

1 Unified neural dynamics of decisions and actions in the cerebral 2 cortex and basal ganglia

3 David Thura[†], Jean-François Cabana[‡], Albert Feghaly^{*}, and Paul Cisek

4 *Groupe de recherche sur le système nerveux central*

5 *Department of Neuroscience*

6 *Université de Montréal*

7 *Montréal, QC, CANADA*

8

9 Abstract

10 Several theoretical models suggest that deciding about actions and executing them are not completely
11 distinct neural mechanisms but instead two modes of an integrated dynamical system. Here, we
12 investigate this proposal by examining how neural activity unfolds during a dynamic decision-making
13 task within the high-dimensional space defined by the activity of cells in monkey dorsal premotor
14 (PMd), primary motor (M1), and dorsolateral prefrontal cortex (dlPFC) as well as the external and
15 internal segments of the globus pallidus (GPe, GPi). Dimensionality reduction shows that the four
16 strongest components of neural activity are functionally interpretable, reflecting a state transition
17 between deliberation and commitment, the transformation of sensory evidence into a choice, and the
18 baseline and slope of the rising urgency to decide. Analysis of the contribution of each population to
19 these components shows differences between regions but no distinct clusters within each region.
20 During deliberation, cortical activity unfolds on a two-dimensional “decision manifold” defined by
21 sensory evidence and urgency, and falls off this manifold at the moment of commitment into a choice-
22 dependent trajectory leading to movement initiation. The structure of the manifold varies between
23 regions: In PMd it is curved, in M1 it is nearly perfectly flat, and in dlPFC it is almost entirely
24 confined to the sensory evidence dimension. In contrast, pallidal activity during deliberation is
25 primarily defined by urgency. We suggest that these findings reveal the distinct dynamics of different
26 regions, supporting a unified recurrent attractor model of action selection and execution.

27

28 Present affiliations:

29 †: CRNL – ImpAct team, Inserm U1028 – University of Lyon 1, 69675 Bron (France) – david.thura@inserm.fr

30 ‡: Integrated Regional Center of Cancerology (CRIC), Lévis, QC (Canada) - jfcabana@outlook.com

31 *: Institute for Research in Immunology and Cancer of the University of Montréal, Montréal, QC (Canada) -
32 albert.feghaly@umontreal.ca

33 Corresponding author:

34 Paul Cisek, paul.cisek@umontreal.ca

35

36 Introduction

37 During natural behavior, we are continuously interacting with a complex and dynamic world^{1,2}. That
38 world often does not wait for us to make up our minds about perceptual judgments or optimal choices,
39 and inaction can lead to lost opportunities, or worse. Furthermore, we must often make decisions
40 while we're already engaged in an action, such as while navigating through our environment or
41 playing a sport³. These considerations suggest that the neural mechanisms involved in selecting and
42 executing actions should be closely integrated within a unified sensorimotor control system⁴. Indeed,
43 many neural studies have shown considerable overlap between the brain regions involved in action
44 selection and sensorimotor control⁵⁻¹⁰.

45 However, while the anatomical overlap between the distributed circuits of decision-making and
46 sensorimotor control is well-established, theoretical models of these processes remain largely
47 separate. Decision-making is often modeled as the accumulation of evidence until a threshold is
48 reached¹¹⁻¹⁹, at which time a target is chosen. Models of movement control usually begin with that
49 chosen target, toward which the system is guided through feedback and feedforward mechanisms²⁰⁻²².
50 But if the neural circuits involved in action selection and sensorimotor control are truly as unified as
51 neural data suggests, then theories of these processes should be similarly unified. One promising
52 avenue toward an integrated account of selection and control is to consider both as aspects of a single
53 distributed dynamical system, which transitions from a biased competition between actions²³⁻²⁸ into
54 an “attractor” that specifies the initial conditions for implementing the chosen action through feedback
55 control²⁹⁻³¹.

56 Here, we test whether neural activity in key cortical and subcortical regions exhibits properties that
57 would be expected from a unified dynamical system for action selection and sensorimotor control. We
58 focus on cells recorded in monkey dorsal premotor (PMd) and primary motor cortex (M1), which are
59 implicated in both selection and control^{7,9,32-36}, as well as the dorsolateral prefrontal cortex (dlPFC),
60 which is implicated in representing chosen actions³⁷. In addition, we examine activity in the output
61 nuclei of the basal ganglia, the globus pallidus externus (GPe) and internus (GPi), whose role in
62 selection and/or motor control is under vigorous debate³⁸⁻⁴¹. Importantly, we examine the activity of
63 all of these regions recorded in the same animals performing the same reach selection task, making it
64 possible to quantitatively compare activities in different brain areas using the same metrics.

65 To disentangle neural activity related to deliberation, commitment, and movement, we trained
66 monkeys to perform the “tokens task” (Figure 1) (see Methods). In the task, the subject must guess
67 which of two targets will receive the majority of tokens jumping randomly from a central circle every
68 200ms (Figure 1a). The subject does not have to wait until all tokens have jumped, but can take an
69 early guess, and after a target is reached the remaining tokens jump more quickly (every 150ms or
70 50ms in separate “Slow” and “Fast” blocks of trials). Thus, subjects are faced with a speed-accuracy
71 trade-off (SAT) – to either wait to be confident about making the correct choice, or to take an early
72 guess and save some time, potentially increasing their overall reward rate. If we assume that
73 commitment occurs shortly before movement onset, then we can delimit within each trial a period of
74 deliberation (Figure 1b) during which neural activity should correlate with the sensory evidence
75 related to token jumps as well as to subjective policies related to the speed-accuracy trade-off.
76 Furthermore, because we can precisely quantify the success probability (SP) associated with each
77 choice after every token jump, we can compute for each trial a temporal profile of the sensory
78 evidence and categorize trials into similarity classes (Figure 1c), including “easy trials”, “ambiguous
79 trials”, and “misleading trials” (see Methods for details).




Figure 1. The “tokens task”. **A.** During each trial, 15 tokens jump, one every 200ms, from the central circle to one of two outer target circles. The subject’s task is to move the cursor (black cross) to the target that will ultimately receive the majority of the tokens. **B.** Temporal profile of the “success probability” that a given target is correct. Once a target is reached, the remaining token jumps accelerate to one every 150ms (“Slow” block) or 50ms (“Fast” block). We subtract from movement onset the mean reaction time (RT), measured in a separate delayed-response task, to estimate commitment time (purple bar) and the success probability at commitment time (dotted red horizontal line). **C.** Success probability for choosing the right target in trial types defined on the basis of the success probability profile, here computed after aligning to movement onset (see Methods). Solid: correct target on the right; Dashed: correct target on the left.

80 Previous studies have shown that both human and monkey behavior in the tokens task is well-
 81 explained by the “urgency-gating model” (UGM)^{14,42}, which suggests that during deliberation, the
 82 sensory “evidence” about each choice (provided by the token distribution) is continuously updated
 83 and combined with a non-specific “urgency signal”, which grows over time in a block-dependent
 84 manner, and commitment to a given choice is made when the product of these reaches a threshold.
 85 Recent neural recordings largely supported these proposals (Extended Data Figure 2). For example,
 86 some cells in PMd (29%) and M1 (45%) were significantly tuned during deliberation, reflecting both
 87 the evidence favoring their preferred target and the growing urgency signal^{36,43}. In contrast, cells in
 88 GPe and GPi did not show tuning until after commitment, and about 50% exhibited either increasing
 89 or decreasing activity that resembled the context-dependent urgency signal⁴⁴.

90 While these and other studies reveal important properties of activity in these regions, many questions
 91 remain. Are “decision-related” neurons part of a module for choosing a target, which sends its output
 92 to a separate module of “movement-related” neurons? Do the basal ganglia contribute to the
 93 deliberation process^{39,45–49} or do they simply reflect a choice taken in cortical regions^{50,51} and
 94 contribute only to movement execution⁵²? Answering these questions is difficult given the
 95 heterogeneity of cell properties^{7,32,53–56} and their apparently continuous distribution along rostrocaudal
 96 gradients^{54,57} or layers⁵⁸. This leads one to consider whether, instead of serial modules, action
 97 selection and execution are two modes of a unified recurrent system distributed across the
 98 frontoparietal cerebral cortex and associated basal ganglia/thalamic loops. According to this model⁵⁹,
 99 action selection occurs through a competition within the regions of PMd and M1 associated with the
 100 relevant effector, biased by signals arriving from dlPFC related to evidence in favor of specific
 101 targets. As time passes, that competition is invigorated by an urgency signal coming from the basal
 102 ganglia, which gradually amplifies the competitive dynamics in PMd/M1. As the contrast develops
 103 between the activity of cortical cell groups associated with different candidate actions, selectivity is
 104 gradually induced in the striatum and the pallidum, leading to a positive feedback that further favors
 105 the winning cells and suppresses the others, constituting volitional commitment and launching the
 106 dynamics of execution^{29–31}.

107 Here, we test this proposal by examining activity across all of the regions we recorded in the tokens
 108 task (PMd, M1, GPe, GPi, and dlPFC), but without *a priori* classifying cells into putative functional
 109 categories. Instead, we use the neural space approach pioneered in recent years^{60–70}, in which the
 110 entire system is described as a point in a very high-dimensional space defined by the activity of all

111 recorded cells, and then reduced into a lower-dimensional representation that reveals the main factors
112 governing cell activity across the system. We then perform specific analyses to characterize how
113 activity unfolds during deliberation and commitment, comparing the dynamics of different regions,
114 and quantify to what extent cell properties cluster into distinct functionally interpretable roles. Some
115 of these results have previously appeared in abstract form⁷¹⁻⁷⁴.

116

117 Results

118 Neural data

119 We recorded spiking activity from a total of 736 well-isolated individual neurons in the cerebral
120 cortex and basal ganglia of two monkeys (S and Z), recorded at the locations shown in Extended Data
121 Figure 1. Of these, 356 were recorded in PMd (237 from monkey S), 211 in M1 (79 from monkey S),
122 62 in dlPFC (60 from monkey S), 51 in GPe (19 from monkey S), and 56 in GPi (22 from monkey S).
123 The properties of some of these neurons have been reported in previous publications, focusing on
124 tuned activity in PMd and M1 and the basal ganglia^{36,43,44,75}, as summarized in Extended Data Figure
125 2. Here, we additionally include neurons recorded in dlPFC, which for technical reasons were almost
126 exclusively obtained in monkey S and to date only described in abstract form^{71,72}. While a more
127 complete description of their properties awaits more recordings in additional animals, they are
128 included in the analyses described below, in part to provide a contrast to the other cortical and
129 subcortical areas.

130 Dimensionality reduction of neural activity into principal components

131 To examine how the activity of this entire population evolves over time during the task, we performed
132 Principal Components Analysis (PCA) on all cells recorded in all five brain regions, including any
133 well isolated neuron that was recorded in both Slow and Fast blocks. This included a total of 637
134 neurons, including 277 in PMd, 191 in M1, 52 in dlPFC, 41 in GPe and 46 in GPi. For the PCA, we
135 used data from only four conditions (right and left choices in the fast and slow blocks) and counted
136 each neuron once. This means that the variance explained was dominated by the PMd and M1
137 populations, in which we recorded the largest number of neurons. However, reducing the number of
138 cells to 35 in each region did not change the results apart from making them noisier and changing the
139 percentage of variance explained by individual principal components (PCs) (See Extended Data
140 Figure 6a). An alternative approach would be to perform PCA on each population of neurons
141 separately, but this would yield region-specific PCs that make quantitative comparisons between
142 regions impossible. Thus, we elected to perform PCA on all cells together, counting each neuron
143 once, producing a “loading matrix” of coefficients (from each neuron to each PC) that then allows us
144 to “project” each population into the space of the same PCs. As shown below, this allows us to
145 directly compare the components of different regions and to infer how each region contributes to the
146 same distributed dynamical system. We imposed symmetry on our population by following the “anti-
147 neuron” assumption, which assumes that for every neuron we recorded there exists a similar neuron
148 with the opposite relationship to target direction (even for cells that are not tuned), effectively
149 doubling the number of neurons. See the Methods section for the justification and motivation for this
150 approach.

151 The first 20 PCs together explained 97.9% of the variance in activity over time across the four groups
152 of trials (slow/fast blocks x left/right choices) used for the Principal Components Analysis. Figure 2
153 shows the temporal profile for the first 7 PCs constructed as a weighted average of all cells, separately
154 for easy, ambiguous, and misleading trials in both slow and fast blocks, for both left and right choices.
155 Here, the PCs are calculated after aligning the data to movement onset. See Extended Data Figure 3
156 for the same PCs calculated after aligning to the beginning of the trial.

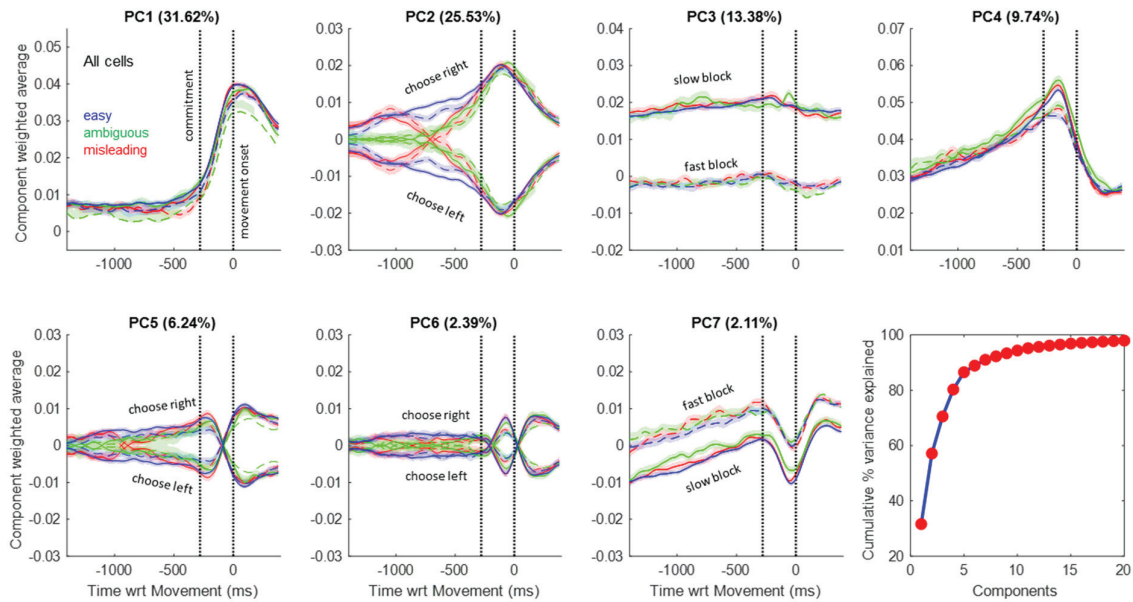


Figure 2. Components produced by PCA. The cumulative variance explained by the first 20 components is shown at the bottom right, and the temporal profiles of the top 7 PCs are shown in the rest of the figure. Each of those 7 panels shows the average activity of all cells, weighted by their loading coefficient onto the given PC, for 12 trial groups, including combinations of 3 trial types: easy (blue), ambiguous (green), and misleading (red); two blocks: Slow (solid) and Fast (dashed); for choices made to the left or right (indicated for the components where they differ). Note that the sign is arbitrary because the loading matrix can have positive or negative values. Shaded regions indicate 95% confidence intervals. In each panel, the second vertical dotted line indicates movement onset and the first indicates the estimated moment of commitment, 280ms earlier, based on our prior studies. Each panel is scaled to have the same range in the y-axis. PCs for the Slow trials are built from 557 cells that have all trial types in Slow blocks, while PCs for Fast trials are built from 452 cells that have all trial types in Fast blocks.

157 The first 4 components, which together explain 80.3% of total variance, are clearly interpretable in
 158 terms of the key elements of the urgency gating model. The first PC (31.62% of variance) is nearly
 159 identical across all conditions and reflects the transition between deliberation (prior to commitment)
 160 and action (after movement onset). It is similar to the main condition-independent component
 161 reported in other neural space studies in both primates^{76–80} and rodents^{81–83}. The second PC (25.53%)
 162 exhibits two phases. Prior to commitment, it reflects the time course of the sensory evidence on which
 163 the monkey made his choice, distinguishing easy, ambiguous, and misleading trials. After
 164 commitment, it simply reflects the choice made without distinguishing trial types. The third PC
 165 (13.38%) reflects the block-dependent aspect of urgency, distinguishing between the slow and fast
 166 blocks (even before the start of the trial, as shown in Extended Data Figure 3). The fourth PC (9.74%)
 167 reflects the time-dependent aspect of urgency until just before movement onset. The remaining
 168 components are similar to PCs 2 and 4 during deliberation, but capture some of the heterogeneity
 169 across cell activity patterns after commitment and movement onset. We discuss these higher PCs in
 170 the Supplemental Materials.

171 PC2 warrants special attention. Note that until the moment of commitment, it correlates very well
 172 with the evidence provided by the token movements (compare to Figure 1c). Indeed, as shown in
 173 Extended Data Figure 4, the correlation between the time-delayed evidence and the value of PC2 is
 174 highly significant ($p < 10^{-100}$) with a correlation coefficient of $R = 0.92$. This is notable because the PCA
 175 algorithm was not given any information about these different trial types (easy, ambiguous,
 176 misleading, etc.) but was simply given data averaged across four large trial groups that only
 177 distinguished left versus right choices and slow versus fast blocks. Nevertheless, when the resulting

178 temporal profiles of PC2 are calculated for specific trial types they clearly reflect how the evidence
179 dynamically changes over the course of deliberation in those trials.

180 Could this finding be a trivial consequence of overall cell tuning? To test this possibility, we used the
181 Tensor Maximum Entropy method of Elsayed & Cunningham⁸⁴ to generate synthetic data sets that
182 retain primary features such as tuning, but are otherwise random (see Methods). As shown in
183 Extended Data Figure 4c, when PCA is applied to such synthetic data sets it does not produce
184 components that are as well correlated with evidence as PC2 from our true data ($p < 0.01$). The
185 implication is that the emergence of PC2 in the real neural data requires a consistent relationship
186 between how cells reflect the final choice (left vs. right) and how they reflect the evidence that leads
187 to that choice during deliberation (easy vs. ambiguous vs. misleading, etc.).

188 Figure 3 shows the trajectories of the different trial types, separately for the Slow and Fast blocks,
189 plotted in the space of PCs 1, 2, and 4. While a quantitative comparison between the Slow and Fast
190 blocks is made difficult because a slightly different subset of cells is included in each (see Methods),
191 the qualitative shape of the trajectories is very similar. As indicated in Figure 3a (dotted black
192 arrows), in both block types the neural activity evolves in a clockwise manner in the space of PC1 and

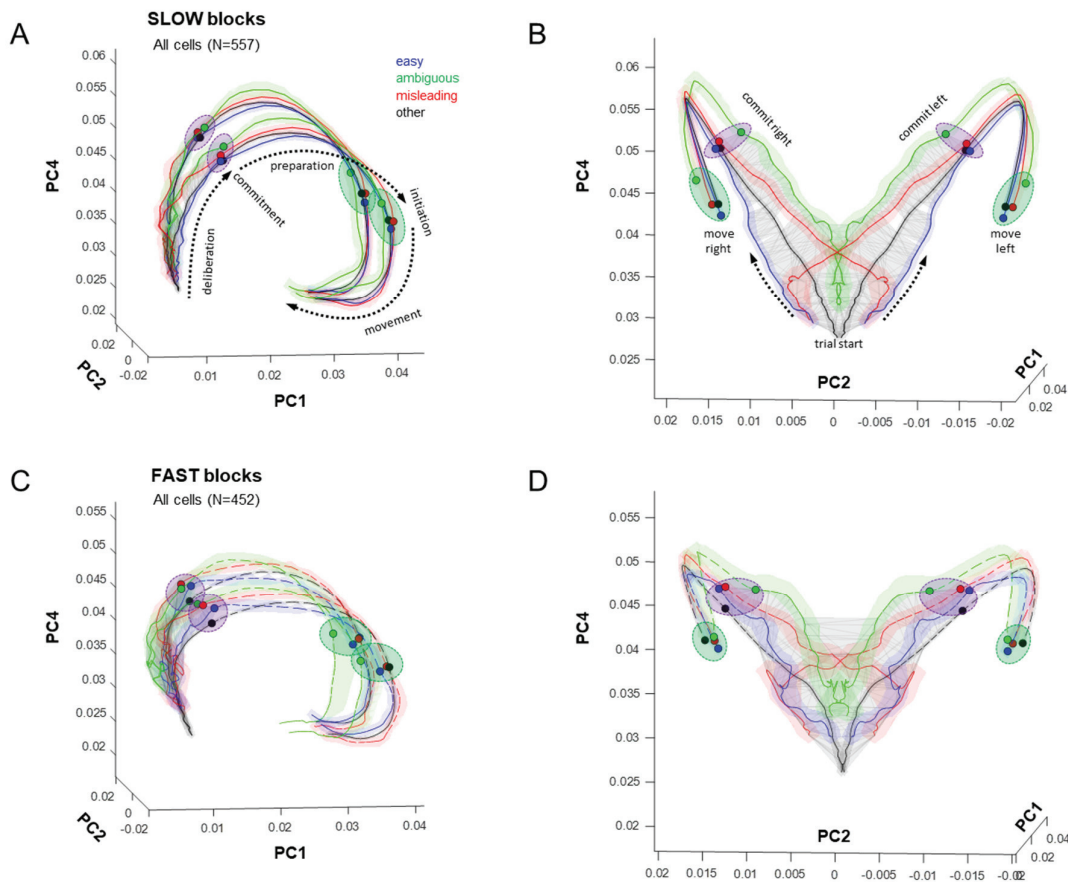


Figure 3. Neural trajectories in the space of PCs 1, 2 and 4, averaged across easy, ambiguous, and misleading trials as well as the other (unclassified) trials. Panels A and B show two views of data from Slow blocks, and panels C and D show the same views of data from the Fast blocks. Shaded colored regions around each trajectory indicate the 95% confidence interval. Dotted arrows in A indicate how the state of activity evolves over time. The gray wireframe encloses all states visited during the deliberation epoch, across all trial types. Purple ellipses indicate the region in which commitment occurs (indicated for individual trial types by small colored circles) and green ellipses indicate the point at which movement is initiated. For clarity, the neural state after movement initiation is only shown in panels A and C.

193 PC4, passing over a region of deliberation until reaching a commitment state (purple ellipses),
194 whereupon it rapidly moves to a movement-specific initiation subspace (green ellipses), and then
195 turns back toward the starting point during movement execution. Some of these phenomena have
196 previously been reported using neural space analyses of preparatory and movement-related activity in
197 cortical regions during instructed reaching tasks with a single target^{76–79,85–87}. Here, our task allows us
198 to examine in more detail what happens during the process of prolonged deliberation when subjects
199 are selecting among multiple targets.

200 In each panel of Figure 3, we’ve drawn a gray wireframe around all of the points from the beginning
201 of the trial until commitment time (280ms before movement onset), across all trial types, thus defining
202 the subspace within which deliberation occurs. This subspace resembles a triangular surface that is
203 extended mostly in PC2 and PC4 and curved slightly into PC1. It is quite thin – for example, in the
204 slow block the value of Ψ (see Methods) is 0.204, which is roughly equivalent to a triangular sheet
205 whose thickness is $1/57^{\text{th}}$ of the length of each side. We call this the “decision manifold”.

206 The flow of neural states upon the decision manifold is quite orderly, proceeding from bottom to top
207 as time elapses and shifting left and right with the sensory evidence. For example, consider the
208 misleading trials, in red, which clearly reveal the switch in sensory evidence. In effect, the flow of the
209 neural state during deliberation resembles the temporal profile of evidence (Figure 1c) mapped onto
210 that curved wireframe surface. The neural state continues to flow along the decision manifold until it
211 reaches one of two edges (purple ellipses) at the time of commitment, and then turns into PC1 and
212 accelerates to rapidly flow along one of two paths, each corresponding to the choice taken, until
213 movement initiation (green ellipses).

214 Region-specific dynamics

215 While the structure of the neural space computed across all neurons is interesting, it is still more
216 informative to compare that structure across the different brain regions in which we recorded. Because
217 the loading matrix produced by PCA provides coefficients that map each individual neuron’s
218 contribution to each PC, we can “project” the activity of any subset of neurons into the space of these
219 same PCs (see Extended Data Figure 5), and then plot region-specific neural space trajectories. Figure
220 4a shows this for the dorsal premotor cortex, where we see a structure that is quite similar to what was
221 shown for all neurons (not surprisingly since the PMd population is the largest). As before, we see a
222 triangular decision manifold that is relatively thin ($\Psi=0.343$) and extends along PC2 and PC4.
223 However, note that it initially strongly leans in the negative PC1 direction and then curves around just
224 before commitment (see the side view shown in Figure 4a, right). Interestingly, the PMd decision
225 manifold is curved as if it lies on the surface of a sphere (see inset, spherical fit $R^2 = 0.65$), a point to
226 which we will return below.

227 In contrast, the decision manifold of primary motor cortex (M1, Figure 4b) is remarkably flat and thin
228 ($\Psi=0.252$) and leans into the positive PC1 direction. Nevertheless, the evolution of the neural state
229 along the surface of the decision manifold in both regions obeys the same pattern seen in Figure 3,
230 proceeding from bottom to top as time elapses and shifting left and right with the sensory evidence,
231 always lying within the same subspace (curved for PMd, planar for M1). Note that Figure 4 shows
232 data aligned on movement onset, as in previous figures, but also superimposes neural space
233 trajectories computed on the basis of data aligned on the start of the trial (see Extended Data Figure 3)
234 and then projected into the same PC space using the same loading matrix. As is clear from the figures,
235 regardless of how the trajectories are computed, they always fall within the same decision manifold in
236 PC space, for both PMd and M1.

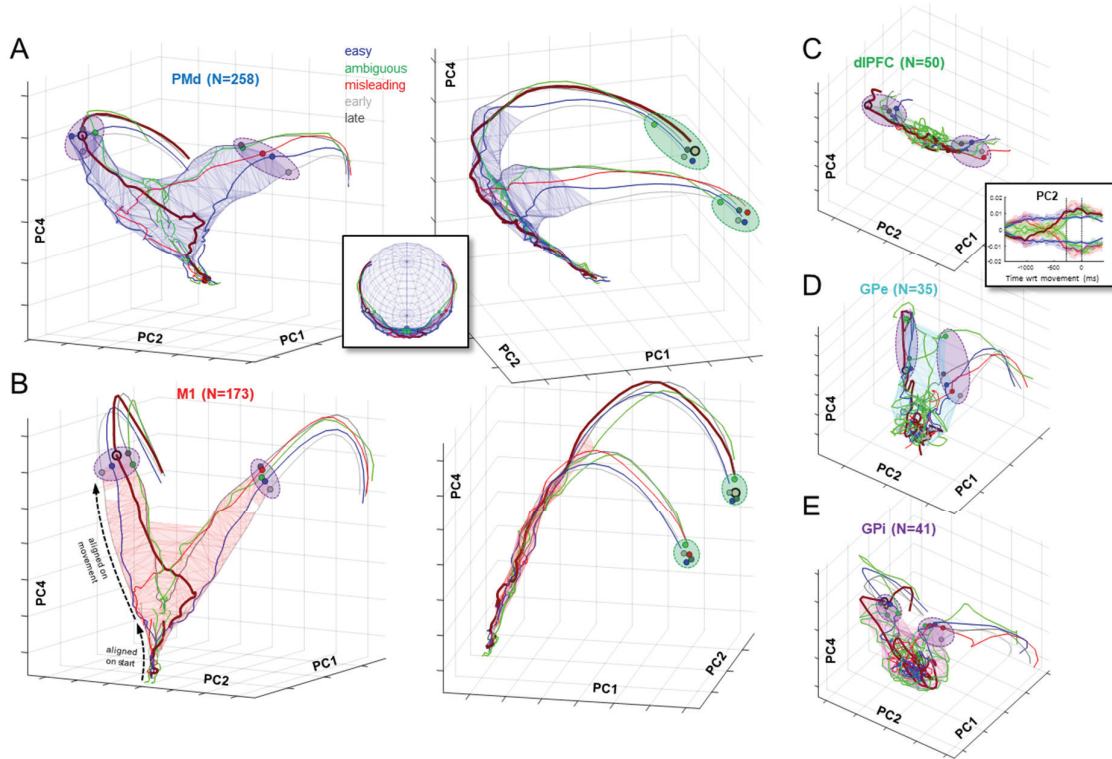


Figure 4. Neural trajectories in the space of PCs 1, 2 and 4, plotted using activity during Slow blocks, separately for PMd (A), M1 (B), dlPFC (C), GPe (D) and GPi (E). Separate trajectories are plotted for easy, ambiguous, and misleading trials, as well as all trials in which decisions were shorter than 1400ms (early) and in which decisions were longer than 1400ms (late). For clarity, confidence intervals have been omitted. As in Figure 3, trajectories are computed from data aligned on movement onset, and extend from 1400ms prior. Blue (PMd) and red (M1) wireframes enclose all states before commitment (280ms before movement). Here, we also superimpose trajectories computed from data aligned on the start of the trial and until 500ms later (projected through the same loading matrix). Dashed black arrows in panel B, left, indicate where these separate trajectories can be seen in the M1 space. In all panels, the trajectory of misleading trials in which the monkey correctly chose the right target is highlighted with a thicker line. Purple ellipses emphasize the time of commitment and green ellipses emphasize movement onset. The inset in A shows how PMd activity tends to remain on the surface of a sphere, particularly during deliberation. The inset in C shows PC2 computed from the dlPFC population for easy (blue), ambiguous (green), and misleading trials (red).

237 It is noteworthy that in both PMd and M1, the state reached at the moment of commitment (purple
 238 ellipses in Figure 4a,b) shows an orderly relationships with the reaction time in each trial type
 239 (shortest in “early” trials, and longest in “late” trials). This is in agreement with the observation that
 240 even at a single trial level, a consistent relationship between neural state and reaction time can be
 241 observed during a simple instructed reach task⁷⁶. Furthermore, in M1 the trajectories after that point
 242 converge to arrive in a relatively compact subspace (green ellipses) at movement onset – what
 243 Churchland et al.⁸⁸ called an “optimal subspace”.

244 In contrast to PMd and M1, the deliberation manifold in dorsolateral prefrontal cortex (dlPFC, Figure
 245 4c) is almost exclusively extended along PC2 ($\Psi=0.564$, like a cylinder whose length is 22 times its
 246 radius). Like PMd and M1, the neural state along PC2 shifts left and right with sensory evidence (see
 247 inset), but after commitment it exhibits only a small excursion into PC1. This suggests that neural
 248 activity in dlPFC primarily reflects the sensory evidence used to make decisions in the task, consistent
 249 with many previous studies^{89–96}.

250 A strong contrast to the cortical data is seen when we examine the neural state of cells recorded in the
251 globus pallidus (GPe and GPi). In both regions, the neural state during deliberation is confined within
252 a subspace that is not a thin manifold but instead resembles a ball compressed along PC2 and
253 extended along PC4 (Figure 4d,e). During deliberation, activity in these regions does not evolve in an
254 orderly fashion as seen in cortex. This could be partly due to the lower number of cells recorded in
255 GPe and GPi as compared to PMd and M1, although similar results hold when we restrict *all* regions
256 to 35 cells (Extended Data Figure 6a). Nevertheless, by the time commitment occurs, the state of both
257 GPe and GPi lies in a choice-specific subspace (purple ellipses) and then evolves quickly to a
258 corresponding initiation subspace. These findings are consistent with our previous report of GPe/GPi
259 activity in the tokens task, in which we suggested that these regions do not determine the choice but
260 rather contribute to the process of commitment⁴⁴.

261 Interpreting the shape of decision manifolds

262 As shown in Figure 4, the decision manifolds computed from PMd and M1 possess distinct shapes.
263 While M1 is almost completely flat, the PMd manifold is curved, as if lying on a surface of a sphere.
264 Do these shapes reveal differences in the neural dynamics in these regions? To address this question,
265 we consider the shape of the PMd decision manifold in terms of two separate phenomena: its
266 curvature in the PC1-PC4 plane (Figure 4a, right), and its curvature in the PC1-PC2 plane (inset in
267 Figure 4a).

268 First, we consider why the PMd manifold initially leans in the negative PC1 direction and then bends
269 toward the positive PC1 direction prior to commitment (its curvature in the PC1-PC4 plane). One
270 potential explanation suggests that an inhibitory influence prevents premature movement before
271 selection is complete⁹⁷ by keeping PMd away from commitment, but it is gradually overcome by
272 positive feedback in the recurrent circuit between PMd and GPi^{44,59,98}. According to this hypothesis, as
273 the cortical activity becomes increasingly biased in favor of one target over another, it gradually
274 begins to produce the emergence of choice selectivity in the GPe, about 200ms before commitment
275 (see Extended Data Figure 2b, row 3). When that becomes strong enough to engage selectivity in the
276 GPi, it in turn strengthens emerging selectivity in the thalamus, which then further strengthens
277 selectivity in the cortex. Thus, a positive feedback is established leading to a winner-take-all process
278 that overcomes inhibition and constitutes volitional commitment. This hypothesis predicts a
279 relationship between how tuning emerges in GPi and how and when the PMd state begins to flow
280 toward commitment.

281 To test this prediction, we examined the correlation between the flow toward commitment in PMd (as
282 reflected in the rate of change of PC1 during deliberation) with the depth of directional selectivity in
283 GPi (as reflected in the absolute value of PC2). Figure 5a shows this comparison for 6 large trial
284 groups: all rightward choices in Slow and Fast blocks, early rightward choices in Slow and Fast
285 blocks, and late rightward choices in Slow and Fast blocks (see Methods for trial definitions).
286 Leftward decision trials were not considered because they are symmetric with rightward decision
287 trials and thus redundant. As can be seen for the different trial types, except for a constant scaling
288 factor the match between these two very different variables is strong, particularly in the epoch
289 between commitment and movement onset. Even if the data is restricted to the period *before*
290 commitment, the correlation for each trial group shown in the individual panels of Figure 5a is
291 significant at $p < 10^{-10}$ with Pearson's R above 0.86, and for all trials together with $R = 0.84$ (Figure 5b).
292 Furthermore, across these six conditions there was a significant correlation ($p = 0.0029$, $R = 0.9556$)
293 between the time when the PMd manifold began to tilt toward commitment (derivative of PC1
294 became consistently positive) and the time when tuning became significant in GPi (95% CI of PC2 no
295 longer included 0). Similar results were not obtained when comparing GPi PC2 against the derivative
296 of PC1 computed from other regions (although there was a weaker but significant correlation with the
297 derivative of PC1 in M1). It is important to note that this particular prediction – a relationship
298 between the derivative of one component in PMd and the absolute value of another in GPi – is not

299 arbitrary. It is motivated by the specific proposal that the timing of the bend in the PMd decision
 300 manifold in the PC1-PC4 plane (Figure 4a, right) is related to the timing of how tuning emerges in the
 301 basal ganglia, which is itself motivated by the hypothesis that both of these variables reflect positive
 302 feedback in a recurrent attractor circuit. Of course, like any correlation analysis, it cannot conclusively
 303 prove a specific causal relationship. An alternative, though not mutually exclusive hypothesis is that
 304 both of these regions are influenced by a common source of inhibition that is gradually released as
 305 commitment approaches.

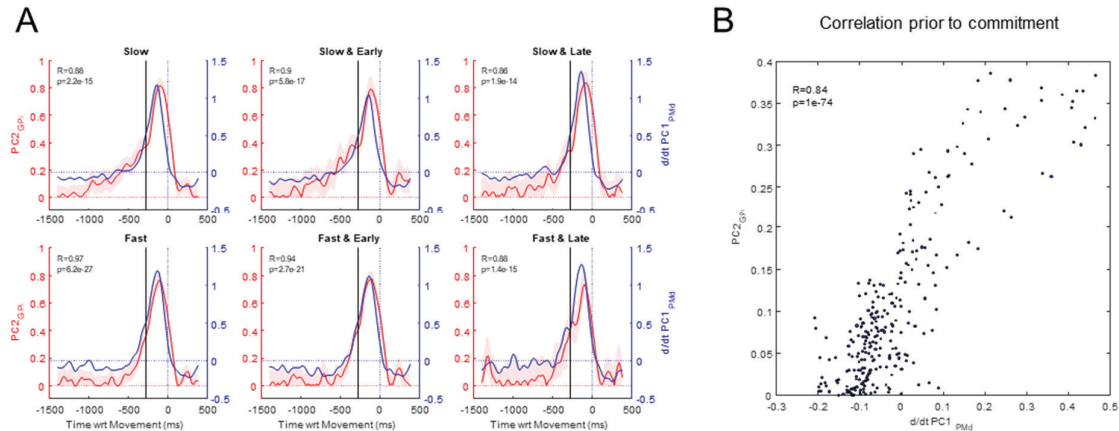


Figure 5. Comparison of selectivity in GPi with the slope of the PMd decision manifold. **A.** Comparison within six conditions: all Slow block trials, Slow block trials where $DT < 1400$, Slow block trials where $DT > 1400$, all Fast block trials, Fast block trials where $DT < 950$, Fast block trials where $DT > 950$. Each panel plots the absolute value of PC2 computed from GPi (red, left y-axis, shaded region indicates 95% CI) along with the derivative of PC1 computed from PMd (blue, right y-axis). The black vertical line indicates the estimated time of commitment. At the top left corner are Pearson's R and the p -value of the correlation between these two signals from 1000ms to 280s before movement. **B.** The correlation of these signals, from 1000 to 280ms before movement onset, across all 6 conditions.

306 In addition to its curvature in the PC1-PC4 plane, the PMd manifold is also curved in the PC1-PC2
 307 plane, as if it lies upon a surface of a sphere (see inset in Figure 4a). This is strikingly different than
 308 the M1 manifold, which is nearly perfectly flat (compare Figure 4a with b). What could explain this
 309 difference in shapes?

310 Here, we consider one possible explanation related to the dynamics of recurrent attractor networks.
 311 We illustrate this using a very simple system consisting of two neurons that compete against each
 312 other through recurrent inhibition. Note that this minimal model is not intended to simulate our data,
 313 but simply to demonstrate the possibility that some features of our data (e.g. shape of the decision
 314 manifold) could be the result of very general properties of recurrent non-linear dynamical systems.

315 Let's consider two neurons whose activity is denoted as x_1 and x_2 (Figure 6a) and governed by the
 316 following differential equation²³

$$317 \quad \frac{dx_i}{dt} = (1 - x_i)(E_i + U + f(x_i)) - x_i(A + f(x_j)) \quad (1)$$

318 On the right-hand side of equation (1), the first term is excitation and the second is inhibition. E_i is the
 319 input evidence for choice i , U is the urgency signal, A is a passive decay rate, $f(x_i)$ is the recurrent
 320 excitation of each cell to itself, and $f(x_j)$ is the recurrent inhibition from the other cell, j . Note that
 321 equation (1) ensures that cell activities are always in the interval from 0 to 1. The function $f(x)$ is
 322 sigmoidal of the form

$$323 \quad f(x) = \frac{G}{1 + e^{-S(x-0.7)}} \quad (2)$$

324 where G is the gain and S is the steepness of the sigmoid. Figure 6 schematizes this simple model and
 325 illustrates some of its dynamics.

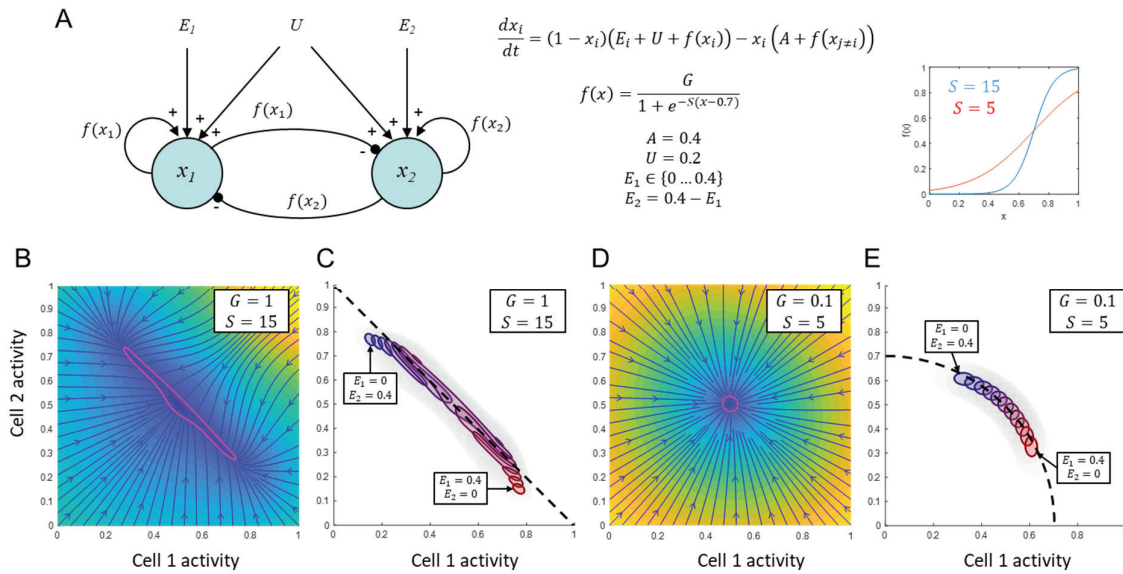


Figure 6. Simulation of a simple 2-neuron recurrent competitive attractor. **A**. The structure of the model is shown at left, next to the governing equations and parameter settings. Two forms of the interaction function $f(x)$ are shown in the inset at right. **B**. The flow field for a system with high gain ($G=1$) and steep slope ($S=15$) when the input evidence is balanced such that $E_1=E_2=0.2$. Blue arrows depict the flow and shading indicates the speed (yellow=fastest, dark blue=slowest). The purple outline indicates a region around the stable equilibrium, in which the state will tend to remain even with substantial noise. Note that the dynamics flow quickly toward that region and then move more slowly within it. **C**. The stable regions of the same system for several input patterns, ranging from strong evidence in favor of choice 2 (top left, blue) to strong evidence in favor of choice 1 (bottom right, red). Note that these align in what is approximately a straight line across the state space (compare to dashed black line). **D**. Same as **B** but for a system in which the interaction function has low gain ($G=0.1$) and a shallow slope ($S=5$). **E**. Stable regions of the system in **D**. Note that now, the stable regions for the same range of input patterns do not form a straight line but rather fall on a curve roughly equivalent to a circle (dashed black curve).

326 Figure 6b shows the flow field of the system in the (x_1, x_2) -plane assuming a function $f(x)$ with high
 327 gain and a steep slope, when the evidence is balanced. Note how the flow field quickly pushes the
 328 neural state to a central stable region, outlined in purple, within which the flow is slower. In the
 329 presence of noise, the system will be strongly constrained to remain inside this region, but can shift
 330 within it relatively easily. As the balance of evidence changes, the stable region shifts in the plane
 331 (Figure 6c) but always lies oriented along that straight line. As a result, activity is normalized such
 332 that x_1+x_2 is approximately constant (L1-normalization). In contrast, Figure 6d,e shows a system
 333 where the interaction function has low gain and a shallow slope. Note that the stable region for
 334 different input patterns still shifts in the (x_1, x_2) -plane, but not along a straight line (panel e). Instead,
 335 the different stable regions now lie along a circular path, such that $x_1^2 + x_2^2$ is approximately constant
 336 (L2-normalization). In both systems, the state is strongly confined into a narrow subregion of the full
 337 space (a “decision manifold”), but can shift within it due to changes of evidence as well as noise. If
 338 the value of U is increased, this subspace will shift toward the upper right (not shown) until the
 339 system bifurcates into two stable attractors, each corresponding to choosing either x_1 or x_2 .

340 This simple model proposes a straightforward candidate explanation for the difference in shapes of
 341 the decision manifolds in PMd versus M1 (Figure 4a,b). To summarize, relatively gentle competitive
 342 dynamics between candidate options can produce a curved quasi-spherical manifold, as seen in PMd,

343 while strong and steep winner-take-all dynamics can produce a flat one, as seen in M1. This may
344 imply that PMd tolerates multiple competing potential actions more readily than does M1, a
345 conjecture made in previous modeling work²⁵. However, at present these are just conjectures. A more
346 complete analysis would require a more sophisticated model, in which different neural populations are
347 dynamically coupled and interact in more complex ways. It is possible that in such a model, other
348 parameter settings and other features of dynamics may better explain the shapes of decision
349 manifolds, but exploring those possibilities is beyond the scope of the present paper.

350 Analyses of the loading matrix

351 Another approach for inferring the putative functional contributions of different brain regions is to
352 examine the distribution of the loading coefficients for cells in each region. For example, if a given
353 population of cells is strongly related to the sensory evidence provided by token jumps, then cells in
354 that population should tend to have higher loading onto PC2 than cells from another population that is
355 less sensitive to evidence. As described in Methods, we characterized the distribution of loading
356 coefficients for the first 11 PCs (which capture 95% of variance) by fitting the 11-dimensional space
357 of points for each brain region with gaussian mixture models (GMMs). The results are shown in
358 Figure 7.

359 The distribution of loading coefficients for cells in PMd (Figure 7a) was highly distributed but not
360 without structure, and was best fit with two gaussians. The first (1: 68% contribution to the fit) was
361 dominated by cells only weakly contributing to PCs 1, 2, and 4, but strongly to PC3. The second (2:
362 32%) included cells that strongly contributed to PCs 1 and 2 but more weakly to PCs 3 and 4. The
363 distribution for M1 was also best fit with two gaussians, one (1: 56%) contributing mostly to PC3 and
364 another (2: 44%) contributing to 1, 2, and 4 but not 3. Thus, in both PMd and M1, there was a trend
365 for cells that most strongly reflect the animal's speed-accuracy trade-off (PC3) to be less strongly
366 tuned to direction (PC2), and vice-versa (Figure 7a,c, middle panels).

367 In contrast, populations in dIPFC, GPe, and GPi were well fit with a single gaussian (Figure 7e,f),
368 though perhaps more structure would have been seen with larger populations. Importantly, there were
369 significant and potentially functionally relevant differences between these regions. In particular, the
370 distribution of loadings from dIPFC and GPi were nearly orthogonal in the space of PCs 1, 2, and 4
371 (Figure 7f). The dIPFC was extended along PC2 and not along the other components, consistent with
372 the proposal that it primarily carries information on the sensory evidence provided by token
373 movements. In contrast, GPi was relatively narrow in PC2 and instead extended along PCs 3 and 4,
374 related to the block and time-dependent aspects of urgency. Furthermore, there was a significant
375 negative correlation ($R=-0.32$, $p<0.001$) between GPi loadings on PC1 versus PC4 (Figure 7e, right
376 panel, purple). In other words, cells that build-up over time (positive on PC4) tend to reduce their
377 activity after commitment (negative on PC1), while cells that decrease over time (negative on PC4)
378 tend to increase after commitment (positive on PC1).

379 In summary, our analyses of loading matrices suggest that while different regions do appear to
380 contribute to different subsets of PCs (e.g. the orthogonal relationship of dIPFC and GPi in Figure 7f),
381 the population within each region does not contain distinct clusters. In other words, while there is
382 some structure in the loading matrix (e.g. Figure 7a,b, middle panel), the distributions of properties
383 within each region are continuous.

384 Could that continuity of properties be an artifact resulting from dimensionality reduction? That is, if
385 distinct categories of cells with different functional roles really did exist, would our analyses be able
386 to identify them? To address this question, we created a variety of synthetic populations of neuron-
387 like units and applied to them the same analyses we used to examine real data, including PCA and
388 GMM analyses of the resulting loading matrix. As described in the Supplemental Materials (see
389 Extended Data Figure 8), this yielded three conclusions: It confirmed that PCA does correctly identify

390 all of the components from which our synthetic neural populations were constructed. However, it also
 391 showed that some of the higher-order components (like PC5 in Figure 2) can result from PCA
 392 “cancelling out” some of the firing patterns already captured by lower-order components (e.g. PC2),
 393 in order to explain things like neurons tuned only during movement. Nevertheless, in all cases the
 394 GMM analysis of the loading matrix correctly identified any real categories of neurons in the
 395 population. This suggests that the lack of distinct clusters we found in our data indeed reflects the
 396 absence of separate categories in the real neural populations.

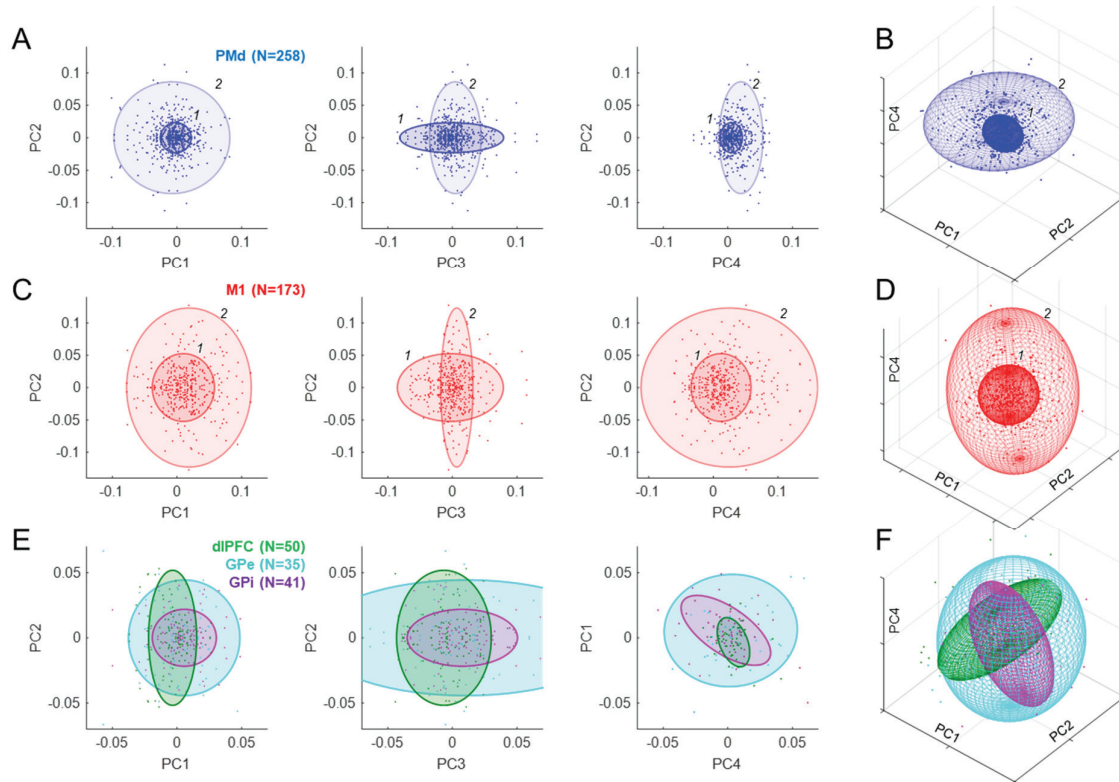


Figure 7. Analysis of loading coefficients. **A.** Each point indicates the weight of the contribution of a given PMd cell to the principal component indicated on the axes. Colored ellipses indicate the centroid and 3 times standard deviation of each of the two Gaussians (labeled 1 and 2) that provided the best fit to the distribution of these populations. **B.** The same PMd Gaussians in the 3-d space of PC1, PC2, and PC4. **C.** Same as A for M1. **D.** Same as B for M1. **E.** Same as A for cell populations in dlPFC (green), GPe (cyan), and GPi (purple), each of which was best fit with a single Gaussian. Note that unlike A and C, here the right panel shows PC4 vs PC1. **F.** Same as B for dlPFC, GPe, and GPi.

397 Discussion

398 Many neurophysiological studies have suggested that decisions between actions unfold within the
 399 same sensorimotor regions responsible for the control of those actions^{1,99–101}. This includes FEF and
 400 LIP for gaze choices^{6,91,102–105} and PMd and MIP for reaching choices^{7,36,105–107}. Several computational
 401 models of the decision mechanism suggest that it behaves like a “recurrent attractor system”^{23–28},
 402 where reciprocally competing groups of neurons tuned for the available choices compete against each
 403 other until one group wins and the system falls into an attractor corresponding to a specific choice.

404 The results reported here provide strong support for this class of models. In particular, the cells we
 405 recorded in PMd and M1 do not appear to belong to separate categories related to decision-making
 406 versus movement preparation or execution, but instead behave like part of a unified dynamical system
 407 that implements a biased competition and transitions to commit to a choice through a winner-take-all
 408 process. During deliberation, the pattern of cell activity in these regions is confined to a highly

409 constrained subspace in the shape of a thin manifold, and is shifted around within that manifold by the
410 decision variables pertinent to the task (here, the sensory evidence and the rising urgency). When
411 commitment occurs, the same group of cells now transitions from the decision manifold to a roughly
412 orthogonal tube-shaped subspace corresponding to a specific choice⁸¹ and quickly flows to a subspace
413 related to movement initiation^{77,88}. This is precisely the kind of “winner-take-all” phase transition that
414 occurs in recurrent attractor models.

415 Additional insights can be obtained by examining the low-dimensional components produced by
416 dimensionality reduction. In particular, it is noteworthy that during deliberation, the four strongest
417 components of neural activity (which together account for just over 80% of the variance) capture the
418 key elements of the urgency-gating model¹⁴: The momentary evidence (in PC2), the urgency signal (a
419 context-dependent baseline in PC3 and time-dependent ramping in PC4), and the transition to
420 commitment (PC1). Perhaps even more important is how these components are differently expressed
421 in the different cell populations. In particular, the neural state in dlPFC varies almost entirely along
422 PC2 while, conversely, the state in GPe/GPi during deliberation is primarily determined by PC3 and
423 PC4, and not at all by PC2. This suggests that information about sensory evidence is provided by
424 prefrontal cortex^{89–96} while the urgency signal is coming from the basal ganglia^{44,108,109}.

425 The presence of the evidence-related component PC2 in the cortical data is particularly remarkable
426 because the dimensionality reduction algorithm was not provided with any information about the
427 variety of trial types (easy, ambiguous, etc.) but was merely given data averaged across four very
428 large groups of trials: left or right choices during slow or fast blocks. Nevertheless, the difference of
429 activity related to left vs. right choices led the algorithm to assign a choice-related component that
430 also happens to capture the evolving evidence for that choice. A neural population control method⁸⁴
431 verified that this is not a simple consequence of tuning. Additional analyses (Extended Data Figure
432 6b) show that a unified evidence/choice component is obtained even if the PCA algorithm is only
433 given data restricted to activity *after* movement onset. This suggests consistency between cell
434 properties before versus after movement onset⁵⁴, and argues against a categorical distinction between
435 movement selection vs. execution circuits and in favor of a unified dynamical system, at least among
436 the cells we recorded.

437 The emergence of the other components is also highly robust. If the PCA algorithm is only given data
438 from the slow block, PC3 is lost (unsurprisingly), but the other components remain (Extended Data
439 Figure 6c). Also not surprisingly, if we provide PCA with data from all 28 of our trial classes, the
440 components are even more clearly distinguished, even though the number of cells that possess all of
441 the required trials is reduced by 37% (Extended Data Figure 6d). In fact, any reasonable subset of data
442 we have tried leads the PCA algorithm to identify the same functionally relevant components, albeit
443 sometimes in a slightly different order depending on how much variance is captured by each. Finally,
444 similar features are obtained if we leave out the cell duplication step (see Methods), although the
445 resulting decision manifolds become asymmetrically distorted and some of the structure of loading
446 matrices is more difficult to see (Extended Data Figure 7).

447 Finally, it is noteworthy that many of our findings reliably reproduce previous observations from very
448 different tasks, including ones in which monkeys did not have to make any decision and were simply
449 instructed to reach out to a single target. This includes the general flow of the neural state from target
450 presentation to movement onset and offset⁷⁷ (Figure 3a,c), the orderly relationship between reaction
451 times and the neural state in PMd/M1⁷⁶ (Figure 4a,b, purple ellipses), the compact subspace in M1 at
452 movement onset⁸⁸ (Figure 4a,b, green ellipses), and the presence of condition-independent
453 components related to state transitions and elapsing time⁷⁹ (Figure 2, PC1, PC4). Although most of
454 our data comes from a prolonged period of deliberation that is not present in those other studies, the
455 phenomena related to preparation and execution, which are shared between paradigms, are
456 nevertheless robustly reproduced. This further strengthens the proposal that action selection and
457 sensorimotor control are two modes of a single unified dynamical system.

458 Indeed, analyses of the loading matrix suggest that among the cells we recorded in PMd and M1, there
459 is no categorical distinction between those involved in selection and those responsible for movement.
460 Further analyses of synthetic populations (e.g. Extended Data Figure 8c) demonstrate that some of the
461 higher order PCs we observed (PC5 and PC6, see Figure 2) may result from the heterogeneity of
462 properties across our cell population, which includes purely decision-related and purely movement-
463 related cells. However, analysis of the loading matrix concluded that these properties are not clustered
464 into distinct categories as in synthetic data (Extended Data Figure 8d), but are instead distributed
465 along a continuum (Figure 7).

466 While some of these findings could have been anticipated from analyses of individual cells, one
467 important observation that would not have been possible concerns the shape of the decision subspace
468 in the cortical populations. In particular, it is highly consistent across all of the different trial types and
469 always resembles a thin manifold. This suggests strong normalization dynamics that are an inherent
470 feature of recurrent attractor models. That is, the state of neural activity is pushed to lie on a surface
471 that conserves some quantity (e.g. total neural activity with respect to some baseline) but is then free
472 to move upon that surface under the influence of evidence, urgency, or simply noise. Furthermore, the
473 particular shape of the manifold reveals consistent differences in the dynamics of different neural
474 populations. In particular, regardless of what data we provide to PCA, the M1 manifold is always
475 almost perfectly flat while the PMd manifold always exhibits a characteristic curvature. Interestingly,
476 a very similar curvature was observed in preliminary analyses of PMd data in a very different decision
477 task¹¹⁰. In contrast, the decision subspaces in the globus pallidus are much more compact and nothing
478 like the thin manifolds in cortex.

479 We believe that the shape of the decision manifolds computed from different regions reveals
480 properties of the underlying dynamics of neural activity in those regions. The planar manifold seen in
481 M1 would be expected from a dynamical system with a steep interaction function governing mutual
482 inhibition between competing groups of neurons (Figure 6c). In contrast, the curved PMd manifold
483 would be expected from a system with a shallower interaction function (Figure 6e). In addition, we
484 observed that the PMd manifold bends, approximately 200ms prior to commitment, along the
485 component associated with the transition from deciding to acting. We suggest that this bend is the
486 signature of a gradually emerging positive feedback in the cortico-striatal-thalamo-cortical circuit,
487 which gradually overcomes inhibitory signals preventing premature selection⁹⁷. Indeed, the analysis
488 shown in Figure 5 reveals a strong correlation between how directional selectivity begins to emerge in
489 GPi and how the PMd state begins to flow toward commitment, consistent with positive feedback
490 between these regions^{59,98,111}. Of course, establishing a causal relationship will require future studies,
491 including simultaneous microstimulation in one region and recording in the other.

492 In conclusion, our analyses support the hypothesis that decisions between actions emerge as a
493 competition within the sensorimotor system^{7,9,25,91,102,104,105}, which is governed by recurrent attractor
494 dynamics^{23,24,26–28}. That competition is biased by sensory information coming at least in part from the
495 prefrontal cortex⁹¹ and is gradually amplified by an urgency signal from the basal ganglia^{44,108,109}.
496 Commitment to a choice occurs through a positive feedback between premotor cortex and the basal
497 ganglia^{44,59,98,111}, leading to a winner-take-all process. That process then brings the cortical system to a
498 state suitable for initiating the selected action^{88,112}, setting into motion the “first cog” of a dynamical
499 machine that controls our actions in the world^{29–31}.

500 Acknowledgements

501 The authors wish to thank Terrence Sanger for valuable observations about our data, and John
502 Kalaska and Andrea Green for additional observations and helpful comments on the manuscript. PC
503 and DT designed the study, DT collected the data; DT, JFC, AF, and PC analyzed the data; DT and
504 PC wrote the manuscript; all authors approved the final version.

505 The data that support the findings of this study are available from the corresponding author upon
506 reasonable request.

507 **Methods**

508 **Subjects and apparatus**

509 Neural recordings were performed in two male rhesus monkeys (*Macaca mulatta*; *monkey S*: 4-9
510 years old, 5-9kg; *monkey Z*: 4-6 years old, 4-7kg). Animals were implanted, under anesthesia and
511 aseptic conditions, with a titanium head fixation post and recording chambers. The University of
512 Montreal animal ethics committee approved surgery, testing procedure and animal care.

513 Monkeys sat head-fixed in a custom primate chair and performed two planar reaching tasks using a
514 vertically oriented cordless stylus whose position was recorded by a digitizing tablet (*CalComp*,
515 125Hz). Their non-acting hand was restrained on an arm rest with Velcro bands. In most sessions,
516 unconstrained eye movements were recorded using an infrared camera (*ASL*, 120Hz). Stimuli and
517 continuous cursor feedback were projected onto a mirror suspended between the monkey's gaze and
518 the tablet, creating the illusion that they are in the plane of the tablet. Neural activity was recorded
519 from the hemisphere contralateral to the acting hand with tungsten microelectrodes (1.0 – 1.5 M Ω ,
520 *Frederic Haer* and *Alpha-Omega Eng.*) moved with a computer-controlled microdrive (*NAN*
521 *Instruments*) and electrophysiological signals were acquired with the AlphaLab system (*Alpha-Omega*
522 *Eng.*). Spikes were sorted offline using Plexon Offline Sorter (*Plexon Inc.*). Electrodes were targeted
523 based on 3D reconstructions (Brainsight, Rogue Research) using structural MRI images (Siemens 3.0
524 T). Behavioral data were collected with the software used to display the task (LabView, *National*
525 *Instruments*).

526 **Behavioral task and classification of trial types**

527 Monkeys were trained to perform the “tokens” task (Figure 1a) in which they are presented with one
528 central starting circle (1.75cm radius) and two peripheral target circles (1.75cm radius, arranged at
529 180° around a 5cm radius circle). The monkey begins each trial by placing a handle in the central
530 circle, in which 15 small tokens are randomly arranged. The tokens then begin to jump, one-by-one
531 every 200ms (“pre-decision interval”), from the center to one of the two peripheral targets always
532 oriented at 180° to each other with respect to the center. The monkey's task is to move the handle to
533 the target that he believes will ultimately receive the majority of tokens. The monkey is allowed to
534 make the decision as soon as he feels sufficiently confident, and has 500ms to bring the cursor into a
535 target after leaving the center. When the monkey reaches a target, the remaining tokens move more
536 quickly to their final targets (“post-decision interval”, which was either 150ms in “Slow” blocks or
537 50ms in “Fast” blocks (in a few sessions, the post-decision interval was reduced to 20ms in fast
538 blocks). Once all tokens have jumped, visual feedback is provided to the monkey (the chosen target
539 turns green for correct choices or red for error trials) and a drop of water or fruit juice is delivered for
540 choosing the correct target. A 1500ms inter-trial interval precedes the following trial. We alternated
541 between Slow and Fast blocks after about 75-125 trials, typically several times in each recording
542 session.

543 The monkeys were also trained to perform a delayed reach (DR) task (usually 30-48 trials per
544 recording session). In this task, the monkey again begins by placing the cursor in the central circle
545 containing the 15 tokens. Next, one of six peripheral targets is presented (1.75cm radius, spaced at 60°
546 intervals around a 5cm radius circle) and after a variable delay (500±100ms), the 15 tokens
547 simultaneously jump into that target. This “GO signal” instructs the monkey to move the handle to the
548 target to receive a drop of juice. This task is used to determine cell's task response and tuning as well
549 as the animal's mean reaction time (RT), used as an estimate of the total delays attributable to sensory
550 processing and response initiation. This quantity was then used to estimate the time of commitment in

551 the tokens task. That is, the “decision time” (DT) in the tokens task was quantified as movement onset
552 minus the mean RT from the DR task.

553 The tokens task allows us to calculate, at each moment in time, the “success probability” (SP)
554 associated with choosing each target. To characterize the success probability profile for each trial, we
555 calculated this quantity (with respect to the target ultimately chosen by the monkey) for each token
556 jump (Figure 1b). For example, with a total of 15 tokens, if at a particular moment in time the right
557 target contains N_R tokens, the left contains N_L tokens, and N_C tokens remain in the center, then the
558 probability that the target on the right will ultimately be the correct one (i.e., the success probability of
559 guessing right) is:

$$560 \quad p(R|N_R, N_L, N_C) = \frac{N_C!}{2^{N_C}} \sum_{k=0}^{\min(N_C, 7-N_L)} \frac{1}{k!(N_C-k)!} \quad (3)$$

561

562 Although each token jump in every trial was completely random, we could classify *a posteriori* some
563 specific classes of trials embedded in the fully random sequence (e.g. Figure 1c). In previous studies,
564 we defined “easy”, “ambiguous”, and “misleading” trials on the basis of their success probability
565 profile defined with respect to the first token jump. In contrast, because here we were primarily
566 interested in examining activity with respect to commitment, we defined these trial types according to
567 the success probability with respect to commitment time, estimated to be 280ms before movement
568 onset. A trial was classified as “easy” if the SP was above 0.5 five tokens before commitment, above
569 0.55 three tokens before, and above 0.65 at the time of commitment. A trial was classified as
570 “ambiguous” if SP was between 0.35 and 0.65 at five and three tokens before commitment as well as
571 at the time of commitment. A trial was classified as “misleading” if SP was below 0.5 five tokens
572 before commitment and then above 0.5 at the time of commitment. A trial was classified as “other” if
573 it didn’t meet any of these criteria. Note that these four classes are non-overlapping. In the “Slow”
574 block, we defined “early” trials as those where $DT < 1400$ and “late” as those where $DT \geq 1400$ ms. In
575 the “Fast” block, “early” trials were defined as those where $DT < 950$ and “late” as those where
576 $DT \geq 950$ ms.

577 In all tokens task trials, the targets were presented in opposite directions from the center, in two of six
578 possible locations around the circle. Their placement was chosen according to the tuning of recorded
579 cells, and in cases where a cell was not tuned, the leftmost and rightmost targets were used. Thus, in
580 all cases there was a target to the left of the center and one to the right (sometimes at an oblique
581 angle). Here, we grouped all of these into two groups: We defined the three targets between 90° and
582 270° as “left” targets, and the other three as “right” targets, and did not examine directional tuning in
583 any more detail.

584 For the analyses in the present report, we defined 28 task conditions as follows. First, we separated
585 trials into those recorded during the “Slow” block and those recorded during the “Fast” block, and in
586 each block we split trials into those in which a “left” target was chosen by the monkey and those in
587 which a “right” target was chosen. This yields 4 main groups: Slow-Left, Slow-Right, Fast-Left, and
588 Fast-Right. Next, we split each of these four main groups into Easy, Ambiguous, Misleading, and
589 Other, yielding another 16 conditions (Slow-Easy-Left, Slow-Easy-Right, Slow-Ambiguous-Left,
590 etc.). Finally, we again split our four main groups into Early and Late, yielding another 8 conditions
591 (Slow-Early-Left, Slow-Early-Right, Slow-Late-Left, etc.).

592 [Neural recording](#)

593 Detailed methods for neural recording in PMd and M1 are described in Thura & Cisek³⁶, and for
594 recording in the globus pallidus in Thura & Cisek⁴⁴. Recordings in dlPFC used the same methods as
595 PMd/M1, and focused on the region just dorsal to the caudal end of the principal sulcus (Extended
596 Data Figure 1). We used 2-4 independently moveable electrodes for cortical recordings, and one

597 electrode at a time for recordings in the globus pallidus. Thus, most of the cells whose activity is
598 analyzed and reported here were not recorded simultaneously.

599 In all sessions, we focused on cells showing any change of activity in the tokens task, and monkeys
600 were usually performing the task while we were searching for cells. When one or more task related
601 cells were isolated, we ran a block of 30-48 trials of the DR task to determine spatial tuning and select
602 a preferred target (PT) for each cell (i.e. the target associated with the highest firing rate during one or
603 more task epochs). Next, we ran blocks of tokens task trials using the PT of an isolated cell and the
604 180° opposite target (OT). We sometimes simultaneously recorded several task-related cells showing
605 different spatial preferences, and since we always selected a single pair of targets, the actual best
606 direction for all of the recorded cells was not always among these two. We usually started recording
607 cells in the slow block because monkeys were more conservative in this condition. It was thus easier
608 to assess cell properties online and more convenient to search for cells because fewer rewards were
609 spent. When possible, cells were tested with multiple repetitions of slow and fast blocks to control for
610 potential confounds related to evolving signals, elapsing time, and the monkey's fatigue or satiation
611 (see Thura and Cisek⁴³ for control analyses on this question).

612 Data analysis

613 A neuron was included in the analysis if it was recorded in both Slow and Fast blocks, thus including
614 trials in each of the four main conditions: Slow-Left, Slow-Right, Fast-Left, and Fast-Right. This
615 constraint was satisfied by 637 neurons out of the total 736 recorded across all brain regions (277 in
616 PMd, 191 in M1, 52 in dlPFC, 41 in GPe, and 46 in GPi). For each neuron and each trial, neural
617 activity was aligned to movement onset and the firing frequency was computed using partial spike
618 intervals in ninety 20ms-bins from 1400ms before to 400ms after movement onset. The firing rate was
619 then square root transformed and smoothed using a 25ms Gaussian kernel. Finally, we imposed
620 symmetry on our neural population by duplicating each of our neurons with an identical "anti-neuron"
621 that has the same activity in all trials, but with the trial labels switched between left and right choices.
622 We describe our rationale for including this step below in the section "Rationale for duplicating
623 neurons".

624 To calculate the principal components of neural activity, we then grouped trials into our four main
625 classes (Slow-Left, Slow-Right, Fast-Left, Fast-Right). In each group, neural activity was averaged
626 together for each individual neuron regardless of the type of trial (easy, ambiguous, misleading, or
627 other) and regardless of whether the choice was correct or not, or the reaction time was early or late.
628 For each of these four classes, we constructed a 60x1274 matrix where rows are time bins (from
629 1000ms before to 200ms after movements onset) and columns are individual neurons. We then
630 concatenated these four matrices into a 240x1274 matrix on which we performed standard Principal
631 Component Analysis (using the *pca* function in Matlab 2019b), yielding a matrix of weights ("loading
632 coefficients") between cells and PCs as well as the variance explained by each PC. For further
633 analysis, we only kept the top 20 PCs, which explained 97.9% of the total variance, though most of
634 the interpretation will focus on the first four (total 80.3% of variance).

635 In addition to standard PCA, we also tried other dimensionality reduction methods, including factor
636 analysis and Gaussian Process Factor Analysis (GPFA)^{68,70}, but these yielded almost identical results
637 (not shown). This is not surprising. For example, the major advantage of GPFA is that it jointly
638 provides filtering with dimensionality reduction, which is critical when analyzing data from individual
639 trials in which many neurons were recorded simultaneously. Our neurons were not recorded
640 simultaneously, so we had to combine trials into similarity classes under the assumption that specific
641 trial types were associated with similar neural activities on different occasions. This meant that we
642 averaged activity across multiple similar trials, yielding much of the noise reduction that GPFA would
643 have otherwise provided. Given the similarity of our results with different dimensionality reduction

644 techniques, we chose to use PCA because it is the simplest, most widely known and well-understood,
645 and ultimately easiest to interpret.

646 Using the 1274x20 loading matrix, we computed the average neural activity profile along each PC for
647 a given trial condition and a given population of neurons by multiplying each neuron's activity by its
648 loading coefficient for that PC, summing all of these together, and dividing by the number of neurons
649 in that population. We calculated confidence intervals around this average neural activity by randomly
650 resampling trials, with replacement within a given trial condition, 1000 times for each neuron and
651 time bin. We performed this process for all neurons together as well as for each of our neural
652 populations separately (PMd, M1, dlPFC, GPe, and GPi).

653 Because not every neuron possessed a trial in each condition, potentially making it inappropriate to
654 compare "PC space" trajectories from different conditions, we defined the following two groups of
655 neurons designed to facilitate the analyses of primary interest. Group 1 consisted of neurons that
656 possessed all conditions in the Slow block (Easy, Ambiguous, Misleading, Other, Early, and Late), for
657 both Left and Right choices, as well as Early and Late conditions in the Fast block. This group was
658 the focus of most of our analyses (e.g. Figure 3a,b; Figure 4; Figure 5), and included a total of 557
659 cells (258 in PMd, 173 in M1, 50 in dlPFC, 35 in GPe, and 41 in GPi). To facilitate comparisons
660 between trial types in the Fast block (e.g. Figure 3c,d), we also defined Group 2, which consisted of
661 neurons that possessed Easy, Ambiguous, Misleading, and Other conditions in the Fast block. This
662 group included 452 cells (226 in PMd, 126 in M1, 46 in dlPFC, and 27 each in GPe and GPi). When
663 calculating the temporal profile along the principal components (e.g. Figure 2), we used Group 1
664 neurons for all conditions except Easy, Ambiguous, Misleading, and Other trials in the Fast block.
665 Although these definitions were used to make quantitative comparisons most accurate, they did not
666 have a strong impact on the qualitative aspects of our data, which were similar even if we restricted all
667 analyses to the 402 neurons that possessed trials in all 28 conditions (Extended Data Figure 6d).

668 Once we reconstructed the temporal profile along each PC for given conditions, we plotted each of
669 them across time (Figure 2) and with respect to one another (e.g. Figure 3). To analyze the subspace
670 visited by a neural population during the deliberation process, we examined PCs 1, 2, and 4 (see
671 Figure 2) and constructed a 3-dimensional concave hull enclosing all of the neural states until the
672 moment of commitment, for all conditions within a given block (using the *alphaShape* function in
673 Matlab 2019b). For some of our brain regions this subspace resembles a thin 2-dimensional sheet,
674 which we call the "decision manifold". To quantify its thickness, we compared its surface area /
675 volume ratio with that of a perfect sphere with equal volume, using the equation $\Psi = \frac{\pi^{1/3}(6V)^{2/3}}{A}$,
676 where A is the surface area and V is the volume¹³. The quantity Ψ is called "sphericity", and is a
677 dimensionless scalar that ranges from 0 for a 2D surface to 1 for a perfect sphere (the geometric shape
678 with the smallest surface area / volume ratio). For the reader's intuition, we note that an equal-sided
679 cube of any size has $\Psi=0.806$ while a square sheet whose thickness is 1/50th of the length of each side
680 has $\Psi=0.171$.

681 To characterize the loading matrix, we considered each cell as a point in a high-dimensional space
682 defined by its loading coefficients for the top 11 PCs (which accounted for 95.2% of the total
683 variance). We then fit the resulting distribution of points for cells in a given brain region with
684 Gaussian Mixture Models (GMMs) using the Expectation-Maximization algorithm (*fitgmdist* function
685 in Matlab 2019b). For each fit, we tried GMMs consisting of anywhere from one to six 11-
686 dimensional gaussians, performing 100 randomly initialized fits for each, and then used the Bayesian
687 Information Criterion (BIC) to select the best fitting model. We found that with 100 randomly
688 initialized fits, we reliably found the same best solution each time.

689 Rationale for duplicating neurons

690 The rationale for duplicating the neurons stems from the fact that our population is inevitably a highly
691 sparse under-sampling of the millions of neurons in the regions where we recorded. Importantly, we
692 can assume that this under-sampling is not symmetric with respect to the proportions of neurons that
693 contribute to movements to the left versus right. For example, we might have in our sample many
694 high-firing cells preferring leftward movements, but fewer high-firing cells preferring rightward
695 movements. Applying PCA to such asymmetrically under-sampled data will force the algorithm to try
696 to capture this variance by aligning PCs to that asymmetry. We believe this is not informative. We can
697 assume that our sampling is unlikely to be symmetric with respect to leftward or rightward
698 movements, so it will not reflect any real asymmetry that may or may not exist in the brain. It will
699 only distort our data in ways that will make interpretations more difficult.

700 For this reason, just before performing PCA we imposed symmetry on our data by using the “anti-
701 neuron” approach classically used to produce population histograms of neural activity. In short, we
702 assumed that for every given neuron that we recorded, there exists in the brain another neuron that we
703 didn’t record, which has the same properties (same firing rate profile, same sensitivity to relevant
704 variables, etc.) but has the opposite relationship with respect to the direction of movement. To
705 implement this assumption, for each of our neurons we created a “sister” neuron with the same firing
706 data except with the trial labels switched between left and right choices. We did this for all neurons
707 regardless of whether they are tuned or not. All of the analyses shown in the main paper apply this
708 anti-neuron duplication, but in the supplemental data we show analogous results obtained without that
709 step. As can be seen (Extended Data Figure 7), the only difference is that the PCs produced without
710 anti-neuron duplication are rotated and skewed versions of the ones produced after anti-neuron
711 duplication, and their properties are slightly mixed between adjacent PCs. However, the general
712 conclusions remain unchanged.

713 Control analyses

714 To determine whether the PCs we found in the population data are a trivial consequence of properties
715 of single neurons (e.g. directional tuning), we applied a control based on the Tensor Maximum
716 Entropy (TME) method of Elsayed & Cunningham⁸⁴, using code they provide at
717 <https://github.com/gamaleldin/TME>. Briefly, this method preserves the primary first and second order
718 covariance in the data along the temporal, neural, and condition-dependent dimensions but is
719 otherwise maximally random. For example, tuning is preserved but not correlated across time. Any
720 metric applied to analyzing the real data can then be compared to that same metric applied to
721 surrogate data sets. If the metric computed from the real data lies within the distribution of that metric
722 computed from surrogate data, then whatever is measured by the metric is a simple consequence of
723 first or second order covariance.

724 Here, we were interested to know whether our finding of PCs that covary with evidence is a simple
725 consequence of cell tuning. To that end, we computed the correlation between the evidence provided
726 by token jumps (e.g. Figure 1c) and the temporal profiles of the components generated by PCA when
727 applied to surrogate data. Our metric was equal to the maximum absolute value of the correlation
728 coefficient for any of the top 10 components. We calculated this metric for the components produced
729 using the real data (in which the best correlation was with PC2) and compared it to the distribution of
730 the metric for 100 surrogate data sets generated using the TME method. That is, for each of the 100
731 surrogate data sets we performed PCA on the four main trial classes, then used these to calculate the
732 temporal profile of the top 10 PCs for all trial types, correlated each of these with evidence, and used
733 the highest correlation value. We consider the emergence of an evidence-related PC as a non-trivial
734 consequence of tuning if the correlation metric of fewer than 5 of these 100 surrogate data sets is
735 equal to or higher than the metric applied to the real data (i.e. $p < 0.05$).

736

737 References

- 738 1. Cisek, P. & Kalaska, J. F. Neural Mechanisms for Interacting with a World Full of Action
739 Choices. *Annual Review of Neuroscience* **33**, 269–298 (2010).
- 740 2. Gomez-Marin, A. & Ghazanfar, A. A. The Life of Behavior. *Neuron* **104**, 25–36 (2019).
- 741 3. Michalski, J., Green, A. M. & Cisek, P. Reaching decisions during ongoing movements. *Journal*
742 *of Neurophysiology* **123**, 1090–1102 (2020).
- 743 4. Cisek, P. Cortical mechanisms of action selection: the affordance competition hypothesis.
744 *Philosophical Transactions of the Royal Society B: Biological Sciences* **362**, 1585–1599 (2007).
- 745 5. Platt, M. L. & Glimcher, P. W. Neural correlates of decision variables in parietal cortex. *Nature*
746 **400**, 233–238 (1999).
- 747 6. Roitman, J. D. & Shadlen, M. N. Response of Neurons in the Lateral Intraparietal Area during a
748 Combined Visual Discrimination Reaction Time Task. *J. Neurosci.* **22**, 9475–9489 (2002).
- 749 7. Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex:
750 specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814
751 (2005).
- 752 8. Kim, B. & Basso, M. A. Saccade target selection in the superior colliculus: a signal detection
753 theory approach. *J. Neurosci.* **28**, 2991–3007 (2008).
- 754 9. Hernández, A. *et al.* Decoding a Perceptual Decision Process across Cortex. *Neuron* **66**, 300–
755 314 (2010).
- 756 10. Klaes, C., Westendorff, S., Chakrabarti, S. & Gail, A. Choosing goals, not rules: deciding
757 among rule-based action plans. *Neuron* **70**, 536–548 (2011).
- 758 11. Ratcliff, R. A Theory of Memory Retrieval. *Psychological Review* **85**, 59–108 (1978).
- 759 12. Ratcliff, R., Smith, P. L., Brown, S. D. & McKoon, G. Diffusion Decision Model: Current
760 Issues and History. *Trends in Cognitive Sciences* **20**, 260–281 (2016).
- 761 13. Ditterich, J. Evidence for time-variant decision making. *European Journal of Neuroscience* **24**,
762 3628–3641 (2006).
- 763 14. Cisek, P., Puskas, G. A. & El-Murr, S. Decisions in Changing Conditions: The Urgency-Gating
764 Model. *Journal of Neuroscience* **29**, 11560–11571 (2009).
- 765 15. Gluth, S., Rieskamp, J. & Buchel, C. Deciding When to Decide: Time-Variant Sequential
766 Sampling Models Explain the Emergence of Value-Based Decisions in the Human Brain.
767 *Journal of Neuroscience* **32**, 10686–10698 (2012).
- 768 16. Thura, D., Beaugard-Racine, J., Fradet, C.-W. & Cisek, P. Decision making by urgency
769 gating: theory and experimental support. *J. Neurophysiol.* **108**, 2912–2930 (2012).
- 770 17. Murphy, P. R., Boonstra, E. & Nieuwenhuis, S. Global gain modulation generates time-
771 dependent urgency during perceptual choice in humans. *Nat Commun* **7**, 13526 (2016).
- 772 18. Farashahi, S., Ting, C.-C., Kao, C.-H., Wu, S.-W. & Soltani, A. Dynamic combination of
773 sensory and reward information under time pressure. *PLoS Comput Biol* **14**, e1006070 (2018).
- 774 19. Palestro, J. J., Weichart, E., Sederberg, P. B. & Turner, B. M. Some task demands induce
775 collapsing bounds: Evidence from a behavioral analysis. *Psychon Bull Rev* **25**, 1225–1248
776 (2018).
- 777 20. Kawato, M. Internal models for motor control and trajectory planning. *Current Opinion in*
778 *Neurobiology* **9**, 718–727 (1999).
- 779 21. Shadmehr, R. & Krakauer, J. W. A computational neuroanatomy for motor control. *Exp Brain*
780 *Res* **185**, 359–381 (2008).
- 781 22. Franklin, D. W. & Wolpert, D. M. Computational mechanisms of sensorimotor control. *Neuron*
782 **72**, 425–442 (2011).
- 783 23. Grossberg, S. Contour Enhancement, Short Term Memory, and Constancies in Reverberating
784 Neural Networks. *Studies in Applied Mathematics* **52**, 213–257 (1973).
- 785 24. Amari, S. Dynamics of pattern formation in lateral-inhibition type neural fields. *Biol Cybern* **27**,
786 77–87 (1977).
- 787 25. Cisek, P. Integrated neural processes for defining potential actions and deciding between them:
788 a computational model. *J. Neurosci.* **26**, 9761–9770 (2006).
- 789 26. Wang, X.-J. Decision Making in Recurrent Neuronal Circuits. *Neuron* **60**, 215–234 (2008).

- 790 27. Wu, S., Hamaguchi, K. & Amari, S.-I. Dynamics and computation of continuous attractors.
791 *Neural Comput* **20**, 994–1025 (2008).
- 792 28. Standage, D., You, H., Wang, D.-H. & Dorris, M. C. Gain Modulation by an Urgency Signal
793 Controls the Speed–Accuracy Trade-Off in a Network Model of a Cortical Decision Circuit.
794 *Front. Comput. Neurosci.* **5**, (2011).
- 795 29. Bullock, D. & Grossberg, S. Neural dynamics of planned arm movements: emergent invariants
796 and speed-accuracy properties during trajectory formation. *Psychol Rev* **95**, 49–90 (1988).
- 797 30. Bullock, D., Cisek, P. & Grossberg, S. Cortical networks for control of voluntary arm
798 movements under variable force conditions. *Cereb. Cortex* **8**, 48–62 (1998).
- 799 31. Churchland, M. M., Cunningham, J. P., Kaufman, M. T., Ryu, S. I. & Shenoy, K. V. Cortical
800 preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*
801 **68**, 387–400 (2010).
- 802 32. Wallis, J. D. & Miller, E. K. From rule to response: neuronal processes in the premotor and
803 prefrontal cortex. *J. Neurophysiol.* **90**, 1790–1806 (2003).
- 804 33. Michelet, T., Duncan, G. H. & Cisek, P. Response competition in the primary motor cortex:
805 corticospinal excitability reflects response replacement during simple decisions. *Journal of*
806 *Neurophysiology* **104**, 119–127 (2010).
- 807 34. Pastor-Bernier, A. & Cisek, P. Neural correlates of biased competition in premotor cortex. *The*
808 *Journal of Neuroscience: The Official Journal of the Society for Neuroscience* **31**, 7083–7088
809 (2011).
- 810 35. Klein-Flügge, M. C. & Bestmann, S. Time-dependent changes in human corticospinal
811 excitability reveal value-based competition for action during decision processing. *The Journal of*
812 *Neuroscience: The Official Journal of the Society for Neuroscience* **32**, 8373–8382 (2012).
- 813 36. Thura, D. & Cisek, P. Deliberation and commitment in the premotor and primary motor cortex
814 during dynamic decision making. *Neuron* **81**, 1401–1416 (2014).
- 815 37. Cai, X. & Padoa-Schioppa, C. Contributions of orbitofrontal and lateral prefrontal cortices to
816 economic choice and the good-to-action transformation. *Neuron* **81**, 1140–1151 (2014).
- 817 38. Mink, J. W. The basal ganglia: focused selection and inhibition of competing motor programs.
818 *Progress in Neurobiology* **50**, 381–425 (1996).
- 819 39. Redgrave, P., Prescott, T. J. & Gurney, K. The basal ganglia: a vertebrate solution to the
820 selection problem? *Neuroscience* **89**, 1009–1023 (1999).
- 821 40. Turner, R. S. & Desmurget, M. Basal ganglia contributions to motor control: a vigorous tutor.
822 *Current Opinion in Neurobiology* **20**, 704–716 (2010).
- 823 41. Dudman, J. T. & Krakauer, J. W. The basal ganglia: from motor commands to the control of
824 vigor. *Current Opinion in Neurobiology* **37**, 158–166 (2016).
- 825 42. Thura, D., Cos, I., Trung, J. & Cisek, P. Context-dependent urgency influences speed-accuracy
826 trade-offs in decision-making and movement execution. *J. Neurosci.* **34**, 16442–16454 (2014).
- 827 43. Thura, D. & Cisek, P. Modulation of Premotor and Primary Motor Cortical Activity during
828 Volitional Adjustments of Speed-Accuracy Trade-Offs. *J. Neurosci.* **36**, 938–956 (2016).
- 829 44. Thura, D. & Cisek, P. The Basal Ganglia Do Not Select Reach Targets but Control the Urgency
830 of Commitment. *Neuron* **95**, 1160-1170.e5 (2017).
- 831 45. Ding, L. & Gold, J. I. Separate, causal roles of the caudate in saccadic choice and execution in a
832 perceptual decision task. *Neuron* **75**, 865–874 (2012).
- 833 46. Tai, L.-H., Lee, A. M., Benavidez, N., Bonci, A. & Wilbrecht, L. Transient stimulation of
834 distinct subpopulations of striatal neurons mimics changes in action value. *Nat. Neurosci.* **15**,
835 1281–1289 (2012).
- 836 47. Santacruz, S. R., Rich, E. L., Wallis, J. D. & Carmena, J. M. Caudate Microstimulation
837 Increases Value of Specific Choices. *Curr. Biol.* **27**, 3375-3383.e3 (2017).
- 838 48. Amemori, K.-I., Amemori, S., Gibson, D. J. & Graybiel, A. M. Striatal Microstimulation
839 Induces Persistent and Repetitive Negative Decision-Making Predicted by Striatal Beta-Band
840 Oscillation. *Neuron* **99**, 829-841.e6 (2018).
- 841 49. Wang, L., Rangarajan, K. V., Gerfen, C. R. & Krauzlis, R. J. Activation of Striatal Neurons
842 Causes a Perceptual Decision Bias during Visual Change Detection in Mice. *Neuron* **97**, 1369-
843 1381.e5 (2018).

- 844 50. Arimura, N., Nakayama, Y., Yamagata, T., Tanji, J. & Hoshi, E. Involvement of the Globus
845 Pallidus in Behavioral Goal Determination and Action Specification. *Journal of Neuroscience*
846 **33**, 13639–13653 (2013).
- 847 51. Hoshi, E. Cortico-basal ganglia networks subserving goal-directed behavior mediated by
848 conditional visuo-goal association. *Front. Neural Circuits* **7**, (2013).
- 849 52. Desmurget, M. & Turner, R. S. Testing Basal Ganglia Motor Functions Through Reversible
850 Inactivations in the Posterior Internal Globus Pallidus. *Journal of Neurophysiology* **99**, 1057–
851 1076 (2008).
- 852 53. Kalaska, J. F., Cohen, D. A., Hyde, M. L. & Prud'homme, M. A comparison of movement
853 direction-related versus load direction-related activity in primate motor cortex, using a two-
854 dimensional reaching task. *J. Neurosci.* **9**, 2080–2102 (1989).
- 855 54. Crammond, D. J. & Kalaska, J. F. Prior information in motor and premotor cortex: activity
856 during the delay period and effect on pre-movement activity. *J. Neurophysiol.* **84**, 986–1005
857 (2000).
- 858 55. Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron
859 activity in premotor and motor cortex. *J. Neurophysiol.* **97**, 4235–4257 (2007).
- 860 56. Hoshi, E. & Tanji, J. Distinctions between dorsal and ventral premotor areas: anatomical
861 connectivity and functional properties. *Curr. Opin. Neurobiol.* **17**, 234–242 (2007).
- 862 57. Johnson, P. B., Ferraina, S., Bianchi, L. & Caminiti, R. Cortical networks for visual reaching:
863 physiological and anatomical organization of frontal and parietal lobe arm regions. *Cereb.*
864 *Cortex* **6**, 102–119 (1996).
- 865 58. Chandrasekaran, C., Peixoto, D., Newsome, W. T. & Shenoy, K. V. Laminar differences in
866 decision-related neural activity in dorsal premotor cortex. *Nat Commun* **8**, 614 (2017).
- 867 59. Thura, D. & Cisek, P. A computational model of cortico-basal ganglia circuits for deciding
868 between reaching actions. Program No. 606.13. 2019 Neuroscience Meeting Planner. Chicago,
869 IL: Society for Neuroscience, 2019. Online. (2019).
- 870 60. Smith, A. C. & Brown, E. N. Estimating a state-space model from point process observations.
871 *Neural Comput* **15**, 965–991 (2003).
- 872 61. Stopfer, M., Jayaraman, V. & Laurent, G. Intensity versus identity coding in an olfactory
873 system. *Neuron* **39**, 991–1004 (2003).
- 874 62. Brown, S. L., Joseph, J. & Stopfer, M. Encoding a temporally structured stimulus with a
875 temporally structured neural representation. *Nat. Neurosci.* **8**, 1568–1576 (2005).
- 876 63. Levi, R., Varona, P., Arshavsky, Y. I., Rabinovich, M. I. & Selverston, A. I. The role of sensory
877 network dynamics in generating a motor program. *J. Neurosci.* **25**, 9807–9815 (2005).
- 878 64. Briggman, K. L., Abarbanel, H. D. I. & Kristan, W. B. From crawling to cognition: analyzing
879 the dynamical interactions among populations of neurons. *Curr. Opin. Neurobiol.* **16**, 135–144
880 (2006).
- 881 65. Broome, B. M., Jayaraman, V. & Laurent, G. Encoding and decoding of overlapping odor
882 sequences. *Neuron* **51**, 467–482 (2006).
- 883 66. Churchland, M. M., Yu, B. M., Sahani, M. & Shenoy, K. V. Techniques for extracting single-
884 trial activity patterns from large-scale neural recordings. *Curr. Opin. Neurobiol.* **17**, 609–618
885 (2007).
- 886 67. Sasaki, T., Matsuki, N. & Ikegaya, Y. Metastability of active CA3 networks. *J. Neurosci.* **27**,
887 517–528 (2007).
- 888 68. Yu, B. M. *et al.* Gaussian-process factor analysis for low-dimensional single-trial analysis of
889 neural population activity. *J. Neurophysiol.* **102**, 614–635 (2009).
- 890 69. Shenoy, K. V., Sahani, M. & Churchland, M. M. Cortical control of arm movements: a
891 dynamical systems perspective. *Annu. Rev. Neurosci.* **36**, 337–359 (2013).
- 892 70. Cunningham, J. P. & Yu, B. M. Dimensionality reduction for large-scale neural recordings. *Nat*
893 *Neurosci* **17**, 1500–1509 (2014).
- 894 71. Thura, D. & Cisek, P. Monkey frontal cortex reflects the time course of changing evidence for
895 reach decisions. Program No. 805.9. 2010 Neuroscience Meeting Planner. San Diego, CA:
896 Society for Neuroscience, 2010. Online. (2010).

- 897 72. Thura, D. & Cisek, P. A cortico-basal ganglia network for dynamic decision-making. Program
898 No. 359.06. 2018 Neuroscience Meeting Planner. San Diego, CA: Society for Neuroscience,
899 2018. Online. (2018).
- 900 73. Cisek, P., Cabana, J.-F., Thura, D. & Feghaly, A. The dynamics of neural population activity
901 during decision-making. Program No. 176.22. 2015 Neuroscience Meeting Planner. Chicago,
902 IL: Society for Neuroscience, 2015. Online. (2015).
- 903 74. Cisek, P., Thura, D., Cabana, J.-F. & Feghaly, A. Neural dynamics of cortical and basal ganglia
904 circuits during dynamic decision-making. Program No. 543.07. 2016 Neuroscience Meeting
905 Planner. San Diego, IL: Society for Neuroscience, 2016. Online. (2016).
- 906 75. Thura, D., Guberman, G. & Cisek, P. Trial-to-trial adjustments of speed-accuracy trade-offs in
907 premotor and primary motor cortex. *J. Neurophysiol.* **117**, 665–683 (2017).
- 908 76. Afshar, A. *et al.* Single-trial neural correlates of arm movement preparation. *Neuron* **71**, 555–
909 564 (2011).
- 910 77. Churchland, M. M. *et al.* Neural population dynamics during reaching. *Nature* **487**, 51–56
911 (2012).
- 912 78. Michaels, J. A., Dann, B., Intveld, R. W. & Scherberger, H. Predicting Reaction Time from the
913 Neural State Space of the Premotor and Parietal Grasping Network. *J. Neurosci.* **35**, 11415–
914 11432 (2015).
- 915 79. Kaufman, M. T. *et al.* The Largest Response Component in the Motor Cortex Reflects
916 Movement Timing but Not Movement Type. *eneuro* **3**, ENEURO.0085-16.2016 (2016).
- 917 80. Michaels, J. A. & Scherberger, H. Population coding of grasp and laterality-related information
918 in the macaque fronto-parietal network. *Sci Rep* **8**, (2018).
- 919 81. Rappoport, D., Kaufman, M. T. & Churchland, A. K. A category-free neural population supports
920 evolving demands during decision-making. *Nat. Neurosci.* **17**, 1784–1792 (2014).
- 921 82. Miri, A. *et al.* Behaviorally Selective Engagement of Short-Latency Effector Pathways by
922 Motor Cortex. *Neuron* **95**, 683-696.e11 (2017).
- 923 83. Sauerbrei, B. A. *et al.* Cortical pattern generation during dexterous movement is input-driven.
924 *Nature* **577**, 386–391 (2020).
- 925 84. Elsayed, G. F. & Cunningham, J. P. Structure in neural population recordings: an expected
926 byproduct of simpler phenomena? *Nat. Neurosci.* **20**, 1310–1318 (2017).
- 927 85. Churchland, M. M. *et al.* Stimulus onset quenches neural variability: a widespread cortical
928 phenomenon. *Nat Neurosci* **13**, 369–378 (2010).
- 929 86. Ames, K. C., Ryu, S. I. & Shenoy, K. V. Neural dynamics of reaching following incorrect or
930 absent motor preparation. *Neuron* **81**, 438–451 (2014).
- 931 87. Gallego, J. A. *et al.* Cortical population activity within a preserved neural manifold underlies
932 multiple motor behaviors. *Nat Commun* **9**, 4233 (2018).
- 933 88. Churchland, M. M., Yu, B. M., Ryu, S. I., Santhanam, G. & Shenoy, K. V. Neural variability in
934 premotor cortex provides a signature of motor preparation. *J. Neurosci.* **26**, 3697–3712 (2006).
- 935 89. Hoshi, E., Shima, K. & Tanji, J. Task-dependent selectivity of movement-related neuronal
936 activity in the primate prefrontal cortex. *J. Neurophysiol.* **80**, 3392–3397 (1998).
- 937 90. Rainer, G., Asaad, W. F. & Miller, E. K. Selective representation of relevant information by
938 neurons in the primate prefrontal cortex. *Nature* **393**, 577–579 (1998).
- 939 91. Kim, J. N. & Shadlen, M. N. Neural correlates of a decision in the dorsolateral prefrontal cortex
940 of the macaque. *Nat. Neurosci.* **2**, 176–185 (1999).
- 941 92. Quintana, J. & Fuster, J. M. From perception to action: temporal integrative functions of
942 prefrontal and parietal neurons. *Cereb. Cortex* **9**, 213–221 (1999).
- 943 93. Fuster, J. M., Bodner, M. & Kroger, J. K. Cross-modal and cross-temporal association in
944 neurons of frontal cortex. *Nature* **405**, 347–351 (2000).
- 945 94. Miller, E. K. The prefrontal cortex and cognitive control. *Nat. Rev. Neurosci.* **1**, 59–65 (2000).
- 946 95. Rowe, J. B., Toni, I., Josephs, O., Frackowiak, R. S. & Passingham, R. E. The prefrontal cortex:
947 response selection or maintenance within working memory? *Science* **288**, 1656–1660 (2000).
- 948 96. Barraclough, D. J., Conroy, M. L. & Lee, D. Prefrontal cortex and decision making in a mixed-
949 strategy game. *Nat. Neurosci.* **7**, 404–410 (2004).
- 950 97. Duque, J., Lew, D., Mazzocchio, R., Olivier, E. & Ivry, R. B. Evidence for two concurrent
951 inhibitory mechanisms during response preparation. *J. Neurosci.* **30**, 3793–3802 (2010).

- 952 98. Leblois, A. Competition between Feedback Loops Underlies Normal and Pathological
953 Dynamics in the Basal Ganglia. *Journal of Neuroscience* **26**, 3567–3583 (2006).
- 954 99. Schall, J. D. On building a bridge between brain and behavior. *Annu Rev Psychol* **55**, 23–50
955 (2004).
- 956 100. Gold, J. I. & Shadlen, M. N. The Neural Basis of Decision Making. *Annu. Rev. Neurosci.* **30**,
957 535–574 (2007).
- 958 101. Shadlen, M. N., Kiani, R., Hanks, T. D. & Churchland, A. K. Neurobiology of decision making:
959 An intentional framework. in *Better than conscious? Decision making, the human mind, and*
960 *implications for institutions* 71–101 (MIT Press, 2008).
- 961 102. Thompson, K. G., Hanes, D. P., Bichot, N. P. & Schall, J. D. Perceptual and motor processing
962 stages identified in the activity of macaque frontal eye field neurons during visual search. *J.*
963 *Neurophysiol.* **76**, 4040–4055 (1996).
- 964 103. Sato, T. R. & Schall, J. D. Effects of stimulus-response compatibility on neural selection in
965 frontal eye field. *Neuron* **38**, 637–648 (2003).
- 966 104. Stanford, T. R., Shankar, S., Massoglia, D. P., Costello, M. G. & Salinas, E. Perceptual decision
967 making in less than 30 milliseconds. *Nat. Neurosci.* **13**, 379–385 (2010).
- 968 105. de Lafuente, V., Jazayeri, M. & Shadlen, M. N. Representation of accumulating evidence for a
969 decision in two parietal areas. *J. Neurosci.* **35**, 4306–4318 (2015).
- 970 106. Hernández, A., Zainos, A. & Romo, R. Temporal evolution of a decision-making process in
971 medial premotor cortex. *Neuron* **33**, 959–972 (2002).
- 972 107. Coallier, É., Michelet, T. & Kalaska, J. F. Dorsal premotor cortex: neural correlates of reach
973 target decisions based on a color-location matching rule and conflicting sensory evidence. *J.*
974 *Neurophysiol.* **113**, 3543–3573 (2015).
- 975 108. Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U. & Nieuwenhuis, S. The neural basis of the
976 speed–accuracy tradeoff. *Trends in Neurosciences* **33**, 10–16 (2010).
- 977 109. Forstmann, B. U. *et al.* Cortico-striatal connections predict control over speed and accuracy in
978 perceptual decision making. *Proceedings of the National Academy of Sciences* **107**, 15916–
979 15920 (2010).
- 980 110. Nakahashi, A. & Cisek, P. Dynamics of the neural state in premotor and parietal cortex during
981 multi-attribute decision-making. Program No. 228.01. 2019 Neuroscience Meeting Planner.
982 Chicago, IL: Society for Neuroscience, 2019. Online. (2019).
- 983 111. Brown, J. W., Bullock, D. & Grossberg, S. How laminar frontal cortex and basal ganglia circuits
984 interact to control planned and reactive saccades. *Neural Netw* **17**, 471–510 (2004).
- 985 112. Churchland, M. M., Afshar, A. & Shenoy, K. V. A central source of movement variability.
986 *Neuron* **52**, 1085–1096 (2006).
- 987 113. Wadell, H. Volume, Shape, and Roundness of Quartz Particles. *The Journal of Geology* **43**,
988 250–280 (1935).
- 989

1 **Unified neural dynamics of decisions and actions in the cerebral**
2 **cortex and basal ganglia**

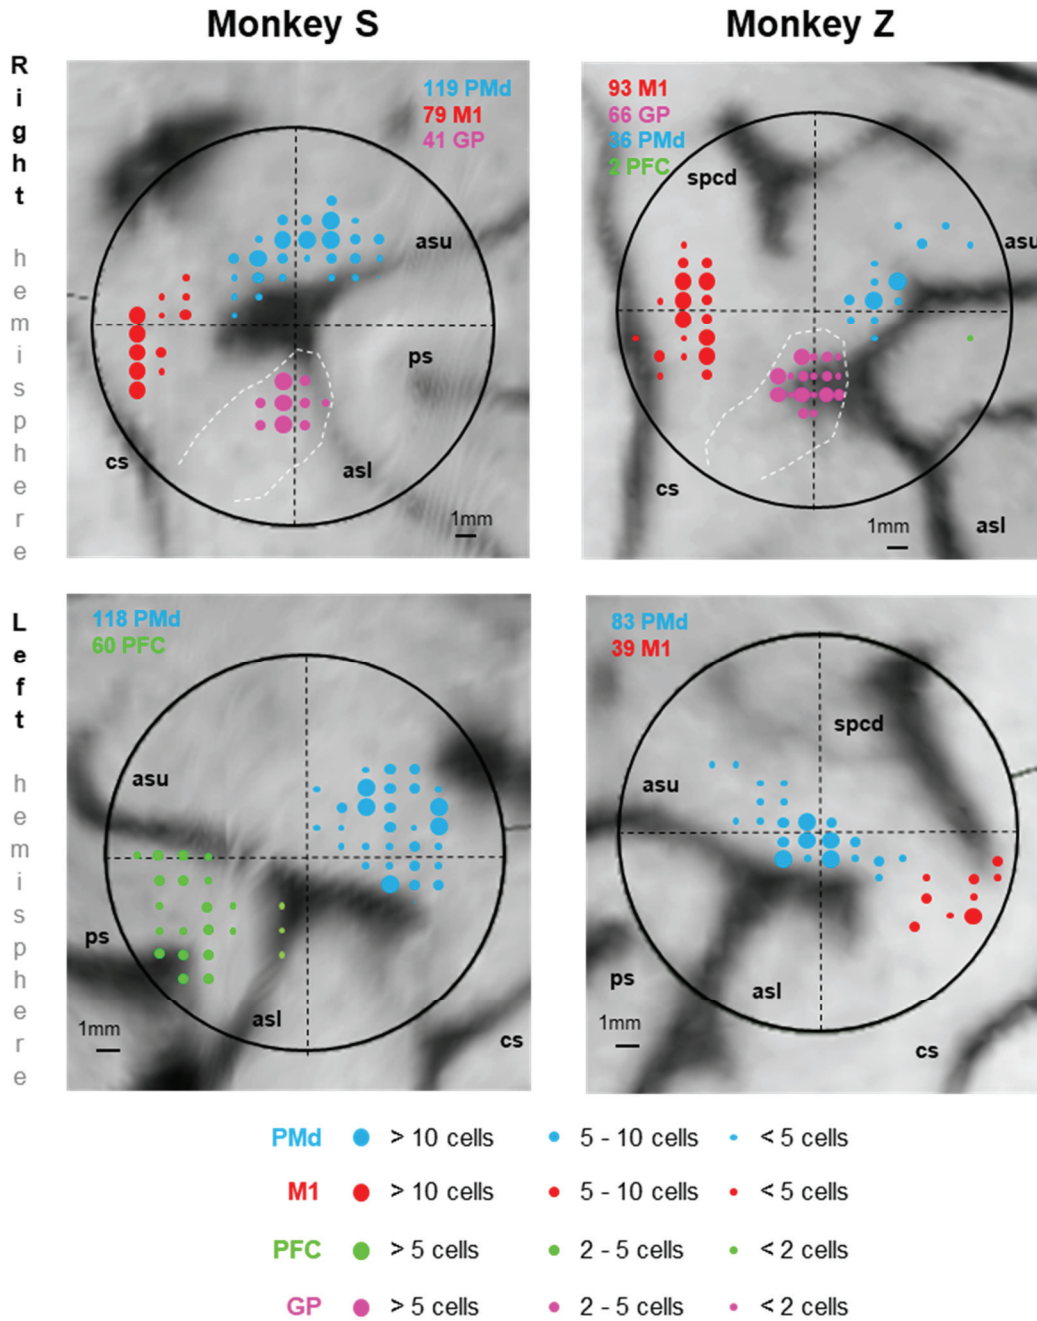
3 David Thura, Jean-François Cabana, Albert Feghaly, and Paul Cisek

4 Extended Data Figures 1-8

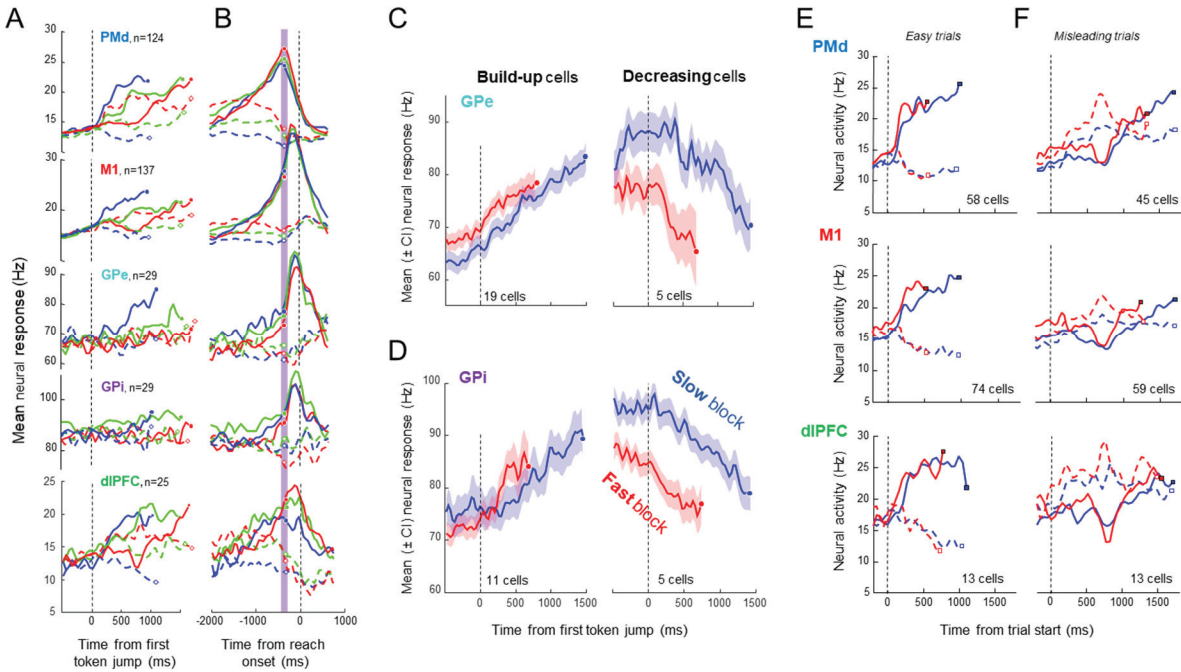
5 Supplemental analysis: Distinguishing gradients versus clusters of cell properties

6 Supplemental references


7




Extended Data Figure 1: Recording locations. Colored circles indicate the penetration locations of recorded neurons, superimposed over the reconstruction of the brain surface using Brainsight (Rogue Research). Large black circles indicate the extent of the recording chamber. Medial is up and anterior is to the right in the top panels and to the left in the bottom panels. Abbreviations: cs – central sulcus; ps – principal sulcus; asu – upper limb of the arcuate sulcus; asl – lower limb of the arcuate sulcus; spcd – superior precentral dimple.



Extended Data Figure 2: Neural activity during the tokens task. **A.** Average activity of directionally tuned neurons in five brain regions, computed separately for choices in each neuron's preferred (solid) and opposite direction (dashed) for easy (blue), ambiguous (green), and misleading trials (red), during Slow blocks. Activity is aligned on the start of token jumps and truncated at decision commitment (circles), and only cells exhibiting tuning prior to movement onset are included. PMd: dorsal premotor cortex; M1: primary motor cortex; GPe: globus pallidus externus; GPi: globus pallidus internus; dlPFC: dorsolateral prefrontal cortex. **B.** Average activity of the same cells aligned on movement onset. The purple vertical bar indicates the estimated time of commitment, 280ms prior to movement onset. **C.** Average activity of build-up cells (left) and decreasing cells (right) in the GPe during Slow (blue) and Fast blocks (red), aligned on the first token jump. Shading indicates the 95% confidence interval. **D.** Same as C for cells in the GPi. **E.** Average neural activity in PMd, M1, and dlPFC during easy trials in Slow (blue) and Fast blocks (red), when the monkey chose the target in each neuron's preferred (solid) or opposite direction (dashed). Here, only cells recorded in both blocks are included. **F.** Same as E for activity during misleading trials. Except for the PFC results, data in A-D is reproduced [with permission] from Thura & Cisek (2017) and E-F from Thura & Cisek (2016).

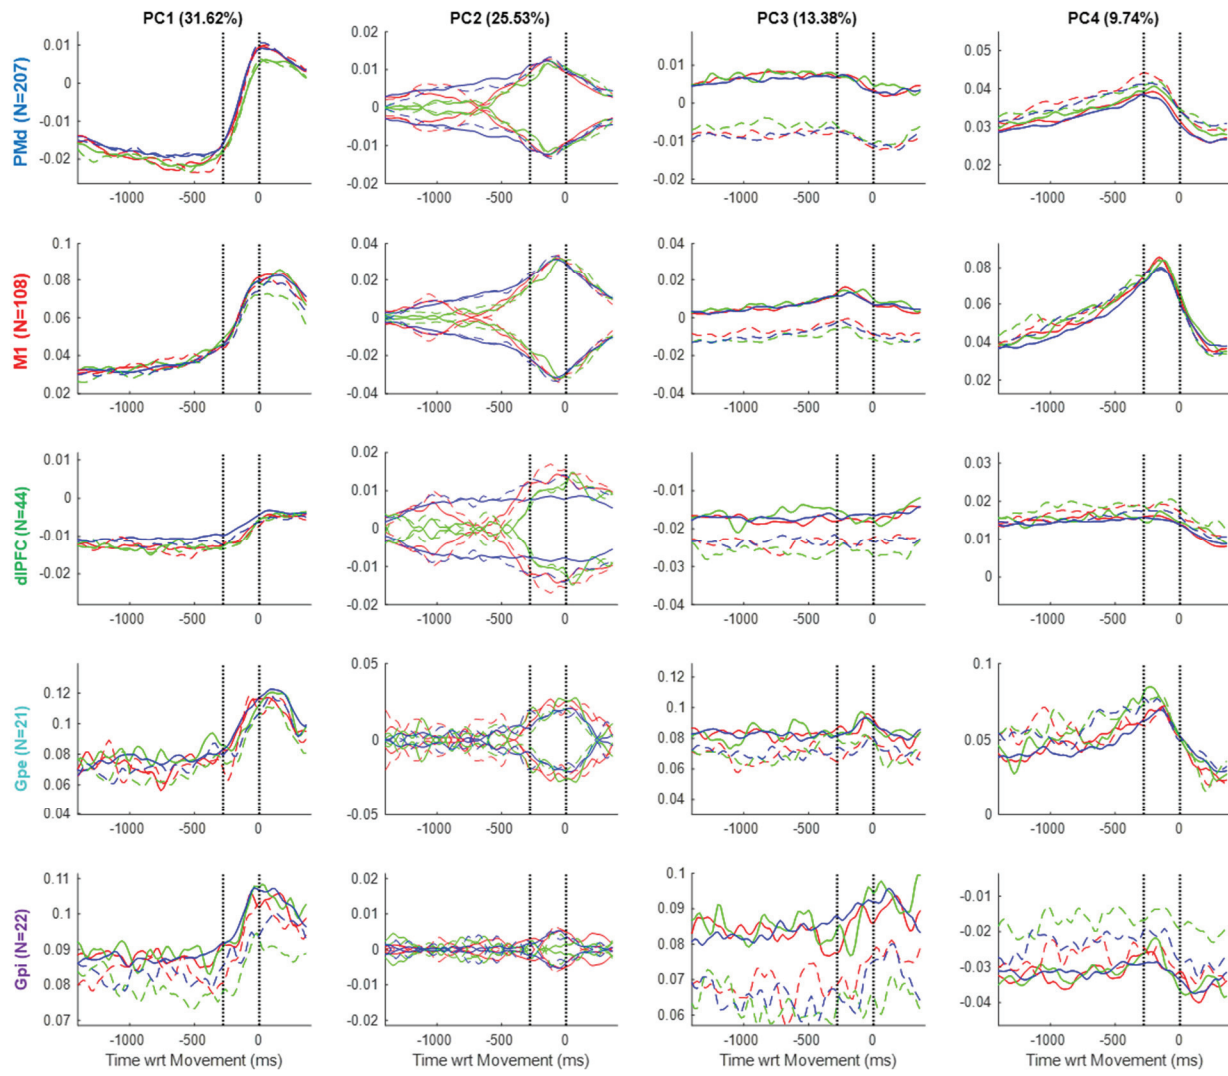


Extended Data Figure 3: The temporal profile of the principal components (as in Figure 2) computed using data aligned on the first token jump (dotted vertical line).



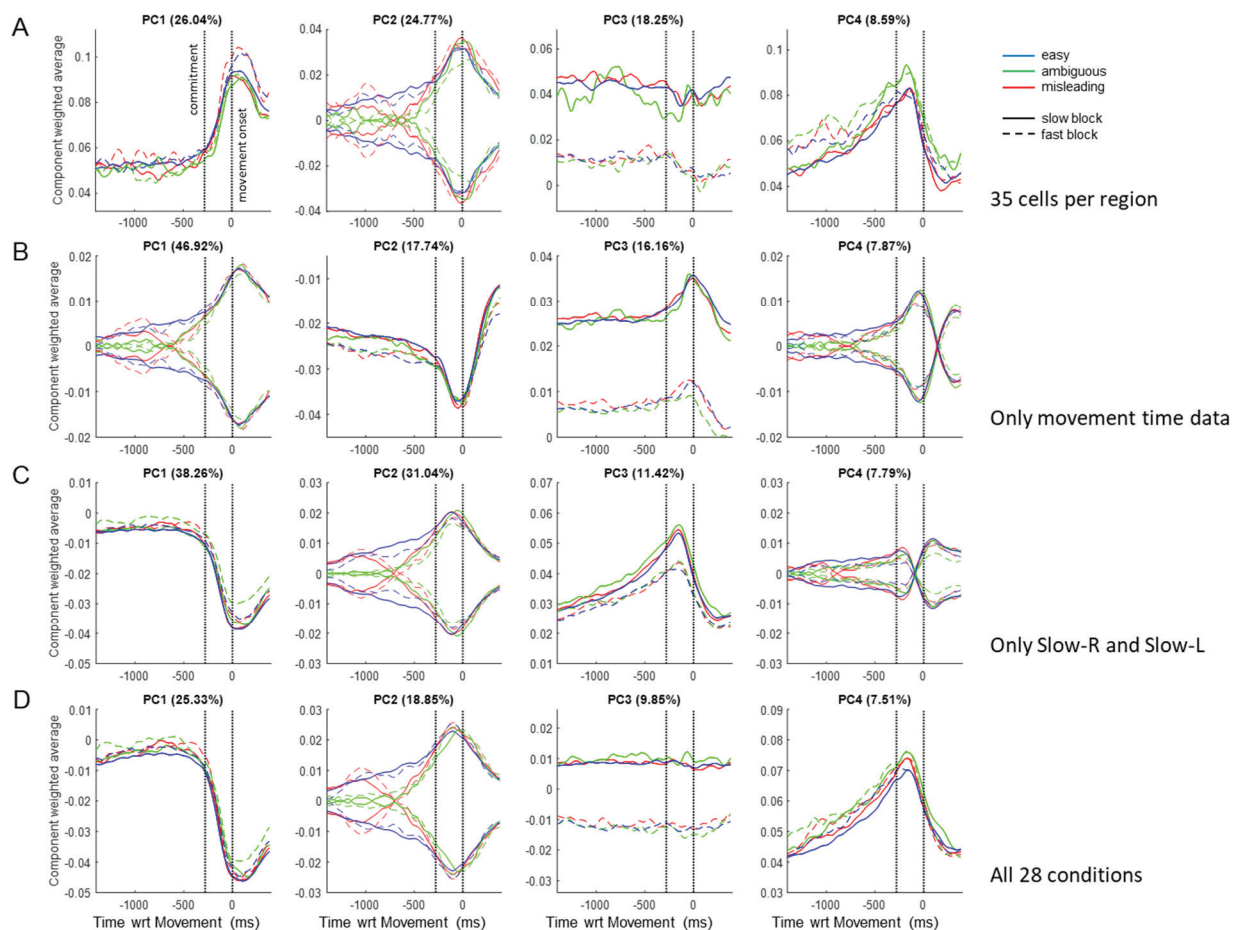
Extended Data Figure 4: Comparison between PC2 and sensory evidence. **A.** Dotted lines show the sensory evidence in easy (blue), ambiguous (green), and misleading trials (red) calculated as the difference between success probability for the right target and 0.5. Shaded ribbons show the mean and 95% confidence interval of PC2 in the same trials (Slow block). The first vertical dotted line indicates commitment and the second indicates movement onset, on which all data is aligned. The evidence trace is delayed by 300ms, which provides the best fit. Note that until the moment of commitment, the pattern of PC2 closely resembles the evidence, except for diverging toward one of the choices even in the absence of evidence during ambiguous trials). **B.** The same data, *prior to commitment*, plotted as evidence versus PC2 in these six trial conditions. The correlation coefficient is $R=0.9234$ and p -value is well below 0.001. **C.** The distribution ($N=100$) of correlation coefficients obtained by performing the same analysis on surrogate data sets generated using the Tensor Maximum Entropy approach (see Methods). Here, each surrogate data set is represented as the highest correlation coefficient of any of the top 10 PCs against the profile of evidence. The mean R is 0.5294 (s.t.d.=0.1615). For comparison, the red line shows the R value from the real data (panel B), and it is higher than all of the R values from surrogate data ($p<0.01$).

12

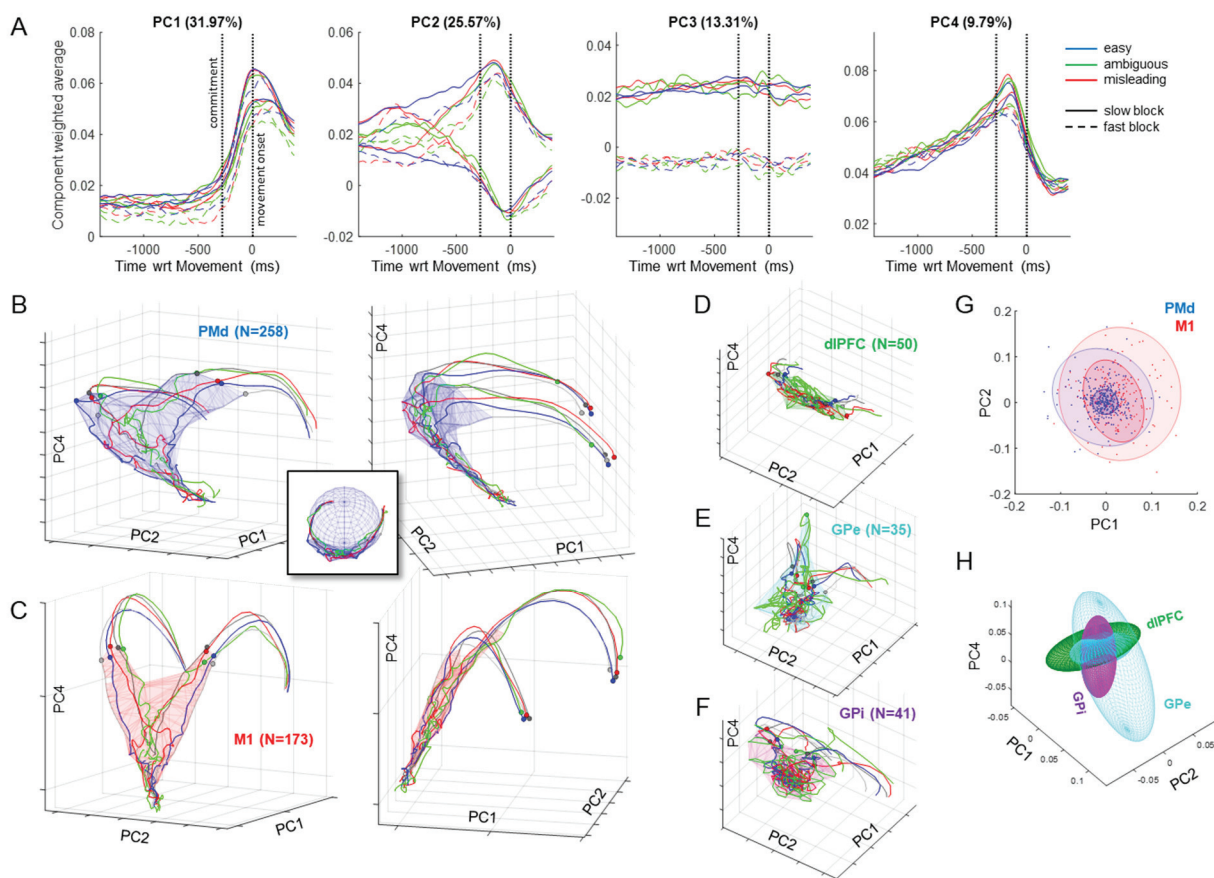


Extended Data Figure 5: The top four principal components shown in Figure 2, but here computed separately from each of the five brain regions. Note: To facilitate comparison between Slow and Fast blocks, here we only include cells that were recorded in all trial conditions.

13



Extended Data Figure 6: The first four principal components produced by providing PCA with different subsets of the data and then using the resulting loading matrix to compute PC profiles for all trials. **A.** PCs generated using data from only 35 neurons per region. Note that all of the general features of the PCs are similar to those in Figure 2, although they are more noisy. **B.** PCs generated using data only from the 400ms immediately following movement onset. The transition from deliberation to movement is now captured by PC2, the evidence by PC1, and the SAT by PC3. In contrast with the results reported in the main text, here there is no component related to elapsing time like PC4 in Figure 2. **C.** PCs generated using only data from Slow-Left and Slow-Right conditions. In comparison with Figure 2, there is now no SAT-related component and PC1 is inverted (note that the sign is arbitrary in PCA components). **D.** PCs generated using all 28 conditions, including activity from the 402 cells that possessed all of the 28 trial types. Note the similarity with the PCs shown in Figure 2, except for the inversion of PC1.



Extended Data Figure 7: Results produced without the cell duplication step. **A.** The first four principal components. Note the similarity to the PCs shown in Figure 2, except without the imposed symmetry. Thus, PCs 1, 3 and 4 are now slightly different for the two choices, and PC 2 is no longer perfectly symmetric. **B.** Two views of the 3-D trajectories in PMd. Dots indicate the moment of commitment on the left panel, and movement initiation on the right panel. Note that the decision manifold (blue wireframe) is still curved as in Figure 4a, and relatively well fit by a sphere (inset). **C.** Two views of M1, same format as A. Note that the decision manifold (red wireframe) is still flat as in Figure 4b. **D.** A view of dIPFC. **E.** GPe. **F.** GPi. For all populations, the patterns are similar as data processed using the cell duplication step, except that symmetry is lost. **G.** Analysis of loading matrices for PMd (blue) and M1 (red). Both are still best fit by a mixture of two gaussians similar to those shown on Figure 7a-d, but their orientations are now rotated. **H.** The distribution of loadings in dIPFC (green), GPe (cyan), and GPi (purple). While the dIPFC distribution is still similar to that in Figure 7f, GPe is now best fit by a 2-gaussian mixture and GPi is rotated. Nevertheless, despite the rotations the distributions of dIPFC and GPi are still roughly orthogonal.

16 Supplemental analysis: Distinguishing gradients versus clusters of cell properties

17 One of our key results is the apparent absence of distinct neural categories in our populations (Figure 7),
18 arguing for a unified dynamical system of cells with continuous properties. However, is this continuity
19 real or could it be an artifact of dimensionality reduction? To address this question, we created a variety
20 of synthetic populations of neuron-like units, in which we deliberately created specific categories, and
21 then applied to them the same analyses we used to examine real data. If our analyses fail to find these
22 categories, then we cannot be confident that real categories do not exist in the real data. We attempted a
23 variety of different ways to construct synthetic populations, but here report on two, each of which
24 constructs synthetic units using combinations of the PCs we obtained from the real data (Figure 2).

25 The first population consisted of three distinct groups of cells. Group 1 consisted of 200 neurons with
26 loading coefficients of 1 on PC1 and PC2, and zero for all others. Group 2 consisted of 200 neurons with
27 loading of 1 on PC2 and PC4 and zero for all others. Group 3 consisted of 200 neurons with loading of 1
28 on PC1 and PC4 and zero on others. These loading coefficients were multiplied by $1+0.25N$, where N is a
29 normally distributed random variable with mean 0 and standard deviation 1. The synthetic activity of each
30 cell was then computed as a sum of the temporal profile of each PC weighted by its loading coefficient
31 for that cell. Finally, noise was added by multiplying the value of each bin of activity by a uniformly
32 distributed random variable between 0 and 1.

33 We then subjected these synthetic units to exactly the same analyses we used for our real data, including
34 square-root transformation, smoothing, and duplication. PCA analysis was then used to compute loading
35 matrices of “synthetic” principal components (SPCs). The temporal profiles of the SPCs were then
36 computed using a weighted sum of all units, and the loading matrix was analyzed using GMMs. For the
37 synthetic data, we used 4-dimensional Gaussians to analyze the loading matrix.

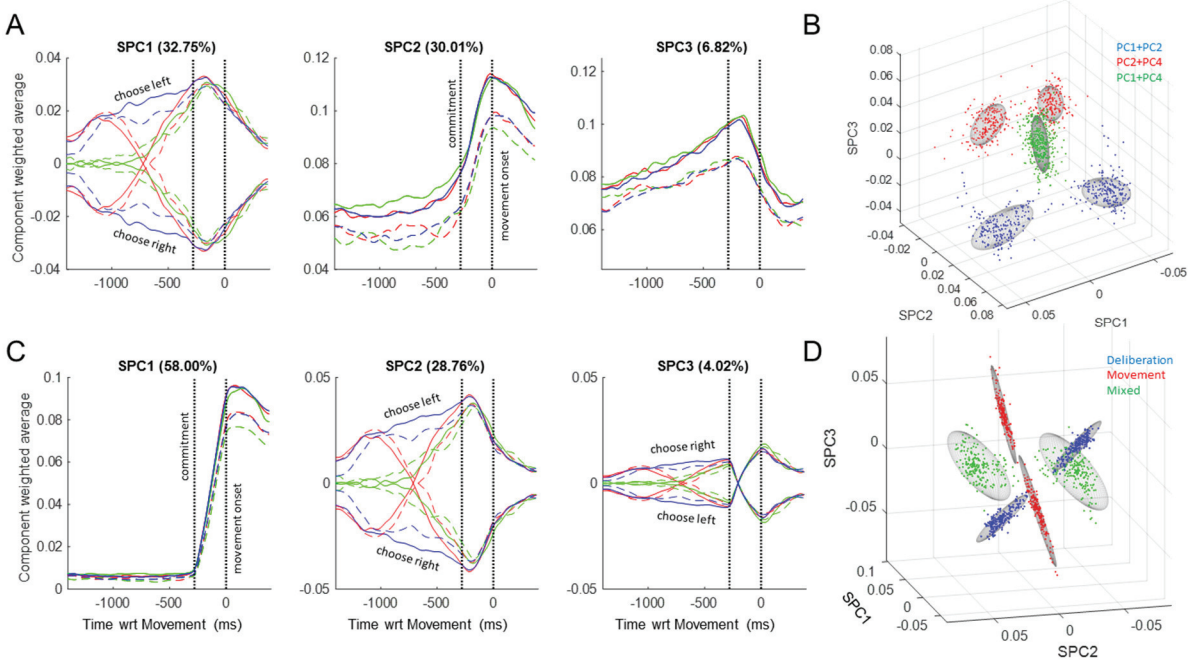
38 As shown in Extended Data Figure 8a, the “synthetic” principal components (SPCs) almost perfectly
39 capture the original PCs from which the cells were built, albeit in a different order (SPC1 is like PC2,
40 SPC2 is like PC1, and SPC3 is like PC4). Furthermore, although the GMM analysis was applied to all
41 600 units together, without information on how the different groups were built, it correctly identified the
42 relevant clusters (Extended Data Figure 8b).

43 The second synthetic population consisted of 600 neurons that all had loading coefficients of 1 on PC1
44 and PC2 and zero for others, again multiplied by $1+0.25N$, turned into activity profiles using a weighted
45 sum of the real PCs 1 and 2, and multiplied by a random variable uniformly distributed between 0 and 1.
46 Next, these 600 units were split into three groups: For “deliberation” units, activity before commitment
47 was multiplied by 1, activity between commitment and movement was scaled by a number linearly
48 dropping from 1 to zero, and activity after movement onset was set to zero. For “movement” units,
49 activity before commitment was set to zero, activity between commitment and movement was scaled by a
50 number linearly rising from zero to 1, and activity after movement onset was multiplied by 1. For
51 “mixed” units, all activity was kept unchanged.

52 We then ran PCA and obtained the synthetic components shown in Extended Data Figure 8c. Note that
53 SPC1 is similar to PC1, SPC2 is like PC2, but now we also find an SPC3 that looks like PC2 except for a
54 switch of activity between deliberation and movement. This is reminiscent of some of the higher
55 components we found in the neural data (see PC5 and PC6 in Figure 2). Importantly, *no such component*
56 *was used to build the synthetic units*. So where did it come from? The answer lies in the PCA algorithm,
57 which sequentially identifies components on the basis of variance explained. After finding SPC1 it
58 discovered SPC2, which together explain a large proportion of activity for both “deliberation” and
59 “mixed” units. Since no linear combination of these two components could explain “movement” units, the

60 next component (SPC3) captures some of the remaining unaccounted variance. In particular, note that a
 61 linear combination of SPC2 plus 3 times SPC3 can cancel out deliberation activity and produce a
 62 “movement” unit.

63 These observations on synthetic data suggest that the higher principal components we found in our real
 64 data (PCs 5, 6, 7, etc.) also do not necessarily reveal additional “higher order” features of neural
 65 dynamics, but simply result from the type of heterogeneity that has long been observed in sensorimotor
 66 cortical regions¹⁻⁶. But this then revives the question of functional clusters of real neurons – might distinct
 67 categories exist in these regions? If they did, would our analysis of loading matrices identify them
 68 correctly? Extended Data Figure 8d shows the result of GMM fits to the entire population of 600
 69 synthetic units, again performed without any information on the underlying groups. Clearly, the relevant
 70 clusters were found. In particular, “movement” units (red) were identified with two Gaussians (one for
 71 right-tuned and one for left-tuned units), loaded positively onto SPC1, and onto both SPC2 and SPC3
 72 with 3 times larger coefficients for the latter – consequently cancelling out their deliberation-time activity.
 73 In contrast, “deliberation” units (blue) were identified with two oriented Gaussians orthogonal to those of
 74 “movement” units and negatively loaded on SPC1. Finally, “mixed” units (green) were identified as
 75 clusters lying in-between. This demonstrates that if distinct categories of neurons really did exist in the
 76 real populations we recorded, the GMM analysis would have identified them correctly.



Extended Data Figure 8. Analyses of synthetic populations. **A.** The top 3 Synthetic PCs (SPCs) obtained from a population of 600 neurons constructed using combinations of the original PCs 1, 2, and 4 (see Figure 2). **B.** The loading coefficients of all 600 neurons, fitted with GMMs (gray ellipsoids). Colors indicate the three groups of neurons built with different combinations of those original PCs (see legend). **C.** The top 3 SPCs obtained from a population of 600 neurons constructed only using the original PC1 and PC2, but separated into three distinct categories: cells that are only active during deliberation, cells only active during movement, and cells active during both epochs. **D.** The loading coefficients of all 600 neurons, fitted with GMMs (gray ellipsoids). Colors indicate the three categories of neurons.

77

78

79 Supplemental references

- 80 1. Kalaska, J. F., Cohen, D. A., Hyde, M. L. & Prud'homme, M. A comparison of movement direction-
81 related versus load direction-related activity in primate motor cortex, using a two-dimensional
82 reaching task. *J. Neurosci.* **9**, 2080–2102 (1989).
- 83 2. Crammond, D. J. & Kalaska, J. F. Prior information in motor and premotor cortex: activity during the
84 delay period and effect on pre-movement activity. *J. Neurophysiol.* **84**, 986–1005 (2000).
- 85 3. Wallis, J. D. & Miller, E. K. From rule to response: neuronal processes in the premotor and prefrontal
86 cortex. *J. Neurophysiol.* **90**, 1790–1806 (2003).
- 87 4. Cisek, P. & Kalaska, J. F. Neural correlates of reaching decisions in dorsal premotor cortex:
88 specification of multiple direction choices and final selection of action. *Neuron* **45**, 801–814 (2005).
- 89 5. Churchland, M. M. & Shenoy, K. V. Temporal complexity and heterogeneity of single-neuron activity
90 in premotor and motor cortex. *J. Neurophysiol.* **97**, 4235–4257 (2007).
- 91 6. Hoshi, E. & Tanji, J. Distinctions between dorsal and ventral premotor areas: anatomical connectivity
92 and functional properties. *Curr. Opin. Neurobiol.* **17**, 234–242 (2007).
- 93