Direct observation of the neural computations underlying a single decision

Natalie A Steinemann[1,2,‡], Gabriel M Stine[1,2,‡], Eric M Trautmann[1,2], Ariel Zylberberg[1,2], Daniel M Wolpert[1,2] & Michael N Shadlen[1,2,3,4,*]

[1] Department of Neuroscience, Columbia University, New York, NY 10027, USA
[2] Zuckerman Mind Brain Behavior Institute, Columbia University, New York, NY 10027, USA
[3] Howard Hughes Medical Institute, Columbia University, New York, NY 10027, USA
[4] Kavli Institute for Brain Science, Columbia University, New York, NY 10027, USA
[‡]These authors contributed equally to this work.
[*]corresponding author

Neurobiological investigations of perceptual decision-making have furnished the first glimpse of a flexible cognitive process at the level of single neurons[1,2]. Neurons in the parietal and prefrontal cortex[3–6] are thought to represent the accumulation of noisy evidence, acquired over time, leading to a decision. Neural recordings averaged over many decisions have provided support for the deterministic rise in activity to a termination bound[7]. Critically, it is the unobserved stochastic component that is thought to confer variability in both choice and decision time[8]. Here, we elucidate this stochastic, diffusion-like signal on individual decisions by recording simultaneously from hundreds of neurons in the lateral intraparietal cortex (LIP). We show that a small subset of these neurons, previously studied singly, represent a combination of deterministic drift and stochastic diffusion—the integral of noisy evidence—during perceptual decision making, and we provide direct support for the hypothesis that this diffusion signal is the quantity responsible for the variability in choice and reaction times. Neuronal state space and decoding analyses, applied to the whole population, also identify the drift diffusion signal. However, we show that the signal relies on the subset of neurons with response fields that overlap the choice targets. This parsimonious observation would escape detection by these powerful methods, absent a clear hypothesis.

Introduction

Difficult decisions require our brains to acquire many samples of evidence before making a choice. Evidence from psychophysics[9–11] and neurophysiology[8] support the idea that perceptual decisions are explained by a quantity, termed a decision variable (DV), which evolves as the sum of accumulated signal and noise—that is, deterministic drift plus stochastic diffusion. Critically, this formulation predicts that each decision is explained by a unique *drift-diffusion* trajectory that determines the outcome of the decision and the time of its completion. Recordings from single neurons in parietal and prefrontal association cortex in monkeys provide indirect support for this quantity, based on second order statistical properties from many decisions[12] or successful application of machine learning to predict decision outcome[13]. However, to date, drift-diffusion has not been observable on single trials, making it impossible to assess whether individual decisions are governed by this process.

Here we provide the first direct evidence for a drift-diffusion process underlying single decisions. We recorded simultaneously from up to 203 neurons in the lateral intraparietal area (LIP) while monkeys made perceptual decisions about moving random dots. Using a variety of dimensionality reduction techniques, we show that a drift-diffusion signal can be detected in such populations on individual trials, and that this signal satisfies the criteria for a decision variable. Notably, the signal of interest is dominated by a subpopulation of neurons whose response fields overlap one of the choice targets, consistent with earlier single neuron studies[1,7,8,12].

Results

Two monkeys made perceptual decisions, reported by an eye movement, about the net direction of dynamic random dot motion (RDM; Fig. 1a). We measured the speed and accuracy of their decisions as a function of motion strength (Fig. 1b). The choice probabilities and the distribution of reaction times (RT) are well described by a bounded drift diffusion model (Fig. 1b,c). The fits of this model also specify the evolving probability distributions of the decision variable. They will guide our interrogation of the neural data.

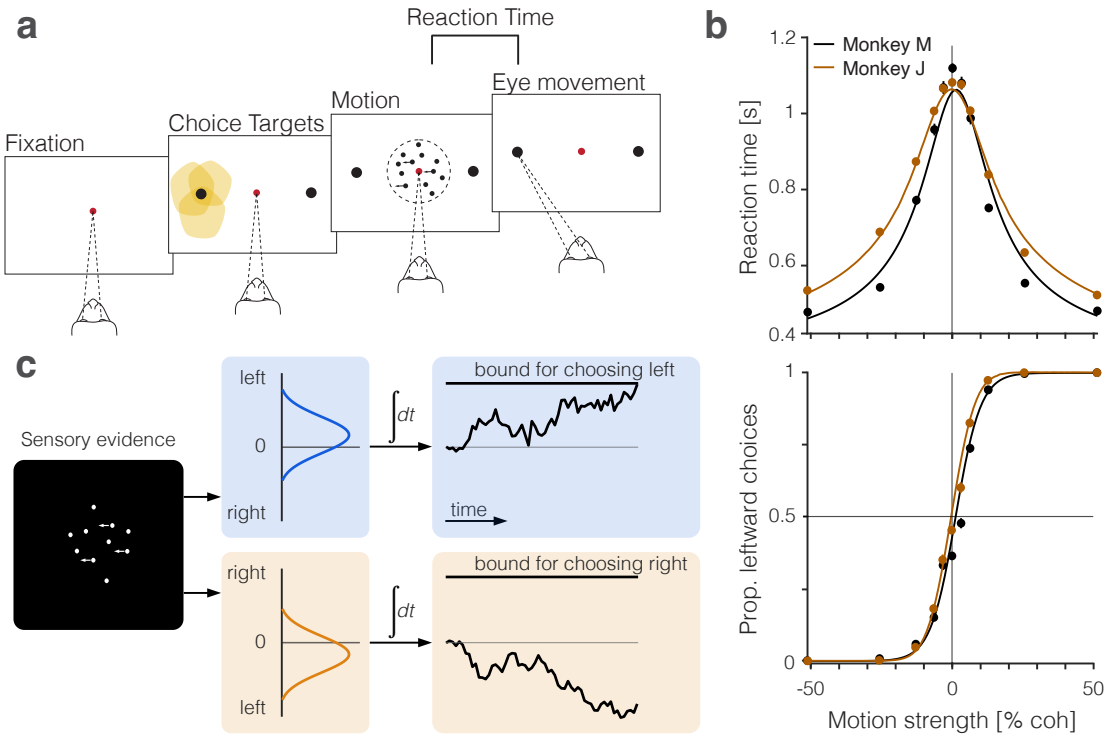

Figure 1: **Perceptual decisions are explained by the accumulation of noisy evidence to a stopping bound. a,** Random dot motion discrimination task. The monkey fixates a central point. After a delay, two peripheral targets appear, followed by the random dot motion. When ready, the monkey reports the net direction of motion by making an eye movement to the corresponding target. Yellow shading indicates the response fields of a subset of neurons in area LIP that we refer to as T_{in} neurons. **b,** Mean reaction times (top) and proportion of leftward choices (bottom) plotted as a function of motion strength and direction, indicated by the sign of the coherence (positive is leftward, toward T_{in}). Data are from all sessions from monkey M (black, 9684 trials) and monkey J (brown, 8142 trials). Solid lines are fits of a bounded drift-diffusion model. **c,** Drift-diffusion model. The decision process is modelled as a race between two negatively-correlated ($\rho = -0.71$) accumulators: one integrating momentary evidence towards T_{in} and another towards T_{out} choices. The decision is terminated when one accumulator hits its positive bound.

During the task, we used Neuropixels probes to record simultaneously from 54–203 neurons in area LIP, representing a random and unbiased sample of neurons within the targeted recording location (Fig. 2a). This location, deep in the intraparietal sulcus was previously inaccessible using existing high density recording devices. We conducted eight sessions in two monkeys (1696–2894 trials per session). In what follows, we highlight one representative session from monkey M.

Our initial analyses focus on the subset of neurons with response fields that overlap the contralateral choice target (termed T_{in}) and show spatially selective persistent activity on an oculomotor delayed response task[14]—the type

of neuron previously selected for study in single-cell recording experiments. Fig. 2b shows the activity of 17 such neurons, averaged across neurons and trials of the same motion coherence in this representative session. The average firing rates from this ensemble replicate earlier findings: ramp-like trajectories during decision formation with buildup rates that depend on the strength and direction of motion and stereotyped firing rates just preceding saccadic choices for T_{in}. The motion-dependent separation of responses is thought to mark the beginning of decision formation in LIP, $\sim$200 ms after onset of the RDM. While the build-up is thought to reflect the deterministic component of the accumulation of momentary evidence, the diffusion process is elusive. This is because noise is suppressed through averaging across trials, and for a single neuron, the spikes are too sparse to permit inferences of time-varying rate over the appropriate scale. With a pool of 17 neurons, however, it is now possible to measure the rate on single trials.

Each trace in Fig. 2c shows the average firing rate of the 17 T_{in} neurons, $S_{T_{in}}(t)$, on single trials. They are redolent of drift-diffusion. The traces contain many alternations in the sign of their derivative (even after smoothing), consistent with an accumulation of independent samples of positive and negative numbers. The traces spread apart by $\sim\sqrt{t}$, because the variance of a running sum of independent identically distributed random samples is the sum of the individual variances ($\sigma^2 \propto t$). This relationship holds for $\sim$100 ms for the weakest motion strength. As shown below, the attenuation of the spreading at later times is a sign of a termination bound that limits the range of the accumulation. It thus appears that the ramp-like averages in Fig. 2b belie drift-diffusion processes on single trials.

We hypothesized that the single trial average firing rate, $S_{T_{in}}$, approximates the decision variable[8]. If so, the diffusion component should explain the variability of choice and reaction time for trials sharing the same direction and motion strength. Specifically, (*i*) a sample of $S_{T_{in}}(t)$ should be predictive of choice and correlate inversely with the RT on trials that result in T_{in} choices; (*ii*) later samples ought to predict choice better and correlate more strongly (negatively) with RT than earlier samples, and (*iii*) later samples should contain most, if not all, of the information present in the earlier samples. As we will show, these predictions are borne out.

As shown in Fig. 2d, a sample of $S_{T_{in}}(t)$ taken near the beginning of decision formation is negatively correlated with the RT on that trial (black line), even though we are only considering trials with long RTs (≥ 670 ms). The correlation is statistically significant for $t \geq 250$ ms, which is only 50 ms after the responses begin to exhibit a dependency on motion strength and direction. The magnitude of the correlation is stronger with later samples, and critically, the later samples contain most of the information in the earlier sample, thereby rendering a reduced partial correlation, given the later sample (dashed line). For example, $r[S_{T_{in}}(0.35), RT] = -0.20$, whereas 200 ms later $r[S_{T_{in}}(0.55), RT] = -0.32$, and the partial correlation $r[S_{T_{in}}(0.35), RT \mid S_{T_{in}}(0.55)] = -0.12$. The later sample mediates 66% of the earlier sample's coefficient of determination (R^2). The same pattern holds for the choice (Fig. 2e). Early samples of $S_{T_{in}}(t)$ exert positive leverage (β) on the probability of a T_{in} choice; later samples have more leverage and explain away the leverage of earlier samples from logistic regression (GLM). We compute a simple index of mediation using the change in β when the later time point is absent or included in the regression (see Methods). In the example in Fig. 2e, the mediation at $t = 0.35$ s is complete (across-session mean $\pm$ sem $= 74 \pm 9\%$; $p < 0.001$, Fig. 2f, Fig. S1). Thus $S_{T_{in}}(t)$ is an evolving representation that controls the choice and decision time on a single trial—the decision variable (DV), consistent with the hallmarks enumerated above.

The representation of the DV is concentrated

Up to now, our focus has been on the same neurons we would have selected in single-neuron recording experiments, in this case screened *post hoc* for spatially selective persistent activity associated with the contralateral choice-target. We pursued several strategies to assess whether the representation is more broadly distributed among the population of simultaneously recorded neurons. The first is a targeted approach to find the weighting of all

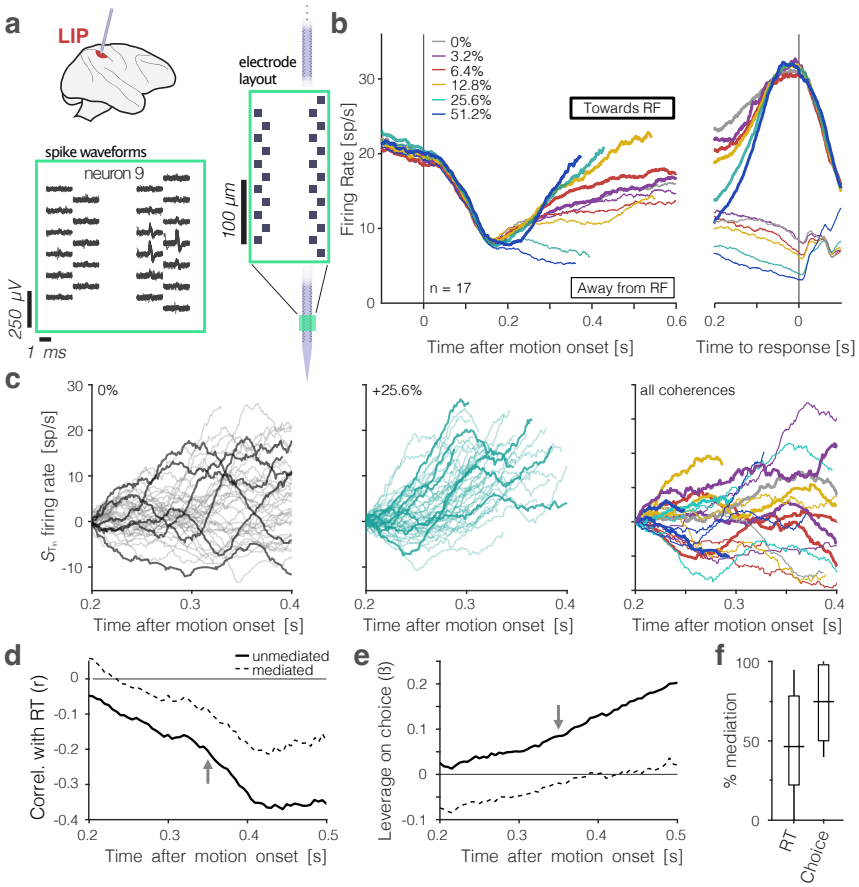

Figure 2: **Representation of the decision variable on single trials. a,** Neural recordings were obtained from area LIP in the intraparietal sulcus of the right hemisphere, using a prototype 45 mm Neuropixels probe capable of recording simultaneously from 384 of 4416 contacts (see Methods). The large green rectangle contains voltage traces from a single neuron, which typically appear on several contacts. In the example session, 17 neurons had response fields that overlapped the left (contralateral) choice target (T_{in}). **b,** Trial-averaged mean firing rates of 17 T_{in} neurons grouped by motion strength (color) and direction (thickness). *left*, responses aligned to stimulus onset; *right*, responses aligned to saccade (error trials are excluded from the means). The traces reflect the direction and strength of motion from $\sim$200 ms after motion onset. **c,** Single trial firing rates, $S_{T_{in}}$, approximate drift-diffusion; *left*, 100 samples of $S_{T_{in}}$ baseline corrected for 0% coherence trials. Thick lines are intended to aid tracing of single trajectories; *middle*, same as left for 25.6% coherence; *right* Examples of two trials of $S_{T_{in}}$ for each coherence and direction of motion after subtracting the within-coherence mean activity. **d,** Correlation of neural activity, $S_{T_{in}}$, with RT (solid). Only trials with RT$\geq$670 ms resulting in T_{in} choices are include in the analysis. This correlation was partially mediated (dashed; partial correlation) when accounting for activity at a later time point ($t = 0.55$ s). **e,** Leverage of neural activity, $S_{T_{in}}$, on choice (β) also show mediation as in **d**. **f,** Percent mediation on all sessions using the values at $t = 0.35$ s (gray arrows in **d** and **e**). Boxes show median and interquartile range; whiskers end at minimum and maximum.

191 neurons that best approximates the bounded drift-diffusion process inferred from behavior, while remaining agnostic about the role of the 17 T_{in} neurons.

We fit the choice and RTs on the session using a bounded drift-diffusion model, structured as a race between two partially anticorrelated drift-diffusion processes (Fig. 1c). The architecture approximates the neural organization—populations of neurons that accumulate evidence for right vs. left and populations that accumulate evidence for left vs. right. Each has an upper stopping bound[8]. The fits furnish an estimate of the time-varying distribution of the DV associated with each motion strength and direction. The estimates incorporate the fact that trials terminate

by reaching a positive bound in either of the competing processes. This gives rise to nonlinear functions of the expectation and variance as a function of time (Fig. 3a). We optimized the weights for each neuron such that the time-varying distribution of the weighted average firing rates best matched the model's predicted distribution (see Methods). The vector of weights, **d**, define the direction in the 191-dimensional neuronal state space that best capture the one-dimensional drift-diffusion decision variable. As shown in Fig. 3b, the projection of activity on **d**, $S_{\mathbf{d}}(t)$, exhibits time dependent mean and variance similar to the model predictions.

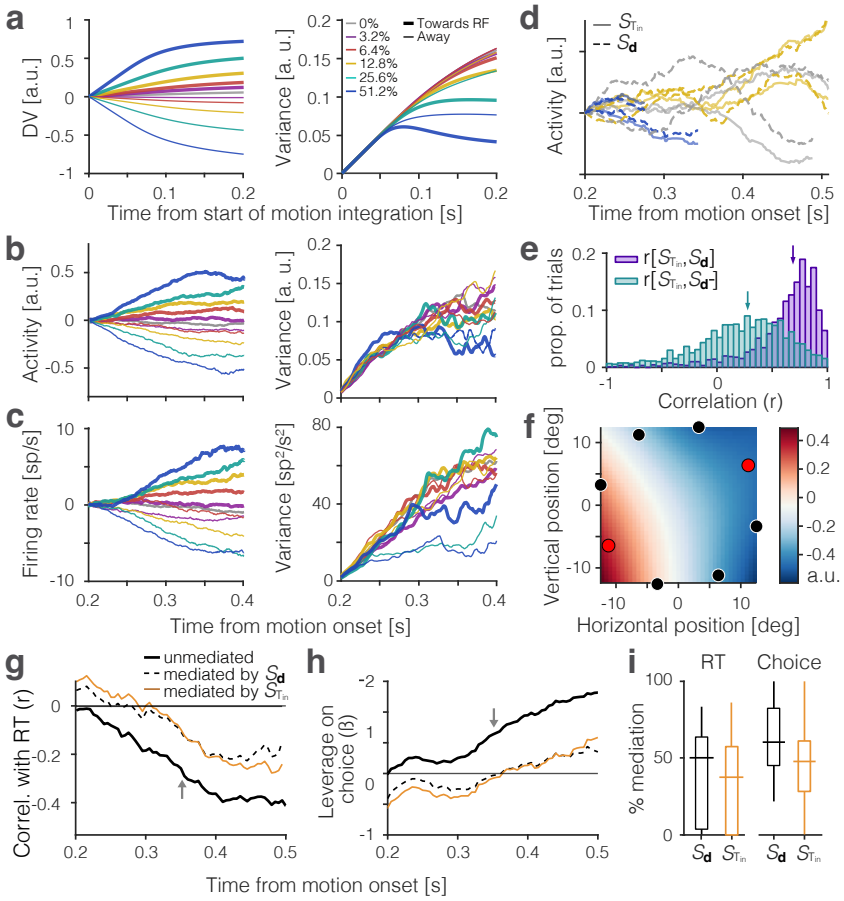

Figure 3: **Neuronal state space representation of diffusion. a,** Predicted mean and variance of single-trial diffusion paths derived from the fits to behavior (Fig. 1b). The conditionalization imposed by the stopping bounds are responsible for the departure from linear functions of time (both columns). For plotting purposes only we incorporate a lower bound on DV in the model as spike rate cannot be negative (see Methods). **b,** Mean and variance of population single-trial responses, projected on the diffusion vector. **c,** Mean and variance of the $S_{\mathrm{T_{in}}}(t)$ single trial responses. **d,** Example pairs of single-trial diffusion paths rendered as $S_{\mathrm{T_{in}}}(t)$ (solid) and $S_{\mathbf{d}}(t)$ (dashed). **e,** Histogram of within-trial correlations between pairs of trials like those in panel **d**. Correlations between $S_{\mathrm{T_{in}}}$ and $S_{\mathbf{d}}$ (purple) and between $S_{\mathrm{T_{in}}}$ and $S_{\mathbf{d}}$ with the $\mathrm{T_{in}}$ neurons removed ($S_{\mathbf{d}-}$; green). Arrows indicate median correlations. **f,** Heat map shows the population response projected on the diffusion vector, determined during the waiting period of memory saccades. Circles show the location of targets in the oculomotor delayed response task. Red circles show the location of the choice targets. This shows a response field overlapping the left choice target. **g,h,** Correlation of $S_{\mathbf{d}}(t)$ with RT and choice on trials with RT>670 ms. Solid and dashed black curves show simple and partial correlations, as in Fig. 2d,e. The third curve (copper) shows mediation by $S_{\mathrm{T_{in}}}(0.55)$ **i,** Mediation by $S_{\mathbf{d}}$ and $S_{\mathrm{T_{in}}}$ across eight sessions. Boxes show median and interquartile range; whiskers end at minimum and maximum.

The theoretical means and variances are also approximated by the 17 $\mathrm{T_{in}}$ neurons (Fig. 3c). Indeed, the unweighted average firing rates, $S_{\mathrm{T_{in}}}(t)$ are similar to $S_{\mathbf{d}}(t)$ on individual trials (Fig. 3d). The histograms in Fig. 3e show correlations for all trials from this session. They show a strong correlation between $S_{\mathbf{d}}$ and $S_{\mathrm{T_{in}}}$ (median

$r[S_{\mathrm{T_{in}}}, S_{\mathbf{d}}] = 0.70$), which is reduced significantly when the $\mathrm{T_{in}}$ neurons are excluded from **d** to construct the 174-dimensional $\mathbf{d}^-$ (median $r[S_{\mathrm{T_{in}}}, S_{\mathbf{d}^-}] = 0.27$, $p < 10^{-6}$, KS-test; Table S1), indicating a large contribution of the $\mathrm{T_{in}}$ neurons to **d**. In fact, the population activity projected onto **d** renders spatially selective persistent activity to targets in the same region of the visual field as the $\mathrm{T_{in}}$ neural response fields (Fig. 3f; 7 of 8 sessions, Fig. S2). It is therefore unsurprising that early samples of $S_{\mathbf{d}}(t)$ exhibit correlations with choice and RT. Nor is it surprising that a sample of $S_{\mathbf{d}}(0.55)$ mediates the correlation of earlier samples of $S_{\mathbf{d}}(t)$ (dashed curve). Surprisingly, however, a sample of $S_{\mathrm{T_{in}}}(0.55)$ also mediates earlier samples of $S_{\mathbf{d}}(t)$ (copper curve). We observe this pattern of mediation across measures in all but one session for which a later sample of $S_{\mathrm{T_{in}}}$ does not mediate earlier activity of $S_{\mathbf{d}}$ (Fig. S3).

Clearly, the 17 $\mathrm{T_{in}}$ neurons play an important role in rendering the decision. We assume they are but a subset of the larger pool of neurons with response fields that overlap the same choice target—hence the partial mediation. Yet, they do not stand out on the basis of their weighting coefficients (Fig. S7)—not even when the weights are penalised so as to be sparse (L1 regularization). Thus without prior knowledge, we might not have recognized the simple logic of the decision process represented by neurons that inform the saccadic report. Taken together these analyses support the hypothesis that the DV is a one dimensional process[15] that is largely explained by the 17 $\mathrm{T_{in}}$ neurons.

Data driven approaches

Multineuron recordings present an opportunity to discover unanticipated features of data, beyond the imagination of the scientist and existent hypotheses. Such efforts are aided by a growing set of tools, devised to facilitate discovery of neural computations realized in the neuronal state space. Here we ask whether three data driven methods would discover diffusion and the importance of the $\mathrm{T_{in}}$ neurons had we not known about the latter and looked for the former. The answers to these questions may be relevant to many types of experiments that seek to exploit large scale neural recordings.

We first trained a logistic decoder to predict the monkey's choices using the full population of 191 neurons in the sample session. We used all coherences and allowed a unique set of decoding weights at each time step (similar to Ref. 13; see Methods). As shown in Fig. 4a, the cross-validated accuracy of this *optimal* decoder increased as a function of time from motion onset. To estimate the upper bound of decoding, we performed the same analysis on data simulated with the diffusion model fit to the behavior. The peak performance of this decoder is 72% (Fig. 4a). The population decoder rivals this performance, and remarkably, so does the application of a threshold to $S_{\mathbf{d}}(t)$ and $S_{\mathrm{T_{in}}}(t)$ to predict choice. Indeed, removal of the 17 $\mathrm{T_{in}}$ neurons from the population reduces decoding performance far in excess of what would be expected by removing a random 17 neurons (average reduction in decoding accuracy of 8% and $1 \pm 1\%$, mean $\pm$ sd, respectively).

While the decoder assigns a new set of weights to each time point, Fig. 4b shows that weights derived as early as 300 ms after motion onset (100 ms from the beginning of decision-related activity) perform nearly as well when those same weight are applied at all other times[16]. This observation is consistent with the characterization of the process as one-dimensional, supported by the mediation analyses (*e.g.*, Fig. 3g,h). We applied decoder weights derived at $t = 0.35$ s to define a direction in the neuronal state space to produce time dependent signals on single trials. The across-trial means exhibit coherence-dependent drift and the single-trials bear semblance to diffusion. The latter have leverage on choice, by design, and they correlate inversely with RT. Both effects (at $t = 0.35$ s) are mediated by a later sample of the same signal (*self*, RT: 67.4%; choice: 62.8%) and also by $S_{\mathrm{T_{in}}}(0.55)$ (RT: 81.1%; choice: 66.9%). The top row of Fig. 4f shows the mean percent mediation for the eight sessions (see also Fig. S4).

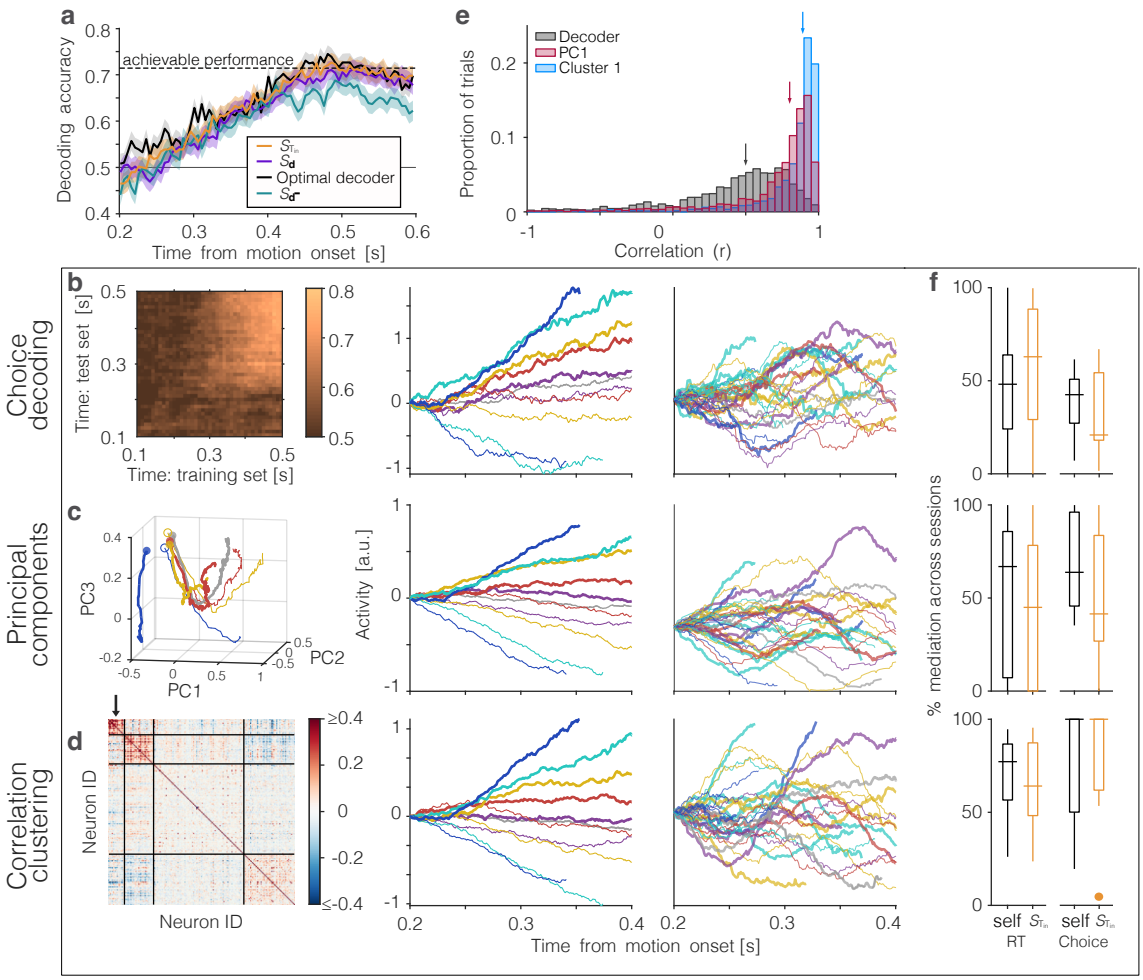

Figure 4: **Data driven approaches. a,** Logistic choice-decoder. The decoder is trained to predict the monkey's choice from a weighted sum of firing rates from each neuron, using a random half of the trials. Decoding accuracy is established on the left out trials (black trace). The best achievable performance (horizontal dashed line) is estimated from simulation of a drift-diffusion model. Colored traces are accuracy of predictions using the 17 T_{in} neurons ($S_{T_{in}}$), $S_{\mathbf{d}}$ and $S_{\mathbf{d}-}$ (re-derived diffusion direction without the T_{in} neurons). Shading indicates s.e. **b,** The decoder defines a direction in neuronal state space. Heat map shows the degree to which a decoder trained at one time (abscissa) performs at other times (ordinate). The weights at $t = 0.35$ s are used to generate single-trial responses analyzed in columns 2 & 3. *2nd column*, Trial averaged signals grouped by motion strength (color) and direction (line style) with same convention as Fig. 2a. *3rd column*, Examples of detrended single-trial responses. **c,** Principal components analysis. The image shows trial average activity along the first 3 PCs (200 ms after motion onset until 100 ms before median RT) for 4 motion strengths (both directions for non-zero coherence). Filled and empty circles indicate the beginning of the trials for thick and thin traces, respectively (same convention as Fig. 2b). This depiction is less interpretable than those in the next columns. Middle and right columns use single-trial signals rendered by the PC-1. Same as in **b**. **d,** Correlation clustering. *Column 1*, Pairwise correlations between the neurons in the example session. Vertical and horizontal lines separate four spectral clusters based on r-values. Columns 2–3 use single-trial signals rendered by Cluster-1 (arrow in *1st column*) as in **b**. **e,** Similarity of single-trial signals from $S_{T_{in}}(t)$ and the three methods in **b**–**d**. Arrows indicate median correlations. **f,** Summary of mediation by a later sample of the signal in the corresponding row (*self*) or $S_{T_{in}}$ on all sessions using the same RT restriction and sample time points as Fig. 3i. The stray point in the box-whisker plot (bottom) is an outlier.

We also applied Principal Components Analysis (PCA) to trial-averaged standardized firing rates, using the epoch from 200 ms to 500 ms from motion onset (as in Fig. 2b). The first three PCs (Fig. 4c) capture 57% of the total

variance. We use just the first PC (39% of variance) to render single-trial signals. The trial averaged means show clear dependence on motion strength and direction and single-trial traces are redolent of diffusion. They are positively correlated with $S_{\mathrm{T_{in}}}$ (Fig. 4e), and their correlation with behavior is mediated by $S_{\mathrm{T_{in}}}$ on most sessions (Fig. 4f, Fig. S5). We obtain similar results using demixed PCA[17] (Fig. S5).

We conclude that PCA has the potential to reveal drift-diffusion on single trials in most of our experiments. This might seem surprising initially, as the PCs are derived from trial-averaged firing rates, which suppress the diffusion component. However, the PCs assign weight to the neurons that represent the drift, and those averages comprise drift-diffusion on single trials. However, like the choice decoder, which also mimics diffusion, PCA (and dPCA) do not expose the simplicity of the organization around the $\mathrm{T_{in}}$ neurons (Fig. S7). An inquisitive experimenter could discover this organization by noticing a visual response to the targets and a perisaccadic response associated with $\mathrm{T_{in}}$ choices.

Is there any feature of the data that might provide a clue to the data analyst about the importance of the $\mathrm{T_{in}}$ neurons? Although it is not obvious on single sessions, pooling the eight sessions reveals that the largest weights of **d**, the choice-decoder and PC1 are assigned to $\mathrm{T_{in}}$ neurons (Fig. S7). However, this does not reveal the underlying logic. Inspired by Kiani *et al.*[18], we conducted a spectral analysis of the pairwise correlation matrix using spike counts in a short epoch before the onset of RDM (Fig. 4d). Clusters of positively correlated neurons are apparent along the main diagonal. We define Cluster-1 (arrow) as the cluster with the strongest single-trial correlations with $S_{\mathrm{T_{in}}}$. Thus, it is unsurprising that the single-trial traces from Cluster-1 correlate strongly with $S_{\mathrm{T_{in}}}$ (Fig. 4e). In the example session, this cluster also stands out as the one with the largest average correlation ($\bar{r}$), but this was not a reliable indicator across sessions. That said, 9 of the 16 neurons in this cluster are $\mathrm{T_{in}}$ neurons, and across all sessions, $43 \pm 11\%$ of $\mathrm{T_{in}}$ neurons are in Cluster-1 (Table S2). Thus, without prior knowledge of the $\mathrm{T_{in}}$ neurons, spectral clustering might bring them to our attention. Many are in the clusters with the stronger $\bar{r}$.

Taken together, we conclude that even without a hypothesis, it is possible to discover drift-diffusion and its status as a decision variable. Absent knowledge of $\mathrm{T_{in}}$, one could recognize that it is represented by a low (one) dimensional subspace of the neuronal state space. While accurate, the characterization would obfuscate the simple biological organization.

Discussion

We have directly observed the neural representation of a stochastic decision variable—the accumulation of noisy evidence that determines choice and RT on single decisions. The signal is evident in the mean firing rate of small pools of neurons with response fields that overlap a choice target ($\mathrm{T_{in}}$ neurons) and also in the weighted sum of the full population of 54–203 sampled neurons, the neuronal state space (NSS), which includes these $\mathrm{T_{in}}$ neurons. It has the hallmarks of a decision variable: (*i*) it is correlated with both choice and RT, (*ii*) the magnitude of the correlations increase over time, as the decision is forming, (*iii*) later samples of the signal mediate correlations of earlier samples, (*iv*) its time-varying mean and variance (for each motion strength) are consistent with a drift-diffusion process, and (*v*) although not shown here, the same signals lead to termination of the decision process[19]. We focused on the early epoch of decision formation because this is where the drift-diffusion signal is least distorted by the stopping bound, trial dropout, and response features associated with saccade initiation[19].

The finding validates the conclusions from single-neuron studies, which inferred one-dimensional diffusion dynamics using indirect methods[12,15]. It also complements the choice-decoding strategy applied in Peixoto et al. (Ref. 13) to population recordings from the prefrontal and motor cortex. Choice decoders, like ours and Peixoto's, are machine learning tools designed to predict the choice, but they are not the same as the actual DV. Unlike a choice-decoder, a real DV does not represent the probability of a choice but the signal that *determines* the choice and

RT. The choice and RT are stochastic because the drift-diffusion signal is the accumulation of noisy evidence. While the stochastic drift-diffusion cannot be detected in trial-averaged firing rates, we can now recognize that the coding of a stochastic quantity conforms to the same principles as the encoding of simpler, deterministic quantities, such as direction selectivity—that is, pooling of spikes from many weakly correlated neurons with similar selectivity[20,21] (but see Ref. 22). This explains why the sample of T_{in} neurons mediates strongly and also why it does not mediate completely, as these are only a sample of the full population.

A decision variable is just one of many neural representations that underlies a single instance of a mental process. The capacity to record simultaneously from an ever increasing number of neurons opens the possibility of studying the mechanisms underlying such instances. In this spirit we are encouraged by the capacity of purely data driven methods to "discover" diffusion. However, none of the methods we tried would identify the important T_{in} neurons. That is, they do not reveal the spatial organization of the response fields. Based on the experience here, one could imagine a path to their discovery by testing neurons identified by correlation-based clustering or large weights in PC1, say. For example one might notice that the responses of such neurons are correlated with the single trial diffusion signal rendered by PC1, and in addition they respond to both the onset of the choice targets and just preceding contraversive saccades. It would be natural to map the neural response fields and recapitulate the single neuron studies that discovered spatially selective persistent activity[14]. Thus, while they do not guarantee discovery of the underlying neural organization, data driven approaches may point to observations that lead to discovery through hypothesis-driven inquiry and experiments.

These hypothesis-driven steps are essential, in our view, if we desire an explanation of brain function in the language of biology, in addition to mathematics. The present findings identify the DV with a particular functional cell type in LIP. The ability to observe the idiosyncratic, time-varying brain signal that gives rise to one decision at a moment in time opens a new lens to view neural computations that are only revealed on single instances. For example, the capacity to observe a decision variable on a single trial opens the possibility of elucidating interactions between LIP and other nodes of the network involved in decision making. As we show in the companion paper[19], single-trial firing rates in functionally connected T_{in} neurons in LIP and the superior colliculus perform distinct computations, which form and terminate the decision, respectively. One wonders what biological parsimony might hide in other neuronal state spaces.

Methods

Ethical approval declarations. Two adult male rhesus monkeys (*Macacca mulatta*) were used in the experiments. All training, surgery, and experimental procedures complied with guidelines from the National Institutes of Health and were approved by the Institutional Animal Care and Use Committee at Columbia University. A head post and recording chamber were implanted using aseptic surgical procedures.

Behavioral tasks. The monkeys were trained to interact with visual stimuli presented on an CRT video monitor (Vision Master 1451, iiyama; viewing distance 57 cm; frame rate 75 Hz). They were trained to control their gaze and make saccadic eye movements to peripheral targets to receive a liquid reward (juice). The direction of gaze was monitored by an infrared camera (EyeLink 1000; SR Research, Ottawa, Canada; 1 kHz sampling rate). The tasks involve stages separated by random delays, distributed as truncated exponential distributions

$$f(t) = \begin{cases} \frac{\alpha}{\lambda} e^{-\frac{t-t_{\text{min}}}{\lambda}} & t_{\text{min}} \leq t \leq t_{\text{max}} \\ 0 & \text{otherwise} \end{cases} \tag{1}$$

where t_{min} and t_{max} define the range, λ s is the time constant, and α is chosen to ensure the total probability is unity. Below, we report the range and the exponential parameter λ. Note that because of truncation, the expectation $\mathbb{E}(t) < t_{\text{min}} + \lambda$.

In the *random dot motion* (RDM) task, monkeys are trained to make decisions about the net direction of motion in a dynamic random dot motion display. The RDM is confined to a circular aperture (diameter 5 dva; degrees visual angle) centered on the fixation point (dot density 16.7 dots·dva^{-2}·s^{-1}). The task flow is shown in Fig. 1a. The random wait from onset of the choice targets to onset of the RDM is (0.25–0.7 s, λ =0.15). The direction (bottom left or top right) and strength of motion are determined randomly from $\pm\{0, 3.2, 6.4, 12.6, 25.6, 51.2\}$% coherence (coh). The sign of the coherence indicates direction (positive for leftward). The values control the probability that a dot plotted on frame n will be displaced by Δx on frame $n+3$ (Δt =40 ms), as opposed to randomly replaced. The displacement is consistent with velocity 5 dva·s^{-1} (see Ref. 7 for additional details). The monkey is rewarded for choosing the correct target (trials with 0% coh were rewarded randomly). Errors are punished by extending the intertrial interval by up to 3 s (see Ref. 19 for additional details). On approximately half of the trials, a 100 ms pulse of weak motion ($\pm$4% coh) was added to the RDM stimulus at a random time (0.1–0.8 s, λ =0.4) relative to RDM onset (similar to Ref. 23). Monkey M performed 9684 trials (5 sessions); monkey J performed 8142 trials (3 sessions). The data are also analyzed in a companion paper[19].

In the *oculomotor delayed response* task (ODR;[14,24]), one target was flashed briefly (200 ms at a pseudorandom location in the visual field. After a variable delay (0.4–1.1 s, λ =0.3 for monkey M, 0.5–1.5 s, λ =0.2 for monkey J), the fixation point was extinguished and the monkey made a saccade to the remembered location of the target ± 2.5 dva to receive a reward. This task was conducted to provide a rough characterization of the neural response fields during the visual, perisaccadic and delay epochs. They were also used to assess the stability of the recording over the session. Neurons were designated T_{in} if they exhibited spatially selective persistent activity at the location of the response target in the visual hemifield contralateral to the recorded hemisphere. The example session contained 17 T_{in} neurons, and we refer to the unweighted mean firing rate as $S_{\text{T}_{\text{in}}}$. These analyses were conducted *post hoc*, after spike sorting.

Behavioral analyses. We fit a variant of the drift-diffusion model (Fig. 1c) to the choice-RT data from each session. Details of the model and the fitting method are described in Ref. 25. The model constructs the decision process as a race between two accumulators: one accumulating evidence for the T_{in} choice and against the T_{out}

choice (*e.g.*, left minus right) and one accumulating evidence for a T_{out} choice and against a T_{in} choice (*e.g.*, right minus left). The decision (T_{in} or T_{out}) is determined by the accumulator that first exceeds its positive decision bound, at which point the decision is terminated. The races are negatively correlated with one another, owing to the common source of noisy evidence. We assume they share half the variance, $\rho = -\sqrt{0.5} \approx -0.71$, but the results are robust to a wide range of reasonable values. The decision bounds are allowed to collapse linearly as a function of time, such that

$$B(t) = B_0 - \alpha t, \quad \alpha \geq 0 \tag{2}$$

We used the method of images to compute the probability density of the accumulated evidence (x) for each accumulator as a function of time (t) using a time-step of 1 ms. We assumed that $x = 0$ at $t = 0$. The decision time distributions rendered by the model were convolved with a Gaussian non-decision time distribution, which summarizes sensory and motor delays, to generate the predicted RT distributions. The model has 6 parameters in total: κ, B_0, α, μ_{nd}, σ_{nd}, and C_0, where κ determines the scaling of motion strength to drift rate, C_0 implements bias in units of signed coherence[26], μ_{nd} is the mean non-decision time and σ_{nd} is its standard deviation.

For analytic tractability using the method of images, the model has no lower bound. This leads to a distortion of the mean and variance for the strongly negative coherences (given that spike rates can not go negative). Therefore, for plotting purposes only (Fig. 3a), we implemented a simple version of a lower bound (at $-B_0$). That is, having obtained the distribution of DV for all time points, any density below the lower bound was placed at the lower bound before we calculated the means and variances across time.

Neurophysiology. We used prototype "alpha" version Neuropixels1.0-NHP45 probes (IMEC/HHMI-Janelia) to record the activity of multiple isolated single-units from the ventral subdivision of area LIP (LIP_v[27]). We used anatomical MRI to identify LIP_v and confirmed its physiological hallmarks with single-neuron recordings (Thomas Recording GmbH) before proceeding to multi-neuron recordings. Neuropixels1.0-NHP45 enable recording from 384 out of 4416 total electrical contacts distributed along the 45 mm long shank. All data presented here were recorded using the 384 contacts closest to the tip of the probe (Bank 0), spanning 3.84 mm. Reference and ground signals were directly connected to each other, and connected to the monkey's headpost. A total of 1084 neurons were recorded over eight sessions (54–203 neurons per session). (Table 1).

Session	1*	2	3	4	5	6	7	8
monkey	M	M	M	M	M	J	J	J
trials	1797	1696	2256	1859	2076	2449	2799	2894
neurons	191	90	54	140	107	138	161	203
excluded	5	0	1	2	2	0	1	0
T_{in}	17	13	9	21	19	23	34	27
mean FR (T_{in})	12	8	9	10	16	14	14	14

Table 1: **Information about individual experimental sessions.** *Example session; FR, mean firing rate (sp/s) 75–125 ms after motion onset, across trials and T_{in} neurons; *excluded*, refers to neurons with weights set to zero in some state space analyses (see Methods)

The Neuropixels1.0-NHP45 probe is connected to a standard commercially-available headstage for the Neuropixels1.0 probes, connected via the standard Neuropixels1.0 5m cable to the PCI eXtensions for Instrumentation (PXIe) hardware (PXIe-1071 chassis and PXI-6141 and PXIe-8381 I/O modules, National Instruments). Raw data were acquired using the SpikeGLX software (http://billkarsh.github.io/SpikeGLX/), and single-units were identified offline using the Kilosort 2.0 algorithm[28,29], followed by manual curation using Phy (https://github.com/cortex-lab/phy). The spike times were then synchronized with task events acquired by the experimental control system (Rex[30]) and OmniPlex (Plexon Inc.).

Neural data analysis. The spike times from each neuron are represented as delta functions of discrete time (1 kHz) and convolved with a 50 ms boxcar filter to achieve a coarse representation of firing rate as a function of time, $s_{ij}(t)$, on each trial i and each neuron j. We used an 80 ms boxcar filter to calculate correlation between neuron firing rate and behavior (Fig. 2d–f) because it led to more robust estimates of summary statistics (*e.g.*, %-mediation for different choice of early and late time points).

For analyses that use standardized signals—derivation of **d**, PCA, and choice-decoding—we expressed the rates as z-scores based on the mean and standard deviation of the firing rate 100 ms after motion onset. In some sessions, this led to the exclusion of a small number of very low firing-rate neurons which did not show any activity on any trial in the normalization window (*i.e.*, 75–125 ms after motion onset; Table 1). Those neurons were assigned zero weight and do not contribute to the degrees of freedom in statistical tests. In the example session, this eliminated 5 neurons, thereby reducing the effective number of neurons from 191 to 186. The same uniformly-weighted 50 ms filter was used for data visualization with the exception of examples of single-trial activity (Fig. 2c, Fig. 3d, Fig. 4b–d) for which the raw data was convolved with a truncated Gaussian filter generated using the Matlab function, *gausswin* (width $= 80$ ms, width-factor $= 1.5$, $\sigma \approx 26$ ms).

The expected mean and variance of the decision variable as a function of time and coherence ($\mu_c(t)$ and $\sigma_c^2(t)$, respectively) was furnished by the fits of the behavioral model (see above). The model parameters govern the probability density of the DV for each coherence and time, and $\mu_c(t)$ and $\sigma_c^2(t)$ were computed using the density unabsorbed by the decision bounds. We then found the direction in neural state space (defined by **d**) for which projections onto it—rendering $S_{\mathbf{d}}(t)$—best fit $\mu_c(t)$ and $\sigma_c^2(t)$. Goodness of fit was quantified using the K-L divergence between the predicted and fitted distributions of the DV. For simplicity, we calculated K-L divergence under the assumption that these distributions are Gaussian:

$$KL(p,q) = \log \frac{\sigma_q}{\sigma_p} + \frac{\sigma_p^2 + (\mu_p - \mu_q)^2}{2\sigma_q^2} - \frac{1}{2} \tag{3}$$

where p is the DV distribution given by the behavioral model fit and q is the DV distribution given by $S_{\mathbf{d}}$. The weights comprising **d** were optimized using a *quasi-newton* search algorithm (*fminunc*, MATLAB). The weights were applied to the normalized activity of each neuron and the fits were evaluated within the epoch 200–400 ms after motion onset. Because it is always the case that $\mu_p(0) = 0$ and $\sigma_p^2(0) = 0$, we enforced that the same be true for $\mu_q(0)$ and $\sigma_q^2(0)$ by subtracting $S_{\mathbf{d}}(0)$ from $S_{\mathbf{d}}(t)$ on each trial. We initially encouraged sparsity in the weight vector by adding a $\lambda \sum_{i=1}^{n} |w_i|$ term to the cost function (similar to L1 regularization), where λ was optimized through cross-validation. As results were robust to changes in λ, final weight vectors were fit without a regularization term. To determine to what degree the representation of this diffusion process was driven by $S_{\mathrm{T_{in}}}$ activity, we removed all $\mathrm{T_{in}}$ neurons from the population and then refit the weight vector $\mathbf{d}^{-}$. We refer to the projection of the data onto this NSS direction as $S_{\mathbf{d}^{-}}$. We refer to the projection of the population response on **d** as $S_{\mathbf{d}}$ and use $S_{\mathbf{d}}(t)$ to render single trial firing rates (in arbitrary units, owing to standardization).

Spatial selectivity of **d**. To create the heat map in Fig. 3f we used the mean of $S_{\mathbf{d}}(t)$ during the memory-delay period of the ODR task (300–600 ms after target onset). Equivalent maps for all sessions capturing the visual response to the target (100–300 ms after target onset), the delay period activity and activity just before the saccade (250 to 50 ms before the saccade) are shown in Fig. S2.

Choice decoder. For each experimental session, we trained a linear choice decoder on 50% of trials in sliding 50 ms windows between 100 and 500 ms after motion onset. Decoder accuracy was cross-validated using the activity of held-out trials at the same time point. We additionally computed a logistic regression of $S_{\mathrm{T_{in}}}$, $S_{\mathbf{d}}$ and $S_{\mathbf{d}^{-}}$

activity against choice at each time point and compared their cross-validated choice prediction accuracy against the performance of the choice decoder and the theoretical peak decoding performance determined through simulations (Fig. 4a). To evaluate whether the optimal choice decoder represented a one-dimensional signal, we tested whether decoding accuracy was stable using the same weights over time[16]. Decoder weights generated for one time point were therefore applied to the held-out neural data at all other time points to then recomputed the decoding accuracy. With such time-independence established (Fig. 4b, left), a decoder dimension was generated by applying the decoder weights fit to the activity at $t = 0.35$ s after motion onset to all trials and time points.

Principal Components Analysis. We performed PCA on the normalized within-coherence mean activity from 200 ms after motion onset to the earlier of two events: 500 ms after motion onset or median RT minus 100 ms. Visual inspection of the Principal Components (PC) indicated that this yielded typically one or a small number of PCs that represented an evidence-dependent rise in activity after motion-stimulus onset. To determine which PC was most likely to reflect decision-related activity on single trials, we computed the correlations between projections of the neural data onto the first ten PCs with $S_{\mathrm{T_{in}}}$. In seven out of the eight experimental sessions, the first PC yielded the greatest correlation coefficient, highlighting that not only the average but also trial-to-trial variations in $S_{\mathrm{T_{in}}}$ resembles the main signal embedded in LIP activity in this task. Neural data projected onto this PC was then evaluated against our criteria for a decision variable (see below).

Spectral clustering. We used spectral clustering to group neurons into four clusters, using the correlation matrix between neuron pairs to construct the similarity matrix. For each neuron and trial, we counted the number of spikes in a time window between -200 and 0 ms relative to motion onset. We then calculated the Pearson correlation coefficient (ρ_{ij}), using these spike counts, for every pair of neurons i and j. The similarity (affinity) matrix was defined as $\rho_{ij} - \min(\rho)$, where $\min(\rho)$ is the smallest element of the correlation matrix. *k*-means clustering with 4 components was used to cluster the eigenvalues of the (random-walk normalized) Laplacian matrix[31]. We again computed the correlations between the mean activity in each cluster with $S_{\mathrm{T_{in}}}$ to determine which cluster was the most likely to yield diffusion dynamics. The mean activity of the most-highly correlated cluster was used to generate a spectral clustering dimension (Cluster-1).

Demixed principal component analysis (dPCA). We use dPCA[17] to reduce the dimensionality of the data while attempting to separate the motion-strength dependent components. For each neuron, we calculated the mean spike-counts across trials within coherence (in 10 ms bins) in the epoch 200–500 ms from motion onset. These signals were z-scored within neuron before we derived 5 dPCs using Matlab code provided by Ref. 17 (http://github.com/machenslab/dPCA). We identified the component (dPC1) that rendered coherence-dependent mean signals (Fig. S8a).

Correlation between single-trial activity and behavior. We used measures of correlation between the value of single-trial signals and both the RT and choice on that trial, restricting to trials with reaction times between between 0.670–2 s. *Single-trial signal* refers to the projections of the neural population activity onto vectors of weights (*e.g.*, $S_{\mathbf{d}}(t)$). For $\mathrm{T_{in}}$ neurons and Cluster-1 (Fig. 4d) it is the average firing rate of the neural responses as a function of time. We computed the Pearson correlation, $\rho(t)$, between reaction time and residual neural activity (mean-subtracted within motion coherence in a sliding 80-ms time window) on trials that resulted in a $\mathrm{T_{in}}$ choice. Negative correlation indicate that greater neural activity at a given time point was predictive of earlier $\mathrm{T_{in}}$ choices. We also compute the partial correlation conditional on the residual at $t = 0.55$ s. The reduction in R^2 induced by the later sample is the %-mediation, which we limit to the range 0–100%.

Correlation with a $\mathrm{T_{in}}$ choice was established using the coefficient, β_1, in logistic regression (same restriction on RT):

$$P_{\mathrm{T_{in}}}(t) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 S(t))}}, \quad 0.2 \leq t \leq 0.5. \tag{4}$$

Analogous to partial correlation, we included the later time point and fit

$$P_{\mathrm{T_{in}}}(t) = \frac{1}{1 + e^{-\left(\beta_0 + \beta_1^* S(t) + \beta_2 S(0.55)\right)}} \tag{5}$$

and defined mediation as

$$M = 100 \times \begin{cases} 1 - \frac{\beta_1^*}{\beta_1} & \beta_1 > 0 \text{ and } \beta_1^* > 0 \\ 1 & \beta_1 > 0 \text{ and } \beta_1^* \leq 0 \\ 0 & \beta_1 \leq 0 \end{cases} \tag{6}$$

As represented here, the mediation of the earlier $S(t)$ is by a later sample of the same signal (self). We prefer the simplicity of this definition of mediation to others based on *pseudo-*R^2. For both types of mediation, we also test whether the earlier $S(t)$ is mediated by the $\mathrm{T_{in}}$ neurons by substituting $S_{\mathrm{T_{in}}}(0.55)$ for $S(0.55)$ in Eq. 6.

The box-whisker plots (Figs. 2 to 4) summarize the distribution of %-mediation statistics for $t = 0.35$ s across experimental sessions. To determine statistical significance, we compare the %-mediation of data pooled across sessions against a null distribution generated by permuting, within session and signed coherence, the trials used for the later time point ($N = 1000$).

Similarity to $S_{\mathrm{T_{in}}}$. To quantify the similarity of the stochastic (diffusion) component of $S_{\mathrm{T_{in}}}(t)$ to other single-trial signals, we compute the correlation coefficient between the pair of detrended signals produced on the same trial, from $t = 0.2$ s to $t = RT - 0.1$. Detrending removes the mean response of trials of the same signed motion coherence. The r-values from each trial are summarized by histograms in Figs. 3 and 4. We used a one-tailed Kolmogorov-Smirnov (KS) test to compare histograms.

Acknowledgments

We thank Shushruth for comments on the manuscript, Cornel Duhaney and Brian Madeira for their assistance in the planning and execution of surgeries, animal training and general support, and we thank Columbia University's ICM for the quality of care they provide for our animals, especially during the pandemic and lockdown. We would further like to thank Tanya Tabachnik and her team at the Zuckerman Institute Advanced Instrumentation Core and Tim Harris, Wei-lung Sun, Jennifer Colonell, and Bill Karsh at HHMI Janelia for their continued support with Neuropixels1.0-NHP45 probes development and testing. This research was supported by the Howard Hughes Medical Institute; an R01 grant from the NIH Brain Initiative (M.N.S., R01NS113113); a T32 and F31 grant from the National Eye Institute (G.M.S, T32 EY013933, F31 EY032791); the Grossman center; and the Brain and Behavior Research Foundation.

Author contributions

G.M.S. and M.N.S. designed the experiment. N.A.S., G.M.S. and E.M.T. implemented the experimental set-up. G.M.S. collected the data. N.A.S., G.M.S., A.Z., D.M.W and M.N.S. analyzed the data. All authors wrote the manuscript.

Additional information

Supplementary Information is available for this paper. Correspondence and requests for materials should be addressed to shadlen@columbia.edu.

References

[1] Shadlen, M. N. & Newsome, W. T. Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America* **93**, 628–33 (1996). URL http://www.ncbi.nlm.nih.gov/pubmed/8570606.

[2] Shadlen, M. N. & Kiani, R. Decision making as a window on cognition. *Neuron* **80**, 791–806 (2013). URL http://www.ncbi.nlm.nih.gov/pubmed/24183028.

[3] Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nature neuroscience* **2**, 176–85 (1999). URL http://www.ncbi.nlm.nih.gov/pubmed/10195203.

[4] Romo, R., Hernandez, A. & Zainos, A. Neuronal Correlates of a Perceptual Decision in Ventral Premotor Cortex. *Neuron* **41**, 165 – 173 (2004).

[5] Hernández, A., Zainos, A. & Romo, R. Temporal evolution of a decision-making process in medial premotor cortex. *Neuron* **33**, 959 – 972 (2002).

[6] Ding, L. & Gold, J. I. Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. *Cereb Cortex* **22**, 1052–67 (2012). URL http://www.ncbi.nlm.nih.gov/pubmed/21765183.

[7] Roitman, J. D. & Shadlen, M. N. Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of neuroscience* **22**, 9475–89 (2002). URL http://www.ncbi.nlm.nih.gov/pubmed/12417672.

[8] Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annual review of neuroscience* **30**, 535–74 (2007). URL http://www.ncbi.nlm.nih.gov/pubmed/17600525.

[9] Laming, D. *Information theory of choice reaction time* (Wiley, 1968).

[10] Luce, D. *Response Times: Their Role in Inferring Elementary Mental Organization* (Oxford University Press, 1986).

[11] W Link, S. *The wave theory of difference and similarity* (Routledge, 1992).

[12] Churchland, A. K. *et al.* Variance as a signature of neural computations during decision making. *Neuron* **69**, 818–31 (2011). URL https://www.ncbi.nlm.nih.gov/pubmed/21338889.

[13] Peixoto, D. *et al.* Decoding and perturbing decision states in real time. *Nature* **591**, 604–609 (2021).

[14] Gnadt, J. W. & Andersen, R. A. Memory related motor planning activity in posterior parietal cortex of macaque. *Experimental brain research. Experimentelle Hirnforschung. Experimentation cerebrale* **70**, 216–20 (1988). URL http://www.ncbi.nlm.nih.gov/pubmed/3402565.

[15] Ganguli, S. *et al.* One-dimensional dynamics of attention and decision making in LIP. *Neuron* **58**, 15 – 25 (2008). URL http://linkinghub.elsevier.com/retrieve/pii/S0896-6273(08)00168-2.

[16] King, J.-R. & Dehaene, S. Characterizing the dynamics of mental representations: the temporal generalization method. *Trends in cognitive sciences* **18**, 203–210 (2014).

[17] Kobak, D. *et al.* Demixed principal component analysis of neural population data. *eLife* **5**, 9424 (2016).

[18] Kiani, R. *et al.* Natural grouping of neural responses reveals spatially segregated clusters in prearcuate cortex. *Neuron* **85**, 1359 – 1373 (2015).

[19] Stine, G. M., Zylberberg, A., Trautmann, E. & Shadlen, M. N. A neural mechanism for terminating decisions (forthcoming).

[20] Zohary, E., Shadlen, M. N. & Newsome, W. T. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature* **370**, 140–143 (1994).

[21] Shadlen, M. N. & Newsome, W. T. The Variable Discharge of Cortical Neurons: Implications for Connectivity, Computation, and Information Coding. *Journal of Neuroscience* **18**, 3870–3896 (1998).

[22] Okazawa, G., Hatch, C. E., Mancoo, A., Machens, C. K. & Kiani, R. Representational geometry of perceptual decisions in the monkey parietal cortex. *Cell* **184**, 3748–3761 (2021).

[23] Kiani, R., Hanks, T. D. & Shadlen, M. N. Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of neuroscience* **28**, 3017–29 (2008). URL http://www.ncbi.nlm.nih.gov/pubmed/18354005.

[24] Hikosaka, O. & Wurtz, R. H. Visual and oculomotor functions of monkey substantia nigra pars reticulata. III. Memory-contingent visual and saccade responses. *Journal of Neurophysiology* **49**, 1268–1284 (1983).

[25] van Den Berg, R. *et al.* A common mechanism underlies changes of mind about decisions and confidence. *Elife* **5**, e12192 (2016).

[26] Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E. & Shadlen, M. N. Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *J Neurosci* **31**, 6339–6352 (2011).

[27] Lewis, J. W. & Van Essen, D. C. Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *J Comp Neurol* **428**, 112–37 (2000). URL https://www.ncbi.nlm.nih.gov/pubmed/11058227.

[28] Pachitariu, M., Steinmetz, N., Kadir, S., Carandini, M. & D., H. K. Kilosort: realtime spike-sorting for extracellular electrophysiology with hundreds of channels. *bioRxiv* 061481 (2016).

[29] Pachitariu, M. Kilosort 2.0. https://zenodo.org/record/4147288 (2021).

[30] Hays, A. V., Jr, Richmond, B. J. & Optican, L. M. Unix-based multiple-process system, for real-time data acquisition and control. *WESCON Conf. Proc.* (1982). URL https://www.osti.gov/biblio/5213621.

[31] Shi, J. & Malik, J. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence* **22**, 888–905 (2000).

Extended Data

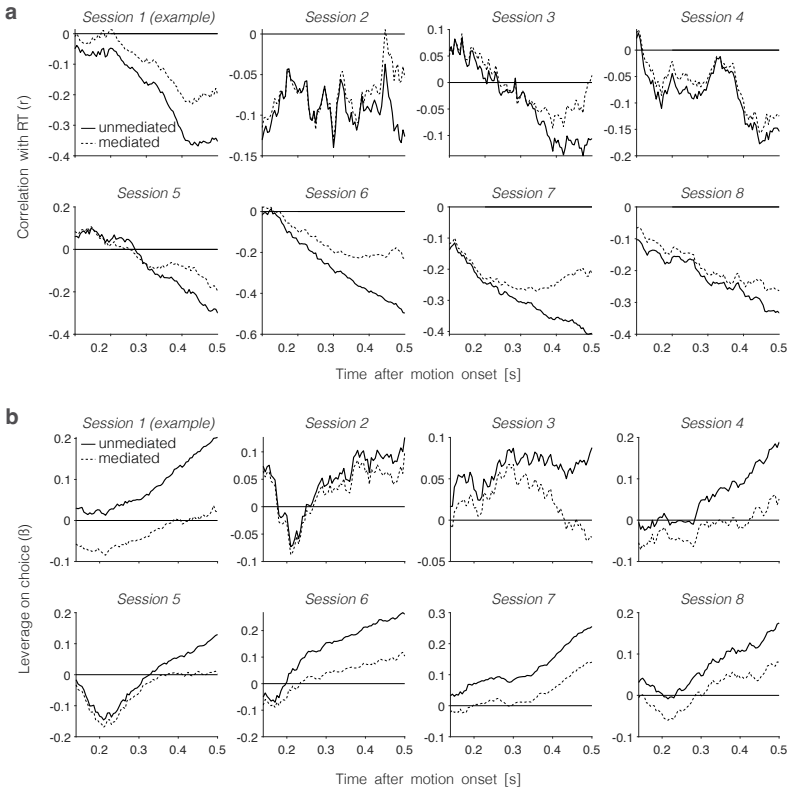

Figure S1: **Mediation analysis in Fig. 2d and e, for all sessions.**

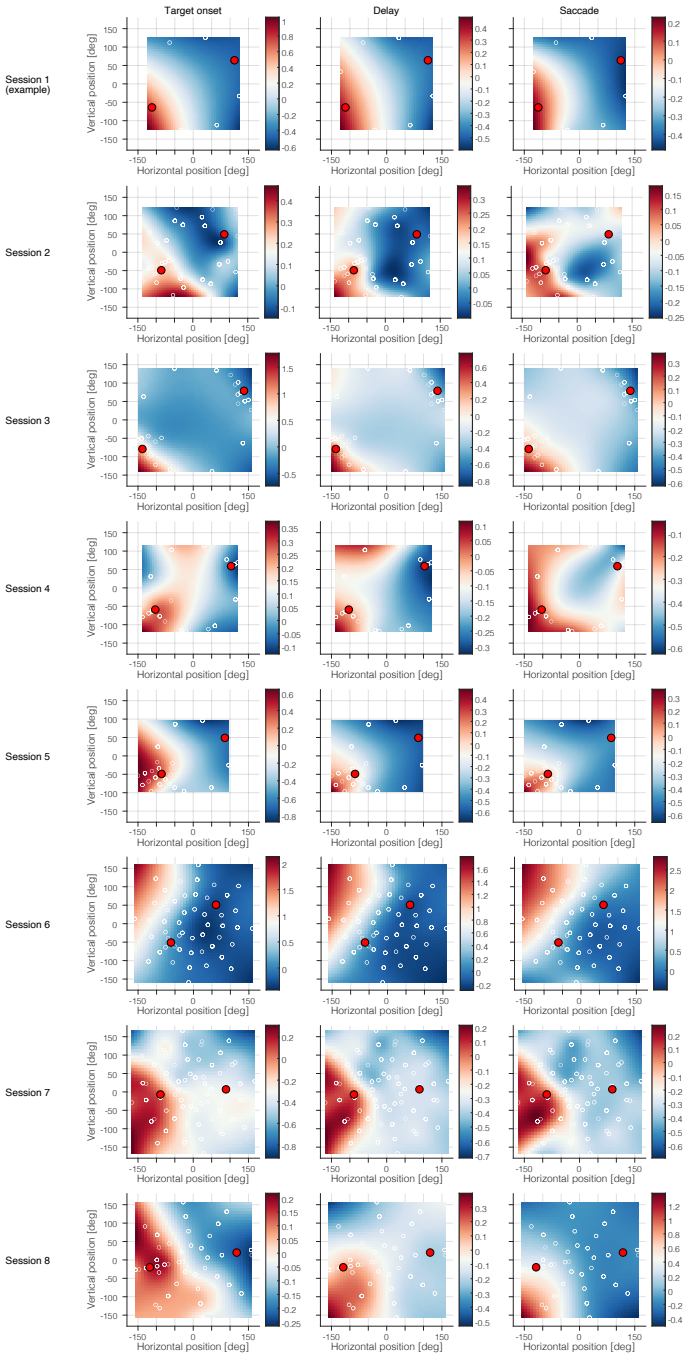

Figure S2: **Projection of neural activity during memory saccades task onto vector d.** Activity in the diffusion direction **d** during three time windows of the memory saccade task is plotted as a heat map as a function of target location. From left to right, each row contains maps of the visual response to the target (100–300 ms after target onset), the delay period activity during which the monkeys are remembering the location of the target no longer shown on the screen (300–600 ms after target onset) and activity just before the saccade (250 to 50 ms before the saccade). Filled red circles indicate the locations or response targets during the RDM task and open white circles the locations of all other target locations probed during the mapping task.

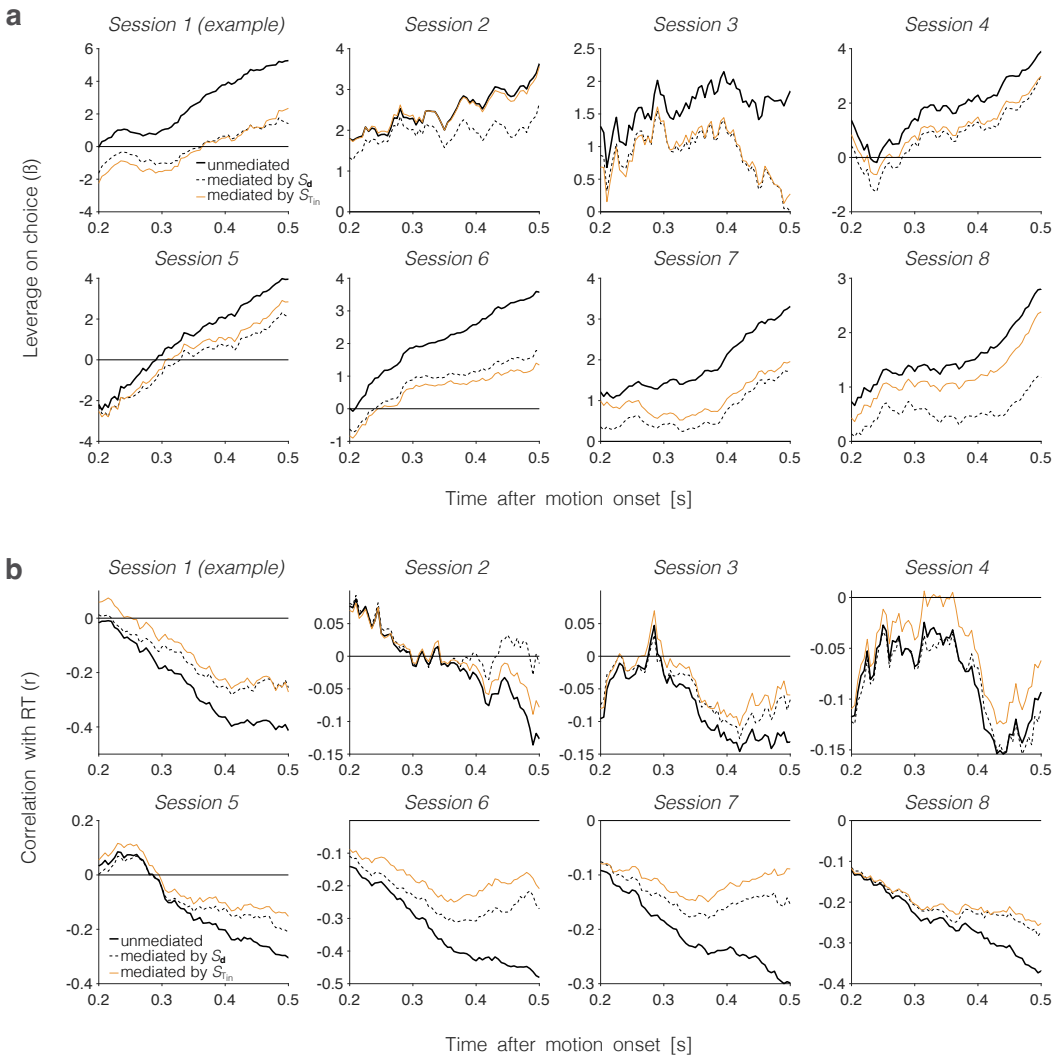

Figure S3: **Mediation analysis in Fig. 3g and h, for all sessions.**

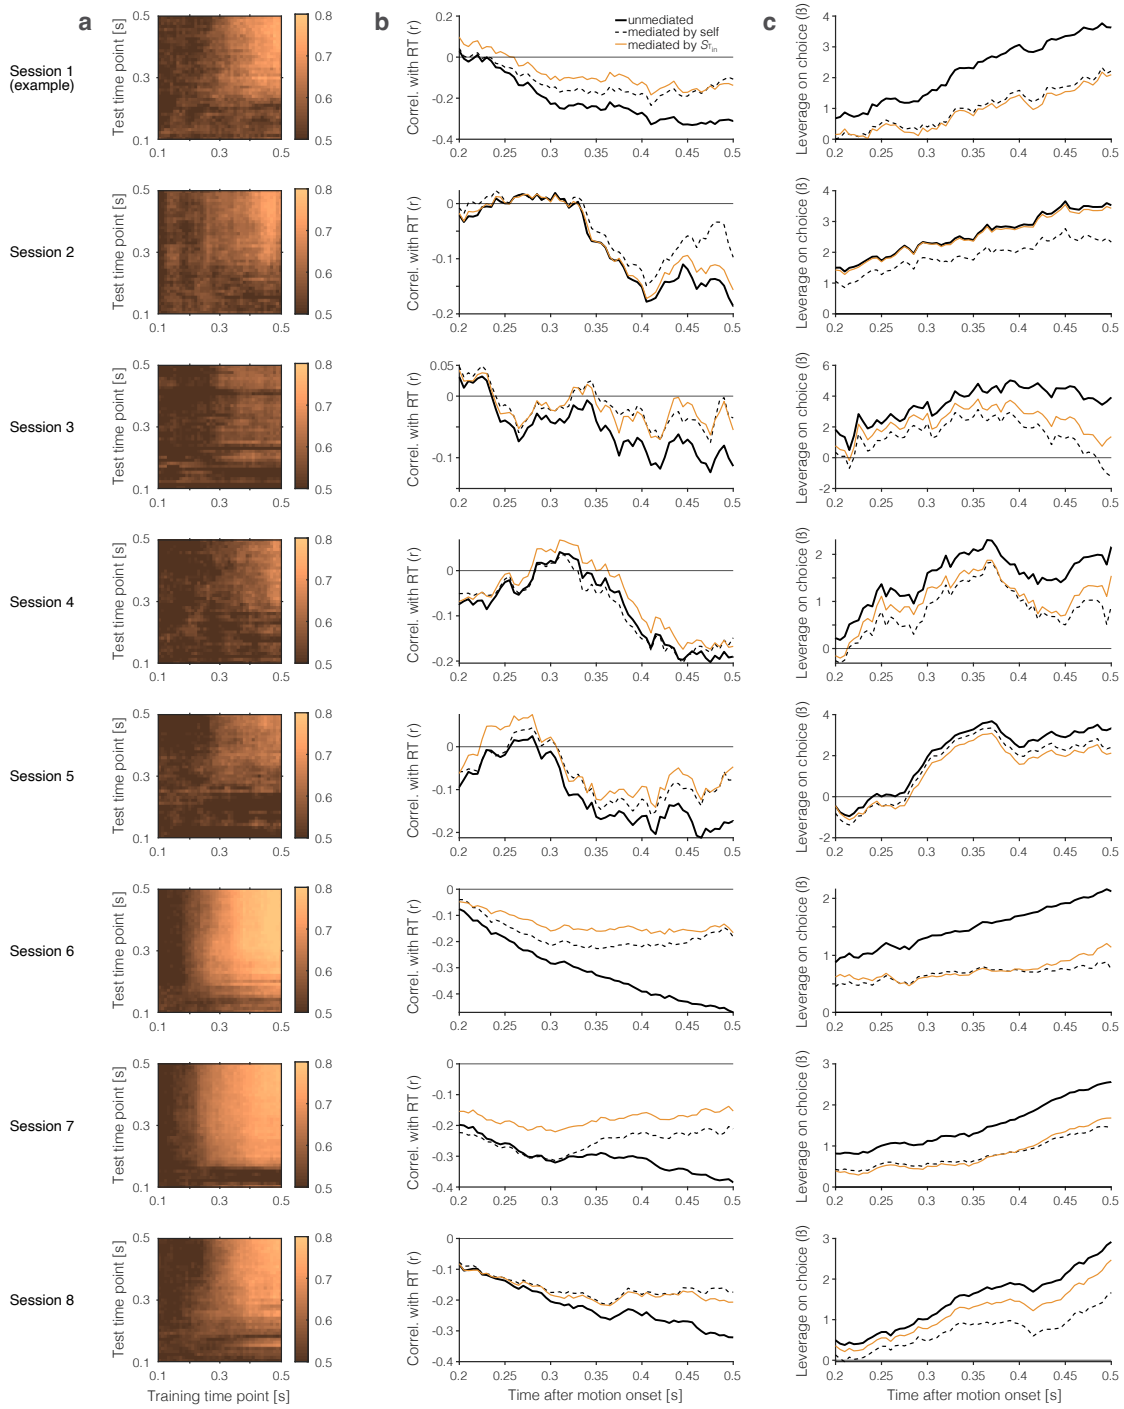

Figure S4: **Stabilty of choice-decoder and mediation analyses for all sessions. a,** Same as in Fig. 4a. The weights at $t = 0.35$ s define a direction in neuronal state space which is used to construct a one-dimensional signal, $S(t)$. *b,* Correlation and partial correlation $S(t)$ with RT for RT>0.67 s, as in Fig. 3g. **c,** Leverage of $S(t)$ on choice, as in Fig. 3h. Partial mediation by later sample of $S(t)$ is apparent in seven sessions; partial mediation by $S_{T_{in}}$ is apparent in six ($p < 0.001$).

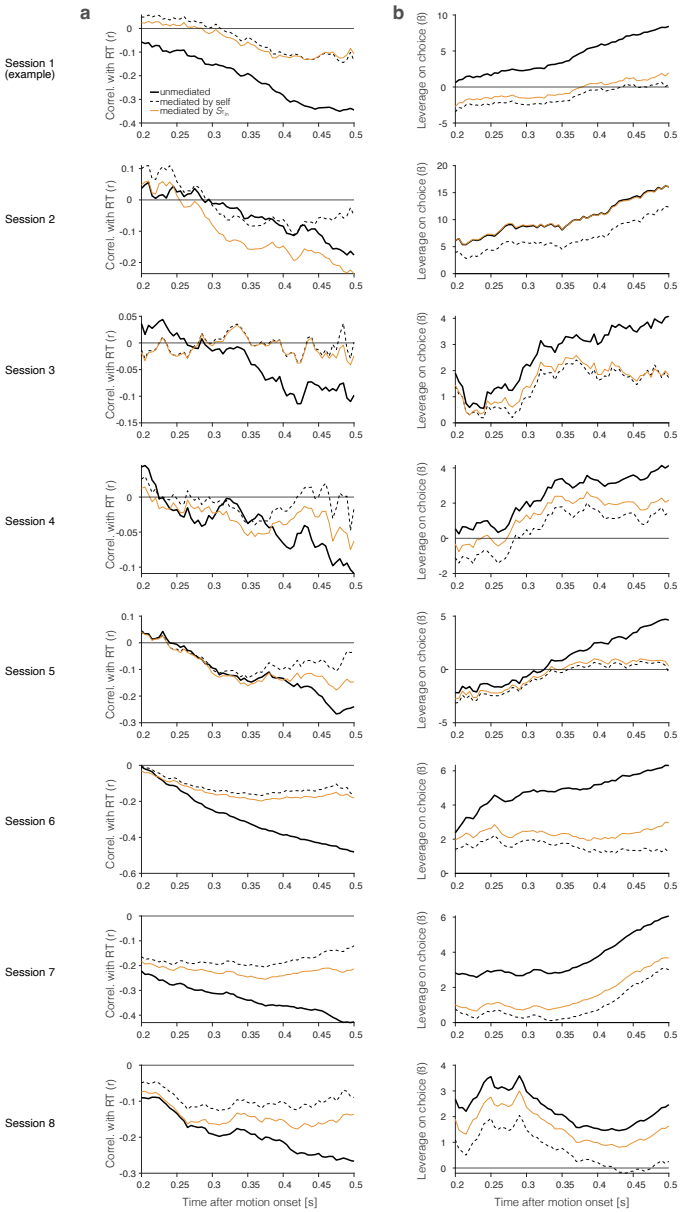

Figure S5: **Mediation analyses using the first Principle Component (Fig. 4c), as in Fig. S4b and c**.

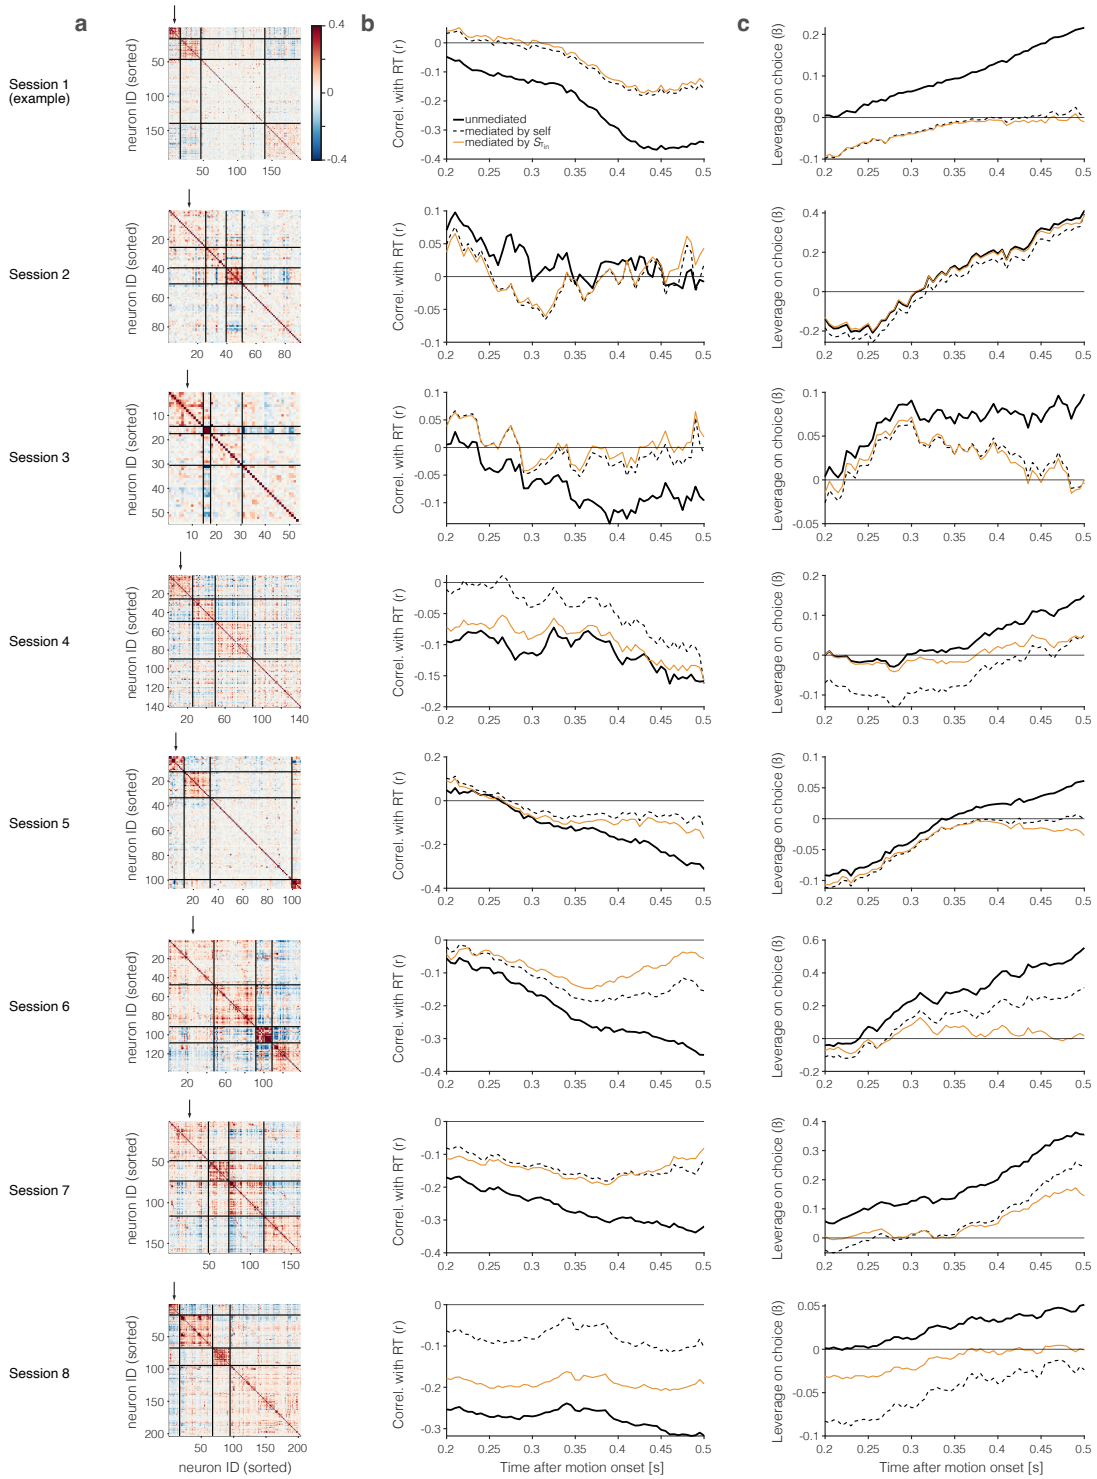

Figure S6: **Mediation analyses using the mean firing rate of Cluster-1 (Fig. 4d), as in Fig. S4.**

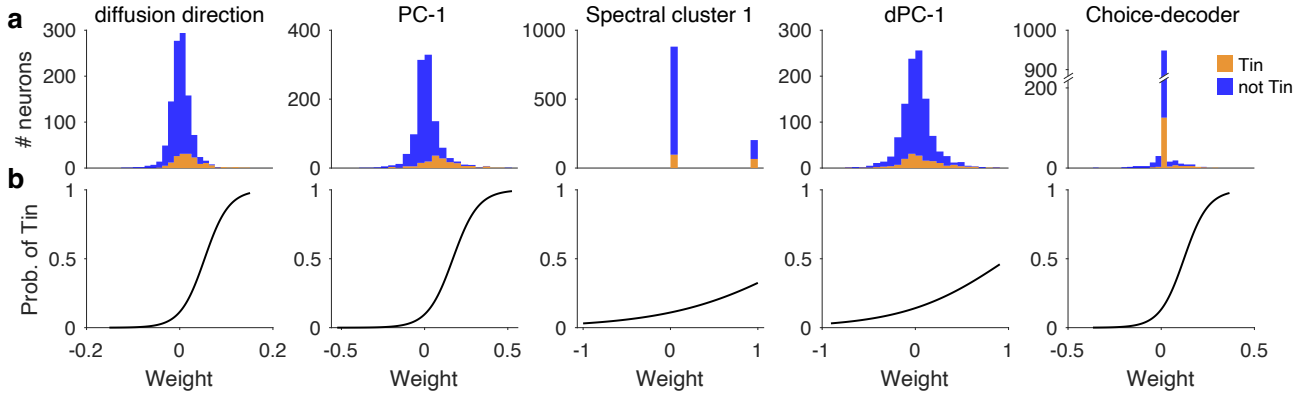

Figure S7: **Characterization of the weights assigned to T_{in} and non-T_{in} neurons.** **a,** Distribution of weights for T_{in} (copper) and non-T_{in} (blue) neurons, combined across all sessions. Columns address the same analysis method: diffusion as in (Fig. 3), first Principal Component (Fig. 4c), first spectral clustering component (Fig. 4d), first demixed dPC (Fig. S8), and weights derived from the logistic choice-decoder (Fig. 4a,b). The weights usually follow a bell-shaped distribution, with the exception of the spectral clustering analysis for which the weights are binary. **b,** Probability that a neuron assigned a weight (abscissa) is a T_{in} neuron. The curve is logistic regression.

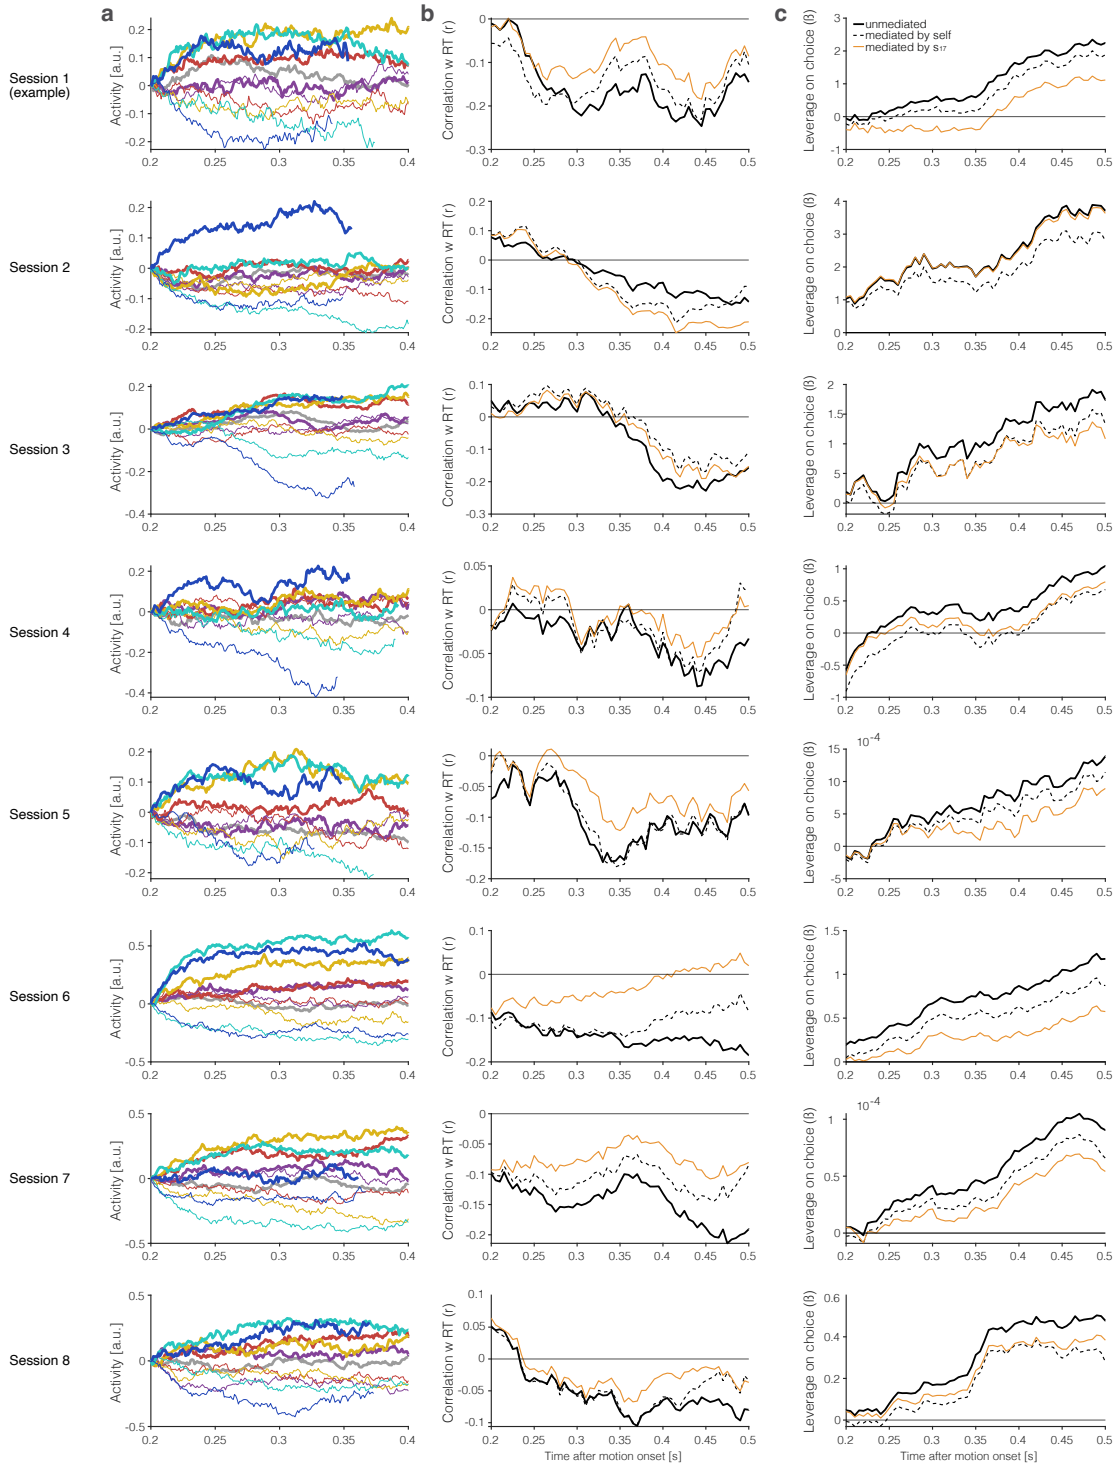

Figure S8: **Mean activity and mediation analyses using the first demixed Principle Component**. **a,** Mean activity projected onto dPC1 aligned to motion stimulus onset sorted by coherence and motion direction. **b-c,** Mediation analyses using the mean activity projected onto dPC1, as in Fig. S4.

Session	$S_{\mathbf{d}}$	$S_{\mathbf{d}^-}$	Decoder	PC1	Cluster-1
1	0.705	0.270	0.522	0.833	0.908
2	0.364	0.087	0.235	0.244	0.344
3	0.710	0.046	0.741	0.571	0.825
4	0.233	0.116	0.225	0.337	0.591
5	0.648	0.314	0.348	0.755	0.899
6	0.637	0.534	0.701	0.804	0.538
7	0.606	0.318	0.690	0.818	0.665
8	0.547	0.290	0.456	0.580	0.658

Table S1: Median within-trial correlation between $S_{\mathrm{T_{in}}}$ and five measures of decision-related activity.

Session	n_N	n_{C1}	% C1	$n_{\mathrm{T_{in}}}$	% $\mathrm{T_{in}}$ in C1	% C1 in $\mathrm{T_{in}}$
1	191	16	8	17	56	53
2	90	25	28	13	20	38
3	54	14	26	9	43	67
4	140	25	18	21	28	33
5	107	12	11	19	67	42
6	138	47	34	23	21	43
7	161	48	30	34	27	38
8	203	16	8	27	50	30
mean	136	25	20	20	39	43
std	47	14	10	7	16	11

Table S2: Intersection between $\mathrm{T_{in}}$ population and Spectral Cluster-1. n_N = number of neurons; n_{C1} = number of cells in cluster; % C1 = percentage of cells in cluster; $n_{\mathrm{T_{in}}}$ = number of $\mathrm{T_{in}}$ cells; %$\mathrm{T_{in}}$ in C1 = of the $\mathrm{T_{in}}$ cells, what percent are in Cluster-1. % C1 in $\mathrm{T_{in}}$ = of the Cluster-1 cells, what percent are $\mathrm{T_{in}}$ neurons;