

Coordinated cross-brain activity during accumulation of sensory evidence and decision commitment

Adrian G. Bondy^{*†1}, Julie A. Charlton^{*†1}, Thomas Zhihao Luo^{*†1,2}, Charles D. Kopec¹, Wynne M. Stagnaro¹, Sarah Jo C. Venditto¹, Laura Lynch¹, Sanjeev Janarthanan¹, Stefan N. Oline¹, Timothy D. Harris^{3,4}, Carlos D. Brody^{†1,2}

¹ Princeton Neuroscience Institute, Princeton University, Princeton NJ, USA

² Howard Hughes Medical Institute, Princeton University, Princeton NJ, USA

³ Janelia Research Campus, Howard Hughes Medical Institute, VA, USA

⁴ Department of Biomedical Engineering, Center for Imaging Science Institute, Kavli Neuroscience Discovery Institute, Johns Hopkins University, MD, USA

* equal contribution, alphabetically ordered

† corresponding authors: {abondy, jc2988, zhihaol, brody} [at] princeton [dot] edu

Abstract

Cognition is produced by the continuous interactions between many regions across the brain, but has typically been studied one brain region at a time. How signals in different regions coordinate to achieve a single coherent action remains unclear. Here, we address this question by characterizing the simultaneous interactions between up to 20 brain regions across the brain (10 targeted regions per hemisphere), of rats performing the “Poisson Clicks” task, a decision-making task that demands the gradual accumulation of momentary evidence. Using 8 Neuropixels probes in each animal, we recorded simultaneously in prefrontal cortex, striatum, motor cortex, hippocampus, amygdala, and thalamus. To assess decision-related interactions between regions, we quantified correlations of each region’s “decision variable”: moment-to-moment co-fluctuations along the axis in neural state space that best predicts the upcoming choice. This revealed a network of strongly correlated brain regions that include the dorsomedial frontal cortex (dmFC), anterior dorsal striatum (ADS), and primary motor cortex (M1), whose decision variables also led the rest of the brain. If coordinated activity within this subnetwork reflects an ongoing evidence accumulation process, these correlations should cease at the time of decision commitment. We therefore compared correlations before versus after “nTc”, a recently reported estimator for the time of internal decision commitment. We found that correlations in the decision variables between different brain regions decayed to near-zero after nTc. Additionally, we found that choice-predictive activity steadily increased over time before nTc, but abruptly stopped growing at nTc, consistent with an evidence accumulation process that has stopped evolving at that time. Assessing nTc from the activity of individual regions revealed that nTc could be reliably detected earlier in M1 than other regions. These results show that evidence accumulation involves coordination within a network of frontal cortical and striatal regions, and suggests that termination of this process may initiate in M1.

Introduction

A fundamental unsolved question in neuroscience is how neural activity is coordinated across the brain to produce cognition. A longstanding obstacle has been the challenge of recording from many individual neurons at once across large swaths of the brain, especially in deep structures and at electrophysiological timescales relevant for many behaviors. Recent technological advances in silicon probe manufacturing (Jun et al., 2017; Steinmetz et al., 2021) have begun to yield the tools needed to overcome these obstacles. A recent set of studies has used these tools to record throughout the brain (although largely not simultaneously), revealing widespread coding for many relevant task variables during perceptual decision making (Steinmetz et al., 2019; International Brain Laboratory et al., 2023; Chen et al., 2024). Perceptual decision making provides a useful behavior for investigating the nature of brain-wide dynamics underlying cognition for three reasons: 1) signatures of the decision making process have long been identified from single-region recordings in frontal and parietal cortices (Shadlen and Newsome, 2001; Thura and Cisek, 2014; Hanks et al., 2015), striatum (Ding and Gold, 2010; Yartsev et al., 2018) and superior colliculus (Horwitz et al., 2004; Jun et al., 2021; Stine et al., 2023), suggesting a broadly distributed computation; 2) precisely timed sensory inputs create a high degree of experimental control; and 3) behavioral choices and neural activity have been parsimoniously modeled using a one-dimensional accumulation-to-bound model (Mazurek et al., 2003; Ratcliff and McKoon, 2008; Steinemann et al., 2022; Luo et al., 2023; DePasquale et al., 2024) which can provide a principled framework for relating neural activity to computation.

The brain-wide studies mentioned above have illuminated brain-wide correlates of the decision making process. However, the moment-to-moment coordination across the brain that gives rise to a decision remains obscure. To address this requires simultaneous sampling of large populations distributed across the brain, as well as methods for understanding these high-dimensional recordings and relating them to computational models of the decision formation process. Here, we report results from experiments that begin to overcome these technical obstacles and shed new light on the nature of coordinated dynamics underlying perceptual decision making.

We combined several critical elements to do this. First, we expanded on our previously reported methods (Luo et al., 2020) for chronic implantation of multiple Neuropixels probes in rats to perform simultaneous recordings from 8 probes per subject, which we targeted to decision-related areas of prefrontal and motor cortex, striatum, thalamus, hippocampus and other forebrain regions. The recordings yielded a median of 2,749 single- and multi-units per session. Recording all brain regions simultaneously means that the representation of the evolving decision could be measured on each single trial under identical conditions across them, providing a fundamentally new window onto decision-related dynamics across the brain. Recordings were performed while 3 rat subjects performed a well-established auditory evidence accumulation task using pseudo-randomly timed pulses of auditory evidence (the “Poisson Clicks” task; Brunton et al., 2013).

Next, we used existing targeted dimensionality reduction techniques to measure brain-wide coordination of the decision formation process. We projected each brain region’s neural activity onto its respective choice-predictive axis in neural state space (Kiani et al., 2014; Kaufman et al., 2015; Chen et al., 2021; Peixoto et al., 2021; Steinemann et al., 2022), sometimes called a “decision variable” or DV (Kiani et al., 2014), and then examined both the structure and timing of across-region DV correlations (Li et al., 2016; Chen et al., 2021, 2024) across all pairs of regions. Importantly, we presented many repeats of each exact stimulus sequence, allowing us to assess the component of DV correlations expected simply from shared coding for choice and stimulus. The remainder reflects correlated fluctuations along the choice dimension beyond what can be explained by these covariates, which we interpret as reflecting the brain’s coordinated, internal decision process.

Using this approach, we could decode the animal's upcoming choice better than chance (and increasingly well over the course of the trial) from each brain region we recorded, although choice prediction accuracy was highest in three interconnected regions in frontal cortex and striatum: anterior dorsal striatum (ADS), primary motor cortex (M1) and dorsomedial frontal cortex (dmFC). DV fluctuations across time were highly correlated within this subset of regions, and anticipated later DV fluctuations in the rest of the brain. We therefore hypothesized that coordinated activity within this fronto-striatal subnetwork is most closely linked to the decision formation process, and it may be the dominant source of choice-related signals found elsewhere.

This hypothesis predicts we should observe dramatic coordinated changes in activity within this subnetwork at the time that decision formation stops and the subject commits to a choice, even if commitment occurs covertly before a movement response. To investigate this, we used a recently developed model (multi-mode drift diffusion model; MMDDM; Luo et al., 2023) that estimates, on each trial, a putative time of decision commitment, and whether or not this time could be detected before the onset of the choice-reporting movement. The estimate is based on the timing of sensory stimuli and on simultaneously recorded firing rates of populations of neurons presumed to be linked to the brain's latent decision process. We refer to it as “nTc”, for “neurally-inferred time of commitment”. This covert state transition happens at highly variable times relative to stimulus start and movement onset. It is interpreted as indicating commitment because, despite the variability in its timing, it coincides with the moment at which sensory evidence ceases having an effect on the subject's decision (Luo et al., 2023).

We found that nTc was accompanied by striking, brain-wide coordinated activity changes consistent with the termination of evidence accumulation. Specifically, we found that although choice prediction accuracy steadily grew before nTc, it abruptly stopped growing around nTc, and in the most choice-predictive areas remained approximately constant thereafter. The time of the plateau was reached in a cascading manner, starting in M1, followed by other fronto-striatal regions and then the rest of the brain. When we estimated nTc separately from each region's neural activity, we observed a similar cascade: M1-derived commitment times significantly preceded those derived from either ADS or dmFC, followed by the rest of the brain. Strikingly, we also found that the correlated DV fluctuations dropped to near-zero around nTc, consistent with their origin being the stochastic dynamics of a shared decision process terminating at a bound.

Taken together, our results provide multiple lines of support for the conclusion that widespread choice-related activity during perceptual decision making is governed by shared dynamics in frontostriatal subcircuits and that dramatic changes in coordinated activity in this subnetwork accompanies the moment of decision commitment, being first apparent in M1. Choice-related activity elsewhere in the brain appears to be driven by this subnetwork, being both weaker and delayed in time. These findings point to the power of combining simultaneous recordings of neural populations throughout the brain and models of their shared latent dynamics to elucidate the coordinated activity underlying cognition.

Results

Simultaneous multi-region neural recordings during auditory evidence accumulation

To understand coordinated brain-wide dynamics underlying perceptual decision making, we first devised novel methods for brain-wide simultaneous recording using chronically-implanted Neuropixels 1.0 probes in freely moving rats. Previous methods (Juavinett et al., 2018; Luo et al., 2020; Steinmetz et al., 2021; Bimbard et al., 2024; Horan et al., 2024) have focused on reusability, surrounding each probe with a bulky enclosure from which it could be removed at the end of the experiment. Here, we sought instead to maximize probe density and flexibility in target selection, forgoing reusability by directly cementing each probe to the skull. We chose 4 penetration sites, targeting a set of cortical and subcortical regions previously shown to be

involved in evidence accumulation, many of which are anatomically interconnected. The penetration sites were mirrored bilaterally, for a total of 8 probe insertions per subject. We report results from 3 subjects, implanted targeting the same set of regions. Results were consistent across the 3 subjects.

The analyses presented here focus on a subset of the regions recorded, chosen for their likely involvement in the decision task the subjects were performing: dorsomedial frontal cortex (dmFC), medial prefrontal cortex (mPFC), primary motor cortex (M1), primary somatosensory cortex (S1), anterior dorsal striatum (ADS), the tail of the striatum (TS), the hippocampus (HPC), basolateral amygdala (BLA), nucleus accumbens (NAc), and medial geniculate body (MGB). A complete list of recorded regions, and description of their atlas equivalents, is in **Table S1**. ADS and dmFC (its strongest cortical input source) have been directly implicated, through recording and causal perturbation, to the decision formation process during the “Poisson Clicks” decision-making task used here (Yartsev et al., 2018; Luo et al., 2023). dmFC is bidirectionally connected (Anastasiades and Carter, 2021) to M1 and mPFC, and ADS receives weak input from them (Hunnicuttt et al., 2016; Luo et al., 2023). MGB is the principal relay for all auditory information reaching the forebrain (Musiek and Baran, 2018), and projects directly to TS (Jiang and Kim, 2018; Chen et al., 2019), which has been shown to be necessary for auditory decision making (Znamenskiy and Zador, 2013; Guo et al., 2018). S1 is reciprocally connected to M1 (Aronoff et al., 2010), projects broadly within striatum (Hunnicuttt et al., 2016), and has recently been established to play a causal role in the formation of perceptual decisions (Buetfering et al., 2022). Decision-related neurons have also been discovered in HPC and NAc during accumulation-of-evidence tasks (Nieh et al., 2021; Luo et al., 2023). BLA contains many auditory-responsive neurons (LeDoux et al., 1991) and projects broadly throughout the striatum (Hunnicuttt et al., 2016).

We used brain clearing and lightsheet imaging to obtain histology for each subject (**Fig. 1D**). We used this volumetric data, in combination with electrophysiological signatures, to confirm targeting and registered the probe tracks to the Princeton RATlas (Dennis et al., 2023) for visualization (**Fig. S1**). The arrangement of the probes on the skulls of the rats was performed using CAD software (**Fig. 1C**), as was the design of a “chassis” that surrounded the probes and held the headstages (**Fig. S2**). Spike sorting was performed using Kilosort 2 (Pachitariu et al., 2023) without manual curation. For inclusion in analysis, we applied criteria based on waveform shape and activity level to exclude non-neural artifacts, fibers of passage and inactive units (see Methods). Multi-units were not excluded. Each recording session (3 rats, 21 sessions, 6-8 sessions per rat) yielded thousands of simultaneously recorded units (median=2,749, 2,385-4,032 range) from tens of brain regions (median = 19, range 13-20) (**Fig. 1E,F**).

Subjects performed the “Poisson Clicks” task (Brunton et al., 2013); **Fig. 1A**). To initiate a trial, the subject had to poke their nose in the center nose port and maintain that posture until a visual “go” cue after 1.5 seconds. With a variable delay after port entry (0.5 - 0.8 s), two sequences of randomly-timed broadband auditory clicks were played from speakers to their left and right until the “go” cue, at which point the rat was free to report its choice by poking its nose in one of two side ports. To receive reward, the subject had to report whether the left or right speaker played the greater number of clicks. After surgery, two of the three rats consistently performed with a low lapse rate (i.e., near perfect at the easiest condition); the third rat’s performance was more variable, and we excluded 4/14 sessions on which it showed a detectable lapse rate and 2/14 sessions on which it completed less than 300 trials (**Fig. 1B**; **Fig. S2F**). We reasoned that the sessions when the rat performs at a minimal lapse rate are those for which we have greatest control over the experimental factors (specifically, the clicks) that influence the rat’s choice, and we therefore focus on those sessions for understanding the cross-brain coordination underlying decision-making. The optimal strategy, which a wide range of evidence supports subjects using in this task, is to gradually accumulate evidence for their choice by integrating the sequence of clicks, a decision process that can be described well using bounded drift diffusion (Brunton et al., 2013; Piet et al., 2018; Boyd-Meredith et al., 2022; Luo et al., 2023; DePasquale et al., 2024; Kopec et al., 2024).

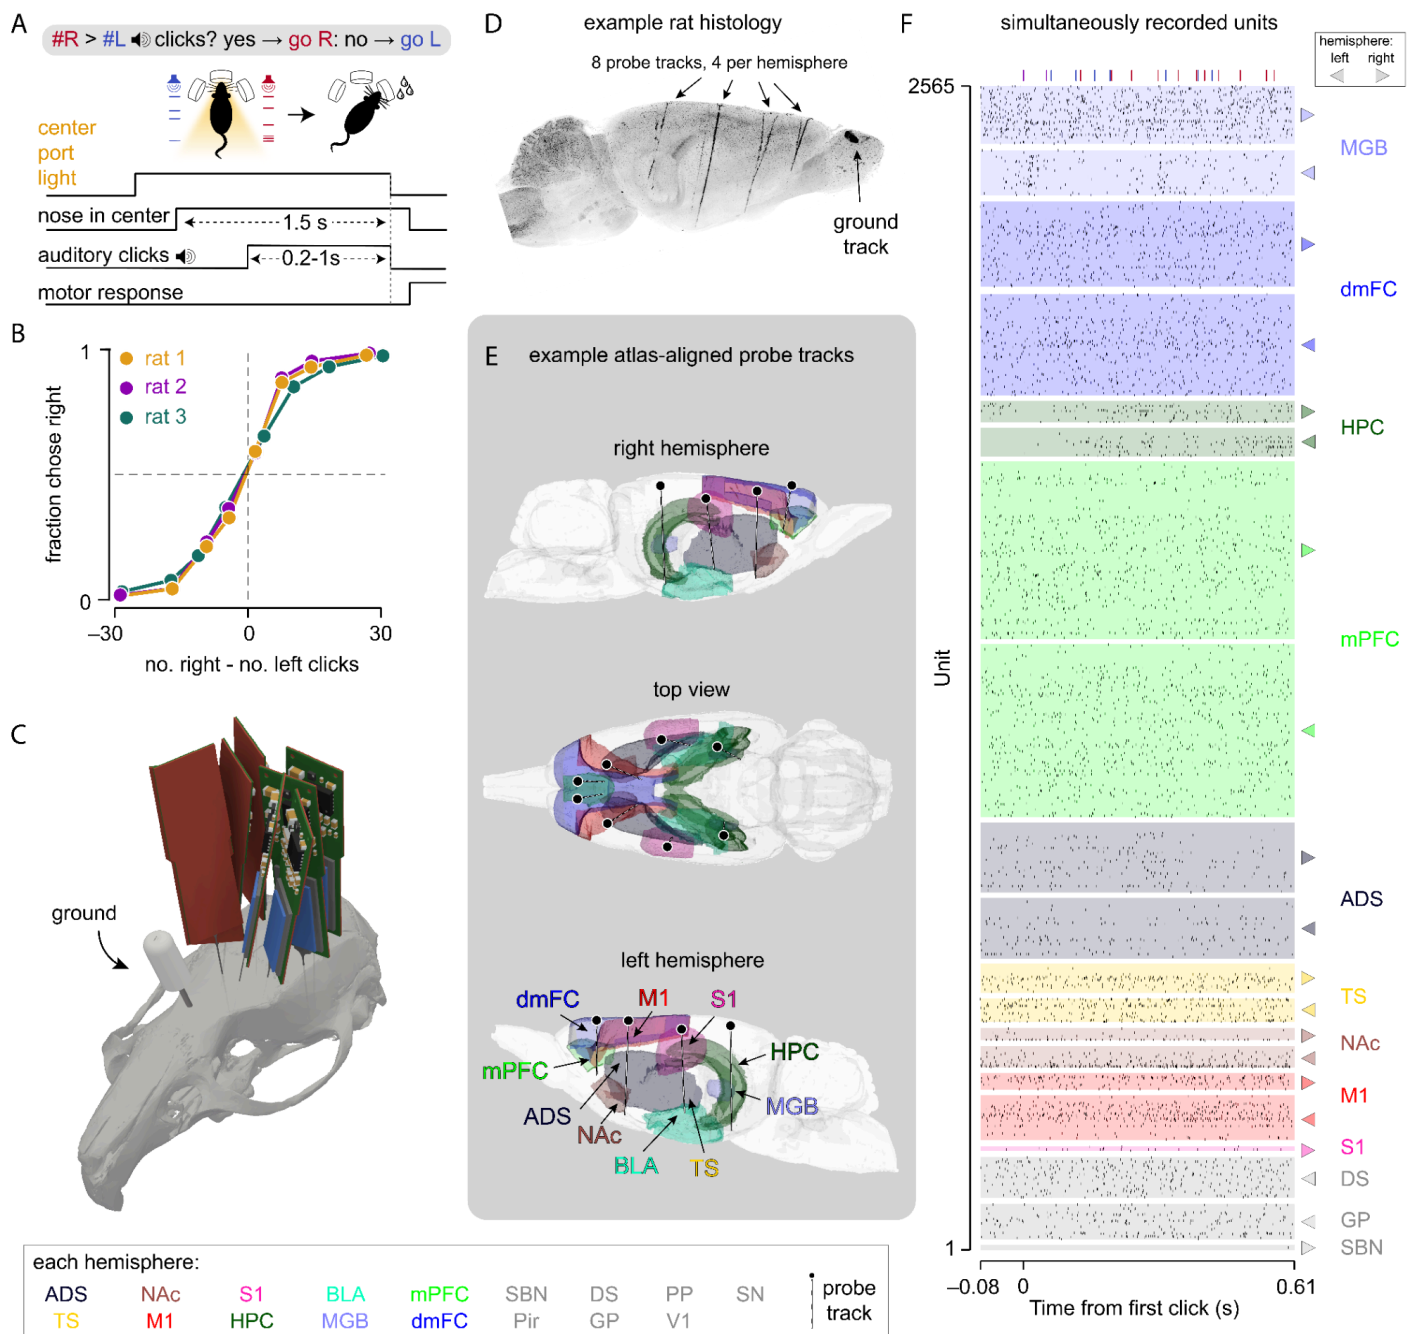


Figure 1. Chronic, simultaneous bilateral recording from 10 targeted regions per hemisphere using 8 Neuropixel probes.

- (A) Task: a rat listens to two simultaneous streams of randomly timed auditory clicks played from loudspeakers on its left and right. At the end of the click train, the rat is rewarded with a drop of water if it turns to the side that played the greater total number of clicks. The end of the click train is always 1.5s from the onset of nose fixation, but the first click varies relative to fixation onset.
- (B) We analyzed recording sessions in which rats had low lapse rates, i.e., were reliably correct on the easy trials when the total click difference is large and when the number of completed trials was at least 300 (3 rats, 21/27 recorded sessions, 6-8 sessions/rat, 10,767 trials).
- (C) 3-dimensional schematic illustrating the position and orientation of the 8 probes on a rat's skull.
- (D) Maximum-intensity projection of full cleared brain volume of rat 2. Black lines are probe tracks labeled by the fluorescent dye cm-Dil.
- (E) 8 probe penetrations for rat 1 registered to a reference 3D rat brain atlas.

(F) Spike raster from a single example trial, showing 2,565 simultaneously recorded putative single units and multiunit clusters, separated by brain region (color) and hemisphere (indicated by arrowhead). Brain regions labeled in light gray indicate areas from which spikes were recorded but no further analysis was performed. Ticks indicate the timing of the left (blue) and right (red) clicks). MGB, medial geniculate body; dmFC, dorsomedial frontal cortex; HPC, hippocampus; mPFC, medial prefrontal cortex; ADS, anterior dorsal striatum; TS, tail of the striatum; NAc, nucleus accumbens; M1, primary motor cortex; S1, primary somatosensory cortex; DS, dorsal subiculum; GP, globus pallidus; SBN, subbrachial nucleus; Pir, piriform cortex; V1, primary visual cortex; SN, substantia nigra; PP, peripeduncular nucleus; BLA, basolateral amygdala.

Coordinated choice-axis fluctuations identify a subnetwork linked to choice formation

To understand how the decision formation process is coordinated across the recorded brain regions on single trials, we used a population decoding analysis to obtain a scalar estimate of the state of the evolving decision from each brain region at each moment in time (illustrated in **Fig. 2A-D** for dmFC on an example session). We initially aligned trials relative to stimulus onset. (Other alignments will be shown in **Fig. 3**.) Following previous approaches (Kiani et al., 2014; Kaufman et al., 2015; Peixoto et al., 2021; Steinemann et al., 2022; Chen et al., 2024), for each brain region on a given recording session, we used logistic regression to define a set of weights across the population (i.e. a vector in neural state space) onto which the projection of the neural state best predicts the upcoming choice (**Fig. 2B**). We found that the vector of weights tended to exhibit rotation across time (**Fig. 2A,B**), and therefore a separate set of weights was fit to each time point. The projection of the neural state onto the weight vector is referred to as the “decision variable” or DV (Kiani et al., 2014) and is monotonically related to the model probability of a rightward choice given the observed neural activity at a given moment in time (see Methods for details). **Fig. 2C** illustrates the evolution of the DV decoded from dmFC on an example session.

Using this approach, we found we could predict the animal’s upcoming choice with growing accuracy across time after stimulus onset in all brain regions we recorded (**Fig. 2E**) (Steinmetz et al., 2019; Steinemann et al., 2022; International Brain Laboratory et al., 2023; Chen et al., 2024). However, we found the highest choice prediction accuracy in M1, dmFC, ADS (peaking at >90%), followed by S1. In fact, decoding accuracy in these regions peaked at a level comparable to accuracy obtained from the entire brain at once, suggesting a high degree of redundancy. The level at which each region’s prediction accuracy peaked depended weakly on population size (**Fig. 2F, Fig. S3B**). However, even when controlling for population size, the peak prediction accuracy was highest in the aforementioned four regions (**Fig. S3D**). We found comparable levels of choice prediction accuracy from the two hemispheres of each region (**Fig. S3A**). These findings are consistent with the widespread representation of the evolving decision observed by others (Steinmetz et al., 2019; Steinemann et al., 2022; International Brain Laboratory et al., 2023; Chen et al., 2024), including the concentration of information about upcoming choice in motor, premotor and somatosensory cortices along with striatum. Notably, choice coding was much lower in cortical area mPFC than in the other frontal cortical regions we recorded (most evident when controlling for population size; **Fig. S3B**) even though it directly borders and receives dense input from dmFC (Ährlund-Richter et al., 2019). Therefore, while choice coding is widespread, it can also display a remarkable degree of anatomical specificity.

Because the entire set of brain regions was recorded simultaneously, we were able to compare the evolution of the decision variables decoded from each region under identical conditions on single trials. **Fig. 2D** illustrates the highly correlated nature of the DVs obtained independently from simultaneously recorded populations in dmFC and ADS on one example session, which was typical for many pairs of regions. However, a great deal of that correlation is due to shared coding for choice, as illustrated by the presence of two discrete choice-conditioned clusters at later time points in the trial. Indeed, DV correlations can, in principle, reflect both shared coding for covariates known by the experimenter (choice, stimulus, etc.) as well as covariability

reflecting shared internal dynamics along each region's choice dimension. We reasoned that the latter is more directly revealing about the brain's internal process of deliberation.

To remove the component of DV correlation driven by stimulus and choice, we used a shuffling procedure (**Fig. S4**) that took advantage of a subset of 5 sessions ("frozen noise" sessions) for which, instead of generating the stimuli independently on each trial, a fixed set of 54 random noise seeds were repeatedly reused across trials (i.e., 540 trials would contain an average of 10 identical repeats of each of the frozen noise click trains). We then randomly shuffled trials, while preserving choice and stimulus seed. Doing this many times provides a distribution of DV correlations for which choice and stimulus are identical but simultaneity is abolished, which can be subtracted from the unshuffled value to correct for the influence of these variables. For some regions (like MGB) the shuffle correction removed virtually all DV correlations with the rest of the brain regions (**Fig. S4D**), implying MGB-derived DVs are nearly conditionally independent from those of the other regions. For other regions, the shuffle correction left intact a significant fraction of DV correlation, suggesting unexplained moment-to-moment co-fluctuations.

In **Fig. 2G**, we show the full matrix of shuffle-corrected DV correlations for all recorded pairs of brain regions, across the set of 5 frozen noise sessions. (That is, these are the residual correlations along the choice axis after accounting for choice and stimulus-induced effects.) The diagonal of this matrix shows the DV correlations between hemispheres for each region. This matrix contains three distinct groups of regions: a first cluster (ADS, M1 and dmFC) that showed the highest correlations ($\rho \approx 0.2$), a second cluster of three regions (S1, mPFC, and TS) showing a slightly lower level of DV correlations ($\rho \approx 0.15$) and the rest (NAc, MGB, HPC, BLA) showing very small, near-zero correlations. We applied multidimensional scaling (MDS), a non-linear dimensionality reduction technique that preserves pairwise distances (**Fig. 2H**), to project the brain regions onto a single dimension using the DV correlations as a similarity metric. ADS, M1 and dmFC cluster along one end of the MDS dimension. Interestingly, the order of regions this yields closely matched the ordering of brain regions given by peak choice prediction accuracy (**Fig. 2E**), showing that those regions that contain the most information about upcoming choice also tend to have the most correlated choice-independent fluctuations along the choice axis.

In **Fig. S3E**, we show the matrix of shuffle-corrected DV correlations with the left and right hemisphere of each region treated as a separate population. The left and right hemisphere of each region show remarkably similar patterns of correlation with the rest of the brain, suggesting that the process of decision formation during evidence accumulation is inter-hemispherically coordinated to a great degree. We also found (**Fig. S5**) that when we projected neural activity onto dimensions predictive of other task-related variables (previous choice, previous rewarded side, and momentary evidence), the pattern of correlations between each population's projection ("decoded variable") was weaker than the pattern observed when projecting onto the choice-predictive dimension, indicating that our results do not reflect generic correlations in activity between regions, but are rather specific to the neural dimension being considered.

We next examined the temporal structure of DV correlations by quantifying them across a large range of time lags, i.e. calculating cross-correlograms. The DV cross-correlograms tended to peak near zero and decay within several hundred milliseconds (**Fig. 2I**), consistent with shared moment-by-moment fluctuations in the decision variable. Next, we quantified the peak of the cross-correlograms, reasoning that, if a subset of brain regions is most closely linked to the decision formation process, then DVs derived from those regions should temporally lead those observed elsewhere. Indeed, the peak lags were highly structured, with some areas leading (like M1, **Fig. 2I**) and others lagging. The full set of pairwise lead-lag relationships are shown in the matrix in **Fig. 2J**, and the average lag of each region relative to the rest of the population is shown in **Fig. 2K**. Supporting the conclusion that ADS, M1 and dmFC are closely linked to decision formation, they were among the group (along with S1 and TS) whose DV fluctuations tended to lead those in other regions. Indeed, the average lags were highly correlated with the manifold coordinate (**Fig. 2H**) obtained from the zero-time-lag DV correlation matrix.

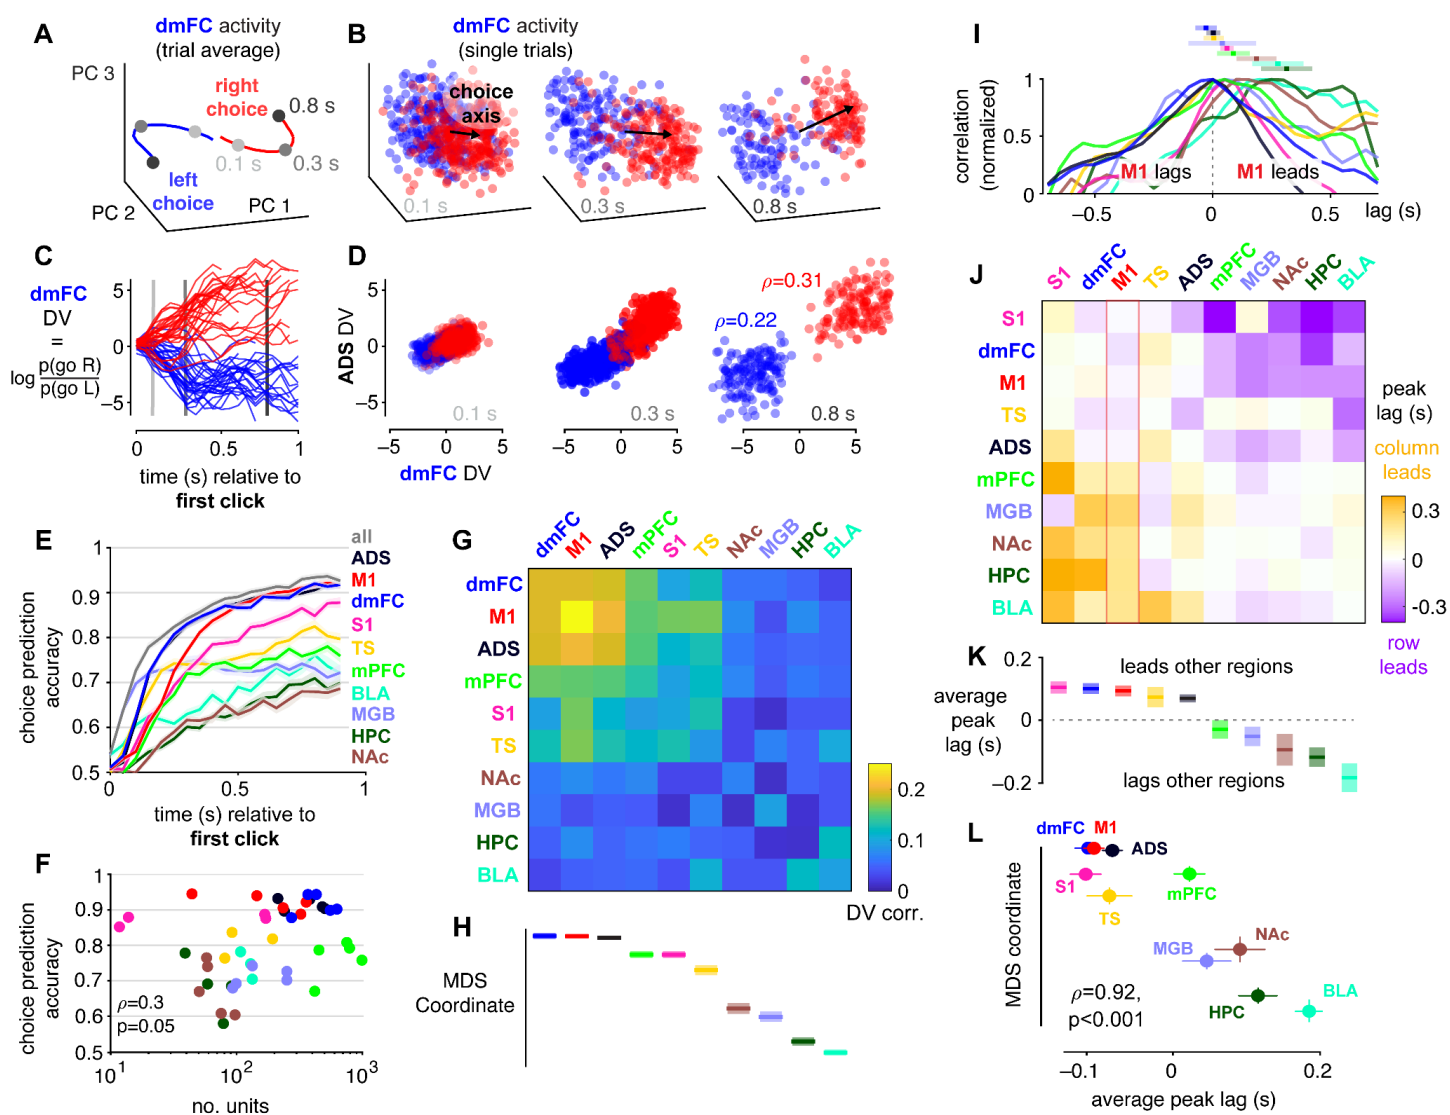


Figure 2. Coordinated choice-axis fluctuations identify a subnetwork linked to perceptual choice formation.

- (A) Schematic illustration of population activity in dmFC on an example session, shown as the average trajectory of neural activity along the top 3 PCs for left and right choices. Mean activity across all trials was first subtracted. Gray dots indicate time in trial relative to the first click.
- (B) Same as in (A) except instead of the average trajectory across time, we show three example time points, with a single dot for each trial in the session. An arrow indicates the vector of weights determined by a logistic regression model for choice. Its direction is referred to as the “choice axis”.
- (C) The dot product of neural activity with the weight vector illustrated in (B) yields the log odds of the model probability of a rightward choice given neural activity. This quantity is referred to as the decision variable (DV) and is shown here for the same set of data in (B).
- (D) Scatter plots illustrating correlation of DVs shown in (C) with those obtained from simultaneous population recordings in ADS at three example time points.
- (E) Prediction accuracy of a logistic regression model of choice given population neural activity for 10 brain regions separately, and for all regions together (gray), as a function of time relative to the first click. Accuracy is assessed using the class-balanced accuracy under 10-fold cross-validation (see Methods). Values indicate the average (\pm 1 s.e.) across three rats and the five frozen noise sessions total.
- (F) Scatter plot showing a weak relationship between the size of the population used to predict choice and choice prediction accuracy. Each dot indicates a single region on a single recording session, with colors as in (E). Here, choice prediction accuracy is assessed as the average accuracy in the later part of the stimulus (0.7s and beyond).

- (G) DV correlations (Pearson's ρ) for all pairs of recorded regions. Diagonal entries in the matrix indicate the correlation between DVs obtained from each hemisphere of the same brain region. A subset of brain regions concentrated on the frontal cortex and anterior striatum show the highest DV correlations. We used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus and choice, described in detail in the Methods and illustrated in **Fig. S4**. Only time points before a neurally-inferred time of decision commitment (see Methods) on each trial are included.
 - (H) DV correlation matrix in (G) was projected down onto a one-dimensional manifold using classical multidimensional scaling (MDS). The position of each region on that manifold (± 1 s.e.) is plotted here.
 - (I) In (D) and (G), DV correlations at zero time lag are shown. Here we plot the full cross-correlogram at a range of time lags between M1 and all other regions. Values are normalized so that 1 defines the value at the peak. The cross-correlogram peaks (± 1 s.e.) are shown as error bars above the main plot.
 - (J) Matrix showing the peak lags in the cross-correlograms for all pairs of regions. The matrix is anti-symmetric along the diagonal, with entry i,j indicating the time lag by which the region of the i th column leads the region of the j th row. The column highlighted in red corresponds to the peaks of the cross-correlograms shown in (I).
 - (K) Plot showing the average (± 1 s.e.) cross-correlogram peak lag of each region relative to the rest of the population. These values correspond to the averages of each column in (J) excluding diagonal entries.
 - (L) Scatter plot illustrating the high degree of correlation between the average peak lag (from K) and the one-dimensional MDS projection of the DV correlations in (G). That is, the subset of regions showing the highest DV correlations tend also to be those whose DVs lead the rest of the brain.
-

Cross-brain activity changes at a neurally-inferred time of decision commitment

The strong choice-related co-fluctuations we observed in ADS, M1 and dmFC are consistent with their joint activity reflecting the stochastic dynamics of a shared latent decision process. If this is the case, we would expect to see dramatic, coordinated changes in their activity after the decision process terminates. To investigate this, we used a recently reported method (Luo et al., 2023) to estimate, based on the timing of sensory stimuli and firing rates of populations of simultaneously recorded neurons, a putative time of decision commitment, “nTc,” for “neurally-inferred time of commitment.” The estimate is based on a model (“MMDDM”, for “multi-mode” or “minimally-modified” drift diffusion model) in which neuronal firing rates are weighted functions of a one-dimensional latent variable, which evolves according to diffusion-to-bound dynamics driven by the sensory stimuli. The model allows for an abrupt change in the weights at the time of reaching the bound. Using spiking data from many simultaneously recorded neurons to estimate the time of this abrupt change allows for high precision in the estimated time. We refer the reader to the Supplementary Information or (Luo et al., 2023) for further details of the MMDDM model. After fitting the model's parameters, the posterior probability of reaching the bound at each time point can be estimated for each trial from the recorded spiking data and the sensory stimulus timing. We considered a trial to have an nTc if this posterior probability reached 0.95 before movement onset (i.e., before the rat withdrew its nose from the center port), and did not decrease thereafter. The time at which the probability reached 0.95 was then defined as the nTc for that trial. By aligning trials to the estimated nTc, Luo et al., 2023 used behavioral analysis to show that, as predicted for a time of decision commitment, sensory evidence presented before nTc affects the subject's choices, but sensory evidence presented after nTc does not.

We first confirmed these results hold for the current data set. The greater the number of simultaneously recorded neurons with choice-predictive activity, the higher the sensitivity to detecting nTc, and the higher the precision of the estimate. The MMDDM model was therefore initially fit, and nTc estimated, from the joint data from dmFC, mPFC, NAc, ADS, and M1. We found that nTc was estimated to occur before movement onset on 54% of trials; subsequent analyses focus on these trials. Trial-averaged responses, conditioned on choice, were reasonably well fit for many neurons (**Fig. 3B-D, Fig. S7E**), as was the behavior of the rats (**Fig. S4D**).

As previously reported (Luo et al., 2023), nTc was broadly distributed with respect to both stimulus onset and movement onset (top panels in **Fig. 3G,H**) and is substantially more precise than behavioral estimates, which use only choice and auditory clicks (**Fig. 3E**). Furthermore, we confirmed in our dataset that nTc coincides with a drop in the behaviorally-estimated weight of auditory clicks on the animal's decision (**Fig. 3F**).

Fig. 3G-I show how choice prediction accuracy evolved over time across the recorded regions for trials aligned to stimulus onset (**Fig. 3G**, as previously shown in **Fig. 2E**), aligned to movement onset (**Fig. 3H**), and aligned to nTc (**Fig. 3I**). The nTc alignment (**Fig. 3I**), enabled by large-scale simultaneous recordings and MMDDM, revealed a striking difference to the stimulus or movement onset alignments: choice prediction accuracy abruptly stopped growing around nTc among the most choice predictive regions (M1, ADS, S1, and dmFC), even though nTc could occur many hundreds of milliseconds before the choice-reporting movement. This is consistent with nTc being a signal that indicates the end of an evolving neural decision process that is coordinated across the brain. It furthermore suggests that the smoothly growing ramps of **Fig. 3G,H** are likely composed of averages over trials that each plateau around the time of the trial's nTc.

A bilinear fit quantitatively identified a point of slope change in the choice prediction accuracy for each brain region (**Fig. 3J,K**). This occurred within tens of milliseconds of nTc for the most choice-predictive regions, and earliest within M1. This point of slope change was tens or even hundreds of milliseconds after nTc for the less choice predictive brain regions (**Fig. 3L**). This result was not guaranteed by the procedure used to estimate nTc. For example, even though mPFC was used in the estimation of nTc, its point of slope change occurred well after nTc. To rule out the possibility that this saturation is a spurious result of the nTc finding a moment on each trial when the sensory evidence itself saturates, we plotted the prediction accuracy of an ideal observer of the clicks aligned to nTc. No abrupt saturation was observed, confirming that the neural effect reflects the termination of an internal decision process not entirely predictable from the stimulus (**Fig. S7G-I**).

We also observed that the change in angle between the choice axis determined at adjacent time points relative to nTc (i.e. its angular velocity) ramped down in anticipation of nTc and then stabilized, consistent with slowing of choice-related rotational dynamics across the brain accompanying decision commitment (**Fig. 3M; Fig. S6**). This was a novel and unanticipated result further confirms that there are profound changes in neural activity around the time of nTc.

We finally sought to revisit the strong DV correlations we identified among a subgroup of regions (**Fig. 2G**). If these correlations can be explained by the stochastic dynamics of a common decision process, they should diminish after the animal has committed to a decision, since at this point the variability of the latent process is abolished. To test this, we computed the DV correlation matrices of **Fig. 2G** separately for the periods of choice formation before and after a trial's nTc. We indeed observed a prominent drop in DV correlation among nearly all pairs of regions (**Fig. 3N**). When we computed the temporal profile of the DV correlations relative to nTc, we saw that the drop was not constant across time, but rather was concentrated around nTc. After nTc, the DV correlations fell to chance levels within 100s of milliseconds (**Fig. 3O**).

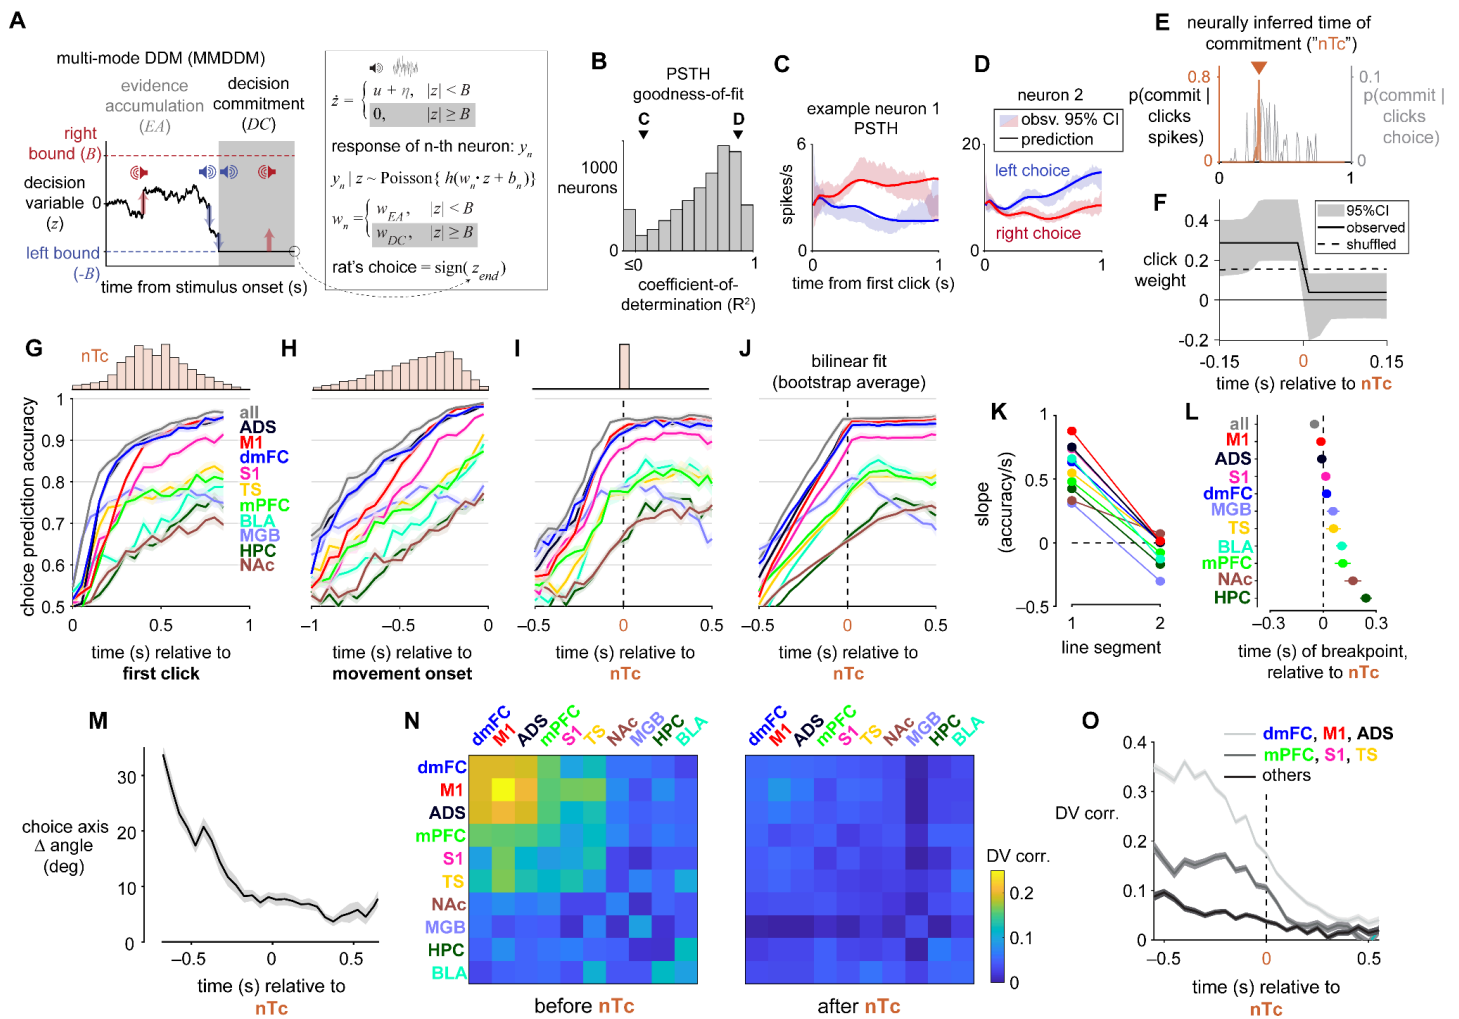


Figure 3. Cross-brain activity changes at a neurally-inferred time of decision commitment.

- (A) A model fit to simultaneously recorded spike trains and behavioral choice. In the multi-mode drift-diffusion model (MMDDM), momentary evidence (u) and noise (η) are accumulated over time in the decision variable (z) until z reaches either $-B$ or $+B$. Abruptly at this moment, the animal commits to a decision: z becomes fixed and insensitive to evidence. Here, also at this moment, each neuron's encoding weight (w), mapping z to the neuron's predicted Poisson firing rate y , abruptly changes from w_{EA} to w_{DC} . The mapping from z to y passes through the softplus nonlinearity h and depends on baseline b .
- (B) Goodness-of-fit of the temporal profile of the trial-averages. The 10th, 50th, and 90th percentiles are 0.11, 0.65, and 0.89.
- (C) Trial-averaged responses of an example neuron with a R^2 of 0.07
- (D) Example neuron with a R^2 of 0.87.
- (E) The inferred time of commitment is far more precise when inferred from neural activity ("nTc") than when inferred from behavior without neural activity.
- (F) Supporting the interpretation of nTc as decision commitment, when trials are aligned to nTc, we find that auditory clicks cease to impact the animal's decision after the nTc. Trials for which the estimated time of commitment occurred at least 0.15s before stimulus offset and 0.15s or more after stimulus onset were included for this analysis (2,544 trials).
- (G-I) Prediction accuracy of a logistic regression model of choice given population neural activity for 10 brain regions separately, and for all regions together (gray), as a function of time relative to the first click (panel G: first click, H: movement onset, I: nTc). When aligned to nTc but not the other task events, the strength of choice coding across the brain sharply rises and then abruptly plateaus, consistent with a coordinated decision termination process. Accuracy is assessed using the class-balanced accuracy under 10-fold cross-validation (see Methods). Only trials for which nTc could be inferred were included. Values indicate the average (\pm 1 s.e.) across the same set of sessions in **Fig. 2**.

Histograms at top indicate the distribution of nTc relative to the respective trial event. Among the 54% of trials (5,818/10,767 across 21 sessions/3 rats) for which the commitment times could be inferred, the nTc occur at variable times relative to first click and movement onset. The fraction of nTc declines for longer stimulus durations because the stimulus duration on each trial is randomly drawn from 0.2-1.0s.

- (J) We found we could capture the choice prediction accuracy curves aligned to nTc (panel I) quite well with a bilinear fit, consistent with the observation of a ramp followed by plateau.
- (K) Slopes of the two line segments of the bilinear fits to the data in (J). The first segment has a positive slope for all regions (an upward ramp) and slopes for the second segment are clustered around zero (plateau).
- (L) Comparison of the breakpoint in the bilinear fits to the data in (J), i.e. the time at which the slope transitions. The slope transition occurs strikingly close to the time of commitment for M1, ADS, S1 and dmFC and follows in other brain regions by tens or hundreds of milliseconds.
- (M) The change in choice axis angle between neighboring timepoints (choice axis Δ angle) dropped in anticipation of nTc and then plateaued, suggesting a stabilizing of the choice axis at the time of decision commitment. Choice axis Δ angle is shown relative to an estimated noise floor, obtained from bootstrapping trials (see Methods). Average of all regions shown here, as results were comparable across them.
- (N) DV correlations (Pearson's ρ) for all pairs of recorded regions, as in **Fig. 2G**, but here separated into two sets of timepoints: those before and after nTc on each trial. A drop to near-zero across virtually all pairs is evident. As in **Fig. 2G**, diagonal entries in the matrix indicate the correlation between DVs obtained from each hemisphere of the same brain region, and we used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus and choice, described in detail in the Methods and illustrated in **Fig. S4**.
- (O) Here we plot DV correlation (computed as in panel N) separately at individual time points relative to nTc, to better illustrate the time course of the reduction in magnitude. Pairs are broken down into three groups based on their pre-nTc strength of DV correlations. A drop is apparent amongst all three groups but is, of course, more prominent amongst dmFC, M1 and ADS which show the strongest pre-nTc DV correlations. The time course shows a stable period of high correlation, a drop over several hundred milliseconds around nTc and then a stabilization near zero at around 300 ms after nTc.

Decision commitment is first detected in M1

To investigate whether there are systematic timing differences across regions in terms of when nTc is detected, we inferred nTc using neural activity from individual regions (region-specific nTc) rather than using the frontal regions (ADS, M1, NAc, dmFC, mPFC) together as before (frontal-nTc). The simultaneous recordings allowed us to compare trial by trial differences between regions. To do this, we fit MMDDM to all the neurons recorded in the forebrain in each session. We excluded neurons in the corpus callosum or neurons that could not be reliably assigned to a brain region. Then, we inferred region-specific nTc using the spikes recorded from for each of the ten regions analyzed (ADS, BLA, MGB, M1, NAc, S1, TS, dmFC, mPFC) and but using the same MMDDM parameters and click times for all region-specific nTc's. For each pair of regions analyzed, we considered only the sets of trials on which nTc could be detected using the spike trains from both brain regions, using the spikes from each brain region alone. The difference between the nTc's of each region was thus well defined in each of the trials considered. We found that nTc inferred from M1 was significantly earlier than nTc inferred from any of the other brain regions from which we recorded (**Fig. 4A-B**; **Fig. S8A**). We then repeated this analysis, but using the parameters of the MMDDM optimized using neural activity from only the frontal regions ADS, M1, NAc, dmFC, mPFC (as in **Fig. 3**), and performing pairwise comparisons among these five frontal regions. We replicated the same results, showing that within this subset, nTc was again estimated earliest in M1 (**Fig. S8E**). We found this result surprising because the mean peak lag of the DV decoded from M1 was similar to that of S1, dmFC, and TS (**Fig. 2I-K**). Because the DV cross-correlograms were computed across the entire trial, it is possible they obscured a leading role for M1 that is confined to the period around decision commitment.

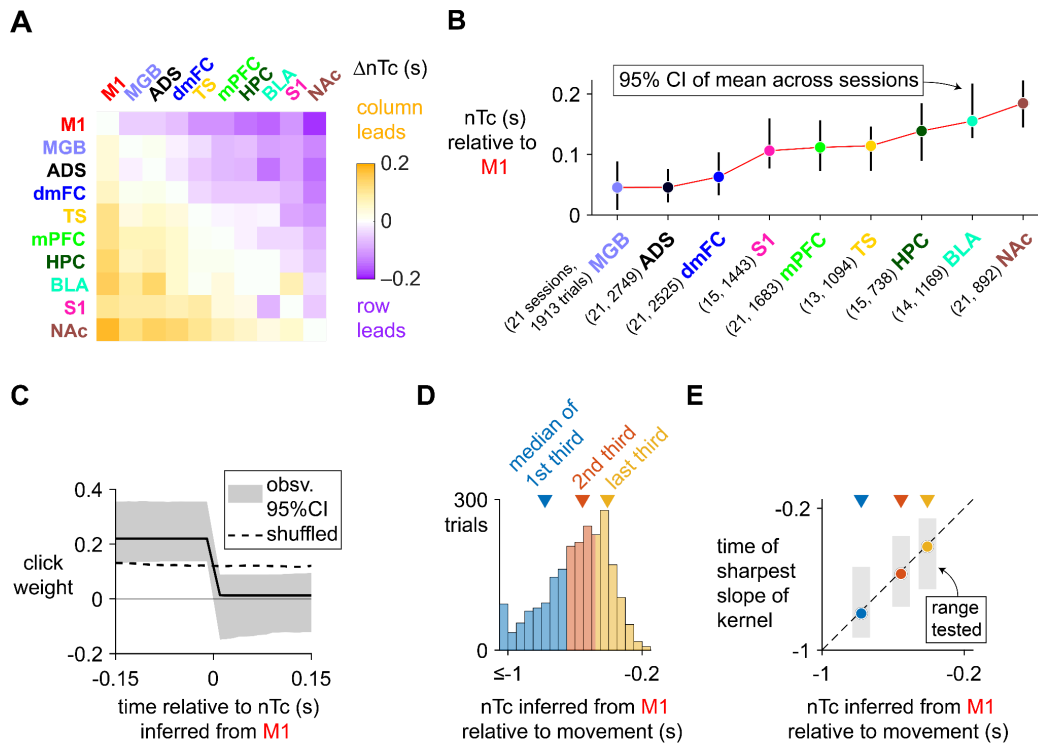


Figure 4. Decision commitment is first detected in M1.

- (A) Comparison between nTc inferred separately from each brain region. For each region A minus region B pairwise comparison, we included only the subset of trials on which nTc could be detected both using spikes from region A alone and using spikes from region B alone.
- (B) Comparison between nTc inferred using M1 and the nTc inferred using each other brain region. Each comparison uses the subset of trials on which commitment could be inferred separately using the spikes M1 and also the other region.
- (C) When trials are aligned to nTc inferred using only spikes from M1, auditory clicks cease to impact the animal's decision after the nTc. Trials for which the estimated time of commitment occurred at least 0.15s before stimulus offset and 0.15s or more after stimulus onset were included for this analysis (2085 trials).
- (D) Trials on which nTc could be inferred from M1 that precede stimulus offset by at least 0.15s were divided into three groups by their time relative to movement onset, and separate psychophysical kernels were estimated for each group. Trials with relative nTc < -1 were included in the leftmost bin.
- (E) The time of the sharpest decrease in the psychophysical kernel for each group of trials is more closely aligned to the nTc than to a particular time relative to movement (which would imply that the markers fall along a horizontal line rather along the diagonal). For each group, the time of sharpest slope of the kernel was estimated relative to the nTc and then added to the median nTc-minus-movement of that group. The range of time intervals tested was, in seconds and relative to the nTc, for each group [-0.15, 0.25], [-0.2, 0.2], and [-0.25, 0.15]. The number of trials used for each group was 541, 614, and 684.

nTc is estimated from neural activity. If it corresponds to decision commitment in the behavior, it must satisfy a key prediction: clicks played after this moment will not influence the animal's upcoming decision. We found that the sensitivity to the clicks abruptly decreased when we aligned trials to the nTc estimated from M1, confirming that this estimate (like frontal-nTc in **Fig. 3F**) satisfies this key prediction (**Fig. 4C**). This was true for M1 across a number of different time intervals (**Fig. S8C**). This drop was not found for all regions when trials were aligned to the nTc estimated from each brain region individually, such as MGB (**Fig. S8D**).

Given the known role of M1 in motor planning, we assessed whether the time of the abrupt decay in the behavioral sensitivity to auditory clicks (the timing of the downward step in **Fig. 4C**) was more closely aligned with the onset of the rat's movement away from the fixation port rather than nTc. To do this, we selected trials on which nTc could be inferred from M1 and also preceded stimulus offset by at least 150 ms. These trials were then divided into three groups based on the time between nTc and movement onset (**Fig. 4D**). We then inferred a separate psychophysical kernel for each group. If the time of the abrupt decay in sensitivity to clicks were more closely aligned to movement onset, then the time of the steepest decrease of the three estimated kernels would be aligned to one specific time relative to movement. Instead, we found that steepest-decrease times of the three kernels were more closely aligned to the nTc than to a particular time relative to movement (**Fig. 4E**; the markers lie on a diagonal rather than a horizontal line). This result supports the interpretation that the nTc estimated from M1 is an internal signal of commitment. Together with the comparison of the nTc's estimated from different regions, the results indicate that decision commitment may be reflected in a cascading series of brain-wide events, and points to M1 as a target for further investigation of the mechanisms of decision commitment.

Discussion

Numerous results from neural recordings and perturbation experiments point to a highly distributed neural basis of perceptual decision making. However, the nature of brain-wide coordination during decision making – or any other cognitive process – has remained unclear, largely due to the difficulty of large-scale simultaneous recordings. Here we addressed this problem through simultaneous, chronic recordings of thousands of neurons distributed throughout anatomically interconnected decision-related areas across the brains of rats performing the “Poisson Clicks” task, an auditory accumulation-of-evidence task. We found that moment-to-moment fluctuations in the decision variables (DVs) decoded from each brain area were most highly correlated between anterior dorsal striatum (ADS) and two reciprocally-connected cortical regions that project to it (M1 and dmFC). Furthermore, DV fluctuations in these regions tended to lead those of other regions by tens or hundreds of milliseconds. These results support the conclusion that the decision formation process is coordinated in a fronto-corticostriatal subnetwork, and that decision-related signals in other regions we recorded are downstream from this. We used a recently developed method for inferring from neural activity the time of decision commitment on individual trials (“neurally-inferred time of commitment”; nTc). This inferred internal state change coincides with striking changes in decision-related neural activity in this same fronto-corticostriatal network, including an abrupt plateau in choice prediction accuracy and a sharp reduction in correlated DV fluctuations, consistent with coordinated activity in these regions reflecting an evidence accumulation process that terminates at nTc. When we aligned trials to stimulus or movement onset, the abrupt change in choice prediction accuracy was obscured. Instead, we found steadily growing choice prediction accuracy (**Fig. 3G,H**), consistent with previous findings (Chen et al., 2024). Finally, when we estimated nTc separately for different brain regions, we found that those derived from M1 occur the earliest, suggesting a possibly critical role for M1 in decision commitment.

A cortico-striatal subnetwork for perceptual decision making

Several aspects of our results point to M1, ADS, dmFC (and arguably S1) comprising a critical subnetwork for decision formation, at least among the set of forebrain regions sampled in this study. First, we found that they contain the most information (overall and on a per-neuron basis) about the upcoming choice. Second, moment-to-moment fluctuations in the decoded decision variable from these regions are the most highly correlated (with the exception of S1), and they tended to precede the DV fluctuations decoded from other brain regions. Choice prediction accuracy in these regions also abruptly plateaued within 50 ms of nTc,

and before the rest of the brain regions we recorded. Lastly, nTc inferred from M1 alone significantly precede those inferred from other regions, followed within 50 ms by ADS and dmFC, and within 100 ms by S1.

Recent studies employing causal inactivation (Yartsev et al., 2018; Luo et al., 2023) have established that two of these regions (dmFC and ADS) play a necessary role in the evidence accumulation process in the context of the task used in this study. S1 and M1 have traditionally been associated with functions upstream and downstream of the decision making process, respectively, and thus this aspect of our results was unexpected to us. Nonetheless, our findings are consistent with several other recent studies of these areas. Experiments using somatosensory decision making tasks in mice (Yang et al., 2015; Kwon et al., 2016; Buettering et al., 2022) have recently identified a role for S1 not just in the representation of somatosensory information but also in the decision making process that depends on those inputs, although it is unclear how far these results can be extended to decision making tasks that involve other sensory modalities. However, two recent studies (Steinmetz et al., 2019; International Brain Laboratory et al., 2023) reporting brain-wide correlates of task encoding during visually-guided decision making in mice found that areas homologous to S1, M1, M2 (secondary motor cortex) and striatum were among those regions containing most information about upcoming choice. Additionally, a recent study (Peixoto et al., 2021) of macaque M1 and PMd (dorsal premotor cortex) found that DVs estimated from the joint recordings displayed statistical features resembling bounded accumulation and found that upcoming choice could be decoded with approximately equal accuracy in both regions. Therefore, a growing body of evidence, including the results we present here, supports the conclusion that M1 and S1 are directly involved in decision formation, not simply its sensory inputs or motor outputs.

Inferring the brain's decision variable

In the present study, we inferred the state of a one-dimensional decision process from neural activity, by projecting neural activity in a given brain region onto the most choice-predictive axis in neural state space. This approach yields a scalar quantity that numerous past studies have empirically demonstrated can provide a useful proxy of the brain's time-varying decision state (Kiani et al., 2014; Peixoto et al., 2021; Steinemann et al., 2022). Separately from this, we used a method (MMDDM) to estimate, from knowledge of the sensory stimuli and population firing rates on each trial, the presence and timing of the putative decision commitment.

The present study suggests that these two approaches give closely related results. In particular, we found that nTc coincides with several striking changes in cross-brain activity projected onto the choice-predictive axis: 1) information along this axis about the upcoming choice abruptly saturates at nTc; 2) inter-regional correlations along this axis dropped to near-chance levels after nTc; and 3) the rotation of the choice axis in neural state space stabilized at nTc. The first two results are directly predicted if the DVs estimated using the first method reflect bounded drift diffusion, and assuming MMDDM provides an accurate estimate of the time that the bound is reached. The stabilization of the choice axis at nTc was not directly predicted, but further supports the conclusion that MMDDM's estimate of decision commitment reflects profound changes in neural state that are apparent in dynamics along the choice axis. MMDDM is a model with highly simplified decision dynamics, describing complex neural population activity in terms of a single one-dimensional latent variable. Nevertheless, regardless of how well or poorly it captures all aspects of the neural activity, it appears to be a good estimator of the internal decision commitment time, both in terms of behavior (Luo et al., 2023) and **Fig. 3F**) and in terms of the evolution of neural activity related to the brain's internal decision process (**Fig. 3I-O**).

Comparison to past studies using DV correlations

Previous studies have used DV correlations to assess cross-brain coordination during a somatosensory decision making task in mice that do not require evidence accumulation. Using simultaneous recording from multiple silicon probes, these studies found significant choice-related co-fluctuations between hemispheres of

ALM (Li et al., 2016; Chen et al., 2021)—a critical region for decision formation in that task (Guo et al., 2014; Li et al., 2015)—and between ALM and downstream regions (Chen et al., 2024). When they examined the strength of these correlations across time, they claimed to observe changes around the time of movement onset.

Our results are consistent with some of these findings but differ markedly from others. First, we also observed coordinated decision-related activity of similar magnitude between subregions of frontal cortex and those with which it is interconnected (**Fig. 2G**). Additionally, we also found that this coordinated activity was strongly interhemispheric (**Fig. S3C**). However, without a trial-by-trial estimate of the putative decision commitment time (nT_c), Chen et al., 2024 could not observe the remarkable drop in correlations after nT_c that is documented here (**Fig. 3N**). Indeed, whether or not mouse subjects make covert decision commitments before movement onset in their task is as yet unclear.

We also note that Chen et al., 2024 did not clearly demonstrate a fall in correlations as time during the trial unfolded, even after movement onset. Although correlations along a choice axis direction computed before movement onset weakened after movement onset, this appeared to be due to a change in choice axis direction at movement onset: after movement onset, correlations along the choice axis direction computed for that time period remained strong. Here we made no assumption about temporal constancy of the choice decoding direction, allowing it to change freely across time points. Therefore, a reduction in DV correlations could be readily distinguished from a rotation in neural dynamics.

Another difference across the two studies is that Chen et al., 2024 were primarily focused on ALM as the central node in the decision making process, and assessed coordination exclusively between ALM and downstream regions whose decision-related activity was thought to be driven by ALM. Regions other than ALM were not always sampled simultaneously. The goal of our study was instead to simultaneously sample as broad a range of putatively decision-related regions across the brain as possible, under the assumption that coordination amongst many regions may underlie the process of evidence accumulation in our task. Indeed, the results we report here support the conclusion that no single brain region is the lone driver of evidence accumulation.

Mechanisms of commitment

The neural correlates of decision commitment that we report here are distinct from those highlighted in another recent study which reports a burst of activity in the superior colliculus (SC), a region of the midbrain (Stine et al., 2023). It is quite possible that bursting activity in SC could contribute to the decision termination in our task. Indeed, midbrain bursting activity might give rise to the commitment-associated changes we detect in the forebrain, in particular in frontal cortex and striatum, given that the midbrain strongly modulates frontal regions during decision-making (Inagaki et al., 2022). Future work can directly address this possibility through simultaneous recording of frontal cortex, striatum and SC, paired with perturbations. Such an experiment may clarify the contribution of SC to not only decision commitment but also evidence accumulation (Jun et al., 2021). Future work addressing the mechanisms of decision termination may wish to take into account a difference in task paradigm that might contribute to the findings. In the task we use, the onset of the animal's motor response must be delayed until a predetermined time, even if it has already committed to a choice ("fixed-duration paradigm") (Brunton et al., 2013), whereas in (Stine et al., 2023), the animal is trained to couple the termination of their decision with the initiation of a motor response ("reaction-time paradigm") (Roitman and Shadlen, 2002). Bursting activity in the superior colliculus (Basso and Wurtz, 1998) and other regions of the midbrain (Inagaki et al., 2022), as well as abrupt population-wide changes in frontal cortex (Chen et al., 2024), correlate with motor initiation even in tasks that do not involve the accumulation of sensory evidence. Contrasting neural correlates in fixed-duration and reaction-time paradigms might help clarify the contribution of the various neural signals that correlate with decision termination.

Even within a single behavioral paradigm, there are likely multiple distinct mechanisms that contribute to termination. Rats' behavior in a reaction-time task is far better captured by a model which includes two distinct latent processes, one that is entirely stimulus-independent, and one that has both stimulus-independent and -dependent components (Hernández-Navarro et al., 2021). The method used here to infer decision commitment (nTc) depends on the auditory stimulus, and in the future may be extended to incorporate a component that is strictly stimulus independent (e.g., adding latent dimensions in MMDDM that are stimulus-independent). This extension may help address whether there are distinct neural signals associated with stimulus-independent and -dependent processes of decision termination, and whether this dichotomy might be useful for capturing not only behavior but also neural activity.

Limitations

Although we provide compelling evidence that coordinated activity in a corticostriatal subnetwork drives the accumulation of evidence, this conclusion certainly provides an incomplete anatomical picture. We did not record from a number of regions with a possible role in evidence accumulation, notably including posterior parietal cortex, the frontal orienting fields (FOF), the superior colliculus (SC), and other midbrain and hindbrain regions known to contain information about upcoming choice (Steinmetz et al., 2019; International Brain Laboratory et al., 2023; Chen et al., 2024). The primary goal of bilateral, simultaneous sampling of all targeted regions imposed geometrical constraints on our implantation strategy that prevented targeting of these regions. We expect that future miniaturization of silicon probes as well as improvements in surgical technique and implant design will alleviate these constraints.

In the present study, we treat evidence accumulation as a covert process that is only expressed behaviorally after the “go” cue when the subject initiates an orienting action to indicate its choice. However, it is likely that subjects make untrained movements throughout the evidence accumulation period that correlate both with neuronal activity and the upcoming choice, as has been observed previously (Musall et al., 2019; Stringer et al., 2019; Wang et al., 2023). One interpretation of the choice-related co-fluctuations we observe is that these reflect shared encoding of these untrained movements. Similarly, it may be the case that postural changes play a role in the neural changes we observe around the neurally-inferred time of commitment. The neural and behavioral recordings performed in this study were accompanied by video from which animal pose can be recovered, and so addressing these possibilities is feasible with future analysis.

Future directions

The work presented here is part of an exciting shift in the field of systems neuroscience toward cross-brain sampling of neural activity during complex behavior, enabled by powerful new electrophysiological tools. This broad sampling has recently revealed decision related activity is much more widely distributed than previously thought (Steinmetz et al., 2019; International Brain Laboratory et al., 2023; Chen et al., 2024). Our recordings have replicated these findings, but also demonstrate the additional power of recording many brain regions at once. We found that, although widespread choice-related activity is found throughout the brain, at the level of single-trials, a small, highly correlated subgroup of regions leads the decision formation process. This suggests that many regions that show choice related activity are not causal to the accumulation computation, but are simply downstream of it. This speculation warrants further experiments to determine what role, if any, these other regions play.

Further, the highly correlated decision-related activity observed among this subgroup on single trials – and the ability of a relatively simple model of the decision making process to explain several aspects of their joint activity – naturally leads to the interpretation that their joint activity implements the accumulation of evidence computation. This contrasts with an alternative view that posits a unique, specialized sub-computation for each node of the decision making circuit, e.g. FOF implementing decision categorization

and the posterior parietal cortex (PPC) accumulating evidence (Hanks et al., 2015). Of course, these views are not mutually exclusive – certain computations may be distributed across multiple regions while others may be local to only one.

How exactly to test the hypothesis that the accumulation computation is jointly implemented by multiple regions is yet unclear. Dissecting a broadly distributed computation will likely require the joint recording and perturbation of the regions thought to compose the causal circuit. Rodent models will allow for the specific inactivation of the axon terminals that interconnect the putative causal circuit (Gupta et al., 2024). The interpretation of such experiments will require mechanistic models of multiple interacting brain regions, constrained by the anatomy. Models of this type have been developed in the context of spatial navigation and movement control (Lo and Wang, 2006; Eliasmith et al., 2012; O'Reilly et al., 2015; Mejias et al., 2016; Grossberg, 2019), but are relatively rare in the field of perceptual decision making (Lo and Wang, 2006; Bogacz and Gurney, 2007). A network model that can accurately predict both behavior and neural data would constitute a quantitative hypothesis of the interregional dynamics that mediate evidence accumulation and decision commitment, generate experimentally testable predictions, and perhaps provide a more accurate estimate of decision commitment than the approach used here.

The main findings reported here could not have been observed without simultaneous recordings. In the near future, we expect that further advances in neural probe technology will allow for even denser simultaneous sampling of neural activity across the brain using many, reusably-implanted probes, to become routine. The experimental, conceptual and modeling advances outlined above together have the potential to rapidly accelerate our understanding of how the brain, as a whole, gives rise to complex cognitive behavior.

Acknowledgments

We thank Ben Lankow and Mark Goldman for suggestions and comments. We also thank Jessica Morrison, Klaus Osorio, Jovanna Teran, Andres Bustos and Emily Valance for technical assistance. We thank Grace Barnett and Jamus MacGuire for veterinary advice. We are grateful to all members of the Brody lab for their support, collegiality and feedback. This work was supported by the Howard Hughes Medical Institute.

Methods

Subjects

Three adult male Long-Evans rats (Hilltop) were used for the experiments presented in this study. All procedures were approved by the Princeton University Institutional Animal Care and Use Committee and were carried out in accordance with National Institutes of Health standards. Rats were pair-housed in Technoplast cages until their implantation surgery and kept in a 12 hour reversed light-dark cycle. All training and testing procedures were performed during the rat's dark cycle. Rats had restricted access to water such that the water consumed daily was at least 3% of their body mass.

Behavioral task

Rats performed the behavioral task in custom-made training enclosures (Island Motion, NY) within sound- and light-attenuated chambers (IAC Acoustics, Naperville, IL). Each enclosure consisted of three straight walls and one curved wall in which three nose ports were embedded (one in the center and one on each side). Each nose port contained one light-emitting diode (LED) as well as an infrared (IR) beam to detect the entrance of the rat's nose into the port. A loudspeaker was mounted above each of the side ports and was used to present auditory stimuli. Each of the side ports also contained a small metal tube that delivered water reward, with the amount of water controlled by valve opening time.

Rats performed an auditory discrimination task in which optimal performance required the gradual accumulation of auditory clicks (Brunton, 2013). At the start of each trial, rats inserted their nose in the central port and maintained this placement for 1.5 s ("fixation period"). After a variable delay of 0.5-1.3 s, two trains of randomly timed auditory clicks were presented simultaneously, one from the left and one from the right speaker. Regardless of onset time, the click trains terminated at the end of the fixation period, resulting in stimuli whose duration varied from 0.2-1 s. The train of clicks from each speaker was generated pseudo-randomly by an underlying Poisson process, with different mean rates for each side. The combined mean click rate was fixed at 40 Hz, and trial difficulty was manipulated by varying the ratio of the generative click rate between the two sides. The generative click rate ratio varied from 39:1 clicks/s (easiest) to 20:20 (most difficult). At the end of the fixation period, rats could orient towards the nose port on the side where more clicks were played and obtain a water reward.

Psychometric functions were computed by dividing trials into eight similarly sized groups according to the total difference in the right and left clicks, and for each group, computing the fraction of trials ending in a right choice. The confidence interval of the fraction of right response was computed using the Clopper-Pearson method. Recording sessions were excluded if the rat completed <300 trials or they show detectable lapse rate in the psychometric function. We excluded 6 sessions from C211. Additionally, 3 sessions were excluded from A327 due to an error in specifying the electrodes to be recorded.

Implant

The probe holder and headstage holder were designed using Autodesk Inventor Professional 2024 software. The probe holders were 3D printed using the Form 3 SLA printer (Formlabs) in Black V4 resin (Formlabs; RS-F2-GPBK-04) and the headstage holder was printed in Tough 1500 (Formlabs; RS-F2-TO15-01). CAD files for these components can be found at https://github.com/Brody-Lab/uberphys_paper/tree/main/CAD_files.

After printing the parts, they were visually inspected and sanded to ensure proper mating. The probe holder parts were secured using two 4 mm M1.2 screws (McMaster; 96817A746). The headstage holder was assembled using 3mm M1 screws (McMaster; 96817A704) and headstages were secured to the headstage holder using 4 mm M1.2 screws (McMaster; 96817A746). To secure the probes in the probe holders, each probe holder was placed in a stereotaxic cannula holder (Kopf, Tujunga, CA, USA; Model 1766-AP Cannula Holder) which was held in place by a vise. The probe was then placed on the probe holder and aligned to the axis of the cannula holder. A small amount of thick-viscosity cyanoacrylate glue (Mercury Adhesives) was applied to the edges of the probe holder using a small wooden dowel.

Surgery

Surgery was performed using similar techniques to those reported previously (Luo et al., 2020). All surgical procedures were performed under isoflurane anesthesia (1.5-2%) using standard stereotaxic technique. Rats were given an intraperitoneal (IP) injection of ketamine (60 mg/kg), ketofen (5 mg/kg) and Ethiqx XR (0.65 mg/kg) to assist induction and provide analgesia. To ensure proper hydration throughout the surgery, rats were given 3 mL saline subcutaneously after induction and every 3 hours afterward.

The dorsal skull was exposed by making an incision along the rostral/caudal orientation along the top of the head. The skull surface between the lambdoid sutures and 20 mm anterior of the frontonasal suture was cleaned and scrubbed. The temporalis muscle was detached from the lateral ridge and retracted to gain access to the tail of the striatum. The sites of nine craniotomies, one for the ground cannula (Protech International, 22G/5mm), and eight for the Neuropixels 1.0 probes were marked with a sterile pen. The craniotomy for the ground had a diameter of approximately two millimeters, and each craniotomy for a Neuropixels 1.0 probe had a diameter of 1 mm. The 3D profile of each craniotomy intended for a Neuropixels 1.0 probes had a conical shape to minimize the amount of dura exposed (to maximize the stability of the chronic recording) while maximizing the range of angles through which the dura can be accessed, thereby facilitating the subsequent durotomy. After completing the nine craniotomies, they were covered with Gelfoam (Pfizer), and then dental cement (C&B Metabond Quick Adhesive Cement System) was applied to the skull surface. Durotomies were made using a 27G needle and fine forceps. After the ground cannula was lowered, the craniotomy was sealed with a silicone adhesive (KWIK-SIL, World Precision Instrument). The cannula was adhered to the skull through a dental composite (Absolute Dentin, Parkell).

Neuropixels probes were stereotaxically inserted into the brain using a motorized micromanipulator (Narshige, MDS-1) at a speed of ~5 $\mu\text{m/s}$. To couple a probe to the motorized manipulator, we designed a custom 3D printed attachment that had two components: one that is glued permanently to a probe, and another that is removed after the probe is anchored to the skull. Each craniotomy in which a probe was inserted was sealed using a silicone gel (Dowsil 3-4680), applied using a micropipette. The Neuropixels probes are bonded to the skull and existing fixtures using dental composite.

After all eight probes had been inserted, the silver wire shorting the ground and reference pad of each probe were twisted together and soldered onto the ground cannula. To reinforce the attachment between the probes and the skull, liquid dental acrylic was applied to the skull surface. To shield the probes and to mount the headstages, a chassis (**Fig. S2**) was attached to the fixtures using dental composite.

Site #	1	2	3	4
<i>Craniotomy coordinates, mm relative to Bregma (AP, ML)</i>	+4.2, 1.0	+1.9, 3.0	[-2.0 : -2.1], [5.0 : 5.2]	[-5.7 : -6.0], 3.7
<i>Insertion angle in sagittal plane (deg)</i>	0	0	[0 : 5]	[0 : 5]
<i>Insertion angle in coronal plane (deg)</i>	-10	-10	5	0
<i>Insertion Depth (mm)</i>	[3.9 : 4.9]	[7.4 : 7.9]	[6.8 : 7.6]	[7.4 : 7.9]
<i>Regions targeted</i>	dmFC, mPFC	M1, ADS, NAc	S1, TS, GP, BLA, Pir	V1, HPC, DS, MGB, SBN, HPC, SN

Table 1. Recording targets

Four insertion targets were used, bilaterally in each subject. Probes were sometimes angled in either or both the sagittal and coronal plane, both to accommodate multiple probes on the subject's head and to target specific combinations of brain regions. A positive angle in the sagittal plane indicates that the probe tip was more anterior than the probe base. A positive angle in the coronal plane indicates the probe tip was more lateral than the probe base. To avoid blood vessels and collisions between probes, some variability in coordinates across subjects was required. In these cases, the range of coordinates is indicated.

Electrophysiological recording

Neural activity was recorded using chronically implanted Neuropixels 1.0 probes that were permanently affixed to the skull using custom-designed 3D-printed probe holders described above. We used acquisition hardware from NI (a PXIe-1071 chassis) in conjunction with SpikeGLX software (<https://github.com/billkarsh/SpikeGLX>) to acquire the data. The reference selected for each probe was a silver wire shorted to the ground wire and penetrating the olfactory bulb. The amplifier gain used during recording was 500. Spikes were sorted offline using Kilosort2 (Pachitariu et al., 2023), using default parameters and without manual curation. In each of three animals, probes bilaterally targeted one of four locations described in detail in Table 1.

Histology

Rats were transcardially perfused with 10% formalin under anesthesia with 0.4 mL ketamine (100 mg/ml) and 0.2 mL xylazine (100 mg/ml) IP. Brains were cleared using modified uDisco, volumetrically imaged using lightsheet microscopy, and aligned to the Princeton RATlas (Dennis et al., 2023). These methods are described in detail elsewhere (Dennis et al., 2023).

While the Princeton RATlas provides a useful tool for visualizing brains in a common coordinate space, adding well validated region annotations to the RATlas remains a work in progress. Therefore, to assign recorded units to brain regions, we used the following procedure. Using the BigDataViewer (Pietzsch et al., 2015) Plugin for Fiji (Schindelin et al., 2012), we dynamically resliced the lightsheet volumes to obtain virtual slices that best visualized each individual probe track. These virtual slices were then segmented into brain regions by visual comparison to the Paxinos and Watson rat atlas. The recording sites on each probe were then assigned to a position within the virtual slice, by converting from image pixels to physical distance given the insertion depth of each probe. We found we could more accurately estimate the insertion depth of each probe from electrophysiology rather than by using the nominal insertion depth recorded during surgery. The

electrophysiological estimate was determined by the most superficial channel on each probe at which multi-unit activity could no longer be clearly observed.

Neuronal selection

Units were only included for analysis if they exceeded predefined thresholds for a number of quality metrics based on waveform shape. These thresholds are defined in Table 2 below, and were designed to exclude units that were a) not of biological origin, i.e. noise artifacts; and 2) not of somatic origin, since axonal spikes could be generated by fibers of passage. These criteria are highly similar to those recently proposed by another group to exclude artifacts and non-somatic units from silicon probe recordings (Fabre et al., 2023). In addition, for decoding analyses and estimation of decision variables, units were only included if they fired at least 1 spike on at least half of trials (i.e. whose “presence ratio” exceeded 0.5). Approximately 65% of units found by Kilosort2 were included given these criteria. Note that no criteria were applied to exclude multi-units.

<i>Quality metric</i>	<i>Description</i>	<i>Allowed values</i>
Spatial spread	Spatial decay constant of an exponential fit to the waveform energy as a function of distance to the peak site	<150 μm
Peak width	Width of main deflection at half height	<1 ms
Peak-trough width	Time from trough to peak	<1 ms
Upward-going spike	Has an upward-going peak deflection	FALSE
uVpp	Peak-to-peak voltage	>50 μV

Table 2. Waveform-shape-based unit inclusion criteria.

Where applicable, metrics are defined for the average waveform on the main channel (i.e. the channel for which the unit had the largest peak-to-peak voltage).

Neural decoding of choice and other behavioral variables

We used logistic regression to decode choice from population neuronal activity. The model probability of a rightward choice on a given trial $p(R)$ was $p(R) = f(X_t \beta_t + \alpha_t)$, where f is the logistic function, X_t is the vector of neuronal firing rates at time t , β_t is the vector of weights applied to each neuron, and α_t is a model constant unique to each timepoint. Neuronal firing rates were estimated by convolving spike times with a 50ms Gaussian smoothing filter and sampled at 50ms intervals. We used 10-fold stratified cross-validation to assess model performance as well as to identify the optimal L1 regularization hyperparameter. The “decision variable” at a given time point DV_t was defined as linear predictor $X_t \beta_t + \alpha_t$, which is equivalent to the log-odds of $p(R)$ (i.e. $\log(p(R)/p(L))$). For decoding other binary variables (such as previous choice) the identical procedure was used. Where choice prediction accuracy is reported, this is calculated as the “balanced accuracy” (the

average of the accuracy in predicting the two choices). This removes upward bias associated with an uneven distribution of choices.

To decode the instantaneous stimulus S_t from population neuronal activity, a similar procedure with linear regression was used. We defined S_t as the difference in right versus left clicks within a 100ms window around time t . Then, under the model, $E(S_t) = X_{t+\tau} \beta_t + \alpha_t$, where τ was a temporal lag to account for a delay in stimulus processing. τ was set to 50ms. To fit the decoding models, we relied on the Glmnet (Friedman et al., 2010) package in MATLAB R2024a (Mathworks, MA, USA).

The choice axis refers to the normalized vector of decoding weights $\beta_t / |\beta_t|$. To calculate its rotation across time, we quantified the angle between neighboring timepoints t and $t+1$ as $\cos^{-1}([\beta_t \cdot \beta_{t+1} / (|\beta_t| |\beta_{t+1}|)])$. Because this angle captures real rotation as well as noise in our estimates of the choice axes, we sought to estimate the noise floor and subtract it from our measurements. The noise floor was estimated as the average of a bootstrap distribution of the angle between choice axes obtained from trial resampling. This was performed separately for each time point in the trial.

Calculating DV Correlations

DV correlations were calculated as the Pearson correlation across trials and timepoints between the decision variables estimated from different brain regions. To estimate the component of DV correlations due to shared coding for stimulus and choice, we used a trial-shuffling procedure as follows (and as illustrated in **Fig. S4**). For one of the two sets of DVs (i.e. for those corresponding to one of the pair of regions) we randomly permuted trial identities while preserving stimulus seed and the subject's choice. The stimulus seed controlled the pseudorandom number generator that determined the stimulus duration and the exact sequence of left and right clicks. For most sessions, stimulus seeds were unique for each trial. For a subset of five sessions across the three animals ("frozen noise" sessions) a small number (54) of unique stimulus seeds were repeatedly presented to the animal throughout the session, spanning the typical distribution of trial difficulties. We performed this permutation 50 times, each time recalculating the DV correlations, to generate a shuffled distribution of DV correlations for which simultaneity between the two sets of DVs had been abolished but for which stimulus seed and choice was identical. Then the mean of this shuffle distribution was subtracted from the unshuffled DV correlation. Unless otherwise stated, this shuffle-corrected value is what is reported throughout the paper.

A similar procedure was used for shuffle-correcting correlations between the decoded variables reported in **Fig. S5** (previous choice, previous reward, previous rewarded side and momentary evidence). For decoded binary variables (previous choice, previous reward and previous rewarded side), the shuffling procedure also preserved the identity of the respective binary variable. For example, correlations along the axis predicting previous choice are shown after subtracting the component predicted by shared coding for previous choice.

Multi-mode drift-diffusion model (MMDDM)

In the Supplementary Materials we provide a detailed description of the multi-mode drift-diffusion model (MMDDM), adapting from (Luo et al., 2023). Briefly, MMDDM consists of a dynamic model governing the time evolution of a 1-dimensional latent variable and measurement models specifying the conditional distributions of the observations (spike counts and behavioral choice) given the value of the latent variable.

In the dynamic model, when the value of the latent variable (z) is not at either absorbing bound $-B$ or B , its value at each time step t depends on the momentary input (u), which is corrupted by multiplicative noise of variance σ_s^2 , and additive noise ϵ :

$$z(t + 1) = z(t) + u(t) + \epsilon$$

When z reaches either bound, it remains at the bound. The dynamic model has three free parameters, the bound height $B=(10, 20)$, variance of the multiplicative noise $\sigma_s^2=(0.1, 20)$ and the mean of the initial state of the latent variable $\mu_0=(-5, 5)$. We chose to fit the input-related noise (rather than other sources of noise) because previous work suggests it to be the dominant source of noise in our task (Brunton et al., 2013). The additive noise on each time step is an i.i.d Gaussian with variance Δt , which is the time step $\Delta t=0.01$ s.

The measurement model of the behavioral choice depends on only the sign of the latent variable z on the last time step of each trial (positive indicating rightward). There is no free parameter in the measurement model of the choice.

The measurement model of the spike count of neuron n at time step t is given by

$$y(n, t) | z(t) \sim \text{Poisson}(\lambda(n, t) * \Delta t)$$

The firing rate λ , which has the unit of spikes/s, is

$$\lambda(n, t) = h\{w(n) * z(t) + b(n, t)\}$$

where h is the softplus activation function to approximate the f-I curve, w the neuron's encoding weight of z , and b a time-varying baseline input that is independent of z , the left or right clicks, or the animal choice (Supplementary Materials). The baseline b accounts for time-varying influences of neural activity aligned to two events in the trial, stimulus onset and movement onset (Park et al., 2014), and slow drifts over minutes across a session (Rabinowitz et al., 2015). The measurement model of each neuron has 19 free parameters.

The parameters of the dynamic model and the measurement models are learned simultaneously. The gradient of the log-likelihood of the model has a closed-form expression and is used to optimize the parameters using the L-BFGS algorithm. Only responsive (> 2 hz) and choice-selective neurons (among the subset selected under the criteria described in "Neuronal selection") were included in the fitting of MMDDM. Choice selectivity was computed using an ideal observer analysis, the receiver operating characteristic (ROC), categorizing between left and right choices using the spike counts during the first 0.5 s from the stimulus onset on each trial (excluding trials ending before 0.5 after stimulus onset). Choice selectivity is defined as $|\text{area under the ROC} - 0.5| > 2^{-5}$ (median choice selectivity among responsive neurons=0.0315).

A separate instance of MMDDM was fit to each of 21 recording sessions (7 sessions performed by A324, 6 by A327, and 8 by C211). For the analysis in Figure 3 and Figure S8D-F, MMDDM was fit to only the frontal corticostriatal brain regions M1, ADS, dmFC, mPFC, and NAc as in the original deployment of the model in (Luo et al., 2023). For the analysis in Figure 4, MMDDM was fit to all neurons except in the corpus callosum or not assigned to a brain region were excluded from model fits.

Neurally inferred time of commitment (nTc)

The time when decision commitment occurred is selected to be when the posterior probability of the latent variable at either the left or the right bound, given the learned MMDDM parameters, click times, spike trains, is greater than 0.95 and remains above 0.95 for the remainder of the trial. The posterior distribution at each time step was computed using the click times and spikes both before and after that time step. For Figure 3F-O and

Figure S7D-F, all the neurons used to fit the model (neurons in M1, ADS, dmFC, mPFC, and NAc) were used in computing the posterior distributions. For Figure 4 and Figure S7A-C, while neurons from across the forebrain were used in optimizing the parameters θ of MMDDM, we computed separate posterior probability distributions using the spikes from neurons from separate brain regions.

$$\begin{aligned} p(z\{t\} | \theta, \text{clicks}, M1 \text{ spikes}) &= p(z\{t\} | \theta, \text{clicks}) p(M1 \text{ spikes} | z\{t\}) / p(M1 \text{ spikes}) \\ p(z\{t\} | \theta, \text{clicks}, ADS \text{ spikes}) &= p(z\{t\} | \theta, \text{clicks}) p(ADS \text{ spikes} | z\{t\}) / p(ADS \text{ spikes}) \\ &\dots \end{aligned}$$

The prior distribution $p(z\{t\} | \theta, \text{clicks})$ is identical across the different posterior distributions computed for separate brain regions.

Psychophysical kernel aligned to nTc

The psychophysical kernel quantifies the weight of the auditory clicks at each moment relative to the neurally estimated time of decision commitment (nTc) on the animal's upcoming choice. The kernel was estimated using a logistic model that regresses the animal's choice against a constant term, the generative (i.e., experimentally specified) difference between the right and left (L) clicks on each trial ($\lambda\Delta t$), and the deviation of the actual difference between (R) and left (L) click times from the generative

$$(\text{choice} | R, L, \lambda\Delta t) \sim \text{Bernoulli}(\text{logistic}\{x\})$$

$$x = w_b + w_{\lambda\Delta t} \lambda\Delta t + \sum_t w(t) [R(t) - L(t) - \lambda\Delta t]$$

The psychophysical kernel is specified by the time-varying weight of of the deviation between the generative and actual difference between right and left clicks is specified as also a logistic function

$$w(t) = a + b / (1 + \exp(-k[t - t_0]))$$

where the parameter a is the weight at a time well before nTc, the parameter b is the weight at a time well after the nTc, k the sharpness of the change in the kernel, and t_0 the point in the kernel with the steep slope. The model has six parameters w_b , $w_{\lambda\Delta t}$, a , b , k , and t_0 . The latter four parameters a , b , k , and t_0 specify the psychophysical kernel. We used the same time step duration $\Delta t = 0.01$ s as in the multi-mode drift-diffusion model (MMDDM). Trials for which clicks occurred at least 0.15s before and also 0.15s after the nTc were included for analysis. The shuffling procedure involves randomly permuting the time of the clicks on each trial and does not change the behavioral choice, the time of decision commitment, the generative right-minus-left click input, or the number of trials.

Statistical tests

Unless otherwise stated, error bars and p-values are calculated non-parametrically, using bootstrap trial resampling.

Author Contributions

Role	AGB	JAC	TZL	CDK	WS	SJV	LL	SJ	SO	TH	CB
Conceptualization	X	X	X							X	X
Surgery	X	X	X	X		X					
Implant Design and Construction	X	X	X								
Chassis Design (Fig. S2)			X								
Recording	X	X	X	X		X					
Spike Sorting	X										
Data Analysis and Figure Preparation (Fig. 1)	X	X	X		X						
Data Analysis and Figure Preparation (Fig. 2; Fig. S3-6)	X										
Data Analysis and Figure Preparation (Fig. 3)	X		X								
Data Analysis and Figure Preparation (Fig. 4; Fig. S7-8)			X								
Brain Clearing							X				
Lightsheet Imaging and Atlas Registration								X			
Unit Brain Region Assignment	X	X									
Probe Track Reconstruction and Fig. S1					X						
Funding										X	X
Manuscript Preparation	X	X	X	X	X						X
Manuscript Comment and Revisions	X	X	X	X	X	X					X
Miscellaneous				X	X	X			X		
Project Supervision											X

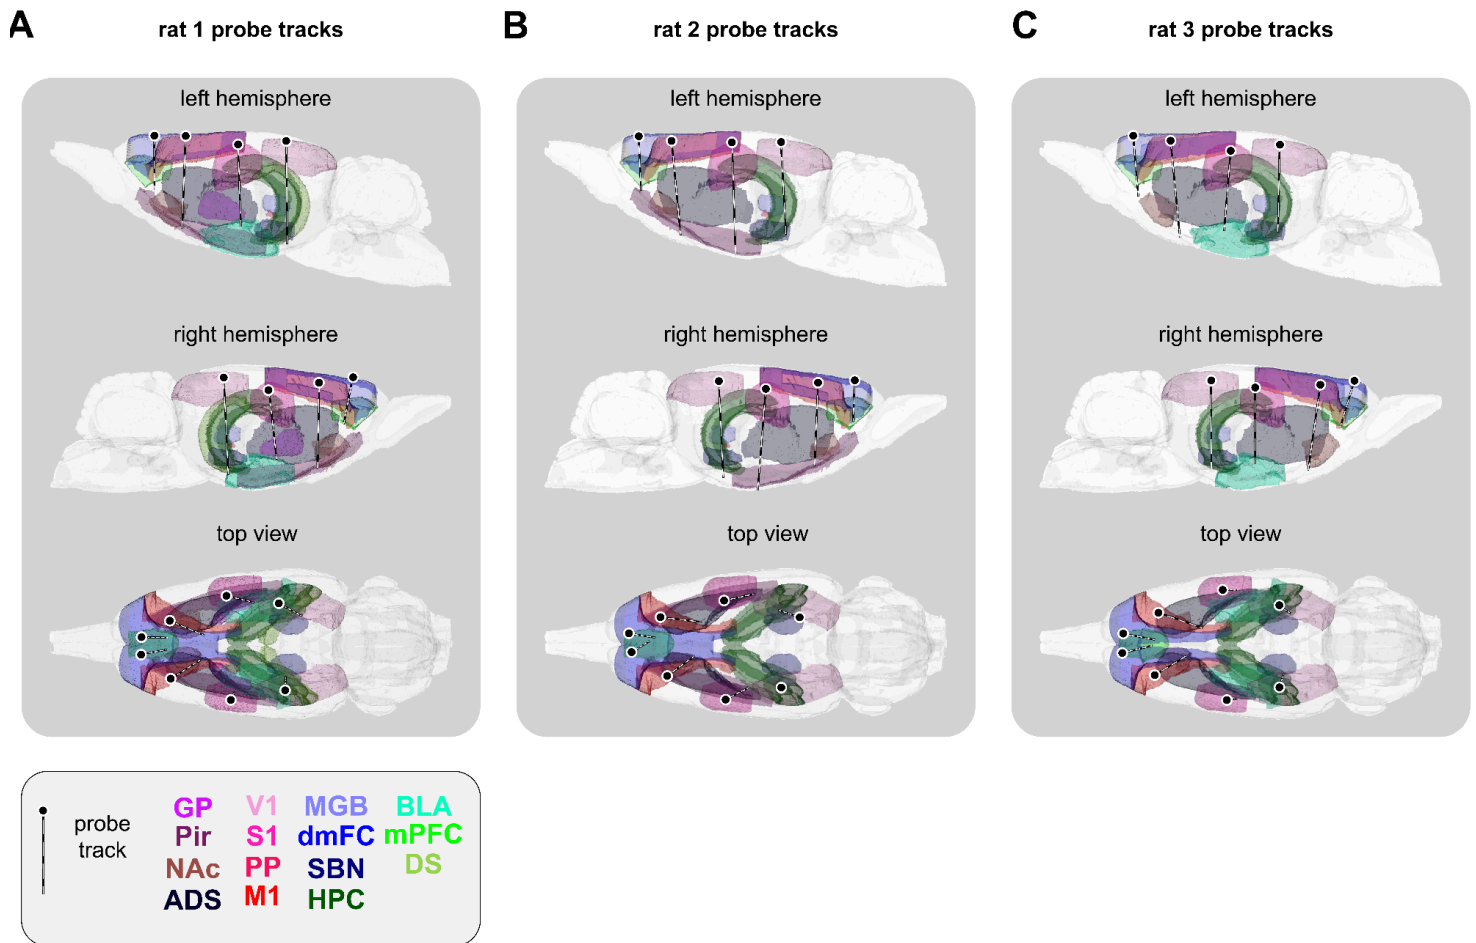


Figure S1. Atlas registered brains for each individual rat.

- (A) The 8 implanted probe tracts of rat 1 reconstructed in the space of a common atlas volume. All regions we recorded units in rat 1 are outlined by annotations from the Waxholm Space rat atlas (Kleven et al., 2023). See Table S1 for full region names and lists for each animal.
- (B) Same as A for rat 2.
- (C) Same as A for rat 3.

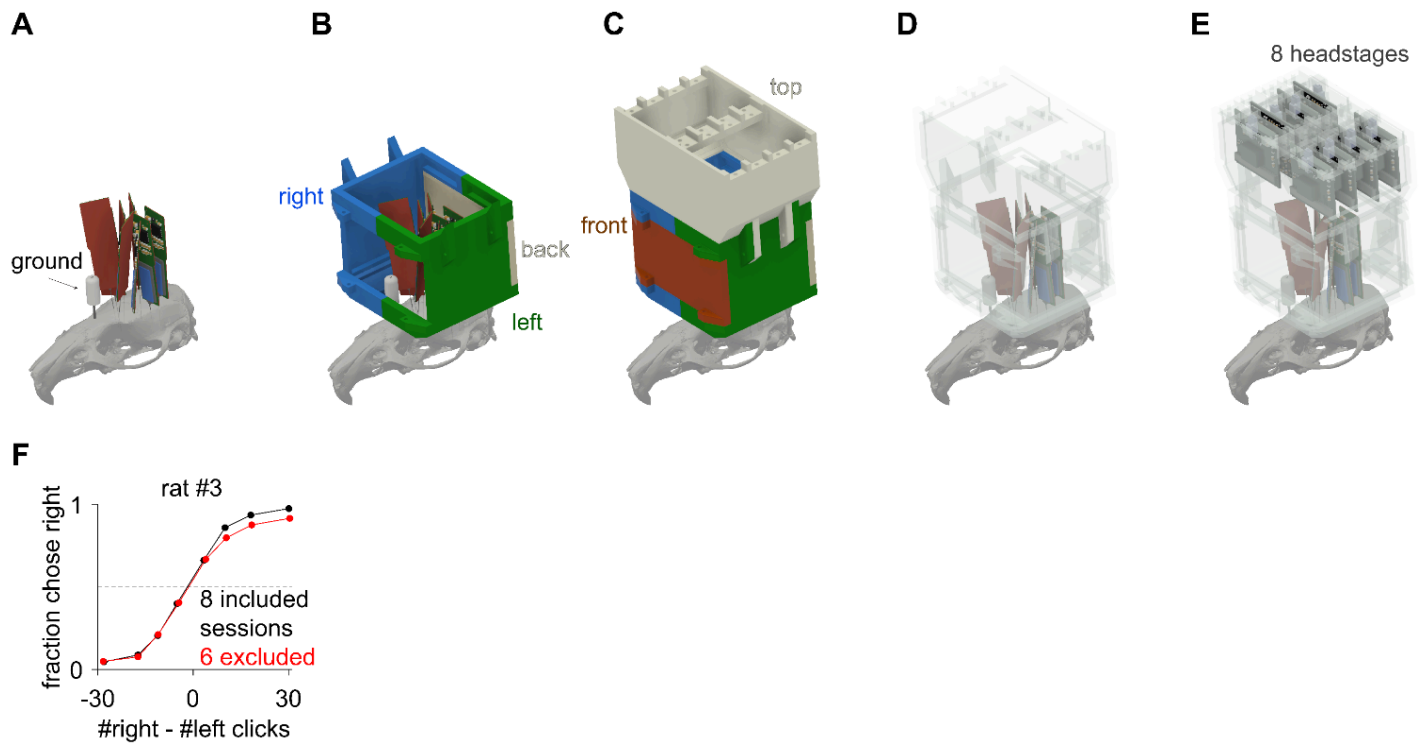


Figure S2. Chassis for protecting the probes and mounting the headstages and behavioral performance of rat 3.

(A) Eight Neuropixels 1.0 probes and a ground cannula.

(B) Three separate 3D-printed components that mate together to form the right, left, and back of the chassis.

(C) The front and top pieces.

(D) View of the probes and ground cannula within the chassis.

(E) The eight headstages are mounted on the top piece.

(F) For rat #3, four recording sessions were excluded on account of a visible lapse, and two excluded due to the number of trials completed being less than 300.

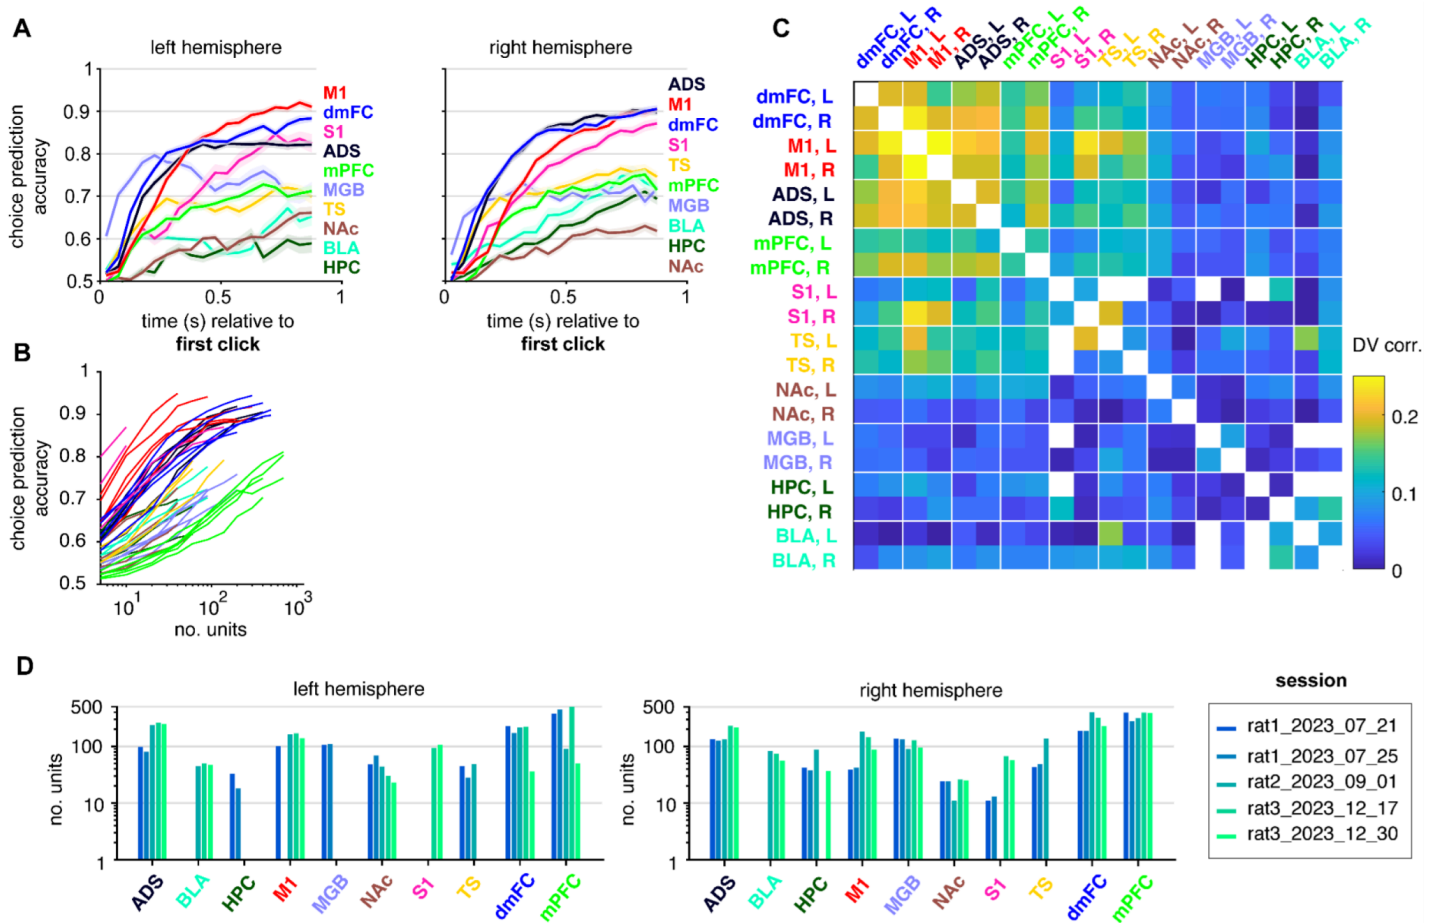


Figure S3. Hemisphere-specific and population size analyses.

- (A) Prediction accuracy of a logistic regression model of choice given population neural activity for 10 brain regions separately, and for all regions together (gray), as a function of time relative to the first click. Accuracy is assessed using the class-balanced accuracy under 10-fold cross-validation (see Methods). Values indicate the average (\pm 1 s.e.) across three rats and the five frozen noise sessions total. This corresponds to the same analysis as in **Fig. 2E**, but with units separated by hemisphere.
- (B) To further understand the influence of population size on choice prediction accuracy, for each of the frozen noise sessions, we repeatedly subsampled units from each region, each time varying the size of the subsampled population. For each population size we plot the peak choice prediction accuracy for each region, estimated as in (A). Each line indicates one region on one session. While this shows a systematic, monotonically increasing effect of population size on choice prediction accuracy, each region traced out its own curve, indicating that the differences in choice prediction accuracy depend both on population size and meaningful differences between regions. Note that the four regions (ADS, M1, dmFC and S1) with highest choice prediction accuracy in (A), when all recorded units are used, correspond to the four regions with the highest per-neuron choice prediction accuracy shown here in (B).
- (C) Matrix of DV correlations (Pearson's ρ) for all pairs of recorded regions (as in **Fig. 2G**), but here with regions separated by hemisphere. We used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus and choice, described in detail in the Methods and illustrated in **Fig. S4**. Only time points before nTc (see Methods) on each trial are included. Note that the pattern of correlations across the brain is highly similar for both hemispheres of each region.
- (D) Histogram showing how many units were recorded from each region/hemisphere in the set of five frozen noise sessions used for choice decoding and DV analysis in **Figs. 2 and 3**.

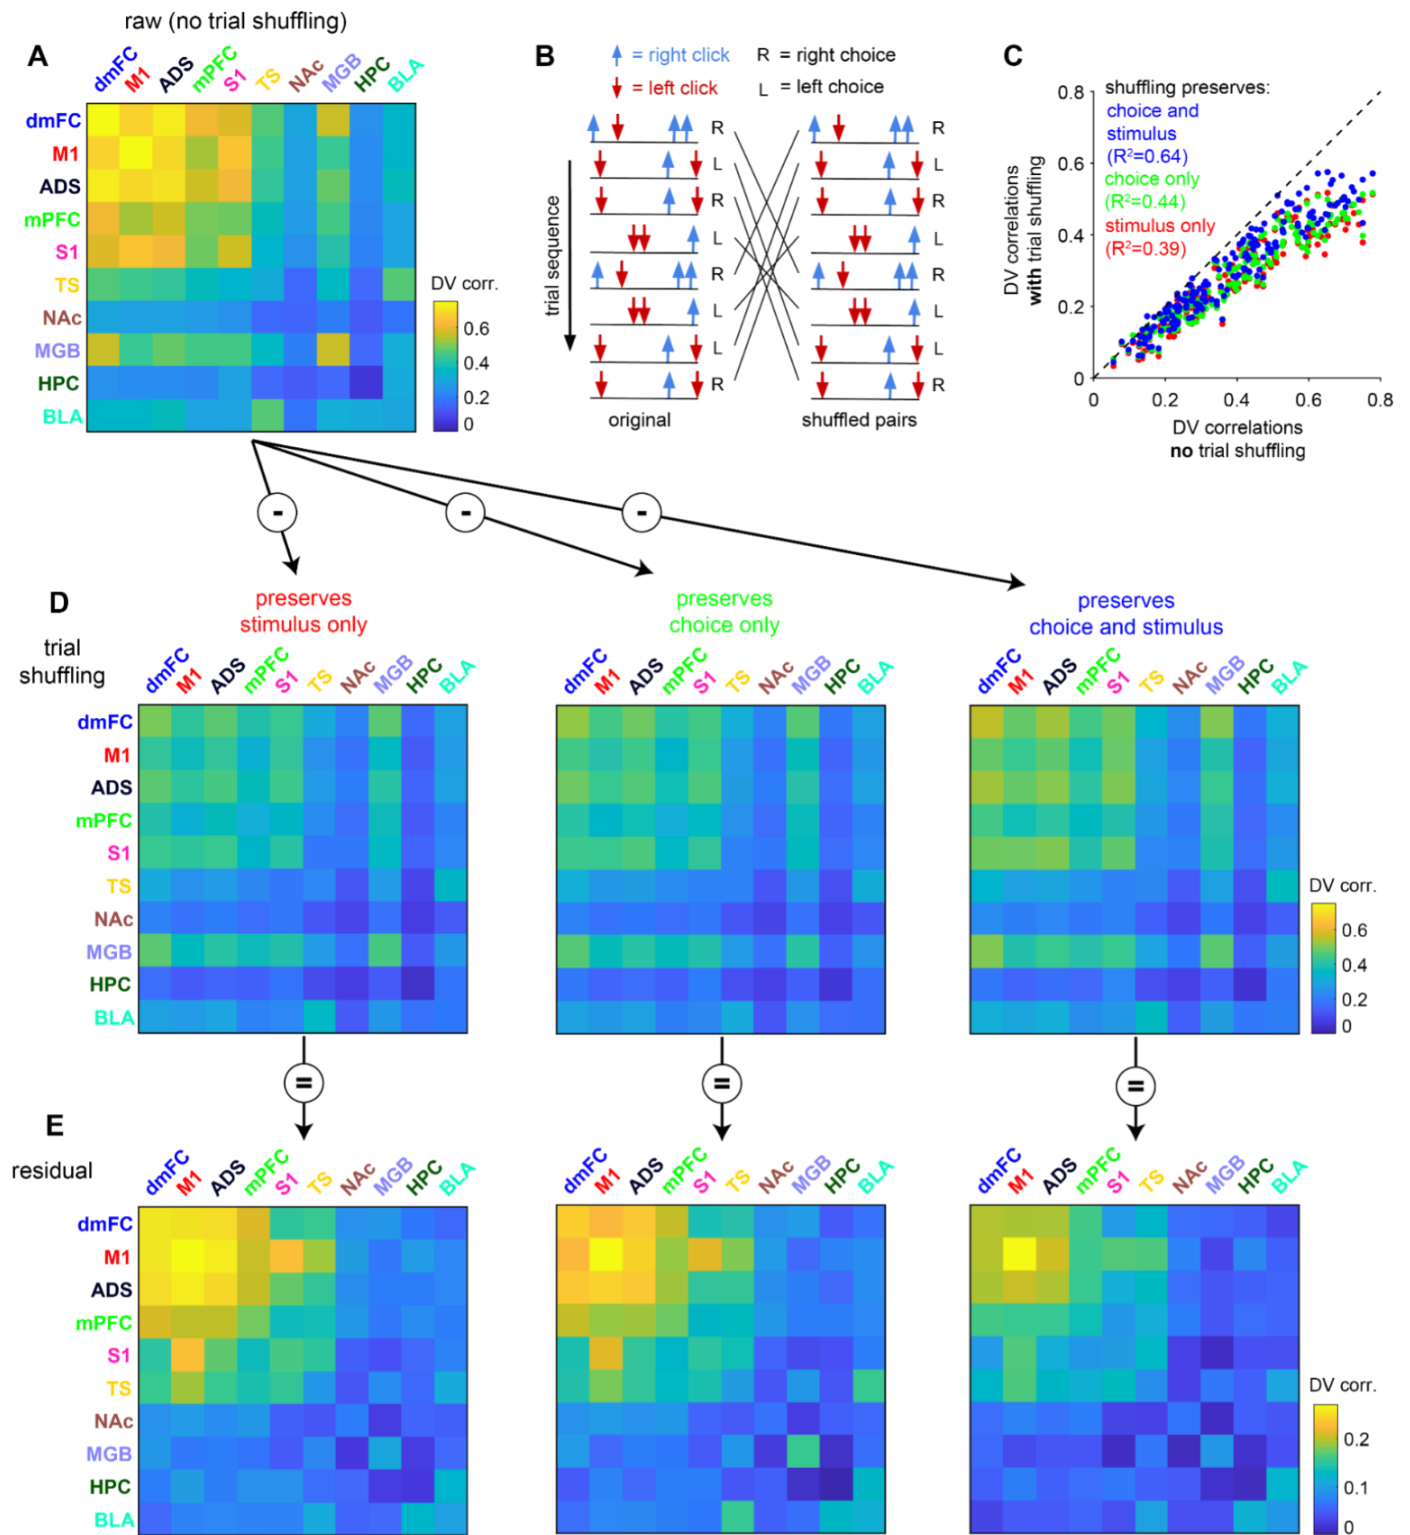


Figure S4. Trial shuffling procedure to remove influence of choice and stimulus on DV correlations.

(A) Matrix of DV correlations (Pearson's ρ) for all pairs of recorded regions (as in Fig. 2G), but here these reflect raw correlations before any correction has been performed. Note how high the values are compared to Fig. 2G. Same set of timepoints as in Fig. 2G used here.

(B) Schematic illustration of the trial shuffling procedure. On a subset of sessions (frozen noise sessions), a fixed set of 54 frozen noise stimuli were used, spanning the normal range of trial difficulties. To obtain a shuffled distribution of DV correlations, we permuted trials while preserving the identity of the stimulus and choice (see Methods).

- (C) Scatter plot illustrating the effect on DV correlation of the trial shuffling procedure. On the x-axis are the values (averaged across sessions) for all pairs of regions, before correction (as in A). On the y-axis are the average of a shuffle distribution. We show the effect of shuffling just stimulus seed, just choice, or both. The shuffle distribution that preserves both captures the largest share of the raw correlations, consistent with the raw DV correlations reflecting shared coding across regions for stimulus and choice.
- (D) Here we show the average of the shuffle distributions (same as y-axis in C) plotted in the form of a matrix.
- (E) Matrix of DV correlations after subtraction of the expectation of the shuffle distributions (i.e. subtracting the matrices in D from the matrix in A). The data in the rightmost column is the remainder after subtracting the correlations expected from shared coding for both stimulus and choice, and corresponds to what is reported throughout the paper.

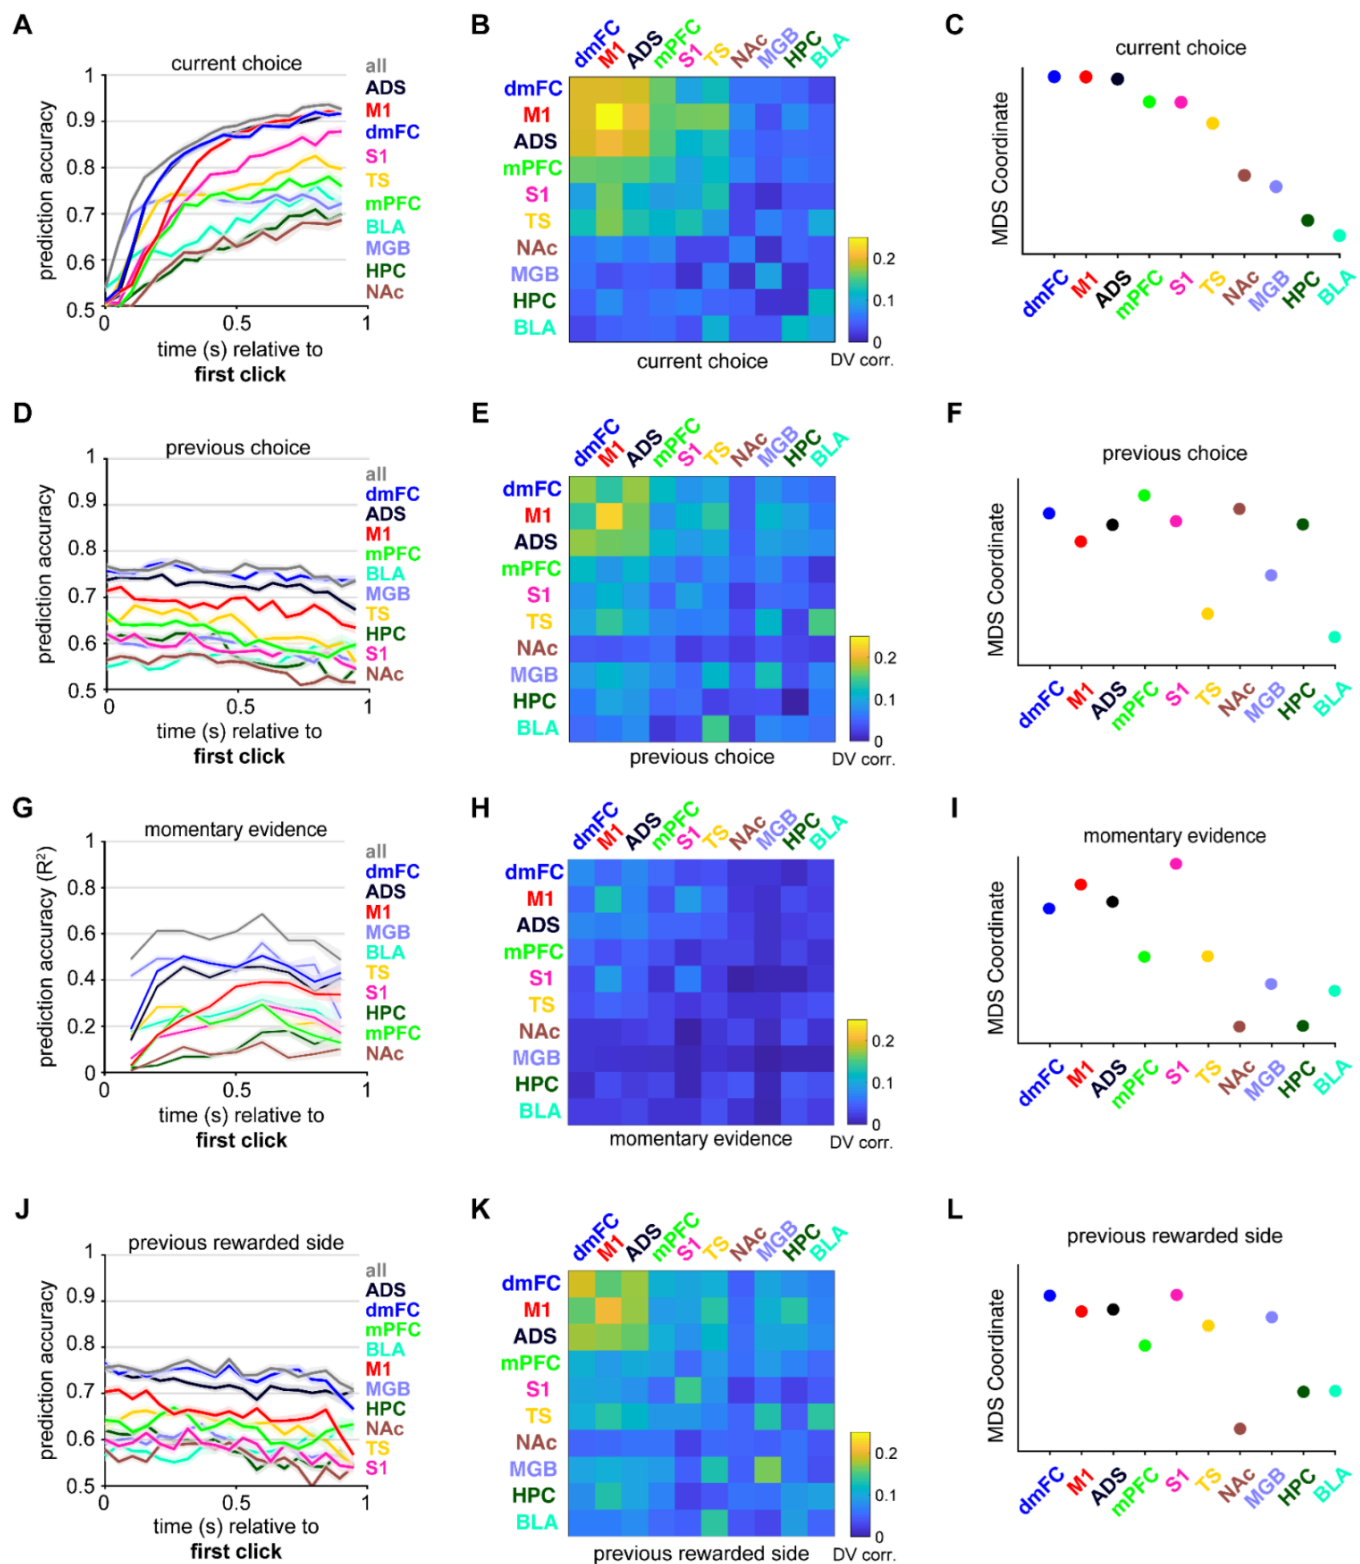


Figure S5. Decoding analyses of other task variables.

(A) Replicated from **Fig. 2E**. Prediction accuracy of a logistic regression model of choice given population neural activity for 10 brain regions separately, and for all regions together (gray), as a function of time relative to the first click. Accuracy is assessed using the class-balanced accuracy under 10-fold cross-validation (see Methods). Values indicate the average (\pm 1 s.e.) across three rats and the five frozen noise sessions total.

(B) Replicated from **Fig. 2G**. Matrix of DV correlations (Pearson's ρ) for all pairs of recorded regions. Diagonal entries in

the matrix indicate the correlation between DVs obtained from each hemisphere of the same brain region. A subset of brain regions concentrated on the frontal cortex and anterior striatum show the highest DV correlations. We used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus and choice, described in detail in the Methods and illustrated in **Fig. S4**. Only time points before nTc on each trial are included.

- (C) Replicated from **Fig. 2H**. DV correlation matrix in (B) was projected down onto a one-dimensional manifold using classical multidimensional scaling (MDS).
- (D) Same as (A) for the previous choice.
- (E) Same as (B) except shows correlations along the dimension in neural state space best predicting previous choice. We refer to these also as DV correlations where DV refers to “decoded variable.” We used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus, current choice and previous choice.
- (F) Previous choice “decoded variable” matrix in (E) was projected onto a one-dimensional manifold using classical MDS.
- (G) Same as (A) except shows accuracy of a linear regression model of momentary evidence using neural population activity, quantified using R^2 . Momentary evidence is defined as the difference in right versus left clicks within a 100 ms window before the time of spiking.
- (H) Same as (B) except shows correlations along the dimension in neural state space best predicting momentary evidence. We used a shuffle-correction procedure to remove the component of correlation due to shared coding for stimulus, current choice and previous choice.
- (I) Momentary evidence “decoded variable” matrix in (H) was projected onto a one-dimensional manifold using classical MDS.
- (J) Prediction accuracy of a logistic regression model of previous rewarded side given population neural activity for 10 brain regions separately, and for all regions together (gray), as a function of time relative to the first click. Accuracy is assessed using the class-balanced accuracy under 10-fold cross-validation (see Methods). Values indicate the average (\pm 1 s.e.) across three rats and the five frozen noise sessions total.
- (K) Same as (E) for previously rewarded side.
- (L) Previous rewarded side “decoded variable” matrix in (K) was projected onto a one-dimensional manifold using classical MDS.

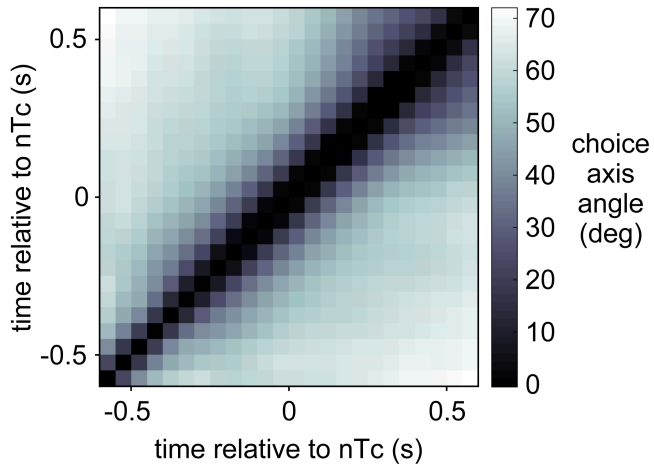


Figure S6. Rotation of choice-related neural dynamics slows around nTc .

In **Fig. 3M**, we show a reduction in the angular velocity of the choice axis (i.e. difference in angle between choice axis computed at neighboring time points). Here we show the difference in angle between all pairs of time points as a matrix. The data in **Fig. 3M** corresponds to the first off-diagonal of this matrix. The full matrix shows that the difference in angle between time points grows continually as time points are separated by more elapsed time, although this growth is slowed after nTc . This is consistent with a slowing of ongoing rotation around nTc , rather than a reduction in the amplitude of temporally-uncorrelated perturbations in the choice axis.

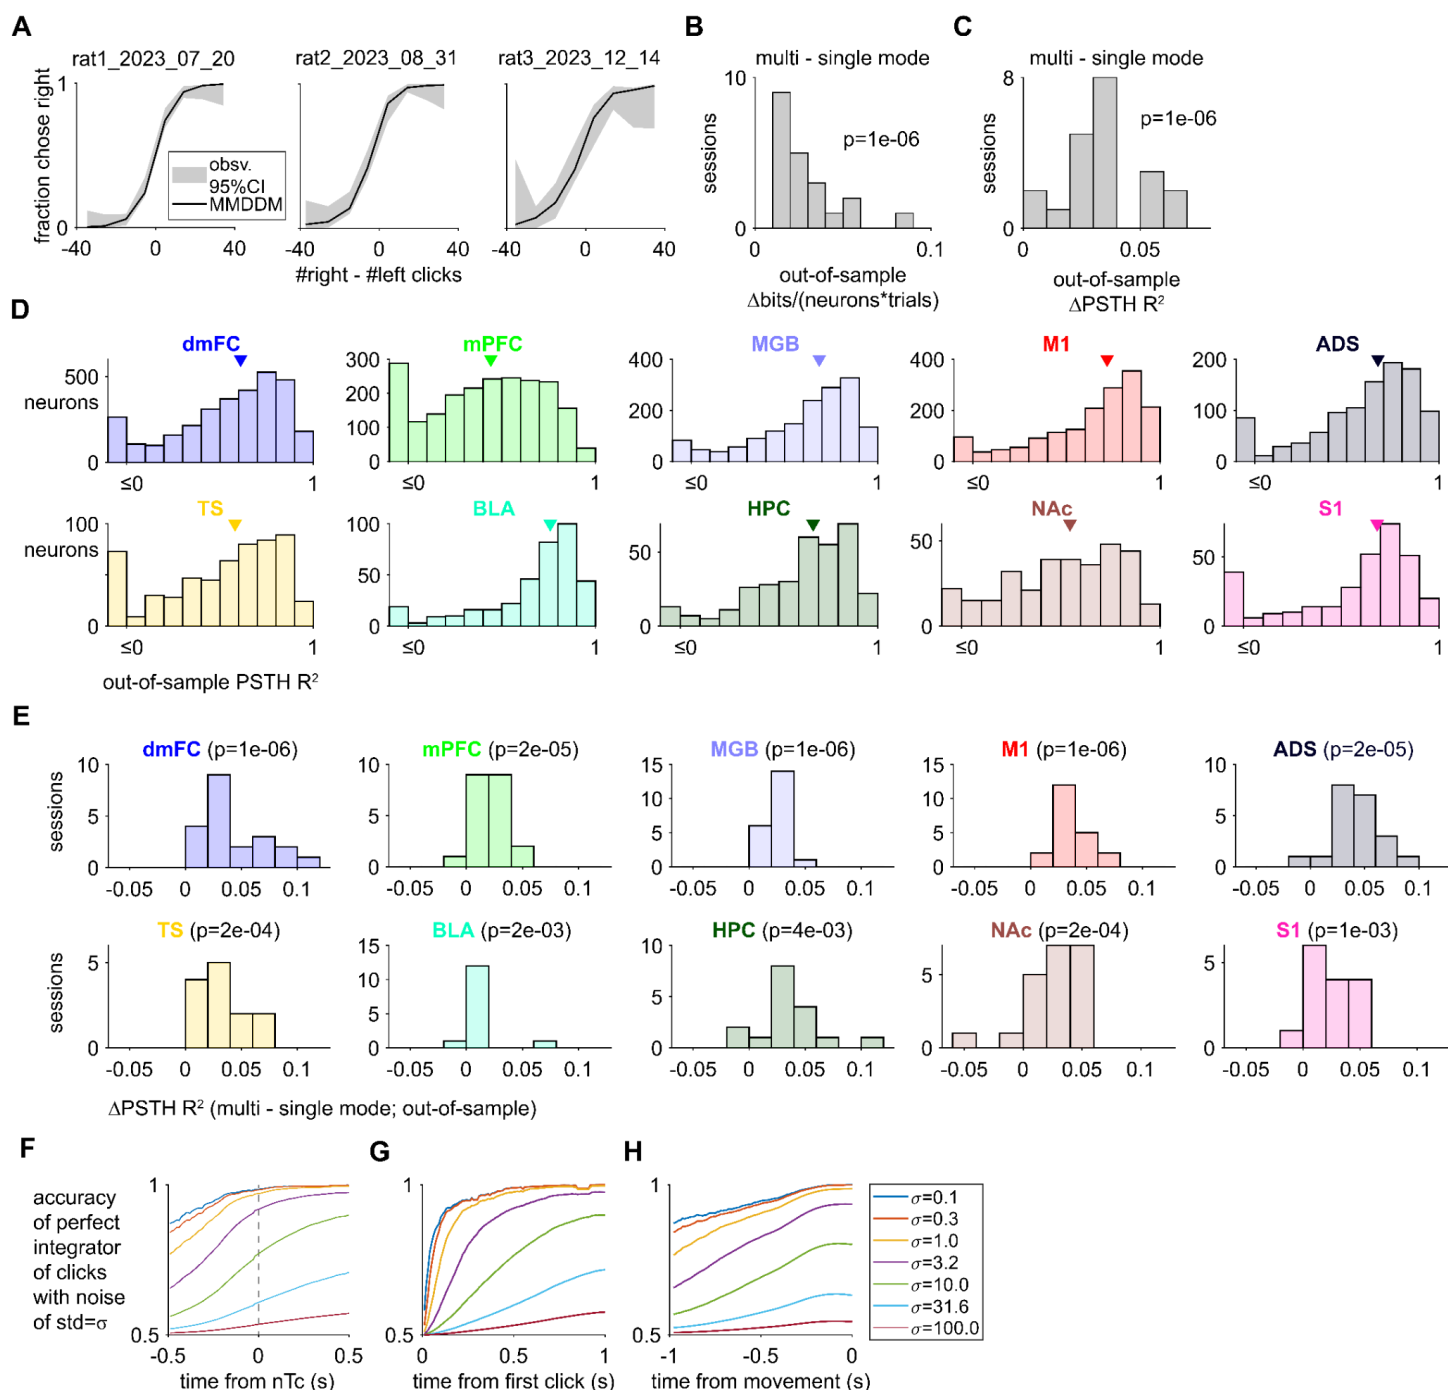


Figure S7. Analyses related to the multi-mode drift diffusion model.

- (A) Observed and MMDDM-predicted psychometric function for a representative session from each of the three rats. Confidence intervals of the observed were computed using the Clopper-Pearson method.
- (B) The data are more likely under MMDDM than the single-mode DDM. P-value is computed from the two-sided sign test, testing the hypothesis that data in x has a continuous distribution with zero median against the alternative that the distribution does not.
- (C) The peri-stimulus time histograms (PSTH; aligned to stimulus onset, conditioned on the behavioral choice) are better captured by the multi-mode DDM than the single-mode DDM (two-sided sign test).
- (D) Goodness-of-fit of the choice-conditioned out-of-sample peristimulus time histogram measured (PSTH) as the coefficient-of-determination (R^2). Marker indicates the median of the distribution. Neurons for whom the PSTH R^2 are less than 0 are in the leftmost bin.

- (E) For each of ten brain regions analyzed, the peri-stimulus time histograms (PSTH; aligned to stimulus onset, conditioned on the behavioral choice) are better captured by the multi-mode DDM than the single-mode DDM (two-sided sign test).
- (F) Accuracy of a perfect integrator with varying levels of noise aligned the actual neurally inferred times of commitment. The perfect integrator has no absorbing bound. Noise is implemented as i.i.d. Gaussian noise added on each time step.
- (G) Accuracy aligned the actual time of the onset of the auditory click trains (always indicated by a simultaneous click from the left and from the right speaker).
- (H) Aligned to the actual times when the rat initiated movement away from the fixation port.

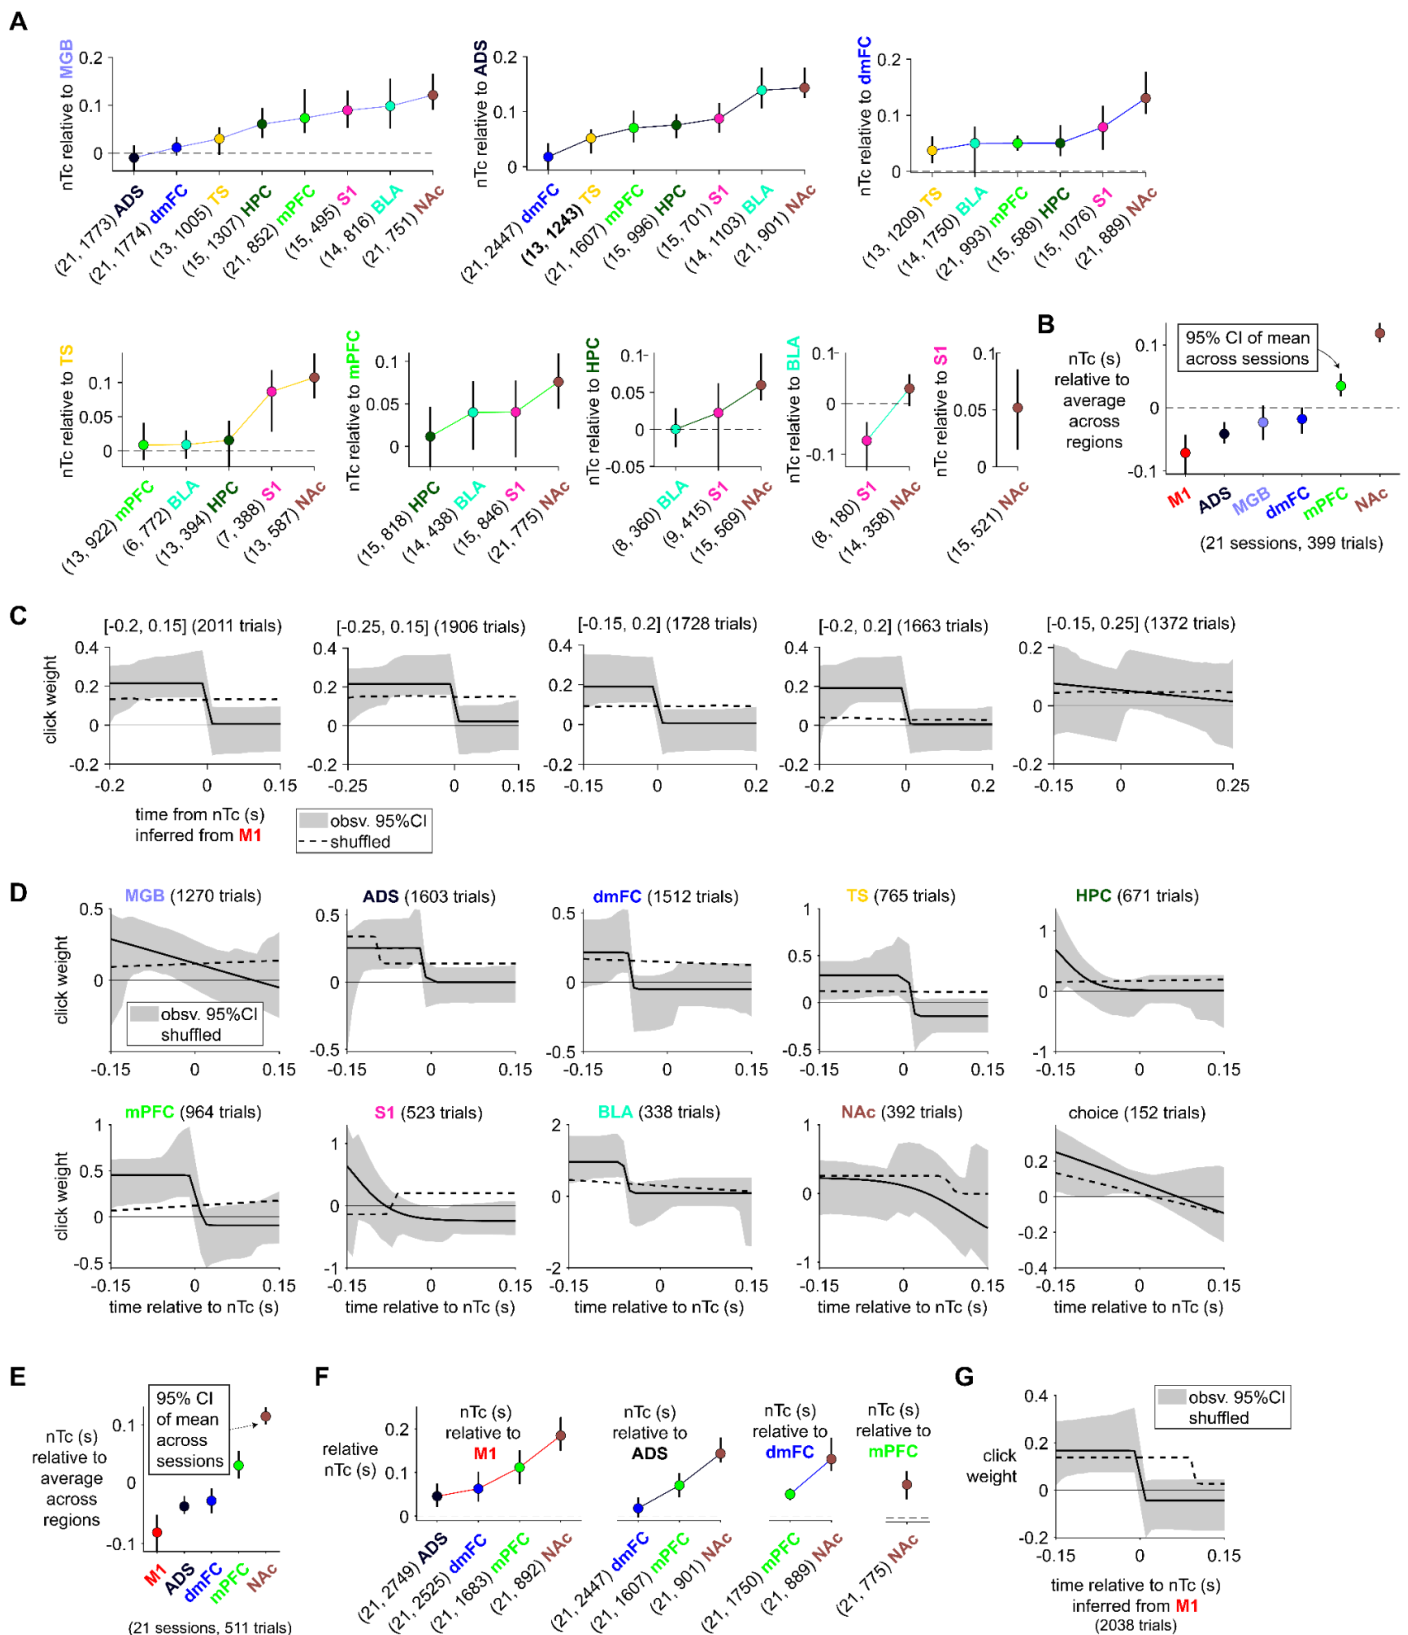


Figure S8. Relative timing of nTc's using spikes from different brain regions and psychophysical kernels aligned to nTc's.

(A) Comparison between the time of decision commitment inferred separately from each pair of brain region, using the subset of trials on which commitment could be inferred using either type of data. The number of sessions and trials

are indicated as (#sessions, #trials). Error bars indicate 95% bootstrapped confidence interval of the mean across sessions.

- (B) For the six brain regions that were recorded simultaneously on each of the 21 sessions, we examined the subset of trials on which a putative time of commitment could be inferred separately using the spikes from each of the six regions (399 trials). Error bar indicates the 95% bootstrapped confidence interval of the mean.
- (C) Psychophysical kernels using different time intervals aligned to the nT_c inferred from M1 activity. As the duration of the time interval increases, there are fewer and fewer trials on which clicks occurred throughout the duration.
- (D) Psychophysical kernel estimated by aligning to the nT_c inferred using neural activity from separate brain regions or the animal's behavioral choice
- (E) The results shown in Figure 4 are not sensitive to the set of the brain regions fit. MMDDM were fit to only brain regions ADS, M1, NAc, and dmFC, mPFC.

<i>(Paxinos and Watson, 2009) region name and abbreviation</i>	<i>Region name and abbreviation in this paper</i>	<i>Recorded in Rat 1</i>	<i>Recorded in Rat 2</i>	<i>Recorded in Rat 3</i>
caudate putamen/striatum (CPu)	anterior dorsal striatum (ADS)	L,R	L, R	L, R
caudate putamen/striatum (CPu)	tail of the striatum (TS)	L, R	L, R	L, R
primary motor cortex (M1)	primary motor cortex (M1)	L, R	L, R	L, R
cingulate cortex, area 1 (Cg1)	dorsomedial frontal cortex (dmFC)	L, R	L, R	L, R
secondary motor cortex (M2)	dorsomedial frontal cortex (dmFC)	L, R	L, R	L, R
prelimbic cortex (PrL)	medial prefrontal cortex (mPFC)	L, R	L, R	L, R
medial orbital cortex (MO)	medial prefrontal cortex (mPFC)	X	L, R	L
primary somatosensory cortex (S1)	primary somatosensory cortex (S1)	L, R	L, R	L, R
dentate gyrus (DG)	hippocampus (HPC)	L, R	L, R	R
field CA1 of the hippocampus (CA1)	hippocampus (HPC)	R	L, R	R
accumbens nucleus (Acb)	nucleus accumbens (NAc)	L, R	L, R	L, R
subbrachial nucleus (SubB)	subbrachial nucleus (SBN)	L, R	R	R
medial geniculate nucleus (MG)	medial geniculate body (MGB)	L, R	L, R	R
basolateral amygdaloid nucleus (BL)	basolateral amygdala (BLA)	X	L, R	L, R
peripeduncular nucleus (PP)	peripeduncular nucleus (PP)	X	L	X
piriform cortex (Pir)	piriform cortex (Pir)	X	L, R	X
globus pallidus (GP)	globus pallidus (GP)	L	X	X
primary visual cortex (V1)	primary visual cortex (V1)	L, R	L, R	R
substantia nigra, reticular part (SNR)	substantia nigra (SN)	X	L	X
dorsal subiculum (DS)	dorsal subiculum (DS)	L	X	X

Supplemental Table 1

Complete listing of recorded brain regions. In some cases, we refer to a brain region differently than what is used in the Paxinos and Watson rat brain atlas. In other cases, we combine two regions named by Paxinos and Watson and refer to both by one name. 'L' indicates that the region was recorded from the left hemisphere of a rat subject; 'R' indicates it was recorded from the right hemisphere. Cells of the table marked 'X' were not recorded from in the given subject.

References

- Ährlund-Richter S, Xuan Y, van Lunteren JA, Kim H, Ortiz C, Pollak Dorocic I, Meletis K, Carlén M (2019) A whole-brain atlas of monosynaptic input targeting four different cell types in the medial prefrontal cortex of the mouse. *Nat Neurosci* 22:657–668.
- Anastasiades PG, Carter AG (2021) Circuit organization of the rodent medial prefrontal cortex. *Trends Neurosci* 44:550–563.
- Aronoff R, Matyas F, Mateo C, Ciron C, Schneider B, Petersen CCH (2010) Long-range connectivity of mouse primary somatosensory barrel cortex. *Eur J Neurosci* 31:2221–2233.
- Basso MA, Wurtz RH (1998) Modulation of neuronal activity in superior colliculus by changes in target probability. *J Neurosci* 18:7519–7534.
- Bimbard C et al. (2024) An adaptable, reusable, and light implant for chronic Neuropixels probes. bioRxiv Available at: <http://dx.doi.org/10.1101/2023.08.03.551752>.
- Bogacz R, Gurney K (2007) The basal ganglia and cortex implement optimal decision making between alternative actions. *Neural Comput* 19:442–477.
- Boyd-Meredith JT, Piet AT, Dennis EJ, El Hady A, Brody CD (2022) Stable choice coding in rat frontal orienting fields across model-predicted changes of mind. *Nat Commun* 13:3235.
- Brunton BW, Botvinick MM, Brody CD (2013) Rats and humans can optimally accumulate evidence for decision-making. *Science* 340:95–98.
- Buetfering C, Zhang Z, Pitsiani M, Smallridge J, Boven E, McElligott S, Häusser M (2022) Behaviorally relevant decision coding in primary somatosensory cortex neurons. *Nat Neurosci* 25:1225–1236.
- Chen G, Kang B, Lindsey J, Druckmann S, Li N (2021) Modularity and robustness of frontal cortical networks. *Cell* 184:3717–3730.e24.
- Chen L, Wang X, Ge S, Xiong Q (2019) Medial geniculate body and primary auditory cortex differentially contribute to striatal sound representations. *Nat Commun* 10:418.
- Chen S, Liu Y, Wang ZA, Colonell J, Liu LD, Hou H, Tien N-W, Wang T, Harris T, Druckmann S, Li N, Svoboda K (2024) Brain-wide neural activity underlying memory-guided movement. *Cell* 187:676–691.e16.
- Dennis EJ, Bibawi P, Dhanerawala ZM, Lynch LA, Wang SS-H, Brody CD (2023) Princeton RAtlas: A Common Coordinate Framework for Fully cleared, Whole Brains. *Bio Protoc* 13:e4854.
- DePasquale B, Brody CD, Pillow JW (2024) Neural population dynamics underlying evidence accumulation in multiple rat brain regions. *Elife* 13 Available at: <https://elifesciences.org/articles/84955> [Accessed August 23, 2024].
- Ding L, Gold JI (2010) Caudate encodes multiple computations for perceptual decisions. *J Neurosci* 30:15747–15759.
- Eliasmith C, Stewart TC, Choo X, Bekolay T, DeWolf T, Tang Y, Rasmussen D (2012) A large-scale model of the functioning brain. *Science* 338:1202–1205.
- Fabre JMJ, van Beest EH, Peters AJ, Carandini M, Harris KD (2023) Bombcell: automated curation and cell classification of spike-sorted electrophysiology data. Zenodo. Available at: <https://zenodo.org/record/8172821>.

- Friedman J, Hastie T, Tibshirani R (2010) Regularization Paths for Generalized Linear Models via Coordinate Descent. *J Stat Softw* 33:1–22.
- Grossberg S (2019) The Embodied Brain of SOVEREIGN2: From Space-Variant Conscious Percepts During Visual Search and Navigation to Learning Invariant Object Categories and Cognitive-Emotional Plans for Acquiring Valued Goals. *Front Comput Neurosci* 13:36.
- Guo L, Walker WI, Ponvert ND, Penix PL, Jaramillo S (2018) Stable representation of sounds in the posterior striatum during flexible auditory decisions. *Nat Commun* 9:1534.
- Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K (2014) Flow of Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* 81:179–194.
- Gupta D, Kopec CD, Bondy AG, Luo TZ, Elliott VA, Brody CD (2024) A multi-region recurrent circuit for evidence accumulation in rats. *bioRxiv* Available at: <http://dx.doi.org/10.1101/2024.07.08.602544>.
- Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD (2015) Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520:220–223.
- Hernández-Navarro L, Hermoso-Mendizabal A, Duque D, de la Rocha J, Hyafil A (2021) Proactive and reactive accumulation-to-bound processes compete during perceptual decisions. *Nat Commun* 12:7148.
- Horan M et al. (2024) Repix: reliable, reusable, versatile chronic Neuropixels implants using minimal components. *bioRxiv:2024.04.25.591118* Available at: <https://www.biorxiv.org/content/biorxiv/early/2024/04/30/2024.04.25.591118> [Accessed August 16, 2024].
- Horwitz GD, Batista AP, Newsome WT (2004) Representation of an abstract perceptual decision in macaque superior colliculus. *J Neurophysiol* 91:2281–2296.
- Hunnicutt BJ, Jongbloets BC, Birdsong WT, Gertz KJ, Zhong H, Mao T (2016) A comprehensive excitatory input map of the striatum reveals novel functional organization. *Elife* 5 Available at: <http://dx.doi.org/10.7554/eLife.19103>.
- Inagaki HK, Chen S, Ridder MC, Sah P, Li N, Yang Z, Hasanbegovic H, Gao Z, Gerfen CR, Svoboda K (2022) A midbrain-thalamus-cortex circuit reorganizes cortical dynamics to initiate movement. *Cell* 185:1065–1081.e23.
- International Brain Laboratory et al. (2023) A Brain-Wide Map of Neural Activity during Complex Behaviour. *bioRxiv:2023.07.04.547681* Available at: <https://www.biorxiv.org/content/10.1101/2023.07.04.547681> [Accessed August 13, 2024].
- Jiang H, Kim HF (2018) Anatomical Inputs From the Sensory and Value Structures to the Tail of the Rat Striatum. *Front Neuroanat* 12:30.
- Juavinett AL, Bekheet G, Churchland AK (2018) Chronically-implanted Neuropixels probes enable high yield recordings in freely moving mice. *bioRxiv:406074* Available at: <https://www.biorxiv.org/content/10.1101/406074v1> [Accessed July 8, 2019].
- Jun EJ, Bautista AR, Nunez MD, Allen DC, Tak JH, Alvarez E, Basso MA (2021) Causal role for the primate superior colliculus in the computation of evidence for perceptual decisions. *Nat Neurosci* 24:1121–1131.
- Jun JJ et al. (2017) Fully integrated silicon probes for high-density recording of neural activity. *Nature* 551:232–236.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2015) Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife* 4:e04677.

- Kiani R, Cueva CJ, Reppas JB, Newsome WT (2014) Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials. *Curr Biol* 24:1542–1547.
- Kleven H, Bjerke IE, Clascá F, Groenewegen HJ, Bjaalie JG, Leergaard TB (2023) Waxholm Space atlas of the rat brain: a 3D atlas supporting data analysis and integration. *Nat Methods* 20:1822–1829.
- Kopec CD, Luo TZ, Bondy AG, Gupta D, Elliott VA, Charlton JA, Breda JR, Stagnaro WM, Reyes EJ, Sirko AI, Bustos AF, Willock JM, Morrison JM, Osorio KL, Brody CD (2024) To integrate or not to integrate: Testing degenerate strategies for solving an accumulation of perceptual evidence decision-making task. *bioRxiv* Available at: <http://biorxiv.org/lookup/doi/10.1101/2024.08.21.609064>.
- Kwon SE, Yang H, Minamisawa G, O'Connor DH (2016) Sensory and decision-related activity propagate in a cortical feedback loop during touch perception. *Nat Neurosci* 19 Available at: <http://www.nature.com/doi/10.1038/nn.4356> <http://www.ncbi.nlm.nih.gov/pubmed/27437910>.
- LeDoux JE, Farb CR, Romanski LM (1991) Overlapping projections to the amygdala and striatum from auditory processing areas of the thalamus and cortex. *Neurosci Lett* 134:139–144.
- Li N, Chen T-W, Guo ZV, Gerfen CR, Svoboda K (2015) A motor cortex circuit for motor planning and movement. *Nature* 519:51–56.
- Li N, Daie K, Svoboda K, Druckmann S (2016) Robust neuronal dynamics in premotor cortex during motor planning. *Nature* 532:459–464.
- Lo C-C, Wang X-J (2006) Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat Neurosci* 9:956–963.
- Luo TZ, Bondy AG, Gupta D, Elliott VA, Kopec CD, Brody CD (2020) An approach for long-term, multi-probe Neuropixels recordings in unrestrained rats. *Elife* 9 Available at: <http://dx.doi.org/10.7554/eLife.59716>.
- Luo TZ, Kim TD, Gupta D, Bondy AG, Kopec CD, Elliott VA, DePasquale B, Brody C (2023) Transitions in dynamical regime and neural mode underlie perceptual decision-making. *bioRxiv:2023.10.15.562427* Available at: <https://www.biorxiv.org/content/10.1101/2023.10.15.562427v2> [Accessed November 21, 2023].
- Mazurek ME, Roitman JD, Ditterich J, Shadlen MN (2003) A Role for Neural Integrators in Perceptual Decision Making. *Cereb Cortex* 13:1257–1269.
- Mejias JF, Murray JD, Kennedy H, Wang X-J (2016) Feedforward and feedback frequency-dependent interactions in a large-scale laminar network of the primate cortex. *Sci Adv* 2:e1601335.
- Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK (2019) Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci* 22:1677–1686.
- Musiek FE, Baran JA (2018) *The Auditory System: Anatomy, Physiology, and Clinical Correlates; Second Edition*. Plural Publishing.
- Nieh EH, Schottdorf M, Freeman NW, Low RJ, Lewallen S, Koay SA, Pinto L, Gauthier JL, Brody CD, Tank DW (2021) Geometry of abstract learned knowledge in the hippocampus. *Nature* 595:80–84.
- O'Reilly RC, Hazy TE, Herd SA (2015) *The Leabra Cognitive Architecture* (Chipman SEF, ed). Oxford University Press.
- Pachitariu M, Sridhar S, Stringer C (2023) Solving the spike sorting problem with Kilosort. *bioRxiv:2023.01.07.523036* Available at: <https://www.biorxiv.org/content/biorxiv/early/2023/01/07/2023.01.07.523036> [Accessed August 17, 2024].

- Park IM, Meister MLR, Huk AC, Pillow JW (2014) Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nat Neurosci* 17:1395–1403.
- Paxinos G, Watson C (2009) *The Rat Brain in Stereotaxic Coordinates*. Elsevier/Academic.
- Peixoto D, Verhein JR, Kiani R, Kao JC, Nuyujukian P, Chandrasekaran C, Brown J, Fong S, Ryu SI, Shenoy KV, Newsome WT (2021) Decoding and perturbing decision states in real time. *Nature* 591:604–609.
- Piet AT, El Hady A, Brody CD (2018) Rats adopt the optimal timescale for evidence integration in a dynamic environment. *Nat Commun* 9:4265.
- Pietzsch T, Saalfeld S, Preibisch S, Tomancak P (2015) BigDataViewer: visualization and processing for large image data sets. *Nat Methods* 12:481–483.
- Rabinowitz NC, Goris RL, Cohen M, Simoncelli EP (2015) Attention stabilizes the shared gain of V4 populations. *Elife* 4:e08998.
- Ratcliff R, McKoon G (2008) The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput* 20:873–922.
- Roitman JD, Shadlen MN (2002) Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22:9475–9489.
- Schindelin J, Arganda-Carreras I, Frise E, Kaynig V, Longair M, Pietzsch T, Preibisch S, Rueden C, Saalfeld S, Schmid B, Tinevez J-Y, White DJ, Hartenstein V, Eliceiri K, Tomancak P, Cardona A (2012) Fiji: an open-source platform for biological-image analysis. *Nat Methods* 9:676–682.
- Shadlen MN, Newsome WT (2001) Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J Neurophysiol* 86:1916–1936.
- Steinemann NA, Stine GM, Trautmann EM, Zylberberg A, Wolpert DM, Shadlen MN (2022) Direct observation of the neural computations underlying a single decision. *bioRxiv* Available at: <https://www.biorxiv.org/content/10.1101/2022.05.02.490321.abstract>.
- Steinmetz NA et al. (2021) Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* 372 Available at: <http://dx.doi.org/10.1126/science.abf4588>.
- Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD (2019) Distributed coding of choice, action and engagement across the mouse brain. *Nature* 576:266–273.
- Stine GM, Trautmann EM, Jeurissen D, Shadlen MN (2023) A neural mechanism for terminating decisions. *Neuron* 111:2601–2613.e5.
- Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD (2019) Spontaneous behaviors drive multidimensional, brainwide activity. *Science* 364:255.
- Thura D, Cisek P (2014) Deliberation and commitment in the premotor and primary motor cortex during dynamic decision making. *Neuron* 81:1401–1416.
- Wang ZA, Chen S, Liu Y, Liu D, Svoboda K, Li N, Druckmann S (2023) Not everything, not everywhere, not all at once: a study of brain-wide encoding of movement. *bioRxiv* Available at: <http://dx.doi.org/10.1101/2023.06.08.544257>.
- Yang H, Kwon SE, Severson KS, O'Connor DH (2015) Origins of choice-related activity in mouse somatosensory cortex. *Nat Neurosci* 19:127–134.

Yartsev MM, Hanks TD, Yoon AM, Brody CD (2018) Causal contribution and dynamical encoding in the striatum during evidence accumulation. bioRxiv:245316 Available at: <https://www.biorxiv.org/content/early/2018/01/10/245316.abstract> [Accessed February 28, 2018].

Znamenskiy P, Zador AM (2013) Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* 497:482–485.