A brain-wide, trial- and time-dependent deterministic drive for self-initiated action decisions

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

[1]Department of Neurobiology, Northwestern University
[2]Current affiliation: Department of Human Genetics, University of Chicago
[3]National Institute for Theory and Mathematics in Biology
[4]Department of Neurology, Northwestern University
[5]Department of Computer Science, Northwestern University

Correspondence: andrewmiri@northwestern.edu

Abstract

Deciding when to act in the absence of external cues is essential for exploration, learning and survival. Yet how the brain makes such decisions remains controversial, with current models favoring either deterministic or stochastic underpinnings. We performed large-scale, single-unit recordings across eight brain areas in mice self-initiating voluntary actions. We found that action timing is predictable from neural activity, up to several seconds before action initiation, across all eight brain areas. This predictive activity reflects a deterministic drive whose initial value and rate vary trial-by-trial, and whose rate accelerates within trials. While noise contributes to timing variability, the deterministic drive is sufficient to trigger action. Thus, self-initiated timing decisions arise from a variable, brain-wide deterministic drive, challenging models where noise is the primary driver.

Introduction

The adaptive value of cognition manifests not only in the decision of how to act, but also of when to act. For instance, a prey increases its chances of survival by fleeing at a moment its predator cannot predict, rather than waiting for the predator to give chase (*1, 2*). This raises a fundamental question: what neural mechanisms underlie decisions about action timing in the absence of external cues? Seminal studies in humans established over 60 years ago the existence of the "readiness potential", a slow buildup in electroencephalographic recordings that precedes the onset of self-initiated actions by up to several seconds (*3, 4*). Correlates of the readiness potential have since been found across numerous mammalian species, and even in fish and invertebrates, perhaps reflecting a broadly conserved neural dynamics (reviewed in *5, 6*).

The apparent monotonicity of the readiness potential led to the claim that its onset reflects the decision to act, with the drift thereafter proceeding deterministically toward action initiation (*3, 4, 6-10*). Yet, more recent studies have challenged this classic interpretation, implicating a prominent stochastic process in setting self-initiated action timing (*6, 11-15*). In these studies, the monotonicity of the readiness potential is reinterpreted as an artifact of trial-averaging – a method commonly used with non-invasive recording techniques, given their relatively low signal-to-noise ratio (*16*). A combined involvement of both stochastic and deterministic processes in timing decisions has been captured using drift-diffusion models (DDMs; *17-21*). Here, drift and diffusion terms model deterministic and stochastic contributions, respectively, to an accumulation up to a threshold, the

crossing of which reflects decision timing. According to this alternative view, the Gaussian noise modeled by the diffusion term plays a prominent and necessary role in threshold crossing. However, while noise is apparent in neural activity (*22, 23*), the extent to which it effectively contributes to self-initiated action decisions has remained controversial. Without measurements that resolve the neural dynamics that underlie the decision on individual trials, the predictions of the deterministic and stochastic interpretations appear equally plausible (*24*).

In addition, the contributions of different brain regions to action timing decisions also remain poorly resolved. Activity modulation preceding the onset of self-initiated actions has been observed in numerous cortical and subcortical regions, and activity perturbation in some regions alters action timing (reviewed in *9*). Previous studies have imputed specific roles to specific regions, suggesting that action timing decisions rely on a modular and potentially sequential process that unfolds across regions. For instance, the medial prefrontal cortex (mPFC) and secondary motor cortex (M2) have been implicated as underlying the deterministic and stochastic components that govern self-initiated action timing, respectively (*25-28*). However, such studies have typically focused on only one or two regions at a time, and/or used activity recording methods that lack the temporal resolution to distinguish the contributions of regions on the timescale at which the decision process unfolds. Thus, self-initiated action timing may instead rely on a broadly distributed and coordinated process, synchronized across multiple brain regions.

Results

Self-initiated action timing is predictable from spiking activity

To address these questions, we sought to capture single-unit, single-trial activity dynamics across multiple brain regions simultaneously in mice initiating actions, without external instruction about when to act. We employed a head-fixed paradigm we have recently developed in which mice self-initiate a naturalistic climbing behavior. Here mice climb with all four limbs across handholds that extend radially from a wheel positioned in front of them (Fig. 1A; *29*). Mice self-initiated ~75 discrete climbing bouts on average per session. We computed climbing onset times using video-based limb tracking (*30*) during epochs surrounding the onset of wheel rotation (Fig. 1B).

To record neural activity, four large-scale multi-electrode arrays (Neuropixels; *31*) were acutely implanted at the beginning of a climbing session (six mice, one 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 (Fig. 1, C to F and fig. S1): mPFC (*25, 26, 28, 32*), M2 (*13, 25*), the primary motor cortex (M1; *10, 33-35*), the mediodorsal nucleus of the thalamus (MD; *36*), the globus pallidus (GP; *26, 27, 37-40*) and, contralaterally, the deep cerebellar nuclei (interposed and dentate nuclei, DCN) and the cerebellar simple lobule (SL; *41-43*). The array targeting cerebellar regions also recorded activity in the primary visual cortex (V1; *44*). Across recording sessions, the number of single units meeting our quality criteria ("qualifying units") in individual regions ranged from 35 (median across sessions, DCN) to 398 (mPFC; fig. S1). In all these regions of interest, trial-averaged single unit activity appears modulated around climbing onset and during ongoing climbing (Fig. 1G).

To test for a deterministic process underlying the action timing decisions mice made in our paradigm, 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 (*45*) to predict the latency until climbing onset (time-to-movement) at time points ranging from the end of the preceding movement (immobility onset) to the beginning of the next (climbing onset). The duration of the "trials", defined as the epochs between immobility onset and climbing onset, governed by each mouse's internal deliberation, ranged from 1.7 ± 0.55 to 3.6 ± 1.2 (mean $\pm$ std) seconds across mice. At each time point, the forecast model used the firing rates of qualifying units averaged in 50 ms bins over the preceding 250 ms. Predicted time-to-movement often showed a ramping time course on individual held-out trials (Fig. 2, A and B), the form of which rules out the confounding explanation that it reflects neural activity tracking the passage of time (supplementary text 1, fig. S2; *46*). Predictions were significantly better than chance for a majority of trials based on their residuals (Fig. 2, C and E), and for nearly all trials based on their slopes (Fig. 2, D and E). Thus, there appears to be a prominent deterministic component to action timing at the single-trial level.

To evaluate when predictive information emerges relative to immobility onset on individual trials, we examined how the residuals (Fig. 2F) and slopes (Fig. 2G) of the forecast model output within sliding 250 ms windows leading up to climbing onset compared to chance expectations. We found that for a majority of trials, significant predictability emerges within the first 500 ms after immobility onset (Fig. 2H), which corresponds to a median time of 1.4 s before movement onset (fig. S3). Collectively, these results indicate that neural activity is predictive of self-initiated action timing on single trials, up to several seconds in advance (fig. S3).

Decision-related activity is synchronized across brain regions

Is the activity used by the forecast model to predict time-to-movement localized in a subset of brain regions? In all eight regions of interest, we found qualifying units that significantly improved forecast accuracy ("forecasting neurons"). When aligned on climbing onset, the trial-averaged activity of these neurons often exhibited a ramping time course leading up to onset (Fig. 3A). This pattern was also often apparent on single trials (fig. S4). Across all qualifying single units, the fraction that qualified as forecasting neurons ranged from 27% (median, mPFC) to 39% (M2; Fig. 3, B and C). Thus, the activity of neuronal ensembles within each region of interest carries information about action timing. This suggests that a deterministic process underlying action timing unfolds across all of these regions.

We next asked whether self-initiated action timing involves a sequential or synchronized process across brain regions. We trained separate forecast models using activity from each individual region. The resulting single-trial forecast traces were remarkably similar across all the regions examined here, often exhibiting coordinated deviations from a linear prediction of time (Fig. 3D). Time-to-movement predictions were positively correlated on individual trials across all pairs of regions (Fig. 3, E and F). Moreover, when these cross-region correlations were calculated at each time step, they remained consistently positive and significant throughout the trials, even from their emergence soon after immobility onset leading up to action initiation (Fig. 3G). This persistent, brain-wide coordination strongly suggests that self-initiated action timing emerges from a synchronized, distributed process rather than a sequential, modular one.

Parameters of the deterministic drive vary across trials

The preceding results demonstrate the existence of a substantial, brain-wide deterministic component underlying self-initiated action timing on individual trials. However, the duration of immobility between climbing bouts varied considerably across trials, and predictions exhibited substantial errors (Fig. 2, A and E). We thus sought to model the decision process in a manner that could account for this timing variability. Based on models applied to perceptual decision-making in humans (*17, 24, 47*), we considered whether variability in action timing might to some extent reflect discrete variation in the parameters (initial value and/or slope) of the deterministic component from one trial to the next.

To test this possibility, we assessed whether trial-by-trial parameter variation could account for some degree of the variability across trials in the output of our time-to-movement forecast models. We plotted the initial value of the model output (i.e., the forecasted time-to-movement at 250 ms after immobility onset), its slope when best fit with a line, and its final value (i.e., the forecasted time-to-movement at action onset) versus immobility duration across trials. The initial values of the forecast traces were negatively correlated with immobility duration in all six mice (Fig. 4A). This suggests that early in the immobility period, the neural activity that predicts action timing is variable from one trial to the next, in a manner that influences action timing. Similarly, we found that the best-fit linear slopes for the forecast traces were also negatively correlated with immobility duration (Fig. 4B). This is consistent with a linear approach to a decision threshold that has a variable slope from one trial to the next, in a manner that also influences action timing. Lastly, the final values of the forecast traces, though substantially variable, were uncorrelated with immobility duration (Fig. 4C). This is consistent with a similar threshold for action initiation across trials.

These observations are compatible with the involvement of two discrete sources of timing variability: one that varies the deterministic drive's initial value and one that varies its slope from one trial to the next. Conversely, these observations are not consistent with a typical DDM, which assumes no such discrete source of variability. With a typical DDM, the initial value and the slope of the drift term are indeed similar across trials, and timing variability arises only from a continuous source of noise, captured by the diffusion term (*11, 24*).

Assessing the contribution of continuous noise to timing variability

Based on these results, we sought to identify an updated model for the decision process and to quantify the relative contribution of a continuous noise source. The typical DDM (Equation 1) consists of a four-parameter model for Δx_i, the change in the decision variable at time t_i: x_0 is the initial value of the decision variable; I is the slope of the deterministic drift term (drift rate); c is the amplitude of the noise from the diffusion term with ξ_i drawn from a standard normal distribution, and k is a leak coefficient (*11*).

$$\Delta x_i = I\Delta t + c\xi_i\sqrt{\Delta t} - kx_i\Delta t \qquad [1]$$

We thus formulated a DDM variant that involves both discrete and continuous sources of variability in action timing (equations 2-4). Here the initial value and drift rate are drawn from normal and log-normal distributions, respectively, on individual trials. Previous studies of stimulus-driven perceptual decision have observed increases in drift rate toward a decision threshold over time within individual trials (*48, 49*). Thus, we also added a growth factor α to the drift term to capture

potential time-dependent changes in drift rate. This yielded a model for Δx_i with seven parameters: $\mu_{intercept}$, $\sigma_{intercept}$, μ_{drift}, σ_{drift}, c, k and α.

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

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

$$\Delta x_i = I_0 e^{\alpha t_i} \Delta t + c\xi_i \sqrt{\Delta t} - k x_i \Delta t \qquad [4]$$

To justify the inclusion of several of the model parameters, we used model selection. The dependence of drift parameters on immobility duration we observed (Fig. 4B) could to some extent reflect a feature of the decision process typically captured by the leak term (supplementary text 2). We thus assessed whether the parameters μ_{drift}, σ_{drift}, and k improve the fit to our time-to-movement forecast traces. We did the same for α since time-dependent growth factors have yet to be tested in models of self-initiated action decisions. We used cross-validated maximum likelihood estimation to measure the fit improvements from including each parameter (Fig. 4D, fig. S5). All four parameters significantly increased the cross-validated likelihood for all mice (Fig. 4, D and E, and fig S6). Best-fit α values were positive, indicating that the drift rate increases over time. Furthermore, simulations revealed that the trial-to-trial variability in the deterministic drive's initial value (Fig. 4A) is substantially larger than what would be expected without a discrete source of variability governing the initial state of the decision process on each trial (supplementary text 3, fig. S7). Thus, our observations are consistent with a new DDM variant in which discrete sources of variability set the deterministic drive's parameters on a trial-by-trial basis.

We then used our variant DDM to assess the relative contribution of a continuous noise source captured by the diffusion term in setting action timing. We first verified that our observations are most consistent with a source of noise that is not autocorrelated (supplementary text 4, fig. S8). We then performed an ablation analysis: we simulated forecast traces using the DDM variant with best-fit parameters for each session, both with and without the diffusion term (Fig. 4F). We found that the decision threshold is reachable both with and without the diffusion term, though its presence contributes to slightly shortening the average duration of the trials (Fig. 4, G and H). In total, these analyses indicate that a trial- and time-dependent deterministic drive plays a prominent role in self-initiated action timing, while noise appears less critical than previously assumed.

Finally, we asked whether our results using multi-region activity hold for single-neuron activity within each region of interest. To address this, we applied the same model selection approach as above to the firing rates of the forecasting neurons previously identified in each region (*50*). This analysis revealed that in every region, most forecasting neurons exhibit a drift rate that varies both across trials and over time within trials (Fig. 5A). Best-fit parameter values were also similar across regions (Fig. 5B, fig. S9 and fig. S10). Furthermore, for most forecasting neurons, the effect of the diffusion term is only to shorten the average time to reach their decision threshold (Fig. 5C). These results demonstrate that a similar decision process is reflected in the activity of individual neurons across all regions of interest. This again suggests a broadly distributed decision process coordinated across regions.

Discussion

Here we have addressed a long-standing debate over the neural basis of self-initiated action decisions, by resolving single-trial, single-unit dynamics across multiple brain regions simultaneously. We identified a deterministic component preceding action onset by up to several seconds on individual trials that was ubiquitously present, from the cortex to the thalamus, basal ganglia and cerebellum. A buildup of activity is a genuine feature of single trials, not an artifact of trial-averaging. Our results indicate that the deterministic drive varies on a trial-by-trial basis, both in its initial value and its drift rate. Moreover, the drift rate is time-dependent, increasing within a trial. This increase is consistent with an urgency signal, previously proposed to underlie perceptual decisions (*48, 49, 51*). Altogether, our findings support a model of self-initiated action timing dominated less by noise and more by a deterministic drive, that resembles a ballistic-like accumulation mechanism with parameters that vary across trials (*24, 47*).

Our finding that noise is not necessary for the decision process to reach its threshold stands in contrast to recent "stochastic accumulator models", where continuous noise is not only the sole source of inter-trial differences in action timing, but is also required to drive the decision process across threshold (Fig. S11; *6, 11*). Our results rather demonstrate that the deterministic drive is sufficient on its own to reach threshold. However, this does not render the diffusion term irrelevant. Indeed, our analysis revealed that while not essential, this term reduces the average time to threshold crossing. Taken together, our results are at odds with both fully deterministic and predominantly stochastic accounts: they suggest that variability in deterministic drive parameters and noise act in concert to make action timing unpredictable.

Our new model includes three distinct sources of timing variability: (1) trial-to-trial variability in the initial value of the deterministic drive, (2) trial-to-trial variability in the drive's drift rate, and (3) continuous, within-trial noise. These do not appear to be manifestations of a single underlying noise source. For instance, the variance of the initial values is substantially larger than what would be generated by the diffusion term (fig. S7). This suggests that distinct neurophysiological mechanisms likely govern these sources. Neuromodulatory systems appear as prime candidates for setting the trial-specific parameters of the deterministic drive. Notably, the basal forebrain cholinergic system has been implicated in self-initiated action timing (*26, 27, 52*) and provides the main source of acetylcholine to the cortex (*53*), including the prefrontal and motor areas where we observed synchronized decision-related activity. Cholinergic input is known to modulate cortical gain (*54*), so cholinergic tone at the start of a trial could set the rate of the deterministic drive.

The highly similar decision-related activity we observed across eight regions of interest suggests that the decision process is implemented collectively by a brain-wide network. This highly similar decision-related activity across regions appears to contrast with previous suggestions of a more modular organization whereby, for example, deterministic and stochastic components might be segregated in different regions (*25-28*). By leveraging the statistical power of large-scale spike recordings together with model selection, a more distributed picture emerged here. This echoes the broad coordination across brain regions reported recently during motor execution and perceptual decision-making (*44, 55-59*).

Main Figures

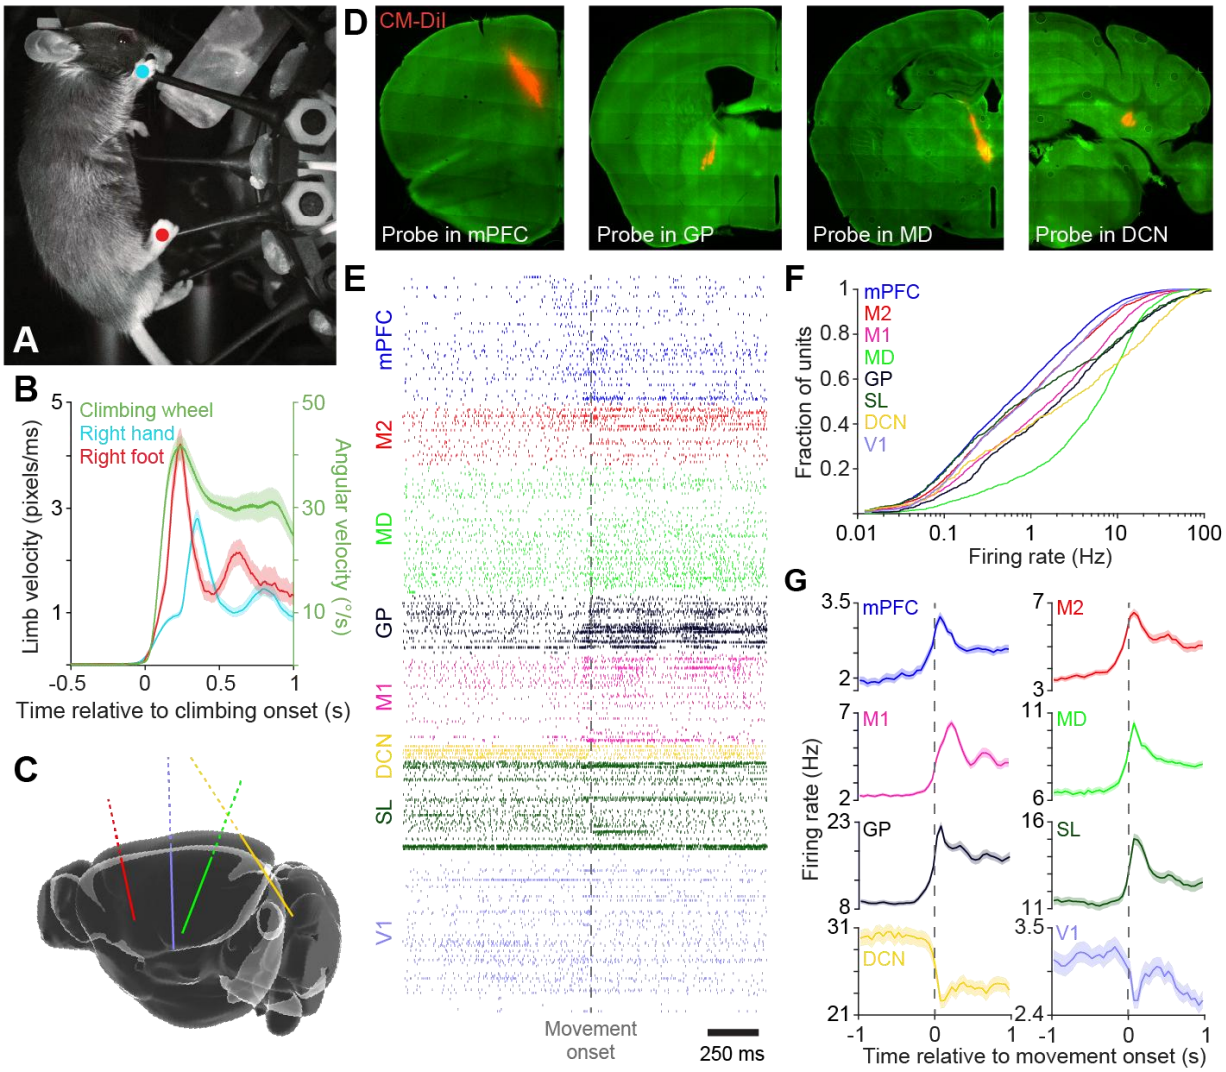

Fig. 1. Large-scale neural recordings during self-initiated climbing. (**A**) Head-fixed mouse performing self-initiated climbing on a wheel with radial handholds. (**B**) Trial-averaged limb velocity for the right hand and right foot (left axis) and angular velocity of the climbing wheel (right axis), aligned to climbing onset. (**C**) 3D schematic of a mouse brain with approximate trajectories and locations of the four acutely implanted Neuropixels probes. (**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). (**E**) Raster plots of single-unit activity from one session, aligned to movement onset, for all eight recorded regions of interest: mPFC, M2 (secondary motor cortex), M1 (primary motor cortex), MD, GP, DCN, SL (simple lobule) and V1 (primary visual cortex). (**F**) Cumulative distribution of firing rates for all units in each region, pooled across all recording sessions. (**G**) Trial-averaged firing rates (mean ± 95% confidence interval) for all qualifying units in each region, aligned to climbing onset.

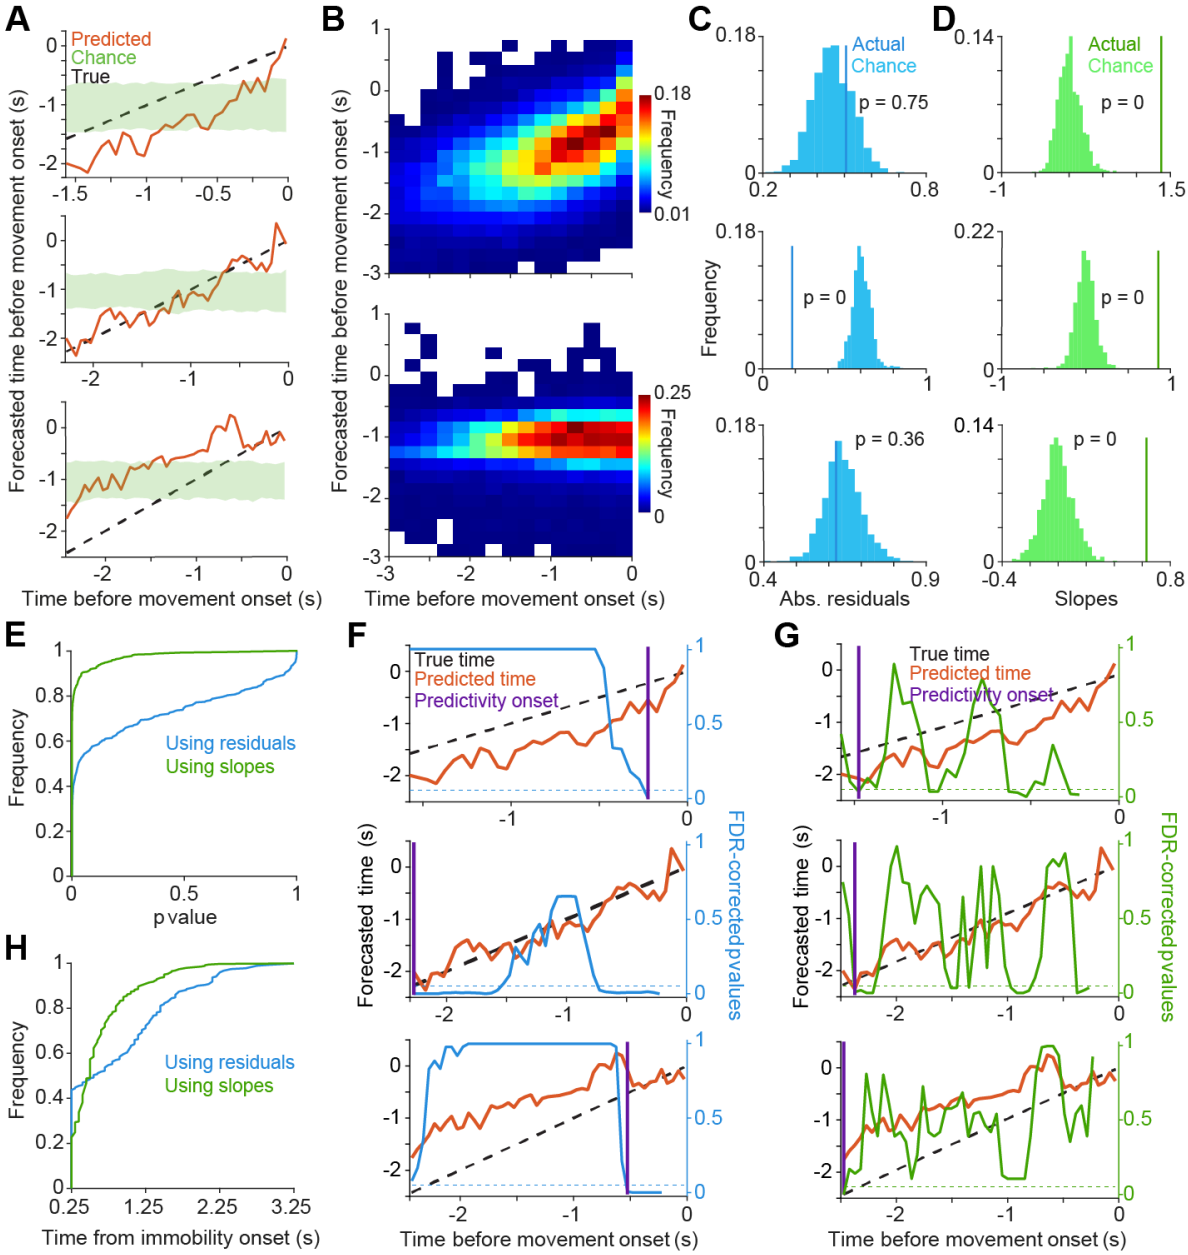

Fig. 2. Self-initiated action timing is predictable from single-trial neural activity. (**A**) Three example single-trial predictions of time-to-movement (red) versus true time-to-movement (dashed black line). The green shaded area represents the 95% confidence interval of predictions from shuffled data. Predictions are made every 50 ms by an elastic net linear model using 250 ms of past neuronal activity. Here and in subsequent analyses, all shown and analyzed data correspond to cross-validated held out data. (**B**) 2D histograms of forecasted against actual time-to-movement for all trials pooled across all recorded sessions; actual data (top) and shuffled data (bottom). (**C, D**) Absolute residual (C) and slope (D) for each example trial from (A), compared with their corresponding chance distributions. Each statistic is calculated as a single value across the entire trial duration. (**E**) Cumulative distribution of p-values from the statistical tests illustrated in (C-D), across all trials. (**F, G**) Identification of predictability onset for the example trials shown in (A). Using

either residuals (F) or slopes (G) calculated in 250-ms sliding windows, predictability onset is marked when the forecast traces (red) significantly deviate from chance (FDR-corrected p-value < 0.05). (**H**) Cumulative distribution of predictability onset times relative to immobility onset for all trials, determined using either residuals or slopes.

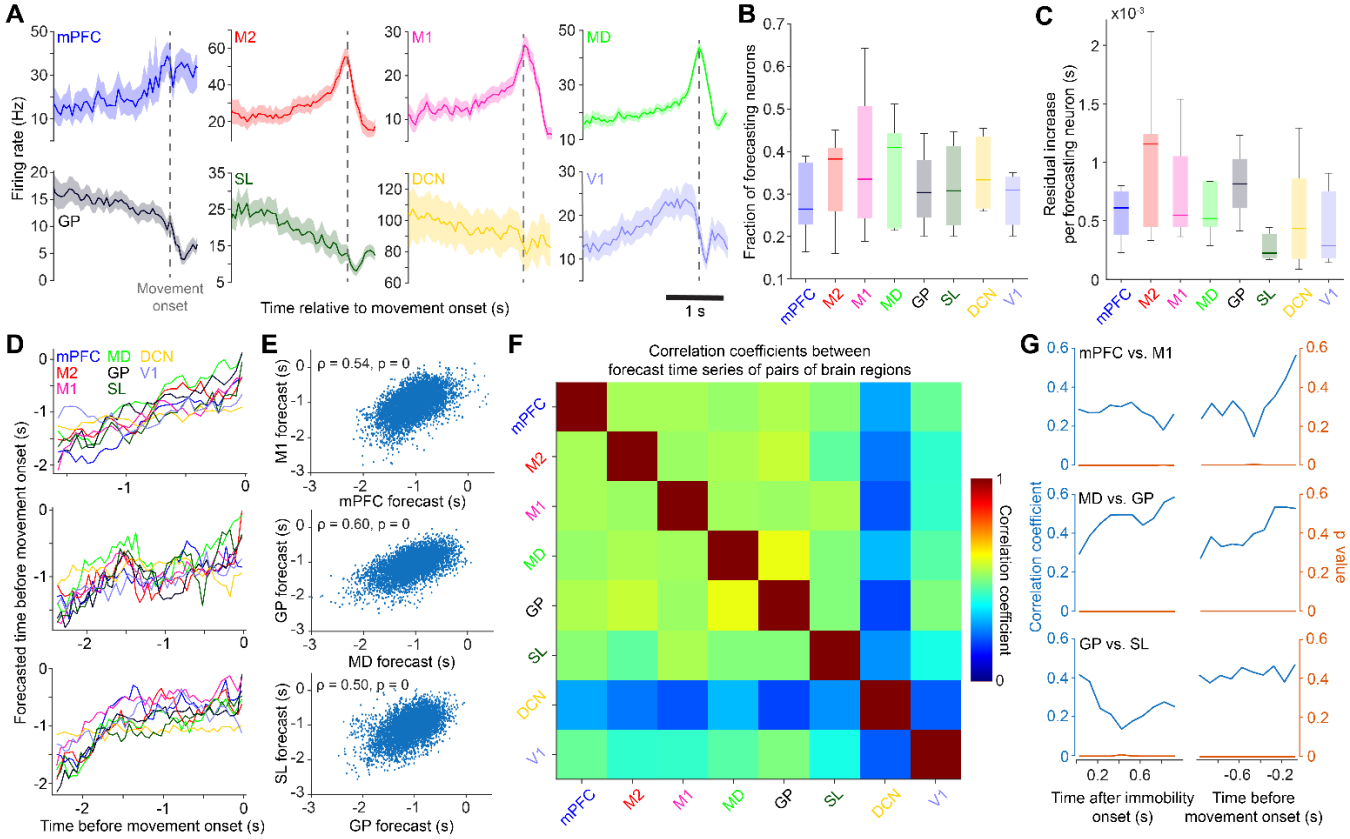

Fig. 3. Decision-related activity is synchronized across brain regions. (**A**) Trial-averaged firing rate of individual forecasting neurons from each of the eight regions, aligned to climbing onset. (**B**) Proportion of forecasting neurons in each region across all sessions. (**C**) Increase in forecast error upon removal of single forecasting neuron from each region across all sessions. (**D**) Per-region time-to-movement predictions for three example single trials. (**E, F**) Time-to-movement predictions are positively correlated across all pairs of regions. The correlations were computed using forecast data from immobility onset up to 250 ms before movement onset, using all trials of a given recording session. (E) Example scatter plots for three pairs of brain regions, comparing their time-to-movement predictions across all time points and all trials for a given recording session. ρ denotes the Pearson correlation coefficient. (F) Correlation matrix of time-to-movement predictions showing significant positive correlations for all pairs of regions. (**G**) Cross-region correlation coefficients (blue) and associated p-values (red) calculated at each time step across all trials of a given recording session.

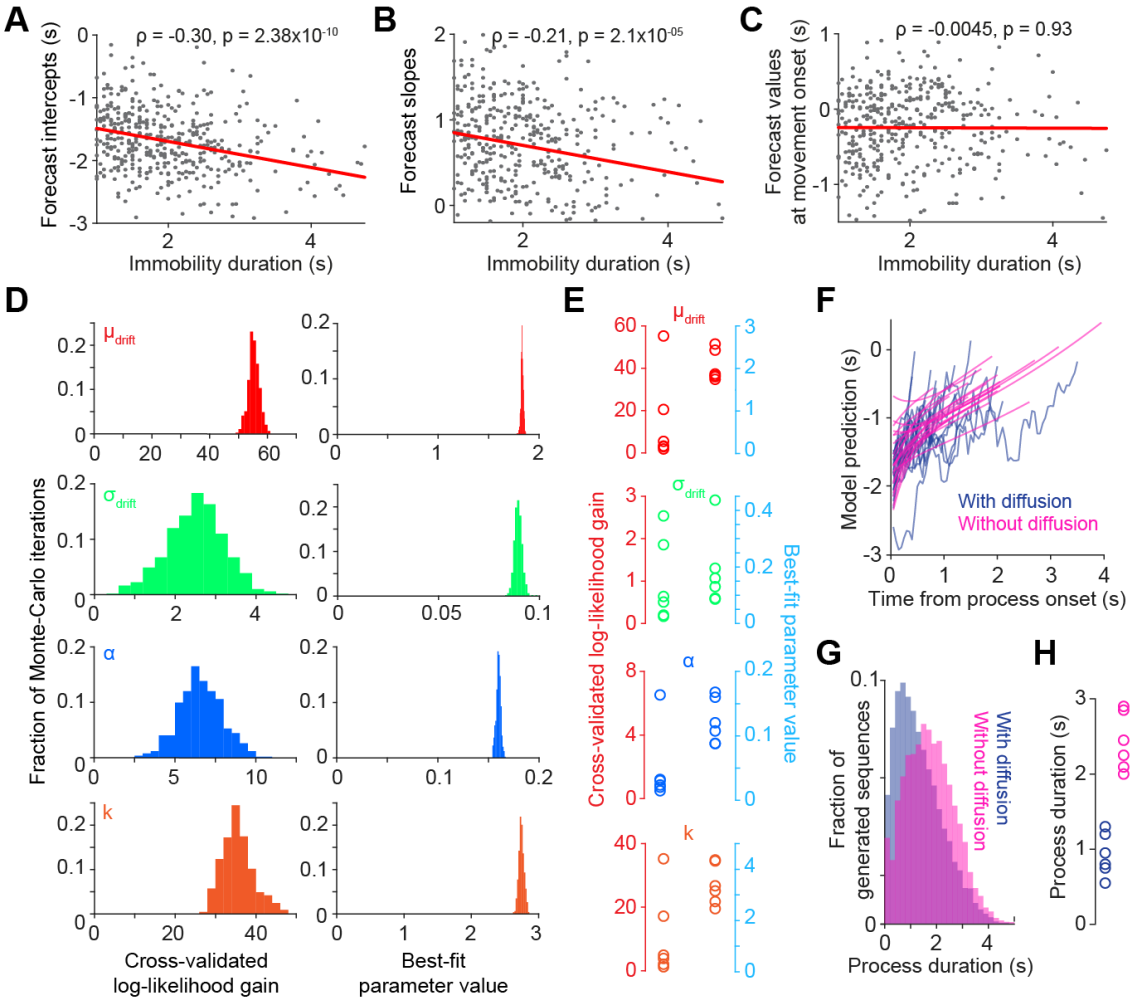

Fig. 4. A trial- and time-dependent deterministic drive underlies self-initiated action timing. (**A**) Negative correlation between the initial value of the forecast traces and immobility duration across trials from all sessions. Each point corresponds to a trial. ρ denotes the Pearson correlation coefficient. (**B**) Negative correlation between the linear slope of the forecast traces and immobility duration across trials. For each trial, the slope is computed over the last second prior to movement onset. (**C**) Lack of correlation between the final forecast value at movement onset and immobility duration. (**D, E**) DDM model selection. (D) DDM model selection results for a given session. Distributions show the cross-validated log-likelihood gain (left column) and the best-fit parameter value (right column). Each parameter's contribution is assessed by comparing a full model to a simpler, nested version lacking that parameter: a mean drift rate (μ_{drift}, red), trial-to-trial variability in drift rate (σ_{drift}, green), a time-dependent drift rate (α, blue), and a leak term (k, orange). Note that μ_{drift} and σ_{drift} correspond to the mean and standard deviation of the underlying normal distribution for $log(I_0)$ (Equation 3). The mean and standard deviation of the log-normal distribution itself are calculated from these parameters and shown in fig. S6. (E) Median cross-validated log-likelihood gains (left) and best-fit parameter values (right) for each of the six mice. Each circle represents a mouse. All cross-validated log-likelihood gains are significant (p ≈ 0). The best-fit value of α implies a ~1.5-fold increase for a 2-s trial, indicating a substantial contribution to action timing

from the growth factor. (**F, G, H**) Testing for the sufficiency of the deterministic drive to reach the decision threshold. (F) Simulated sequences from the best-fit DDM, with (blue) and without (magenta) the diffusion term. (G) Distribution of process durations for simulated trials with and without the diffusion term for a given session. (H) Median process duration from simulations with (blue) and without (magenta) diffusion for each recording session.

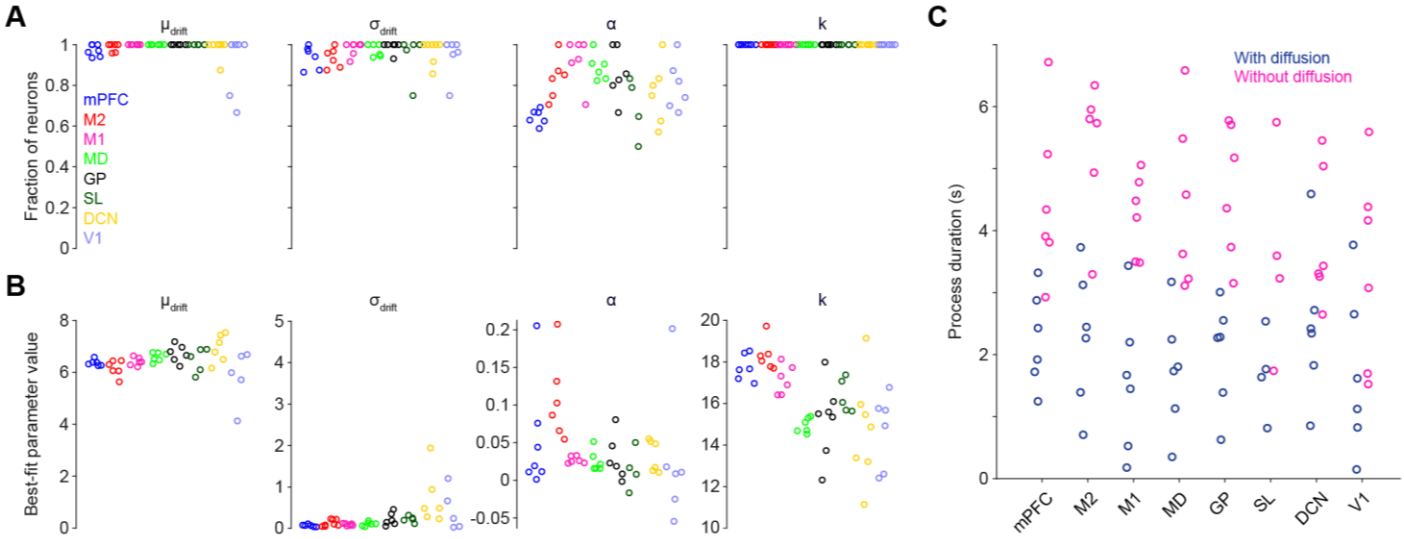

Fig. 5. A similar decision process is reflected in the activity of individual neurons across all recorded regions. (**A**) Per-session fraction of forecasting neurons for which each of the five key DDM parameters significantly improved model fit, color-coded by brain region. (**B**) Per-session median of best-fit parameter values for all forecasting neurons that significantly improved the DDM fit, color-coded by brain region. (**C**) Comparison of median process duration from simulated sequences with (blue) and without (magenta) the diffusion term across forecasting neurons, showing that diffusion shortens the time to threshold crossing. Each circle represents one session.

Supplementary Figures

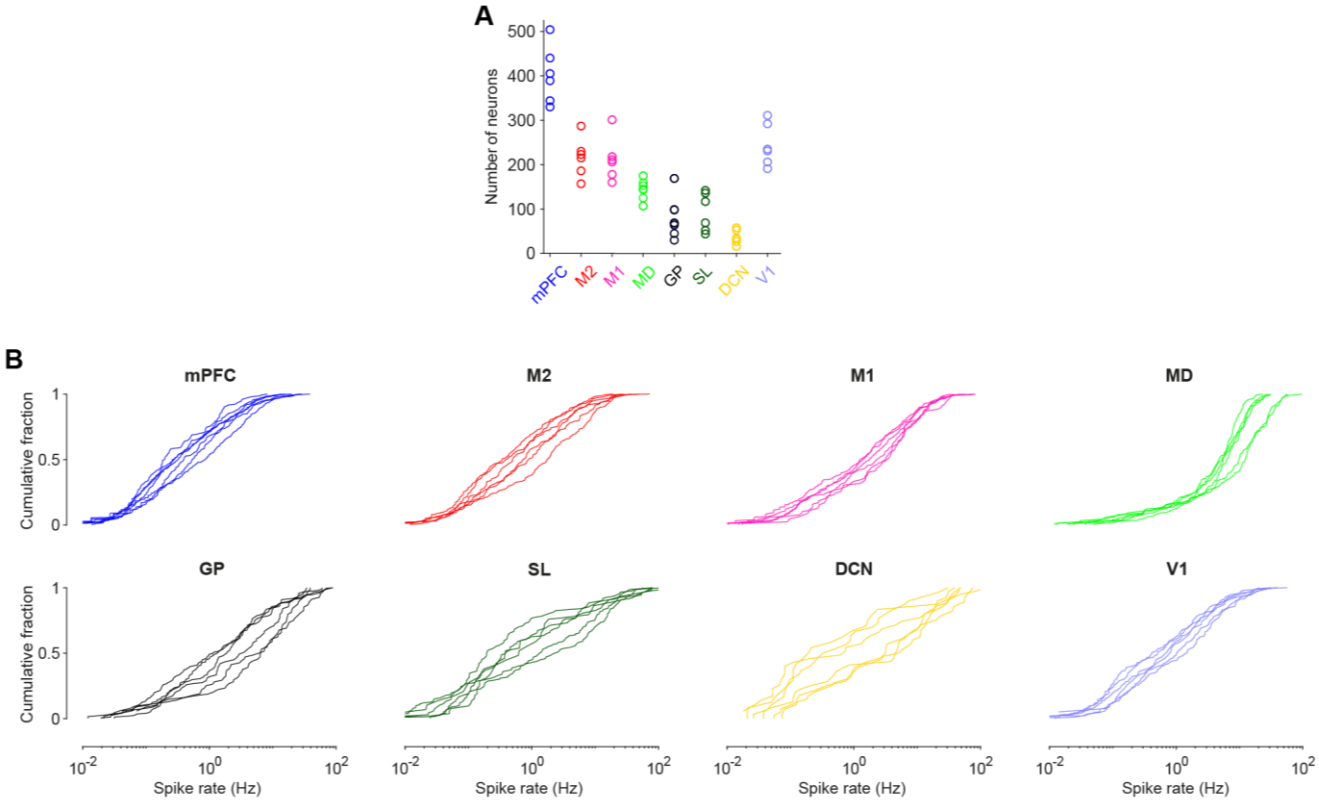

Fig. S1. Unit recordings and firing rate characteristics. (**A**) Number of qualifying single units per recorded region across all six recording sessions, with each circle representing one session. (**B**) Cumulative distributions of spike rates for units within each region, with each trace representing one session.

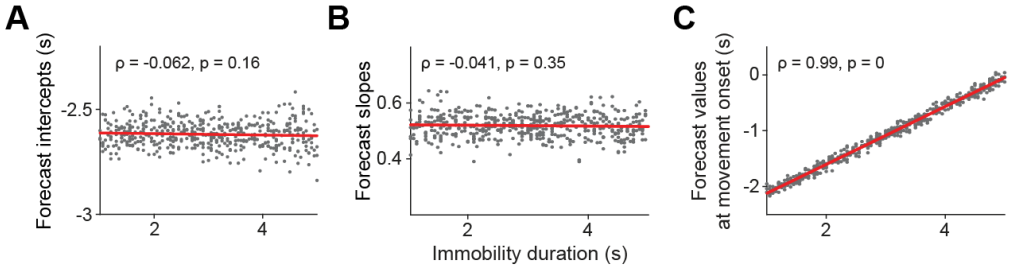

Fig. S2. Properties of forecast traces obtained using synthetic "clock-like" activity. (A) Lack of correlation between the initial forecast values and immobility duration. ρ denotes the Pearson correlation coefficient. **(B)** Lack of correlation between the forecast slopes and immobility duration. **(C)** Positive correlation between the final value of the forecast traces correlated with immobility duration. These properties are qualitatively opposite to those obtained using the actual neural data (Fig. 4A-C).

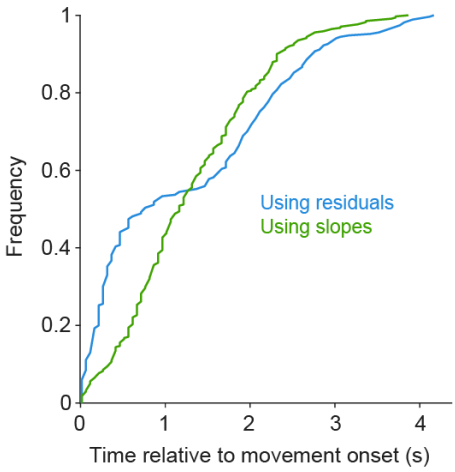

Fig. S3. Emergence of predictability relative to movement onset. Cumulative distribution of predictability onset times relative to movement onset for all trials, determined using either residuals or slopes. Same data as in Fig. 2H but aligned to movement onset rather than immobility onset.

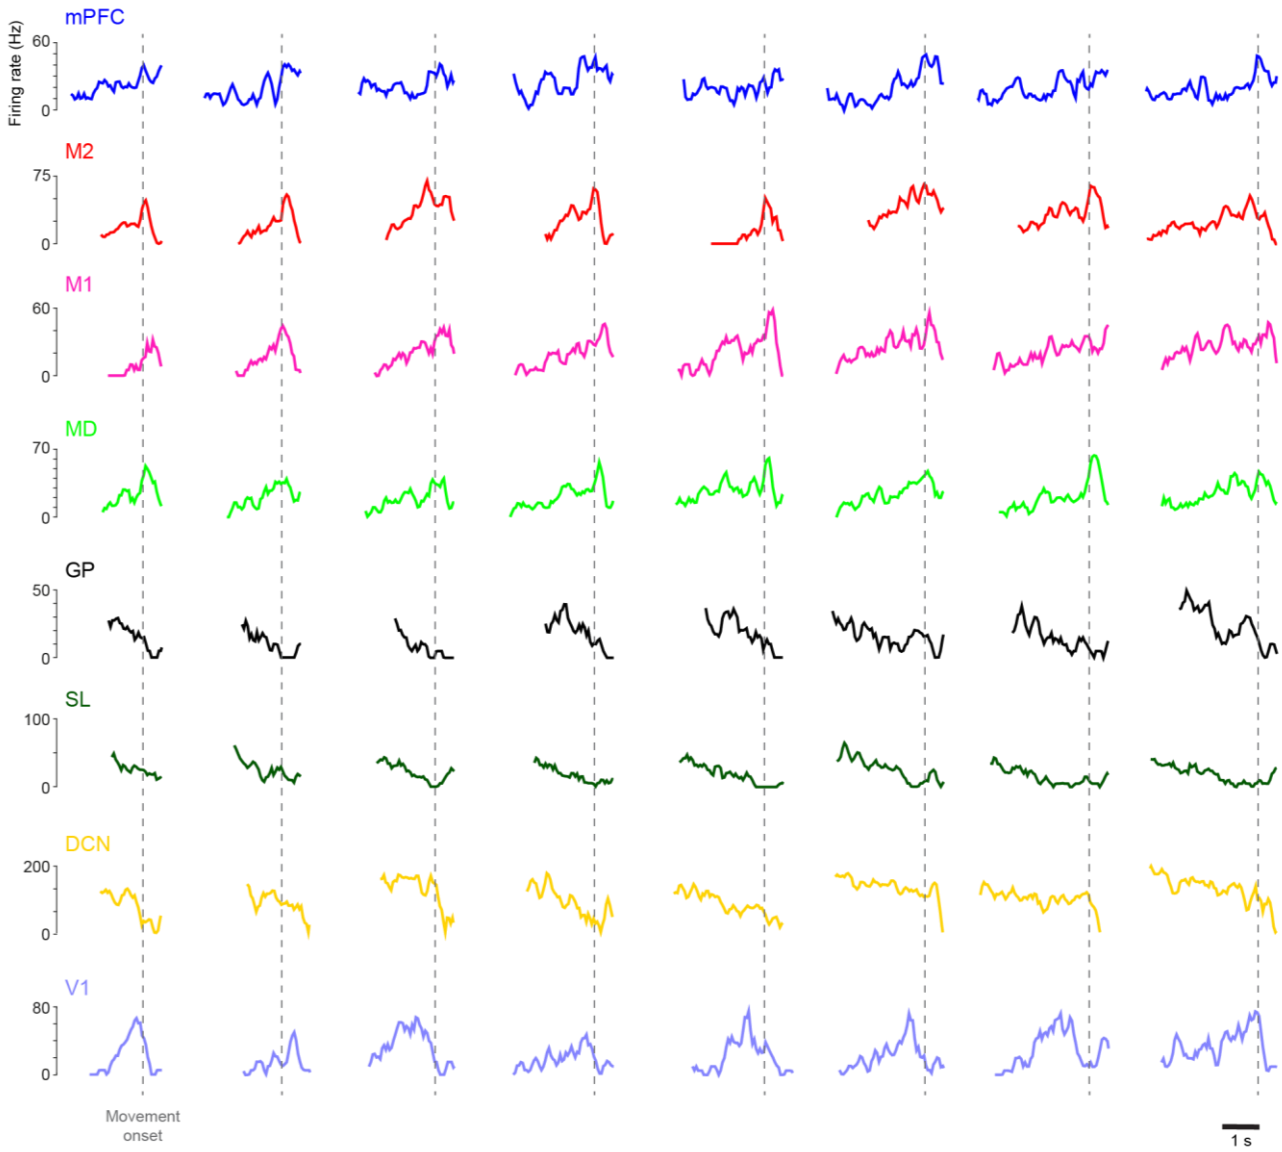

Fig. S4. Single-trial activity of example forecasting neurons. Single-trial firing rate traces from a forecasting neuron for each of the eight recorded regions, aligned to movement onset. Trials are sorted by increasing immobility durations. These example neurons correspond to the ones whose trial-averaged firing rates are shown in Fig. 3A.

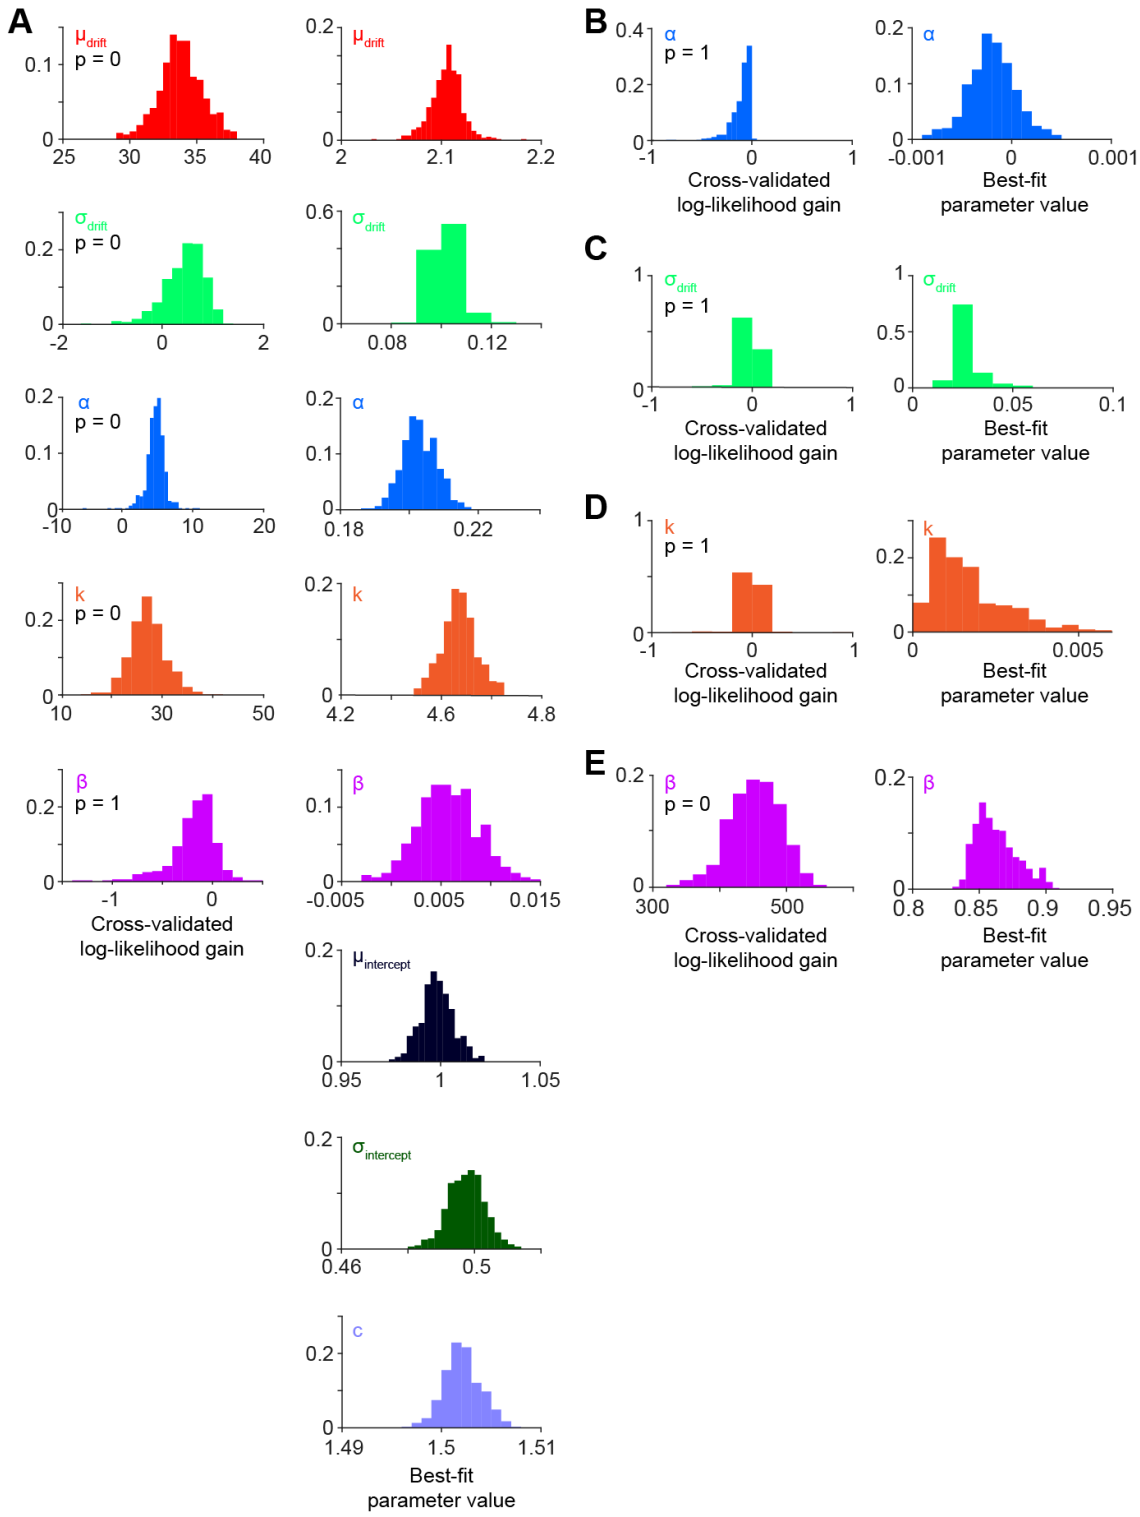

Fig. S5. Validation of the DDM fitting and model selection procedure via parameter recovery. (**A**) Ground-truth parameters for generating sequences are: $\mu_{drift} = 2$, $\sigma_{drift} = 0.1$, $c = 1.5$, $k = 4$, $\mu_{intercept} = 1$, $\sigma_{intercept} = 0.5$, $\alpha = 0.2$, and $\beta = 0$ (white noise). (A) Recovered parameter distributions for sequences generated with the 7-parameter model (β fixed to 0). (**B**) Recovered

parameters for sequences generated with a reduced 6-parameter model ($\alpha = 0$). (**C**) Recovered parameters for sequences generated with a reduced 6-parameter model ($\sigma_{drift} = 0$). (**D**) Recovered parameters for sequences generated with a reduced 6-parameter model ($k = 0$). (**E**) Recovered parameters for sequences generated with the full 8-parameter model ($\beta = 1$).

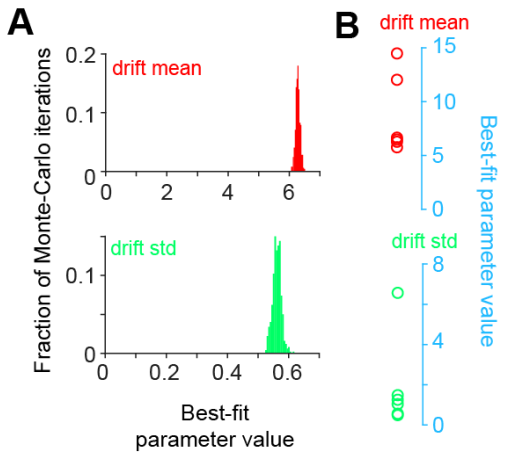

Fig. S6. Best-fit parameter values for the mean and standard deviation of the initial drift rate I_0. (**A, B**) $drift\ mean$ and $drift\ std$ correspond to the mean and standard deviation of the log-normal distribution from which the initial drift, I_0, is drawn on a per-trial basis. These are distinct from, but calculated using, the fitted parameters μ_{drift} and σ_{drift} which define the mean and standard deviation of the underlying normal distribution for $log(I_0)$ (Equation 3 and Fig. 4, D and E). (A) Example from the same session as in Fig. 4D. (B) Median cross-validated best-fit parameter values for each of the six mice. Each circle represents a mouse.

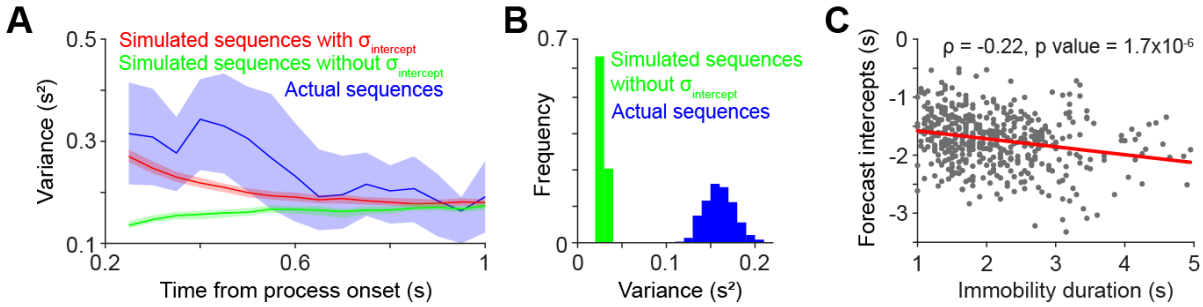

Fig. S7. Contribution of a discrete source of trial-to-trial variability for the initial value of the deterministic drive. (A) Variance of the forecast traces over time for a given representative session. The variance of the actual data (blue) is high from the beginning of the trial and matches simulations from the best-fit DDM that includes a discrete source of initial variability ($\sigma_{intercept}$, red). In contrast, simulations lacking this source (green) show a much lower initial variance that grows over time, indicating that the observed initial variability cannot be explained by a drift-diffusion process with constant initial values. Shaded areas represent the 95% bootstrap confidence interval. **(B)** Distribution of forecast variance at 50 ms after immobility onset, using a memory-less forecast model for the same session as in (A). The variance of the actual data (blue) is significantly larger than that generated by one step of the DDM process without a discrete source of initial variability (green). **(C)** The initial forecast values, at 50 ms after immobility onset, are significantly and negatively correlated with the total immobility duration across trials from all sessions, using a memory-less forecast model. This confirms that the initial state of the process is already predictive of the final action time. ρ denotes the Pearson correlation coefficient.

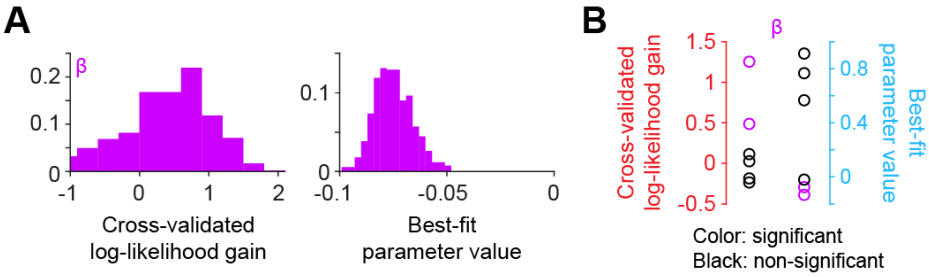

Fig. S8. Model selection provides no evidence for temporally autocorrelated (colored) noise. To test whether the diffusion term had temporal correlations, the DDM was augmented with a spectral exponent β, where $\beta = 0$ corresponds to white noise. (**A**) Distributions of the cross-validated log-likelihood gain (left) and the best-fit β value (right) for the same session as in Fig. 4D. (**B**) Summary of results for all six mice. Median cross-validated log-likelihood gains (left) and best-fit parameter values (right) for each of the six mice. Each circle represents a mouse. The inclusion of β did not significantly improve the fit for 4 of the 6 mice (black circles, $p > 0.05$). For the two mice for which the improvement was statistically significant (purple circles), the best-fit β values were close to zero. Together, these results justify the use of our 7-parameter model with white noise in the Main text.

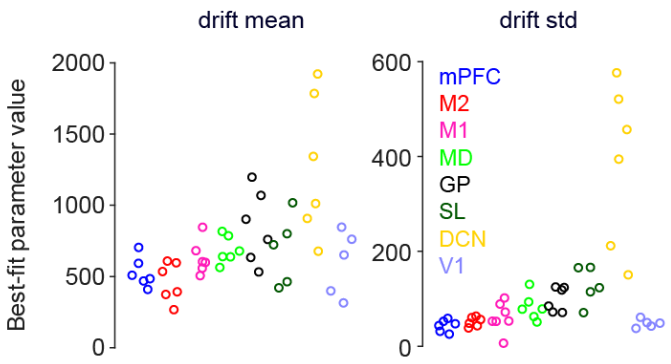

Fig. S9. Best-fit parameter values for the mean and standard deviation of the initial drift rate, I_0, for individual forecasting neurons. $drift\ mean$ and $drift\ std$ correspond to the mean and standard deviation of the log-normal distribution from which the initial drift, I_0, is drawn on a per-trial basis. These are distinct from, but calculated using, the fitted parameters μ_{drift} and σ_{drift} which define the underlying normal distribution for $log(I_0)$ (Equation 3 and Fig. 5B). Each circle represents the median best-fit parameter value for all forecasting neurons within a given brain region for one recording session. Colors denote the different brain regions.

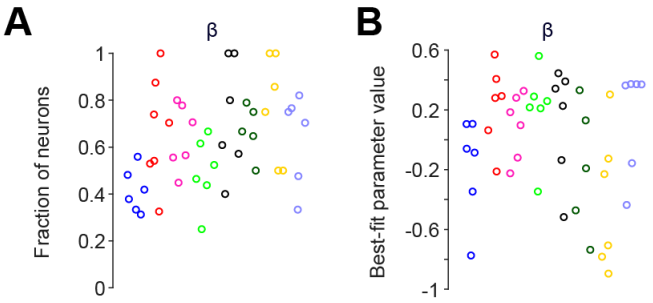

Fig. S10. Model selection for a colored noise parameter at the single-neuron level. To test for temporal correlations in the noise for individual neurons, the DDM was augmented with a spectral exponent β, where a value of 0 corresponds to white noise. (**A**) Per-session fraction of forecasting neurons in each brain region for which including the β parameter resulted in a statistically significant improvement in model fit. (**B**) Per-session median of best-fit β parameter values for all forecasting neurons that significantly improved the DDM fit, color-coded by brain region.

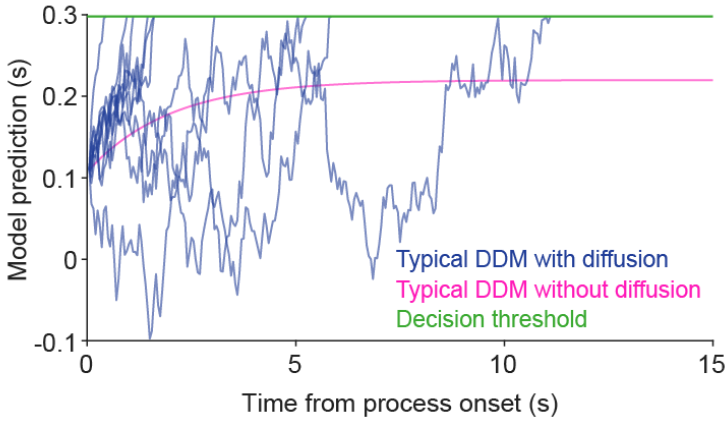

Fig. S11. Necessity of the diffusion term to reach the decision threshold, in previous electroencephalographic-based accounts of self-initiated actions in humans (*11*). Simulations from a typical DDM (Equation 1) using parameters from a previous study about self-initiated action decisions ($final\ value\ =\ 0.298, I = 0.11, c = 0.1, k = 0.5$; *11*). The deterministic component alone (magenta trace) fails to reach the decision threshold (green): the addition of a diffusion term (blue traces) is required for threshold crossing.

Supplementary text

Supplementary text 1. Predictive neural activity reflects a decision process, not a time-tracking mechanism.

We rule out the confounding explanation that neural activities underlying action onset predictability may reflect a neural mechanism for tracking the passage of time (*46*), rather than being decision-related. Indeed, "clock-like" activity would have initial values and slopes that are uncorrelated with immobility duration, while its final values would be positively correlated with immobility duration (fig. S2). Instead, we observed the opposite (Fig. 4, A to C).

Supplementary text 2. Model selection supports trial-to-trial variability in drift rate over confounding effects of drift-leak dynamics.

Our analysis of the forecast traces revealed a negative correlation between their best-fit linear slopes and the duration of immobility across trials (Fig. 4B). One possible explanation is that this reflects a genuine feature of the decision process, where the rate of the deterministic drive toward action varies from one trial to the next. However, an alternative explanation could arise from the inherent properties of a drift-diffusion process including a leak term. In such a model, the deterministic drift is counteracted by a leak term that grows stronger as the decision variable accumulates. Over the course of a long trial, the drift and leak can approach a dynamic equilibrium, causing the rate of accumulation to slow down (*24*). This progressive slowing could create the appearance of a lower overall slope on longer trials, even if the underlying drift rate were constant across all trials.

To distinguish between these possibilities, we used model selection to explicitly test whether the parameter modeling trial-to-trial variability in the drift rate, σ_{drift}, significantly improved the fit to our forecast traces, beyond the contribution of the leak term k. The inclusion of σ_{drift} significantly improved the cross-validated likelihood for all mice (Fig. 4, D and E). This result indicates that the observed negative correlation is not simply a consequence of drift-leak dynamics, but rather reflects a trial-by-trial variation in the rate of the deterministic drive.

Supplementary text 3. Trial-to-trial variability in the deterministic drive's initial value implies a discrete source of variability.

A key feature of our DDM variant is that the initial value of the deterministic drive can vary from trial to trial, a property captured by the $\sigma_{intercept}$ parameter. However, the structure of our analysis prevents a straightforward model selection test for $\sigma_{intercept}$ (i.e., comparing models with and without it). Indeed, our DDM is fit to the forecast traces which begin at 250 ms after immobility onset, as the forecast model uses 250 ms of past neural activity for each prediction. Consequently, the first data point of each forecast trace reflects the outcome of the initial 250 ms (i.e., five 50-ms time steps) of the underlying drift-diffusion process. Thus, to avoid distorting the estimation of the dynamic parameters (c, k, α), our fitting procedure treats the initial forecast values separately, modeling their distribution with just the $\mu_{intercept}$ and $\sigma_{intercept}$ parameters. As a result, the fitted $\sigma_{intercept}$ is a "lumped" parameter, capturing variance from two potential sources: (1) the drift-diffusion process itself over the initial 250 ms after immobility onset, and (2) any discrete variability in the process's starting value at immobility onset, i.e., a discrete source of initial variability.

To disentangle these sources and test for a discrete source of initial variability, we performed a simulation-based analysis. We tested the null hypothesis that the observed variance in the forecast's initial values (i.e., at 250 ms after immobility onset) could be fully explained by the drift-diffusion process running for 250 ms from a fixed starting point. To do this, we simulated our DDM using the best-fit parameters for each session, but set the discrete variability in initial values to zero. We found that the variance of the actual forecast traces over their initial values is substantially larger than that of the simulated sequences with no discrete source of initial variability (Fig. S7, A and B). This indicates that the drift-diffusion process over the first 250 ms is insufficient to account for the observed variability. In contrast, simulations that included the fitted $\sigma_{intercept}$ term closely matched the variance of the actual data from 250 ms after immobility onset to several hundreds of ms later (Fig. S7A).

This result was further strengthened by another analysis using a "memory-less" forecast model that used only the preceding 50 ms of neural data. With this model, the initial forecast values (at 50 ms after immobility onset) are also significantly more variable than what would be expected from a single, 50-ms time step of the DDM process without initial variability (Fig S7B). Finally, these initial forecast values (at 50 ms after immobility onset) showed a significant negative correlation with immobility duration (Fig. S7C), corroborating the finding from our primary forecast model (at 250 ms after immobility onset, Fig. 4A).

Altogether, these results support the existence of a mechanism that sets the initial value of the deterministic drive on a trial-by-trial basis, in a manner that influences action timing.

Supplementary text 4. Testing for temporal autocorrelation in the continuous noise component.

To address the possibility that autocorrelation in the stochastic component captured by the diffusion term may partly explain our results (*22, 60, 61*), we augmented our variant DDM with an eighth parameter, a spectral exponent β, which introduces dependence across time points in the diffusion term. Adding this parameter did not improve the fit for 4 mice out of 6, and for those for which it did, the best-fit values for β imply noise that is nearly white ($\beta \approx 0$) (fig. S8).

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 was implanted in a separate craniotomy over the right frontal lobe, soldered to a connector pin 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 (*62*) along with the Allen Mouse Brain Common Coordinate Framework (*63*), aided by the "Neuropixels trajectory explorer software" developed by Andy Peters. Targets in the left hemisphere included mPFC, M2, M1, MD and GP. Targets in the right hemisphere were V1, SL and DCN. Table S1 recapitulates the coordinates for the four probes 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)	Recording start (AP, ML, DV)	Recording end (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 (*64*), 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 degrees apart. A ratchet mechanism permitted the wheel to rotate only in the downward direction, consistent with the mouse's climbing motion. The wheel's absolute angular position was continuously tracked by a shaft encoder (A2-A-B-D-M-D, U.S. Digital). All experimental events were controlled, and the shaft encoder's signal acquired, via custom MATLAB (MathWorks) scripts and a National Instruments DAQ card (PCI-e-6323, NI).

Behavioral training

Mice were initially acclimated to being handled by the experimenter, to head-fixation, and to the climbing wheel apparatus according to previously described procedures (*64*). 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 scaled to their recent performance to maintain motivation.

Once mice were proficient, the final behavioral protocol also used during recording sessions was initiated. This protocol consisted of two period types controlled by a custom MATLAB script – "go periods" and "no-go periods" – which occurred in a structured, alternating sequence (go, then no-go, then go, etc.). Cue presentation was counterbalanced across animals: for some mice, auditory and visual 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, both go and no-go periods lasted for an average of 3 seconds, with a random jitter. The rules in case of climbing were as follows:

- During a "go period", a single climb would terminate the period and trigger a reward.
- During a "no-go period", mice could initiate climbing bouts ad libitum, at times of their choice. In case of climbing, the period would only conclude after a climbing sequence was followed by a continuous immobility of 3 seconds on average.

The lick rate indicated that mice were not expecting reward during the no-go periods while they were during the go periods, confirming the self-initiated nature of the climbing bouts occurring during no-go-periods. The present study focuses exclusively on the self-initiated climbing sequences, which

occur during no-go periods. A typical behavioral session lasted one hour.

Electrophysiological recording

Prior to each acute recording session, several preparatory steps were taken. To facilitate smooth penetration through the dura mater, the tips of the four Neuropixels 1.0 probes (IMEC) were sharpened using a grinder (EG-45, Narishige), following a protocol established by Nicholas Steinmetz and Adam Davis. A few minutes before insertion, the sharpened tips were coated with a fluorescent lipophilic dye (CM-DiI, Invitrogen) to mark the probe tracks for post-hoc histological verification, as per a protocol established by Liu Lui (https://www.protocols.io/view/painting-neuropixels-probes-and-other-silicon-prob-5jyl8neprl2w).

For the recording session, the mouse was head-fixed in the climbing apparatus and the protective silicone elastomer was removed from the craniotomies. The four Neuropixels probes, mounted on a 3-axis multi-probe micromanipulator system (New Scale Technologies), were then slowly lowered to their target depths. The exposed skull was covered with silicone oil.

Electrode voltages were acquired at 30 kHz using the Open Ephys acquisition software. To ensure precise temporal alignment, all data streams were synchronized using a rhythmic TTL pulse. The first role of this pulse was to synchronize the data streams from each of the four independent Neuropixels headstages, which operate at slightly different sampling rates. This common pulse stream, acquired by each headstage, was also simultaneously recorded by the NI-DAQ card controlled by MATLAB. This allowed for a subsequent, precise alignment of the entire synchronized neural dataset with all behavioral and video data.

Spike sorting and firing rate estimation

Raw electrophysiological data was spike-sorted using Kilosort 3 (*65, 66*). Only units classified as "good" single units were retained for further analysis. To obtain a continuous estimate of the instantaneous firing rate for each neuron, the spike trains were convolved with a Gaussian kernel ($\sigma = 10\ ms$) (*67*). The resulting smoothed signals were then binned into 50 ms windows to compute the firing rates used in subsequent analyses. Finally, we applied a selection criterion to focus on task-relevant neurons. Only units with a mean firing rate of at least 0.5 Hz during the periods from -1 s to -0.25 s relative to movement onset were included. These units are referred to as "qualifying units".

Video recording

Video tracking of the behavior was done using a high-resolution monochrome camera (acA2040-55um, Basler), fitted with a 16 mm fixed focal length lens (33-304, Edmund Optics). The camera was positioned to the right of the head-fixed mouse and videos were acquired continuously throughout the entire electrophysiology session. Recording was performed under near-infrared illumination at 100 frames per second with a resolution of 800 x 700 pixels.

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

The continuous video recordings were analyzed to extract limb kinematics. We used DeepLabCut (*30*) to track the positions of the mouse's right hand and right foot. The resulting marker coordinates were then used to compute kinematic variables, such as limb velocity.

Histological procedure and brain registration

At the conclusion of the recording sessions, mice were deeply anesthetized and transcardially perfused with 3% paraformaldehyde (P6148, MilliporeSigma). Brains were extracted, post-fixed in paraformaldehyde, and subsequently sectioned into 60 µm thick coronal slices using a sliding freezing microtome (SM2000, Leica). The sections were then mounted on glass slides with a mounting medium containing DAPI (00-4959-52, Invitrogen). Whole-slide images were acquired using a slide scanner fluorescence microscope (VS120, Olympus). To precisely verify the final placement of the probe tracks, the resulting images were registered to the Allen Mouse Brain Common Coordinate Framework (CCFv3) using the open-source Python package HERBS (*68*).

Climbing onset detection

To precisely identify the onset of self-initiated climbing bouts, we used a synergistic, two-stage approach that combined wheel rotation data with video-based limb kinematics.

1. First, a coarse detection of climbing periods was performed using the angular displacement of the wheel. The signal from the shaft encoder, acquired at 1 kHz, was smoothed, and a climbing bout was defined as a continuous rotation exceeding an angular displacement of 10 degrees. 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 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 interpolation, and an instantaneous velocity was calculated for each limb. A velocity threshold, empirically set to 0.1 pixel/ms to separate genuine limb movements from baseline noise, was applied to classify time points as either mobile or immobile. To distinguish genuine climbing from twitches, we applied strict temporal filtering: (1) brief periods of mobility lasting less than 100 ms were excluded 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. 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 degrees. Finally, the climbing bout had to be preceded by a continuous period of whole-body immobility (defined by both limbs being still) of at least 1 second.

An "immobility period", which constitutes a single trial in our analyses, was defined as the continuous period of immobility preceding the onset of a climbing bout.

Single-trial firing rate visualization

The single-trial firing rate traces shown in fig. S4 were generated by smoothing the binned spike counts for each trial. This was achieved by applying a causal moving average, where the value at each

time point was calculated by averaging the data from that point with those of the preceding 200 ms window.

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 for its simplicity and interpretability. The model was trained to predict the time remaining until action onset ("time-to-movement") at each time step. Trials, which had variable durations corresponding to the animal's self-determined immobility period, 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 250 ms (i.e., five consecutive 50 ms bins of activity). As a control, we also trained a "memory-less" model using only the single preceding 50 ms bin of activity (fig. S7, B and C). We used a 5-fold cross-validation scheme to train and test the 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 estimate a chance-level prediction, we generated a null distribution by training 1,000 separate models on shuffled data. Specifically, for each null model, the relationship between neural activity and the time-to-movement labels was broken by circularly shuffling the time-to-movement labels of the training data. This procedure preserves the temporal autocorrelation structure within the neural data while disrupting its specific predictive relationship with the timing of the action (*69*). The statistical significance of the forecast for each held-out test trial was then assessed with a permutation test (Fig. 2, C to E). We used two metrics to evaluate the model's performance across the full duration of each trial: 1) the mean absolute residual between the predicted and true time-to-movement ; 2) the slope of the best-fit line of the forecast traces. For each test trial, the true value of these metrics was compared against the null distribution of 1,000 values obtained from the predictions of the shuffled-data models to compute a p-value. This allowed us to determine, on a trial-by-trial basis, whether the model's predictive accuracy was significantly better than chance.

To determine when predictive information emerged on individual trials, we adapted the permutation test to a sliding window analysis (Fig. 2, F to H). For each trial, we calculated the same two metrics – mean absolute residual and slope – within a 250 ms sliding window that advanced in 50 ms increments. This procedure generated a sequence of p-values over time for each trial, where each p-value represented the forecast's significance within that specific temporal window. We applied the Benjamini-Hochberg procedure (*70*) 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.

Identification of forecasting neurons

To identify which individual neurons contributed significantly to the forecast (Fig. 3, A to C), we implemented a permutation-based feature analysis embedded within a repeated 5-fold cross-validation framework. This approach allowed us to assess each neuron's predictive value in the context of the full forecast model.

The entire procedure was repeated 1,000 times, and within each repetition, the following steps were performed for each of the 5 cross-validation folds:

1. An elastic net model was trained on the actual training/validation data.

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 shuffling 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 shuffled test data.

4. The effect of shuffling a neuron was quantified by the change in mean absolute residual (MAR) compared to the MAR on the original, unshuffled 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.

We favored this approach – disrupting a single neuron's activity in the test set – over the alternative of removing each neuron one-by-one and re-training the model from scratch. The re-training method can indeed be confounded by informational redundancy across the neural population. If a predictive neuron is removed, a newly trained model can often compensate for its absence by up-weighting other, correlated neurons that carry similar information. This would mask the true predictive contribution of the removed neuron, leading to an underestimation of its importance. In contrast, our method directly probes the extent to which the original, fully trained model relies on the specific temporal information from each neuron, providing a more sensitive assessment of their individual contribution.

Control analysis with synthetic "clock-like" activity

To test the alternative hypothesis that the neural activity used by our forecast model was tracking the passage of time since immobility onset, rather than a decision-related process, we performed a control analysis using synthetic "clock-like" activity (supplementary text 1, fig. S2). We first generated a synthetic dataset consisting of sequences with variable durations. The duration of each sequence was randomly drawn from a uniform distribution between 1 and 5 seconds. The activity of a single artificial neuron was then generated for each trial. Its activity was modeled as a perfect linear ramp to which we added independent Gaussian noise at each time step. This corresponds to a noisy representation of time elapsed since the start of the sequences. This synthetic dataset was then processed through the exact same 5-fold cross-validated elastic net forecast pipeline used for the real multi-neuron data. Finally, we analyzed the output of the model by comparing the relationship between the forecast output's parameters (initial value, slope, and final value) and trial duration (fig. S2) to the results obtained with the actual neural recordings (Fig. 4, A to C).

Drift-diffusion modelling

To quantitatively model the decision process, we developed and compared a series of nested drift-diffusion models. The models were fit to the time-to-movement forecast traces (Fig. 4) and to the activity of the forecasting neurons (Fig. 5), on a per-session basis.

Model formulations and likelihood functions

Our modeling approach started from an Ornstein-Uhlenbeck process (*71*, *72*) and incrementally added parameters to build towards our most comprehensive model. All models were based on the discretization of a stochastic differential equation (SDE) using the Euler-Maruyama method, a standard numerical procedure for approximating the solution of SDEs (*73*). This was performed with a time step of $\Delta t = 50\,ms$, matching our data's bin size. Below are the key models and their corresponding log-likelihood functions for a single observed trial trajectory $x = [x_1, x_2, \dots, x_T]$. A special consideration was made for the first data point (x_1), whose distribution was modeled separately as $p(x_1) = \mathcal{N}\left(x_1; \mu_{intercept}, \sigma^2_{intercept}\right)$ (supplementary note 3).

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

This baseline model represents a leaky accumulator with no drift term ($dx_t = -kx_t dt + cdW_t$ with $dW_t = \xi\sqrt{dt}$).

The parameters are $\theta_{OUP} = \left[c, k, \mu_{intercept}, \sigma_{intercept}\right]$.

Its log-likelihood is given by:

$$\log \mathcal{L}\left(\theta_{OUP}|x\right) = \log \mathcal{N}\left(x_1; \mu_{intercept}, \sigma^2_{intercept}\right) + \sum_{i=1}^{T-1} \log \mathcal{N}\left(x_{i+1}; x_i - kx_i\Delta t, c^2\Delta t\right)$$

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

This model adds a constant drift rate, I_0, to the OUP ($dx_t = (I_0 - kx_t)dt + cdW_t$).

The parameters are $\theta_{constant\ drift} = \left[c, k, \mu_{intercept}, \sigma_{intercept}, I_0\right]$.

Its log-likelihood function is:

$$\log \mathcal{L}\left(\theta_{constant\ drift}|x\right) = \log \mathcal{N}\left(x_1; \mu_{intercept}, \sigma^2_{intercept}\right) + \sum_{i=1}^{T-1} \log \mathcal{N}\left(x_{i+1}; x_i + (I_0 - kx_i)\Delta t, c^2\Delta t\right)$$

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

This is the primary model described in the Main text, which incorporates trial-to-trial variability in the initial drift rate and a time-dependent urgency signal ($dx_t = (I_0 e^{\alpha t} - kx_t)dt + cdW_t$).

The parameters are $\theta_{7\ params} = \left[c, k, \alpha, \mu_{intercept}, \sigma_{intercept}, \mu_{drift}, \sigma_{drift}\right]$.

The initial drift rate I_0 is a latent variable drawn from a log-normal distribution ($I_0 \sim Lognormal(\mu_{drift}, \sigma^2_{drift})$) to ensure its positivity. Because I_0 is a latent variable, the log-likelihood is calculated by integrating it out:

$$\log \mathcal{L}\left(\theta_{7\ params}|x\right) = \log \mathcal{N}\left(x_1; \mu_{intercept}, \sigma^2_{intercept}\right) + \log \int \left(\prod_{i=1}^{T-1} p(x_{i+1}|x_i, I_0)\right) p(I_0) dI_0$$

with $p(x_{i+1}|x_i, I_0) = \mathcal{N}(x_{i+1}; x_i + (I_0 e^{\alpha t_i} - k x_i)\Delta t, c^2 \Delta t)$

The contribution of the leak (k) and urgency (α) parameters was tested by comparing this full model to nested versions where k or α were respectively fixed to zero.

4. Full model with colored noise (8 parameters)

To test for temporal autocorrelations in the noise (*61*), we extended the full model with a spectral exponent parameter β.

The parameters are $\theta_{8\ params} = \left[c, k, \alpha, \mu_{intercept}, \sigma_{intercept}, \mu_{drift}, \sigma_{drift}, \beta\right]$.

In this case, the noise increments are not independent. The likelihood of the transitions is therefore given by a single multivariate normal distribution over the vector of innovations $\boldsymbol{D}(\boldsymbol{x}, I_0) = [D_1, D_2, \dots, D_{T-1}]$ where each element D_i represents the prediction error at step i:

$$D_i = x_{i+1} - \underbrace{(x_i + (I_0 e^{\alpha t_i} - k x_i)\Delta t)}_{\text{Model's prediction for } x_{i+1}}.$$

The log-likelihood function becomes:

$$\log \mathcal{L}\left(\theta_{8\ params}|x\right) = \log \mathcal{N}\left(x_1; \mu_{intercept}, \sigma^2_{intercept}\right) + \log \int N(D(x, I_0); 0, \Sigma_D) p(I_0) dI_0$$

The covariance matrix Σ_D is determined by the noise autocorrelation function ρ, such that $(\Sigma_{\boldsymbol{D}})_{ij} = c^2 \Delta t \cdot \rho(|i-j|; H)$, where the Hurst exponent H is related to the spectral exponent by $H = (\beta + 1)/2$. The parameter β was constrained to be between -1 and +1 to ensure the stationarity of the noise process.

For both models involving a latent variable (7 and 8 parameters), the integral over I_0 was computed using Laplace approximation as a faster alternative to numerical integration.

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 (*74, 75*) 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 (*76, 77*) to refine the minimal negative likelihood estimate (function *fmincon* in MATLAB). Parameters such as leak, 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 repeated cross-validated maximum likelihood estimation. This was performed within a Monte-Carlo framework (1,000 repetitions) using a 5-fold cross-validation scheme. For each repetition and fold, we computed the difference in the log-likelihood between the full model and the nested model on the held-out test

data, i.e., trials on which the model was not trained. This resulted in a [1000 repetitions x 5 folds] matrix of log-likelihood differences for each parameter tested for each session. To determine statistical significance, we performed a non-parametric sign permutation test on this matrix (the same as the one used to identify forecasting neurons, described above). A parameter was considered to significantly improve the model if the resulting p-value was less than 0.05 (Fig. 4D, E). We confirmed that our model selection approach could accurately recover parameters using simulated data (fig. S5).

Data fitted by the DDM

For each session, we fit the drift-diffusion models to two types of data: 1) the forecast traces, and 2) the single-trial spike rates of individual forecasting neurons. For both data types, the fitting procedure used trials from their beginning up to 250 ms before action onset. The start of the trial was included based on our finding that predictive information typically emerges within the first 500 ms for most trials (Fig. 2H). The end point was chosen to isolate the decision-making process and minimize the influence of potential motor execution-related activity. As a pre-processing step, because our DDM formulation assumes a positive drift rate, the forecast traces were first made non-negative by subtracting from all forecast traces the minimum value observed across all trials within that session.

Calculation of log-normal drift parameters

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 $log(I_0)$. For visualization purposes (figs. S6 and S9), we calculated the mean ($drift\ mean$) and standard deviation ($drift\ std$) of the log-normal distribution itself using the following equations:

$$drift\ mean = \exp\left(\mu_{drift} + \frac{\sigma_{drift}^2}{2}\right)$$

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

Simulations for testing for a source of initial variability

As described in supplementary text 3, we performed two analyses to test the null hypothesis that a dynamic process alone could account for the observed initial variance, one using our primary forecast model, another using a memory-less forecast model.

<u>Primary forecast model (fig. S7A)</u>

This analysis compared the variance of the actual forecast data from 250 ms after immobility onset to two sets of simulations:

1. Full model simulation. To generate trajectories that replicate our fitting procedure, the first point of each simulated forecast trace – corresponding to 250 after immobility onset – was not generated dynamically. Instead, it was drawn directly from the fitted initial distribution, $\mathcal{N}\left(\mu_{intercept}, \sigma_{intercept}^2\right)$. This reflects our fitting method, where the dynamic parameters (c, k, α, and

the drift) do not contribute to this first-time step. The subsequent points of the simulated trajectories were evolving dynamically.

2. Null hypothesis simulation. To generate trajectories under the null hypothesis (no discrete initial variability), each simulation started at t = 0 from a single fixed value. The process then evolved for five 50-ms time steps with all best-fit dynamic parameters (c, k, α, and the drift) to generate a distribution of values at t = 250ms.

Memory-less forecast model (fig. S7B)

This control analysis used a forecast model using a single (50-ms) time bin for each prediction, allowing us to test the hypothesis of a discrete source of initial variability at t = 50 ms after immobility onset. We compared the variance of the actual forecast data at this time point to a null hypothesis simulation. For this simulation, trajectories were generated for a single 50 ms step, starting from a fixed point and evolving with all dynamic parameters (c, k, α, and the drift) applied immediately in the first step.

Testing the role of the diffusion term via ablation analysis

To assess the necessity of the diffusion term for reaching the decision threshold, we performed an ablation analysis. Using the best-fit parameters for each session, we simulated trial 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 the empirical data, the start and end points for each simulated trajectory were drawn from distributions derived from the actual data. Specifically, the starting point of each trajectory was drawn from the fitted initial value distribution, $\mathcal{N}(\mu_{intercept}, \sigma^2_{intercept})$. The decision threshold for that same trajectory was drawn from the normal distribution fitted to the actual data values observed at 250 ms prior to climbing onset. This endpoint was thereby defined as a realistic decision boundary while avoiding contamination from post-decisional, motor execution-related activity. We then compared the distributions of the time taken to cross the same decision threshold between the two simulation conditions. This entire procedure was performed for both the DDM fitted to the multi-neuron forecast traces (Fig. 4, F to H) and the models fitted to the individual forecasting neurons (Fig. 5C).

References

1. M. Broom, G. D. Ruxton, You can run—or you can hide: optimal strategies for cryptic prey against pursuit predators. *Behavioral Ecology* **16**, 534-540 (2005).
2. P. Domenici, J. M. Blagburn, J. P. Bacon, Animal escapology I: theoretical issues and emerging trends in escape trajectories. *J Exp Biol* **214**, 2463-2473 (2011).
3. H. H. Kornhuber, L. Deecke, 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 (1965).
4. B. Libet, C. A. Gleason, E. W. Wright, D. K. Pearl, 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 (1983).
5. F. Di Russo *et al.*, Beyond the "Bereitschaftspotential": Action preparation behind cognitive functions. *Neurosci Biobehav Rev* **78**, 57-81 (2017).
6. A. Schurger, P. Hu, J. Pak, A. L. Roskies, What Is the Readiness Potential? *Trends Cogn Sci* **25**, 558-570 (2021).
7. H. Shibasaki, M. Hallett, What is the Bereitschaftspotential? *Clin Neurophysiol* **117**, 2341-2356 (2006).
8. C. S. Soon, M. Brass, H. J. Heinze, J. D. Haynes, Unconscious determinants of free decisions in the human brain. *Nat Neurosci* **11**, 543-545 (2008).
9. T. Dominik, A. Mele, A. Schurger, U. Maoz, Libet's legacy: A primer to the neuroscience of volition. *Neurosci Biobehav Rev* **157**, 105503 (2024).
10. C. Mitelut *et al.*, Mesoscale cortex-wide neural dynamics predict self-initiated actions in mice several seconds prior to movement. *Elife* **11**, (2022).
11. A. Schurger, J. D. Sitt, S. Dehaene, An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proc Natl Acad Sci U S A* **109**, E2904-2913 (2012).
12. J. Gavenas, U. Rutishauser, A. Schurger, U. Maoz, Slow ramping emerges from spontaneous fluctuations in spiking neural networks. *Nat Commun* **15**, 7285 (2024).
13. M. Murakami, M. I. Vicente, G. M. Costa, Z. F. Mainen, Neural antecedents of self-initiated actions in secondary motor cortex. *Nat Neurosci* **17**, 1574-1582 (2014).
14. C. Moutard, S. Dehaene, R. Malach, Spontaneous Fluctuations and Non-linear Ignitions: Two Dynamic Faces of Cortical Recurrent Loops. *Neuron* **88**, 194-206 (2015).
15. S. Schmidt, H. G. Jo, M. Wittmann, T. Hinterberger, 'Catching the waves' - slow cortical potentials as moderator of voluntary action. *Neurosci Biobehav Rev* **68**, 639-650 (2016).
16. E. Travers, N. Khalighinejad, A. Schurger, P. Haggard, Do readiness potentials happen all the time? *Neuroimage* **206**, 116286 (2020).
17. R. Ratcliff, P. L. Smith, S. D. Brown, G. McKoon, Diffusion Decision Model: Current Issues and History. *Trends Cogn Sci* **20**, 260-281 (2016).
18. P. L. Smith, Stochastic Dynamic Models of Response Time and Accuracy: A Foundational Primer. *J Math Psychol* **44**, 408-463 (2000).
19. J. I. Gold, M. N. Shadlen, The neural basis of decision making. *Annu Rev Neurosci* **30**, 535-574 (2007).
20. M. Usher, J. L. McClelland, The time course of perceptual choice: the leaky, competing accumulator model. *Psychol Rev* **108**, 550-592 (2001).
21. R. Ratcliff, A theory of memory retrieval. *Psychological Review* **85**, 59-108 (1978).
22. B. J. He, J. M. Zempel, A. Z. Snyder, M. E. Raichle, The temporal structures and functional significance of scale-free brain activity. *Neuron* **66**, 353-369 (2010).

23. A. A. Faisal, L. P. Selen, D. M. Wolpert, Noise in the nervous system. *Nat Rev Neurosci* **9**, 292-303 (2008).
24. C. Bogler, B. Grujičić, J.-D. Haynes, Clarifying the nature of stochastic fluctuations and accumulation processes in spontaneous movements. *Frontiers in Psychology* **14**, (2023).
25. M. Murakami, H. Shteingart, Y. Loewenstein, Z. F. Mainen, Distinct Sources of Deterministic and Stochastic Components of Action Timing Decisions in Rodent Frontal Cortex. *Neuron* **94**, 908-919 e907 (2017).
26. N. Khalighinejad *et al.*, A Basal Forebrain-Cingulate Circuit in Macaques Decides It Is Time to Act. *Neuron* **105**, 370-384 e378 (2020).
27. N. Khalighinejad, L. Priestley, S. Jbabdi, M. F. S. Rushworth, Human decisions about when to act originate within a basal forebrain-nigral circuit. *Proc Natl Acad Sci U S A*, (2020).
28. N. Khalighinejad, N. Garrett, L. Priestley, P. Lockwood, M. F. S. Rushworth, A habenula-insular circuit encodes the willingness to act. *Nat Commun* **12**, 6329 (2021).
29. N. Koh *et al.*, Selective direct motor cortical influence during naturalistic climbing in mice. *bioRxiv*, 2023.2006.2018.545509 (2025).
30. A. Mathis *et al.*, DeepLabCut: markerless pose estimation of user-defined body parts with deep learning. *Nat Neurosci* **21**, 1281-1289 (2018).
31. J. J. Jun *et al.*, Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232-236 (2017).
32. N. S. Narayanan, M. Laubach, Top-down control of motor cortex ensembles by dorsomedial prefrontal cortex. *Neuron* **52**, 921-931 (2006).
33. B. Hyland, Neural activity related to reaching and grasping in rostral and caudal regions of rat motor cortex. *Behav Brain Res* **94**, 255-269 (1998).
34. Y. Isomura, R. Harukuni, T. Takekawa, H. Aizawa, T. Fukai, Microcircuitry coordination of cortical motor information in self-initiation of voluntary movements. *Nat Neurosci* **12**, 1586-1593 (2009).
35. R. Romo, W. Schultz, Neuronal activity preceding self-initiated or externally timed arm movements in area 6 of monkey cortex. *Exp Brain Res* **67**, 656-662 (1987).
36. S. U. Okoro *et al.*, Organization of Cortical and Thalamic Input to Inhibitory Neurons in Mouse Motor Cortex. *J Neurosci* **42**, 8095-8112 (2022).
37. A. Klaus, J. Alves da Silva, R. M. Costa, What, If, and When to Move: Basal Ganglia Circuits and Self-Paced Action Initiation. *Annu Rev Neurosci* **42**, 459-483 (2019).
38. W. Schultz, R. Romo, 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 (1992).
39. I. H. Lee, J. A. Assad, Putaminal activity for simple reactions or self-timed movements. *J Neurophysiol* **89**, 2528-2537 (2003).
40. J. Catanese, D. Jaeger, 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 (2021).
41. M. A. Gaffield, B. A. Sauerbrei, J. M. Christie, Cerebellum encodes and influences the initiation, performance, and termination of discontinuous movements in mice. *Elife* **11**, (2022).
42. J. S. Blouin, C. Bard, J. Paillard, Contribution of the cerebellum to self-initiated synchronized movements: a PET study. *Exp Brain Res* **155**, 63-68 (2004).
43. T. Tsujimoto, H. Gemba, K. Sasaki, Effect of cooling the dentate nucleus of the cerebellum on hand movement of the monkey. *Brain Res* **629**, 1-9 (1993).

44. C. Stringer *et al.*, Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
45. H. Zou, T. Hastie, Regularization and Variable Selection Via the Elastic Net. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **67**, 301-320 (2005).
46. G. M. Stine, M. Jazayeri, Control Principles of Neural Dynamics Revealed by the Neurobiology of Timing. *Annu Rev Neurosci* **48**, 43-63 (2025).
47. S. D. Brown, A. Heathcote, The simplest complete model of choice response time: linear ballistic accumulation. *Cogn Psychol* **57**, 153-178 (2008).
48. D. Thura, J. Beauregard-Racine, C. W. Fradet, P. Cisek, Decision making by urgency gating: theory and experimental support. *J Neurophysiol* **108**, 2912-2930 (2012).
49. J. Ditterich, Computational approaches to visual decision making. *Novartis Found Symp* **270**, 114-126; discussion 126-118, 164-119 (2006).
50. B. A. Purcell *et al.*, Neurally constrained modeling of perceptual decision making. *Psychol Rev* **117**, 1113-1143 (2010).
51. M. Milosavljevic, J. Malmaud, A. Huth, C. Koch, A. Rangel, 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 (2023).
52. N. Khalighinejad, S. Manohar, M. Husain, M. F. S. Rushworth, Complementary roles of serotonergic and cholinergic systems in decisions about when to act. *Curr Biol* **32**, 1150-1162 e1157 (2022).
53. K. C. Allaway *et al.*, Cellular birthdate predicts laminar and regional cholinergic projection topography in the forebrain. *Elife* **9**, (2020).
54. V. Minces, L. Pinto, Y. Dan, A. A. Chiba, Cholinergic shaping of neural correlations. *Proc Natl Acad Sci U S A* **114**, 5725-5730 (2017).
55. S. Musall, M. T. Kaufman, A. L. Juavinett, S. Gluf, A. K. Churchland, Single-trial neural dynamics are dominated by richly varied movements. *Nature Neuroscience* **22**, 1677-1686 (2019).
56. A. Khilkevich *et al.*, Brain-wide dynamics linking sensation to action during decision-making. *Nature* **634**, 890-900 (2024).
57. N. A. Steinmetz, P. Zatka-Haas, M. Carandini, K. D. Harris, Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266-273 (2019).
58. L. International Brain *et al.*, A brain-wide map of neural activity during complex behaviour. *Nature* **645**, 177-191 (2025).
59. C. Findling *et al.*, Brain-wide representations of prior information in mouse decision-making. *Nature* **645**, 192-200 (2025).
60. A. Schurger, 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**, (2018).
61. R. Guevara Erra, M. Arbotto, A. Schurger, An integration-to-bound model of decision-making that accounts for the spectral properties of neural data. *Sci Rep* **9**, 8365 (2019).
62. G. Paxinos, K. B. J. Franklin, *Paxinos and Franklin's The mouse brain in stereotaxic coordinates*. Mouse brain in stereotaxic coordinates (Academic Press, an imprint of Elsevier, London, ed. Fifth edition., 2019).
63. Q. Wang *et al.*, The Allen Mouse Brain Common Coordinate Framework: A 3D Reference Atlas. *Cell* **181**, 936-953 e920 (2020).
64. N. Koh *et al.*, Selective direct motor cortical influence during naturalistic climbing in mice. *bioRxiv*, (2025).

65. M. Pachitariu, N. Steinmetz, S. Kadir, M. Carandini, H. Kenneth D., Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv*, 061481 (2016).
66. M. Pachitariu, S. Sridhar, J. Pennington, C. Stringer, Spike sorting with Kilosort4. *Nat Methods* **21**, 914-921 (2024).
67. P. Dayan, L. F. Abbott, *Theoretical neuroscience : computational and mathematical modeling of neural systems*. Computational neuroscience (Massachusetts Institute of Technology Press, Cambridge, Mass, 2001).
68. J. G. Fuglstad, P. Saldanha, J. Paglia, J. R. Whitlock, Histological E-data Registration in rodent Brain Spaces. *Elife* **12**, (2023).
69. G. F. Elsayed, J. P. Cunningham, Structure in neural population recordings: an expected byproduct of simpler phenomena? *Nat Neurosci* **20**, 1310-1318 (2017).
70. Y. Benjamini, Y. Hochberg, 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 (1995).
71. G. E. Uhlenbeck, L. S. Ornstein, On the Theory of the Brownian Motion. *Physical Review* **36**, 823-841 (1930).
72. I. Karatzas, S. E. Shreve, *Brownian motion and stochastic calculus*. Graduate texts in mathematics ; 113 (Springer-Verlag, New York, ed. 2nd, 1991).
73. P. E. Kloeden, E. Platen, *Numerical Solution of Stochastic Differential Equations*. (1992).
74. D. E. Goldberg, *Genetic algorithms in search, optimization, and machine learning*. (Addison-Wesley Pub. Co., Reading, Mass, 1989).
75. A. Pétrowski, S. Ben-Hamida, *Evolutionary algorithms*. Metaheuristics set ; volume 9 (ISTE, London, 2017).
76. R. A. Waltz, J. L. Morales, J. Nocedal, D. Orban, An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Mathematical Programming* **107**, 391-408 (2005).
77. R. H. Byrd, J. C. Gilbert, J. Nocedal, A trust region method based on interior point techniques for nonlinear programming. *Mathematical Programming* **89**, 149-185 (2000).