

# Bayesian inference by visuomotor neurons in the prefrontal

cortex

Thomas A. Langlois<sup>a,1</sup>, Julie A. Charlton<sup>b,1</sup>, and Robbe L. T. Goris<sup>c,2</sup>

Affiliations are included on p. 10.

Edited by John Maunsell, The University of Chicago, Chicago, IL; received October 9, 2024; accepted February 20, 2025

Perceptual judgments of the environment emerge from the concerted activity of neural populations in decision-making areas downstream of the sensory cortex. When the sensory input is ambiguous, perceptual judgments can be biased by prior expectations shaped by environmental regularities. These effects are examples of Bayesian inference, a reasoning method in which prior knowledge is leveraged to optimize uncertain decisions. 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 isolate 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. We show that these results resemble a specific variant of Bayesian inference known as approximate hierarchical inference. Our findings expand our understanding of the mechanisms by which prefrontal circuits can execute Bayesian inference.

perceptual decision-making | prefrontal cortex | 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 (1, 2). This inference strategy typically manifests as perceptual interpretations that are biased toward the prior expectation. Such biases may reflect implicit knowledge of statistical regularities that are stable features of the environment (3–6), such as the tendencies of sunlight to come from above (4), image velocities to be slow (5, 6), and cardinal orientations to be overrepresented in visual scenes (7). However, biased perceptual interpretations can also reflect knowledge of statistical regularities that are context-specific and short-lived in nature (8–12). The diversity of experimental settings under which prior-induced perceptual biases occur suggests that a general neural mechanism may underlie these effects (13–15).

Prior expectations about sensory stimulation are known to 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 (16), during stimulus presentation (17–20), and while the perceptually driven behavior is being produced (19, 20). At the population level, these effects conspire to bias neural representations toward the prior expectation (20). 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 (16, 18–20). 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. Thus, the biases in perceptual reports coincided with the biases in motor responses. Given that decision-making circuits are typically involved in action planning (21–23), it is unclear 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 (23). Monkeys judged whether a visual stimulus was oriented clockwise or counterclockwise from vertical and communicated their decision with a saccadic eye movement toward one of two visual targets (Fig. 1*A*). The meaning of each response option was signaled by

# Significance

#### Perceptual

judgments of ambiguous stimuli are often biased by prior expectations. These biases may offer a window into the neural computations that give rise to perceptual interpretations of the environment. However, isolating the neural correlates of perceptual decisions is challenging. Here, we studied neural population activity in the prefrontal cortex of macaque monkeys while they judged ambiguous visual stimuli. The task design enabled us to isolate the component of neural activity that solely pertains to the formation of a perceptual choice (the decision variable). We found that prior expectations impact the decision variable both before and during stimulus presentation. Our analysis suggests that the neural computation at work is a variant of Bayesian inference known as approximate hierarchical inference.

Author contributions: T.A.L., J.A.C. and R.L.T.G. designed research; T.A.L., and J.A.C. performed research; T.A.L., J.A.C., and R.L.T.G. analyzed data; T.A.L., J.A.C., and R.L.T.G. developed the theory and the statistical modeling; and T.A.L., J.A.C., and R.L.T.G. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2025 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

Published March 27, 2025.

<sup>&</sup>lt;sup>1</sup>T.L. and J.A.C. contributed equally to this work.

 $<sup>^2 \</sup>mbox{To}$  whom correspondence may be addressed. Email: robbe.goris@utexas.edu.

This article contains supporting information online at https://www.pnas.org/lookup/suppl/doi:10.1073/pnas. 2420815122/-/DCSupplemental.

Downloaded from https://www.pnas.org by MASSACHUSETTS INSTITUTE OF TECHNOLOGY; MIT LIBRARIES on July 4, 2025 from IP address 18.29.215.45.

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, and the other was placed on the opposite side of the fixation mark; see Materials and 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 (SI Appendix, Fig. S1), an area of the prefrontal cortex (PFC) involved in the selection of saccadic eye movements (24, 25) that represents visuomotor deliberation (26, 27). We previously found that the population activity initially represents the formation of a perceptual choice, before transitioning into the representation of the motor plan (23). Here, we study how prior expectations and sensory signals shape the perceptual decisionmaking 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. Our analysis shows that these effects are consistent with a variant of Bayesian inference known as approximate hierarchical inference, and culminate in more biased and less sensitive perceptual decisions.

### Results

Task and Behavior. Two rhesus macaques were trained to perform a visual orientation discrimination task in which they communicated perceptual decisions under two different stimulus-response mapping rules (Fig. 1A). The data were previously described in detail in ref. 23. Task difficulty was controlled by manipulating stimulus orientation and contrast (Fig. 1B). Monkeys received a reward if they selected the appropriately oriented choice target. They performed the task similarly well under both mapping rules (median performance: monkey F = 79.7% correct; monkey J = 79.6% correct; median difference in performance across mapping rules: monkey F = 2.58%, P = 0.002; monkey J = 0.5%, P = 0.57;Wilcoxon signed-rank test). We manipulated the animals' prior expectations by sampling stimulus orientation from one of two skewed distributions (Fig. 1C). On each trial, the shape of the fixation mark signaled the current prior condition. The prior condition and stimulus contrast were switched across short blocks of trials (Materials and Methods). The prior manipulation selectively biased the animals' perceptual judgments of ambiguous stimuli (median difference in choice bias for vertical gratings: monkey  $F = 1.22 \log \text{ odds}, P < 0.001; \text{ monkey } J = 1.49$ log odds, P < 0.001), but not the motor responses used to communicate these decisions (median difference in motor bias: monkey  $F = 0.03 \log \text{ odds}$ , P = 0.31; monkey  $J = 0.03 \log \log 1000$ odds, P = 0.06; Fig. 1D). Together, these results suggest that our paradigm engages brain mechanisms that specifically bias perceptual judgments of ambiguous sensory inputs.

What is the nature of the inference process that underlies these biased judgments of visual stimuli? A prominent hypothesis is that these biases naturally arise under an inference strategy that seeks to make the best possible decision given ambiguous sensory measurements and prior experience (1, 2). This notion is formalized in the framework of Bayesian inference. An ideal Bayesian decision-maker computes the posterior evidence in support of each response option by multiplying the prior probability of each stimulus orientation with the likelihood that a specific stimulus orientation gave rise to the current sensory measurement (Fig. 1 E, Top). As a consequence, perceptual decisions are biased for all stimulus orientations, manifesting as a horizontal shift of the psychometric function (Fig. 1 E, Bottom, gray vs. orange line). The impact of the prior on the decision depends on the ambiguity of the sensory response. The more ambiguous the sensory response, the broader the likelihood function, and the larger the impact of the prior on the posterior (Fig. 1 E, Bottom Left vs. Right panel). A Bayesian inference strategy maximizes choice accuracy, even though it results in decision biases and systematic errors (Fig. 1F, gray vs. orange lines). Does this strategy explain why the monkeys' judgments of ambiguous stimuli were biased? In keeping with this hypothesis, the monkeys' decisions were biased for all stimulus orientations, not just vertical gratings (Fig. 1G, red vs. blue symbols). Moreover, they tended to make more biased decisions in task conditions associated with higher sensory uncertainty. We estimated decision bias by measuring the separation between the prior-conditioned psychometric functions, and sensory uncertainty by calculating the slope of the psychometric function (Materials and Methods; Fig. 1G). For both monkeys, decision bias and sensory uncertainty were significantly correlated (Spearman rank correlation: Monkey J = -0.42, P = 0.017; Monkey F = -0.60, P = 0.0015; Fig. 1H), though note that this effect is largely driven by stimulus contrast (High contrast only: Monkey J = -0.07, P = 0.8; Monkey F = -0.45, P = 0.13; Low contrast only: Monkey J = 0.12, P = 0.66; Monkey F = -0.64, P = 0.022). This data pattern suggests that the animals used a perceptual decisionmaking strategy that resembles Bayesian inference.

Linking PFC Population Activity to Perceptual Inference at the Single Trial Level. The animals' choice behavior was variable. For most experimental conditions, the overall proportion of clockwise choices was neither zero nor one (Fig. 1*G*). This choice variability may either arise from cross-trial variability in sensory measurements (28–31) or from fluctuations in prior expectations (11). Given that choice variability was minimal or absent for the most extreme stimulus orientations (Fig. 1*G*), less likely sources are noise in the choice–response mapping process or attentional "lapses." Because the choice variability has no obvious external origin, identifying the neural factors that determine the outcome of individual decisions ultimately requires a moment-to-moment analysis of neural population activity within single trials (32, 33).

We obtained a trial-by-trial measurement of the animals' evolving decision state by decoding a time-varying decision variable (DV) from jointly recorded neural responses using linear discriminant analysis (Materials and Methods). The DV indicates how well the subject's upcoming perceptual choice can be predicted from a 50 ms bin of neural population activity. In previous work, we showed that trial-averaged DV trajectories exhibit key signatures of a decision-making process (23). Most importantly, grouping trials by stimulus strength and choice accuracy revealed a graded representation of sensory evidence (the more the stimulus orientation deviates from vertical, the higher the sensory evidence). In contrast, the action-planning component of neural activity did not show these effects. We concluded that the decision-making process was implemented as a competition between possible perceptual interpretations, independent of the ensuing action. Here, we build on these findings to develop an analysis that aims to uncover how decisionmaking circuits integrate perceptual expectations and sensory signals.



Fig. 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 toward 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 (Materials and 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 (Materials and 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 (Materials and 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 are 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 (Materials and 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. (A, G, and H): Adapted from ref. 23, which is licensed under CC BY 4.0.

Consider the DV trajectories of two example trials, measured from the same set of neurons in the same experimental condition (Fig. 2A, symbols). Both trials yielded identical overt choice behavior (a clockwise choice indicated with a rightward saccade). The raw DV trajectories are noisy, but share some clear commonalities. On both trials, the DV value hovers around zero before stimulus onset. Following stimulus onset, the DV ramps up to a peak value, after which it decays steadily. We speculate that the peak occurs around the time of choice commitment, after which the animal starts to prepare the corresponding motor response (23). Closer inspection of the trajectories reveals that the build-up toward the peak differs in a number of ways. The excursions seem to start from a different baseline level, and the ramping phase appears to differ in slope and amplitude. We hypothesize that these features capture an important aspect of the temporal evolution of the animal's decision state. We therefore sought to obtain quantitative estimates of these trajectory features for each trial.

We estimated key features of the DV trajectories by fitting a model in which trajectories evolve smoothly to the raw DV values (Fig. 2A, lines; see *Materials and Methods*). This way, we obtained a single "model DV-trajectory" for each trial. Due to the irregular nature of the raw DV trajectories, the correlation between model DV-trajectory and raw DV value was modest (median Pearson correlation: Monkey J = 0.63; Monkey F = 0.58; Fig. 2B). Nevertheless, the model captured the systematic structure of the trial-averaged data well. This was evident at the level of individual recording sessions (example shown in Fig. 2C). To test whether the model DV-trajectories afford additional insight into the animals' decision state, we conducted a logistic regression analysis of the choice behavior. For each recording session, we first measured the association between the experimentally controlled variables (stimulus prior, orientation, and contrast) and choice outcome (i.e., correct or incorrect). We then asked whether including the peak value of the model DV-trajectory as an additional regressor helped to better predict choice outcome (Materials and Methods). We computed a standard measure for prediction error (Akaike's Information Criterion, AIC) and found that including the peak value of the model DV-trajectories systematically improved prediction quality (manifesting as a difference in AIC that is larger than 0 in Fig. 2D).

We have argued that the monkeys' behavior suggests that they combined prior expectations and sensory inputs in a manner that resembles Bayesian inference. We therefore hypothesized that the peak value of the neurally decoded DV should exhibit key signatures of a Bayesian posterior. To test this, we investigated how the model DV-trajectories' peak value depended on the stimulus prior, orientation, and contrast. In the vast majority of recording sessions, the peak DV value exhibited the hypothesized structure, as can be seen for an example session (Fig. 2*E*; more examples shown in *SI Appendix*, Fig. S2). Specifically, the peak



**Fig. 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 = clockwise skew, stimulus orientation = 2.2 deg, stimulus contrast = high, mapping rule = 1, and choice = "clockwise"). Symbols show the raw DV estimates, lines the fit of a model. The average time of stimulus onset is indicated by the black arrow. (*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* and *G*) Comparison of slope (*F*) and offset (*G*) for high vs. low contrast stimuli.

DV grows with stimulus orientation but is biased by the stimulus prior. The impact of the sensory input is greater when stimulus contrast is high, as is evident from the difference in the slope of a linear regression line (Fig. 2 E, Left vs. Right panel, slope = 0.18 for low contrast stimuli and 0.31 for high contrast stimuli). Conversely, the impact of the prior expectation is greater when stimulus contrast is low, as is evident from the difference in the offset of the regression lines (Fig. 2 E, Left vs. Right panel, offset = 0.86 for low contrast stimuli and 0.35 for high contrast stimuli). These effects are representative of both animals' data (median difference 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 signed-rank test; Fig. 2F; median difference in offset: monkey J = -0.28, P = 0.007; monkey F = -0.04, P = 0.05; Fig. 2G). Thus, around the putative time of choice commitment, the neurally estimated decision state exhibits key signatures of a Bayesian posterior. Having established this, we now turn to the question of how prior expectations impact the evolving decision state over the course of a single trial.

**Prior Expectations Bias Neural Population Trajectories.** The effects of prior expectation on the evolving decision state can be appreciated by considering neural responses to low contrast stimuli whose orientation matches the categorization boundary (vertical gratings). These stimuli yield sensory responses that are, on average, completely ambiguous. It follows that any systematic choice biases, summarized in Fig. 1*D*, are purely driven by the subject's prior belief. As can be seen in an example recording session, the neural DV reflects this prior belief well before stimulus onset (Fig. 3*A*). At this point in time, the animal can only leverage knowledge about the skew of the stimulus distribution (i.e., whether a clockwise or counterclockwise orientation is more likely) and the level of sensory uncertainty (i.e., whether a low

or high contrast stimulus is more likely). For both animals, this contextual information biases the DV's early value in the direction of the prior expectation. This effect was prominent for Monkey J, but subtle for Monkey F, as is evident from the DV traces averaged across all sessions (Fig. 3B) and from the model DV-trajectory value before stimulus onset (difference in median value for the -800 to -600 ms epoch: monkey J = 0.1354, P < 0.001; two-sided Wilcoxon rank sum test; monkey F = 0.0111, P < 0.05; Fig. 3C). As can also be seen from the average DV traces, the prior-induced bias is not static, but grows over time. To quantify this, we used the model DV-trajectories to compute the average DV value per recording session for an early and late temporal window (spanning the range from -800 to -600 ms and -500 to -300 ms). The first window precedes any stimulus-driven activity in PFC, while the second one overlaps with the putative time of choice commitment. Plotting the late against the early DV value reveals a relationship that is steeper than the line of unity (Fig. 3D). For both monkeys, the gradient of the first principal component was approximately 72.5 deg, significantly steeper than 45 deg (monkey  $J: 72.0 \deg, P < 0.001$ , bootstrap analysis; see *Materials* and Methods; monkey F: 72.7 deg, P < 0.001).

Why does the influence of the prior expectation grow during deliberation? We hypothesize that both animals integrated incoming sensory evidence over time and gave more weight to noisy orientation estimates that agreed with their prior expectation. Such a confirmation bias deviates from an ideal inference strategy, in which all sensory evidence is given the same weight, resulting in a flat average DV trace for ambiguous stimuli. However, a confirmation bias naturally arises when subjects use an inference strategy that approximates the posterior rather than computing it exactly (34). Under approximate inference, as incoming sensory evidence is accumulated to update the posterior



**Fig. 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 symbol summarizes the trial-averaged values of a single recording session (one point per stimulus prior). (*F*) 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 arrow). (*F*) Late DV value plotted against early DV value for the simulations without (*Left*) and with (*Right*) prior-induced drift.

belief in the state of the world, the prior belief enters into the update multiple times, resulting in a positive feedback loop (34). Critically, this explanation implies that a decision-making model composed of an initial prior-induced offset and a bound that terminates the deliberation process does not suffice to capture the dynamically increasing influence of the prior on the DV. The extra ingredient that is needed is a prior-induced drift (34). To test this explanation of our data, we simulated bounded accumulation of ambiguous sensory evidence and compared early and late DV values (Materials and Methods; Fig. 3E). We varied the strength of the prior-induced effects and the height of the bound, resulting in a range of DV values (Fig. 3F). In the simulations that only included an early offset and lacked a prior-induced drift, early and late DV values were equal on average (Fig. 3 F, Left). We found that inclusion of a priorinduced drift was necessary to produce late DV values that were greater than the early DV values, as seen in the physiological data (Fig. 3 *F*, *Right*). Together, these results suggest that the subjects interpreted neutral sensory evidence in a biased fashion, thus further entrenching their prior expectations.

Relating the Dynamic Range of Neural Trajectories to Decision Bias and Sensory Uncertainty. Our analysis so far suggests that the decision-related component of neural population activity in PFC resembles a bounded accumulation process in which prior expectations influence the DV both before and during accumulation of sensory evidence. This hypothesis is in part motivated by the neural DV's temporal evolution during presentation of ambiguous sensory stimuli (Fig. 3). To test its explanatory power, we first asked whether such a process can capture the behavioral effects of prior expectation for ambiguous as well as nonambiguous stimuli. To this end, we extended the drift-diffusion model simulation to nonambiguous stimulus conditions (Materials and Methods; SI Appendix, Fig. S3A). In our simulation, stimulus orientation determines the mean of the momentary sensory evidence distribution, while stimulus contrast governs its spread. We further assumed that prior expectations influence the DV's initial value and its stimulusdependent drift in exactly the same manner across all 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 at the time of choice commitment determines choice outcome. Consider one specific instantiation of this process (Fig. 4A). As can be seen, prior expectations produce a horizontal shift of the psychometric functions, just like we observed in the animals' behavior (compare with Fig. 1G). Moreover, lowering stimulus contrast makes the psychometric function more shallow and increases the decision bias (Fig. 4A, light vs. dark lines), again recapitulating the animals' behavior. Varying the model parameters can change the magnitude of the decision bias and the steepness of the psychometric functions, but it does not alter this basic pattern (*SI Appendix*, Fig. S3B).

The proposed decision-making process gives rise to a relationship between the dynamic range of the DV trajectory and decision bias. Specifically, simulated trials in which the eventual choice aligns with the prior expectation tend to have a smaller dynamic range (defined as the difference between the simulated DV's peak value and its initial value), while trials in which the choice deviates from the prior expectation tend to have a larger dynamic range (Fig. 4*B*). This relationship arises naturally if a prior expectation influences the DV's initial value while a fixed boundary terminates the deliberation process (or conversely, if a prior expectation influences the bounds without impacting the initial value). To examine the relationship between dynamic range and decision bias, we computed the distribution of the DV's dynamic range across all trials within a simulated experiment and separately analyzed the choice behavior for the trials whose dynamic range was below and above the median (Fig. 4*C*). We observed that the effect of the prior differs substantially across the low- and high-dynamic range trials in the simulated dataset. Specifically, when the DV's dynamic range is low, the decision bias is amplified, and the psychometric function is more shallow (Fig. 4*C*, light vs. dark lines). This pattern of results is specific to a decision-making process in which the prior changes the distance between the DV's starting point and the bounds. Alternative scenarios in which the prior only induces a drift during the accumulation of sensory evidence, or in which there is no terminating boundary, do not produce a relationship between dynamic range and decision bias (*SI Appendix*, Fig. S4).

To test whether these simulated effects are present in the real data, we conducted the same median-split analysis on the neural DVs, using the model DV-trajectories to estimate each trial's dynamic range (Materials and Methods). We identified the same relationship between dynamic range and the animals' decision bias. To quantify this relationship, we computed the difference between the 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, the mean value of  $\Delta B$  was 0.45 (P < 0.001, Wilcoxon signed rank test; Fig. 4 E, Left). Considering the data at a more granular level (per monkey and per contrast level) reveals the robustness of these results, though not every individual case reached statistical significance (Fig. 4 E, Left). We also measured the relationship between dynamic range and the steepness of the psychometric function, quantified as  $\Delta S = \Delta \sigma_L - \Delta \sigma_H$ , and again found a pattern of results consistent with the model simulation. Lower dynamic range trials tended to be associated



**Fig. 4.** Shorter neural excursions result in more biased perceptual decisions. (*A*) We simulated bounded accumulation of ambiguous and nonambiguous 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 (gray trace). The dynamic range of the excursion measures the distance traveled 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 (gray symbols), and split out by monkey and stimulus contrast (nongray 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, lotice a s function of the difference in Akaike's Information Criterion. (*G*) (*Left*) second biase lotted as a function of the difference and 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. (*Right*) Association between the AIC difference and the difference in decision bias. (*Right*) Association between the AIC difference in the slope of the psychometric function. Error bars show  $r \pm 1$  s.e.m



**Fig. 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 (gray symbol), and split out by monkey and stimulus contrast (nongray symbols). Error bars show  $r \pm 1$  s.d. \* P < 0.05, \*\* P < 0.01, and \*\*\* P < 0.001.

with a shallower psychometric function (mean value of  $\Delta S$  was 0.09, P < 0.01; Fig. 4 E, Right). For some sessions, these effects were substantial; for others, they were small. Having established earlier that the peak DV values exhibit key signatures of a Bayesian posterior, we asked whether the variability in effect size over sessions might be predicted by how strongly the peak DV values correlated with the animal's behavior. Interestingly, the effects tended to be larger for datasets where the model DV-trajectories' peak values better predicted the animal's choice behavior (Fig. 4F). To summarize these results, we computed the correlation between  $\Delta B$  and  $\Delta S$  on the one hand, and the previously introduced measure for model prediction error (AIC) on the other. Across all experiments, the correlation with  $\Delta B$ was 0.62 (P < 0.001; Fig. 4 G, Left panel), while the correlation with  $\Delta S$  was 0.54 (P < 0.001; Fig. 4 G, Right panel). Thus, the stronger the neural correlate of the decision-making process, the more the empirical data resemble the predictions of an idealized model of how prior expectations shape this process.

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

#### Discussion

In this study, we investigated neural population activity in the prefrontal cortex of macaque monkeys while they judged ambiguous visual stimuli under different prior stimulus statistics. We sought to understand how decision-making circuits combine sensory signals and prior knowledge to form a perceptual decision. We used a task that requires flexible reporting of perceptual decisions and found that prior expectations exert numerous influences on neural population activity during decision formation. In our task, the neural correlate of decision formation resembles a tug-of-war style competition between candidate perceptual interpretations, independent of the ensuing motor response used to communicate the decision (23). Prior expectations shift the starting point of the competition toward the more probable stimulus interpretation (Fig. 3 A–C). Additionally, prior expectations selectively change the impact of each piece of sensory evidence on the competition, favoring the evidence that is congruent with the expectation (Fig. 3 D–F). Finally, prior expectations shorten the distance between the starting point and end point of the decision formation trajectory (Fig. 4).

Several previous studies also used biased perceptually driven behavior as a gateway to study the neural implementation of Bayesian computation in decision-making areas (16-20, 35). These studies revealed that expectations can modulate the activity of single neurons before stimulus onset (16, 17), while the stimulus is presented (16, 18-20), and while the behavior is produced (19, 20). However, due to the experimental design, it is impossible to say whether these effects reflect a neural correlate of perceptual inference or of motor planning. Here, we built on this work by using a task in which the perceptual and motor components of the decision process are orthogonalized (Fig. 1A). This enabled us to isolate the component of neural activity that solely pertains to the formation of a perceptual choice. Our analysis revealed population level effects that echo elements of the aforementioned findings but that are clearly situated in the domain of perceptual inference. Specifically, we showed that neurons in the prearcuate gyrus can collectively represent abstract perceptual expectations before stimulus onset and confirmation biases during the evaluation of sensory evidence. These results suggest that neural mechanisms that perhaps originally evolved to select the best candidate action (36) over time acquired sufficient flexibility such that they can be repurposed to select the best candidate interpretation of the state of the environment in light of incoming sensory signals and prior experience.

Similarities across studies that investigate different brain areas in different perceptual tasks are helpful to identify key properties of a general neural mechanism. We wish to highlight three striking similarities. First, Hanks et al. conducted single-cell recordings in the lateral intraparietal cortex (LIP) during a nonorthogonalized visual motion discrimination task (18). They found that prior expectations bias responses of LIP neurons in a time-dependent manner, similar to the time-dependent effects we observed for a population-level signal in PFC (Fig. 3 A-D). Second, Sohn et al. studied neural population activity in the dorsomedial frontal cortex during a time-interval reproduction task (20). They found that the neural implementation of Bayesian computation was evident in the dynamic evolution of population activity, which is also a key finding of this study (Fig. 4). Finally, Serrano-Fernandez et al. investigated neural population activity in PFC during a vibrotactile frequency-discrimination task. They found that the state-space trajectories encoded a warped representation of the stimulus, in line with the predictions of a Bayesian estimator (12), just as we observed for a visual orientation discrimination task (Fig. 3E). Together, these similarities suggest that biases in perceptual judgments arise from dynamically distorted population representations in decision-making areas downstream of the sensory cortex. However, by themselves, these similarities do not reveal which area plays which role in Bayesian inference processes. An important goal for future work is to directly compare multiarea recordings to elucidate this issue.

We extracted the DV from neural population recordings by creating two groups of trials based on the animals' categorical decisions and applying Linear Discriminant Analysis. Given this, is it not guaranteed that neural DVs will reveal a bounded accumulation process? It is not. First, it is not guaranteed in advance that neural correlates of a decision reflect a decision formation process. For example, decision-related activity in the early visual cortex is thought to arise at least in part after the decision-making process is completed (37-39). Second, it is not obvious in advance that the brain relies on a bounded evidence accumulation computation to perform a fine orientation discrimination task. Two reasonable alternatives are a snapshot decision process (stimulus orientation is a static aspect of the environment in our task, so this would represent an economic strategy that saves spikes), or unbounded integration of noisy sensory responses (this strategy would be costly in terms of spikes, but it would harvest more information contained in the sensory inputs). However, we can reject these alternatives since they predict no relation between the DV's dynamic range and the strength of the bias, and thus are incommensurate with our data (Fig. 4). In sum, detailing the dynamic trajectories of candidate causal neural correlates of a decision and studying their relation with the choice behavior is necessary to discover which neural circuits make decisions and how they do so.

Our work is not without limitations. First, while we found our binary decision-making task effective in exposing key characteristics of neural computations, it is not rich enough to distinguish between inference schemes that involve full probability distributions (Fig. 1E) and simpler heuristic approximations that involve point estimates. Second, the prior condition and stimulus contrast were held constant for short blocks of trials. Doubtlessly, the monkeys figured out this temporal structure of the task. It is therefore possible that the animals' attention level varied across these blocks-perhaps they preferred one type of block over others. If this were the case, it could result in neural effects beyond the influences of expectation and sensory evidence. We cannot exclude the presence of such effects in our data. Finally, our populations consisted of 5 to 35 units-a very small number compared to the number of neurons involved in the computations under consideration. As is evident from recent work, recording from substantially larger populations can provide deep insight into the neural basis of individual decisions (40), and reveal functionally distinct subpopulations (41).

Figurative language used to describe biased reasoning often invokes inertia. A biased person can be said to be entrenched in a position, immovable by evidence. In the prefrontal cortex, we found a literal manifestation of the figurative relation between bias and inertia. When the neurally decoded decision variable covered less distance, monkeys gave a more biased account of the state of the sensory environment. This finding establishes a direct link between Bayesian computation and the dynamic evolution of neural activity during decision-formation (also see refs. 12 and 20). We suggest that this link exists because the brain's implementation of Bayesian 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 the sensory cortex (29, 42, 43), and thus the quality of perceptual estimates informed by these messages (44). Temporal integration averages out some of this noise, thereby improving the quality of perceptual estimates (43, 45-48). Under an approximate hierarchical inference process, prior expectations change the initial starting point and drift of the accumulator (34). 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 nonambiguous 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 the 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 (49-52). 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 toward uncovering the neurobiological basis of perception and cognition.

## **Materials and Methods**

All analysis and simulation code for this work are included in a public GitHub repository associated with this manuscript.\*

**Experimental Methods and Behavioral Analysis.** Experimental methods were previously described in ref. 23. In brief, the experiments were conducted on two adult male macaque monkeys (*Macaca mulatta*, ages 8 to 9 y old over the course of the experiments). The animals were trained to perform a memory-guided saccade task and an orientation-discrimination task. All training, surgery, and recording procedures conformed to the National Institute of Health Guide for the Care and Use of Laboratory Animals and were approved by The University of Texas at Austin Institutional Animal Care and Use Committee. In order to identify recording sites where neurons exhibited neural activity that selectively

<sup>\*</sup>Project GitHub Repository link: https://github.com/andrefrancois22/Bayesian\_Inference\_ PAG.

preceded an upcoming eye movement, we used a variant of the classic memoryguided saccade task (23, 53). Once we identified a recording site with a well-defined motor response field (see SI Appendix, Fig. S1 for locations of prearcuate gyrus recordings for both animals), we conducted several additional orientation discrimination training sessions with one choice target placed within the estimated response field location and one on the opposite side of the fixation mark. We used exactly the same task for these training sessions as for the subsequent recording sessions. Once psychophysical performance reached a high level, we proceeded with physiological data collection. The orientationdiscrimination task involved two distinct prior contexts, each associated with differently skewed distributions from which stimulus orientations were drawn (see Fig. 1C for the two prior context distributions). Blocks of both prior contexts alternated randomly (with 80 trials per block). Stimuli were drawn from a set of seven possible drifting gratings with orientations evenly spaced over a narrow range of orientations, which was calibrated to each monkey's orientation sensitivity (for monkey F, the seven stimulus orientations ranged from -3.75deg to 3.75 deg, and for monkey J, the stimulus orientations ranged from —3.3 deg to 3.3 deg). Stimuli were presented at either a high or a low contrast (Michelson contrast: 100% or 4%). Blocks of high and low contrast stimuli alternated randomly (trials per block: monkey F = 100, monkey J = 80). Correct choices were followed by a liquid reward. Incorrect choices were followed by an error sound. Vertically oriented stimuli received random feedback. Finally, the reward size was constant across trials and recording sessions for each monkey. We measured observers' choice bias for ambiguous stimuli by computing the log odds of "clockwise" choice under each prior context (Fig. 1D). We measured the ability of the animals to discriminate stimulus orientation by fitting the relationship between stimulus orientation and the probability of a "clockwise" choice with a psychometric function composed of a lapse rate and a cumulative Gaussian function. Model parameters were optimized by maximizing the likelihood of the model over the observed data, assuming responses arise from a Bernoulli process. Each recording session was analyzed independently. To quantify the decision bias of the animals across all stimulus orientations, we fit one psychometric function per stimulus prior and contrast level (assuming that the slope of the psychometric function was the same across both prior conditions; see Fig. 1H). We computed the decision bias by taking the difference between the means of both cumulative Gaussians (i.e., the magnitude of the horizontal displacement of both psychometric functions). Error bars of model-based statistics are based on a 100-fold nonparametric bootstrap of the behavioral choice data.

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

**Electrophysiological Recordings.** As previously reported in ref. 23, we performed 13 successful recordings from monkey *F* and 16 from monkey *J* (the average number of trials per session was 3,171 for monkey *J*, and 1,593 for monkey *F*). For each subject, we used a chronically implanted recording chamber to measure extracellular spiking activity from populations of PFC neurons. We used a microdrive (Thomas recording) to mechanically advance a linear electrode array (Plexon S-probe; 32 contacts) into the brain at an angle approximately perpendicular to the cortical surface at the beginning of each recording session (*SI Appendix*, Fig. S1). We targeted recording sites that had well-defined motor response fields in a memory-guided saccade task the animals had previously performed in the days leading up to the recordings in the main

task. We positioned the linear arrays so that they approximately spanned the cortical sheet and retracted them after each recording session. Continuous neural data from each channel were acquired and saved at a sampling rate of 30 kHz (Plexon Omniplex System). In order to isolate the responses of individual units, we performed offline spike sorting. First, we spike-sorted the data automatically with Kilosort (54). This was followed by a manual merging and splitting as needed. Because the electrode's position could not be optimized for all contact sites, most of our units likely consist of multineuron clusters (we estimate that the fraction of well-isolated single cells is at most 25%). For all analyses reported here, we included all units whose mean firing rate during the task exceeded 3 ips.

**Decision Variable Analysis.** All decoding analyses were performed on simultaneously recorded, single session data. For each trial, we obtained moment-to-moment measurements of the decision variable by projecting 50 ms bins of population activity onto a linear decoder optimized to distinguish the activity patterns associated with both choice options ("left" vs. "right" choices for the motor DV, and "clockwise" vs. "counterclockwise" choices for the categorical DV, respectively). Specifically, we first individually z-scored each unit's spike counts within every time bin. We then used these z-scored responses to estimate the set of linear weights,  $\mathbf{w} = (w_1, ..., w_n)$ , that best separate the choice-conditioned z-scored response patterns, assuming a multivariate Gaussian response distribution:

$$w = \frac{s}{\Sigma'}$$
[1]

where **s** is the mean difference of the choice-conditioned z-scored responses and  $\Sigma$  is the covariance matrix of the z-scored responses. The decoder weights are calculated from observed trials. To avoid double-dipping, we excluded the trial under consideration from the calculation and solely used all other trials to estimate the weights. This way, we obtained "cross-validated" DV estimates for each time bin:

$$DV_j = \sum w_{ij} Z_{ij},$$
 [2]

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 example single trial DV trajectories. Further information about the decoders can be found in ref. 23, including the temporal evolution of the categorical and motor projection planes.

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

$$AIC = 2k + nln(\hat{\sigma}^2) - 2C,$$
[3]

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 maximum likelihood estimate for the variance of a model's residuals distribution given by the residual sum of squares divided by the degree of freedom. Because only differences in AIC are meaningful, the constant C can be ignored when comparing models, yielding a statistic known as  $\Delta AIC$  (shown in Fig. 2D).

We examined how the average peak DV value depended on the stimulus prior, orientation, and contrast (Fig. 2 *E-G*). For each recording session, we

quantified the effect of orientation with the slope of a linear regression line. We allowed these lines to have a prior-specific intercept and conducted this analysis separately for the high and low contrast trials (example shown in Fig. 2E).

We studied the average temporal evolution of the DV trajectories for the zerosignal stimulus condition (Fig. 3 A and B). To this end, we plotted the average DV around the putative time of choice commitment (-500 to -300 ms) computed from the model trajectory against the average value right before the onset of stimulus-driven activity (-800 to -600 ms). To quantify the growing influence of the prior expectation over the course of a trial, we calculated the first principal component of the resulting scatter plot (Fig. 3D). We evaluated the significance of the PC's gradient by conducting a 1,000-fold bootstrap-analysis and computing the 99.9% CI from the resulting gradient estimate distribution.

Simulated Bounded Accumulation Process. We investigated the dynamic evolution of the decision variable in a bounded accumulation process and considered various scenarios. In all cases, the variable that was integrated was composed of one term representing the momentary sensory evidence, and one term representing the prior expectation. We modeled the momentary sensory evidence, s(t), at each time point t as 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 effects of stimulus orientation), and  $\sigma$  inversely proportional to the reliability of the evidence (capturing the effects of stimulus contrast). We reserved positive values to represent evidence for a clockwise stimulus orientation, and negative values evidence for a counterclockwise orientation. We modeled the time-varying prior expectation as a step function, p(t), that either only had energy at the first time point (capturing a pure stimulus expectation), or decayed to a lower but sustained level for the remaining time points (capturing an additional biased interpretation of sensory evidence). In addition, we introduced trial-bytrial variability to p(t) by multiplying this term with a scaling factor sampled from a log-normal distribution with a mean of one. Each trial consisted of 200 time steps, and each simulation consisted of 3,000 trials. We simulate a decision variable in the following way. We started by adding the prior expectation component p(t) to the momentary sensory evidence component s(t), and then computed the cumulative sum over t. For an example of p(t), s(t), and the resulting DV, see SI Appendix, Fig. S3. Formally, the simulated decision variable is defined as follows:

$$DV = \int_{1}^{T} s(t) + p(t) dt,$$
 [4]

where: p(t) is the step-wise prior expectation with a constant initial offset  $\xi$  at time t = 1 followed by a constant value that depended on the variant of the model. In the "early offset and drift" case (the main model illustrated in Fig. 4 and *SI Appendix*, Fig. S4), this constant value was the initial offset  $\xi$  multiplied by a factor  $\alpha \in [0, 1]$  which typically decreased the strength of the initial "pure" prior expectation at time t = 1 to a lesser sustained level for t > 1. In the case of the "no early offset" model (illustrated in the rightmost panel of SI Appendix, Fig. S3), the initial offset parameter  $\xi$  was set to zero. p(t) is defined as follows:

$$\rho(t) = \begin{cases} \xi & \text{if } t = 1\\ \alpha \cdot \xi & \text{if } t > 1 \end{cases}$$
[5]

- H. von Helmholtz, Handbuch der physiologischen Optik (Leopold Voss, 1867), vol. III.
- D. Knill, A. Pouget, The Bayesian brain: The role of uncertainty in neural coding and computation. 2. Trends Neurosci. 27, 712-719 (2004).
- W. S. Geisler, J. S. Perry, B. J. Super, D. P. Gallogly, Edge co-occurrence in natural images predicts contour grouping performance. Vis. Res. 41, 711-724 (2001).
- 4. P. Mamassian, R. Goutcher, Prior knowledge on the illumination position. Cognition 81, B1-B9 (2001).
- Y. Weiss, E. P. Simoncelli, E. H. Adelson, Motion illusions as optimal percepts. Nat. Neurosci. 5, 598-604 (2002).
- A. A. Stocker, E. P. Simoncelli, Noise characteristics and prior expectations in human visual speed 6. perception. Nat. Neurosci. 9, 578-585 (2006).
- A. R. Girshick, M. S. Landy, E. P. Simoncelli, Cardinal rules: Visual orientation perception reflects 7. knowledge of environmental statistics. Nat. Neurosci. 14, 926-932 (2011).
- W. J. Adams, E. W. Graf, M. O. Ernst, Experience can change the "light-from-above" prior. Nat. 8 Neurosci. 7, 1057-1058 (2004).
- 9 D. C. Knill, Learning Bayesian priors for depth perception. J. Vis. 7, 13 (2007).

When the simulated process contained decision bounds, the deliberation process was typically terminated early (at the time when one of the bounds was crossed). Fig. 3F illustrates an analysis in which we only included zero-signal stimuli and varied the strength of the prior-induced effects and the height of the bounds across simulations. Fig. 4A shows an analysis of the choice behavior for a simulation that included all of our experimental conditions as well as the effects of stimulus expectation, biased interpretation of sensory evidence, and terminating decision bounds.

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

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

$$\Delta B = \Delta \beta_L - \Delta \beta_{H'}$$
 [6]

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

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

$$\Delta U = \Delta v_L - \Delta v_{H'}$$
<sup>[7]</sup>

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

Data, Materials, and Software Availability. Previously published data were used for this work (23).

ACKNOWLEDGMENTS. This work was supported by US NIH grants T32 EY021462 (T.A.L. and J.A.C.), EY032999 (R.L.T.G.), and NSF CAREER award #2146369 (R.L.T.G.).

Author affiliations: <sup>a</sup>Massachusetts Institute of Technology, Department of Brain and Cognitive Sciences, Cambridge, MA 02139, <sup>b</sup>Princeton Neuroscience Institute, Princeton University, Princeton, NJ 08540; and <sup>c</sup>Center for Perceptual Systems, The University of Texas at Austin, Austin, TX 78712

- 10. M. Jazayeri, M. N. Shadlen, Temporal context calibrates interval timing. Nat. Neurosci. 13, 1020-1026 (2010).
- 11. J. A. Charlton, W. F. Młynarski, Y. H. Bai, A. M. Hermundstad, R. L. Goris, Environmental dynamics shape perceptual decision bias. PLoS Comput. Biol. 19, e1011104 (2023).
- 12. L. Serrano-Fernández, M. Beirán, R. Romo, N. Parga, Representation of a perceptual bias in the prefrontal cortex. Proc. Natl. Acad. Sci. U.S.A. 121, e2312831121 (2024).
- C. Summerfield, F. P. De Lange, Expectation in perceptual decision making: Neural and computational mechanisms. Nat. Rev. Neurosci. 15, 745-756 (2014).
- 14. W. F. Mlynarski, A. M. Hermundstad, Adaptive coding for dynamic sensory inference. eLife 7, e32055 (2018).
- 15. M. Hahn, X. X. Wei, A unifying theory explains seemingly contradictory biases in perceptual
- estimation. *Nat. Neurosci.* **27**, 793–804 (2024). 16. V. Rao, G. C. DeAngelis, L. H. Snyder, Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. J. Neurosci. 32, 10063–10074 (2012).
  17. M. A. Basso, R. H. Wurtz, Modulation of neuronal activity in superior colliculus by changes in target
- probability. J. Neurosci. 18, 7519-7534 (1998).

- 18. T. D. Hanks, M. E. Mazurek, R. Kiani, E. Hopp, M. N. Shadlen, Elapsed decision time affects the weighting of prior probability in a perceptual decision task. J. Neurosci. 31, 6339-6352 (2011).
- 19. T. R. Darlington, J. M. Beck, S. G. Lisberger, Neural implementation of Bayesian inference in a sensorimotor behavior. Nat. Neurosci. 21, 1442-1451 (2018).
- H. Sohn, D. Narain, N. Meirhaeghe, M. Jazayeri, Bayesian computation through cortical latent 20. dynamics. Neuron 103, 934-947 (2019).
- 21. V. P. Ferrera, M. Yanike, C. Carlos, Frontal eye field neurons signal changes in decision criteria. Nat. Neurosci. 12, 1458-1462 (2009).
- P. Cisek, J. F. Kalaska, Neural mechanisms for interacting with a world full of action choices. Annu. Rev. Neurosci. 33, 269-298 (2010).
- J. A. Charlton, R. L. T. Goris, Abstract deliberation by visuomotor neurons in prefrontal cortex. Nat. 23. Neurosci. 27, 1167–1175 (2024). J. D. Schall, "Visuomotor areas of the frontal lobe" in *Extrastriate Cortex in Primates*, K. S. Rockland,
- 24. J. H. Kaas, A. Peters, Eds. (Springer, Boston, MA, 1997), pp. 527-638.
- 25. E. K. Miller, J. D. Cohen, An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. 24, 167-202 (2001).
- R. Kiani, C. J. Cueva, J. B. Reppas, W. T. Newsome, Dynamics of neural population responses in 26. prefrontal cortex indicate changes of mind on single trials. Curr. Biol. 24, 1542-1547 (2014).
- V. Mante, D. Sussillo, K. V. Shenoy, W. T. Newsome, Context-dependent computation by recurrent 27. dynamics in prefrontal cortex. Nature 503, 78-84 (2013).
- 28. W. P. Tanner, J. A. Swets, A decision-making theory of visual detection. Psychol. Rev. 61, 401-409 (1954).
- 29. P. Heggelund, K. Albus, Response variability and orientation discrimination of single cells in striate cortex of cat. Exp. Brain Res. 32, 197-211 (1978).
- M. N. Shadlen, K. Britten, W. Newsome, J. Movshon, A computational analysis of the relationship 30. between neuronal and behavioral responses to visual motion. J. Neurosci. 76, 1486-1510 (1996).
- 31. R. L. T. Goris, J. A. Movshon, E. P. Simoncelli, Partitioning neuronal variability. Nat. Neurosci. 17, 858-865 (2014).
- 32. D. Peixoto et al., Decoding and perturbing decision states in real time. Nature 591, 604-609 (2021).
- 33. G. Stine, E. M. Trautmann, D. Jeurissen, M. N. Shadlen, A neural mechanism for terminating decisions. *Neuron* **111**, 2601-2613.e5 (2023).
- R. Lange, A. Chattoraj, J. Beck, J. Yates, R. M. Haefner, A confirmation bias in perceptual decision-34. making due to hierarchical approximate inference. PLoS Comput. Biol. 17, e1009517 (2021).
- T. B. Crapse, H. Lau, M. A. Basso, A role for the superior colliculus in decision criteria. Neuron 97, 181-194 (2018).
- P. Cisek, Cortical mechanisms of action selection: The affordance competition hypothesis. Philos. 36. Trans. R. Soc. B Biol. Sci. 362, 1585-1599 (2007).

- 37. H. Nienborg, B. G. Cumming, Decision-related activity in sensory neurons reflects more than a neuron's causal effect. Nature 459, 89-92 (2009).
- R. M. Haefner, P. Berkes, J. Fiser, Perceptual decision-making as probabilistic inference by neural 38 sampling. Neuron 90, 649-660 (2016).
- R. L. Goris, C. M. Ziemba, G. M. Stine, E. P. Simoncelli, J. A. Movshon, Dissociation of choice 39. formation and choice-correlated activity in macaque visual cortex. J. Neurosci. 37, 5195-5203 (2017).
- 40. N. Steinemann et al., Direct observation of the neural computations underlying a single decision. eLife 12, RP90859 (2024).
- 41. A. Zylberberg, M. N. Shadlen, Population representation of the confidence in a decision in the lateral intraparietal area of the macaque. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2024. 08.15.608159 (Accessed 2 February 2025).
- D. J. Tolhurst, J. A. Movshon, A. F. Dean, The statistical reliability of signals in single neurons in cat and monkey visual cortex. *Vis. Res.* 23, 775–785 (1983).
- R. Vogels, G. A. Orban, How well do response changes of striate neurons signal differences in orientation: A study in the discriminating monkey. *J. Neurosci.* **10**, 3543–3558 (1990).
   Z. Boundy-Singer, C. M. Ziemba, O. J. H. Henaff, R. L. T. Goris, How does v1 population activity
- inform perceptual certainty? J. Vis. 24, 1-17 (2024).
- T. Uka, G. C. DeAngelis, Contribution of middle temporal area to coarse depth discrimination: 45 Comparison of neuronal and psychophysical sensitivity. J. Neurosci. 23, 3515-3530 (2003).
- 46. L. C. Osborne, W. Bialek, S. G. Lisberger, Time course of information about motion direction in visual area MT of macaque monkeys. J. Neurosci. 24, 3210-3222 (2004).
- 47. Y. Chen, W. Geisler, E. Seidemann, Optimal temporal decoding of neural population responses in a reaction-time visual detection task. J. Neurosphysiol. 99, 1366-1379 (2008).
- 48. R. L. Goris, C. M. Ziemba, J. A. Movshon, E. P. Simoncelli, Slow gain fluctuations limit benefits of temporal integration in visual cortex. J. Vis. 18, 8-8 (2018).
- 49. T. Z. Luo et al., Transitions in dynamical regime and neural mode underlie perceptual decisionmaking. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2023.10.15.562427 (Accessed 2 February 2025).
- S. Bitzer, J. Bruineberg, S. J. Kiebel, A Bayesian attractor model for perceptual decision making. PLoS Comput. Biol. 11, e1004442 (2015).
- A. Dubreuil, A. Valente, M. Beiran, F. Mastrogiuseppe, S. Ostojic, The role of population structure in computations through neural dynamics. *Nat. Neurosci.* 25, 1–12 (2022).
- 52. M. Khona, I. R. Fiete, Attractor and integrator networks in the brain. Nat. Rev. Neurosci. 23, 744-766 (2022).
- S. Funahashi, C. J. Bruce, P. S. Goldman-Rakic, Mnemonic coding of visual space in the monkey's 53. dorsolateral prefrontal cortex. J. Neurophysiol. 61, 331-349 (1989).
- 54. M. Pachitariu et al., Kilosort: Realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv [Preprint] (2016). https://doi.org/10.1101/061481 (Accessed 2 February 2025).