A brain-wide, trial- and time-dependent deterministic drive synergizes with within-trial noise to time self-initiated actions

Michaël A. Elbaz[1], Kole Butterer[1,2], Sara A. Solla[3,4,5], Joshua I. Glaser[5,6,7], Andrew Miri[1,5]

[1]Department of Neurobiology, Northwestern University
[2]Current affiliation: Department of Human Genetics, University of Chicago
[3]Department of Neuroscience, Northwestern University
[4]Department of Physics and Astronomy, Northwestern University
[5]NSF-Simons National Institute for Theory and Mathematics in Biology
[6]Department of Neurology, Northwestern University
[7]Department of Computer Science, Northwestern University

Correspondence: andrewmiri@northwestern.edu, mielbaz@gmail.com

Summary

Deciding when to act in the absence of external cues is essential for exploration, learning, and survival. Yet the neural mechanisms underlying such decisions remain controversial, with current views favoring either deterministic or stochastic underpinnings. We simultaneously recorded from large neuronal populations in cortical, thalamic, pallidal, and cerebellar regions as mice self-initiated voluntary actions. Action onset timing was predictable from firing patterns up to several seconds in advance with predictions correlated across regions, demonstrating a prominent deterministic drive that spans regions. Computational modeling indicated that this drive has an initial value and rate that vary trial-by-trial, and the rate increases within trials. Although the deterministic drive is sufficient to trigger action, noise within trials also contributes to setting action timing. Therefore, discrete (across-trial) and continuous (within-trial) sources of variability synergize to time self-initiated actions. This synergy is observed brain-wide, suggesting a distributed decision-making process rather than a hierarchical, modular one.

Introduction

The adaptive value of cognition manifests not only in deciding how to act, but also when to act. For instance, prey increase their chances of survival by fleeing at a moment predators cannot predict, rather than waiting for predators to give chase[1,2]. This raises a fundamental question: what neural mechanisms underlie action timing decisions in the absence of external cues? Seminal studies in humans over 60 years ago established the existence of the "readiness potential", a slow buildup in electroencephalographic recordings (EEG) that precedes the onset of self-initiated actions by up to several seconds[3-5]. Correlates of the readiness potential have since been found across numerous mammalian species[6-9], and even in fish[10] and invertebrates[11,12], perhaps reflecting a broadly conserved neural dynamics[13,14].

The apparent monotonicity of the readiness potential, and classic experiments suggesting the potential begins before subjects are aware of their desire to act[4,15-18], led to the claim that its onset reflects the decision to act, with the buildup thereafter proceeding deterministically toward action initiation[3,7,14,19-22]. However, more recent studies have challenged this classic interpretation by

implicating a prominent stochastic process in setting self-initiated action timing[14,23-28]. Accordingly, the monotonicity of the readiness potential has been reinterpreted as an artifact of trial-averaging – a method commonly used with noninvasive recording techniques that yield a low signal-to-noise ratio[29]. Indeed, if action initiation requires crossing an activity threshold[7,14], averaging neural activity aligned on action onset will produce an approximately monotonic trajectory, even if the underlying activity is dominated by random fluctuations on individual trials.

This alternative, stochastic interpretation has been formalized using drift-diffusion models (DDM), which were previously used to describe perceptual decision-making[30-36]. DDMs combine contributions from deterministic and stochastic components in a balance that, in perceptually guided tasks, varies depending on the strength of sensory evidence: the stronger the sensory evidence, the larger the deterministic component[32]. In the case of self-initiated actions, the deterministic component was assumed to be minor, insufficient on its own to reach the activity threshold beyond which an action is initiated[23,37]. Conversely, the stochastic component, thought to derive from random fluctuations in the firing rates of neurons underlying the decision[38,39], was posited to play a prominent role, being necessary to reach the activity threshold and thereby setting action timing[14]. This view well accounts for the variable and seemingly unpredictable timing of self-initiated actions. Nevertheless, other models not involving stochastic fluctuations within trials can also account for this timing variability, such as a deterministic process that starts at random times or whose parameters vary randomly from trial to trial. Here the timing decision is made whenever a pre-movement buildup starts, consistent with the classic interpretation of the readiness potential. Without measurements that resolve the decision-related neural dynamics on individual trials, the predictions of these conflicting models appear equally plausible[37], leaving the field at a conceptual impasse.

Compounding matters, the contributions of different brain regions to action timing decisions remain poorly resolved. Activity modulation preceding the onset of self-initiated actions has been observed in numerous cortical and subcortical regions[40-48]. Lesions[49-58] and other activity perturbations[59-64] in some regions alter action timing. Previous studies have promoted a hierarchical view, in which timing decisions rely on distinct modular contributions from different regions carried out in sequence[16,40,45-47,59,61,65,66]. In line with this, functional imaging in humans has shown a continuous sequence of activations across different cortical regions as action onset approaches[16]. However, this interpretation of hierarchical, modular contributions is based on activity measurements with low temporal resolution or that involve a limited number of neurons recorded in one or two regions at a time. Thus, it is possible that self-initiated action decisions instead rely on a shared, distributed process to which many brain regions simultaneously contribute.

To address these issues, we resolved single-unit, single-trial dynamics across large neuronal populations in eight brain regions simultaneously, in mice making self-initiated action decisions. We demonstrate that action onset timing is predictable up to several seconds prior to onset from neural activity. This predictive activity reflects a prominent deterministic drive that is temporally aligned across cortical, thalamic, pallidal, and cerebellar regions. By developing a DDM variant that accounts for the decision process at a single trial resolution, we show that the deterministic drive is parametrized differently across trials, contributing to action timing variability, and is on its own sufficient to trigger action. Yet, *in silico* ablation analyses reveal that within-trial noise also contributes to action timing and its variability across trials. The same synergistic decision process is

observable across disparate brain regions. In total, our results support a hybrid model in which discrete (trial-to-trial) and continuous (within-trial) sources of timing variability act synergistically to set action timing. This resolves the prevailing dichotomy in which timing decisions were seen as either fully deterministic once a buildup process starts or primarily stochastic based on within-trial noise.

Results

A new behavioral paradigm to elicit self-initiated action decisions in mice

Self-initiated decisions have typically been elicited in human subjects by asking them to initiate a motor action "whenever they wish" with no experimental constraints that reinforce a specific timing[3,4,15,16]. To study such decisions in mice, we leveraged a behavioral paradigm we recently developed in which mice perform a naturalistic climbing behavior in the absence of external cues[67]. Here, head-fixed mice climb with all four limbs across handholds that extend radially from a wheel positioned in front of them, thereby rotating the wheel (Figure 1A). In the original version of the task, mice can initiate climbing at any time and are rewarded after climbing. This self-paced paradigm therefore implicitly reinforces a rapid succession of climbing bouts that maximizes the reward rate. To elicit actions whose timing is not influenced by external factors like the pursuit of rewards, we developed a variant of this task.

Our task variant consists of alternating "GO" and "NO-GO" periods, which are distinguishable by the presence or absence of sensory cues. When no movement occurs, both periods last 3 seconds on average (exact lengths are jittered). If the mouse climbs during a GO period, a reward is delivered, the period ends and a NO-GO period begins. If the mouse climbs during a NO-GO period, the period timer is reset, delaying the start of the next GO period. Climbing onset times were computed using video-based limb tracking[68] during epochs surrounding the onset of wheel rotation (Figure 1B). All animals trained on this task learned to climb promptly after the start of a GO period (Figures S1A-S1B). All trained animals also occasionally initiated climbing during NO-GO periods (Figure S1C), with an average of 75 ± 27 (mean ± sem) NO-GO climbing bouts per session. Analysis of anticipatory licking behavior indicated that mice are less likely to expect a reward during NO-GO periods than during GO periods (Figure S1D). This argues against the possibility that climbing bouts occurring during NO-GO periods are reward-motivated errors, wherein mice mistake a NO-GO period for a GO period. Thus, NO-GO climbing bouts do not appear reward-driven but self-initiated. We focus on these events in what follows.

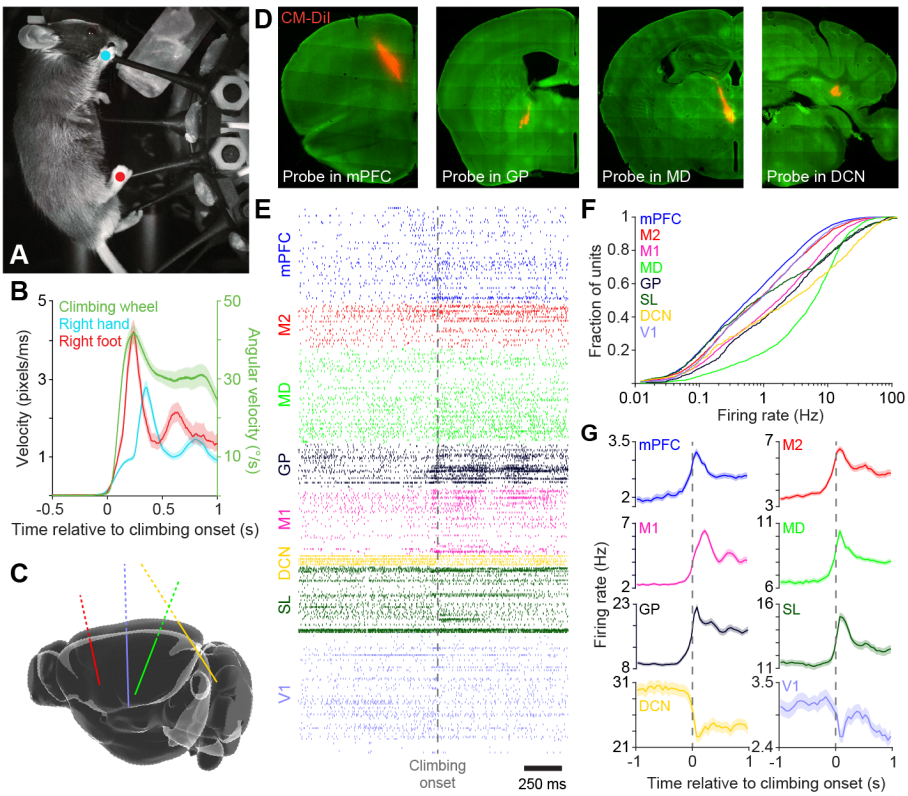

Figure 1. Large-scale neural recordings during self-initiated climbing.
(A) Head-fixed mouse climbing a wheel equipped with radial handholds. **(B)** Trial-averaged velocity of the right hand and foot (left axis) and angular velocity of the climbing wheel (right axis), aligned to climbing onset (mean ± 95% confidence interval estimated from bootstrapping). **(C)** 3D schematic of a mouse brain depicting the trajectories of acutely implanted Neuropixels arrays targeting cortical, thalamic, pallidal, and cerebellar regions. **(D)** Histological sections where probe tracks appear marked with CM-DiI in the medial prefrontal cortex (mPFC), globus pallidus (GP), mediodorsal nucleus of the thalamus (MD), and deep cerebellar nuclei (DCN). Background tissue autofluorescence is shown in green. **(E)** An example spike raster surrounding climbing onset, for all eight regions of interest: mPFC, M2 (secondary motor cortex), M1 (primary motor cortex), MD, GP, DCN, SL (simple lobule) and V1 (primary visual cortex). **(F)** Empirical cumulative distribution functions of firing rates for all qualifying units in each region of interest, pooled across all recording sessions. **(G)** Trial-averaged firing rates (mean ± 95% confidence interval estimated from bootstrapping) across all qualifying units in each region, aligned to climbing onset.

Self-initiated action timing is predictable from spiking activity

We next sought to examine spiking activity across the brain that precedes self-initiated climbing onsets. Four large-scale multielectrode arrays (Neuropixels 1.0[69]) were acutely implanted at the beginning of climbing sessions (six mice, one recording session per mouse). We targeted a range of brain regions in the left hemisphere that have previously been implicated in self-initiated action timing, or which receive direct projections from implicated regions (Figures 1C-1F, S1E and S1F, Table S1): the medial prefrontal cortex (mPFC)[46,48,59,61], the forelimb primary motor cortex (M1)[6,8,20,70], the forelimb secondary motor cortex (M2)[25,46], the mediodorsal nucleus of the thalamus (MD)[71], the globus pallidus (GP)[44,47,59,64,72,73] and, contralaterally, the deep cerebellar nuclei (interposed and dentate nuclei, DCN) and the cerebellar simple lobule (SL)[42,43,51,62]. The array targeting cerebellar regions also recorded activity in the primary visual cortex (V1) which we included among our regions of interest, given recent evidence that this region exhibits movement-related activity[74,75]. The number of single units meeting our quality criteria ("qualifying units") in individual regions ranged from 35 (median across sessions, DCN) to 398 (mPFC; Figure S1E). In all regions of interest, trial-averaged single unit activity appeared modulated around climbing onset and during ongoing climbing (Figure 1G).

To test for a deterministic process underlying the action timing decisions mice made, we investigated whether the timing of climbing onset could be predicted from the preceding spiking activity across regions. For each recording session, we trained a cross-validated elastic net linear regression model[76] to predict the latency to climbing onset (time-to-movement) from the activity of qualifying units between immobility onset and climbing onset. The average duration of epochs from immobility onset to climbing onset ("trials") ranged from 1.7 to 3.6 s across mice (Figure S1G), after ignoring here and in all subsequent analysis trials lasting less than 1 s. Firing rates were averaged in 50 ms bins, and the forecast model used the rates from the preceding five bins (250 ms) to predict time-to-movement at each time bin. Predictions for held-out trials often showed a ramping time course (Figures 2A and 2B), but were not consistent with neural activity merely tracking the passage of time (Figures S2A-S2C)[77].

To determine whether the model's predictive accuracy was significantly better than chance, we computed two performance metrics for each held-out trial: the mean absolute difference between the forecasted and true times-to-movement (residuals), and the slope of a linear fit to the forecasted time-to-movement over time. To compute a *p*-value for each metric on each trial, actual values were compared against null distributions generated using the forecast values of models fit to permuted data. Forecast values were significantly better than chance for a majority of trials based on their residuals, and for nearly all trials based on their slopes (Figures 2C-2E). Importantly, when either performance metric was examined as a function of the time elapsed between the last GO to NO-GO period transition and immobility onset, no significant correlation was observed (Figures S2D and S2E). This argues against the possibility that models exploit lingering sensory-related activity associated with GO to NO-GO transitions. Together, these results indicate the presence of a prominent deterministic component in action timing on individual trials.

To evaluate when predictive information emerges on individual trials, we computed the same two performance metrics within 250-ms sliding windows leading up to climbing onset and compared the results to chance expectations (Figures 2F-2G). For each metric, significant predictivity emerged

171 within the first 500 ms after immobility onset on a majority of trials (Figure 2H). This corresponds to
172 1.4 s $\pm$ 0.66 (median $\pm$ median absolute deviation) before action onset across all trials (Figure S2F),
173 and to 2.78 s $\pm$ 0.58 for trials with durations $\geq$ 3 s (Figures S2G and S2H). This indicates that neural
174 activity is predictive of self-initiated action timing on single trials, up to several seconds before action
175 onset. This identification of activity predictive of action timing through our regression-based
176 approach extends prior studies that relied on classification to predict whether an action would occur
177 within pre-defined temporal windows[16,20].

178

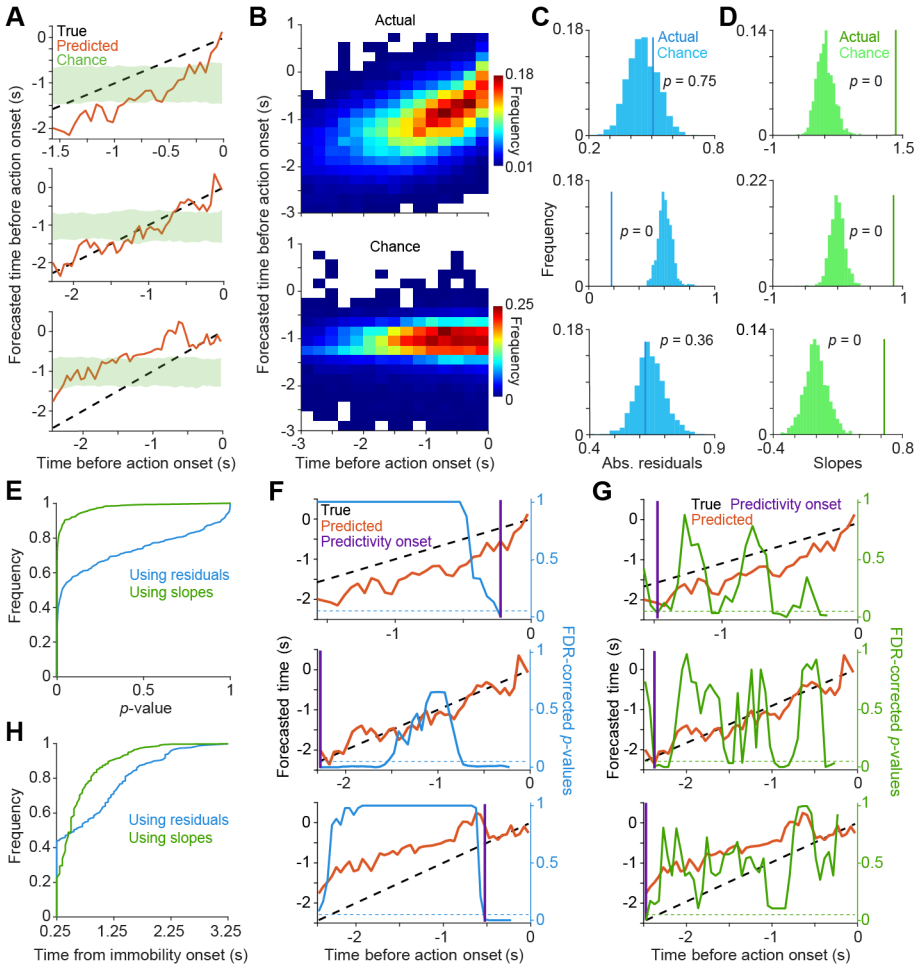

Figure 2. Self-initiated action timing is predictable from neural activity on individual trials. **(A)** Three example single-trial predictions of time-to-movement (red) versus true time-to-movement (dashed black). The green shaded area represents the 95% confidence interval of predictions from permuted data. Here and in subsequent analyses, all figures illustrate performance on held-out data. **(B)** 2D histograms of forecasted versus actual time-to-movement for all trials pooled across all sessions (top) and after shuffling (bottom). **(C, D)** Absolute residual (C) and slope (D) for each example trial from (A), compared with their corresponding null distributions computed with shuffling. **(E)** Empirical cumulative distribution functions of *p*-values from tests on absolute residuals (blue) and slopes (green) across all trials and sessions. **(F, G)** Identification of predictability onset for the example trials shown in (A) using either the absolute residuals (F) or slopes (G) calculated in 250-ms sliding windows. Predictability onset was defined as the first time step when the forecast traces drop below an FDR-corrected *p*-value of 0.05. **(H)** Empirical cumulative distribution functions of predictability onset times relative to immobility onset for all trials, determined using either absolute residuals (blue) or slopes (green).

Decision-related, temporally aligned activity across disparate brain regions

Is the activity used by the forecast model to predict time-to-movement localized in a subset of brain regions? To address this, we tested whether each unit's activity improved time-to-movement predictions compared to circularly permuted controls. In all eight regions of interest, we found neurons that significantly improved forecast accuracy ("forecasting neurons"; Figure 3A), ranging from 27% (median, mPFC) to 39% (M2; Figure 3B) of units in each region. When aligned on climbing onset, the trial-averaged activity of these neurons often exhibited a ramping time course leading up to action onset (Figure 3C). This pattern was also often apparent on single trials (Figure S3A). Thus, the activity of neurons within each region of interest carries information about future action timing.

We then examined how decision-related activity in different regions correlates on individual trials. We trained and tested separate forecast models using activity from each individual region. The resulting single-trial forecast traces were remarkably similar across all regions, often exhibiting coordinated deviations from an ideal linear prediction of time-to-movement (Figure 3D). Predictions were positively correlated across all region pairs when pooled across all test trials and all time points up to 250 ms before action onset to avoid activity dynamics related to movement execution[15,78] (Figures 3E and 3F). Moreover, when these across-region correlations were calculated for each 50-ms time bin, they remained consistently positive and significant throughout trials, from their emergence soon after immobility onset leading up to action onset (Figure 3G). To test whether predictive activity in one region consistently led predictive activity in another region, we computed the cross-correlations of the single-trial predictions for all region pairs. The maximum cross-correlation peaked at or near a time lag of zero for all pairs (Figures S3B-S3E). This temporal alignment held even when predictions used firing rates in shorter, 10-ms bins and did not rely on additional preceding bins (Figures S3C and S3E). This brain-wide temporal alignment throughout the decision epoch suggests that self-initiated action timing emerges from a unified, distributed process rather than a modular, sequential one.

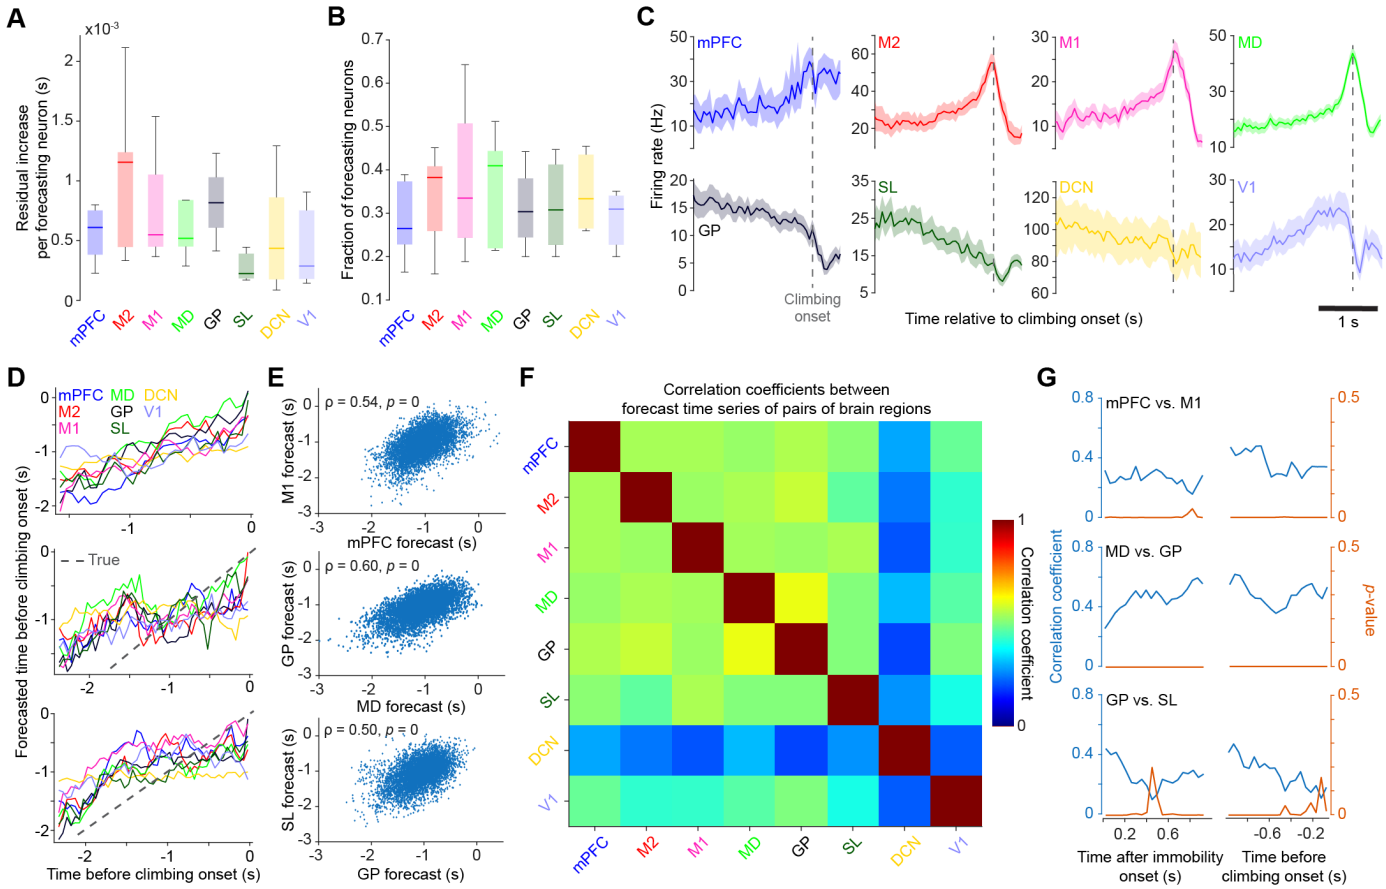

Figure 3. Decision-related activity is present and temporally aligned brain-wide. **(A)** Increase in forecast error upon removal of individual forecasting neurons from each region across all sessions. **(B)** Proportion of forecasting neurons in each region across all sessions. **(C)** Trial-averaged firing rate of example forecasting neurons from each of the eight regions, aligned to climbing onset (mean ± 95% confidence interval estimated from bootstrapping). **(D)** Single-region time-to-movement predictions for three example held-out single trials. **(E)** Scatterplots of time-to-movement predictions for two regions plotted against one another across all time points and trials from one session, illustrating the correlation seen between regions. ρ denotes the Pearson correlation coefficient. **(F)** Correlation matrix for single-region time-to-movement predictions. All correlations are significantly positive (Student's t-test, FDR-corrected *p*-value < 0.05). **(G)** Cross-region correlation (blue) and associated *p*-values (red) calculated at each time step for the single-region time-to-movement predictions of three region pairs across all trials from one session.

The deterministic drive varies within and across trials

The preceding results demonstrate the existence of a substantial, brain-wide deterministic drive underlying self-initiated action timing, in contrast to recent models that center on the contribution of a stochastic process within trials. However, we also found that trial durations varied greatly (Figure S1G) and predictions exhibited substantial errors (Figures 2A and 2E-2G). This variability and unpredictability of action timing could reflect a substantial additional contribution from a stochastic process within trials, but could also arise from features of a deterministic process. To clarify the nature and relative prominence of the deterministic drive, we thus sought to assess different potential sources of variability in action timing.

We considered four distinct mechanisms. First, stochastic fluctuations in neural activity during individual trials (within-trial noise) might influence the decision process, as in the typical DDM. Second, timing variability could be attributed to trial-to-trial variation in the parameters of a deterministic drive. For instance, timing variability could arise from a ramping process whose initial value and slope vary across trials[37,79]. Third, the decision process might start at variable latencies after immobility onset, generating timing variability even if the subsequent dynamics are similar across trials. Finally, timing variability could also arise from trial-to-trial differences in the activity threshold.

To test these possibilities, we leveraged the output of our multi-region, time-to-movement forecast model as a single-trial readout of the decision variable, i.e., the evolution of the decision process over time. If deterministic drive parameters vary across trials, trial duration should correlate with the forecast output's initial value and/or its slope. If the decision process starts at variable latencies after immobility onset, trial duration should be largely uncorrelated with the forecast output's initial value. If the activity threshold varies across trials, trial duration should correlate with the forecast output's final value.

Thus we first plotted three variables versus trial duration: the forecast output's initial value (250 ms after immobility onset because of the window used for prediction), its slope when best fit with a line, and its final value (250 before action onset or at action onset). We found that the initial values and the slopes were negatively correlated with trial duration in all six mice (Figures 4A and 4B). By contrast, the final values, though highly variable, were uncorrelated with trial duration (Figures 4C and 4D). To ensure that the negative correlation between initial value and trial duration was robust across trials, and not due to a subset of trials with an exceptionally early decision process start time, we confirmed that the negative correlation holds for subsets of trials (Figure S4). This also further supports an early engagement of the decision process soon after immobility onset. Together, these results suggest that variability in action timing stems from variation in deterministic drive parameters across trials, but not in the start time or activity threshold of the decision process.

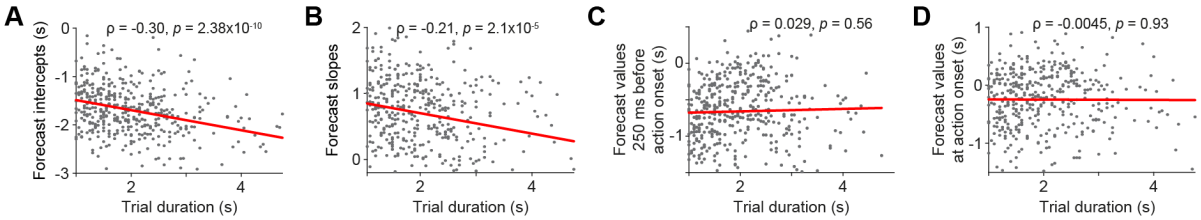

Figure 4. Action timing correlates with the initial value and slope of the forecast model output, but not its final value.

(A-D) Scatterplots of the initial value (A), slope of a linear fit (B), values 250 ms prior to climbing onset (C), and final value at climbing onset (D) for forecast traces versus trial duration across held-out trials from all sessions. ρ denotes the Pearson correlation coefficient, with *p*-values determined by a Student's t-test. Red lines show a best-fit line.

To further probe for trial-to-trial variation in the parameters of a deterministic drive, we used model selection. In the typical DDM, the decision variable x has an initial value x_0, and three other parameters govern Δx_n, the change in the decision variable at time t_n [Equation 1]: I is the drift rate, the slope of a linear ramp; k is the amplitude of a leak process that acts as a restoring force towards baseline[33]; and c is the amplitude of a diffusion process that models the accumulation of stochastic fluctuations within trials, with ξ_n drawn from a standard normal distribution. Thus, in this model, the approach of the decision variable to the activity threshold effectively depends on deterministic drift and leak processes, as well as the accumulation of within-trial noise.

$$\Delta x_n = I\Delta t - kx_n\Delta t + c\xi_n\sqrt{\Delta t} \qquad [1]$$

In this model, since the values of the four parameters are constant across trials, timing variability across trials can arise only from the diffusion process. To capture potential trial-to-trial variability in deterministic drive parameters, we formulated a DDM variant [Equations 2-5]. Here, the initial value and drift rate are drawn from normal and log-normal distributions, respectively, on individual trials. Additionally, previous studies of stimulus-driven perceptual decisions have described an apparent increase over time in drift rate within individual trials[80-88]. Thus, we also incorporated either a linear or exponential factor into the drift term to capture potential time-dependent changes in the drift rate[82,86,87]. This yielded a model with seven parameters: $\mu_{intercept}$, $\sigma_{intercept}$, μ_{drift}, σ_{drift}, α, k and c.

$$x_0 \sim \mathcal{N}(\mu_{intercept}, \sigma^2_{intercept}) \qquad [2]$$

$$I_0 \sim Lognormal(\mu_{drift}, \sigma^2_{drift}) \qquad [3]$$

$$\Delta x_n = I(t)\Delta t - kx_n\Delta t + c\xi_n\sqrt{\Delta t} \qquad [4]$$

$$I(t) = \begin{cases} I_0 + \alpha t_n \\ I_0 e^{\alpha t_n} \end{cases} \qquad [5]$$

We then tested whether the inclusion of individual parameters improved fits to the forecast output. Specifically, we used maximum likelihood estimation to quantify the fit improvements on held-out forecast traces from including $\sigma_{intercept}$, μ_{drift}, σ_{drift}, α and k (Figure S5). All models included $\mu_{intercept}$ since zero is not the average but the lowest possible initial value of our fitted forecast traces (see Methods). Our model selection analysis revealed that each tested parameter significantly improved the fit quality for each mouse (Figures 5A, 5B, S6A and S6B). This further supports the presence of trial-to-trial variability in the deterministic drive parameters ($\sigma_{intercept}$ and σ_{drift}).

In addition, we found that best-fit α values were positive, indicating that the drift rate increases over time within a trial, with a greater likelihood gain achieved with a linear rather than an exponential growth (Figures 5A and 5B). The best-fit values of α imply a ~1.5-fold increase for a 2-s trial, indicating a substantial contribution to action timing from this factor. Thus, the deterministic drive effectively increases in strength as time passes, like an increasing urgency signal.

As a control, we assessed whether variability in the initial value ($\sigma_{intercept}$) might actually reflect within-trial noise accumulated prior to immobility onset, during the execution of the previous

movement[29]. We calculated the theoretical variance limit of the system ($c^2/2k$), which represents the equilibrium between noise accumulation and leak decay[89,90]. We found that the observed initial variance $\sigma^2_{intercept}$ consistently and substantially exceeded this limit (Figure S6C). This implies that $\sigma_{intercept}$ is not a byproduct of pre-trial noise accumulation, but instead reflects a state at immobility onset that varies across trials.

Altogether, these observations are consistent with the decision process involving a deterministic drive that varies (a) discretely from trial to trial in initial value and rate, and (b) over time within a trial, having a drift rate that increases over time.

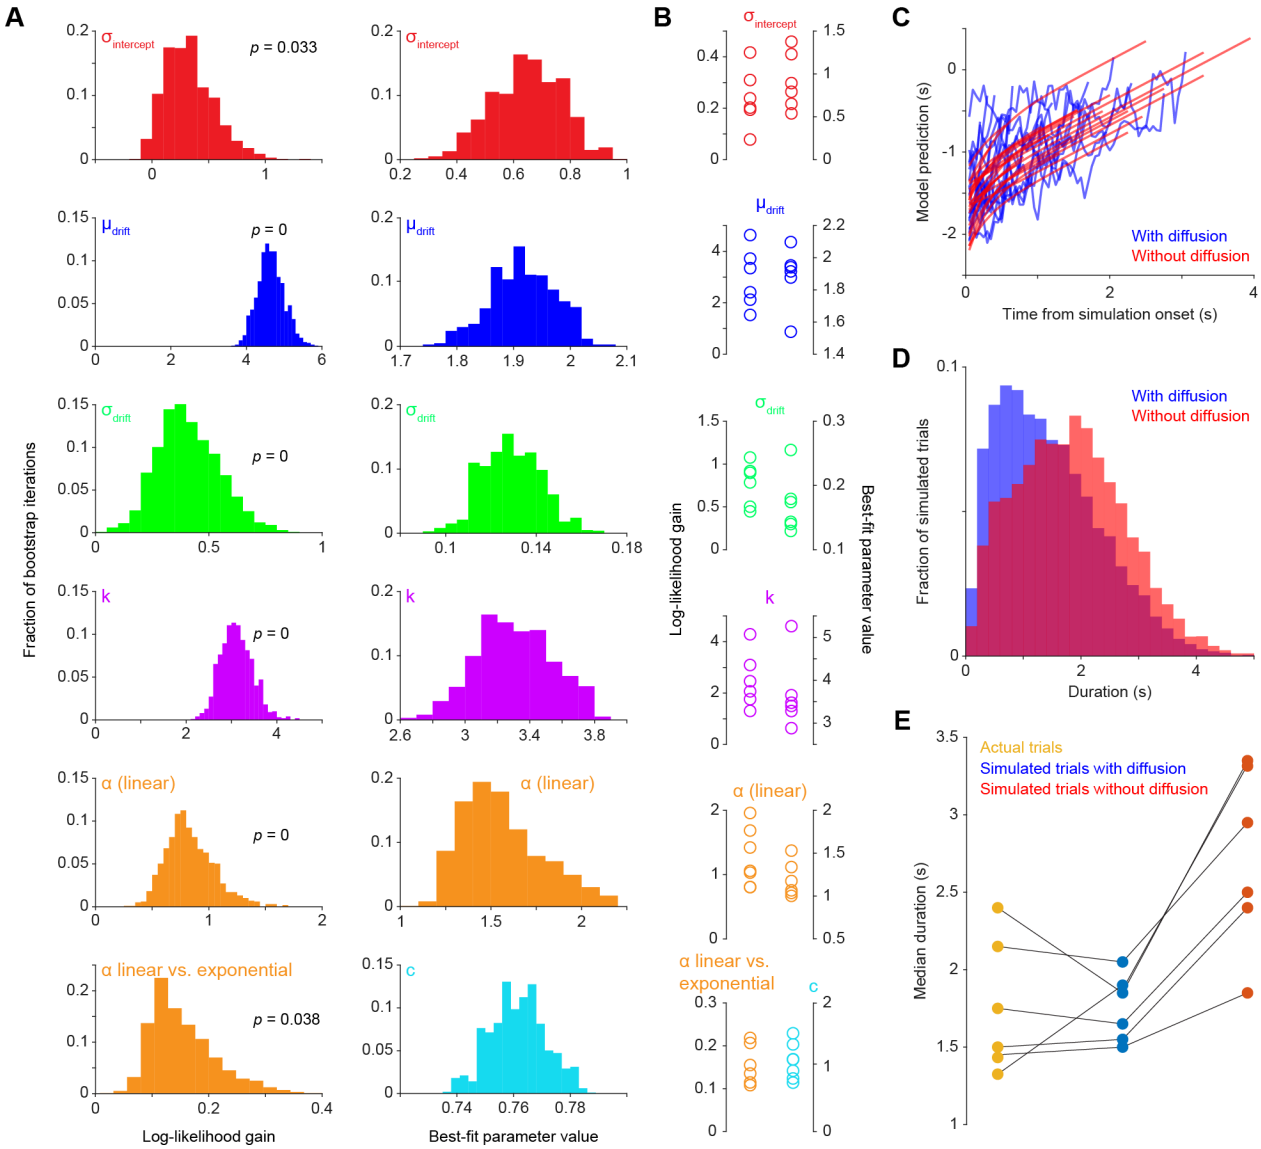

Figure 5. Modeling reveals a trial- and time-dependent deterministic drive and trial-shortening within-trial noise.

(A) *Top five rows*: Histograms of the log-likelihood gain from including each parameter (rows) in the model (left column) and the best-fit parameter value (right column) across bootstrap iterations for one example session. *Bottom row*: for the same session, log-likelihood gain for linear versus exponential drift rate growth (left) and the best-fit values for the diffusion amplitude c (right) across bootstrap iterations. Note that μ_{drift} and σ_{drift} correspond to the mean and standard deviation of the underlying normal distribution for $\log(I_0)$ (see Figures S6A and S6B for conversion to the log-normal distribution scale, and Methods). **(B)** Median log-likelihood gain (left column) and best-fit parameter values (right column) for individual mice (circles). **(C)** Simulations of the decision variable over time using the best-fit DDM for one example session, either with (blue) or after excluding (red) the diffusion term. **(D)** Distribution of simulated trials' duration (time to threshold) using the best-fit DDM for one example session, either with (blue) or after excluding (red) the diffusion term. **(E)** Median

duration of actual trials (yellow), and simulated trials using the best-fit DDM for each session, either with (blue) or after excluding (red) the diffusion term. The duration of actual trials excludes 0.2 s after immobility onset and 0.25 s before action onset as the corresponding time-to-movement predictions, which are fit to the DDMs, exclude these epochs (Methods). Accordingly, the first 0.2 s of simulated trials are discarded before computing their duration.

Within-trial noise is not required for action initiation but contributes to action timing

Having established the trial- and time-dependent nature of the deterministic drive, we then sought to examine how within-trial noise contributes to action timing. Previous work using the typical DDM to account for trial-averaged EEG data has suggested that within-trial noise plays a prominent role, being necessary to reach the activity threshold (Figure S6D)[23,37]. Since noisy fluctuations are inherent to neural activity and, consequently, to our forecast traces, a diffusion term will necessarily improve the fit of our models to forecast model output. Thus we did not test whether including the diffusion term improved fits, but rather investigated its contribution to timing variability and whether it is necessary to reach the activity threshold.

To do this, we performed an *in silico* ablation analysis: we simulated trials using our DDM variant fit to data from individual sessions, and then repeated these simulations without the diffusion term. We found that the activity threshold was reached in both cases (Figure 5C). However, including the diffusion term did shorten the duration of the trials by 22 to 45% across mice (Figures 5D and 5E). As expected, the presence of the diffusion term also increased the variance in trial duration, by 16 to 36% across mice. These results suggest that while not required to self-initiate actions, within-trial noise contributes substantially to setting action timing.

As a control, we assessed the possibility that autocorrelation in the within-trial noise may further improve our model[38,91,92]. We augmented our DDM variant with an eighth parameter, a spectral exponent β, which introduces dependence across time points in the diffusion term. However, adding this parameter did not yield a significantly better fit compared to using temporally uncorrelated noise (Figures S6E and S6F).

Trial-dependent deterministic drive and trial-shortening within-trial noise are found brain-wide

Finally, we asked whether our modeling results using multi-region activity were also locally preserved within individual brain areas. We repeated the model selection with our DDM variant fit to forecast output derived from each region individually. This analysis revealed that for all eight regions of interest, model fits were significantly improved by incorporating trial-to-trial variability in the decision variable's initial value ($\sigma_{intercept}$) and drift rate (σ_{drift}), as well as a within-trial increase of the drift rate (α; Figure 6A). Moreover, the best-fit parameter values were remarkably similar across regions (Figures 6B, S6G and S6H). We again found that the diffusion term hastened action onset in each region (Figure 6C). These results demonstrate that a common decision process is reflected across all regions we examined. This further supports a broadly distributed decision process coordinated across regions.

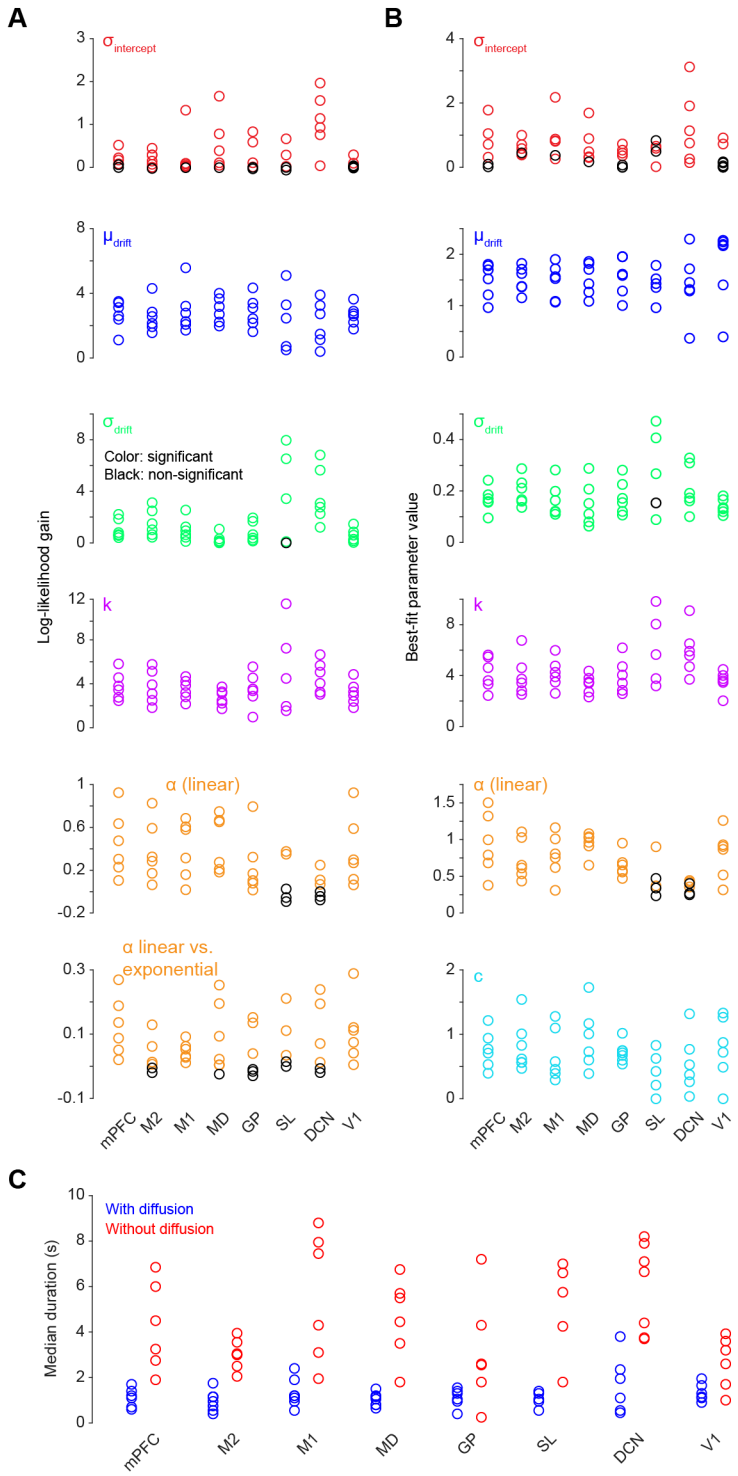

Figure 6. A similar decision process is reflected across brain regions.
(A) For individual regions (columns) in each mouse (circles), the median log-likelihood gain across bootstrap iterations from including each parameter (rows) in the model using the forecast traces (top five rows) and from linear versus exponential drift rate growth (bottom row). Colored circles in (A) and (B) indicate a significant improvement in model fit. **(B)** For individual regions (columns) in each

mouse (circles), the median best-fit parameter values across bootstrap iterations. Note that μ_{drift} and σ_{drift} correspond to the mean and standard deviation of the underlying normal distribution for $\log(I_0)$ (see Figures S6G for conversion to the log-normal distribution scale, and Methods). **(C)** Median duration of simulations (time to threshold crossing) generated with the single-region best-fit DDMs, with (blue) and after excluding (red) the diffusion term, across sessions (circles).

Discussion

In this work, we sought to shed light on the neural basis of self-initiated timing decisions by resolving single-trial, single-unit activity across multiple brain regions simultaneously. We found that action onset timing can be predicted from spiking activity up to several seconds in advance, and predictions are temporally aligned throughout the decision epoch across cortex, thalamus, basal ganglia, and cerebellum. This implies that a prominent, brain-wide deterministic drive underlies such decisions.

By assessing different potential sources of variability in action timing, we found that the deterministic drive varies in both its initial value and its rate across trials, and its rate increases over time within trials. In contrast to recent models, within-trial noise is not required for crossing the activity threshold, but does shorten the average time to threshold crossing. Similar decision-related activity dynamics were observed in all brain regions we examined, indicating a broadly shared decision process. Altogether, our results offer a unifying account in which trial-by-trial variability in deterministic drive parameters and within-trial noise act in concert to render the timing of self-initiated actions unpredictable.

Our results imply that a buildup of neural activity preceding action onset is a genuine feature of the decision process, not an artifact of trial-averaging. Yet, they stand in contrast to classic views of a purely deterministic process in which action onset is set whenever the buildup starts[3,4,7,14-17,19-22]. Our results also stand in contrast to more recent accounts whereby noisy neural activity within trials is the sole driver[23,28], or principal driver[25], of timing decisions. This aligns with a previous proposal[37] that the early predictability of action timing shown in reference[25] is more consistent with a strong deterministic drive than with accumulation of within-trial noise.

Our ability to resolve a long-standing debate regarding the deterministic versus stochastic underpinnings of self-initiated timing decisions stems from two key features of our approach. First, the vast majority of neural data recorded in relation to self-initiated decisions has been obtained in humans, using methods whose resolution necessitate trial-averaging. Yet single-trial resolution is required to arbitrate between competing models. For instance, in the "stochastic accumulator model", developed by fitting the typical DDM to such trial-averaged data, within-trial noise is both the sole source of timing variability across trials and is required for threshold crossing (Figure S6D)[14,23]. However, these same data are equally well fit by models such as the linear ballistic accumulator, which lacks within-trial noise, and in which the decision process is fully deterministic once its parameters are set when the pre-movement buildup starts[37,79]. Second, previous models of self-initiated decisions have treated within-trial noise and trial-to-trial variation in deterministic drive parameters as mutually exclusive[14,37], though the same has not been true for models of perceptual decisions[30,36]. By combining these potential mechanisms within a single model, we resolved a dichotomy and established that they act synergistically.

Here we uncovered a temporal alignment in decision-related activity across eight brain areas, from soon after immobility onset to action onset. This alignment was best illustrated by the striking similarity of the errors in time-to-movement predictions (Figure 3D). Key to this observation was our simultaneous spike-resolution recording of large neuronal populations in multiple regions, which allowed the comparison of predictive dynamics across different regions on the same single trials. Such brain-wide coordination appears to contrast with previous suggestions of a more modular

organization whereby, for example, deterministic and stochastic components originate in mPFC and M2 respectively, with activities predictive of action timing confined to M2[46]. A more distributed picture emerges from our results, with predictive activity detectable in both of these regions, as well as several others. Contrary to the classical view of a hierarchical organization in which regions perform distinct roles in a feedforward sequence, our results are more consistent with a unified decision-making process that unfolds synchronously across disparate brain regions, relying on recurrent interactions between regions. These brain-wide dynamics we found in relation to self-initiated decisions echo the broad coordination reported recently during motor execution[74,75,93,94] and perceptual decision-making[95-101], including regions historically seen as inherently motor (e.g., M1) or sensory (e.g., V1).

Our results also reveal additional parallels with perceptual decisions, which remain far more extensively studied than self-initiated decisions. In tasks that require waiting, neural activity that builds up at rates that scale with the waiting time has been observed across cortical regions and the basal ganglia[30,73,102-106]. Moreover, our findings of a drift rate that increases within a trial are reminiscent of urgency signals proposed to underlie perceptual decisions[36,80-88]. These urgency signals have been modeled by either a time-varying drift rate or a collapsing threshold. Thus, urgency may be a general organizing principle of decision dynamics, even in the absence of reward or other incentives to act quickly. This mechanistic similarity between sensory-guided and self-initiated decisions appears well-suited for negotiating complex natural environments, where behavior is rarely dictated purely by internal or external factors, but rather emerges from a dynamic integration of both.

However, while mechanistic features appear conserved between self-initiated and sensory-guided decisions, the nature of the decision process differs fundamentally in its input and adaptive function. In perceptual tasks, the deterministic drive reflects the strength of external sensory evidence, and stochasticity primarily represents perceptual noise that can be detrimental to choice accuracy[107]. In contrast, in the absence of external cues, the deterministic drive we observed likely reflects an internal imperative to act[14] ensuring that an action will be executed, while stochasticity gives rise to an advantageous unpredictability of behavior.

An important question unaddressed by our results concerns the neurophysiological underpinnings of the sources of action timing variability we have described. Our analysis implicates three such sources: trial-to-trial variability in the initial value and drift rate of the deterministic drive, and noise within trials. While within-trial noise is likely attributable to aspects of neuronal physiology, like ion channel gating and synaptic transmission[39], the origins of discrete sources of variability are more obscure. Neuromodulatory systems may be one possible source. Notably, the basal forebrain cholinergic system has been causally implicated in self-initiated action timing[47,59,60] and serves as the primary source of acetylcholine to the cortex[108], including the prefrontal and somatomotor regions where we observed predictive activity. Cholinergic input can modulate correlations across a neuronal population[109], and acetylcholine appears capable of modulating cortical function on the timescale of individual trials[110-116]. Thus, the basal forebrain cholinergic system could participate in regulating the gain of the deterministic drive (i.e., the drift rate) on individual trials.

Supplementary Figures

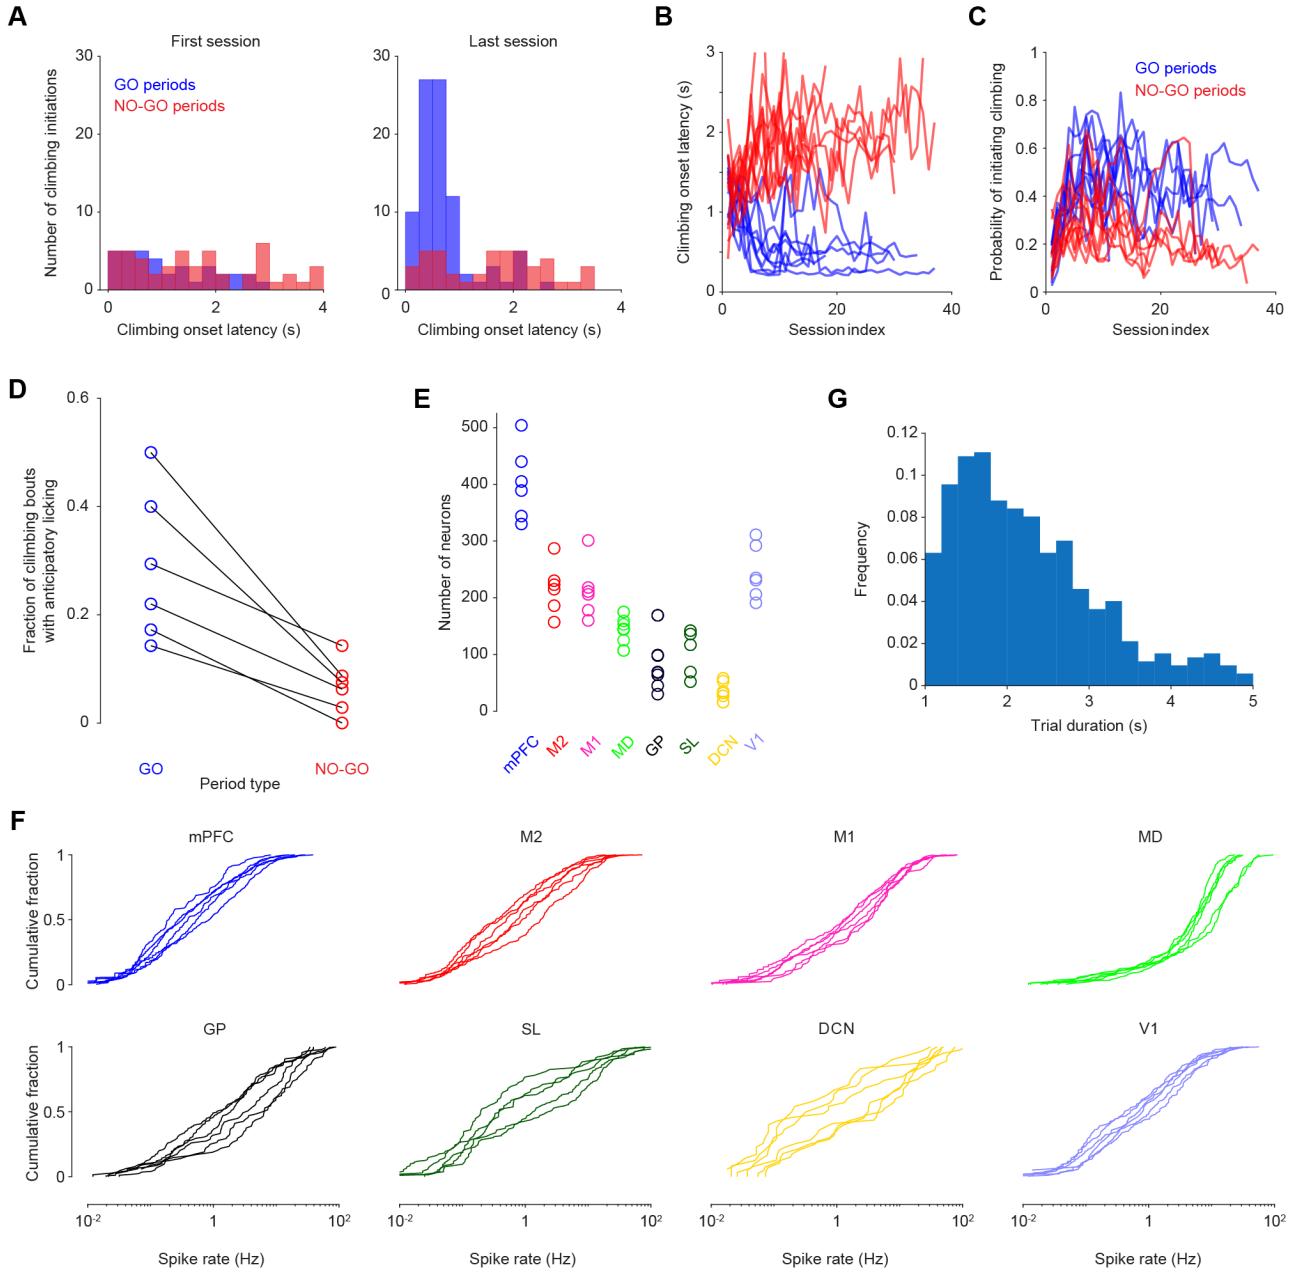

Figure S1. Behavioral characterization and electrophysiological yield.

(A-C) Evolution of climbing behavior across learning. Over the course of training, animals learn to rapidly initiate climbing upon GO onset, while maintaining a baseline rate of self-initiated climbing during NO-GO periods. These panels show data from all trained animals (besides the ones for which we carried out electrophysiological recordings) across all learning sessions. Because video monitoring was only performed during electrophysiological recording sessions, climbing onset times were here determined exclusively using wheel data. For NO-GO periods, the latency of the first climbing initiation only is considered in the event of multiple initiations. (A) Histograms of climbing onset latency during GO (blue) and NO-GO (red) periods for a single representative mouse, for the

first training session (left) and the last training session (right). (B) Median climbing onset latency across training sessions for all mice. Each line represents an individual mouse for GO (blue) and NO-GO (red) periods. (C) Probability of initiating climbing across training sessions for GO (blue) and NO-GO (red) periods. Each line represents an individual mouse. **(D)** Fraction of climbing bouts that included anticipatory licking during GO and NO-GO periods. This analysis is restricted to the recording sessions, where video data was recorded and used to track licking behavior. Each connected line represents data from a single recording session. **(E)** Number of qualifying single units per recorded region across all six recording sessions, with each circle representing one session. **(F)** Cumulative distributions of spike rates for single units within each region, with each trace representing one session. **(G)** Histogram of trial durations pooled across all recording sessions.

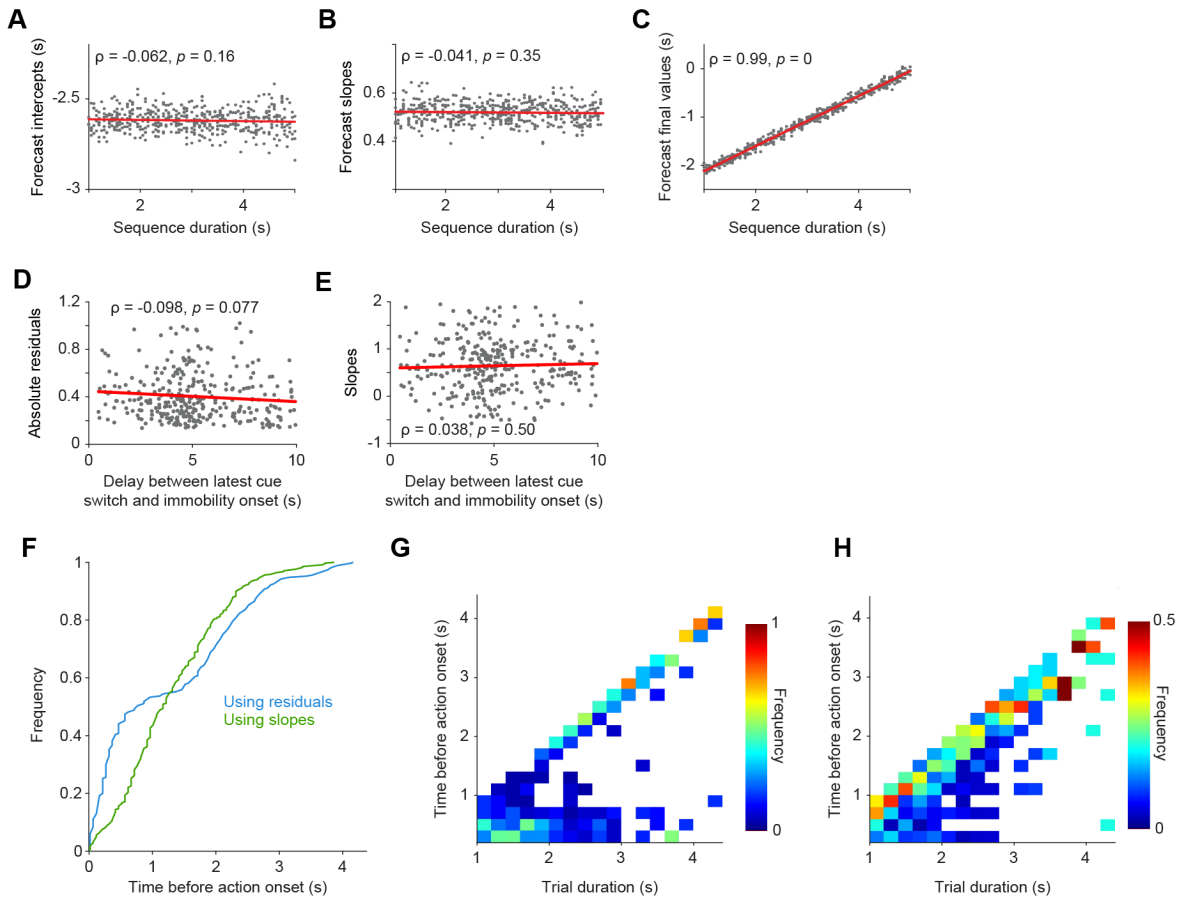

Figure S2. Forecast controls and emergence of single-trial predictability.
(A-C) Properties of forecast traces obtained using synthetic "clock-like" activity. (A) Lack of correlation between the initial forecast values (intercepts) and synthetic trial duration. ρ denotes the Pearson correlation coefficient. (B) Lack of correlation between the forecast slopes and synthetic trial duration. (C) Positive correlation between the final forecast values and synthetic trial duration. These properties are qualitatively opposite to those obtained using the actual neural data (Figures 4A-4D). **(D, E)** Lack of correlation between the delay from the latest cue switch to immobility onset and either the absolute residuals (D) or slopes (E) of the forecast traces. For each trial, metrics were computed over the first second of the forecast traces. All data shown are from cross-validated, held-out test trials. **(F)** Cumulative distribution of predictability onset times relative to climbing onset across all trials. Figure 2H shows the same data relative to immobility onset. **(G, H)** 2D histograms depicting the emergence of predictability relative to climbing onset as a function of trial duration. The time of predictability emergence is calculated using either residuals (G) or slopes (H).

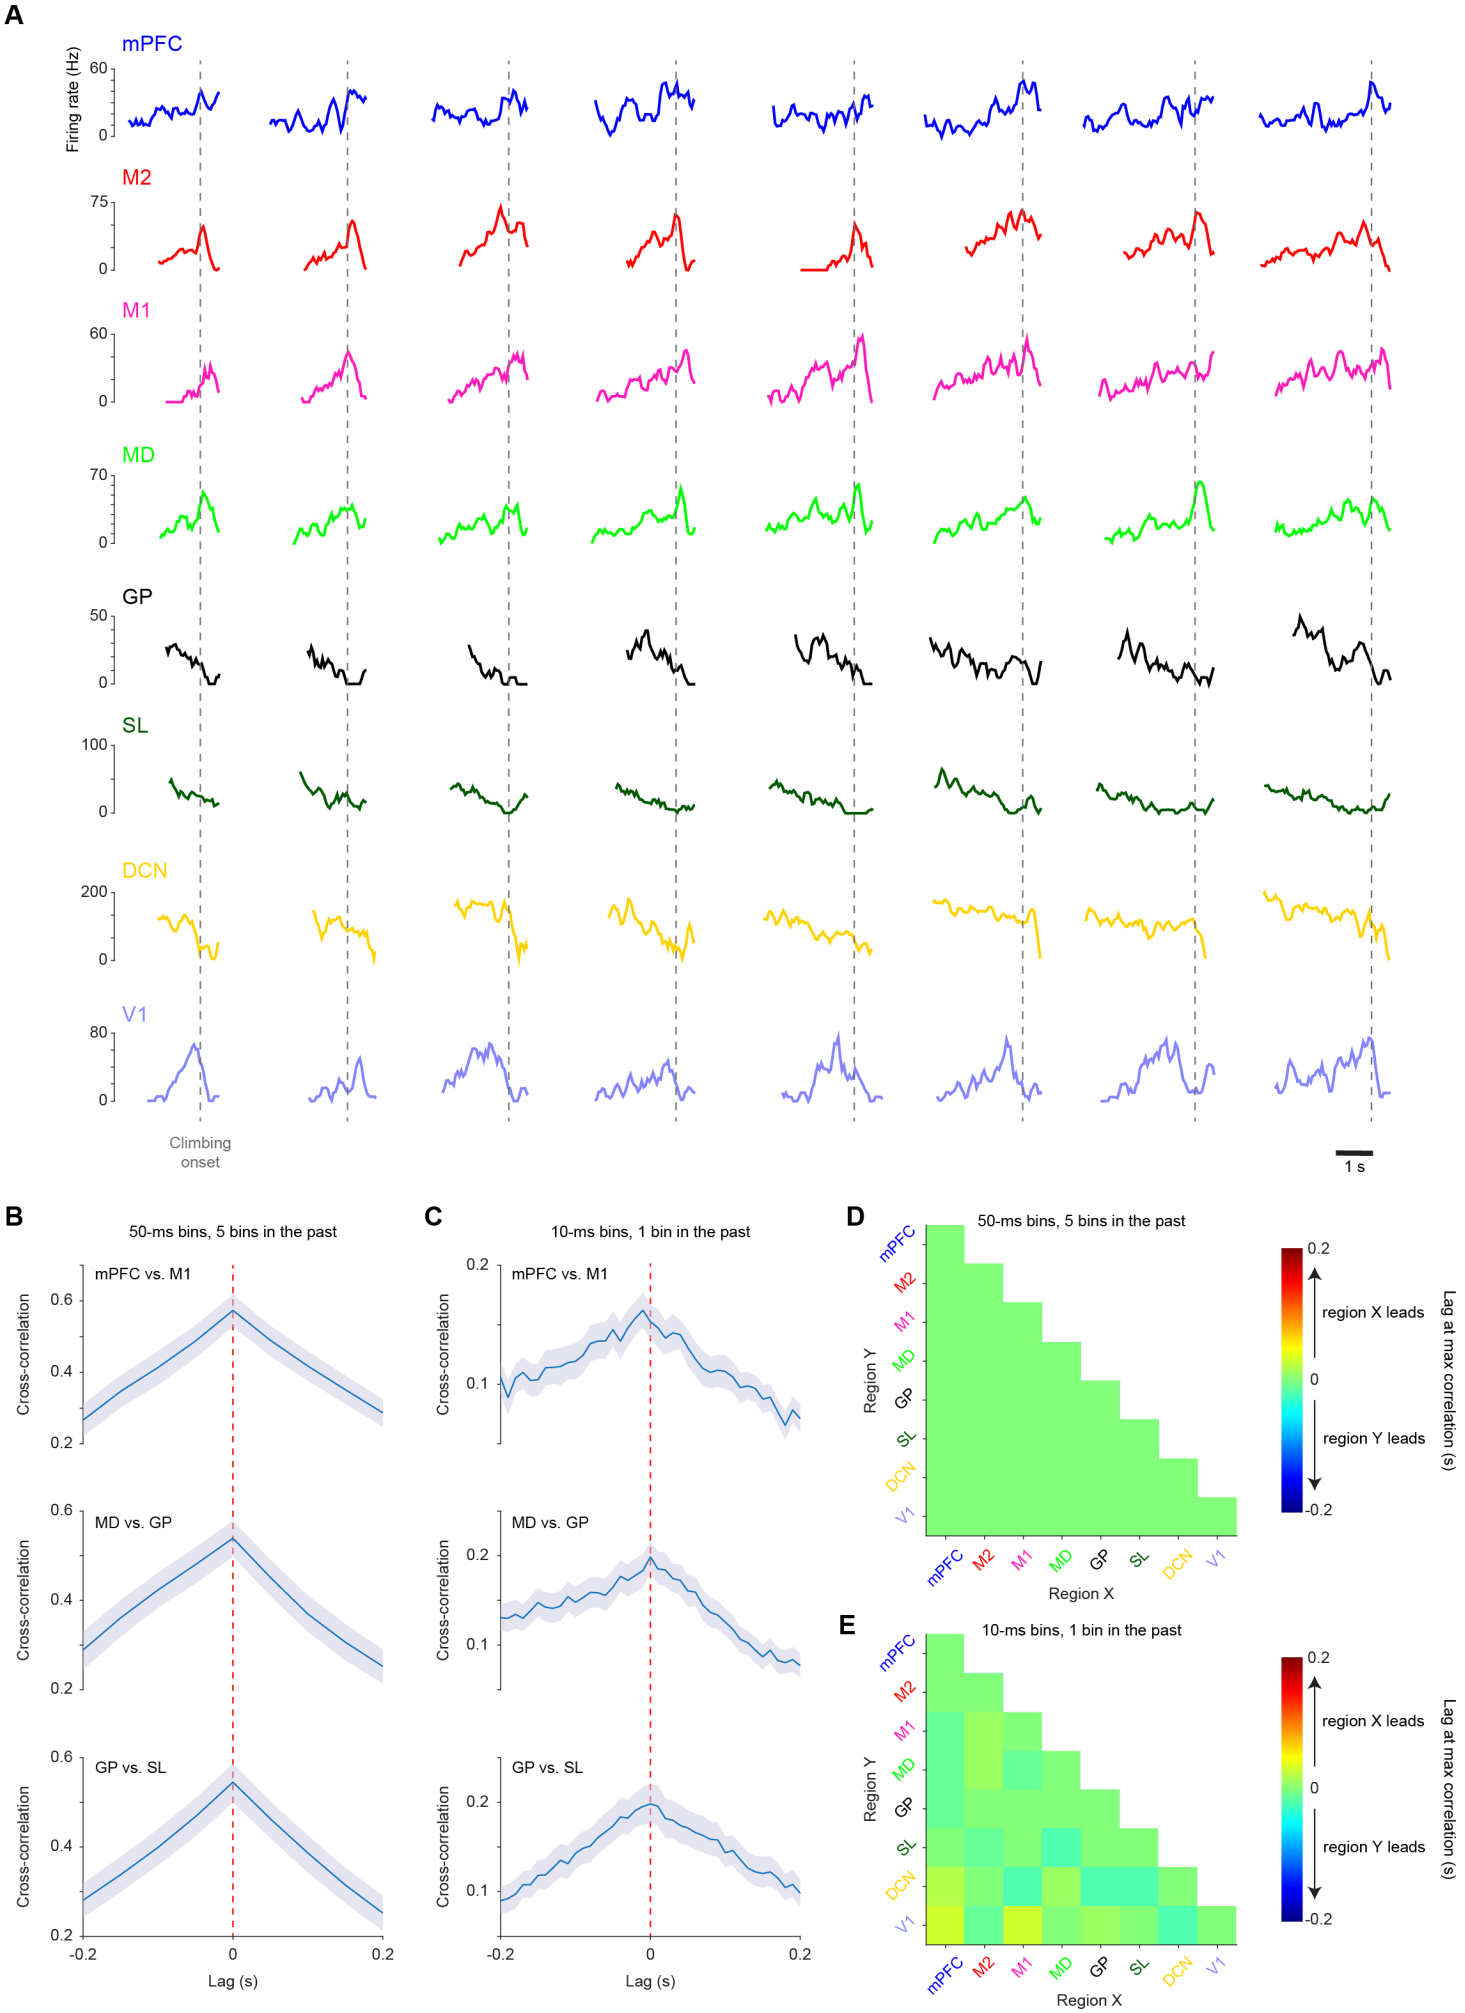
A
mPFC
M2
M1
MD
GP
SL
DCN
V1
Firing rate (Hz)
Climbing onset
1 s
B
50-ms bins, 5 bins in the past
mPFC vs. M1
MD vs. GP
GP vs. SL
Cross-correlation
Lag (s)
C
10-ms bins, 1 bin in the past
mPFC vs. M1
MD vs. GP
GP vs. SL
Cross-correlation
Lag (s)
D
50-ms bins, 5 bins in the past
Region Y
Region X
region X leads
region Y leads
Lag at max correlation (s)
E
10-ms bins, 1 bin in the past
Region Y
Region X
region X leads
region Y leads
Lag at max correlation (s)

Figure S3. Predictive activity across brain regions.
(A) Single-trial firing rate of an example forecasting neuron from each of the eight regions of interest, aligned to climbing onset. Causal moving average was applied: the value at each time point was calculated by averaging the data from that point with those of the preceding 200 ms window. Trials are sorted by increasing immobility durations. These example neurons correspond to the ones whose trial-averaged firing rates are shown in Figure 3C. **(B-C)** Cross-correlation of time-to-movement predictions across brain regions, averaged across trials (mean ± 95% confidence interval estimated from bootstrapping). A positive lag at maximum cross-correlation indicates that the first region listed (i.e., region X in "region X vs. region Y") is leading. (B) Cross-correlation for three example region pairs using our standard forecast models: trained and tested with spike trains convolved with a Gaussian kernel (σ = 10 ms), binned into 50-ms intervals, and using 5 past bins for each prediction. (C) Cross-correlation for the same three region pairs using forecast models trained and tested with spike trains without convolution, binned into 10-ms intervals, using 1 past bin for each prediction. **(D)** Matrix of lags at maximum cross-correlation for all region pairs using the standard forecast models described in B. **(E)** Corresponding lag matrix using the forecast models without convolution described in C. No lag was significantly different from 0 (non-parametric bootstrap test, FDR-corrected p-value > 0.05).

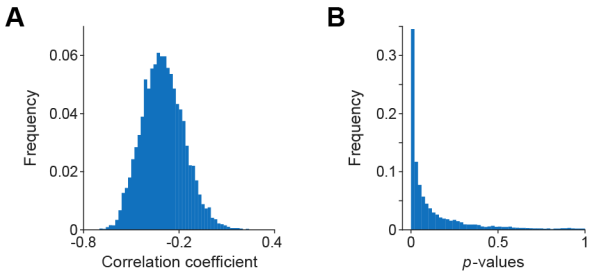

Figure S4. Robustness of the negative correlation between initial forecast values and trial duration (established in Figure 4A) assessed via random subsampling.
(A) Distribution of Pearson correlation coefficients across 10,000 iterations, where 10% of the trials were randomly retained, without replacement, in each iteration. The resulting distribution remains predominantly negative and centered around the mean value obtained from the full dataset. **(B)** The corresponding distribution of *p*-values. Together, these distributions confirm the stability of the relationship even with substantially reduced statistical power.

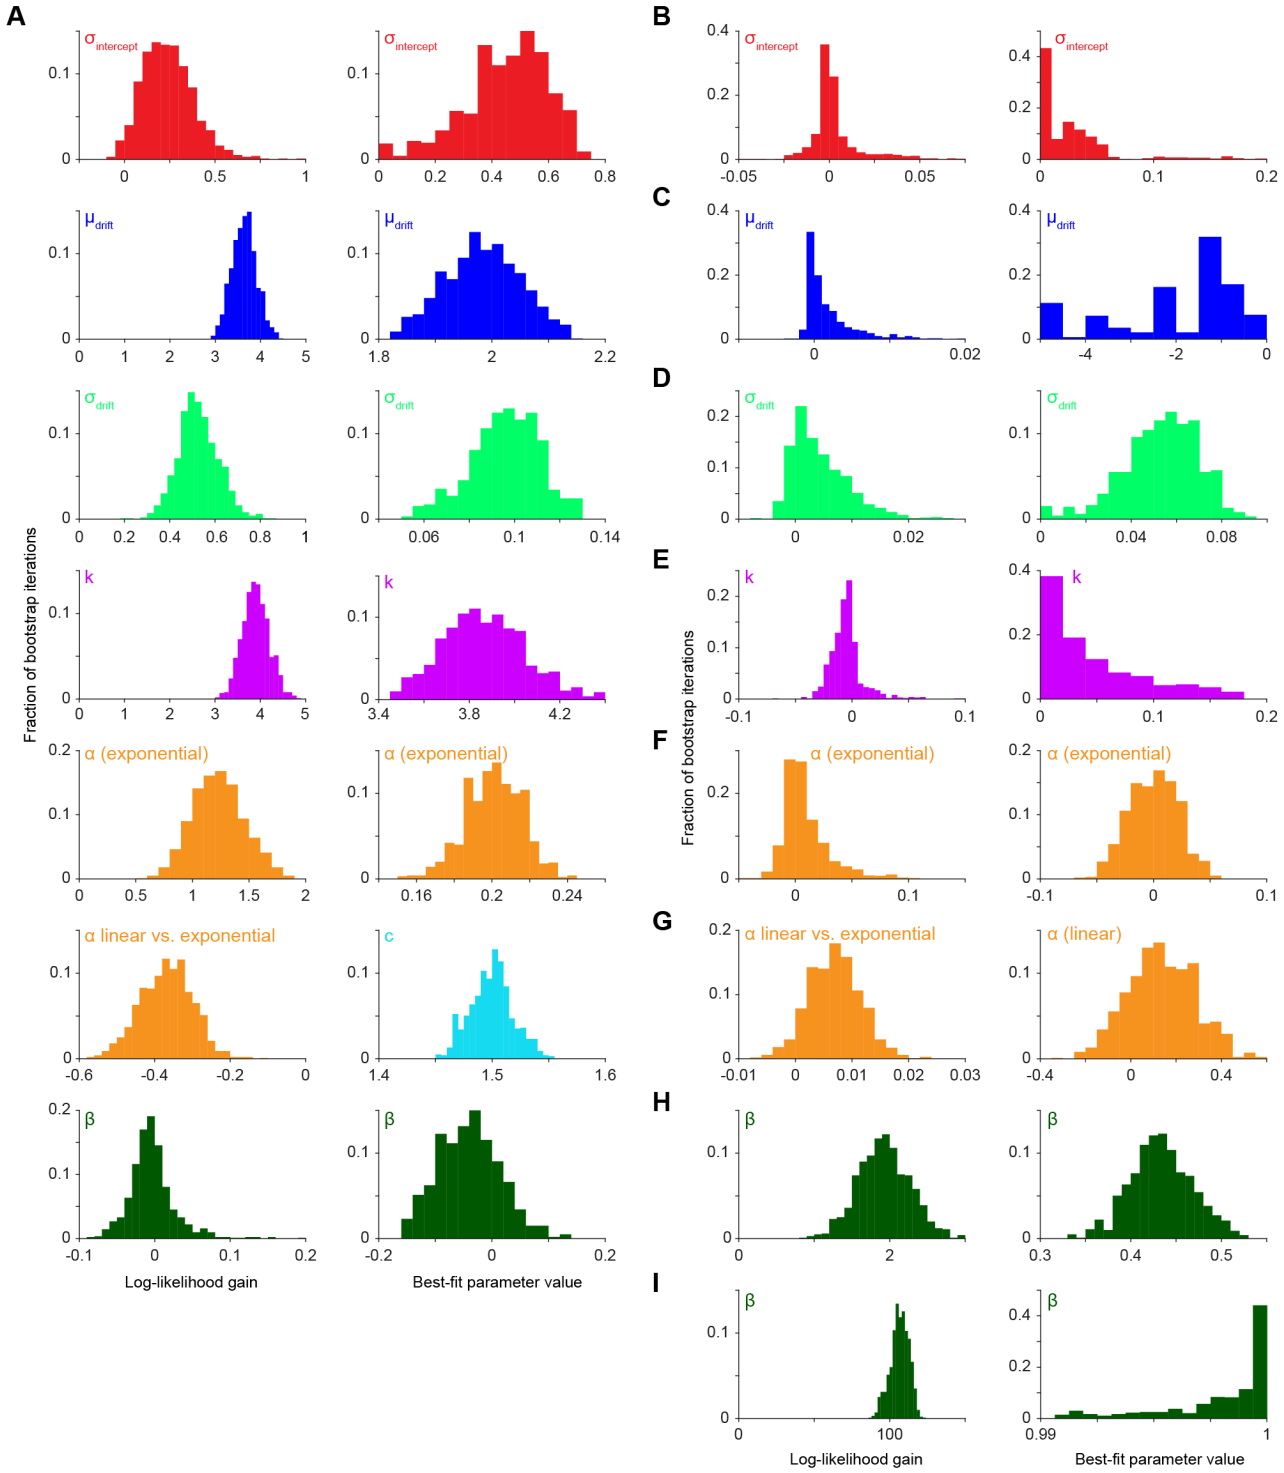

Figure S5. Validation of the DDM fitting and model selection procedure via parameter recovery. To mimic the fitting conditions of forecast traces, whose first point corresponds to 250 ms after immobility onset, all simulated sequences were fit after truncating their first 200 ms. Distributions show the results across bootstrap iterations. For each panel, the left column shows the log-likelihood gain, and the right column shows the best-fit parameter value. A positive likelihood value

for the "α linear vs. exponential" comparison denotes a higher likelihood for the model with a linear growth rate. **(A)** Recovered parameter distributions for sequences generated using our 7-parameter DDM variant ($\beta = 0$). Ground-truth parameters for generating the sequences were: diffusion $c = 1.5$, leak $k = 4$, time-dependent drift $\alpha = 0.2$ (exponential growth), mean initial value $\mu_{intercept} = 1$, trial-to-trial variability in initial value $\sigma_{intercept} = 0.5$, mean drift rate $\mu_{drift} = 2$, and trial-to-trial variability in drift rate $\sigma_{drift} = 0.1$. The procedure accurately recovers the ground-truth values for the 7 parameters and correctly identifies β as non-significant (log-likelihood gain ≈ 0). **(B-F)** Parameter recovery using reduced 6-parameter generative models. Sequences were generated with one specific parameter omitted (set to 0) to confirm that the model selection procedure correctly assigns a log-likelihood gain near zero for the missing parameter: (B) $\sigma_{intercept} = 0$ (similar initial value across trials), (C) $\mu_{drift} = 0$ (no drift), (D) $\sigma_{drift} = 0$ (similar drift across trials), (E) $k = 0$ (no leak), and (F) $\alpha = 0$ (constant drift within a trial). **(G)** Parameter recovery for sequences generated with a linear rather than exponential drift growth ($\alpha = 0.2$). **(H, I)** Recovered parameter distributions for sequences generated with a full 8-parameter model, demonstrating the procedure's ability to successfully extract the additional parameter β when set to (H) $\beta = 0.5$ and (I) $\beta = 1$.

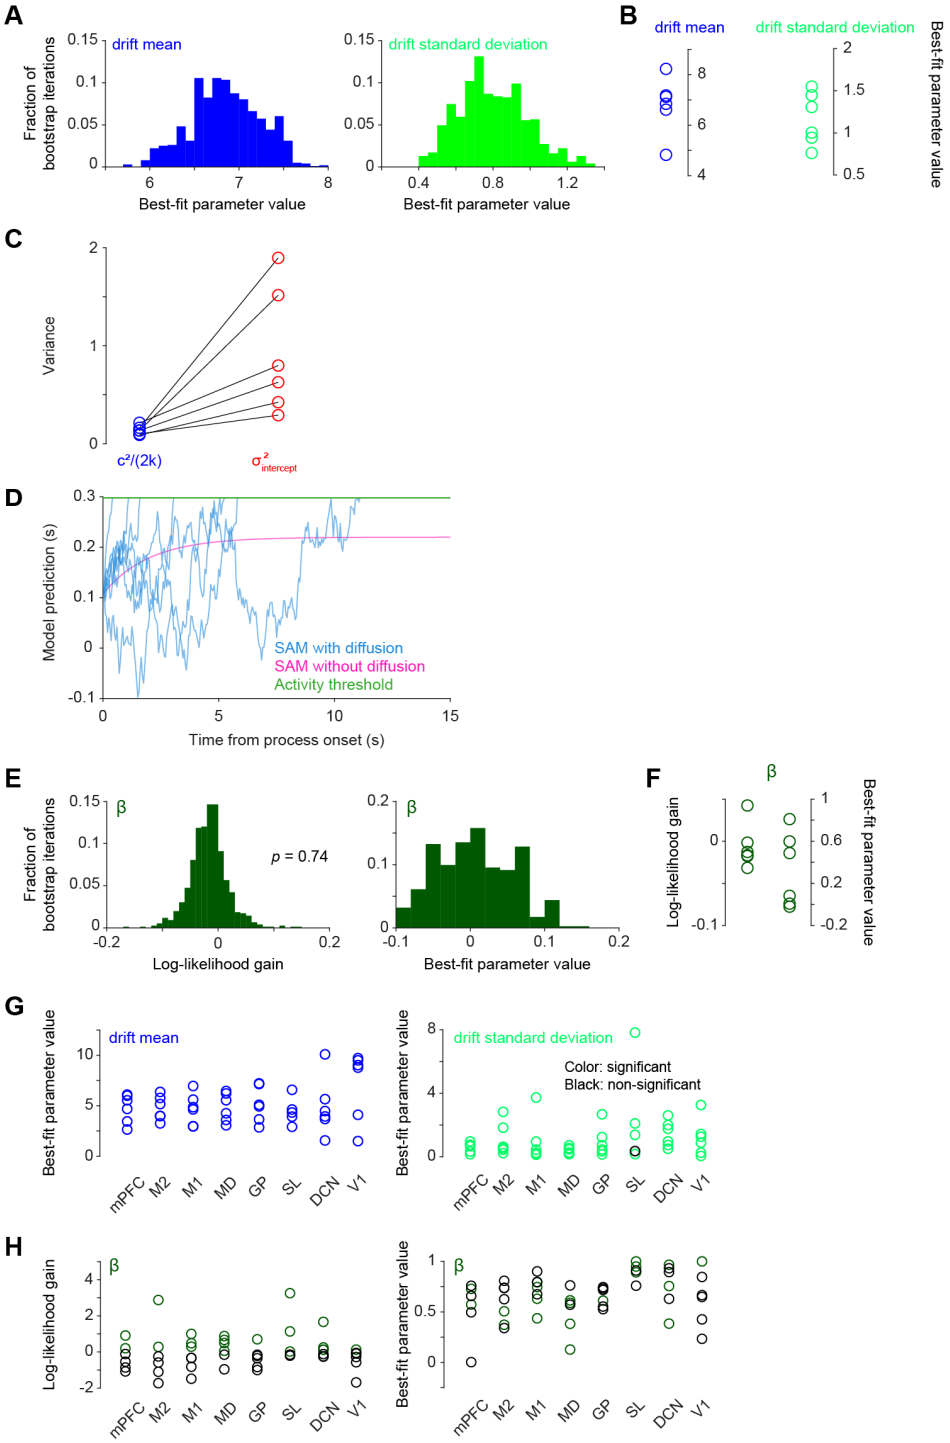

Figure S6. DDM controls and comparison with the "stochastic accumulator model".
(A, B) Best-fit parameter values for the mean and standard deviation of the initial drift rate I_0. As noted in Figure 5, the fitted parameters μ_{drift} and σ_{drift} define the mean and standard deviation of the underlying normal distribution for $\log(I_0)$ [Equation 3]. Here, drift mean (blue) and drift standard deviation (green) represent the corresponding mean and standard deviation of the log-normal distribution itself, from which the initial drift I_0 is drawn on a per-trial basis. (A) Parameter

distributions across bootstrap iterations for the same illustrative session as in Figure 5A. (B) Median best-fit parameter values for each of the six mice (corresponding to the data in Figure 5B). Each circle represents a mouse. **(C)** Comparison of steady-state variance and initial variance. Comparison between the theoretical asymptotic variance limit of the leak process ($c^2/(2k)$, blue) and the trial-to-trial variance in initial value ($\sigma^2_{intercept}$, red) for each of the six mice. Lines connect data from the same recording session. **(D)** Simulations based on a typical DDM [Equation 1] using parameters from a previous study[23] which define the "stochastic accumulator model" (SAM; activity threshold = 0.298, $I = 0.11, c = 0.1, k = 0.5$). The deterministic component alone (SAM without diffusion, magenta trace) plateaus and fails to reach the activity threshold (green line). In this framework, the addition of a diffusion term (SAM with diffusion, blue trace) is strictly required for threshold crossing. **(E, F)** Model selection provides no evidence for temporally autocorrelated (colored) noise. To test whether the diffusion term had temporal correlations, the baseline DDM was augmented with a spectral exponent β, where $\beta = 0$ corresponds to white noise. (E) Distributions of the log-likelihood gain (left) and the best-fit β value (right) for the same illustrative session as in Figure 5A. The inclusion of β did not yield a significant gain. (F) Summary of colored noise testing for all six mice. Shown are the median log-likelihood gains (left) and best-fit parameter values (right). The inclusion of β did not significantly improve the fit for any of the mice. Together, these results justify the use of our 7-parameter model with white noise in the Main text. **(G, H)** Region-specific parameter conversions and colored noise controls. (G) Best-fit parameter values for the mean (left) and standard deviation (right) of the initial drift rate I_0 for each individual brain region across all sessions, corresponding to the models evaluated in Figure 6. (H) Region-specific evaluation of temporally autocorrelated noise (β). The inclusion of β did not significantly improve the fit for the majority of individual sessions across regions.

Material and Methods

Experimental subjects

All experiments and procedures involving animals were performed in accordance with NIH guidelines and approved by the Institutional Animal Care and Use Committee of Northwestern University. A total of 26 adult C57BL/6J male mice (Jackson Laboratories stock #000664) were used in reported experiments and early pilot studies to establish methodology. All mice were individually housed under a 12-hour light/dark cycle during the course of the experiment. At the time of reported measurements, animals were 12-24 weeks old and weighed approximately 24-30 g. All animals were used for the first time in this experiment and were not previously exposed to pharmacological substances or altered diets.

Surgical procedures

All surgeries were performed under isoflurane anesthesia (1-3%) in a stereotaxic frame. The surgical workflow consisted of two main procedures separated by a recovery and behavioral training period.

Headplate Implantation

First, mice were implanted with a custom 3D-printed plastic headplate. The headplate, which features an open-center design for cranial access, was affixed to the skull using dental cement (Metabond, Parkell). The position of bregma relative to marks on either side of the headplate was measured to facilitate the positioning of craniotomies during the subsequent surgery. The skull within the open window was then covered with a protective layer of dental cement. After a one-week recovery period, mice were placed on a water schedule (1-1.5 ml per day) and began behavioral training.

Craniotomies

After the completion of behavioral training, a second surgery was performed. The protective dental cement was removed to expose the skull, and four craniotomies (0.5-1 mm diameter) were drilled over the target regions. A platinum-iridium ground wire soldered to a connector pin was implanted in a separate craniotomy over the right frontal lobe and cemented to the headplate implant. Silicone elastomer (Kwik-Cast, World Precision Instruments) was used to cover the craniotomy after surgery.

Probe insertion trajectory planning

Probe insertion trajectories were planned using a reference mouse brain atlas[117] and the Allen Mouse Brain Common Coordinate Framework (CCFv3)[118], aided by the "Neuropixels trajectory explorer software" (https://github.com/petersaj/neuropixels_trajectory_explorer). Targets in the left hemisphere included mPFC, M2, M1, MD, and GP. Targets in the right hemisphere were V1, SL and DCN. Table S1 lists the coordinates for the four probe trajectories (AP: antero-posterior; ML: medio-lateral; DV: dorso-ventral).

Table S1. Stereotaxic coordinates for craniotomies and Neuropixels probe trajectories.

Trajectory	Hemisphere	Target regions	Craniotomy coordinates (AP, ML)	Top electrode (AP, ML, DV)	Bottom electrode (AP, ML, DV)
1	Left	mPFC, M2	2.25, 1.04	2.35, 1.27, 0.38	1.92, 0.24, 4.06
2	Left	M1, GP	-0.60, 1.11	-0.60, 1.20, 0.69	-0.60, 1.70, 4.50
3	Left	MD	-1.70, 1.80	-1.61, 1.57, 0.45	-1.04, 0.08, 3.94
4	Right	V1, SL, DCN	-3.91, -3.28	-4.15, -3.20, 0.58	-6.49, -2.34, 3.50

Climbing apparatus

The behavioral apparatus, described in detail previously[67], was housed inside a sound-attenuating chamber (H10-24A, Coulbourn). The core of the apparatus was a custom, 3D-printed cylindrical wheel featuring handholds positioned 12° apart. A ratchet mechanism permitted the wheel to rotate only in one direction, downward as the mouse climbed. The wheel's absolute angular position was continuously tracked by a shaft encoder (A2-A-B-D-M-D, U.S. Digital). Experimental control utilized custom MATLAB (MathWorks) scripts and a DAQ card (PCI-e-6323, National Instruments).

Behavioral training

Mice were initially acclimated to being handled by the experimenter, to head-fixation, and to the climbing apparatus according to previously described procedures[67]. Briefly, following acclimation to handling and head-fixation in front of the wheel, mice underwent several daily training sessions to associate rotating the wheel with a water reward. Initially, rewards were delivered manually for any wheel rotation. Subsequently, an automated protocol rewarded mice for performing progressively longer climbing bouts, with the reward amount for each bout adapted based on the mouse's performance on recent bouts. This helped maintain motivation.

Once mice were proficient, they were transitioned to the behavioral protocol also used during recording sessions. This protocol consisted of two period types – "GO periods" and "NO-GO periods" – occurring alternately. So that mice could distinguish between periods, auditory (9 kHz pure tone; Murata Electronics 7BB-12-9) and visual (green led; Visual Communications Company CNX722C50005B) cues were presented during one of the two period types. Which period involved cues was varied across animals: for some mice, cues were active during the GO periods, while for others, they were active during the NO-GO periods. The duration of the GO and NO-GO periods was contingent on the mouse's behavior. If the animal remained immobile, the durations of both GO and NO-GO periods were randomly selected from a continuous uniform distribution ranging from 2 to 4 seconds. If the mouse climbed, period duration was adjusted as follows:

- During a GO period, a climbing bout would terminate the period and trigger a reward.
- During a NO-GO period, a climbing bout would forestall the end of the period and reset the timer from the moment immobility is detected.

The present study focuses exclusively on the self-initiated climbing bouts, which occurred during NO-GO periods. A typical behavioral session lasted one hour.

Video recording

Video tracking of the behavior was done using a high-resolution monochrome camera (acA2040-55um, Basler), fit with a 16-mm fixed focal length lens (33-304, Edmund Optics). The camera was

660 positioned to the right of the head-fixed mouse and videos were acquired continuously throughout
661 behavioral sessions. Recording was performed under near-infrared illumination at 100 frames per
662 second with a resolution of 800 x 700 pixels.

663 To ensure precise temporal alignment, video acquisition was hardware-triggered. A 100 Hz TTL pulse
664 train was generated by the NI-DAQ card, with each pulse triggering the acquisition of a single camera
665 frame. This pulse train was recorded by the NI-DAQ system, providing a common temporal reference
666 for the subsequent synchronization of the video frames with all other data streams, including
667 electrophysiology and behavioral signals like wheel position.

668 We used DeepLabCut[68] to track the positions of the mouse's right hand and foot from videos. The
669 resulting marker coordinates were then used to compute kinematic variables, such as limb velocity.

670 Electrophysiological recording

671 Prior to each acute recording session, several preparatory steps were taken. To facilitate smooth
672 penetration through the dura mater, the tips of the four Neuropixels 1.0 probes (IMEC) were
673 sharpened using a grinder (EG-45, Narishige). A few minutes before insertion, the sharpened tips
674 were coated with a fluorescent lipophilic dye (CM-DiI, Invitrogen)[119] to mark the probe tracks for post-
675 hoc histological verification, according to the following protocol:
676 https://www.protocols.io/view/painting-neuropixels-probes-and-other-silicon-prob-5jyl8neprl2w.

677 For the recording session, the mouse was head-fixed in the climbing apparatus and the silicone
678 elastomer was removed from the craniotomies. The four Neuropixels, mounted on 3-axis
679 micromanipulators (New Scale Technologies), were then slowly lowered to their target depths. The
680 exposed skull was covered with silicone oil. Electrode voltages were acquired at 30 kHz using the
681 Open Ephys acquisition software[120]. A 1 Hz TTL pulse train was used to align Open Ephys and Matlab
682 recordings, including video.

683 Spike sorting and firing rate estimation

684 Raw electrophysiological data was spike-sorted using Kilosort 3[121,122]. Only units classified as "good"
685 single units by Kilosort were retained for further analysis. To estimate the instantaneous firing rate for
686 each neuron, the binary spike trains (1 kHz) were convolved with a Gaussian kernel (σ = 10 ms)[123].
687 The resulting smoothed signals were then binned into 50-ms windows to compute the firing rate time
688 series used in subsequent analyses. An exception was made for the analyses presented in Figures
689 S3C and S3E: to capture dynamics at a faster timescale, these firing rate time series were derived by
690 binning raw spike trains into 10-ms windows without prior Gaussian smoothing. Finally, we applied a
691 selection criterion to focus on task-relevant neurons: only units with a mean firing rate of at least 0.5
692 Hz averaged across all periods from -1 s to -0.25 s relative to self-initiated climbing onset were
693 included. These units are referred to as "qualifying units".

694 Histological procedure and brain registration

695 At the conclusion of the recording sessions, mice were deeply anesthetized and transcardially
696 perfused with 3% paraformaldehyde (P6148, MilliporeSigma). Brains were extracted, post-fixed in
697 paraformaldehyde, and subsequently sectioned into 60-µm thick coronal slices using a sliding
698 freezing microtome (SM2000, Leica). The sections were then mounted on glass slides with a
699 mounting medium containing DAPI (00-4959-52, Invitrogen). Whole-slide images were acquired

using a slide scanner fluorescence microscope (VS120, Olympus). To verify the final position of probe tracks, the resulting images were registered to CCFv3[118] using the open-source Python package HERBS[124].

Climbing onset detection

To precisely identify the onset of self-initiated climbing bouts, we used a two-stage method that combined wheel rotation and limb kinematics.

1) First, during behavioral sessions, climbing bouts were detected using the angular displacement of the wheel. The shaft encoder signal, acquired at 1 kHz, was smoothed, and a climbing bout was defined as a continuous rotation exceeding an angular displacement of 10°. Immobility, in this context, was defined as less than 0.25° of rotation over 400-ms windows. This step ensured that detected movements corresponded to genuine locomotion that displaced the wheel.

2) Second, this detection was refined offline using the kinematic data from the right hand and foot to determine the precise moment of climbing onset. The DeepLabCut marker coordinates were upsampled to 1 kHz via Akima spline interpolation[125], and instantaneous velocity was calculated for each limb. A velocity threshold, empirically set to 0.1 pixel/ms to separate limb movements from baseline noise, was applied to classify time points as either mobile or immobile. Among periods of mobility, to distinguish genuine climbing from twitches, we applied the following criteria: 1) brief periods of mobility lasting less than 100 ms were classified as twitches, and 2) short pauses of immobility (< 400 ms) within a larger movement were reclassified as part of the ongoing movement. The onset of a potential climb was defined as the first moment a limb transitioned from an immobile to a mobile state that was not a twitch. For a potential climbing bout to be validated, the kinematics-defined onset had to be followed by a wheel-defined onset within 500 ms, ensuring the limb movement resulted in a productive wheel displacement exceeding 10°. Finally, the climbing bout had to be preceded by a continuous period of immobility of at least 1 second.

A single trial was defined as the continuous period of immobility preceding the onset of a climbing bout initiated during a NO-GO period.

Quantification of anticipatory licking behavior

To quantify anticipatory licking behavior, we tracked individual licks during climbing from video recordings. First, using a custom MATLAB routine, we identified potential licks by detecting contrast changes within a region of interest between the water spout and the mouse's mouth. Second, to eliminate false positives, which could occur when the mouse's paw crossed the region of interest during climbing, every potential lick was manually curated by visualizing the corresponding frames.

We compared the anticipatory lick rate between GO and NO-GO periods by calculating the fraction of climbing bouts during which at least one lick occurred. To ensure an unbiased comparison between the two conditions, we enforced two criteria:

1) To prevent reward-triggered licking from contaminating the GO-period data, climbing bouts during GO periods were truncated immediately prior to the moment of bout completion, before the animal could perceive the reward or any reward-predicting cue. This truncation point corresponded either

to the offset of the cues (for mice trained with active auditory and visual cues during GO periods) or to the moment of completion of a climbing bout (for mice trained with active cues during NO-GO periods, where GO period reward attainment is uncued prior to water delivery).

2) Because climbing bouts initiated during NO-GO periods are not rewarded, they tend to be longer than those during GO periods. To prevent this duration disparity from biasing the lick probability, we defined a session-specific "evaluation window" time-locked to climbing onset. Licking was only assessed within this window, and only climbing bouts (truncated as per criterion 1, if applicable) that met or exceeded this duration were included in the analysis. The length of this evaluation window was determined per session to resolve a trade-off: longer evaluation windows increase the cumulative probability of observing a lick, but simultaneously decrease the pool of eligible climbing bouts. To optimize this balance, we set the evaluation window to the median duration of all climbing bouts (truncated as per criterion 1, if applicable) in a given session. Consequently, our analysis retained half of the climbing bouts for each session, with licking evaluated over a period equal to this median duration.

Cross-validated predictions of time-to-movement

To test whether the timing of self-initiated actions was predictable from neural activity on single trials, we trained a linear forecast model. We chose an elastic net linear regression model[76] for its simplicity and interpretability. The model was trained to predict the time remaining until action onset ("time-to-movement") at 50-ms time steps during trials. Trials, which had variable durations, were aligned to action onset. The input features for the model at each 50 ms time step consisted of the firing rates of all qualifying units, over the preceding five time steps (250 ms). We used a 5-fold cross-validation scheme to train and test the forecast model. The regularization parameter for the elastic net was optimized for each fold using an internal 5-fold cross-validation loop performed exclusively on the training set. To assess the correlation of single-trial forecast predictions across brain regions (Figures 3D-3G), the same procedure was applied using qualifying units from individual regions.

To estimate chance-level predictions, we circularly permuted[126] the time-to-movement relative to the firing rates for the training data 1,000 times for each cross-validation fold, and tested the resulting null models on the held-out test trials. We could thus compute a null distribution for each performance metric (mean absolute residuals and best-fit slopes) for each test trial. Actual performance metric values were compared against the null distributions to compute a *p*-value for each performance metric (Figures 2C-2E).

These calculations were repeated for 250-ms windows that began every 50 ms during a trial. This produced a sequence of *p*-values over time for each trial, where each *p*-value represented the forecast's significance within that specific time window. We applied the Benjamini-Hochberg procedure[127] to control the False Discovery Rate (FDR) across all time windows within a given trial. The onset of predictability for a trial was then defined as the first time point at which the FDR-corrected *p*-value fell below the significance threshold of 0.05.

Control analysis to exclude "clock-like" activity

To test the alternative hypothesis that the neural activity used by our forecast model was merely tracking the passage of time from immobility onset, rather than reflecting a decision-related process,

we performed a control analysis using synthetic "clock-like" activity (Figures S2A-S2C). We generated a synthetic dataset consisting of trials with variable durations, randomly drawn from a uniform distribution between 1 and 5 s. For each trial, we simulated a one-dimensional latent variable modeled as a linear ramp to which we added independent Gaussian noise. This signal mimics the readout of a neural population encoding the elapsed time since immobility onset. This synthetic dataset was then processed through the exact same 5-fold cross-validated elastic net forecast procedure used with the actual neuronal data. We then compared the relationships between forecast parameters (initial value, slope, final value) and trial duration in the synthetic data (Figures S2A-S2C) versus in the actual neural recordings (Figures 4A-4D).

Random subsampling validation

We performed a random subsampling analysis to verify that the negative correlation observed between the forecast initial value and trial duration (Figure 4A) was a robust feature of our data, rather than being driven by a small subset of trials where the decision process started very early after immobility onset. We performed 10,000 iterations in which we randomly sampled 10% of the total trials without replacement. For each iteration, we computed the Pearson correlation coefficient and the associated *p*-value between the initial forecast values and the trial durations (Figure S4).

Identification of forecasting neurons

To identify which individual neurons contributed significantly to time-to-movement predictions (Figures 3A-3C), we implemented a permutation-based feature importance analysis. The procedure was repeated 1,000 times; within each iteration, the following steps were performed for each of the five cross-validation folds:

1. For each recording session, an elastic net model was trained on training trials with all qualifying units.

2. On the held-out test data, we assessed the importance of each qualifying unit that had been assigned a non-zero coefficient by the model. To do this, we created a test set where the temporal information of a single neuron was destroyed by circularly permuting its activity time-series, while leaving the activity of all other neurons intact.

3. The original, already-trained model was used to predict time-to-movement from this partially permuted test data.

4. The effect of permuting a neuron was quantified by the change in mean absolute residual (MAR) compared to the MAR on the original test data. A neuron that improves the forecast should cause the MAR to increase when its activity is shuffled.

This process yielded a [1000 repetitions x 5 folds] matrix of MAR differences for each neuron. To determine statistical significance, we performed a non-parametric sign permutation test on this matrix. The null hypothesis of this test is that there is no consistent effect of shuffling a neuron's activity, meaning the sign of the MAR difference is random across folds and repetitions. The observed statistic was compared against a null distribution generated by randomly flipping the signs of the differences in the matrix (10,000 permutations). Finally, the resulting *p*-values for all neurons were

corrected for multiple comparisons using the Benjamini-Hochberg FDR procedure. A neuron was classified as a "forecasting neuron" if its FDR-corrected *p*-value was less than 0.05.

Cross-correlation analyses

To quantify temporal offset between brain regions in their predictive activity, we performed cross-correlation analyses. For each pair of brain regions, cross-correlograms were computed on each trial's time-to-movement predictions (including data up to 250 ms before action onset). The empirical delay for a given pair of regions (Figures S3B and S3C) was defined as the time lag that yielded the maximum positive correlation within this trial-averaged cross-correlogram.

To assess whether the observed temporal delay between regions was statistically significantly different from zero, we employed a non-parametric bootstrap procedure. For each pair of regions, we generated 1,000 bootstrap samples. Each sample was constructed by randomly drawing N trials with replacement from the original pool of N trials. For every bootstrap iteration, we averaged the cross-correlograms of the resampled trials and extracted the time lag corresponding to the new peak correlation. This procedure generated an empirical distribution of 1,000 expected lags, from which we defined the 95% confidence intervals of the true lag. Statistical significance was determined using a two-tailed test on the bootstrap lag distribution. Finally, to account for multiple testing across all regional pairs, raw p-values were adjusted using the Benjamini-Hochberg FDR procedure, with an alpha threshold set at 0.05.

Drift-diffusion modelling

DDMs were fit to time-to-movement forecast traces derived from the entire neural population (Figure 5) and from individual regions (Figure 6). Separate models were fit for each recording session.

Model formulations and likelihood functions

Our modeling approach is aimed at describing the evolution of a decision variable x over time. It starts from an Ornstein-Uhlenbeck process[89,128] and incrementally adds parameters to build towards the models reflected in Equations 1 and 2-5. All continuous-time stochastic differential equations (SDEs) were discretized using the Euler-Maruyama method, a standard numerical procedure for approximating the solution of SDEs[129], with a time step of $\Delta t = 50\ \mathrm{ms}$ matching the bin size of firing rate time series.

Let x_n denote the value of the decision variable at time step n, corresponding to time $t_n = n\Delta t$. The process begins at immobility onset ($n = 0, t = 0$) with a latent initial state x_0. The first observed data point in our forecast traces corresponds to $t = 250$ ms after immobility onset, which is time step $n = 5$. Thus, a single trial trajectory is observed as the sequence $\boldsymbol{x_{obs}} = [x_5, x_6, \dots, x_N]$. We explicitly modeled the unobserved values of x between immobility onset (x_0) and the first observation (x_5). The process was initialized at $n = 0$ as $x_0 \sim \mathcal{N}\left(\mu_{intercept}, \sigma^2_{intercept}\right)$. The mean and variance were recursively propagated forward from the initial state x_0 to the first observed time step x_5 ($n = 5$, $t = 250$ ms) according to each model's dynamics. This yielded μ_{x_5} and $\sigma^2_{x_5}$, the conditional moments for the first observed point ($\mu_{x_5}, \sigma^2_{x_5}$) against which the model likelihood for x_5 was computed. For the remainder of the trajectories ($t > 250\ \mathrm{ms}$), the likelihood of Models 1 to 3 (see below) was evaluated

step-by-step by computing the probability of each new observation given the model dynamics propagated from the previous empirical data point.

To understand the derivation of the likelihood functions below, consider a general SDE of the form $dx_t = A(x_t, t)dt + cdW_t$. The Euler-Maruyama discretization yields the recursive update $x_{i+1} = x_i + A(x_i, t_i)\Delta t + c\eta_i$, where the noise increment η_i is drawn from a normal distribution $\mathcal{N}(0, \Delta t)$. Consequently, the probability of transitioning to the next state is normally distributed: $P(x_{i+1}|x_i) = \mathcal{N}(x_{i+1}; x_i + A(x_i, t_i)\Delta t, c^2\Delta t)$. Because standard white noise has independent increments, the process satisfies the Markov property: the state at t_{i+1} will depend only on the state at t_i. The joint likelihood of a full trajectory is therefore the product of the initial state probability and all subsequent step-by-step transition probabilities, which translates to a sum of log-normal densities when computing the log-likelihood.

This step-by-step multiplication is only valid if the noise increments are independent. In Model 4, where noise is colored (fractional Gaussian noise), the increments are autocorrelated, violating the Markov property. In that case, the probability of transitioning to x_{i+1} depends not only on x_i but on the entire history of the noise, necessitating a multivariate normal evaluation of the entire trajectory rather than a product of local transitions.

<u>Model 1: Ornstein-Uhlenbeck process (OUP; 4 parameters)</u>

This baseline model represents a leaky accumulator with no drift term. Its parameters are parameters are $\theta_{Model1} = [c, k, \mu_{intercept}, \sigma_{intercept}]$. The dynamics are governed by the SDE:

$$dx_t = -kx_t dt + cdW_t$$

where dW_t represents the increments of a standard Wiener process (white noise).

Its discretized form is:

$$x_{n+1} = x_n - kx_n\Delta t + c\Delta W_n$$

The log-likelihood for a single trajectory x is given by:

$$\log L(\boldsymbol{x_{obs}}|\theta_{Model1}) = \log \mathcal{N}\left(x_5; \mu_{x_5}, \sigma_{x_5}^2\right) + \sum_{n=5}^{N-1} \log \mathcal{N}(x_{n+1}; x_n(1 - k\Delta t), c^2\Delta t)$$

where μ_{x_5} and $\sigma_{x_5}^2$ are obtained by propagating the initial state parameters $(\mu_{intercept}, \sigma_{intercept})$ through the leak and noise accumulation for the first five time steps ($t = 0 \rightarrow 250$ ms).

<u>Model 2: Constant drift model (5 parameters)</u>

This model adds a constant drift rate, I, to the OUP, with parameters $\theta_{Model2} = \left[c, k, \mu_{intercept}, \sigma_{intercept}, I\right]$. Its SDE is:

$$dx_t = (I - kx_t)dt + cdW_t$$

Its discretized form is:

$$x_{n+1} = x_n + (I - kx_n)\Delta t + c\Delta W_n$$

The log-likelihood for a single trajectory is given by:

$$\log L(\boldsymbol{x_{obs}}|\theta_{Model2}) = \log \mathcal{N}\left(x_5; \mu_{x_5}, \sigma_{x_5}^2\right) + \sum_{n=5}^{N-1} \log \mathcal{N}\left(x_{n+1}; x_n + (I - kx_n)\Delta t, c^2 \Delta t\right)$$

where μ_{x_5} and $\sigma_{x_5}^2$ are obtained by propagating the initial state parameters $(\mu_{intercept}, \sigma_{intercept})$ through the model for the first five time steps. The contribution of a constant drift was tested by comparing the likelihoods of Models 1 and 2.

<u>Model 3: Full model with white noise (7 parameters)</u>

This is our primary model described in the main text (Equations 2-5), which incorporates trial-to-trial variability in the initial drift rate and a time-dependent drift rate $I(t)$. Model parameters are $\theta_{Model3} = \left[c, k, \alpha, \mu_{intercept}, \sigma_{intercept}, \mu_{drift}, \sigma_{drift}\right]$.

Its SDE is:

$$dx_t = (I(t) - kx_t)dt + cdW_t$$

To account for trial-to-trial variability, the initial drift magnitude I_0 is treated as a latent random variable following a lognormal distribution $\log I_0 \sim \mathcal{N}\left(\mu_{drift}, \sigma_{drift}^2\right)$. This distribution was chosen to constrain the drift rate to positive values. The drift can evolve monotonically over time according to either an exponential function $I(t) = I_0 e^{\alpha t}$ or a linear function $I(t) = I_0 + \alpha t$.

The discretized form is:

$$x_{n+1} = x_n + (I_n - kx_n)\Delta t + c\Delta W_n$$

with $I_n = I_0 e^{\alpha(n\Delta t)}$ (exponential) or $I_n = I_0 + \alpha(n\Delta t)$ (linear).

The total log-likelihood across all trials $\mathbf{X}_{\boldsymbol{all}}$ is computed by marginalizing over the initial drift magnitude I_0:

$$\log L_{total}\left(\mathbf{X}_{\boldsymbol{all}}|\theta_{Model3}\right) = \sum_{m=1}^{M} \log \int_0^{\infty} P\left(\boldsymbol{x}_{\mathbf{obs}}^{(m)} \mid I_0, \theta_{7\,params}\right) \cdot p\left(I_0 \mid \mu_{drift}, \sigma_{drift}\right) dI_0$$

where M denotes the total number of trials and $\boldsymbol{x}_{\mathbf{obs}}^{(m)}$ represents the observed trajectory of the m-th trial. $P(\boldsymbol{x} \mid \ldots)$ denotes the conditional likelihood of a trajectory, while $p(I_0 \mid \ldots)$ represents the probability density function of the initial drift magnitude.

The conditional log-likelihood of a single trajectory given a specific I_0 is:

$$\log P(\boldsymbol{x_{obs}}|I_0, \theta_{Model3}) = \log \mathcal{N}\left(x_5; \mu_{x_5|I_0}, \sigma_{x_5}^2\right) + \sum_{n=5}^{N-1} \log \mathcal{N}\left(x_{n+1}; x_n + (I_n - kx_n)\Delta t, c^2 \Delta t\right)$$

Here, $\mu_{x_5|I_0}$ and $\sigma_{x_5}^2$ are derived by propagating the initial state parameters $(\mu_{intercept}, \sigma_{intercept})$ through the model for the first five time steps. The contribution of the initial state variability $(\sigma_{intercept})$, leak rate (k) and urgency (α) parameters was tested by comparing this full model to nested versions where either $\sigma_{intercept}$, k, or α was fixed to zero.

Model 4: Full model with colored noise (8 parameters)

To test for temporal autocorrelations in the noise, we extended the full model by replacing the standard white noise process with fractional Gaussian noise, characterized by a spectral exponent parameter β. This model's parameters are $\theta_{Model4} = [c, k, \alpha, \mu_{intercept}, \sigma_{intercept}, \mu_{drift}, \sigma_{drift}, \beta]$. The corresponding SDE is:

$$dx_t = (I_t - kx_t)dt + cdW_t^H$$

where dW_t^H represents the increments of a fractional Brownian motion. The parameter β is defined as the spectral exponent of the noise power spectrum ($S(f) \propto 1/f^{\beta}$ where $S(f)$ denotes the power spectral density at frequency f), which relates to the Hurst exponent by $H = (\beta + 1)/2$.

Unlike the white noise models where transitions are independent, the presence of colored noise introduces correlations between time points within a trajectory. Consequently, the conditional likelihood of a full observed trajectory $\boldsymbol{x_{obs}}$ given a drift I_0 is evaluated using a single multivariate normal distribution:

$$\log P\left(\boldsymbol{x_{obs}} | I_0, \theta_{Model4}\right) = \log \mathcal{N}\left(\boldsymbol{x_{obs}}; \mu_{\boldsymbol{x_{obs}}|I_0}, \Sigma_{colored}\right)$$

Any observed trajectory remains defined from the 5th step onward (i.e., $\boldsymbol{x_{obs}} = [x_5, x_6, \dots, x_N]$). The mean vector $\mu_{\boldsymbol{x_{obs}}|I_0}$ is constructed by propagating the initial state mean $\mu_{intercept}$ and the trial-specific drift I_0 through the deterministic dynamics of the model from immobility onset ($n = 0$) up to the observed steps.

To evaluate the likelihood, we first derived a relation for the covariance matrix $\Sigma_{colored}$, which captures the dependency structure resulting from the discrete leaky accumulation of colored noise. We analytically derived its elements from the recursive equation for the model. First, since the noise is additive (i.e., the diffusion amplitude c is independent of the state x_t), the continuous-time SDE can be integrated pathwise and straightforwardly discretized using an Euler-Maruyama scheme[130,131]. With a time step Δt, this yields:

$$x_{n+1} = x_n + (I_n - kx_n)\Delta t + c\Delta W_n^H$$

Defining the discrete-time leak decay factor as $\phi = 1 - k\Delta t$ and the discrete noise increment as $\eta_n = c\Delta W_n^H$, the recursive update equation becomes:

$$x_{n+1} = \phi x_n + I_n \Delta t + \eta_n$$

Unrolling this recursive equation from immobility onset ($n = 0$) to an arbitrary time step n yields x_n as the sum of the leak-attenuated initial state, the drift, and the accumulated noise:

$$x_n = \phi^n x_0 + \sum_{u=0}^{n-1} \phi^{n-1-u} I_u \Delta t + \sum_{u=0}^{n-1} \phi^{n-1-u} \eta_u$$

With this decomposition of x_n, we can now derive the exact dependency structure between any two time points. Specifically, the elements of the covariance matrix $\Sigma_{colored}(n_a, n_b)$ represent the covariance between the states at any two time steps n_a and n_b (where $n_a, n_b \geq 5$, corresponding to the observed time steps) conditioned on a specific trial drift I_0:

$$\Sigma_{colored}(n_a, n_b) = \mathrm{Cov}\left(x_{n_a}, x_{n_b} \middle| I_0\right)$$

Because the drift component is deterministic for a trial-specific initial value I_0, it strictly shifts the mean of the trajectory ($\mu_{x_{obs}|I_0}$) without contributing to the covariance of the fluctuations. Consequently, the covariance structure depends only on the attenuated initial state variability and the history of the colored noise increments. Furthermore, the initial state $x_0 \sim \mathcal{N}\left(\mu_{intercept}, \sigma^2_{intercept}\right)$ is statistically independent of the subsequent within-trial noise increments η. Applying the covariance operator $\mathrm{Cov}\left(x_{n_a}, x_{n_b} \middle| I_0\right)$ thus simplifies to the sum of two independent covariance terms:

$$\Sigma_{colored}(n_a, n_b) = \mathrm{Cov}(\phi^{n_a} x_0, \phi^{n_b} x_0) + \mathrm{Cov}\left(\sum_{u=0}^{n_a-1} \phi^{n_a-1-u} \eta_u, \sum_{v=0}^{n_b-1} \phi^{n_b-1-v} \eta_v\right)$$

For the first term, factoring out the scalar multipliers, $\mathrm{Cov}(\phi^{n_a} x_0, \phi^{n_b} x_0)$ becomes $\phi^{n_a+n_b}\mathrm{Var}(x_0)$, which equals $\phi^{n_a+n_b}\sigma^2_{intercept}$. For the second term, the covariance of the sums expands into a double summation of the covariance between individual noise increments:

$$\mathrm{Cov}\left(\sum_{u=0}^{n_a-1} \phi^{n_a-1-u} \eta_u, \sum_{v=0}^{n_b-1} \phi^{n_b-1-v} \eta_v\right) = \sum_{u=0}^{n_a-1} \sum_{v=0}^{n_b-1} \phi^{n_a+n_b-2-u-v} \mathrm{Cov}(\eta_u, \eta_v)$$

In discrete time, the covariance between two fractional Gaussian noise increments separated by a lag $\tau = |u - v|$ is proportional to the autocorrelation function ρ, such that $\mathrm{Cov}(\eta_u, \eta_v) = c^2 \Delta t \rho(|u - v|)$[132].

Substituting these relationships yields:

$$\Sigma_{colored}(n_a, n_b) = \phi^{n_a+n_b}\sigma^2_{intercept} + c^2 \Delta t \sum_{u=0}^{n_a-1} \sum_{v=0}^{n_b-1} \phi^{n_a+n_b-2-u-v} \rho(|u - v|)$$

The first term represents the propagated variance from the initial state, while the double summation represents the convolution of the noise autocorrelation with the system's impulse response (that is, the discrete-time leak decay). Analytically, the autocorrelation function of the fractional Gaussian noise is given by[132]:

$$\rho(\tau) = \frac{1}{2}(|\tau + 1|^{2H} - 2|\tau|^{2H} + |\tau - 1|^{2H})$$

where $\tau = |u - v|$ is the lag between noise increments.

Finally, the total log-likelihood across all trials $\boldsymbol{X_{all}}$ is computed by integrating this multivariate conditional likelihood over the lognormal prior distribution of I_0:

$$\log L_{total}(\boldsymbol{X_{all}}|\theta_{Model4}) = \sum_{m=1}^{M} \log \int_0^{\infty} P\left(\boldsymbol{x_{obs}^{(m)}} \middle| I_0, \theta_{Model4}\right) \cdot p\left(I_0 \middle| \mu_{drift}, \sigma_{drift}\right) dI_0$$

The parameter β was constrained to be between -1 and +1 to ensure the stationarity of the noise process[132]. Its contribution was tested by comparing the likelihoods of Models 3 and 4.

Data fit by the DDM

DDM parameters were estimated by fitting forecast traces derived from held-out test trials, spanning from their beginning up to 250 ms before action onset. The forecast traces were fit from their initial value (250 ms after immobility onset) based on our finding that this initial value is correlated with trial duration (Figures 4A and S4), and that predictive information typically emerges within the first 500 ms for most trials (Figure 2H), consistent with a decision-making process starting soon after immobility onset. The end point (250 ms before action onset) was chosen to isolate the decision-making process and avoid contamination from potential post-decision, motor execution-related activity. This choice should be conservative: in humans, the point of no return beyond which a self-initiated action can no longer be aborted occurs approximately 200 ms before action onset[15,78]. This delay is likely shorter in mice, given the much lower limb inertia.

As a pre-processing step, because our DDM formulation assumes a positive drift rate and leak rate, the forecast traces were first made non-negative by subtracting from all forecast traces the minimum value observed across all traces within that session.

Optimization and parameter fitting

Model parameters were fit by minimizing the negative log-likelihood of the data. We used a two-stage global optimization procedure to avoid local minima. First, a genetic algorithm[133,134] was run five times with different random seeds to broadly explore the parameter space (function *ga* in MATLAB). The best-fit parameters from these runs were then used as the starting point for an interior-point local optimizer[135,136] to refine the minimal negative likelihood estimate (function *fmincon* in MATLAB). Parameters such as leak rate, noise amplitude, and distribution widths ($k, c, \sigma_{intercept}, \sigma_{drift}$) were constrained to be positive, while the drift's temporal rate α was allowed to be positive or negative.

Statistical procedure for model selection

The improvement in model fit for each parameter was assessed using a non-parametric bootstrap procedure[137]. We performed 1,000 bootstrap iterations. For each iteration, we constructed a training set by sampling N trials with replacement from the trials for the given session, where N is the total number of trials for that session. The remaining trials that were not selected in the training set served as the held-out test set. We fit full and nested models on the training set and evaluated their goodness-of-fit on the test set by computing the total log-likelihood for all held-out trials. To ensure comparability between bootstrap iterations and sessions, the total log-likelihood was normalized by the number of data points in the given test set. For each iteration, we computed the log-likelihood gain (ΔLL) as the difference between the full and nested models; these values are reported in Figures 5, 6, and S6. A positive ΔLL thus indicates that the additional parameter improved the model fit. Statistical significance was determined via an empirical *p*-value, defined as the fraction of bootstrap iterations where the full model failed to outperform the nested model ($\Delta LL \leq 0$). A parameter was considered to significantly improve the model if $p < 0.05$. This model selection approach was validated using parameter recovery on simulated data (Figure S5).

Derivation of drift rate moments

In our DDM variant, the initial drift rate I_0 is drawn from a log-normal distribution for each trial. The parameters directly estimated by our fitting procedure, μ_{drift} and σ_{drift}, correspond to the mean and standard deviation of the underlying normal distribution for $log(I_0)$ (Figures 5 and 6). For visualization purposes (Figures S6A, S6B and S6G), we also calculated the arithmetic mean ($drift\ mean$) and standard deviation ($drift\ std$) of the log-normal distribution itself using the following equations:

$$drift\ mean = e^{\mu_{drift}+\frac{\sigma_{drift}^2}{2}}$$

$$drift\ std = \sqrt{\left(e^{\sigma_{drift}^2} - 1\right) e^{2\mu_{drift}+\sigma_{drift}^2}}$$

Testing the necessity of within-trial noise via ablation analysis

To assess whether the noise accumulation is necessary for reaching the activity threshold, we performed an ablation analysis on the diffusion term. Using the best-fit parameters obtained with the 7-parameter model (Equations 2-5), we simulated decision variable trajectories under two conditions: with the diffusion term (c set to its best-fit value) and without it (c = 0). To ensure the simulations were realistically constrained by observations, the threshold beyond which each simulated trajectory ended was drawn from a normal distribution fit to the actual forecast traces at 250 ms prior to climbing onset. This time point corresponds to the end of the DDM fitting window and was chosen to define a realistic activity threshold. To guarantee an unbiased comparison between the duration of simulations and actual trials, we discarded simulations shorter than the minimum permitted trial duration fit by the DDMs. The duration of the forecast traces fit by the DDMs equals the total trial (immobility) duration minus 0.45 s. This deduction accounts for a 0.2 s initial offset because the forecast model uses five time steps, and a 0.25 s truncation prior to climbing onset. Since the minimum trial duration is 1 s, the minimum fitted trial duration is 0.55 s. Accordingly, the duration of simulated sequences was calculated starting at 250 ms after their onset, and only simulated sequences lasting at least 0.55 s were retained for this comparison.

Assessment of distinct variability sources via dynamical constraints

To test that trial-to-trial variability in the decision process's initial state ($\sigma_{intercept}$) and the within-trial continuous noise arise from distinct sources of variability, we assessed whether the initial variance ($\sigma_{intercept}$) could be explained by the noise accumulation during the period preceding immobility onset. For a process governed by leak rate k and noise amplitude c (i.e., an Ornstein-Uhlenbeck process), the variance derivative is $\frac{dV}{dt} = -2kV + c^2$, which yields the time-dependent solution $V_t = \frac{c^2}{2k}\left(1 - e^{-2kt}\right)$[89,90]. This function implies a theoretical upper bound, $\frac{c^2}{2k}$, which represents the maximum variability the noise source can sustain against the leak, even if allowed to accumulate over an infinite duration. We compared the fitted initial variance ($\sigma_{intercept}^2$) obtained from our 7-parameter model to this theoretical maximum calculated using the best-fit c and k parameters for each session (Figure S6C). Our finding of $\sigma_{intercept}^2 \gg V_\infty$ indicates that the initial variability exceeds the variability that could be attributed to noise accumulation.

References

1. Broom, M., and Ruxton, G.D. (2005). You can run—or you can hide: optimal strategies for cryptic prey against pursuit predators. Behavioral Ecology *16*, 534-540. 10.1093/beheco/ari024.
2. Domenici, P., Blagburn, J.M., and Bacon, J.P. (2011). Animal escapology I: theoretical issues and emerging trends in escape trajectories. J Exp Biol *214*, 2463-2473. 10.1242/jeb.029652.
3. Kornhuber, H.H., and Deecke, L. (1965). Hirnpotentialänderungen bei Willkürbewegungen und passiven Bewegungen des Menschen: Bereitschaftspotential und reafferente Potentiale. Pflüger's Archiv für die gesamte Physiologie des Menschen und der Tiere *284*, 1-17. 10.1007/BF00412364.
4. Libet, B., Gleason, C.A., Wright, E.W., and Pearl, D.K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential). The unconscious initiation of a freely voluntary act. Brain *106 (Pt 3)*, 623-642. 10.1093/brain/106.3.623.
5. Deecke, L., Scheid, P., and Kornhuber, H.H. (1969). Distribution of readiness potential, pre-motion positivity, and motor potential of the human cerebral cortex preceding voluntary finger movements. Exp Brain Res *7*, 158-168. 10.1007/BF00235441.
6. Romo, R., and Schultz, W. (1987). Neuronal activity preceding self-initiated or externally timed arm movements in area 6 of monkey cortex. Exp Brain Res *67*, 656-662. 10.1007/BF00247297.
7. Maimon, G., and Assad, J.A. (2006). A cognitive signal for the proactive timing of action in macaque LIP. Nat Neurosci *9*, 948-955. 10.1038/nn1716.
8. Isomura, Y., Harukuni, R., Takekawa, T., Aizawa, H., and Fukai, T. (2009). Microcircuitry coordination of cortical motor information in self-initiation of voluntary movements. Nat Neurosci *12*, 1586-1593. 10.1038/nn.2431.
9. Seki, T., Gemba, H., Matsuzaki, R., and Nakao, K. (2005). Readiness potential and movement initiation in the rat. Jpn J Physiol *55*, 1-9. 10.2170/jjphysiol.R2073.
10. Melanson, A., Mejias, J.F., Jun, J.J., Maler, L., and Longtin, A. (2017). Nonstationary Stochastic Dynamics Underlie Spontaneous Transitions between Active and Inactive Behavioral States. eNeuro *4*. 10.1523/ENEURO.0355-16.2017.
11. Kagaya, K., and Takahata, M. (2011). Sequential synaptic excitation and inhibition shape readiness discharge for voluntary behavior. Science *332*, 365-368. 10.1126/science.1202244.
12. Kagaya, K., and Takahata, M. (2010). Readiness discharge for spontaneous initiation of walking in crayfish. J Neurosci *30*, 1348-1362. 10.1523/JNEUROSCI.4885-09.2010.
13. Di Russo, F., Berchicci, M., Bozzacchi, C., Perri, R.L., Pitzalis, S., and Spinelli, D. (2017). Beyond the "Bereitschaftspotential": Action preparation behind cognitive functions. Neurosci Biobehav Rev *78*, 57-81. 10.1016/j.neubiorev.2017.04.019.
14. Schurger, A., Hu, P., Pak, J., and Roskies, A.L. (2021). What Is the Readiness Potential? Trends Cogn Sci *25*, 558-570. 10.1016/j.tics.2021.04.001.
15. Libet, B. (1985). Unconscious cerebral initiative and the role of conscious will in voluntary action. Behavioral and Brain Sciences *8*, 529-539. 10.1017/s0140525x00044903.
16. Soon, C.S., Brass, M., Heinze, H.J., and Haynes, J.D. (2008). Unconscious determinants of free decisions in the human brain. Nat Neurosci *11*, 543-545. 10.1038/nn.2112.
17. Dominik, T., Mele, A., Schurger, A., and Maoz, U. (2024). Libet's legacy: A primer to the neuroscience of volition. Neurosci Biobehav Rev *157*, 105503. 10.1016/j.neubiorev.2023.105503.

18. Braun, M.N., Wessler, J., and Friese, M. (2021). A meta-analysis of Libet-style experiments. Neurosci Biobehav Rev *128*, 182-198. 10.1016/j.neubiorev.2021.06.018.
19. Shibasaki, H., and Hallett, M. (2006). What is the Bereitschaftspotential? Clin Neurophysiol *117*, 2341-2356. 10.1016/j.clinph.2006.04.025.
20. Mitelut, C., Zhang, Y., Sekino, Y., Boyd, J.D., Bollanos, F., Swindale, N.V., Silasi, G., Saxena, S., and Murphy, T.H. (2022). Mesoscale cortex-wide neural dynamics predict self-initiated actions in mice several seconds prior to movement. Elife *11*. 10.7554/eLife.76506.
21. Fried, I., Mukamel, R., and Kreiman, G. (2011). Internally generated preactivation of single neurons in human medial frontal cortex predicts volition. Neuron *69*, 548-562. 10.1016/j.neuron.2010.11.045.
22. Maimon, G., and Assad, J.A. (2006). Parietal area 5 and the initiation of self-timed movements versus simple reactions. J Neurosci *26*, 2487-2498. 10.1523/JNEUROSCI.3590-05.2006.
23. Schurger, A., Sitt, J.D., and Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. Proc Natl Acad Sci U S A *109*, E2904-2913. 10.1073/pnas.1210467109.
24. Gavenas, J., Rutishauser, U., Schurger, A., and Maoz, U. (2024). Slow ramping emerges from spontaneous fluctuations in spiking neural networks. Nat Commun *15*, 7285. 10.1038/s41467-024-51401-x.
25. Murakami, M., Vicente, M.I., Costa, G.M., and Mainen, Z.F. (2014). Neural antecedents of self-initiated actions in secondary motor cortex. Nat Neurosci *17*, 1574-1582. 10.1038/nn.3826.
26. Moutard, C., Dehaene, S., and Malach, R. (2015). Spontaneous Fluctuations and Non-linear Ignitions: Two Dynamic Faces of Cortical Recurrent Loops. Neuron *88*, 194-206. 10.1016/j.neuron.2015.09.018.
27. Schmidt, S., Jo, H.G., Wittmann, M., and Hinterberger, T. (2016). 'Catching the waves' - slow cortical potentials as moderator of voluntary action. Neurosci Biobehav Rev *68*, 639-650. 10.1016/j.neubiorev.2016.06.023.
28. Recanatesi, S., Pereira-Obilinovic, U., Murakami, M., Mainen, Z., and Mazzucato, L. (2022). Metastable attractors explain the variable timing of stable behavioral action sequences. Neuron *110*, 139-153 e139. 10.1016/j.neuron.2021.10.011.
29. Travers, E., Khalighinejad, N., Schurger, A., and Haggard, P. (2020). Do readiness potentials happen all the time? Neuroimage *206*, 116286. 10.1016/j.neuroimage.2019.116286.
30. Ratcliff, R., Smith, P.L., Brown, S.D., and McKoon, G. (2016). Diffusion Decision Model: Current Issues and History. Trends Cogn Sci *20*, 260-281. 10.1016/j.tics.2016.01.007.
31. Smith, P.L. (2000). Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer. J Math Psychol *44*, 408-463. 10.1006/jmps.1999.1260.
32. Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu Rev Neurosci *30*, 535-574. 10.1146/annurev.neuro.29.051605.113038.
33. Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. Psychol Rev *108*, 550-592. 10.1037/0033-295x.108.3.550.
34. Ratcliff, R. (1978). A theory of memory retrieval. Psychological Review *85*, 59-108. 10.1037/0033-295x.85.2.59.
35. Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput *20*, 873-922. 10.1162/neco.2008.12-06-420.
36. Ratcliff, R. (1985). Theoretical interpretations of the speed and accuracy of positive and negative responses. Psychological Review *92*, 212-225. 10.1037/0033-295x.92.2.212.
37. Bogler, C., Grujičić, B., and Haynes, J.-D. (2023). Clarifying the nature of stochastic fluctuations and accumulation processes in spontaneous movements. Frontiers in Psychology *14*. 10.3389/fpsyg.2023.1271180.

38. He, B.J., Zempel, J.M., Snyder, A.Z., and Raichle, M.E. (2010). The temporal structures and functional significance of scale-free brain activity. Neuron *66*, 353-369. 10.1016/j.neuron.2010.04.020.
39. Faisal, A.A., Selen, L.P., and Wolpert, D.M. (2008). Noise in the nervous system. Nat Rev Neurosci *9*, 292-303. 10.1038/nrn2258.
40. Deecke, L. (1987). Bereitschaftspotential as an indicator of movement preparation in supplementary motor area and motor cortex. Ciba Found Symp *132*, 231-250. 10.1002/9780470513545.ch14.
41. Paradiso, G., Cunic, D., Saint-Cyr, J.A., Hoque, T., Lozano, A.M., Lang, A.E., and Chen, R. (2004). Involvement of human thalamus in the preparation of self-paced movement. Brain *127*, 2717-2731. 10.1093/brain/awh288.
42. Gaffield, M.A., Sauerbrei, B.A., and Christie, J.M. (2022). Cerebellum encodes and influences the initiation, performance, and termination of discontinuous movements in mice. Elife *11*. 10.7554/eLife.71464.
43. Blouin, J.S., Bard, C., and Paillard, J. (2004). Contribution of the cerebellum to self-initiated synchronized movements: a PET study. Exp Brain Res *155*, 63-68. 10.1007/s00221-003-1709-9.
44. Schultz, W., and Romo, R. (1992). Role of primate basal ganglia and frontal cortex in the internal generation of movements. I. Preparatory activity in the anterior striatum. Exp Brain Res *91*, 363-384. 10.1007/BF00227834.
45. Suzuki, S., Lawlor, V.M., Cooper, J.A., Arulpragasam, A.R., and Treadway, M.T. (2021). Distinct regions of the striatum underlying effort, movement initiation and effort discounting. Nat Hum Behav *5*, 378-388. 10.1038/s41562-020-00972-y.
46. Murakami, M., Shteingart, H., Loewenstein, Y., and Mainen, Z.F. (2017). Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. Neuron *94*, 908-919 e907. 10.1016/j.neuron.2017.04.040.
47. Khalighinejad, N., Priestley, L., Jbabdi, S., and Rushworth, M.F.S. (2020). Human decisions about when to act originate within a basal forebrain-nigral circuit. Proc Natl Acad Sci U S A *117*, 11799-11810. 10.1073/pnas.1921211117.
48. Khalighinejad, N., Garrett, N., Priestley, L., Lockwood, P., and Rushworth, M.F.S. (2021). A habenula-insular circuit encodes the willingness to act. Nat Commun *12*, 6329. 10.1038/s41467-021-26569-1.
49. Wiese, H., Stude, P., Nebel, K., Osenberg, D., Volzke, V., Ischebeck, W., Stolke, D., Diener, H.C., and Keidel, M. (2004). Impaired movement-related potentials in acute frontal traumatic brain injury. Clin Neurophysiol *115*, 289-298. 10.1016/s1388-2457(03)00348-1.
50. Wessel, K., Zeffiro, T., Lou, J.S., Toro, C., and Hallett, M. (1995). Regional cerebral blood flow during a self-paced sequential finger opposition task in patients with cerebellar degeneration. Brain *118 (Pt 2)*, 379-393. 10.1093/brain/118.2.379.
51. Sasaki, K., Gemba, H., Hashimoto, S., and Mizuno, N. (1979). Influences of cerebellar hemispherectomy on slow potentials in the motor cortex preceding self-paced hand movements in the monkey. Neurosci Lett *15*, 23-28. 10.1016/0304-3940(79)91523-4.
52. Tsang, E.W., Hamani, C., Moro, E., Mazzella, F., Poon, Y.Y., Lozano, A.M., and Chen, R. (2010). Involvement of the human pedunculopontine nucleus region in voluntary movements. Neurology *75*, 950-959. 10.1212/WNL.0b013e3181f25b35.
53. Thaler, D., Chen, Y.C., Nixon, P.D., Stern, C.E., and Passingham, R.E. (1995). The functions of the medial premotor cortex. I. Simple learned movements. Exp Brain Res *102*, 445-460. 10.1007/BF00230649.

54. Jahanshahi, M., Jenkins, I.H., Brown, R.G., Marsden, C.D., Passingham, R.E., and Brooks, D.J. (1995). Self-initiated versus externally triggered movements. I. An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson's disease subjects. Brain *118 (Pt 4)*, 913-933. 10.1093/brain/118.4.913.
55. Nixon, P.D., and Passingham, R.E. (1998). The striatum and self-paced movements. Behav Neurosci *112*, 719-724. 10.1037//0735-7044.112.3.719.
56. Cohen, R.A., Kaplan, R.F., Zuffante, P., Moser, D.J., Jenkins, M.A., Salloway, S., and Wilkinson, H. (1999). Alteration of intention and self-initiated action associated with bilateral anterior cingulotomy. J Neuropsychiatry Clin Neurosci *11*, 444-453. 10.1176/jnp.11.4.444.
57. Dick, J.P., Rothwell, J.C., Day, B.L., Cantello, R., Buruma, O., Gioux, M., Benecke, R., Berardelli, A., Thompson, P.D., and Marsden, C.D. (1989). The Bereitschaftspotential is abnormal in Parkinson's disease. Brain *112 (Pt 1)*, 233-244. 10.1093/brain/112.1.233.
58. Laplane, D., Talairach, J., Meininger, V., Bancaud, J., and Orgogozo, J.M. (1977). Clinical consequences of corticectomies involving the supplementary motor area in man. J Neurol Sci *34*, 301-314. 10.1016/0022-510x(77)90148-4.
59. Khalighinejad, N., Bongioanni, A., Verhagen, L., Folloni, D., Attali, D., Aubry, J.F., Sallet, J., and Rushworth, M.F.S. (2020). A Basal Forebrain-Cingulate Circuit in Macaques Decides It Is Time to Act. Neuron *105*, 370-384 e378. 10.1016/j.neuron.2019.10.030.
60. Khalighinejad, N., Manohar, S., Husain, M., and Rushworth, M.F.S. (2022). Complementary roles of serotonergic and cholinergic systems in decisions about when to act. Curr Biol *32*, 1150-1162 e1157. 10.1016/j.cub.2022.01.042.
61. Narayanan, N.S., and Laubach, M. (2006). Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. Neuron *52*, 921-931. 10.1016/j.neuron.2006.10.021.
62. Tsujimoto, T., Gemba, H., and Sasaki, K. (1993). Effect of cooling the dentate nucleus of the cerebellum on hand movement of the monkey. Brain Res *629*, 1-9. 10.1016/0006-8993(93)90473-z.
63. Tecuapetla, F., Jin, X., Lima, S.Q., and Costa, R.M. (2016). Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution. Cell *166*, 703-715. 10.1016/j.cell.2016.06.032.
64. Klaus, A., Alves da Silva, J., and Costa, R.M. (2019). What, If, and When to Move: Basal Ganglia Circuits and Self-Paced Action Initiation. Annu Rev Neurosci *42*, 459-483. 10.1146/annurev-neuro-072116-031033.
65. Passingham, R.E., Bengtsson, S.L., and Lau, H.C. (2010). Medial frontal cortex: from self-generated action to reflection on one's own performance. Trends Cogn Sci *14*, 16-21. 10.1016/j.tics.2009.11.001.
66. Hoffstaedter, F., Grefkes, C., Zilles, K., and Eickhoff, S.B. (2013). The "what" and "when" of self-initiated movements. Cereb Cortex *23*, 520-530. 10.1093/cercor/bhr391.
67. Koh, N., Ma, Z., Sarup, A., Kristl, A.C., Agrios, M., Young, M., and Miri, A. (2025). Selective direct influence of motor cortex on limb muscle activity during naturalistic climbing in mice. Nat Neurosci *28*, 2537-2549. 10.1038/s41593-025-02093-z.
68. Mathis, A., Mamidanna, P., Cury, K.M., Abe, T., Murthy, V.N., Mathis, M.W., and Bethge, M. (2018). DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. Nat Neurosci *21*, 1281-1289. 10.1038/s41593-018-0209-y.
69. Jun, J.J., Steinmetz, N.A., Siegle, J.H., Denman, D.J., Bauza, M., Barbarits, B., Lee, A.K., Anastassiou, C.A., Andrei, A., Aydin, C., et al. (2017). Fully integrated silicon probes for high-density recording of neural activity. Nature *551*, 232-236. 10.1038/nature24636.

70. Hyland, B. (1998). Neural activity related to reaching and grasping in rostral and caudal regions of rat motor cortex. Behav Brain Res *94*, 255-269. 10.1016/s0166-4328(97)00157-5.
71. Okoro, S.U., Goz, R.U., Njeri, B.W., Harish, M., Ruff, C.F., Ross, S.E., Gerfen, C., and Hooks, B.M. (2022). Organization of Cortical and Thalamic Input to Inhibitory Neurons in Mouse Motor Cortex. J Neurosci *42*, 8095-8112. 10.1523/JNEUROSCI.0950-22.2022.
72. Lee, I.H., and Assad, J.A. (2003). Putaminal activity for simple reactions or self-timed movements. J Neurophysiol *89*, 2528-2537. 10.1152/jn.01055.2002.
73. Catanese, J., and Jaeger, D. (2021). Premotor Ramping of Thalamic Neuronal Activity Is Modulated by Nigral Inputs and Contributes to Control the Timing of Action Release. J Neurosci *41*, 1878-1891. 10.1523/JNEUROSCI.1204-20.2020.
74. Stringer, C., Pachitariu, M., Steinmetz, N., Reddy, C.B., Carandini, M., and Harris, K.D. (2019). Spontaneous behaviors drive multidimensional, brainwide activity. Science *364*, 255. 10.1126/science.aav7893.
75. Musall, S., Kaufman, M.T., Juavinett, A.L., Gluf, S., and Churchland, A.K. (2019). Single-trial neural dynamics are dominated by richly varied movements. Nat Neurosci *22*, 1677-1686. 10.1038/s41593-019-0502-4.
76. Zou, H., and Hastie, T. (2005). Regularization and Variable Selection Via the Elastic Net. Journal of the Royal Statistical Society Series B: Statistical Methodology *67*, 301-320. 10.1111/j.1467-9868.2005.00503.x.
77. Stine, G.M., and Jazayeri, M. (2025). Control Principles of Neural Dynamics Revealed by the Neurobiology of Timing. Annu Rev Neurosci *48*, 43-63. 10.1146/annurev-neuro-091724-015512.
78. Schultze-Kraft, M., Birman, D., Rusconi, M., Allefeld, C., Gorgen, K., Dahne, S., Blankertz, B., and Haynes, J.D. (2016). The point of no return in vetoing self-initiated movements. Proc Natl Acad Sci U S A *113*, 1080-1085. 10.1073/pnas.1513569112.
79. Brown, S.D., and Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. Cogn Psychol *57*, 153-178. 10.1016/j.cogpsych.2007.12.002.
80. Thura, D., Beauregard-Racine, J., Fradet, C.W., and Cisek, P. (2012). Decision making by urgency gating: theory and experimental support. J Neurophysiol *108*, 2912-2930. 10.1152/jn.01071.2011.
81. Smith, P.L., and Ratcliff, R. (2022). Modeling evidence accumulation decision processes using integral equations: Urgency-gating and collapsing boundaries. Psychol Rev *129*, 235-267. 10.1037/rev0000301.
82. Cisek, P., Puskas, G.A., and El-Murr, S. (2009). Decisions in changing conditions: the urgency-gating model. J Neurosci *29*, 11560-11571. 10.1523/JNEUROSCI.1844-09.2009.
83. Ditterich, J. (2006). Stochastic models of decisions about motion direction: behavior and physiology. Neural Netw *19*, 981-1012. 10.1016/j.neunet.2006.05.042.
84. Ditterich, J. (2006). Evidence for time-variant decision making. Eur J Neurosci *24*, 3628-3641. 10.1111/j.1460-9568.2006.05221.x.
85. Hawkins, G.E., Wagenmakers, E.J., Ratcliff, R., and Brown, S.D. (2015). Discriminating evidence accumulation from urgency signals in speeded decision making. J Neurophysiol *114*, 40-47. 10.1152/jn.00088.2015.
86. Hawkins, G.E., Forstmann, B.U., Wagenmakers, E.J., Ratcliff, R., and Brown, S.D. (2015). Revisiting the evidence for collapsing boundaries and urgency signals in perceptual decision-making. J Neurosci *35*, 2476-2484. 10.1523/JNEUROSCI.2410-14.2015.
87. Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., and Rangel, A. (2023). The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high

and low time pressure. Judgment and Decision Making *5*, 437-449. 10.1017/s1930297500001285.

88. Thura, D., and Cisek, P. (2017). The Basal Ganglia Do Not Select Reach Targets but Control the Urgency of Commitment. Neuron *95*, 1160-1170 e1165. 10.1016/j.neuron.2017.07.039.

89. Uhlenbeck, G.E., and Ornstein, L.S. (1930). On the Theory of the Brownian Motion. Physical Review *36*, 823-841. 10.1103/PhysRev.36.823.

90. Gardiner, C.W. (2004). Handbook of stochastic methods : for physics, chemistry, and the natural sciences, 3rd Edition (Springer).

91. Schurger, A. (2018). Specific Relationship between the Shape of the Readiness Potential, Subjective Decision Time, and Waiting Time Predicted by an Accumulator Model with Temporally Autocorrelated Input Noise. eNeuro *5*. 10.1523/ENEURO.0302-17.2018.

92. Guevara Erra, R., Arbotto, M., and Schurger, A. (2019). An integration-to-bound model of decision-making that accounts for the spectral properties of neural data. Sci Rep *9*, 8365. 10.1038/s41598-019-44197-0.

93. Wang, Z.A., Kurgyis, B., Chen, S., Kang, B., Chen, F., Liu, Y., Liu, D., Svoboda, K., Li, N., and Druckmann, S. (2026). Brain-wide analysis reveals movement encoding structured across and within brain areas. Nat Neurosci *29*, 147-158. 10.1038/s41593-025-02114-x.

94. Chen, S., Liu, Y., Wang, Z.A., Colonell, J., Liu, L.D., Hou, H., Tien, N.W., Wang, T., Harris, T., Druckmann, S., et al. (2024). Brain-wide neural activity underlying memory-guided movement. Cell *187*, 676-691 e616. 10.1016/j.cell.2023.12.035.

95. Khilkevich, A., Lohse, M., Low, R., Orsolic, I., Bozic, T., Windmill, P., and Mrsic-Flogel, T.D. (2024). Brain-wide dynamics linking sensation to action during decision-making. Nature *634*, 890-900. 10.1038/s41586-024-07908-w.

96. Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266-273. 10.1038/s41586-019-1787-x.

97. International Brain, L., Angelaki, D., Benson, B., Benson, J., Birman, D., Bonacchi, N., Bougrova, K., Bruijns, S.A., Carandini, M., Catarino, J.A., et al. (2025). A brain-wide map of neural activity during complex behaviour. Nature *645*, 177-191. 10.1038/s41586-025-09235-0.

98. Findling, C., Hubert, F., International Brain, L., Acerbi, L., Benson, B., Benson, J., Birman, D., Bonacchi, N., Buchanan, E.K., Bruijns, S., et al. (2025). Brain-wide representations of prior information in mouse decision-making. Nature *645*, 192-200. 10.1038/s41586-025-09226-1.

99. Oryshchuk, A., Sourmpis, C., Weverbergh, J., Asri, R., Esmaeili, V., Modirshanechi, A., Gerstner, W., Petersen, C.C.H., and Crochet, S. (2024). Distributed and specific encoding of sensory, motor, and decision information in the mouse neocortex during goal-directed behavior. Cell Rep *43*, 113618. 10.1016/j.celrep.2023.113618.

100. Rosen, M.C., and Freedman, D.J. (2026). How distributed is the brain-wide network that is recruited for cognition? Nat Rev Neurosci *27*, 138-150. 10.1038/s41583-025-00992-5.

101. Chandrasekaran, C., Gupta, D., Javadzadeh, M., Wang, T., Vivar-Lazo, M., Engel, T., Cisek, P., and Fetsch, C.R. (2025). No Central Executive? Decision Formation Through Multi-Area Population Dynamics. J Neurosci *45*. 10.1523/JNEUROSCI.1633-25.2025.

102. Yang, Z., Inagaki, M., Gerfen, C.R., Fontolan, L., and Inagaki, H.K. (2026). Integrator dynamics in the cortico-basal ganglia loop for flexible motor timing. Nature *649*, 1244-1253. 10.1038/s41586-025-09778-2.

103. Emmons, E.B., De Corte, B.J., Kim, Y., Parker, K.L., Matell, M.S., and Narayanan, N.S. (2017). Rodent Medial Frontal Control of Temporal Processing in the Dorsomedial Striatum. J Neurosci *37*, 8718-8733. 10.1523/JNEUROSCI.1376-17.2017.

104. Kunimatsu, J., Suzuki, T.W., Ohmae, S., and Tanaka, M. (2018). Different contributions of preparatory activity in the basal ganglia and cerebellum for self-timing. Elife *7*. 10.7554/eLife.35676.
105. Wang, J., Narain, D., Hosseini, E.A., and Jazayeri, M. (2018). Flexible timing by temporal scaling of cortical responses. Nat Neurosci *21*, 102-110. 10.1038/s41593-017-0028-6.
106. Bakhurin, K.I., Goudar, V., Shobe, J.L., Claar, L.D., Buonomano, D.V., and Masmanidis, S.C. (2017). Differential Encoding of Time by Prefrontal and Striatal Network Dynamics. J Neurosci *37*, 854-870. 10.1523/JNEUROSCI.1789-16.2016.
107. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791-806. 10.1016/j.neuron.2013.10.047.
108. Allaway, K.C., Munoz, W., Tremblay, R., Sherer, M., Herron, J., Rudy, B., Machold, R., and Fishell, G. (2020). Cellular birthdate predicts laminar and regional cholinergic projection topography in the forebrain. Elife *9*. 10.7554/eLife.63249.
109. Minces, V., Pinto, L., Dan, Y., and Chiba, A.A. (2017). Cholinergic shaping of neural correlations. Proc Natl Acad Sci U S A *114*, 5725-5730. 10.1073/pnas.1621493114.
110. Luo, J.K., Tan, J., Myhre, E., Salvino, P., Hark, A., and Pinto, L. (2025). A cholinergic mechanism orchestrating task-dependent computation across the cortex. bioRxiv. 10.1101/2025.11.26.690825.
111. Munoz, W., and Rudy, B. (2014). Spatiotemporal specificity in cholinergic control of neocortical function. Curr Opin Neurobiol *26*, 149-160. 10.1016/j.conb.2014.02.015.
112. Zaborszky, L., Gombkoto, P., Varsanyi, P., Gielow, M.R., Poe, G., Role, L.W., Ananth, M., Rajebhosale, P., Talmage, D.A., Hasselmo, M.E., et al. (2018). Specific Basal Forebrain-Cortical Cholinergic Circuits Coordinate Cognitive Operations. J Neurosci *38*, 9446-9458. 10.1523/JNEUROSCI.1676-18.2018.
113. Yu, B., Yue, Y., Ren, C., Yun, R., Lim, B., and Komiyama, T. (2025). Cholinergic feedback for modality- and context-specific modulation of sensory representations. Science *388*, 1324-1329. 10.1126/science.ads9152.
114. Parikh, V., Kozak, R., Martinez, V., and Sarter, M. (2007). Prefrontal acetylcholine release controls cue detection on multiple timescales. Neuron *56*, 141-154. 10.1016/j.neuron.2007.08.025.
115. Sarter, M., Lustig, C., Howe, W.M., Gritton, H., and Berry, A.S. (2014). Deterministic functions of cortical acetylcholine. Eur J Neurosci *39*, 1912-1920. 10.1111/ejn.12515.
116. Pinto, L., Goard, M.J., Estandian, D., Xu, M., Kwan, A.C., Lee, S.H., Harrison, T.C., Feng, G., and Dan, Y. (2013). Fast modulation of visual perception by basal forebrain cholinergic neurons. Nat Neurosci *16*, 1857-1863. 10.1038/nn.3552.
117. Paxinos, G., Franklin, K.B.J., and Ebscohost. (2019). Paxinos and Franklin's The mouse brain in stereotaxic coordinates, Fifth edition. Edition (Academic Press, an imprint of Elsevier).
118. Wang, Q., Ding, S.L., Li, Y., Royall, J., Feng, D., Lesnar, P., Graddis, N., Naeemi, M., Facer, B., Ho, A., et al. (2020). The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. Cell *181*, 936-953 e920. 10.1016/j.cell.2020.04.007.
119. Song, Z., Alpers, A., Warner, K., Iacobucci, F., Hoskins, E., Disterhoft, J.F., Voss, J.L., and Widge, A.S. (2024). Chronic, Reusable, Multiday Neuropixels Recordings during Free-Moving Operant Behavior. eNeuro *11*. 10.1523/ENEURO.0245-23.2023.
120. Siegle, J.H., Lopez, A.C., Patel, Y.A., Abramov, K., Ohayon, S., and Voigts, J. (2017). Open Ephys: an open-source, plugin-based platform for multichannel electrophysiology. J Neural Eng *14*, 045003. 10.1088/1741-2552/aa5eea.

121. Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M., and Kenneth D., H. (2016). Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. bioRxiv, 061481. 10.1101/061481.
122. Pachitariu, M., Sridhar, S., Pennington, J., and Stringer, C. (2024). Spike sorting with Kilosort4. Nat Methods *21*, 914-921. 10.1038/s41592-024-02232-7.
123. Dayan, P., and Abbott, L.F. (2001). Theoretical neuroscience : computational and mathematical modeling of neural systems (Massachusetts Institute of Technology Press).
124. Fuglstad, J.G., Saldanha, P., Paglia, J., and Whitlock, J.R. (2023). Histological E-data Registration in rodent Brain Spaces. Elife *12*. 10.7554/eLife.83496.
125. Akima, H. (1974). A method of bivariate interpolation and smooth surface fitting based on local procedures. Communications of the ACM *17*, 18-20. 10.1145/360767.360779.
126. Elsayed, G.F., and Cunningham, J.P. (2017). Structure in neural population recordings: an expected byproduct of simpler phenomena? Nat Neurosci *20*, 1310-1318. 10.1038/nn.4617.
127. Benjamini, Y., and Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. Journal of the Royal Statistical Society Series B: Statistical Methodology *57*, 289-300. 10.1111/j.2517-6161.1995.tb02031.x.
128. Karatzas, I., and Shreve, S.E. (1991). Brownian motion and stochastic calculus, 2nd Edition (Springer-Verlag).
129. Kloeden, P.E., and Platen, E. (1992). Numerical Solution of Stochastic Differential Equations 10.1007/978-3-662-12616-5.
130. Neuenkirch, A., and Shalaiko, T. (2016). The Order Barrier for Strong Approximation of Rough Volatility Models. Cornell University Library, arXiv.org.
131. Neuenkirch, A. (2006). Optimal approximation of SDE's with additive fractional noise. Journal of Complexity *22*, 459-474. 10.1016/j.jco.2006.02.001.
132. Mandelbrot, B.B., and Van Ness, J.W. (1968). Fractional Brownian Motions, Fractional Noises and Applications. SIAM review *10*, 422-437. 10.1137/1010093.
133. Goldberg, D.E. (1989). Genetic algorithms in search, optimization, and machine learning (Addison-Wesley Pub. Co.).
134. Pétrowski, A., and Ben-Hamida, S. (2017). Evolutionary algorithms (ISTE).
135. Waltz, R.A., Morales, J.L., Nocedal, J., and Orban, D. (2005). An interior algorithm for nonlinear optimization that combines line search and trust region steps. Mathematical Programming *107*, 391-408. 10.1007/s10107-004-0560-5.
136. Byrd, R.H., Gilbert, J.C., and Nocedal, J. (2000). A trust region method based on interior point techniques for nonlinear programming. Mathematical Programming *89*, 149-185. 10.1007/pl00011391.
137. Efron, B., and Tibshirani, R.J. (1994). An Introduction to the Bootstrap 10.1201/9780429246593.

Acknowledgments

We are grateful to Daniel Dombeck, Christian Ethier, John-Dylan Haynes, Yevgenia Kozorovitskiy, Uri Maoz and Aaron Schurger for helpful discussions and feedback on an earlier version of the manuscript; to Carlos Brody for a helpful discussion; and to Daniel Greenberg for assistance with limb tracking. The computational analyses for this work were largely performed on Quest, Northwestern University's high-performance computing cluster. This research was supported by a Searle Scholar Award, a Sloan Research Fellowship, a Simons Collaboration on the Global Brain Pilot Award and NIH grant DP2 NS120847 awarded to A.M.; a NIH grant R00 NS119787 awarded to J.G.; grants from the NSF (DMS-2235451) and Simons Foundation (MPS-NITMB-00005320) to the NSF-Simons National Institute for Theory and Mathematics in Biology.

Declaration of interests

The authors declare no competing interests.