

# Brain-wide coordination of internal signals during decision-making

Adrian G. Bondy\*<sup>†1</sup>, Julie A. Charlton\*<sup>†1</sup>, Thomas Zhihao Luo\*<sup>†1,2</sup>, Charles D. Kopec<sup>1</sup>, Wynne M. Stagnaro<sup>1</sup>, Sarah Jo C. Venditto<sup>1</sup>, Laura Lynch<sup>1</sup>, Sanjeev Janarthanan<sup>1</sup>, Stefan N. Oline<sup>1</sup>, Timothy D. Harris<sup>3,4</sup>, Carlos D. Brody\*<sup>1,2</sup>

<sup>1</sup> Princeton Neuroscience Institute, Princeton University, Princeton NJ, USA

<sup>2</sup> Howard Hughes Medical Institute, Princeton University, Princeton NJ, USA

<sup>3</sup> Janelia Research Campus, Howard Hughes Medical Institute, VA, USA

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

<sup>†</sup> equal contribution, alphabetically ordered

\* corresponding authors: {abondy, jc2988, zhihao1, brody} [at] princeton [dot] edu

**Neural activity is often analyzed with respect to external referents, such as the onset of a sensory stimulus or an overt motor action. Simultaneous recordings allow referencing neurons' activity to each other and thus detecting signals that are internal to the organism. Further, multi-region simultaneous recordings allow observing how these internal signals are coordinated across the brain. Following this logic in rats performing a perceptual decision-making task, we recorded simultaneously from thousands of neurons across up to 20 brain regions at once. Here we report two internal signals which we found to profoundly shape decision-related neural dynamics and brain states. First, we decoded the continuously evolving decision state separately from each region, and found surprisingly large magnitude co-fluctuations in these measures. Dimensionality analysis showed these to be dominated by a single state variable, suggesting that only a single decision-making computation, not multiple parallel computations, are being carried out during the analyzed period. Second, we found that the precise time the subject commits to a decision – a covert event that we decoded from large-scale neural activity in primary motor cortex – was accompanied by a coordinated change, across the brain, from a decision formation to a post-commitment state. The two states differ substantially in their choice-predictive neural dynamics and in their inter-region correlations. Therefore, knowing the time of this state change on single trials is needed to correctly parse fundamentally different phases of decision-making. Overall, our data suggest that internally-referenced signals and state changes, not timelocked to external events but detectable through simultaneous recordings, are major features of neural activity during cognition.**

Cognitive processes, unlike sensation and action, are often poorly correlated with external events, instead reflecting signals and states that are *internal* to the organism. While traditional neuroscience methods correlate neuronal activity to stimuli or overt motor actions, large-scale recording tools allow us to correlate neuronal activity with itself—offering the potential to reveal coordinated patterns of neural activity that reflect these internal processes. Furthermore, recording populations of neurons across multiple regions simultaneously allows observation of how internally-generated signals co-evolve across the brain in real time, which provides greater precision in characterizing these signals as well as powerful constraints on the multi-region circuits that generate them. Comparing the relative timing of shared cognitive signals across regions on single trials allows assessing their origin, and the patterns of cross-region correlation in these signals allows assessing the number of independent subcircuits generating them (Fig. 1a). Importantly, large-scale surveys across regions that are conducted serially<sup>1,2</sup>, rather than simultaneously, do not provide such constraints, as they only allow comparisons between regions based on referencing to external events.

Recent advances in silicon probe technology and two-photon imaging have ushered in an age where simultaneous, large-scale and multi-region population recordings have become more common. In the context of decision-making, which we pursue here, some studies have reported co-fluctuations between regions in the

single-trial evolution of decision-related signals<sup>3–6</sup>. Here we expand on these studies by first, significantly enlarging the number and anatomical distribution of brain regions recorded simultaneously at large-scale cellular resolution during performance of a cognitive task. And second, we take advantage of a recently described measure that estimates, based on neural activity of single trials, the presence and timing of the moment in which the subject commits to a decision<sup>7</sup> (i.e., the moment when they make up their mind). These two developments combined allowed observing two cross-brain internal signals which we now report to be substantial in magnitude and to profoundly influence decision-making neural activity. Many cellular resolution studies implicitly treat internal fluctuations as noise to be averaged away. But our results suggest that coordinated internal fluctuations reflect internal signals that are integral to cognitive processing. The measurement and elucidation of such internal signals may thus be an indispensable step towards developing an understanding of the neural basis of cognition.

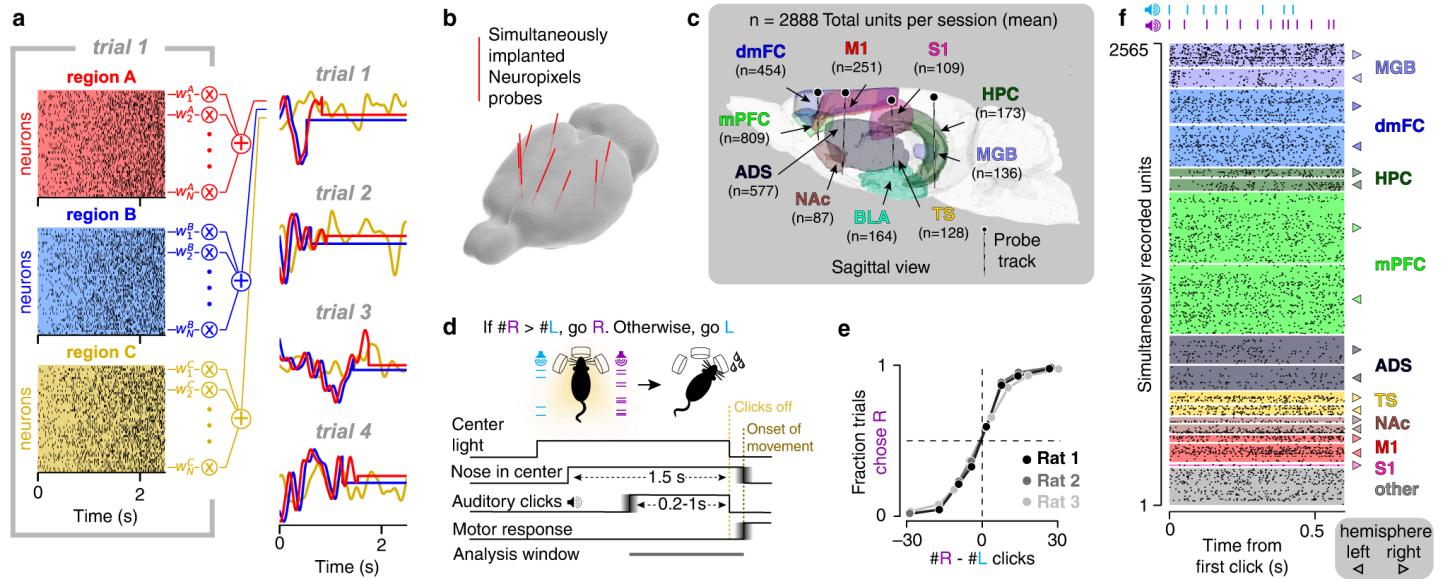
## Brain-wide simultaneous electrophysiology during perceptual decision-making

Tools for observing neural activity at scale have revealed that, during perceptual decision-making tasks, a subject's upcoming choice can be decoded from numerous and widespread cortical and subcortical brain regions (Fig. 1a)<sup>1,5,8–10</sup>. Thus perceptual decisions, like other cognitive computations<sup>11,12</sup>, are represented broadly in the activity of neuronal populations throughout the brain, but how this distributed activity is coordinated to generate a single behavioral choice is not understood. Perceptual decisions are also thought to be the result of two linked, sequential processes: evidence accumulation and decision commitment. First, because sensory evidence is often noisy or unreliable, decision formation is thought to involve integrating evidence over time to improve accuracy and performance<sup>13,14</sup>. Second, the end of decision formation is often marked by a moment in which subjects "make up their mind" in the sense of committing to a choice<sup>3,7,13</sup>. In the task we describe below<sup>15</sup>, as well as others<sup>2,16,17</sup>, decision commitment is often a covert event that cannot be inferred from a subject's overt actions. Recent work has nevertheless shown it is possible to infer the moment of covert decision commitment directly from spiking activity in large neuronal populations<sup>7</sup>.

To assay internal signals related to both evidence accumulation and decision commitment, we developed new methods for large-scale electrophysiological recording using Neuropixels 1.0 probes<sup>18</sup> chronically implanted in freely moving rats. We used brain clearing and lightsheet imaging to visualize the implanted probe tracks (Fig. 1b,c) and combined this with electrophysiological signatures to confirm probe placement. For visualization, we registered the probe tracks to the Princeton RAtlas<sup>19</sup> (Extended Data Fig. 1). For each brain hemisphere, we chose 4 penetration sites for the probes, targeting a set of cortical and subcortical regions previously shown to be involved in evidence accumulation. The penetration sites were mirrored bilaterally, for a total of 8 probe insertions per subject. Successful implantation of probes in this dense configuration required the design of novel implantation hardware, an accurate CAD model of the entire assembly in relation to the skull, and new protocols for the lengthy and invasive surgery (see **Methods** and **Extended Data Fig. 2**).

Each recording session (3 rats, 21 sessions, 6–8 sessions per rat) yielded thousands of simultaneously recorded units (median=2,888, range 2,385–4,032) from tens of brain regions (median = 19, range 13–20). Our analyses focus on a subset of the regions recorded, chosen for their known or potential involvement in the perceptual decision-making task: dorsomedial frontal cortex (dmFC)<sup>7</sup>, medial prefrontal cortex (mPFC), primary motor cortex (M1), primary somatosensory cortex (S1), anterior dorsal striatum (ADS)<sup>20</sup>, the tail of the striatum (TS)<sup>21</sup>, 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 **Extended Data Table 1**.

ADS and dmFC (the strongest cortical input source to ADS) have been directly linked, through recording and causal perturbation, to the decision formation process in the task used here<sup>7,20</sup>. dmFC is bidirectionally connected<sup>22</sup> to M1 and mPFC, and ADS receives some input from those two regions<sup>7,23</sup>. MGB is the principal relay for all auditory information reaching the forebrain<sup>24</sup>, and projects directly to TS<sup>25,26</sup>, which has



**Figure 1. Assaying coordination of decision signals across the brain.** **a**, Synthetic data depicting how multi-region neuronal population recordings allows assaying coordination of internal signals. In the left panel, each row is a neuron and each black tick is the time of a spike from the neuron. Regions A, B, and C all have firing rates that, averaged over neurons, ramp upwards after the start of each trial at  $t=0$ . But weighted sums of the firing rates of individual neurons can reveal otherwise hidden internal signals: here, the weighted sums (right panel) show that regions A and B both also represent a continuously evolving signal that is not time locked to  $t=0$  and that abruptly changes to a non-fluctuating state. Across trials, the continuously evolving signal region in A consistently leads region B (red leads blue), while the state change consistently occurs first in region B (blue goes flat before red). Region C's weighted sum signal is unrelated to that of regions A and B. These observations, which constrain the neural circuit architecture generating the signals, could not be made if the regions were not recorded simultaneously. **b**, Schematic of simultaneously implanted Neuropixels 1.0 probes in a rat brain. **c**, Simultaneous placement of 8 chronically implanted Neuropixels probes, with average recorded unit counts per region across sessions (medial prefrontal cortex, mPFC; primary motor cortex, M1; somatosensory cortex, S1; dorsomedial frontal cortex, dmFC; anterior dorsal striatum, ADS; nucleus accumbens, NAc; tail of striatum, TS; basolateral amygdala, BLA; hippocampus, HPC; medial geniculate body, MGB). **d**, Schematic of a single trial in the “Poisson Clicks” behavioral task. On each trial, at the end of two simultaneous trains of randomly timed auditory clicks, subjects are rewarded for orienting to the side that had the greater total number of clicks. Gradients represent variable timing in the task timelines. **e**, Psychometric functions. **f**, Spike raster of simultaneously recorded neural responses across the brain on a single trial.

been shown to be necessary for auditory decision-making<sup>27,28</sup>. S1 is reciprocally connected to M1<sup>29</sup>, projects broadly within striatum<sup>23</sup>, and has recently been established to play a causal role in the formation of perceptual decisions<sup>30</sup>. Decision-related neurons have also been discovered in HPC and NAc during accumulation-of-evidence tasks<sup>7,31</sup>. BLA contains many auditory-responsive neurons<sup>32</sup> and projects broadly throughout the striatum<sup>23</sup>.

Before the implantation surgery, the subjects were trained to perform an established auditory decision-making task requiring accumulation of evidence presented in the form of randomly-timed pulsatile auditory clicks<sup>15, 15,33</sup> (the “Poisson Clicks” task; **Fig. 1d**). This paradigm allows precise control of the stimulus evidence contributing to the subject’s decision and inferring, from each region on each trial, the state of the evolving decision. Subjects initiate a trial by poking their nose into a center nose port. While the animal has its nose in the center port, and after a variable delay (0.5 - 1.3 s), two randomly-timed trains of broadband auditory clicks are played from speakers to the subject’s left and right, until a visual “go” cue, 1.5 s after the initial center poke. The subject is then free to withdraw from the center port and poke its nose into one of two side ports. Reward is delivered if the subject pokes into the side that played the greater total number of clicks. The optimal strategy thus requires gradually accumulating clicks over time on each trial. After surgery, all three rats continued to exhibit good behavioral performance. Two consistently maintained low lapse rates (i.e., the error rate on the easiest trials), each with a median lapse rate of 2% and maximum of 5% in each session, and completed a large number of trials (median 623 per session; range: 370-806). The third rat exhibited greater variability across sessions; we excluded from analysis 4 of 14 sessions with lapse >8% and 2 sessions with fewer than 300 trials (**Fig. 1e**; **Extended Data Fig. 2f**).

## Evidence for a single brain-wide decision formation computation

Using these brain-wide recordings during perceptual decision-making, we first used an approach established in previous studies of decision formation, which leveraged large-scale recordings of many neurons within a single region (sometimes more regions) to estimate the moment-by-moment evolution of gradually emerging decision signals<sup>34-37</sup>. We reasoned that these simultaneous, brain-wide measures of the evolving decision state would allow us to distinguish hypotheses regarding how activity is coordinated across regions to generate a coherent decision: namely, do different subcircuits reflect independent processes that compete for behavioral expression, as has been proposed for other behaviors<sup>38,39</sup>? Or, do decision-related signals in all these regions reflect a single shared computation? (**Fig. 2a**)

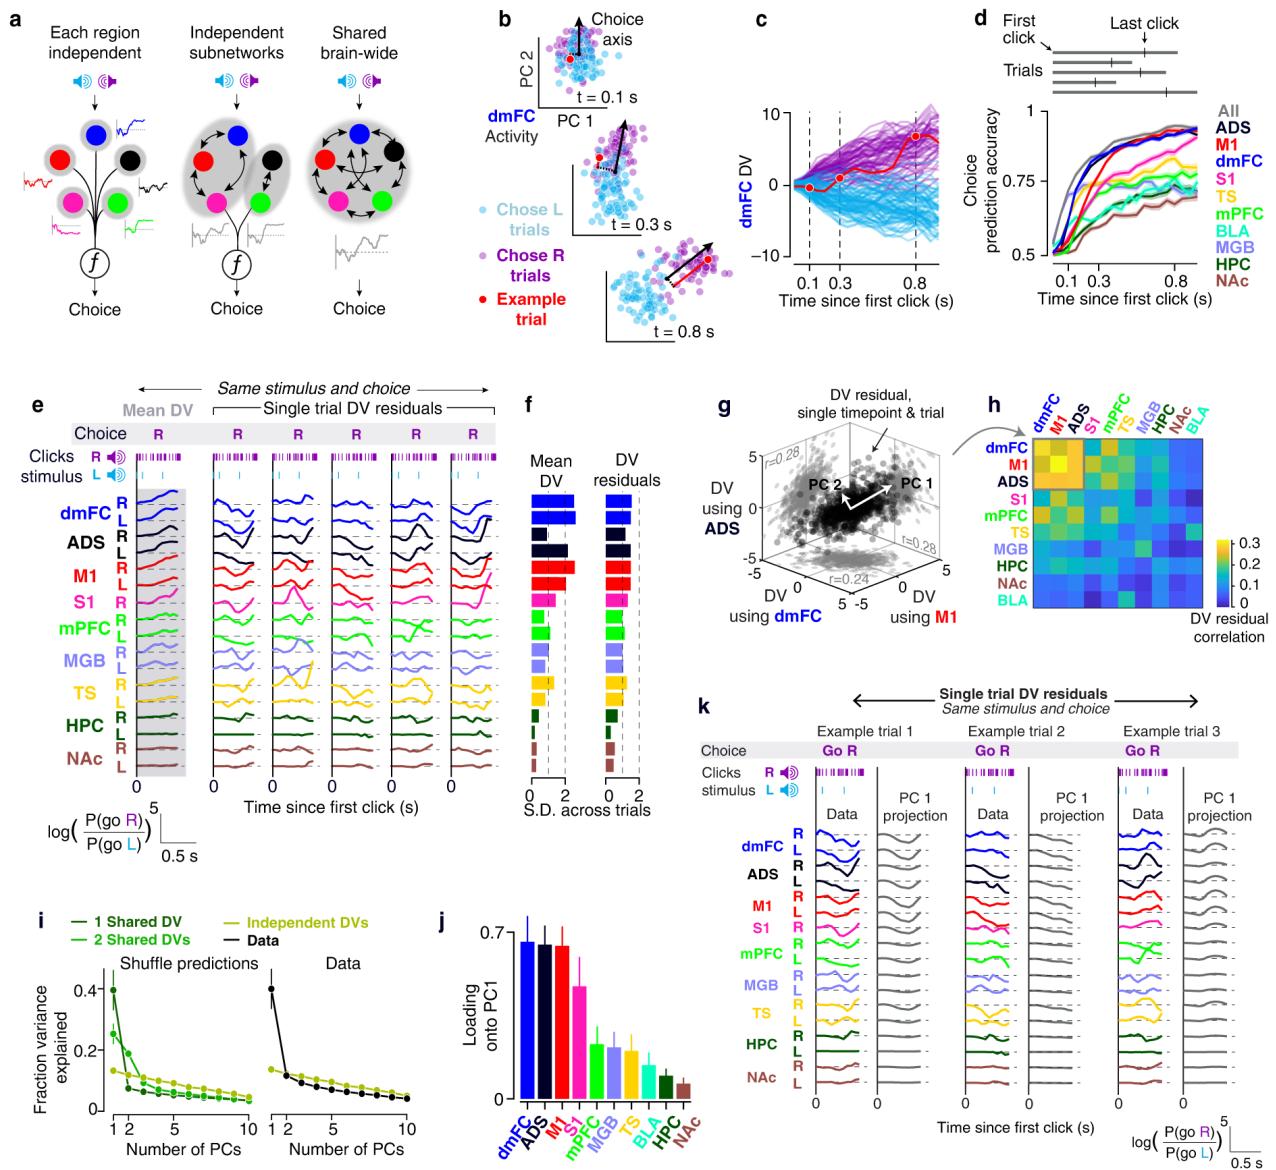
The approach takes firing rate data from all neurons in a region, time-aligned on each trial to the start of the sensory stimulus, and then uses logistic regression to determine, for each time point, the weighted sum of neural firing rates that best predicts, across all trials, the animal's choice (**Fig. 2b**). On each trial, the weighted sum yields a time-varying scalar variable referred to as the "decision variable",  $DV(t)$ <sup>35</sup> (**Fig. 2c**; see **Methods**). For each region,  $DV(t)$  can be thought of as a 1-dimensional summary of decision-related activity in that region, and its evolution over time can be traced for each trial. In the formulation used here, the more positive (negative) the  $DV(t)$ , the more confident the logistic regression model is that the subject will decide to go right (left). The trial-averaged probability that the logistic regression model, at a given time in the trial, would accurately predict the subject's behavioral choices can be computed from the  $DV(t)$ 's for each trial, and is the "choice prediction accuracy" (**Fig. 2d**). Scaling the approach up to simultaneously study many regions across the brain<sup>4,5</sup>, and consistent with the widespread representation of the evolving decision observed by others<sup>1,2,5,36</sup>, we found that for all recorded regions, the trial-averaged choice prediction accuracy was significantly greater than chance and increased gradually over time (**Fig. 2d**). The choice prediction accuracy for each region was similar across the left and right hemispheres (**Extended Data Fig. 3a**).

A group of two frontal cortical regions (M1 and dmFC), together with anterior dorsal striatum (ADS), showed the highest choice prediction accuracy. While choice prediction accuracy showed a dependence on population size (**Extended Data Fig. 3b**), these three regions did not typically yield the largest populations (**Extended Data Fig. 3D**), indicating that choice prediction accuracy differences between regions were not trivially due to population size differences amongst regions.

Simultaneous recordings allow the observation of how  $DV(t)$ 's evolve on single trials, not only within a region, but across all recorded regions. In a subset of 5 sessions distributed across all 3 rats, we used a "frozen noise" task design in which only 54 fixed click trains were presented, in random order across trials; each of the 54 trains was thus repeated over multiple trials (median of 10) within each of five sessions. This allowed assessing how the variable evolution of the  $DV(t)$ 's across trials was correlated across regions, independent of shared variability trivially driven by a changing stimulus. Within groups of trials with the same click train, multiple trials with the same choice further allowed us to assess correlations in the  $DV(t)$ 's across regions independent of those introduced by signals yoked to choice, which may reflect shared motor preparatory signals not of central interest here.

The left column of **Fig. 2e** shows the temporal evolution of  $DV(t)$ 's from 9 brain regions, in each of the two hemispheres (18 regions total), averaged over repeats of one of these 54 click trains, which is shown at the top. For this click train, there were many more Right clicks than Left clicks, making this an easy "go Right" trial. On average for this stimulus, the  $DV(t)$ 's in all the brain regions gradually grew increasingly positive, corresponding to an increasingly strong signal predictive of a Right choice. But how did  $DV(t)$ 's across the brain evolve on single trials?

The five columns to the right show the temporal evolution of the  $DV(t)$ 's in each of five representative single trials that were part of the average to their left. The plots show the single-trial residuals around the left column's average (that is, with the stimulus- and choice-conditioned mean in the left column subtracted). Several striking features, typical of the data as a whole, are immediately apparent: First, the residuals are substantial in magnitude. These are not small fluctuations around the mean, but instead have magnitudes



**Figure 2. Single-trial decision variable residuals reveal a single, shared decision process.** **a.** Competing hypotheses regarding the degree of independence of the decision formation processes of each region, assessed using the DV(t)'s, for a set of regions, all of which have decision-predictive activity. Colored circles indicate individual regions; larger gray ellipses demarcate independent networks. Traces depict a schematic time course of decision-related activity. The function  $f$  represents an unknown mechanism for combining decision-related activity from different regions to generate a behavioral choice. **b.** Single trial dmFC population activity projected onto the first two principal components of the set of choice axis directions obtained across time points, shown separately at three time points relative to the first click. Note that firing rates were obtained by smoothing the spike times, so data from neighboring time points are not independent (see **Methods**). Blue (purple) dots depict neural activity on chose L (chose R) trials. Black arrows indicate the weight vectors that best predict the subject's choice using logistic regression, whose direction defines the "choice axis". Red lines indicate projection onto the weight vector (the "decision variable") for an example trial. Data shown for an example session. **c.** Decision variable time course – DV(t) – across all trials; the example trial is highlighted in red. **d.** Time course of choice prediction accuracy using neural activity from each brain region bilaterally and from all regions combined. Mean  $\pm$  1 s.e. across sessions shown. **e.** Simultaneously obtained DV(t)'s for all recorded regions in an example "frozen noise" session, shown for a subset of trials that shared the same click train and choice. The left column shows the average DV(t) across this set of trials for each region, and the remaining columns show the residual fluctuations (DV(t) residuals) for five example trials around that mean. **f.** Standard deviation of the mean and residual (i.e. stimulus- and choice-conditioned mean subtracted) DV(t)'s across all trials in the session whose data is shown in b, c and e. **g.** State space representation of the coordinated variability in the moment-to-moment, single-trial DV(t) residuals across 3 brain regions. Each axis corresponds to the DV(t) residuals across timepoints and trials computed using firing rates from a different region. Pearson correlations between the data shown for each pair. Units are log-odds; shown for an example session. Arrows shown indicating directions explaining maximum variance. To focus on the trial period when the subjects are still actively forming their decision, only trial timepoints before the inferred moment of decision commitment using M1 spikes (nTc-M1) are included (see **Methods**). **h.** Matrix of DV(t) residual correlations (Pearson's  $\rho$ ) between pairs of regions, across timepoints and trials within a session. Average across the frozen noise sessions (and across the two hemispheres per session) shown. Diagonal entries correspond to the inter-hemispheric correlation for a given region. **i.** Fraction of variance explained per principal

component (PC) of the brain-wide residual DV(t)'s. The left panel displays predictions from the three competing models in (a), and the right panel compares the recorded data with the “Independent DVs” prediction, which provides an estimate of the noise floor (see **Methods**). Error bars indicate 1 s.e. obtained from a hierarchical bootstrap across 5 frozen noise sessions. **j**, Plot showing the loading (averaged across hemispheres and sessions) of each region onto the first PC. Error bars indicate 1 s.e. obtained from a hierarchical bootstrap across 5 frozen noise sessions. **k**, DV(t) residuals for each region and their projection onto PC1 for three example trials with identical click trains and behavioral choices (same set of trials as in Fig. 2e).

comparable to the mean itself (**Fig. 2f**). Across regions, the standard deviation of the DV(t) residuals was similar in magnitude to the average DV(t) traces (99% as large as, +/- 5% s.e. bootstrap across sessions). Second, the fluctuations are significantly correlated across regions, with different time courses across trials but similar time courses across regions within a trial. We quantify this aspect below, but it is strong enough that in the example of **Fig. 2e** it can be seen by simple visual inspection. Importantly, the co-fluctuations across regions are not due to shared fluctuations in the stimulus driving the DV(t)'s or in the subject's final choice, as here the stimulus is held constant across trials, and all trials used in **Fig. 2e** resulted in the same choice (Right). Whatever is being so markedly shared across brain regions appears to reflect internally generated, moment-by-moment variability not determined by the sensory stimulus, nor the overt choice report. Within single regions, DV(t) fluctuations have previously been linked to meaningful aspects of the subject's internal decision formation process, such as changes of mind<sup>34</sup>. Indeed, for sets of trials in which the stimulus was held fixed but the animal's choice varied across trials, we found that we could predict choice from the DV(t) fluctuations well above chance for all recorded regions, even though the mean DV(t) for each given click train had been removed (**Extended Data Fig. 4**). This confirms that these apparently internally-generated signals are related to the subject's decision-making process. If regions are independent—given a particular stimulus and choice—then the residuals would not be expected to co-fluctuate. If regions form  $N$  mutually-independent groups, then we would expect co-fluctuations within but not across groups. Finally, if regions form a single group representing a single shared computation, we would expect co-fluctuations shared across all regions. More generally, separately from how computations are mapped onto subcircuits, the number of state variables underlying the neural activity places an upper bound on the number of computations being carried out. The analysis and differentiation of hypotheses are strengthened by a task design with multiple repeats of a given stimulus<sup>40–43</sup>, as used in **Fig. 2e**, to enable controlling for the co-fluctuations that stimulus and choice can themselves generate. This is done by assessing residual fluctuations around the mean trajectory, for each set of trials with a given stimulus and a given choice. As we show in depth later, subjects in our task do not accumulate evidence over the entire stimulus presentation, but instead often commit to a choice before the end of the full click train. To focus on the trial period when the subjects are still actively forming their decision, we confine the present analysis of inter-regional co-fluctuations in the DV(t) residuals to time points that occur before the estimated moment of decision commitment on each trial. We took the set of DV(t) residuals over time and considered each time point in each trial as a point in high-dimensional space, where each axis corresponds to the DV(t) residuals of one brain region (pooled across hemispheres). **Fig. 2g** illustrates three of those dimensions for an example session. Pairwise correlations between any two given regions are shown in **Fig. 2g** as correlation coefficients on the corresponding two-dimensional “face” of the space spanned by each region's DV(t) residuals. The set of all pairwise correlations across regions is shown in **Fig. 2h**. As can be seen, residual DV(t) correlations between most pairs of regions were well above zero, with the correlations between three regions (dmFC, ADS and M1) clearly standing out from the rest as strongest ( $\rho \approx 0.3$ ). This pattern of inter-regional correlations was highly specific to activity along each region's choice axis: correlations between other one-dimensional summaries of population activity, such as mean firing rate, were much lower overall and did not share the same structure (**Extended Data Fig. 5**). Considering the left and right hemispheres separately (**Extended Data Fig. 3c**) revealed strikingly similar patterns of correlation across hemispheres, demonstrating a remarkable degree of inter-hemispheric coordination of decision formation. Given the similarity across hemispheres, we pooled across the two hemispheres for subsequent analyses, from **Fig. 3** onwards.

Patterns of co-fluctuations across regions can be characterized by the principal components (PCs) and PC spectrum of the data in the space illustrated in **Fig. 2g**. We first shuffled region labels across neurons, to

generate surrogate data in which fluctuations in all regions are equivalent to each other. In this surrogate data, a single shared DV(t) describes co-fluctuations across all regions, and a single principal component will dominate the PC spectrum (**Fig. 2i left**, dark green; see **Methods**). Next, we divided regions into two groups and shuffled trial labels within but not across groups. This shuffle corresponds to the hypothesis that co-fluctuations across regions are described by two separate DV(t)'s (**Fig. 2i left**, medium green). Finally, we shuffled trial labels independently within each region to simulate the null hypothesis that all regions are independent of each other. While this would produce a flat PC spectrum for infinite numbers of trials, here, with finite data, the results of this shuffle can be thought of as the noise floor, i.e. what we expect from independent regions but finite data (**Fig. 2i left**, lightest green).

Applying the analysis to the experimental data (**Fig. 2i right**) revealed that the data's PC spectrum was remarkably similar to the prediction from the hypothesis with a single DV(t) across regions. The first principal component alone accounted for 40% of the variance in the dataset, which was 3.8 times ([2.5-6.4] 95% coverage interval using a hierarchical bootstrap test across sessions) greater than the next principal component, which fell below the noise floor. The prominence of the first PC in the data, consistent with the prediction given a single DV(t), was apparent in each individual frozen noise session (**Extended Data Fig. 6**). While this does not preclude the possibility that a larger dataset would allow further principal components to rise above the noise floor, a single component clearly dominates strongly above all others, implying that brain-wide co-fluctuations in DV(t) residuals can be well described by fluctuations in a single scalar variable, namely the projection onto that first PC. This single scalar variable, reflected in each region's DV(t) residuals, is consistent with a single, internally-generated decision formation process, shared across regions.

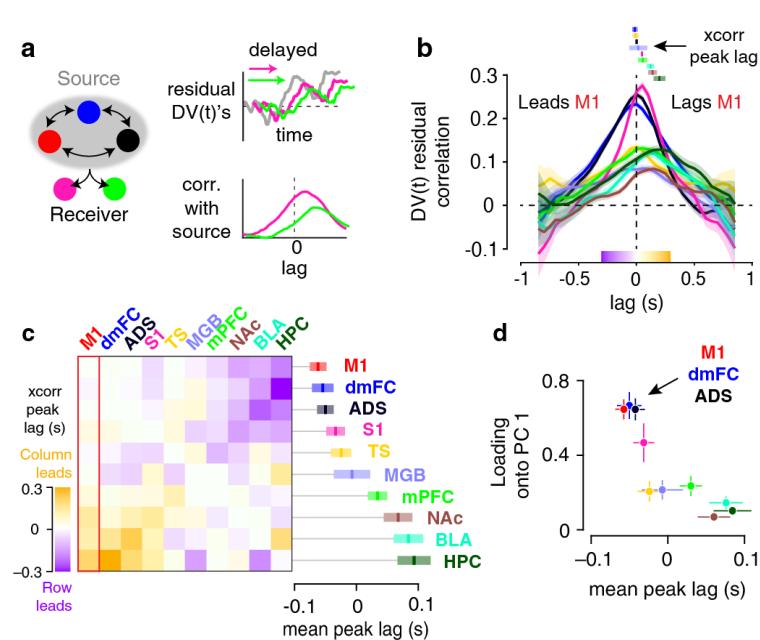
## A frontostriatal subnetwork temporally leads decision formation

Three regions (dmFC, ADS, and M1) had the strongest loading onto the first PC (**Fig. 2j**), meaning that the DV(t) residuals in these regions covaried most strongly with the single, dominant mode. In **Fig. 2k** we show the DV(t) residuals across each brain region for three example trials, and the projection of the brain-wide residuals onto the first PC. This illustrates how this single dimension captures a large fraction of the fluctuations across brain regions, and how the DV(t) residuals in dmFC, ADS and M1 most strongly covary with this projection. Recall that dmFC and M1 are reciprocally connected frontal cortical regions, and project directly to ADS. ADS, in turn, likely projects indirectly to these cortical regions through the cortico-striatal-thalamic loop<sup>44</sup>. The strength with which these 3 brain regions represent the brain's single, shared DV(t) suggests that the decision formation process shared across the brain may emerge within this anatomically-interconnected frontostriatal subnetwork, a hypothesis we explicitly test below.

If DV(t) residuals reflect the ongoing decision process, as their link to behavior suggests (**Extended Data Fig. 4**), and if those signals originate in a subset of regions and are then relayed to other regions, we would expect the single-trial fluctuations in the DV(t) residuals of "receiver" regions to lag those in the "source" regions (**Fig. 3a**). Comparing the latency with which trial-averaged decision-related activity appears across regions has often been used to infer where decisions are first formed<sup>2,5,10,45</sup>. However, the computation could be initially triggered in one set of regions and yet be subsequently led by others. Here, we leverage our simultaneous recordings, along with the prolonged decision formation period afforded by our task, to probe temporal relationships among the ongoing internally-generated decision formation signals that we found to be so prominent on single trials, and having removed the component of decision-related activity driven by the stimulus or yoked to the upcoming response. We computed, for each pair of regions, the cross-correlation function of their DV(t) residuals, by measuring DV(t) residual correlations at different time lags between the two regions. **Fig. 3b** shows the resulting cross-correlation function between DV(t) residuals in primary motor cortex (M1) and all other regions. We found that they peak near zero and decay to negligible values within about 0.5 s, reflecting relatively short timescale co-fluctuations, less than the duration of a trial. The peaks of the cross-correlation functions are shifted to the right of zero for most pairs, indicating that M1 DV(t) fluctuations

lead those of most other regions. Note again that because we are using DV(t) residuals, these correlations exclude those introduced by shared coding for stimulus or choice.

The left panel of **Fig. 3c** shows the results of the peak time lag analysis of **Fig. 3b**, but now for all pairs of regions. (The first column contains the data from **Fig. 3b**). Each region's peak lag, averaged over all other regions, is shown in the right panel of **Fig. 3c**. This shows that during decision formation, dmFC, ADS and M1 – the same set of three interconnected frontostriatal regions that most strongly represented the shared internal decision process – tended to have DV(t) fluctuations that led the rest of the brain by the largest amount (~50 ms), and were statistically indistinguishable from each other in this measure ( $p>0.05$  in all cases, using hierarchical bootstrap test across sessions). We found that, across all regions, the average peak lag was highly correlated with the loading of each region onto the first PC (**Fig. 3d**). In other words, the strength with which the shared decision process is represented is closely correlated with the latency with which it is observed. Taken together, these results show that decision formation is a single, shared brain-wide process, and within the set of regions recorded here, it is led by dmFC, M1, ADS.



**Figure 3. A frontostriatal subnetwork temporally leads decision formation.** **a**, Schematic of the motivating problem, namely identifying where signals related to decision formation originate. The residual DV(t)'s in source regions are predicted to lead those in recipient regions, leading to shifted cross-correlogram peaks. **b**, Pearson correlation between the residual DV(t)'s of each brain region with M1 as a function of different relative time lags, with regions considered bilaterally. Average across trials across all frozen noise sessions is shown. The time lag at which the maximum value occurs (the cross-correlogram peak lag) is shown at the top for each pair, with error bars indicating 1 s.e. obtained using a bootstrap across trials. **c**, At left, lag of the cross-correlogram peaks, computed as in panel b, but here shown for all pairs of brain regions. The first column corresponds to the cross-correlogram peaks in (b). Negative values (purple) indicate that the region on the row leads the region on the column. Note that the matrix is therefore antisymmetric by construction. At right, the average of each row in the matrix, with error bars indicating 1 s.e. obtained by bootstrap across trials. Negative values indicate a region's residual DV(t) tends to lead other regions. **d**, Scatter plot of each region's residual DV(t) loading on PC1 (as in from Fig. 2e) versus its mean peak lag in (c).

## Brain-wide estimates of the time of covert decision commitment

Decision formation comes to an end when the subject commits to a choice (“makes up their mind”). The timing of this moment can be driven by processes that appear to be internally generated and not timelocked to external events such as a stimulus or an overt motor action. Recent work found that, during a subset of trials in the Poisson Clicks task, subjects covertly commit to a choice before the end of the stimulus presentation, disregarding further evidence. This moment, referred to as the “neurally-inferred time of commitment” or “nTc” occurs at a precise time point—not time-locked to the stimulus or motor report—that can be detected and identified from the neural activity of tens of simultaneously recorded choice-selective neurons from frontal cortex and striatum. Here, the brain-wide simultaneous recordings present an opportunity to determine how this internal state change manifests at the level of the whole brain.

We first describe the method of nTc estimation and validate that the main conclusions from the study presenting nTc<sup>7</sup> also apply to the current dataset, which includes a much broader range of brain regions. We then examine the relationship between nTc estimates from different regions. We find that, as with the DV(t)'s of the previous section, nTc estimates are highly correlated across the brain.

The existing method to identify nTc relies on a model of spiking data that is an expansion of the classic behavioral drift-diffusion model (DDM)<sup>13,46</sup> of accumulation of evidence for decision-making. In the DDM, rightwards (leftward) evidence adds (subtracts) to a scalar latent variable  $z$ , such that as  $z$  evolves, it represents the evidence accumulated over time. Noise corrupts this process, and in addition, if  $|z|$  reaches a bound  $B$ , the subject is taken to have committed to the decision indicated by the sign of  $z$ . The “multi-mode” drift-diffusion model<sup>7</sup> used to infer nTc (MMDDM; **Fig. 4a**) maps the DDM onto spiking data in a manner that, critically, is different during the decision formation phase (i.e., before hitting the bound) versus the committed phase (after reaching the bound). Neuron  $n$ ’s firing rate depends on two components added together: a time-varying term  $b^n(t)$  that captures the overall temporal shape of the choice-independent firing rate of the neuron and a decision-related term, formed by the DDM’s latent variable  $z$  multiplied by a weight parameter  $w^n$ . The sum is passed through a rectifying nonlinearity to produce the neuron’s firing rate (**Fig. 4a**). Critically, the neuron’s parameter  $w^n$ , which is fitted to the data, is allowed to be different before hitting the bound ( $w_{EA}^n$ , evidence accumulation phase) versus after ( $w_{DC}^n$ , decision commitment phase): this implies that the model firing rates lie in a different direction in neural space (a different “mode”) before versus after decision commitment, abruptly shifting from one to the other at the time of commitment. This abrupt mode change serves as a precise signature of commitment timing. The model’s parameters were fit to maximize the probability of observing the experimentally-observed spiking and choice data under the model.

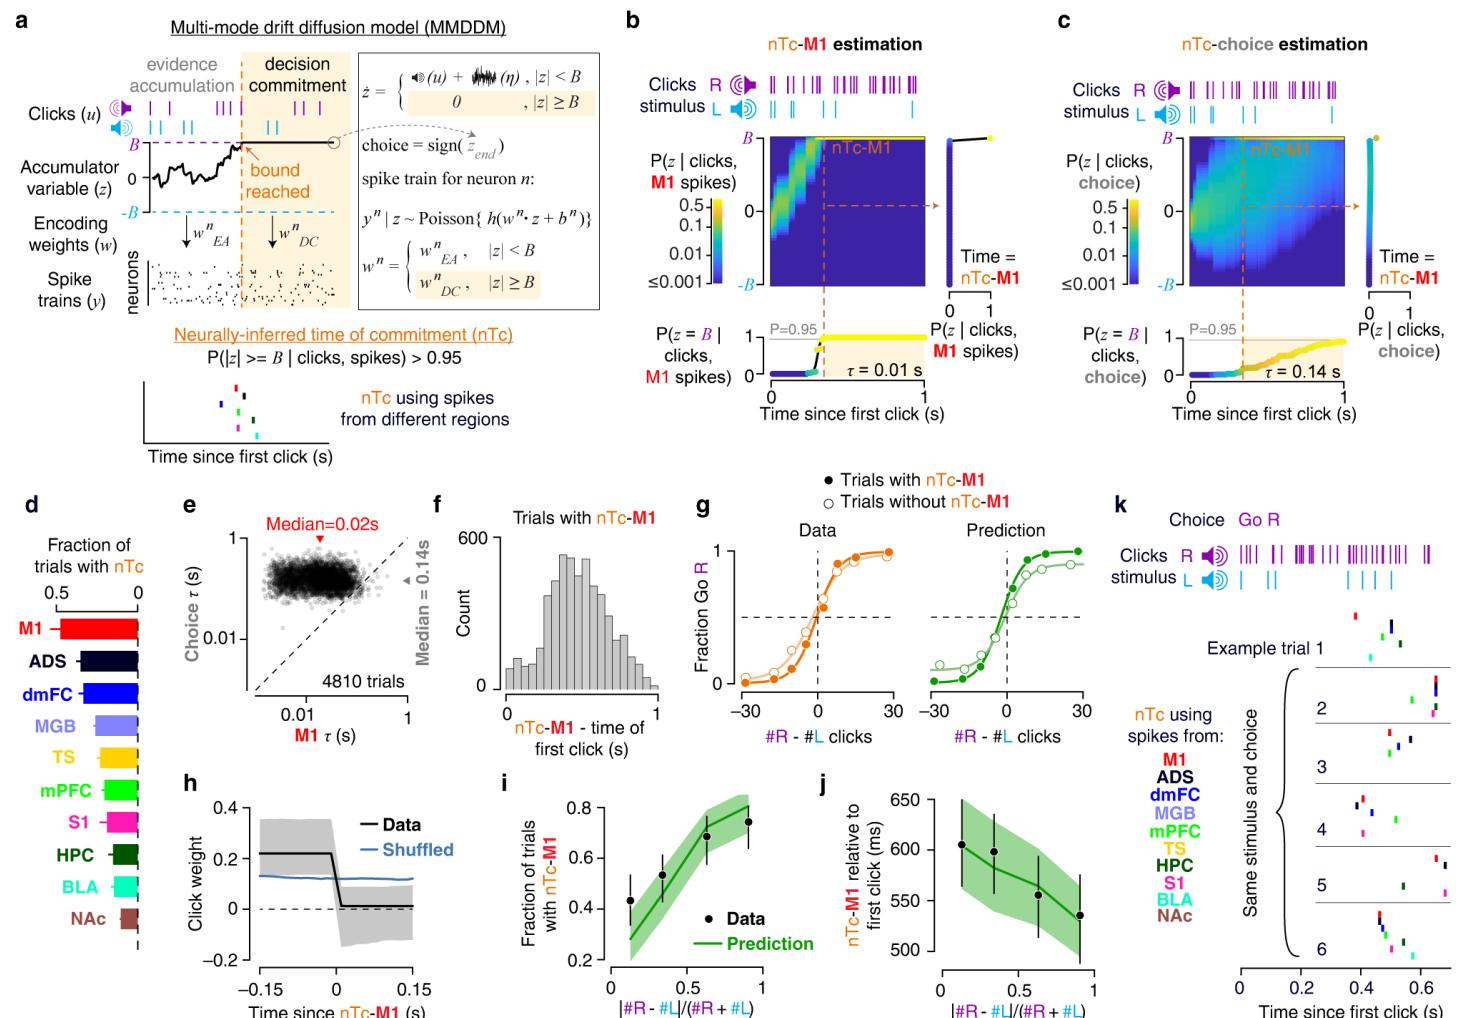
When  $w_{EA}^n$  and  $w_{DC}^n$  are constrained to be equal, the model reduces to a straightforward map of the DDM onto firing rates along a single mode in neural space. Ref. <sup>7</sup> showed that for anterior regions, MMDDM fit cross-validated experimental data better than the single-mode DDM, and was also better able to account for heterogeneity of single neuron temporal profiles. This held true in the current dataset, including posterior regions not studied in ref. <sup>7</sup> (**Extended Data Fig. 7**). The mode shift, rather than a global change in firing rate (**Extended Data Fig. 8a**), defines commitment time.

After fitting the model’s parameters, which are then fixed and the same for all trials, we computed, for each trial, the posterior distribution of the latent variable  $z$  as a function of time during the trial, given the entire trial’s click train and spiking activity from each individual region. **Fig. 4b** shows the temporal evolution of this posterior distribution, for one particular trial and using neural activity from region M1. The bottom panel shows the evolution of the posterior probability of reaching the commitment bound; if this reached 0.95 before the subject began its choice-reporting movement and remained above this threshold, then the earliest time this threshold was reached was defined as nTc for that trial. **Fig. 4c** shows the same estimates for the same trial, when the subject’s behavioral choice was used instead of the neural spiking data to constrain the posterior distributions. As in ref. <sup>7</sup>, the spiking data produces a far more precise estimate (compare **Fig. 4b** and **4c**; **Fig. 4e**).

Considering all recorded brain regions, nTc could be detected most frequently in M1 (48% of trials), ADS (35%), and dmFC (33%; **Fig. 4d**), indicating that decision commitment signals are strongest in the same set of brain regions that most robustly represent the shared process of decision formation (**Fig. 2**). Since nTc could be detected most reliably from M1, we focused on the M1-derived estimates (“nTc-M1”) for the remainder of the paper. Consistent with findings in ref <sup>7</sup>, nTc-M1 was broadly distributed with respect to stimulus onset (**Fig. 4f**; s.t.d.=196ms), movement onset (**Fig. 6a**; s.t.d.=239 ms) and stimulus offset (**Extended Data Fig. 8b**; s.t.d.=225 ms). Moreover, again confirming findings in ref. <sup>7</sup> and as predicted for a marker for the moment when a subject “makes up their mind”, the influence of the auditory stimulus clicks on the subject’s behavioral choice ended abruptly after nTc-M1 (**Fig. 4h; Extended Data Fig. 8c**), although this could not be consistently observed for nTc estimates from other brain regions (**Extended Data Fig. 8d**). Consistent with predictions from the MMDDM model fits, trials with an nTc-M1 had higher behavioral accuracy than those without, even when controlling for evidence strength (**Fig. 4g**). Additionally, nTc-M1 occurred more frequently (**Fig. 4i**; logistic regression coefficient=1.8, 95%CI [1.6, 2.0],  $t(3587)=16.0$ ,  $p=9 \times 10^{-58}$ ) and earlier (**Fig. 4j**; linear regression

coefficient=-0.05, 95%CI [-0.08, -0.02],  $t(2170)=-3.6$ ,  $p=3 \times 10^{-4}$ ) on easier trials than harder trials. We note that here we do not treat MMDDM as a generative model underlying the DV(t)'s, which may emerge from more complex dynamics, but only as a tool for inferring commitment timing.

Our brain-wide recordings allowed us to simultaneously detect nTc from a large number of brain regions, across the same set of trials. As with the DV(t) signals related to gradual decision formation, we found that even when considering trials with identical click trains and behavioral choices, the time at which nTc occurred varied widely across trials, indicating that the particular timing on each trial was internally generated and not determined by the click stimuli or the subject's choice. Moreover, and again as with DV(t)s, this variability was strikingly correlated across regions, which we quantify below but which is visually apparent on example trials (Fig. 4k). This demonstrates a remarkable coordination in signaling of covert decision commitment across brain regions.

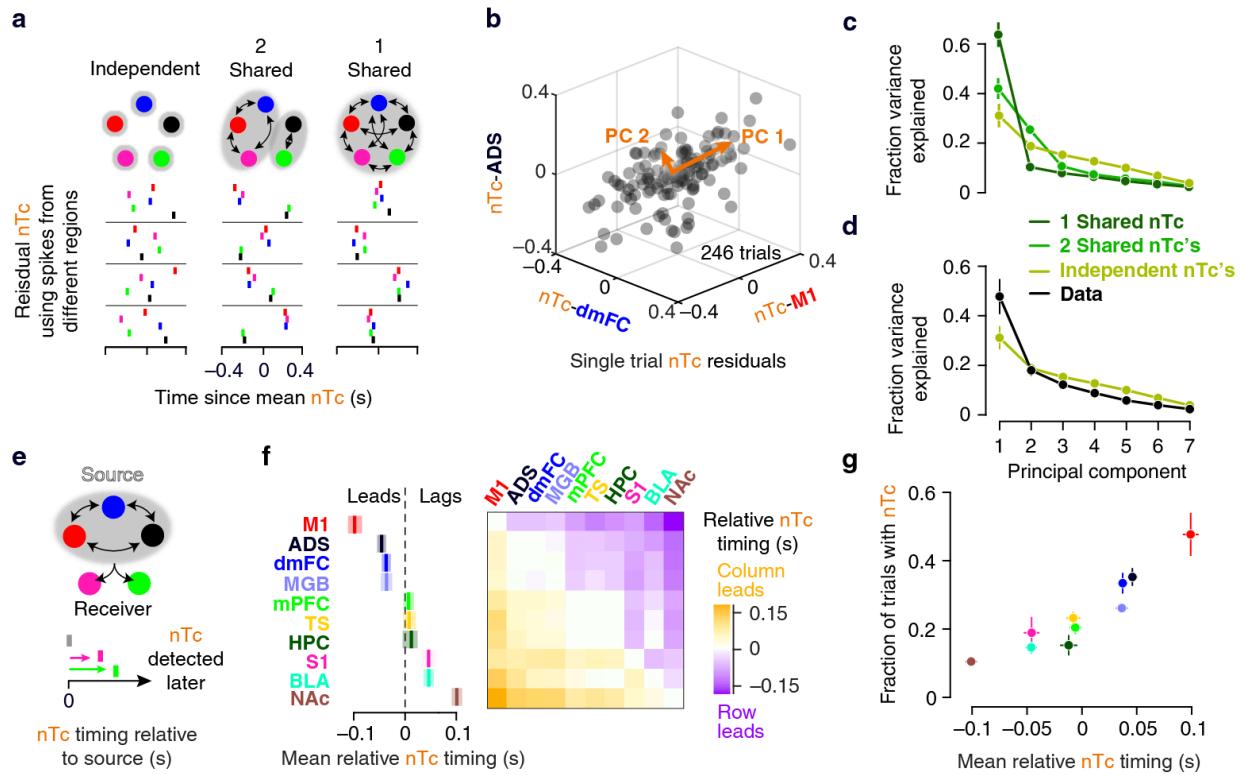


**Figure 4. Brain-wide estimates of the time of covert decision commitment.** **a**, Schematic illustrating the procedure used to infer region-specific neurally inferred time of commitment (nTc) based on a model assuming accumulation to a decision bound with an abrupt change in neural encoding at commitment. **b**, Posterior probability of commitment over time for an example trial, computed using M1 spike data. **c**, Same trial and format as b, with posterior probability computed using choice. **d**, Detection rates of nTc across brain regions. Error bars indicate bootstrapped standard error computed from across 21 sessions. **e**, Time constants of sigmoid curves fitted to the posterior probability traces. **f**, Distribution of nTc across trials on which detection was possible using M1 spikes.  $N=5,135$ . **g**, Psychometric curves on trials with (n=5,135) and without (n=5,632) detectable nTc-M1, along with corresponding curves from simulated data. **h**, Psychophysical kernel showing the influence of stimulus clicks on choice before and after nTc-M1.  $N=2,544$  trials. Shading indicates 95% bootstrapped confidence intervals. **i-j**, Detection rates (**h**) and timing (**i**) for neurally-inferred commitment times (nTc-M1) across levels of stimulus evidence strength. Error bars and shading indicate 95% bootstrapped confidence intervals across 21 sessions. **k**, Colored dots indicate nTc's estimated from different brain regions on six example trials with identical click trains and choices, highlighting correlated variability in commitment timing across regions.

## Decision commitment is shared across the brain and is led by M1

Similar to our previous analysis of the DV(t)s – which were used to constrain models of the multi-region circuits underlying decision formation – how nTc estimates from different regions relate to each other provides similar constraints to how decision commitment arises from distributed circuits. We first assessed whether the correlated nTc timing we observed across regions is consistent with a single, shared process of decision commitment or with multiple independent processes (Fig. 5a). We again took advantage of the frozen noise task design to quantify correlations that by definition, cannot be explained by shared coding for the stimulus or choice-specific preparatory signals. For each region, we separated trials into groups with identical click trains and choices, computed the mean nTc within each group, and subtracted it from the corresponding trials. We then examined how the remaining “nTc residuals” were coordinated across regions. We considered the brain-wide nTc residuals as points in a space whose axes are the nTc residuals from each region (Fig. 5b). The dimensionality of this point cloud serves as a quantitative measure of the number of latent processes generating the brain-wide nTc’s, which we confirmed using MMDDM simulations. (Fig. 5c; see **Methods**). Because nTc was not detected on every trial from each region, we used probabilistic PCA (PPCA) to assess dimensionality in the presence of missing data (see **Methods**). We found that one dimension explained approximately 50% of the variance of the residual nTc’s, and the rest of the dimensions explained a fraction of variance that was below the noise floor (Fig. 5d). While this was most consistent with simulations that assumed a single, shared moment of decision commitment, we do note that the dominance of the first principal component here is not as strong as for the DV(t)’s in Fig. 2b. This may be due to the fact that nTc provides only one data point per trial, whereas the DV(t) residuals provide many points per trial, implying less statistical power in our ability to distinguish a 1-dimensional from a 2-dimensional system.

Next, similar to our approach using DV(t) residuals to assess the origin of gradual decision formation signals, we measured lead-lag relationships between nTc’s from different regions (Fig. 5e) to assess where the highly coordinated decision commitment signal first emerges. We calculated the average difference in nTc timing between each pair of regions, using only those trials in which nTc was detected from both regions. We then summarize these pairwise measurements for each region by calculating the average nTc timing difference of that region with all others (Fig. 5f). As with the earlier finding that nTc could be detected most frequently from M1, followed by ADS and dmFC (Fig. 4d), this timing analysis revealed that decision commitment could be first detected from M1, with nTc-M1 leading nTc obtained from all other regions by approximately 100 ms, on average (Fig. 5f, left; **Extended Data Fig. 8b**). Decision commitment signals in M1 were most closely followed by ADS and dmFC. Across regions, here was a strong correlation (Pearson’s  $r=0.93$ ) between nTc detection frequency and nTc timing (Fig. 5g). This is predicted under the model used to infer nTc (MMDDM), where stronger encoding of decision commitment provides a higher signal-to-noise ratio for determining whether and when the latent variable  $z$  reaches the decision bound. More robust decoding of bound-crossing leads to both a greater fraction of trials in which nTc could be detected (i.e., the posterior probability of  $z$  at the bound reaches 0.95 on a given trial) and earlier estimates. In sum, analyses of nTc timing across regions reveal a low dimensional, highly coordinated representation consistent with a single, shared decision commitment process that can be detected earliest in M1.



**Figure 5. Decision commitment is shared across the brain and is led by M1.** **a**, Schematic of competing hypotheses: the timing of decision commitment may be independent across brain regions or reflect one or more shared processes. Rows show four example trials with the same click train and behavioral choice. Tick marks indicate single-trial residual nTc for each region, relative to the average nTc for trials with that click train and choice. **b**, State-space representation of the coordinated variability in single-trial nTc residuals across 3 brain regions, shown for one session with 246 trials that had nTc estimates for all three regions. Orange arrows indicate directions explaining the maximum variance. **c**, Plot of variance in across-region nTc residuals explained by each principal component (PC), as predicted by the competing hypotheses in (a) and tested using MMDDM simulations (see **Methods**). Error bars indicate 1 s.e. using a hierarchical bootstrap across sessions. **d**, Same as (c) but using the observed data. The prediction assuming independence across regions is reproduced for comparison, as it provides a measure of the noise floor. Error bars indicate 1 s.e. using a hierarchical bootstrap across sessions. **e**, Schematic of the approach for identifying where decision commitment signals first emerge. **f**, Matrix representing the trial-averaged difference in nTc timing between all pairs of brain regions. Left, each region's average difference relative to all others, corresponding to the mean of each row in the matrix on the right. Each pairwise comparison includes only trials with nTc detected in both regions. Negative values indicate that a region's nTc tends to lead other regions. **g**, Scatter plot of the mean relative nTc timing from (f) the fraction of trials in which nTc could be detected from each region. **f-g**, Values indicates the mean +/- 1 s.e. across 21 sessions.

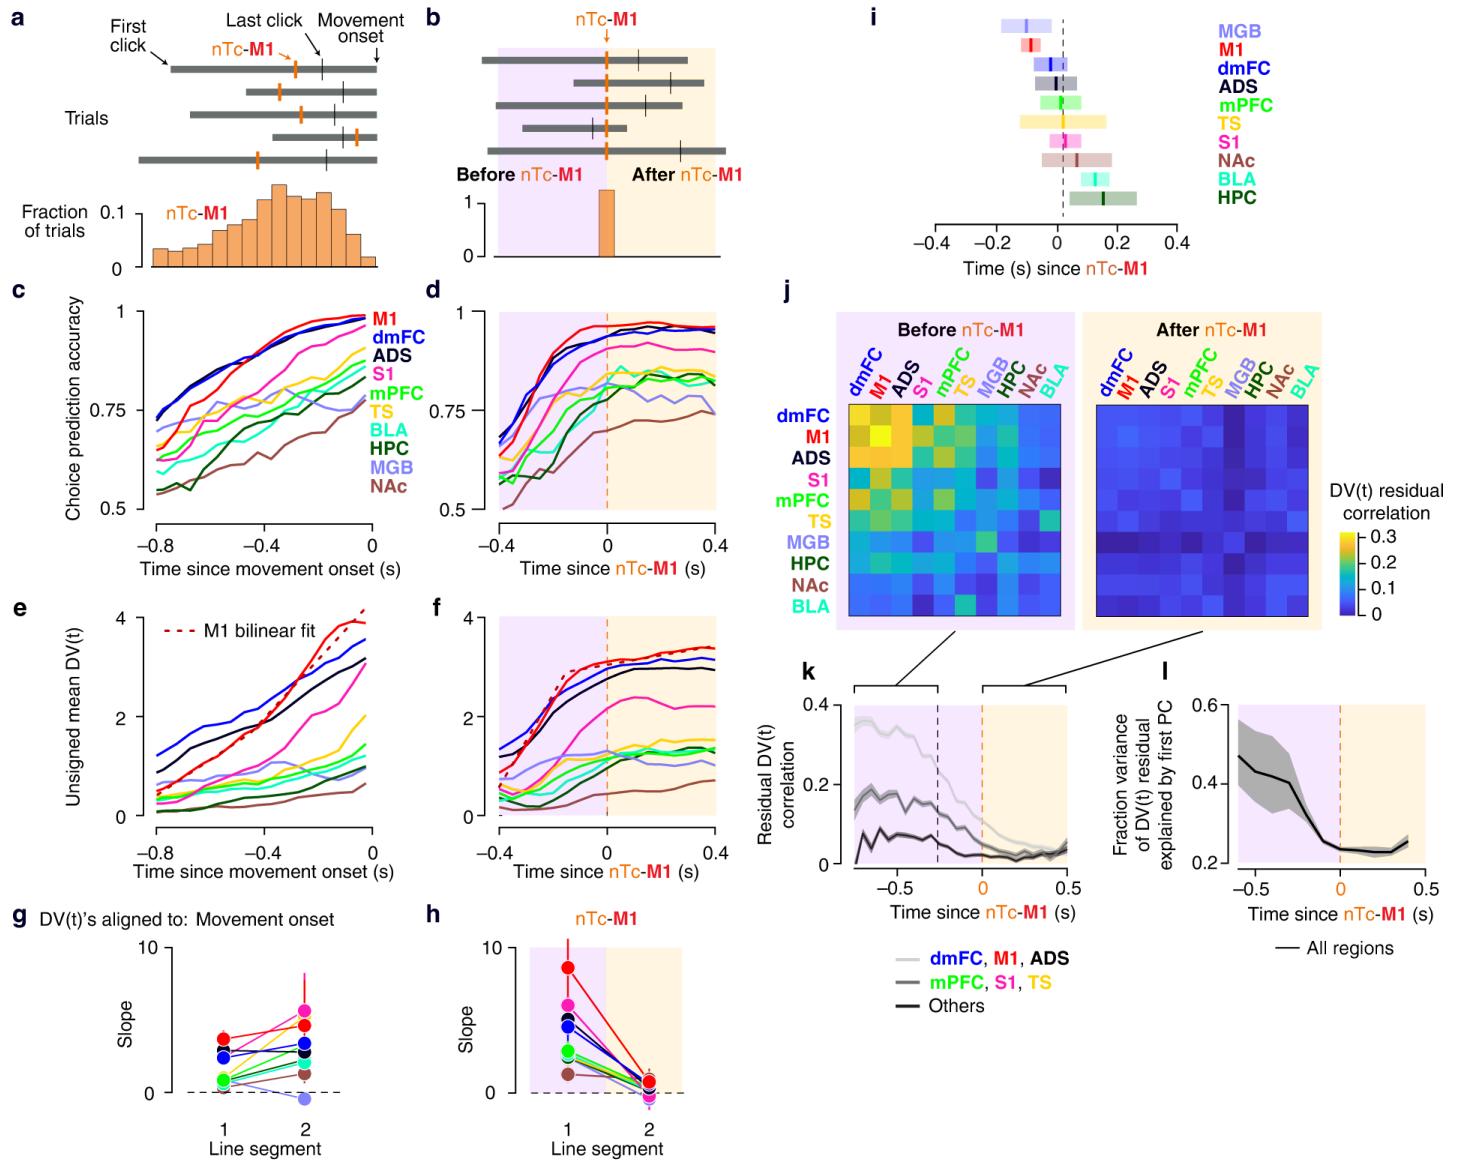
## Decision commitment marks a coordinated, brain wide state change in decision-related activity

Conceptually, decision commitment marks a change of state in the decision-making process, from (1) accumulating evidence as the decision is being formed, to (2) maintaining the decision the subject has committed to until an action reporting it is performed. But whether it marks a corresponding state change in brain-wide neural activity has not been previously examined. The simultaneous recordings enabled us to address this question by aligning neural data from all recorded regions to each individual trial's nTc-M1 (**Fig. 6b**; we remind the reader that "nTc-M1" indicates nTc estimated from neural activity in region M1). We characterized two measures of decision-related activity. First, we examined choice prediction accuracy (**Fig. 1h**), which has been observed to have a very regular property across multiple tasks, brain regions, and species: when trials are aligned to externally-observable task events, such as stimulus onset or the onset of the movement that reports the animal's choice, choice-prediction accuracy from neural activity ramps steadily upwards over time<sup>2,5,47</sup>. Indeed, using this alignment, we found the same pattern in our data (**Fig. 6c**). However,

aligning trials to nTc-M1 revealed a dramatically different pattern: for almost all regions, choice prediction accuracy ramped upwards until nTc-M1 but then abruptly plateaued (**Fig. 6d**; the two exceptions are MGB, where it falls instead of plateauing, and NAc, where it continues growing). This held even for trials in which nTc-M1 occurred many hundreds of milliseconds before movement onset. This surprising finding is consistent with choice prediction accuracy largely reflecting an internal estimate of accumulated evidence that stops evolving when the subject commits to a decision. It suggests that monotonically increasing choice prediction accuracy when aligning to movement onset (**Fig. 6c**) is due to averaging data across two different states.

To rule out that the observed plateauing is artefactually generated by choice prediction accuracy nearing its maximum value of 1, we also analyzed the sign-corrected trial-averaged DV(t)'s for each region: to pool across both right- and left-choice trials, we first flipped the sign of DV(t) on left-choice trials. The resulting DV(t) is an unbounded measure of the strength with which neural activity predicts the upcoming choice (**Fig. 1; Methods**). These region-specific DV(t)'s confirmed the observations found with choice prediction accuracy: steadily upwards ramping when aligning on movement onset (**Fig. 6e**) but upon aligning on nTc-M1, a plateauing that begins around the time of that alignment point, for data across all regions (**Fig. 6f**). To quantify this observation, we performed bilinear regression on each of the DV(t) traces in **Fig. 6e,f**. For visibility, we show here the bilinear fit for the M1 data only (red dashed lines in **Fig. 6e,f**; see **Extended Data Fig. 9** for all bilinear fits and for alignment to stimulus onset). For data aligned to movement onset, the slopes of both the first and second line segments in the bilinear fits were positive for all regions (**Fig. 6g**). However, for data aligned to nTc-M1, the bilinear fits confirmed brain-wide plateauing around nTc-M1: for all regions, the slopes of the first line segment were significantly positive ( $p < 0.05$  using hierarchical bootstrap test across sessions) while none were significantly positive for the second line segment ( $p > 0.05$  using hierarchical bootstrap test across sessions; **Fig. 6h**). For all regions, the time of the breakpoint from the first to the second line in the fits was closely aligned to the time of nTc-M1 (**Fig. 6i**). The temporal order of breakpoint timing between regions (**Fig. 6i**) was similar to that of nTc inferred separately from each region (**Fig. 5f**), with the exception of MGB, which here exhibited the earliest breakpoint but has a much later relative nTc timing in **Fig. 5f**. We interpret the fast initial rise in MGB's choice prediction accuracy (see also **Fig. 1h**) as likely due to MGB's robust and low latency responses to auditory clicks<sup>48</sup>, which are expected from an early region in the auditory pathway that strongly encodes momentary evidence (**Extended Data Fig. 5c**). Thus, the brain-wide plateauing in choice-related activity around the time of nTc-M1 provides independent evidence for the coordinated nature of the decision commitment signal observed across the brain. The relative timing of the plateauing across regions supports the cascading nature of the decision commitment signal observed across the brain, even while highlighting nTc as a more precise measure of neural encoding of decision commitment.

In **Fig. 2**, we showed that the DV(t) residuals strongly covary across regions in the time period before nTc-M1 on each trial, consistent with a single shared decision formation process. Here we directly compare coordination in the DV(t) residuals in the period before commitment (from stimulus start to nTc-M1) and after commitment (from nTc-M1 to the onset of the decision-reporting movement; see grey schematic bars in **Fig. 6b**). We found that after nTc-M1, correlations fell precipitously by  $\sim 3x$  compared to the period before nTc-M1 (**Fig. 6j**). This result held separately for each of the five frozen noise sessions (**Extended Data Fig. 10**). To provide a time-resolved view of the transition, we computed residual correlations for individual time bins around nTc-M1, averaging across region pairs from within three groups of regions (**Fig. 6k**). We found that the decrease in correlations was not uniform over the course of the trial but instead began approximately 250 milliseconds before nTc-M1. The fraction of variance explained by the first principal component of the residuals provides another measure of the degree of correlation. Computing this for each time bin exhibited a similar drop  $\sim 250$  ms before nTc-M1 (**Fig. 6l**), suggesting that the drop in correlations may predict a commitment event that is about to occur. We excluded the period when DV(t) residual correlations were rapidly dropping (from 250ms before nTc-M1 to nTc-M1) from calculations of the average magnitude of coordination in the pre-commitment period (**Fig. 2, Fig. 6j and Extended Data Fig. 10**).



**Figure 6. A coordinated, brain-wide state change at the time of decision commitment.** **a**, Timing of nTc-M1 relative to movement onset. Top, sequence of the first click, nTc-M1, last click and movement onset for 5 example trials. Bottom, distribution of nTc-M1 across trials aligned to movement onset. **b**, Same as (a) with trials aligned to nTc-M1. Purple shading highlights the period before nTc-M1. Yellow shading highlights the time window after nTc-M1. **c**, Choice prediction accuracy using each brain region, as in Fig. 1h, but here aligned to movement onset. Mean across sessions shown. **d**, Same as (c), with trials aligned to nTc-M1. **e**, Trial-averaged DV(t) from each brain region, aligned to movement onset. The sign of the DV(t)s on left choice trials was flipped before averaging. Note that neuronal firing rates were smoothed with a 50-ms symmetric Gaussian filter before the logistic regression was performed to generate the DV(t)s. **f**, Same as (e), with trials aligned to nTc-M1. **g**, Slopes of the first and second line segments of bilinear fits of the mean DV(t) traces in (e). **h**, Same as (g), for the bilinear fits of the mean DV(t) traces in (f). **i**, Timing of the breakpoint from the bilinear fits for each region relative to nTc-M1. Data represented as mean  $\pm$  1 s.e. across 5 frozen noise sessions. **j**, Matrix of DV(t) residual correlations for all pairs of brain regions, across all time points and trials within a session here shown separately for time points before (left, reproduced from Fig. 2c) and after (right) nTc-M1. The diagonal shows the interhemispheric correlation for each region. Values show mean across the 5 frozen noise sessions. The time period when DV(t) residual correlations were rapidly changing (from 250ms before nTc-M1 to nTc-M1) were excluded from these averages. **k**, Time course of mean ( $\pm$  1 s.e.) DV(t) residual correlation relative to nTc-M1, using 50 ms bins. **l**, Time course of fraction of across-region DV(t) residual variance explained by the first PC, similar to Fig. 2d but here shown in 200 ms bins relative to nTc-M1. Mean  $\pm$  1 s.e. shown. All panels, only trials with detected nTc-M1 in the 5 frozen noise sessions are included (n=1656).

Taken together, the results in **Fig. 6** show that nTc-M1, the time of decision commitment as inferred from neural activity in region M1, marks a substantial change across the brain between two very different phases of decision-related activity. Since nTc-M1 varies substantially between trials (**Fig. 5c,j**) accurate trial-by-trial estimates of it will be important for correctly parsing and analyzing brain activity during decision-making. Without this trial-by-trial estimate, data from two very different states, each with very different properties, may be inadvertently mixed and blurred together.

## Discussion

Neural activity has traditionally been described with reference to external markers such as the onset of a stimulus or a motor output. Yet cognitive processes need not be time-locked to external referents, and may instead involve signals that are internal to the organism. Given spikes from only a single neuron, variability across trials in these internal signals may be indistinguishable from noise. However, simultaneous recordings of many neurons can provide the statistical power to uncover them – revealing internal coordinate systems along which cognitive processes evolve. Our work here demonstrates a further evolution of this experimental logic: recording populations of neurons across multiple regions simultaneously allows observation of how internal signals co-fluctuate across brain regions, which not only illuminates these signals in greater detail but also provides powerful constraints on the multi-region circuits that generate them. Surveys of neuronal activity that record serially across regions do not provide this ability. Brain-wide simultaneous recordings that do not provide cellular resolution (e.g. fMRI or EEG) cannot reveal internal signals that occupy the many neural subspaces orthogonal to the recorded pooled signal (**Fig. 1a**).

We developed new methods using chronically implanted Neuropixels probes that allow simultaneous recordings of thousands of neurons across up to 20 brain regions in rats performing an auditory evidence accumulation task. This approach provides both cross-brain simultaneous coverage and cellular resolution, making it suitable for studying coordination of internal signals. We identified two such signals—one reflecting the continuous evolution of a single decision formation computation, and another marking the moment of covert decision commitment—that together dominate decision-related activity across the brain. Both signals were highly variable across trials – even across trials with identical stimuli and choices – but remarkably unified across regions. Aligning neural activity across the brain to the decision commitment signal obtained from M1, we observed a coordinated transition in decision-related dynamics: a dramatic reorganization of brain-wide neural activity at an internally defined time. Our results suggest that methods for uncovering internal signals and state changes are indispensable tools for understanding the neural basis of cognition

To obtain gradually evolving decision formation signals in each brain region, we used logistic regression to determine, for each time point, the weighted sum of neural firing rates that best predicts the animal's choice. Within the neural space of that region, these weights define a one-dimensional decision-related subspace. Position on that subspace defines what is referred to as that region's "decision variable", or DV(t)<sup>34–37</sup> (**Fig. 2c**). The temporal evolution of a region's DV(t) on single trials is a one-dimensional summary of the evolution of decision-related neural activity in that region. To avoid trivial co-fluctuations that would result from cross-trial changes in the sensory stimulus or the subject's choice, we used a "frozen noise" task design and conducted our analyses within sets of trials that had a given choice and click train stimulus, examining co-fluctuations in these signals relative to the mean within each of these groups. This allowed us to remove co-fluctuations driven by shared coding of the stimulus, or yoked to choice, instead specifically focusing on those reflecting internally-generated decision signals. We found that, despite having subtracted the mean DV(t) trajectory for each given stimulus and behavioral choice, the resulting single-trial residuals were substantial in magnitude and significantly correlated between regions (**Fig. 2f,h**).

Using nTc, a neural marker of the moment of decision commitment<sup>7</sup> on each trial, allowed examining co-fluctuations that are specific to ongoing decision formation, before commitment (**Fig. 2**). Focusing on lead-lag relationships amongst the ongoing internally-generated decision formation signals we found that both

hemispheres of three regions—anterior dorsal striatum (ADS), primary motor cortex (M1), and dorsomedial frontal cortex (dmFC), a group that notably includes a subcortical region—led other regions by ~50 ms or more (**Fig. 3**). Ongoing decision formation signals may thus arise in these three regions, and from there spread to the rest of the brain. The idea that signals related to an upcoming decision are locally computed but widely broadcast, as we provide evidence for here, has been suggested before in experimental<sup>14</sup> and theoretical studies<sup>49,50</sup>. How and when such global broadcasting occurs and what function it may serve is an important question for future research. Further simultaneous recordings will be required to determine lead-lag relationships with other regions associated with perceptual decision formation but not yet recorded here, including subcortical regions like superior colliculus<sup>1,3,51,52</sup> the midbrain reticular nucleus<sup>1,2</sup>, and the gigantocellular reticular nucleus<sup>2</sup>. Previous studies have measured the trial-average latency of the first decision-predictive signals to indicate in which region the decision formation computation is first triggered, with various results<sup>2,5,10,45</sup>. However, the computation could be initially triggered in one set of regions and yet be subsequently led by other regions. Our measurement, which focused on lead-lag relationships of ongoing co-fluctuations, not initial triggering decision-predictive signals, is thus distinct from and complementary to previous latency measurements. Consistent with ADS, M1, and dmFC being part of a set of regions where ongoing decision formation signals could originate, those three temporally-leading regions were also found to be the three regions with the strongest choice-predictive activity (**Fig. 2d**), as well as having the strongest pairwise correlations between regions (**Fig. 2h**).

Importantly, the large-scale recordings from many regions (median=19 regions per session) facilitated dimensionality analysis of decision-related co-fluctuations that are specific to the decision formation period (**Fig. 2**). Under the assumption that decision variables DV(t) reflect the state of the evolving decision process, dimensionality analysis can delineate subcircuits involved. This is because distinct computations in distinct subcircuits would require distinct state variables. By placing an upper bound on the number of variables required to describe the process across all recorded regions, dimensionality analysis places an upper bound on the number of state variables and therefore number of subcircuits that are involved. The idea that distinct strategies and subcircuits may compete for behavioral expression is prominent in theoretical and experimental accounts of operant behaviors<sup>38,53–55</sup>. Within perceptual decision-making, previous behavioral<sup>56</sup>, theoretical<sup>57</sup>, and neural activity<sup>58</sup> analyses have suggested that decision formation may involve separate and competing accumulators. For right-versus-left decisions in particular, the two hemispheres have been proposed to separately accumulate evidence for the two choices. Nevertheless, to our surprise, the pattern of co-fluctuations during decision formation in our bilateral brain-wide recordings indicated no sign of distinct subcircuits (**Fig. 2 and Extended Data Fig. 3c**). Strikingly, PCA dimensionality analysis indicated that a single scalar sufficed to describe residual DV(t) co-fluctuations across all recorded regions, in both hemispheres. This single scalar accounted for nearly 4 times as much variance as would be accounted for by a second scalar (which fell below the noise floor). This result was remarkably close to what would be predicted if all regions were participating in a single integrated decision formation computation (**Fig. 2i**), and it places strong constraints on multi-region models of decision-making.

Turning to decision commitment, the nTc measure<sup>7</sup> uses each trial's sensory stimulus, together with spiking data from a population of neurons, to provide an estimate of whether and when decision commitment occurred on that trial. Thus, on any given trial, differences in nTc presence or timing estimated using neurons from different regions is due only to differences in each region's neural activity. We found that nTc could be detected in the largest proportion of trials (~50%) from the neural activity of primary motor cortex (M1). Furthermore, nTc-M1 (i.e. nTc inferred using M1 activity alone) stood out as temporally leading other regions. In particular, nTc-M1 led nTc-ADS and nTc-dmFC by ~50ms (**Fig. 5f**). This contrasted with the decision formation phase, during which all three of M1, ADS, and dmFC were found to have roughly temporally coincident DV(t) fluctuations (**Fig. 3c**), and suggests that temporal relationships across regions change over different phases of decision-making. We once again used the frozen noise design to examine co-fluctuations not due to the stimulus or choice, now examining nTc residuals around the mean with a given stimulus and choice. As with DV(t) residuals in decision formation, nTc residuals had substantial magnitude and were

significantly correlated across brain regions (**Fig. 4k**). PCA dimensionality analysis indicated that their co-fluctuations could be described by fluctuations in a single scalar, with a second dimension falling below the noise floor (**Fig. 5d**; although the dominance of the first dimension was not as marked as for decision formation signals). Overall, these data suggest that a single integrated decision commitment computation is carried out in the brain, perhaps initiated in M1, and that this commitment signal then spreads from M1 to other brain regions.

The superior colliculus (SC), which was not recorded here, has been reported to display a neural signature of decision commitment that, in contrast to nTc, is overt in the sense of immediately preceding and being tightly linked to the subject's overt decision-reporting motor act<sup>3</sup>. Determining the nature of the relationship between the SC's overt commitment signature and the brain-wide covert nTc commitment signature will be an important topic for future study.

One of our most striking findings is that nTc, the neurally-inferred estimate of the moment of decision commitment on each trial, appears to mark a major and previously unknown state change in decision-related activity across the brain. We focused on nTc-M1, since nTc could be most reliably detected in M1 (**Fig. 4d**). An almost ubiquitous observation in perceptual decision-making, across tasks, species, and brain regions, is that the accuracy with which neural activity predicts the subject's choice increases monotonically over time, steadily growing as one approaches the moment in which the subject will begin a motor act to report their choice<sup>2</sup>. In the absence of a trial-by-trial internal time marker (such as nTc-M1), studies have time-aligned trials to externally determined or observable timepoints, such as the start of the sensory stimulus, or the onset of the subject's decision report. Indeed, aligned to such timepoints our data showed the very same pattern (**Fig. 6c**). However, aligning trials to nTc-M1 revealed something dramatically different: for almost all brain regions, choice prediction accuracy increased monotonically before nTc-M1, but abruptly stopped growing and stayed at a roughly constant level after it (**Fig. 6d**). This surprising finding is consistent with decision commitment being a unified, brain-wide signal that marks the end of the temporal evolution of the decision variables. Overall firing rates did not fall after nTc-M1 (**Extended Data Fig. 8a**): what changed before versus after nTc-M1 was the temporal evolution of the decision-related activity.

A second measure that was profoundly different before versus after nTc-M1 was the magnitude of cross-region co-fluctuations in DV(t) residuals: after nTc-M1, average correlation magnitude fell by a factor of 3 (**Fig. 6j,k**). This remarkable drop suggests that upon completion of a unified process of decision formation and a transition to a post-commitment state, brain regions stop coordinating their decision-related activity. Interestingly, the reduction in co-fluctuation magnitude began approximately 250 ms before nTc-M1. This suggests the existence of neural signatures that predict an upcoming decision commitment event even before it occurs. An important observation about the state change is that it occurs at a time that is highly variable across trials. nTc-M1 is not timelocked to overt external timepoints such as the start of the trial, start or end of the stimulus, or onset of the subject's choice-reporting movement (**Extended Data Fig. 8b**). Consequently, analyses of neural data that are timelocked to such external timepoints and are unaware of each trial's commitment time, as most current analyses are, risk blurring data across two very different states.

We note that this state change strongly resembles what would be expected if brain-wide decision-related activity reflected a single, coherent mechanism approximated by bounded diffusion, as described by the well-known drift diffusion model (DDM)<sup>13,46</sup>. Not only do the DV(t)'s from each region plateau at nTc-M1, consistent with an absorbing bound being reached simultaneously at that time, but the DV(t)'s become more independent across regions around the time of nTc, consistent with a reduction in variability of a shared latent variable. Future studies may investigate in greater depth to what degree these resemblances are borne out quantitatively. Nonetheless, it is remarkable that a simple one-dimensional behavioral model has proved so enduring in offering a useful framework for understanding neural responses, first at the level of single neurons<sup>16</sup>, then at the level of neural populations<sup>34-36</sup>, and here, at the level of coordinated, brain-wide activity.

If we consider a set of trials with a given sensory stimulus and subject choice, then by construction, within- and across-trial fluctuations around the average neural trajectory of that set cannot be explained by

changes in the stimulus or choice. Such fluctuations instead appear to be generated internally by the organism. A very striking feature of our data was the large magnitude of such internally generated fluctuations. Decision variable residual fluctuations had a magnitude comparable to that of the mean trajectory itself (**Fig. 2f**). Similarly, nTc times varied widely across trials, over a range comparable to the duration of the trials (**Fig. 4f**). Strong covariations across regions in these fluctuations confirm that they are not single-region measurement noise. They are instead very salient features of neural activity on single trials, as also demonstrated by the profound state change we observed around nTc-M1. Previous studies have documented substantial neural variability arising from uninstructed movements or during spontaneous behaviors<sup>59–61</sup>. Our data show that even when focusing specifically on task-related neural subspaces, we find that internally-generated, within- and across-trial fluctuations are highly prominent features of the data. The extent to which uninstructed movements play a part in the organism's internal generation of brain-wide coordinated variability within task-related subspaces remains to be determined.

Our results highlight how the ongoing shift in systems neuroscience toward brain-wide neural recording during complex behavior<sup>1,2,5</sup> can be dramatically enhanced if recordings are performed simultaneously, rather than serially. Referencing neural signals to external events has given us many fundamental insights. However it can make neural population dynamics seem like a calm, stately progression towards logical outcomes (see left column of **Fig. 2e**). In contrast, single trials of multi-neuron, multi-region simultaneous recordings suggest something starkly different: task-related neural dynamics appear to be more like a vibrant storm, with highly varying, internally-generated signals that have a brain-wide reach (**Fig. 2e**, right, **Fig. 3**, and **Fig. 4k**), and abrupt state changes at externally unpredictable, internally-generated times (**Fig. 6**). The fact that single trials can be very substantially different from one another, even when identically prepared and resulting in the same choice outcome, leads us to postulate that understanding the internally-generated signals and events that make single trials so variable is an essential step towards understanding the neural substrates of behavior. Such internally-referenced and internally-timed signals require simultaneous recordings to detect them and analyze them. We speculate that, powered by ever more readily-accessible large-scale simultaneous recording tools, the findings we have described here will be only some of many<sup>62,63</sup> internal signals to be discovered and deciphered. Such discoveries may lead us to see the vividly dynamic tapestry of neural activity as perhaps colored more by the organism's own internal signals and events than by external stimuli or overt motor acts<sup>64,65</sup>.

## References

1. Steinmetz, N. A., Zatka-Haas, P., Carandini, M. & Harris, K. D. Distributed coding of choice, action and engagement across the mouse brain. *Nature* **576**, 266–273 (2019).
2. International Brain Laboratory *et al.* A brain-wide map of neural activity during complex behaviour. *Nature* **645**, 177–191 (2025).
3. Stine, G. M., Trautmann, E. M., Jeurissen, D. & Shadlen, M. N. A neural mechanism for terminating decisions. *Neuron* **111**, 2601–2613.e5 (2023).
4. Ebrahimi, S. *et al.* Emergent reliability in sensory cortical coding and inter-area communication. *Nature* **605**, 713–721 (2022).
5. Chen, S. *et al.* Brain-wide neural activity underlying memory-guided movement. *Cell* **187**, 676–691.e16 (2024).
6. Chen, G., Kang, B., Lindsey, J., Druckmann, S. & Li, N. Modularity and robustness of frontal cortical networks. *Cell* **184**, 3717–3730.e24 (2021).
7. Luo, T. Z. *et al.* Transitions in dynamical regime and neural mode during perceptual decisions. *Nature* (2025) doi:10.1038/s41586-025-09528-4.
8. Pinto, L. *et al.* Task-Dependent Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron* **104**, 810–824.e9 (2019).
9. Khilkevich, A. *et al.* Brain-wide dynamics linking sensation to action during decision-making. *Nature* **634**, 890–900 (2024).
10. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science* **348**, 1352–1355 (2015).
11. Richman, E. B., Ticea, N., Allen, W. E., Deisseroth, K. & Luo, L. Neural landscape diffusion resolves conflicts between needs across time. *Nature* **623**, 571–579 (2023).
12. Allen, W. E. *et al.* Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science* **364**, 253 (2019).
13. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–574 (2007).
14. Brody, C. D. & Hanks, T. D. Neural underpinnings of the evidence accumulator. *Curr. Opin.*

*Neurobiol.* (2016) doi:10.1016/j.conb.2016.01.003.

15. Brunton, B. W., Botvinick, M. M. & Brody, C. D. Rats and humans can optimally accumulate evidence for decision-making. *Science* **340**, 95–98 (2013).
16. Shadlen, M. N. & Newsome, W. T. Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *J. Neurophysiol.* **86**, 1916–1936 (2001).
17. Charlton, J. A. & Goris, R. L. T. Abstract deliberation by visuomotor neurons in prefrontal cortex. *Nat. Neurosci.* **27**, 1167–1175 (2024).
18. Jun, J. J. *et al.* Fully integrated silicon probes for high-density recording of neural activity. *Nature* **551**, 232–236 (2017).
19. Dennis, E. J. *et al.* Princeton RAtlas: A Common Coordinate Framework for Fully cleared, Whole Brains. *Bio Protoc* **13**, e4854 (2023).
20. Yartsey, M. M., Hanks, T. D., Yoon, A. M. & Brody, C. D. Causal contribution and dynamical encoding in the striatum during evidence accumulation. *bioRxiv* 245316 (2018) doi:10.1101/245316.
21. Pisupati, S., Chartarifsky-Lynn, L., Khanal, A. & Churchland, A. K. Lapses in perceptual decisions reflect exploration. *Elife* **10**, (2021).
22. Anastasiades, P. G. & Carter, A. G. Circuit organization of the rodent medial prefrontal cortex. *Trends Neurosci.* **44**, 550–563 (2021).
23. Hunnicutt, B. J. *et al.* A comprehensive excitatory input map of the striatum reveals novel functional organization. *Elife* **5**, (2016).
24. Musiek, F. E. & Baran, J. A. *The Auditory System: Anatomy, Physiology, and Clinical Correlates; Second Edition*. (Plural Publishing, 2018).
25. Jiang, H. & Kim, H. F. Anatomical Inputs From the Sensory and Value Structures to the Tail of the Rat Striatum. *Front. Neuroanat.* **12**, 30 (2018).
26. Chen, L., Wang, X., Ge, S. & Xiong, Q. Medial geniculate body and primary auditory cortex differentially contribute to striatal sound representations. *Nat. Commun.* **10**, 418 (2019).
27. Znamenskiy, P. & Zador, A. M. Corticostriatal neurons in auditory cortex drive decisions during auditory discrimination. *Nature* **497**, 482–485 (2013).

28. Guo, L., Walker, W. I., Ponvert, N. D., Penix, P. L. & Jaramillo, S. Stable representation of sounds in the posterior striatum during flexible auditory decisions. *Nat. Commun.* **9**, 1534 (2018).
29. Aronoff, R. *et al.* Long-range connectivity of mouse primary somatosensory barrel cortex. *Eur. J. Neurosci.* **31**, 2221–2233 (2010).
30. Buetfering, C. *et al.* Behaviorally relevant decision coding in primary somatosensory cortex neurons. *Nat. Neurosci.* **25**, 1225–1236 (2022).
31. Nieh, E. H. *et al.* Geometry of abstract learned knowledge in the hippocampus. *Nature* **595**, 80–84 (2021).
32. LeDoux, J. E., Farb, C. R. & Romanski, L. M. Overlapping projections to the amygdala and striatum from auditory processing areas of the thalamus and cortex. *Neurosci. Lett.* **134**, 139–144 (1991).
33. Kopec, C. D. *et al.* To integrate or not to integrate: Testing degenerate strategies for solving an accumulation of perceptual evidence decision-making task. *bioRxiv* (2024) doi:10.1101/2024.08.21.609064.
34. Peixoto, D. *et al.* Decoding and perturbing decision states in real time. *Nature* **591**, 604–609 (2021).
35. Kiani, R., Cueva, C. J., Reppas, J. B. & Newsome, W. T. Dynamics of Neural Population Responses in Prefrontal Cortex Indicate Changes of Mind on Single Trials. *Curr. Biol.* **24**, 1542–1547 (2014).
36. Steinemann, N. A. *et al.* Direct observation of the neural computations underlying a single decision. *bioRxiv* (2022) doi:10.1101/2022.05.02.490321.
37. Kaufman, M. T., Churchland, M. M., Ryu, S. I. & Shenoy, K. V. Vacillation, indecision and hesitation in moment-by-moment decoding of monkey motor cortex. *Elife* **4**, e04677 (2015).
38. Daw, N. D., Niv, Y. & Dayan, P. Uncertainty-based competition between prefrontal and dorsolateral striatal systems for behavioral control. *Nat. Neurosci.* **8**, 1704–1711 (2005).
39. Balleine, B. W. & Dickinson, A. Goal-directed instrumental action: contingency and incentive learning and their cortical substrates. *Neuropharmacology* **37**, 407–419 (1998).

40. Bondy, A. G., Haefner, R. M. & Cumming, B. G. Feedback determines the structure of correlated variability in primary visual cortex. *Nat. Neurosci.* **21**, 598–606 (2018).
41. Nienborg, H. & Cumming, B. G. Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature* **459**, 89–92 (2009).
42. Neri, P. & Levi, D. M. Receptive versus perceptive fields from the reverse-correlation viewpoint. *Vision Res.* **46**, 2465–2474 (2006).
43. Yates, J. L., Park, I. M., Katz, L. N., Pillow, J. W. & Huk, A. C. Functional dissection of signal and noise in MT and LIP during decision-making. *Nat. Neurosci.* **20**, 1285–1292 (2017).
44. Foster, N. N. *et al.* The mouse cortico-basal ganglia-thalamic network. *Nature* **598**, 188–194 (2021).
45. Duan, C. A. *et al.* Collicular circuits for flexible sensorimotor routing. *Nat. Neurosci.* **24**, 1110–1120 (2021).
46. Ratcliff, R. & McKoon, G. The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* **20**, 873–922 (2008).
47. Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat. Neurosci.* **2**, 176–185 (1999).
48. Rouiller, E., de Ribaupierre, Y., Toros-Morel, A. & de Ribaupierre, F. Neural coding of repetitive clicks in the medial geniculate body of cat. *Hear. Res.* **5**, 81–100 (1981).
49. Mashour, G. A., Roelfsema, P., Changeux, J.-P. & Dehaene, S. Conscious processing and the global neuronal workspace hypothesis. *Neuron* **105**, 776–798 (2020).
50. Cohen, M. X. & Ranganath, C. Reinforcement learning signals predict future decisions. *J. Neurosci.* **27**, 371–378 (2007).
51. Crapse, T. B., Lau, H. & Basso, M. A. A role for the superior colliculus in decision criteria. *Neuron* **97**, 181–194.e6 (2018).
52. Horwitz, G. D., Batista, A. P. & Newsome, W. T. Representation of an abstract perceptual decision in macaque superior colliculus. *J. Neurophysiol.* **91**, 2281–2296 (2004).
53. Miller, K. J., Botvinick, M. M. & Brody, C. D. Dorsal hippocampus contributes to model-based planning. *Nat. Neurosci.* **20**, 1269–1276 (2017).

54. Dezfouli, A. & Balleine, B. W. Habits, action sequences and reinforcement learning. *Eur. J. Neurosci.* **35**, 1036–1051 (2012).
55. Corbit, L. H. Understanding the balance between goal-directed and habitual behavioral control. *Curr. Opin. Behav. Sci.* **20**, 161–168 (2018).
56. Scott, B. B., Constantinople, C. M., Erlich, J. C., Tank, D. W. & Brody, C. D. Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. *Elife* **4**, e11308 (2015).
57. Mazurek, M. E., Roitman, J. D., Ditterich, J. & Shadlen, M. N. A Role for Neural Integrators in Perceptual Decision Making. *Cereb. Cortex* **13**, 1257–1269 (2003).
58. DePasquale, B., Brody, C. D. & Pillow, J. W. Neural population dynamics underlying evidence accumulation in multiple rat brain regions. *Elife* **13**, (2024).
59. Stringer, C. *et al.* Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
60. Musall, S., Kaufman, M. T., Juavinett, A. L., Gluf, S. & Churchland, A. K. Single-trial neural dynamics are dominated by richly varied movements. *Nat. Neurosci.* **22**, 1677–1686 (2019).
61. Manley, J. *et al.* Simultaneous, cortex-wide dynamics of up to 1 million neurons reveal unbounded scaling of dimensionality with neuron number. *Neuron* **112**, 1694–1709.e5 (2024).
62. Volland, A. Z., Gardner, R. J., Moser, M.-B. & Moser, E. I. Left-right-alternating theta sweeps in entorhinal-hippocampal maps of space. *Nature* **639**, 995–1005 (2025).
63. Senzai, Y. & Scanziani, M. The brain simulates actions and their consequences during REM sleep. *bioRxiv* (2024) doi:10.1101/2024.08.13.607810.
64. Arieli, A., Sterkin, A., Grinvald, A. & Aertsen, A. Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. *Science* **273**, 1868–1871 (1996).
65. Fiser, J., Chiu, C. & Weliky, M. Small modulation of ongoing cortical dynamics by sensory input during natural vision. *Nature* **431**, 573–578 (2004).
66. Juavinett, A. L., Bekheet, G. & Churchland, A. K. Chronically-implanted Neuropixels probes enable high yield recordings in freely moving mice. *bioRxiv* 406074 (2018)

doi:10.1101/406074.

67. Luo, T. Z. *et al.* An approach for long-term, multi-probe Neuropixels recordings in unrestrained rats. *Elife* **9**, (2020).
68. Steinmetz, N. A. *et al.* Neuropixels 2.0: A miniaturized high-density probe for stable, long-term brain recordings. *Science* **372**, (2021).
69. Bimbard, C. *et al.* An adaptable, reusable, and light implant for chronic Neuropixels probes. (2024) doi:10.7554/elife.98522.
70. Pachitariu, M., Sridhar, S. & Stringer, C. Solving the spike sorting problem with Kilosort. *bioRxiv* 2023.01.07.523036 (2023) doi:10.1101/2023.01.07.523036.
71. Pietzsch, T., Saalfeld, S., Preibisch, S. & Tomancak, P. BigDataViewer: visualization and processing for large image data sets. *Nat. Methods* **12**, 481–483 (2015).
72. Schindelin, J. *et al.* Fiji: an open-source platform for biological-image analysis. *Nat. Methods* **9**, 676–682 (2012).
73. Fabre, J. M. J., van Beest, E. H., Peters, A. J., Carandini, M. & Harris, K. D. *Bombcell: Automated Curation and Cell Classification of Spike-Sorted Electrophysiology Data*. (Zenodo, 2023). doi:10.5281/ZENODO.8172821.
74. Friedman, J., Hastie, T. & Tibshirani, R. Regularization Paths for Generalized Linear Models via Coordinate Descent. *J. Stat. Softw.* **33**, 1–22 (2010).
75. Tipping, M. E. & Bishop, C. M. Probabilistic principal component analysis. *J. R. Stat. Soc. Series B Stat. Methodol.* **61**, 611–622 (1999).

## Acknowledgments

We thank Maxime Beau, Zoe Boundy-Singer, Bing Brunton, Mark Goldman, Ben Lankow, Maxwell Melin, Hendrikje Nienborg, Bence Olveczky, Maxwell Shinn, and Corey Ziemba for their helpful 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 James MacGuire for veterinary advice. We are grateful to all members of the Brody lab for their support, collegiality and feedback.

This manuscript was supported in part by the National Institutes of Health (NIH) under grants R01MH108358, R01MH138935, and 5U19NS132720. It is subject to the NIH Public Access Policy. Through acceptance of this federal funding, NIH has been given a right to make this manuscript publicly available in PubMed Central upon the Official Date of Publication, as defined by NIH. Additional support was provided by the Howard Hughes Medical Institute Investigator Program.

## Author Contributions

**Conceptualization.** A.G.B., J.A.C., T.Z.L., T.H., and C.D.B.

**Methodology.** A.G.B., J.A.C., and T.Z.L.

**Formal analysis.** A.G.B. led the majority of data analysis, with contributions from T.Z.L.

**Investigation.** A.G.B., J.A.C., T.Z.L., C.D.K., W.S., S.J.V., L.L., and S.O.

**Data Curation.** A.G.B., J.A.C., S.J., and T.Z.L.

**Visualization.** A.G.B., J.A.C., T.Z.L., and W.S.

**Writing - Original Draft.** A.G.B., J.A.C., and T.Z.L.

**Writing - Reviewing & Editing.** A.G.B., J.A.C., and T.Z.L. C.D.K., W.S., S.J.V., and C.D.B.

**Supervision.** C.D.B.

**Project administration.** A.G.B., J.A.C., T.Z.L., and C.D.B.

**Funding acquisition.** T.H. and C.B.

## Data Availability

The datasets generated during the current study are available from the corresponding authors upon reasonable request and will be deposited in a public repository prior to publication.

## Code Availability

All code used in the analyses will be made freely available upon publication.

## Competing Interests

The authors declare no competing interests.

## 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<sup>15</sup>. 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 ( $\gamma$ ) 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 showed a lapse rate ( $\lambda$ ) > 8%, estimated using the following psychometric model:

$$p(go\ R\ | x) = \lambda/2 + (1 - \lambda) \cdot \sigma(\beta(x - \theta))$$

where  $x$  is the total difference in the number of right and left clicks,  $\theta$  is the threshold (i.e., point of subjective equality),  $\beta$  is the sensitivity (i.e., steepness of the curve), and  $\sigma(\cdot)$  is the logistic (sigmoid) function.

### Frozen noise

In the Poisson Clicks task, the click trains for each trial depends on that trial's generative click rate ratio ( $\gamma$ ) and a seed that determines the sample of two pseudorandom variables: the stimulus duration and the Poisson processes determining left and right click times. 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 of unique stimulus seeds (9 for each of the 6 values of  $\gamma$  used) were interleaved throughout the session. Only

these sessions contributed to the analyses in Figs. 2, 3, 5 and 6 and example frozen noise sessions were used for the schematic panels in Figs. 1i,j and 4j.

## Implant

Previous methods<sup>66-69</sup> 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. Only a small probe holder was used, to provide an attachment point for a micromanipulator used to lower the probes into the brain. This probe holder had two components: one that is glued permanently to a probe, and another that is removed after the probe is anchored to the skull, to reduce the volume and weight of the implant. The arrangement of the probes on the skulls of the rats was performed using CAD software, as was the design of the probe holder and a “chassis” that surrounded the probes and held the headstages (**Extended Data Fig. 2a-e**). The probe holders were 3D printed in-house 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-T015-01). CAD files for these components can be found at [https://github.com/Brody-Lab/uberphys\\_paper/tree/main/CAD\\_files](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<sup>67</sup>, but with several important innovations to support the lengthy and invasive nature of the 8-probe implantation. 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 Etiqa 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. Durotomy 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 stereotactically inserted into the brain using a motorized micromanipulator (Narshige, MDS-1) at a speed of ~5  $\mu\text{m/s}$ . 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 (**Extended Data Fig. 2**) was attached to the fixtures using dental composite. The surgeries were considerably longer than typical rat intracranial implant surgeries, lasting >12 hours. A team of five surgeons contributed in rotating shifts. We also provided the rats extended analgesia for two days post-op and ad libitum water access for at least one week after surgery.

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<sup>70</sup>, 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. We excluded 3 sessions from the rat A327 due to an error in specifying the electrodes to be recorded.

## 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<sup>19</sup>. These methods are described in detail elsewhere<sup>19</sup>.

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<sup>71</sup> Plugin for Fiji<sup>72</sup>, 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<sup>73</sup>. 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. Additionally, because multiple sessions were recorded from the same animals, the same neurons might have been sampled across days.

<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 $\mu$ m
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 $\mu$ V

**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 calculation of “decision variables”

We used logistic regression to decode choice from population firing rates. Neuronal firing rates were obtained by convolving spike times with a 50 ms s.d. symmetric Gaussian smoothing filter and sampled at 50 ms intervals. A separate logistic regression model was fit to each 50 ms sample. The model probability of a rightward choice at a given time point  $t$  on trial  $i$  was  $p_{t,i}(R) = f(X_{t,i}\beta_t + \alpha_t)$ , where  $f$  is the sigmoidal logistic function,  $X_{t,i}$  is the vector of neuronal firing rates for time point  $t$  on trial  $i$ ,  $\beta_t$  is the vector of weights for time point  $t$ , and  $\alpha_t$  is a model constant for time point  $t$ . The subject’s choice on trial  $i$  is modeled as

$C_i \sim \text{Bernoulli}(p_{t,i}(R))$ , where  $C = 1$  corresponds to a rightward choice.  $\beta_t$  and  $\alpha_t$  are the model parameters, shared across trials and chosen to maximize the log-likelihood of the subject’s choices: i.e.

$\text{argmax} \sum_i \log p(C_i | p_{t,i}(R))$ . We used 10-fold stratified cross-validation within sessions to assess model performance as well as to identify the optimal L1 regularization hyperparameter for  $\beta_t$ . The “decision variable” for a given time point and trial  $DV_{t,i}$  was defined as the linear predictor  $X_{t,i}\beta_t + \alpha_t$ , which is equivalent to the log-odds of  $p_{t,i}(R)$ , i.e.  $\log(p_{t,i}(R)/p_{t,i}(L))$ :

$$\begin{aligned} p_{t,i}(R) &= \frac{1}{1 + e^{-(X_{t,i}\beta_t + \alpha_t)}} \\ \frac{1}{p_{t,i}(L)} &= \frac{1}{1 - p_{t,i}(R)} = \frac{1}{1 - \frac{1}{1 + e^{-(X_{t,i}\beta_t + \alpha_t)}}} = \frac{1 + e^{-(X_{t,i}\beta_t + \alpha_t)}}{e^{-(X_{t,i}\beta_t + \alpha_t)}} \\ DV_{t,i} &= \log \left( \frac{p_{t,i}(R)}{p_{t,i}(L)} \right) \\ &= \log \left( \frac{1}{1 + e^{-(X_{t,i}\beta_t + \alpha_t)}} \cdot \frac{1 + e^{-(X_{t,i}\beta_t + \alpha_t)}}{e^{-(X_{t,i}\beta_t + \alpha_t)}} \right) \\ &= \log \left( \frac{1}{e^{-(X_{t,i}\beta_t + \alpha_t)}} \right) \\ &= X_{t,i}\beta_t + \alpha_t \end{aligned}$$

To ease comparisons of DVs across timepoints within a trial, we used a constant value of the L1 regularization hyperparameter across timepoints, obtained as the geometric mean of the hyperparameters identified through cross-validation at each timepoint. Choice prediction accuracy was calculated for each session using the “balanced accuracy”. The standard or *unbalanced* choice prediction accuracy (CPA) across a set of  $n$  trials for time point  $t$  can be given as:

$$CPA_t = \frac{1}{n} \sum_{i=1}^n \Delta(DV_{t,i}, C_i)$$

where  $\Delta$  is an indicator function that equals 1 if  $sign(DV_{t,i}) = C_i$  and 0 otherwise. The balanced accuracy is the average of the unbalanced accuracy computed separately for left and right choice trials. This removes upward bias associated with an uneven distribution of choices. To fit the decoding models, we relied on the Glmnet<sup>74</sup> package in MATLAB R2024a (Mathworks, MA, USA).

## Neural decoding of momentary evidence and calculation of “momentary evidence variables”

We used linear regression to decode momentary evidence from population firing rates (**Extended Data Fig. 5**), using a very similar procedure to decoding of choice. Neuronal firing rates were obtained by convolving spike times with a 50 ms s.d. symmetric Gaussian smoothing filter and sampled at 100 ms intervals.

Momentary evidence was defined as the difference in right versus left clicks within a 100 ms window before the time of the firing rate was sampled. A separate regression model was fit to each 100 ms sample. The model prediction of momentary evidence at a given time point  $t$  on trial  $i$  was  $e_{t,i}(R) = X_{t,i}\beta_t + \alpha_t$ , and we refer to this as the “momentary evidence variable” (MEV), and is analogous to the DV in the choice decoding model. It is a 1D summary of the neural population’s estimate of the momentary stimulus evidence. The model parameters were chosen to maximize the squared error of the momentary evidence estimates. As with choice decoding, 10-fold stratified cross-validation within sessions was used to assess model performance as well as to identify the optimal L1 regularization hyperparameter for  $\beta$ . To fit the decoding models, we again relied on the Glmnet<sup>74</sup> package in MATLAB R2024a (Mathworks, MA, USA).

## Calculating DV Residual Correlations

DV residual correlations were calculated as the Pearson correlation across trials and time points between the DVs estimated from different brain regions. To estimate the component of these correlations due to shared coding for stimulus and choice, we used a trial-shuffling procedure as follows (and as illustrated in **Extended Data Fig. 11**). 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. 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. This procedure was also used for shuffle-correcting correlations between the decoded momentary evidence and between mean firing rates reported in **Extended Data Fig. 5**. We note that, unless otherwise stated (as in **Fig. 6**), DV residual correlations were only shown using time points before nTc-M1 on each trial, since we observed a dramatic reduction in coordination of the DV residuals after that time point.

DV residual correlations were also calculated as a function of different lags between pairs of regions (i.e. DV residual cross-correlograms). The peak lag was calculated simply as the lag that produces the maximum value. A bootstrap across trials was used to assess the uncertainty of the peak lag for each pair. To calculate the mean peak lag of a given region relative to all others, a weighted average was used that factored in this uncertainty. The weights were proportional to the inverse of the bootstrap s.e.

## Assessing Dimensionality of DV Residuals

To calculate DV residuals, for each region, we separated trials into groups with identical click trains (frozen noise seed) and behavioral choices, computed the mean DV within each group, and subtracted it from the corresponding trials. We used principal components analysis (PCA) to assess the dimensionality of the DV residuals across regions. To do this, we define a matrix  $X$  where each column corresponds to the DV residual for a single region in each hemisphere (i.e., a feature) and each row corresponds to a single time point in a trial (i.e., a sample). We concatenated time points across trials within a session, and centered the data. PCA transforms the coordinates of  $X$  such that the axes are ordered by maximum variance explained. The principal component scores (i.e. the coordinates of the original data in the transformed space) are given by  $T = XW$ . The columns of  $W$  are equivalent to the eigenvectors of  $X^T X$ . The variance of  $X$  explained by the  $k^{th}$  principal component is equal to  $\lambda_k / \sum_k \lambda$ , where  $\lambda_k$  is the  $k^{th}$  eigenvalue of  $X^T X$ . The loadings  $L$  are given by the columns of  $W$  scaled by the square root of the corresponding eigenvalues (i.e.  $L = W\Lambda$ , where  $\Lambda$  is the diagonal matrix of eigenvalues). Elements of  $L$  give the cross-covariance of the activity of each region with its projection onto each PC, providing a measure of pairwise alignment between the regions and the PCs.

To generate predictions for the PCA dimensionality analysis under assumptions of different numbers of latent decision processes, we used a set of shuffling procedures. To generate predictions for a single latent decision process, we shuffled the anatomical identity of each cell. This creates “pseudo-regions” of the same size as the original populations recorded in each region, but containing cells randomly sampled from all recorded regions. We then obtained DV residuals from these pseudo-regions, and performed PCA as described above. This shuffling procedure tests the idea that the decision-related activity in each region are i.i.d samples from the same underlying distribution.

To generate predictions for the PCA results assuming complete independence across regions, we randomly shuffled time points of  $X$  separately for each region (columns of  $X$ ). To generate predictions assuming two independent decision processes, we first randomly divided regions (columns of  $X$ ) into two groups and randomly shuffled time points of each group. In this way, temporal alignment was abolished between groups but preserved within groups. These shuffling procedures were applied many times to estimate shuffle-based distributions which were then further sampled as part of the hierarchical bootstrap across sessions used to generate confidence intervals in the plots shown in Fig. 2.

We note that the PCA dimensionality described above was performed only on time points before nTc-M1 on each trial, since we observed a dramatic reduction in coordination of the DV residuals after that time point.

## Multi-mode drift-diffusion model (MMDDM)

A detailed description of the multi-mode drift-diffusion model (MMDDM) can be found in Ref.<sup>7</sup>. 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  $\sigma_s^2$ , and additive noise  $\epsilon$ :

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

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<sup>15</sup>. The additive noise on each time step is an i.i.d Gaussian with variance  $\Delta t$ , 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  $\lambda$ , which has the unit of spikes/s, is

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

where  $h$  is the softplus rectifying nonlinearity,  $w$  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. The baseline  $b$  accounts for

time-varying influences of neural activity aligned to two events in the trial, stimulus onset and movement onset and slow drifts over minutes across a session. 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). For each session, a single model was fit to all neurons combined across brain regions, rather than fitting separate models to individual regions. Using the model parameters ( $\theta$ ) optimized from this joint fit, we then computed separate posterior probability distributions based on spikes from individual regions:

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

Each region-specific posterior distribution depends on the evidence provided by the spikes from that region alone. The posterior was computed using spikes both before and after that time step (i.e., it is the smoothed, not filtered, posterior).

After fitting MMDDM, we simulated spike trains by drawing samples from the latent process on the same set of trials used to fit the model. Each simulated trial used the original click train of the corresponding real trial. For each trial, we began by computing the prior distribution over the latent state at the first time step,  $p(z\{1\})$ , and drew a sample  $\hat{z}\{1\}$  from this distribution. We then computed the conditional prior for the next time step,  $p(z\{2\} | \hat{z}\{1\}, u\{1\})$ , where  $u\{1\}$  is the click input at time 1, and drew a sample  $\hat{z}\{2\}$ . This process was repeated sequentially to generate a full sample path  $\hat{z}\{1\}, \hat{z}\{2\}, \dots, \hat{z}\{T\}$  from the latent process for the  $T$  time steps in the trial. Given this latent trajectory, we computed the firing rate of neuron  $n$  on time  $t$  is computed as  $\hat{\lambda}(n, t) | \hat{z}(t) = h\{w(n) \cdot \hat{z}(t) + b(n, t)\}$  and drew a spike count from a Poisson distribution with rate  $\hat{\lambda}(n, t)$ .

## Neurally inferred time of commitment (nTc)

The time of decision commitment is defined as the first moment when the posterior probability of the latent variable in MMDDM being either bound exceeds 0.95 and remains above 0.95 for the remainder of the trial.

## 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 ( $\Delta t$ ), and the deviation of the actual difference between ( $R$ ) and left ( $L$ ) click times from the generative

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

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

The psychophysical kernel is specified by the time-varying weight of 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.

## Assessing Dimensionality of $nTc$ residuals

To calculate  $nTc$  residuals, for each region, we separated trials into groups with identical click trains (frozen noise seed) and behavioral choices, computed the mean  $nTc$  within each group, and subtracted it from the corresponding trials. We define a matrix  $X$  where each column corresponds to the  $nTc$  residuals for a single region in each hemisphere (i.e., a feature) and each row corresponds to a single trial (i.e., a sample). Because a value for  $nTc$  was not detected from every region on every trial,  $X$  has many missing elements. We therefore used probabilistic principal components analysis<sup>75</sup> (PPCA) instead of PCA, for its ability to gracefully handle missing data. PPCA models the data as  $X^T = WT^T + \varepsilon$ , where  $W$  is a matrix of orthonormal coefficients (as in PCA),  $T$  are the latent factors (analogous to the scores in PCA), and  $\varepsilon = \nu * I(k)$  is an isotropic error term that accounts for unexplained variability in the data. For  $\varepsilon$  to be non-zero,  $T$  must have rank  $k$  that is lower than the original data. Otherwise, PPCA simply reduces to regular PCA. Under the model specified by PPCA,  $X^T \sim N(0, W * W^T + \nu * I(k))$ . The parameters are estimated using EM. We used the PPCA function in MATLAB R2024a (Mathworks, MA, USA) to fit the model to the data. Principal component variances and loadings are computed in the same way as described above for PCA applied to the DV residuals, except in this case we use the eigenvalues of the fitted model estimate of  $X$ .

To generate predictions for the dimensionality of the brain-wide  $nTc$  residuals under assumptions of different numbers of latent decision processes, we used simulations from MMDDM (see “Multi-mode drift-diffusion model (MMDDM)” in Methods) and applied the same procedure for computing  $nTc$  from the simulated spike trains as for the real data. Simulations were run on the same set of trials used to fit the model, preserving both the number of trials and the click train stimuli.

To model a single latent process, we simulated spike trains across all regions using a single run of the MMDDM. To model two latent processes, we simulated two independent runs for each trial and randomly assigned each brain region to use either the first or the second run consistently across all trials (i.e., a region assigned to the first run uses only latent trajectories from that run on every trial). To model fully independent latent processes for each region, MMDDM simulations were computationally impractical, since this would require simulating as many independent runs as brain regions. Instead, we used the same shuffling approach used for the DV residuals, randomly shuffling time points of  $X$  separately for each region (columns of  $X$ ).

## Quantification of plateau in decision-related activity at decision commitment

To quantify the rise and then abrupt plateau of the trial-average DVs aligned to nTc-M1 (Fig. 6C), we fit the trial-average DVs for each region on each session using piecewise linear regression using two segments (i.e. a bilinear model). Specifically, a value of the mean DV at time  $t$  was modeled as:

$$DV(t) = a_1(t - t_o) + c, \text{ for } t < t_o$$
$$DV(t) = a_2(t - t_o) + c, \text{ for } t > t_o$$

Where  $t_o$  is the time of the breakpoint between the two line segments,  $a_1$  and  $a_2$  are the slopes of the two line segments and  $c$  is a model constant, in this formulation parameterized as the predicted value of  $DV(t_o)$  (the mean DV at the breakpoint). The four model parameters ( $t_o$ ,  $a_1$ ,  $a_2$ , and  $c$ ) were found that minimized the mean squared error, using the *lsqlcurvefit* function in MATLAB R2024a (Mathworks, MA, USA).

## Statistical tests

Unless otherwise stated, error bars and p-values are calculated non-parametrically, using bootstrap resampling either across trials or hierarchically across sessions. All p-values were computed using two-sided tests and not adjusted for multiple comparisons, unless otherwise specifically stated.

To test whether nTc-M1 was more likely or occurred earlier on trials with stronger evidence ( $[\#R-\#L]/[\#R+\#L]$ ), we performed separate regression analyses restricted to the tertile of trials with longest stimulus durations. Logistic regression was used to model the binary outcome of whether nTc-M1 occurred on a given trial as a function of evidence strength, and for trials where it occurred, linear regression was used to model the latency of nTc-M1 as a function of evidence strength. Each model included only evidence strength and a constant term and no additional covariates.