model prediction error (AIC) on the other. Across all experiments, the correlation with  $\Delta B$   
215 was 0.62 ( $P < 0.001$ ; Fig. 4g, left panel), while the correlation with  $\Delta S$  was 0.54 ( $P < 0.001$ ; Fig. 4g, right panel). Thus,  
216 the stronger the neural correlate of the decision-making process, the more the empirical data resemble the predictions of an  
217 idealized model of how prior expectations shape this process.

218 We hypothesize that the neural DV's dynamic range primarily reflects the strength of the prior expectation. Under a bounded  
219 accumulation process with a prior-induced initial offset and drift, a prior expectation reduces the dynamic range of congruent  
220 choices, but has the opposite effect on incongruent choices. In this manner, a prior expectation simultaneously amplifies the  
221 decision bias and shallows the psychometric function. Varying the strength of the prior expectation under our imagined scenario  
222 impacts both summary statistics of the observable behavior, as is evident from simulations of this process (Fig. 5a). We found  
223 that this co-occurrence was also present in our empirical observations (Fig. 5b). Datasets that exhibited a stronger association  
224 between dynamic range and decision bias (as measured by  $\Delta B$ ) also exhibited a stronger association between dynamic range  
225 and slope of the psychometric functions (correlation between  $\Delta B$  and  $\Delta S$ : 0.63,  $P < 0.001$ ; Fig. 5c). This relationship  
226 remained significant after controlling for the contribution of model prediction error (AIC), a confounding variable that could  
227 have inflated this relationship ( $F = 10.18$ ,  $P = 0.0025$ ; ANCOVA with 4 levels of prediction error). Together, these results  
228 support the conclusion that shorter excursions of the neural DV yield more biased, less sensitive perceptual decisions.

## 229 Discussion

230 In this study, we investigated neural population activity in the prefrontal cortex of macaque monkeys while they judged am-  
231 biguous visual stimuli under different prior stimulus statistics. We sought to understand how decision-making circuits combine  
232 sensory signals and prior knowledge to form a perceptual decision. We used a task that requires flexible reporting of percep-  
233 tual decisions and discovered that prior expectations exert numerous influences on neural population activity during decision  
234 formation. In our task, the neural correlate of decision formation resembles a tug-of-war style competition between candidate  
235 perceptual interpretations, independent of the ensuing motor response used to communicate the decision [3]. Prior expectations  
236 shift the starting point of the competition towards the more probable stimulus interpretation (Fig. 3a-c). Additionally, prior  
237 expectations selectively change the impact of each piece of sensory evidence on the competition, favoring the evidence that is  
238 congruent with the expectation (Fig. 3d-f). Finally, prior expectations shorten the distance between the starting point and end  
239 point of the decision formation trajectory (Fig. 4).

240 Several previous studies also used biased perceptually-driven behavior as a gateway to study the neural implementation of  
241 Bayesian computation in decision making areas [18, 19, 20, 21, 22, 36]. These studies revealed that expectations can modulate  
242 the activity of single neurons before stimulus onset [19, 18], while the stimulus is presented [20, 18, 21, 22], and while the  
243 behaviour is produced [21, 22]. However, due to the experimental design, it is impossible to say whether these effects reflect  
244 a neural correlate of perceptual inference or of motor planning. Here, we built on this work by using a task in which the  
245 perceptual and motor components of the decision process are orthogonalized (Fig. 1a). This enabled us to isolate the component  
246 of neural activity that solely pertains to the formation of a perceptual choice. Our analysis revealed population level effects that  
247 echo elements of the aforementioned findings but that are clearly situated in the domain of perceptual inference. Specifically,  
248 we showed that neurons in the prearcuate gyrus can collectively represent abstract perceptual expectations before stimulus  
249 onset and confirmation biases during the evaluation of sensory evidence. These results suggest that neural mechanisms that  
250 perhaps originally evolved to select the best candidate action [37] over time acquired sufficient flexibility such that they can be  
251 repurposed to select the best candidate interpretation of the state of the environment in light of incoming sensory signals and  
252 prior experience.

253 Figurative language used to describe biased reasoning often invokes inertia. A biased person can be said to be entrenched in a  
254 position, immovable by evidence. In prefrontal cortex, we found a literal manifestation of the figurative relation between bias  
255 and inertia. When the neurally decoded decision variable covered less distance, monkeys gave a more biased account of the state  
256 of the sensory environment. This finding establishes a direct link between Bayesian computation and the dynamic evolution  
257 of neural activity during decision-formation. We suggest that this link exists because the brain's implementation of Bayesian  
258 computation in our task resembles an approximate hierarchical inference process in which evidence is integrated until a bound

is reached. It may seem surprising that estimating a static aspect of the environment involves temporal integration. However, noise in neural representations limits the signal-to-noise ratio of momentary sensory messages communicated by sensory cortex [30, 38, 39], and thus the quality of perceptual estimates informed by these messages [40]. Temporal integration averages out some of this noise, thereby improving the quality of perceptual estimates [39, 41, 42, 43, 44]. Under an approximate hierarchical inference process, prior expectations change the initial starting point and drift of the accumulator [35]. Using model simulations, we found that this process explains several key features of our data. First, it results in decisions that are biased for ambiguous as well as non-ambiguous stimulus orientations. Second, lowering stimulus contrast jointly reduces the slope of the psychometric function and increases the decision bias. Third, under this process, shorter excursions of the decision variable are associated with more biased and less sensitive decisions. We speculate that the neural implementation of Bayesian computation we have revealed is sufficiently general to support a wide range of perceptual and cognitive estimation tasks.

In this work, we provided a detailed description of the temporal evolution of population activity in prefrontal cortex during perceptual inference and we identified a principled computational process that accounts for these observations. It is natural to ask how neural circuits realize this computational logic. Specifically, how can hardwired circuits provided with a response mapping cue, a stimulus prior cue, and a snippet of noisy sensory measurements realize approximate hierarchical inference in a representational space that is orthogonal to the action selection space? We speculate that a recurrent neural network organized as an attractor network may be a fruitful starting place to address this question [45, 46, 47, 48]. Our analysis has provided a rich set of empirical constraints that should prove helpful to distinguish among candidate network organizations and training regimes and thus represents an important step towards uncovering the neurobiological basis of perception and cognition.

## References

- [1] Nikos K Logothetis and Jeffrey D Schall. Neural correlates of subjective visual perception. *Science*, 245(4919):761–763, 1989.
- [2] Michael N Shadlen and William T Newsome. Neural basis of a perceptual decision in the parietal cortex (area lip) of the rhesus monkey. *Journal of neurophysiology*, 86(4):1916–1936, 2001.
- [3] Julie A Charlton and Robbe L T Goris. Abstract deliberation by visuomotor neurons in prefrontal cortex. *Nature Neuroscience*, 27:1167–1175, 2024.
- [4] Pascal Mamassian and Ross Goutcher. Prior knowledge on the illumination position. *Cognition*, 81(1):B1–B9, 2001.
- [5] Wilson S. Geisler, J. S. Perry, Boaz J. Super, and D. F. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. 2001.
- [6] Y. Weiss, E. P. Simoncelli, and E. H. Adelson. Motion illusions as optimal percepts. *Nature Neuroscience*, 5:598–604, 2002.
- [7] Alan A Stocker and Eero P Simoncelli. Noise characteristics and prior expectations in human visual speed perception. *Nature neuroscience*, 9(4):578–585, 2006.
- [8] Wendy J. Adams, Erich W. Graf, and Marc O. Ernst. Experience can change the 'light-from-above' prior. *Nature Neuroscience*, 7:1057–1058, 2004.
- [9] David C. Knill. Learning Bayesian priors for depth perception. *Journal of Vision*, 7(8):13–13, 06 2007.
- [10] Mehrdad Jazayeri and Michael N Shadlen. Temporal context calibrates interval timing. *Nature neuroscience*, 13(8):1020–1026, 2010.
- [11] Julie A Charlton, Wiktor F Młynarski, Yoon H Bai, Ann M Hermundstad, and Robbe LT Goris. Environmental dynamics shape perceptual decision bias. *PLOS Computational Biology*, 19:e1011104, 2023.
- [12] H. von Helmholtz. *Handbuch der physiologischen Optik*, volume III. Leopold Voss, 1867.
- [13] D. Knill and A. Pouget. The bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12):712–719, dec 2004.
- [14] A. R. Girshick, Landy M. S., and E. P. Simoncelli. Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14:926–932, 2011.
- [15] Christopher Summerfield and Floris P De Lange. Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, 15(11):745–756, 2014.
- [16] Wiktor F Młynarski and Ann M. Hermundstad. Adaptive coding for dynamic sensory inference. *Elife*, 7:e32055, 2018.
- [17] Michael Hahn and Xue-Xin Wei. A unifying theory explains seemingly contradictory biases in perceptual estimation. *Nature Neuroscience*, 27(4):793–804, 2024.
- [18] Vinod Rao, Gregory C DeAngelis, and Lawrence H Snyder. Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *Journal of Neuroscience*, 32(29):10063–10074, 2012.
- [19] Michele A Basso and Robert H Wurtz. Modulation of neuronal activity in superior colliculus by changes in target probability. *Journal of Neuroscience*, 18(18):7519–7534, 1998.

- 339 [20] Timothy D. Hanks, Mark E. Mazurek, Roozbeh Kiani,  
340 Elisabeth Hopp, and Michael N. Shadlen. Elapsed de-  
341cision time affects the weighting of prior probability in  
342a perceptual decision task. *The journal of neuroscience*,  
343 31(17):6339–6352, 2011. 390
- 344 [21] Timothy R Darlington, Jeffrey M Beck, and Stephen G  
345 Lisberger. Neural implementation of bayesian infer-  
346ence in a sensorimotor behavior. *Nature Neuroscience*,  
347 21(10):1442–1451, 2018. 391
- 348 [22] Hansem Sohn, Devika Narain, Nicolas Meirhaeghe, and  
349 Mehrdad Jazayeri. Bayesian computation through cortical  
350latent dynamics. *Neuron*, 103(5):934–947, 2019. 397
- 351 [23] Vincent P Ferrera, Marianna Yanike, and Cassanello  
352Carlos. Frontal eye field neurons signal changes in deci-  
353sion criteria. *Nature Neuroscience*, 12:1458–1462, 2009. 398
- 354 [24] Paul Cisek and John F Kalaska. Neural mechanisms for  
355interacting with a world full of action choices. *Annual*  
356*review of neuroscience*, 33:269–298, 2010. 403
- 357 [25] Jeffrey D Schall. Visuomotor areas of the frontal  
358lobe. In *Extrastriate cortex in primates*, pages 527–638.  
359Springer, 1997. 404
- 360 [26] Earl K Miller and Jonathan D Cohen. An integrative  
361theory of prefrontal cortex function. *Annual review of*  
362*neuroscience*, 24(1):167–202, 2001. 405
- 363 [27] Roozbeh Kiani, Christopher J Cueva, John B Reppas,  
364and William T Newsome. Dynamics of neural popula-  
365tion responses in prefrontal cortex indicate changes of  
366mind on single trials. *Current Biology*, 24(13):1542–  
3671547, 2014. 410
- 368 [28] Valerio Mante, David Sussillo, Krishna V Shenoy, and  
369William T Newsome. Context-dependent computation  
370by recurrent dynamics in prefrontal cortex. *nature*,  
371 503(7474):78–84, 2013. 415
- 372 [29] W. P. Tanner and J. A. Swets. A decision-making theory  
373of visual detection. *Psychological Review*, 61:401–409,  
3741954. 419
- 375 [30] P Heggelund and Albus K. Response variability and ori-  
376entation discrimination of single cells in striate cortex of  
377cat. *Exp Brain Res.*, 32(2):197–211, 1978. 422
- 378 [31] Michael N. Shadlen, KH Britten, WT Newsome, and  
379JA Movshon. A Computational Analysis of the Rela-  
380tionship between Neuronal and Behavioral Responses to  
381Visual Motion. *Journal of Neuroscience*, 76(4):1486–  
3821510, 1996. 423
- 383 [32] Robbe L T Goris, J. Anthony Movshon, and Eero P. Si-  
384moncelli. Partitioning neuronal variability. *Nature Neu-*  
385*roscience*, 17(6):858–865, 2014. 425
- 493 [33] Diogo Peixoto, Jessica R Verhein, Roozbeh Kiani,  
494Jonathan C Kao, Paul Nuyujukian, Chandramouli Chan-  
495drasekaran, Julian Brown, Sania Fong, Stephen I Ryu,  
496Krishna V Shenoy, and William T Newsome. Decod-  
497ing and perturbing decision states in real time. *Nature*,  
498 591(7851):604–609, 2021. 426
- 499 [34] Gabe Stine, Eric M Trautmann, Danique Jeurissen, and  
500Michael N Shadlen. A neural mechanism for terminating  
501decisions. *Neuron*, 2023. 427
- 502 [35] Richard Lange, Ankani Chattoraj, Jeffrey Beck, Jacob  
503Yates, and Ralf M Haefner. A confirmation bias in per-  
504ceptual decision-making due to hierarchical approximate  
505inference. *PLoS Comput Biol*, 17(11), 2021. 428
- 506 [36] Trinity B Crapse, Hakwan Lau, and Michele A Basso. A  
507role for the superior colliculus in decision criteria. *Neu-*  
508*ron*, 97(1):181–194, 2018. 429
- 509 [37] Paul Cisek. Cortical mechanisms of action selection:  
510the affordance competition hypothesis. *Philosophical*  
511*Transactions of the Royal Society B: Biological Sci-*  
512*ences*, 362(1485):1585–1599, 2007. 430
- 513 [38] David J Tolhurst, J Anthony Movshon, and Andrew F  
514Dean. The statistical reliability of signals in single neu-  
515rons in cat and monkey visual cortex. *Vision research*,  
516 23(8):775–785, 1983. 431
- 517 [39] Rufin Vogels and G A Orban. How well do response  
518changes of striate neurons signal differences in orienta-  
519tion: A study in the discriminating monkey. *Journal of*  
520*Neuroscience*, 10(11):3543–3558, 1990. 432
- 521 [40] Zoe Boundy-Singer, Corey M Ziemba, O J H Henaff,  
522and Robbe L T Goris. How does v1 population activity  
523inform perceptual certainty? *Journal of Vision*, 24(6):1–  
52417, 2024. 433
- 525 [41] Takanori Uka and Gregory C DeAngelis. Contribution  
526of middle temporal area to coarse depth discrimination:  
527comparison of neuronal and psychophysical sensitivity.  
528*Journal of Neuroscience*, 23(8):3515–3530, 2003. 434
- 529 [42] Leslie C Osborne, William Bialek, and Stephen G Lis-  
530berger. Time course of information about motion direc-  
531tion in visual area mt of macaque monkeys. *Journal of*  
532*Neuroscience*, 24(13):3210–3222, 2004. 435
- 533 [43] Yuzi Chen, Wilson Geisler, and Eyal Seidemann. Op-  
534timal temporal decoding of neural population responses  
535in a reaction-time visual detection task. *Journal of Neu-*  
536*rophysiology*, 99(3):1366–1379, 2008. 436
- 537 [44] Robbe LT Goris, Corey M Ziemba, J Anthony Movshon,  
538and Eero P Simoncelli. Slow gain fluctuations limit ben-  
539efits of temporal integration in visual cortex. *Journal of*  
540*vision*, 18(8):8–8, 2018. 437

- 434 [45] Thomas Zhihao Luo\*, Timothy Doyeon Kim\*, Diksha  
435 Gupta, Adrian G Bondy, Charles D Kopec, Verity A Eliot, Brian DePasquale, and Carlos D Brody. Transitions  
436 in dynamical regime and neural mode underlie perceptual decision-making. *Biorxiv preprint*, 2024. 447
- 439 [46] Sebastian Bitzer, Jelle Bruineberg, and Stefan J Kiebel.  
440 A bayesian attractor model for perceptual decision making. *PLoS Computational Biology*, 11(8), 2015. 448
- 442 [47] Alexis Dubreuil, Adrian Valente, Manuel Beiran, Francesca Mastrogiuseppe, and Srdjan Ostojic. The role  
443 of population structure in computations through neural dynamics. *Nature Neuroscience*, pages 1–12, 2022.
- [48] Mikhail Khona and Ila R Fiete. Attractor and integrator networks in the brain. *Nature Reviews Neuroscience*, 23:744–766, 2022.
- [49] Shintaro Funahashi, Charles J Bruce, and Patricia S Goldman-Rakic. Mnemonic coding of visual space in the monkey’s dorsolateral prefrontal cortex. *Journal of neurophysiology*, 61(2):331–349, 1989.



## 453 METHODS

### 454 Experimental methods and behavioral analysis

455 The experimental methods are described in detail in ref. [3]. In brief, the experiments were performed on two adult male  
456 macaque monkeys (*Macaca mulatta*). The animals were trained to perform a memory-guided saccade task and an orientation  
457 discrimination task. All training, surgery, and recording procedures conformed to the National Institute of Health Guide for the  
458 Care and Use of Laboratory Animals and were approved by The University of Texas at Austin Institutional Animal Care and  
459 Use Committee. We first used a variation of the classical memory-guided saccade task [49] to identify recording sites where  
460 neurons exhibited neural activity indicative of an upcoming eye movement. Following identification of a suitable recording  
461 site (Extended Data Fig 1), we conducted several additional orientation discrimination training sessions with one choice target  
462 placed within the estimated response field location and one on the opposite site of the fixation mark. Once psychophysical  
463 performance reached a high level, physiological data collection began. The orientation-discrimination task involved two distinct  
464 prior contexts, associated with differently skewed distributions of stimulus orientation (see Fig. 1c). Blocks of both priors  
465 alternated randomly (80 trials per block). Stimuli were seven drifting gratings evenly spaced over a small range of orientation,  
466 tailored to each monkey's orientation sensitivity. Vertically oriented stimuli received random feedback. Stimuli were presented  
467 at either high or low contrast (Michelson contrast: 100% or 4%). Blocks of high and low contrast stimuli alternated randomly  
468 (trials per block: monkey F = 100, monkey J = 80). We measured observers' choice bias for ambiguous stimuli by computing  
469 the log odds of a "clockwise" choice under each prior context (Fig. 1d). We measured observers' behavioral capability to  
470 discriminate stimulus orientation by fitting the relationship between stimulus orientation and probability of a "clockwise"  
471 choice with a psychometric function consisting of a lapse rate and a cumulative Gaussian function. Model parameters were  
472 optimized by maximizing the likelihood over the observed data, assuming responses arise from a Bernoulli process. Each  
473 recording session was analyzed independently. For the analysis documented in Fig. 1h, we fit one psychometric function per  
474 stimulus prior and contrast level. Both prior conditions shared the same sensitivity parameter, resulting in two psychometric  
475 functions with identical slope. We defined decision bias as the difference between the means of both cumulative Gaussians (i.e.,  
476 the magnitude of the horizontal displacement of both psychometric functions). Error bars of model-based statistics are based  
477 on a 100-fold non-parametric bootstrap of the behavioral data.

### 478 Simulated observer models

479 We investigated the choice behavior under an ideal Bayesian and a Maximum Likelihood inference strategy (Fig. 1e,f). We used  
480 two context-specific stimulus distributions that matched those used in the animal experiments. Each trial, the model observers  
481 were presented with a noisy sensory measurement, with the noise modeled as a sample from a zero mean Gaussian distribution.  
482 This sensory measurement informed the likelihood function, computed as a Gaussian probability density evaluated at all possible  
483 stimulus values, using a standard deviation that matched the strength of the sensory noise. For the Maximum Likelihood  
484 inference strategy, we selected the mode of the likelihood function as stimulus estimate. For the Ideal Bayesian inference strategy,  
485 we first used Bayes' rule to compute the posterior probability over all possible stimulus values. The perceptual decision  
486 reflected whether a clockwise stimulus orientation was the most likely interpretation (meaning that the cumulative conditional  
487 posterior probability exceeded 50%).

### 488 Electrophysiological recordings and decision variable analysis

489 As described in ref. [3], we conducted 13 successful recordings from monkey F and 16 from monkey J, using linear electrode  
490 arrays (average number of trials per session, monkey J = 3,171; monkey F = 1,593). We positioned the linear arrays so that  
491 they roughly spanned the cortical sheet and removed them after each recording session. To extract responses of individual  
492 units, we performed offline spike sorting. Given that the electrode's position could not be optimized for all contact sites, most  
493 of our units likely consist of multi-neuron clusters. All units whose mean firing rate during the task exceeded 3 ips were  
494 included in the analysis. For each trial, we obtained moment-to-moment measurements of the decision variable by projecting  
495 50 ms bins of population activity onto a linear decoder optimized to distinguish the activity patterns associated with both choice  
496 options ("left" vs "right" choices for the motor DV, and "clockwise" vs "counterclockwise" choices for the categorical DV,  
497 respectively). Specifically, we first individually z-scored each unit's spike counts within every time bin. We then used these  
498 z-scored responses to estimate the set of linear weights,  $\mathbf{w} = (w_1, \dots, w_n)$ , that best separate the choice-conditioned z-scored  
499 response patterns, assuming a multivariate Gaussian response distribution:

$$\mathbf{w} = \frac{\mathbf{s}}{\Sigma} \quad (1)$$

500 where  $\mathbf{s}$  is the mean difference of the choice-conditioned z-scored responses and  $\Sigma$  is the covariance matrix of the z-scored

501 responses. The decoder weights are calculated from observed trials. To avoid double-dipping, we excluded the trial under  
502 consideration from the calculation and solely used all other trials to estimate the weights. This way, we obtained “cross-  
503 validated” DV estimates for each time bin:

$$DV_j = \sum w_{ij} Z_{ij}, \quad (2)$$

504 where  $w_{ij}$  and  $Z_{ij}$  are the weight and z-scored response of unit  $i$  on trial  $j$  for a given time bin. The symbols in Fig. 2a show  
505 example single trial DV trajectories.

506 For each trial, we fit a model with smoothly evolving DV-trajectories to the raw DV values (lines in Fig. 2a). This model has  
507 five free parameters: one captures an initial offset in the DV, one specifies the dynamic range of the DV trajectory, one controls  
508 the speed of the rise, one the time point at which half of the rise is completed, and one captures the decay in strength that follows  
509 the peak of the trajectory. The rising part of the trajectory follows a cumulative Gaussian profile, while the decay follows an  
510 exponential profile that begins at the time at which the cumulative Gaussian reaches the 99.38th percentile. We fit this model to  
511 the data by minimizing the sum of the squared error of each trial’s DV trajectory (correlations between model trajectory and raw  
512 DV trajectory are shown in Fig. 2b). We used these model trajectories to estimate, among other things, the peak DV value for  
513 each trial. To evaluate the usefulness of these estimates, we conducted a logistic regression analysis in which we predicted the  
514 choice behavior using two different sets of regressors. The first set comprised all experimentally controlled variables (stimulus  
515 prior, orientation, and contrast), the second set additionally included the peak DV value estimate. To compare the goodness of  
516 fit of both sets of regressors, we computed an estimator of prediction error based on information theory (Akaike’s information  
517 criterion, hereafter AIC). Specifically, assuming that the residuals under each model are distributed according to independent  
518 identical normal distributions:

$$AIC = 2k + n \ln(\hat{\sigma}^2) - 2C, \quad (3)$$

519 where  $k$  is the number of free parameters,  $n$  the number of data points,  $C$  a constant that only depends on the data, and  $\hat{\sigma}^2$  the  
520 maximum likelihood estimate for the variance of a model’s residuals distribution given by the residual sum of squares divided  
521 by the degree of freedom. Because only differences in AIC are meaningful, the constant  $C$  can be ignored when comparing  
522 models, yielding a statistic known as  $\Delta AIC$  (shown in Fig. 2d).

523 We examined how the average peak DV value depended on the stimulus prior, orientation, and contrast (Fig. 2e-g). For each  
524 recording session, we quantified the effect of orientation with the slope of a linear regression line. We allowed these lines to  
525 have a prior-specific intercept, and conducted this analysis separately for the high and low contrast trials (example shown in  
526 Fig. 2e).

527 We studied the average temporal evolution of the DV trajectories for the zero-signal stimulus condition (Fig. 3a,b). To this  
528 end, we plotted the average DV around the putative time of choice commitment (–500 to –300 ms) computed from the model  
529 trajectory against the average value right before the onset of stimulus-driven activity (–800 to –600 ms). To quantify the growing  
530 influence of the prior expectation over the course of a trial, we calculated the first principle component of the resulting scatter  
531 plot ( Fig. 3d). We evaluated the significance of the PC’s gradient by conducting a 1000-fold bootstrap-analysis and computing  
532 the 99.9% confidence interval from the resulting gradient estimate distribution.

### 533 **Simulated bounded accumulation process**

534 We investigated the dynamic evolution of the decision variable in a bounded accumulation process and considered various  
535 scenarios. In all cases, the variable that was integrated was composed of one term representing the momentary sensory evidence,  
536 and one term representing the prior expectation. We modeled the momentary sensory evidence,  $s(t)$ , at each time point  $t$  as  
537 a sample from a Gaussian distribution  $s(t) \sim \mathcal{N}(\mu_\theta, \sigma)$ , with  $\mu_\theta$  proportional to the strength of the evidence (capturing the  
538 effects of stimulus orientation), and  $\sigma$  inversely proportional to the reliability of the evidence (capturing the effects of stimulus  
539 contrast). We reserved positive values to represent evidence for a clockwise stimulus orientation, and negative values evidence  
540 for a counterclockwise orientation. We modeled the prior expectation as a time-varying signal,  $p(t)$ , that either only had energy  
541 at the first time point (capturing a pure stimulus expectation), or decayed to a lower but sustained level for the remaining time  
542 points (capturing an additional biased interpretation of sensory evidence). In addition, we introduced trial-by-trial variability  
543 to  $p(t)$  by multiplying this term with a scaling factor sampled from a log-normal distribution with a mean of one. Each trial  
544 consisted of 200 time steps, and each simulation consisted of 3,000 trials. When the simulated process contained decision  
545 bounds, the deliberation process was typically terminated early (at the time when one of the bounds was crossed). Fig. 3f  
546 illustrates an analysis in which we only included zero-signal stimuli and varied the strength of the prior-induced effects and

547 the height of the bounds across simulations. Fig. 4a shows an analysis of the choice behavior for a simulation that included  
548 all of our experimental conditions as well as the effects of stimulus expectation, biased interpretation of sensory evidence, and  
549 terminating decision bounds.

## 550 **Dynamic range analysis**

551 We investigated the relationship between the dynamic range of the DV excursion and behavioral signatures of decision bias  
552 and perceptual uncertainty. For each trial, we first computed the dynamic range of the DV by subtracting the DV's initial  
553 value from the peak value. We then used a median split to divide all low-contrast (or high-contrast) trials of a single recording  
554 session (or model simulation) into a low- and high dynamic range group. Finally, we fit two sets of psychometric functions to  
555 the behavioral choices (one psychometric function per prior context, and one set of functions per dynamic range group). This  
556 yielded a total of four psychometric functions per session and contrast level (example for model simulation shown in Fig.4c,  
557 and for real data in Fig.4d).

558 To quantify the relationship between behavioral decision bias ( $\Delta\beta$ ) and the dynamic range of the neural DV, we used a metric,  
559  $\Delta B$ , defined as the surplus in prior-induced decision bias observed for the low dynamic range trials relative to the high dynamic  
560 range trials:

$$\Delta B = \Delta\beta_L - \Delta\beta_H, \quad (4)$$

561 where  $\Delta\beta_L$  is the prior-induced decision bias for the low dynamic range trials (specifically, the horizontal separation between  
562 the pair of psychometric functions fit to both prior contexts).

563 We devised an analogous metric,  $\Delta U$ , to measure the association between perceptual uncertainty and DV dynamic range:

$$\Delta U = \Delta v_L - \Delta v_H, \quad (5)$$

564 where  $\Delta v_H$  is the standard deviation of the cumulative Gaussian function that relates the proportion of clockwise choices to  
565 stimulus orientation for the high dynamic range trials.

## 566 **Acknowledgements**

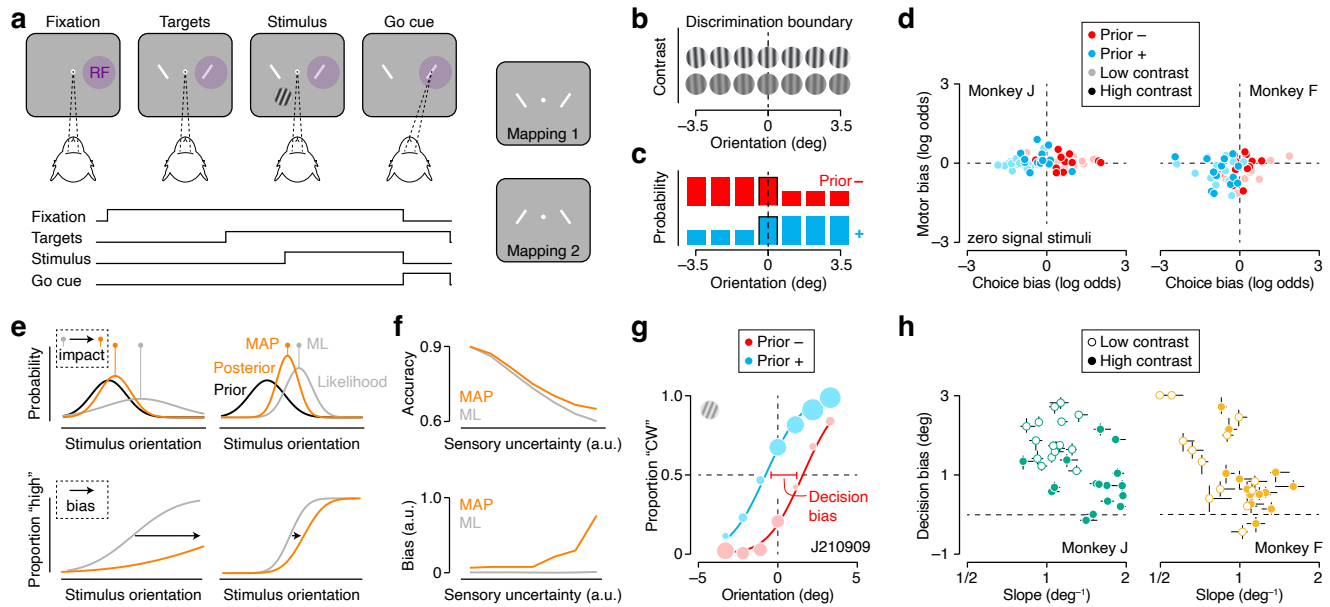
567 This work was supported by US National Institutes of Health grants T32 EY021462 (T.L., J.A.C.), EY032999 (R.L.T.G.), and  
568 National Science Foundation CAREER award #2146369 (R.L.T.G.).

## 569 **Author contributions**

570 J.A.C. and R.L.T.G. conceived and designed the study. J.A.C. collected the data. T.L., J.A.C. and R.L.T.G. analyzed the data.  
571 T.L., J.A.C. and R.L.T.G. wrote the manuscript.

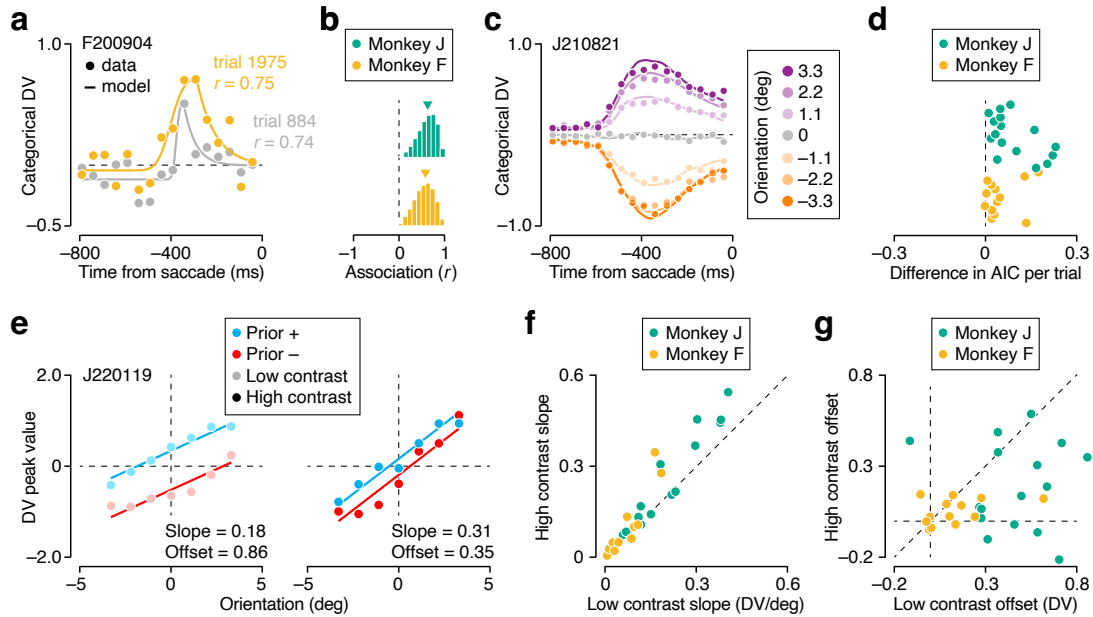
## 572 **Competing Interests**

573 The authors declare no competing interests.

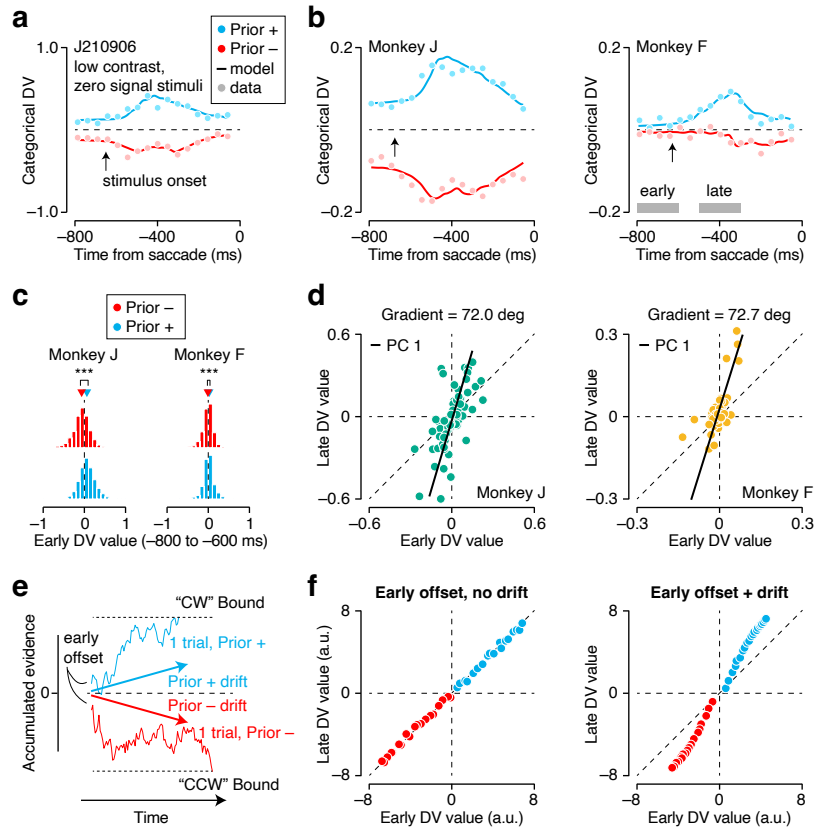


**Figure 1** Flexible orientation discrimination under different prior statistics: task and behavior. **(a)** Orientation discrimination task, task sequence. Each trial begins when the observer initiates fixation. The shape of the fixation mark indicates the prior distribution from which the stimulus will be drawn. After the observer fixates for 500 ms, two choice targets appear, followed by the stimulus. The observer judges whether the stimulus is rotated clockwise or counterclockwise relative to vertical and communicates this decision with a saccade towards the matching choice target. Correct decisions are followed by a juice reward. One of the choice targets is placed in the neurons' presumed motor response field (see Methods). The spatial organization of the choice targets varies randomly from trial-to-trial, giving rise to two stimulus-response mapping rules. **(b)** Stimuli varied in orientation and contrast. **(c)** Stimuli were drawn from one of two orientation distributions. **(d)** Comparison of choice and motor bias for the randomly rewarded ambiguous stimuli (orientation = 0 deg) under both priors for both monkeys (see Methods). **(e)** An ideal Bayesian decision-maker evaluates the likelihood of every possible state of the sensory environment and multiplies this distribution with the prior probability of encountering each state to obtain a posterior belief function. The posterior informs the decision. More ambiguous sensory measurements yield a broader likelihood function, and ultimately more biased decisions. **(f)** Choice accuracy (top) and bias (bottom) in our task under a maximum likelihood (ML) and maximum posterior probability (MAP) inference strategy (see Methods). The performance benefit conferred by accurate prior knowledge grows with sensory uncertainty (top), as does the magnitude of the decision bias (bottom). **(g)** Psychophysical performance for monkey J in an example recording session. Proportion 'clockwise' choices for low contrast stimuli is shown as a function of stimulus orientation under both priors. Symbol size reflects number of trials (total: 1,875 trials). The curves are fits of a behavioral model (see Methods). **(h)** Decision bias plotted as a function of orientation sensitivity for both monkeys (left: Monkey J, right: Monkey F). Each symbol summarizes data from a single recording session. Closed symbols: high contrast stimuli, open symbols: low contrast stimuli. Error bars reflect the IQR of the estimate.

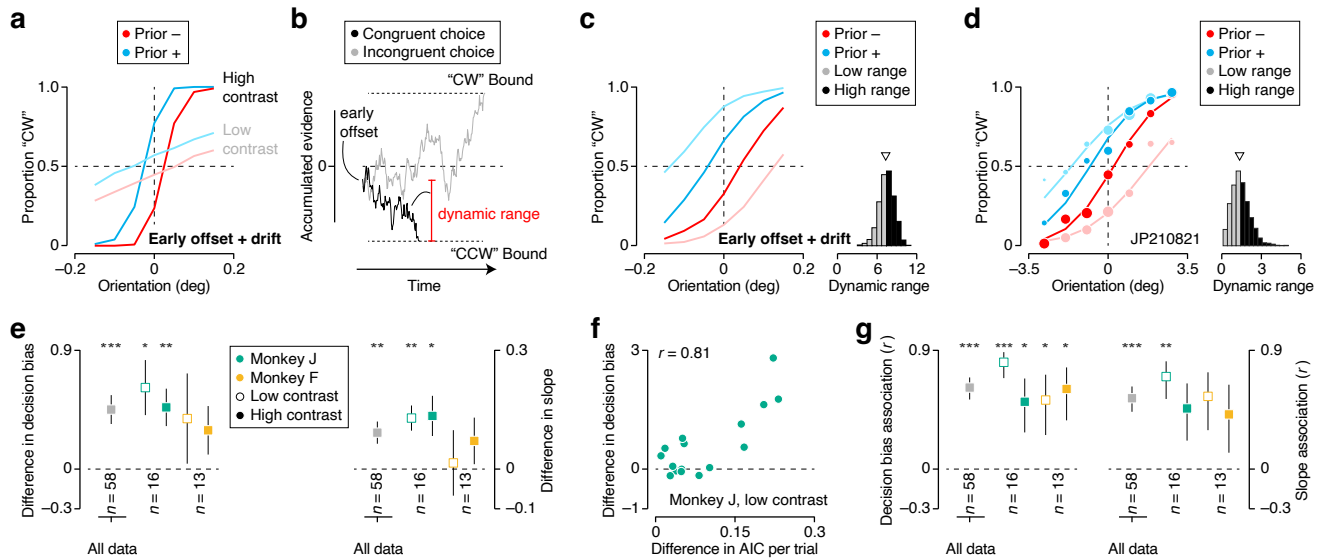





**Figure 2** Extracting neural correlates of Bayesian-like perceptual inference at the single trial level from PFC population activity. **(a)** Temporal evolution of the categorical DV on two different trials from the same recording session with identical experimental parameters and choice behavior (prior = counter-clockwise skew, stimulus orientation = 0 deg, stimulus contrast = low, mapping rule = 1, choice = “clockwise”). Symbols show the raw DV estimates, lines the fit of a model. **(b)** Distribution of the correlation between model DV-trajectories and raw DV estimates across all recording sessions. **(c)** Trial-averaged temporal evolution of the categorical DV, split out by stimulus orientation, for a single recording session. **(d)** Difference in prediction error for two logistic regression analyses of the choice behavior, as measured by Akaike’s Information Criterion. The first analysis used stimulus prior, orientation and contrast as regressors, the second analysis additionally used the model DV-trajectory peak value. Positive values indicate that including the peak DV value improves prediction accuracy. **(e)** The peak value of the categorical DV depends on stimulus orientation (abscissa), stimulus contrast (left vs right panel), and stimulus prior (red vs blue). Symbols show the average value across all trials from a single recording session. Lines show the outcome of a linear regression analysis used to estimate the slope and offset of this relation at low (left panel) and high (right panel) stimulus contrasts. **(f,g)** Comparison of slope **(f)** and offset **(g)** for high vs low contrast stimuli.



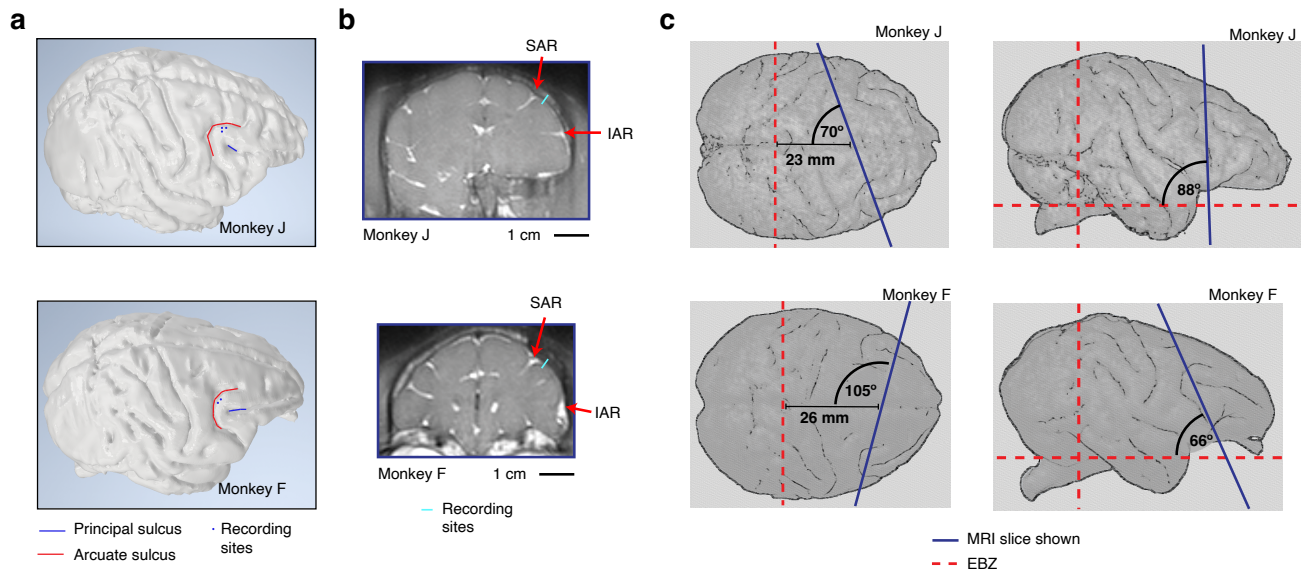
**Figure 3** Prior expectations bias the decision variable before and during stimulus presentation. **(a)** Trial-averaged temporal evolution of the categorical DV, split by stimulus prior, for a single recording session. Only low contrast, zero signal stimuli are included. Symbols show the raw DV estimates, lines the model DV-trajectories. The average time of stimulus onset is indicated by the black arrow. **(b)** The average across all recording sessions, split by monkey (left vs right panel). **(c)** Distribution of the model-estimated early DV value across all recording sessions, split by stimulus prior (blue vs red) and monkey (left vs right). **(d)** Late DV value plotted against early DV value, split by monkey (left vs right). Each symbols summarizes the trial-averaged values of a single recording session (one point per stimulus prior). **(e)** We simulated bounded accumulation of ambiguous sensory evidence under two different prior expectations. In all simulations, the prior changes the starting point of the accumulation process (blue vs red trace, leftmost time point). In some simulations, the prior additionally changes the drift rate of the accumulation process (blue vs red arrow). **(f)** Late DV value plotted against early DV value for the simulations without (left) and with (right) prior-induced drift.



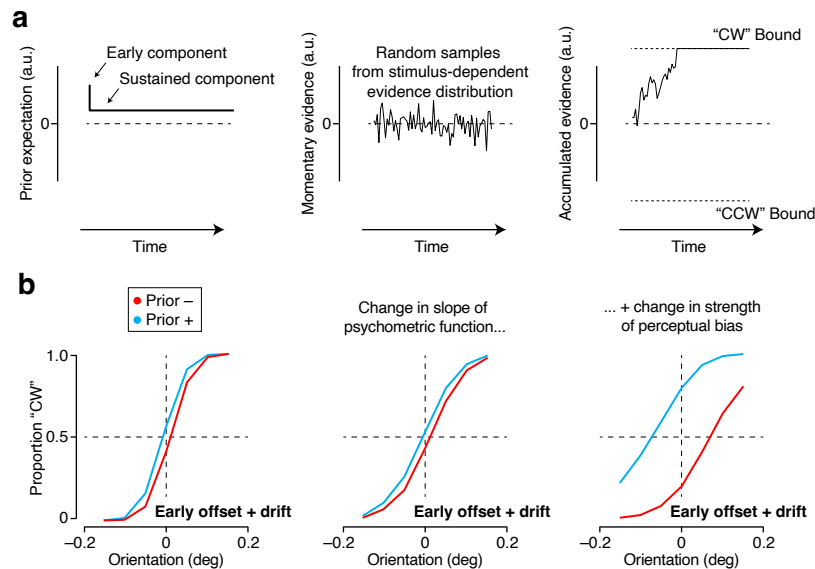
**Figure 4** Shorter neural excursions result in more biased perceptual decisions. **(a)** We simulated bounded accumulation of ambiguous and non-ambiguous sensory evidence under two different prior expectations. In all simulations, the prior influenced the starting point and the drift rate of the accumulation process. Under this process, prior expectations induce a horizontal shift of the psychometric function (blue vs red), that is larger when the sensory evidence is less reliable (dark vs light colors). **(b)** Example of a simulated trial that resulted in a choice that was congruent with the prior expectation (black trace), and a trial that resulted in an incongruent choice (grey trace). The dynamic range of the excursion measures the distance travelled between the beginning and end of the accumulation process. **(c)** Median-split analysis of simulated choice behavior under two different prior expectations (blue vs red). Trials with a smaller dynamic range yield more biased decisions (dark vs light colors). **(d)** Median-split analysis for all low contrast trials in one example recording session. Symbol size reflects number of trials (total: 1,789 trials). **(e)** (Left) The difference in decision bias for low and high dynamic range trials across all data (grey symbols), and split out by monkey and stimulus contrast (non-grey symbols). (Right) The difference in the slope of the psychometric function for low and high dynamic range trials. Error bars show mean  $\pm$  1 s.e.m. **(f)** Difference in decision bias plotted as a function of the difference in Akaike's Information Criterion. **(g)** (Left) Association between the AIC difference and the difference in decision bias. (Right) Association between the AIC difference and the difference in the slope of the psychometric function. Error bars show  $r \pm$  1 s.d. \*  $P < 0.05$ , \*\*  $P < 0.01$ , \*\*\*  $P < 0.001$ .



**Figure 5** Stronger prior expectations yield decisions that are more biased and less sensitive. **(a)** We simulated bounded accumulation of sensory evidence under two different prior expectations and measured the association between the DV's dynamic range and decision bias (abscissa) and between dynamic range and the slope of the psychometric function (ordinate). Each symbol summarizes these measurements for one simulation. We varied the strength of the prior expectation across simulations. **(b)** Difference in slope plotted against difference in decision bias for one monkey. Each symbol summarizes one recording session. Only low contrast trials were included in the analysis. **(c)** Summary of the association between both variables across data (grey symbol), and split out by monkey and stimulus contrast (non-grey symbols). Error bars show  $r \pm$  1 s.d. \*  $P < 0.05$ , \*\*  $P < 0.01$ , \*\*\*  $P < 0.001$ .

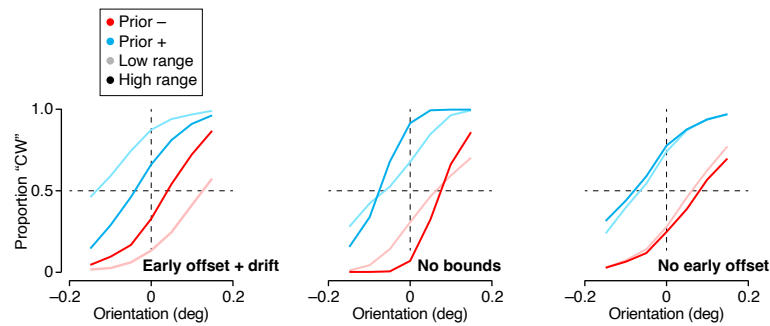


**Extended Data Figure 1** Location of prearcuate gyrus recordings for monkey J and F (top vs bottom). **(a)** 3D reconstruction of the brain based on a structural MRI scan obtained before chamber and post implants. The location of the recording sites are marked by blue dots. **(b)** Structural MRI scan illustrating the approximate recording site. The arrows indicate the superior arcuate sulcus (SAR) and inferior arcuate sulcus (IAR). **(c)** Top and side view (left vs right) of the MRI slice shown in panel **b**.



**Extended Data Figure 2** Simulated bounded accumulation process. **(a)** (left) The simulated bounded accumulation process contained three critical ingredients. First, a prior expectation, modeled as a time-varying signal with an early and late component. (middle) Second, a sensory input, modeled as a time-varying signal composed of random samples drawn from a Gaussian distribution. The mean of the distribution reflects stimulus orientation, the spread of the distribution stimulus contrast. (right) And third, two fixed bounds that terminate the accumulation process when crossed. **(b)** Under this process, prior expectations produce a robust horizontal shift of the psychometric function. The model parameters control the slope of the psychometric function and the magnitude of the decision bias.





**Extended Data Figure 3** The relation between DV dynamic range and decision bias under different model variants. (Left) Median-split analysis of simulated choice behavior under two different prior expectations (blue vs red). The simulated process contained a prior-induced early offset and drift and terminating decision bounds. Trials with a smaller dynamic range yield more biased decisions (dark vs light colors). (Middle) Same analysis for an unbounded evidence accumulation process. Trials with a smaller dynamic range yield less sensitive decisions, as evidenced by the difference in the slope of the dark vs light psychometric functions. However, note that the decision bias, indicated by the horizontal separation of the midpoint of the red and blue psychometric function, does not depend on the dynamic range. (right) Same analysis for a bounded accumulation process in which the prior expectation does not induce an early offset, but only a drift.