1 **Multiple timescales of sensory-evidence accumulation**

2 **across the dorsal cortex**

3 Lucas Pinto^{1,§}, David W. Tank^{1,2,#,*}, Carlos D. Brody^{1,2,3,#,*}

4 ¹ Princeton Neuroscience Institute, ² Department of Molecular Biology, ³ Howard Hughes
5 Medical Institute, Princeton University, Princeton, NJ, 08544 USA

6 [§] Present address: Department of Physiology, Feinberg School of Medicine, Northwestern
7 University, Chicago, IL, 60611 USA

8 [#] these senior authors contributed equally to this work

9 ^{*} corresponding authors:

10 David W. Tank, dwtank@princeton.edu

11 Carlos D. Brody, brody@princeton.edu

12 **Abstract**

13 Cortical areas seem to form a hierarchy of intrinsic timescales, but whether this is causal to
 14 cognitive behavior remains unknown. In particular, decisions requiring the gradual accrual of
 15 sensory evidence over time recruit widespread areas across this hierarchy. Here, we causally
 16 tested the hypothesis that this recruitment is related to the intrinsic integration timescales of
 17 these widespread areas. We trained mice to accumulate evidence over seconds while
 18 navigating in virtual reality, and optogenetically silenced the activity of many cortical areas
 19 during different brief trial epochs. We found that the inactivation of different areas primarily
 20 affected the evidence-accumulation computation per se, rather than other decision-related
 21 processes. Specifically, we observed selective changes in the weighting of evidence over time,
 22 such that frontal inactivations led to deficits on longer timescales than posterior cortical ones.
 23 Likewise, large-scale cortical Ca^{2+} activity during task performance displayed different temporal
 24 integration windows matching the effects of inactivation. Our findings suggest that distributed
 25 cortical areas accumulate evidence by leveraging their hierarchy of intrinsic timescales.

26 **Introduction**

27 The cerebral cortex of both rodents and primates appears to be organized in a hierarchy of
 28 intrinsic integration timescales, whereby frontal areas integrate input over longer time windows
 29 than sensory areas (Cavanagh et al., 2020; Chaudhuri et al., 2015; Gao et al., 2020; Hasson et
 30 al., 2008; Ito et al., 2020; Kiebel et al., 2008; Murray et al., 2014; Runyan et al., 2017; Soltani et
 31 al., 2021; Spitmaan et al., 2020). Although this idea has received increasing attention, there is
 32 still no causal evidence that such timescale hierarchy is relevant for cognitive behavior.

In particular, the decisions we make in our daily lives often unfold over time as we deliberate between competing choices. This raises the possibility that decisions co-opt the cortical timescale hierarchy such that different cortical areas integrate decision-related information on distinct timescales. A commonly studied type of time-extended decision making happens under perceptual uncertainty, which requires the gradual accrual of sensory evidence (Bogacz et al., 2006; Brody and Hanks, 2016; Brunton et al., 2013; Carandini and Churchland, 2013; Gold and Shadlen, 2007; Morcos and Harvey, 2016; Newsome et al., 1989; Odoemene et al., 2018; Stine et al., 2020; Sun and Landy, 2016; Tsetsos et al., 2012; Waskom and Kiani, 2018). Neural correlates of decisions relying on evidence accumulation have been found in a number of cortical and subcortical structures, in both primates and rodents (Brincat et al., 2018; Ding and Gold, 2010; Erlich et al., 2015; Hanks et al., 2015; Horwitz and Newsome, 1999; Kim and Shadlen, 1999; Koay et al., 2020; Krueger et al., 2017; Murphy et al., 2020; Orsolic et al., 2021; Scott et al., 2017; Shadlen and Newsome, 2001; Wilming et al., 2020; Yartsev et al., 2018). Likewise, we have previously shown that, when mice must accumulate evidence over several seconds to make a navigational decision, the inactivation of widespread dorsal cortical areas leads to behavioral deficits, and that these areas encode multiple behavioral variables, including evidence (Pinto et al., 2019). However, we do not understand which aspects of these decisions lead to such widespread recruitment of brain structures.

Here, we hypothesized that the pattern of widespread recruitment of cortical areas during prolonged evidence accumulation can be explained by their underlying timescale hierarchy. To test this, we trained mice to accumulate evidence over seconds towards navigational decisions, and used brief optogenetic inactivation of single or combined cortical areas, restricted to one of six epochs of the behavioral trials. We show that the inactivation of widespread areas in the dorsal cortex affects primarily the evidence accrual process, rather

than other decision-related computations. Further, the inactivation of different areas affects accumulation over distinct timescales, such that to an approximation frontal areas encode sensory evidence over longer temporal windows than posterior areas. In agreement with this, we show that cortical activity during the accumulation task displays a gradient of timescales, which are longer in frontal areas. Our findings thus suggest that evidence is accumulated by distributed cortical regions leveraging an existing hierarchy of temporal integration windows. Further, to our knowledge, they provide the first causal demonstration that this hierarchy is important for cognitive behavior.

Results

Brief inactivation of different cortical areas leads to accumulation deficits on distinct timescales. We trained mice to accumulate evidence over relatively long timescales while navigating in virtual reality (VR)(Figure 1A)(Pinto et al., 2018). The mice navigated a 3 m-long virtual T-maze and during the first 2 m (~4 s) they encountered salient objects, or towers, along the walls on either side, and after a delay of 1 m (~2 s) turned into the arm corresponding to the highest perceived tower count. The towers were visible for 200 ms, and appeared at different positions in each trial, obeying spatial Poisson processes of different underlying rates on the rewarded and non-rewarded side. Compatible with our previous reports (Koay et al., 2020; Pinto et al., 2018, 2019), task performance was modulated by the difference in tower counts between the right and left sides (Figure 1B, $n = 20$). Crucially, beyond allowing us to probe sensitivity to sensory evidence, the task design decorrelated the position of individual towers from the animals' position in the maze across trials. This allowed us to build a logistic

78 regression model that used the net sensory evidence (Δ towers, or #R – #L) from each of four
 79 equally-spaced bins from the cue region to predict the choice the mice made. In other words,
 80 we inferred the weight of sensory evidence from different positions in the maze on the final
 81 decision. While individual mice showed different evidence-weighting profiles, fitting the model
 82 on aggregate data yielded a flat evidence-weighting curve (Figure 1C, $n = 100,787$ trials),
 83 indicating that on average the mice weighted evidence equally from throughout the maze (Pinto
 84 et al., 2018). Note that all of the analyses presented below are performed on aggregate data
 85 (combined across mice, see Materials and Methods), such that our baseline condition is of
 86 even evidence weighting throughout the maze.

87 Our previous results have shown that cortical contributions to the performance of this
 88 task are widespread (Pinto et al., 2019), but our whole-trial inactivations did not allow us to
 89 tease apart the nature of the contributions of different areas. Here, we addressed this by asking
 90 how different dorsal cortical regions contribute to the weighting of sensory evidence in order to
 91 make a perceptual decision. To do this we cleared the intact skull of mice expressing
 92 Channelrhodopsin-2 (ChR2) in inhibitory interneurons (VGAT-ChR2-EYFP, $n = 20$), and used a
 93 scanning laser system to bilaterally silence different cortical regions, by activating inhibitory
 94 cells (Figure 1D) (Guo et al., 2014; Pinto et al., 2019). We targeted 7 different areas – primary
 95 visual cortex (V1), medial secondary visual cortex (mV2, roughly corresponding to area AM),
 96 posterior parietal cortex (PPC), retrosplenial cortex (RSC), the posteromedial portion of the
 97 premotor cortex (mM2), the anterior portion of the premotor cortex (aM2), and the primary
 98 motor cortex (M1) – as well as two combinations of these individual areas, namely posterior
 99 cortex (V1, mV2, PPC and RSC) and frontal cortex (mM2, aM2 and M1). Cortical silencing
 100 occurred in one of six trial epochs: 1st, 2nd or 3rd quarter of the cue region (0 – 50 cm, 50 – 100
 101 cm or 100 – 150 cm, respectively), 1st or 2nd half of the cue region (0 – 100 cm or 100 – 200 cm,

102 respectively), or delay region (200 – 300 cm). We tested all the 54 possible area-epoch
 103 combinations (Figure 1–table supplement 1). This large number of experimental conditions
 104 allowed us to assess how the inactivation of different areas affects the use of current or past
 105 sensory evidence towards a final decision.

106 **Figure 1. Temporally specific inactivation of multiple dorsal cortical regions during performance of**
 107 **a VR-based evidence-accumulation task.**
 108 **(A)** Schematics of the experimental set-up. **(B)** Psychometric functions for control trials, showing the
 109 probability of right-side choice as a function of the strength of right sensory evidence, Δ towers (#R –
 110 #L). Thin gray lines: best-fitting psychometric functions for each individual mouse (n = 20). Black circles:
 111 aggregate data (n = 100,787 trials), black line: fit to aggregate data, error bars: binomial confidence
 112 intervals. **(C)** Logistic regression curves for the weight of sensory evidence from four equally-spaced bins
 113 on the final decision, from control trials. Thin gray lines: individual animals, thick black line: aggregate

114 data, error bars: $\pm$ SD from 200 bootstrapping iterations. **(D)** Experimental design. We bilaterally
 115 inactivated 7 dorsal cortical areas, alone or in combination, while mice performed the
 116 accumulating-towers task. Bilateral inactivations happened during one of six regions in the maze
 117 spanning different parts of the cue region or delay. We thus tested a total of 54 area-epoch
 118 combinations. **(E)** Effects of subtrial inactivations on overall performance during all 54 area-epoch
 119 combinations. Each panel shows inactivation-induced change in overall % correct performance for each
 120 inactivation epoch, for data combined across mice. Error bars: S.D. across 10,000 bootstrapping
 121 iterations. Circles indicate significance according to the captions on the leftmost panel.

122 Compatible with our previous whole-trial inactivation experiments (Pinto et al., 2019),
 123 we found that the inactivation of all tested cortical areas significantly affected behavioral
 124 performance, though to varying degrees (Figure 1E, Figure 1-figure supplement 1).
 125 Furthermore, we observed a variety of effect profiles across regions and inactivation epochs,
 126 as assessed by the difference between the evidence-weighting curves separately calculated
 127 for 'laser off' and 'laser on' trials (Figure 1-figure supplement 2). Different effects were
 128 observed even comparing regions that were in close physical proximity (e.g. V1 and mV2).
 129 Additionally, all tested areas had significant effects in at least a subset of conditions (Figure
 130 1-figure supplement 2, $p < 0.05$, bootstrapping).

131 Most changes in the evidence-weighting curves happened for evidence concomitant to
 132 or preceding laser onset, indicating that the manipulations primarily affected the processing
 133 and/or memory of the evidence, i.e. the accumulation process itself (Figure 1-figure
 134 supplement 2). To quantify this, for each cortical area we aligned the control-subtracted
 135 evidence-weighting curves from different inactivation epochs by the position of laser onset,
 136 and focused on the changes in weights of evidence occurring up to 100 ms in the past (~2s,
 137 Figure 2A). While some variability in these laser-onset-triggered curves suggests that the
 138 effects of inactivation depend somewhat on the exact inactivation epoch, the aligned curves
 139 from different epochs were fairly consistent (Figure 2B, gray lines), providing a concise
 140 summary of the multiple experimental conditions. Interestingly, the effects varied systematically

141 according to the inactivated region. For example, mV2 inactivation led to a significant drop in
142 the weight of evidence occurring while the laser was on ($p = 0.004$, one-sided paired t test), a
143 trend towards affecting the memory of evidence occurring 50 cm in the past (~ 1 s, $p = 0.045$,
144 not significant after false discovery rate correction), and no discernible effect on evidence
145 occurring 100 cm in the past (~ 2 s, $p = 0.41$). Conversely, aM2 inactivation led to significant
146 decreases in weighting all evidence between 100 cm in the past and the time of laser onset (p
147 < 0.05 , one-sided paired t test), with no differences in magnitude between position bins ($F_{2,10} =$
148 0.27 , $p = 0.77$, mixed-model one-way ANOVA). This lack of modulation of effect size across y
149 position bins was also true for the two other frontal areas, mM2 and M1 (Figure 2B, $p > 0.88$).
150 Thus, subsets of cortical areas resembled each other in terms of the effects of their
151 inactivation. Indeed, they could be optimally grouped into three clusters using spectral
152 clustering (Figure 2C). Cluster 1 contained all frontal areas in our dataset: M1, mM2 and aM2.
153 On average, this cluster resembled the effects described for aM2 above. In other words, the
154 inactivation of frontal areas tended to equally and significantly affect weights for evidence
155 occurring up to 100 cm in the past, suggesting that these areas accumulate evidence at fairly
156 long timescales ($p < 0.001$ for all position bins, one-sided paired t test). Cluster 2 contained
157 mV2 and PPC, and on average showed monotonically decreasing effects of inactivation on the
158 weight of evidence as it gets more distal from laser onset ($p < 0.02$ for 0 and 50 cm, $p = 0.47$
159 for 100 cm). Thus, compared to the frontal area cluster, these posterior areas contributed to
160 evidence accumulation on shorter timescales. Finally, cluster 3 contained V1 and RSC, whose
161 inactivation led to non-monotonic changes in evidence weighting, affecting current and
162 long-past evidence ($p < 0.001$), but not evidence occurring in between ($p = 0.07$). This is
163 potentially compatible with findings that multiple timescales of processing can be present
164 within the same cortical regions (Bernacchia et al., 2011; Cavanagh et al., 2020; Scott et al.,

2017; Spitmaan et al., 2020; Wasmuht et al., 2018). Note that, for stimuli occurring while the laser is on, our analysis does not allow us to differentiate between pure visual processing deficits and deficits in evidence accumulation. However, this confound does not affect our main conclusions, since the areas also differ in terms of inactivation effects on evidence occurring prior to laser onset. Importantly, we have previously verified in an identical preparation that our laser parameters lead to robust inactivation and near-immediate recovery of pre-laser firing rates, with little to no rebound (Pinto et al., 2019). Thus, the effects observed here are unlikely to be related to changes in the average population activity levels outside of the nominal inactivation periods, or to different inactivation efficiencies between different epoch durations. Moreover, the effects were not due to increases in the timescale of the behavior itself leading to more forgetting, since we observed no significant laser-induced decreases in running speed (Figure 2–figure supplement 1).

Figure 2. Inactivating different cortical areas leads to evidence-accumulation deficits on distinct timescales.

(A) Illustration of the analysis method presented in panels B and C, using area aM2 as an example. Top six plots: effects of inactivating area aM2 during different epochs on evidence-weighting curves (laser off - laser on). Blue shading: laser-on epoch, error bars: $\pm$ SD across 10,000 bootstrapping iterations, data combined across mice. Bottom plot: Six top plots are aligned by laser onset (gray lines), and combined to include the first data point during 'laser on', and two preceding data points. See panel B for conventions. Error bars, $\pm$ SEM across experimental conditions. **(B)** Laser-onset-aligned changes in evidence-weighting curves for each bilaterally targeted area (laser on - laser off). Thin gray lines, individual inactivation epochs ($n = 6$ for $y = 0$, 4 for $y = 50$ and 3 for $y = 100$). Thick black lines, average across conditions. Error bars, $\pm$ SEM across experimental conditions. Shaded areas indicate laser on. Circles below the lines indicate statistical significance according to the caption on the bottom (one-sided paired t test vs. zero, corrected for multiple comparisons). P-values on top of each panel are from a one-way ANOVA with repeated measures with different y positions as factors. **(C)** Average laser-onset-aligned changes in evidence-weighting curves for each cluster (caption on top). Error bars, $\pm$ SEM across inactivation epochs concatenated for each cluster. Circles below the lines indicate statistical significance for the cluster of corresponding color, according to the caption in panel a (one-sided paired t test vs. zero, corrected for multiple comparisons).

195 Next, we performed a similar analysis to assess changes in evidence weights after laser
196 offset. Confirming our initial impression, the inactivation of most areas did not impact evidence
197 weighting prospectively (Figure 2–figure supplement 2). Interestingly, however, mV2 and mM2
198 inactivation led to moderately but significantly decreased use of evidence occurring 100 cm in
199 the future ($p < 0.05$, one-sided paired t test), perhaps suggesting an additional role for these
200 areas also in post-accumulation decision processes (Hanks et al., 2015). Thus, our inactivation
201 data suggest that the widespread cortical involvement in this task is largely related to the
202 accumulation of sensory evidence, and that different cortical areas accumulate on distinct
203 timescales. We next wondered whether evidence information from different areas is linearly
204 combined, at least from a behavioral standpoint. To do this, we compared the effects of
205 simultaneously inactivating all frontal or posterior areas to that expected by a linear
206 combination of the effects of inactivating areas individually (i.e. their average). Neither posterior
207 nor frontal areas significantly deviated from the linear prediction (Figure 2–figure supplement 3,
208 $p > 0.05$, two-way ANOVA with repeated measures, factors y position and inactivation type).
209 This suggests that signals from the different dorsal cortical areas could be combined by
210 downstream regions in a near-linear fashion. Candidate regions include the medial prefrontal
211 cortex, or subcortical structures such as the striatum and the cerebellum, which have been
212 shown to be causally involved in evidence accumulation (Deverett et al., 2019; Yartsev et al.,
213 2018). Other subcortical candidates are midbrain regions shown to have a high incidence of
214 choice signals in a contrast discrimination task (Steinmetz et al., 2019). A caveat here is that
215 the high variance in the data may have masked small non-linearities that could be revealed
216 with larger sample sizes. The possibility of a downstream integration of cortical signals is
217 agnostic to this limitation, however.

A hierarchy of timescales in large-scale cortical activity during evidence accumulation

Our inactivation results are reminiscent of the findings that cortical areas display a hierarchy of intrinsic timescales, such that primary sensory areas tend to integrate over shorter time windows than frontal and other association areas (Chaudhuri et al., 2015; Hasson et al., 2008; Murray et al., 2014; Runyan et al., 2017). While these are thought to arise in part from intrinsic cellular and circuit properties such as channel and receptor expression, amount of recurrent connectivity and relative proportions of inhibitory interneuron subtypes (Chaudhuri et al., 2015; Duarte et al., 2017; Fulcher et al., 2019; Gao et al., 2020; Wang, 2020), they appear to be modulated by task demands (Gao et al., 2020; Ito et al., 2020). Thus, to confirm whether this timescale hierarchy exists in the mouse cortex during performance of the accumulating-towers task, we reanalyzed previously published data consisting of mesoscale widefield Ca^{2+} imaging of the dorsal cortex through the intact cleared skull of mice expressing the Ca^{2+} indicator GCaMP6f in excitatory neurons (Figure 3A, Emx1-Ai93 triple transgenics, $n = 6$, 25 sessions)(Pinto et al., 2019). To do this, we enhanced our previous linear encoding model (or GLM) of the average activity of anatomically defined regions of interest (ROIs)(Pinto et al., 2019) by including two sets of predictors in addition to task events. First, for each ROI we added the zero-lag activity of other simultaneously imaged ROIs as coupling predictors, similar to previous work (Pillow et al., 2008; Runyan et al., 2017)(Figure 3-figure supplement 1). Crucially, we also included auto-regressive predictors to capture intrinsic activity auto-correlations that are unrelated to behavioral events. In other words, this approach allowed us to estimate within-task auto-correlations while separately accounting for task-induced temporal structure in cortical dynamics (Spitmaan et al., 2020). Adding these new sets of predictors resulted in a large and significant increase in cross-validated model accuracy, as measured by the linear correlation coefficient between the model predictions and a test dataset not used to fit the

242 model (Figure 3B, C; ~ 0.95 vs. ~ 0.3 , $F_{\text{model}(6,2,12)} = 1994.85$, $p = 6.2 \times 10^{-13}$, two-way ANOVA
243 with repeated measures). Note that these values are computed on held-out raw data points
244 rather than averaged activity. Thus, while the original model in our previous work had low
245 cross-validated accuracies in comparison, those values are compatible with other encoding
246 models of cortical activity in the literature that used similarly stringent goodness-of-fit metrics
247 (e.g. Huth et al., 2012; Pinto and Dan, 2015).

248 Motivated by our inactivation findings, we focused our analysis on the auto-regressive
249 coefficients of the model. We observed that across animals the rate of decay of these
250 coefficients over lags slowed systematically from visual to premotor areas, with intermediate
251 values for M1, PPC and RSC (Figure 3D). To quantify this, we fitted exponential decay
252 functions to the auto-regressive coefficients averaged across hemispheres (Figure 3D–F), and
253 extracted decay time constants (τ , Figure 3G). Compatible with our observations, τ differed
254 significantly across cortical areas ($F_{6,30} = 4.49$, $p = 0.002$, one-way ANOVA with repeated
255 measures), being larger for frontal than posterior areas, in particular PPC and mV2. Note that,
256 while it is possible that these coefficients capture auto-correlations introduced by intrinsic
257 GCaMP6f dynamics, there is no reason to believe that this affects our conclusions, as indicator
258 dynamics should be similar across regions. Thus, during the evidence-accumulation task,
259 cortical regions display increasing intrinsic timescales going from visual to frontal areas. This is
260 consistent with previous reports for spontaneous activity and other behavioral tasks (Chaudhuri
261 et al., 2015; Hasson et al., 2008; Murray et al., 2014; Runyan et al., 2017). Moreover, it is in
262 overall agreement with our inactivation findings (Figure 2), and suggests that the different
263 intrinsic timescales across the cortex support evidence integration over time windows of
264 different durations.

Figure 3. A hierarchy of activity timescales during evidence accumulation.

(A) Top: example widefield imaging field of view showing GCaMP6f fluorescence across the dorsal cortex. Bottom: approximate correspondence between the field of view and ROIs defined from the Allen Brain Atlas, ccv3. (B) Distribution of cross-validated accuracies across mice ($n = 6$, sessions for each mouse are averaged) and ROIs ($n = 7$, averaged across hemispheres). (C) Example of actual $\Delta F/F$ (gray) and GLM predictions (colored lines) for the first 5 s of the same held-out single trial, and four simultaneously imaged ROIs. Traces are convolved with a 1-SD gaussian kernel for display only. (D) Auto-regressive GLM coefficients as a function of time lags for an example imaging session and four example ROIs. Gray, coefficient values. Colored lines, best-fitting exponential decay functions. (E) Distribution of R^2 values for the exponential fits across mice ($n = 6$, sessions for each mouse are averaged) and ROIs ($n = 7$, averaged across hemispheres). (F) Exponential decay functions for all seven cortical areas, fitted to the average across mice ($n = 6$). (G) Time constants extracted from the exponential decay fits, for each area. Error bars, $\pm$ SEM across mice ($n = 6$). P-value is from a one-way ANOVA with repeated measures with ROIs as factors.

Discussion

Taken together, our results suggest that distributed cortical areas contribute to sensory-evidence accrual on different timescales. Specifically, brief sub-trial inactivations

282 during performance of a decision-making task requiring seconds-long evidence accumulation
 283 resulted in distinct deficits in the weighting of sensory evidence from different points in the
 284 stimulus stream. This was such that, on average, the inactivation of frontal cortical areas
 285 resulted in decreased use of evidence occurring further in the past from laser onset compared
 286 to a subset of posterior regions (Figure 2). Compatible with this, using an encoding model of
 287 large-scale cortical dynamics, we found that activity timescales vary systematically across the
 288 cortex in a way that mirrors the inactivation results (Figure 3).

289 Our results add to a growing body of literature that has revealed that the cortex of
 290 rodents and primates appears to be organized in a hierarchy of temporal processing windows
 291 across regions (Chaudhuri et al., 2015; Gao et al., 2020; Hasson et al., 2008; Ito et al., 2020;
 292 Murray et al., 2014; Runyan et al., 2017; Spitmaan et al., 2020). Specifically, to the best of our
 293 knowledge, they provide the first causal demonstration that the contributions of different
 294 cortical areas to decision-making computations appear similarly arranged in a temporal
 295 hierarchy. A caveat here is that our inactivation findings do not exactly match the smooth
 296 increases in integration windows going from posterior to frontal areas that we observed in our
 297 neural data. Rather, they appear to reflect a more modular organization, as suggested by our
 298 clustering results (see also Pinto et al., 2019), and one that does not exactly map onto the
 299 expected monotonic effects on accrual timescales. The latter could be due the fact that diverse
 300 timescales exist at the level of individual neurons within each region (Bernacchia et al., 2011;
 301 Cavanagh et al., 2020; Scott et al., 2017; Spitmaan et al., 2020; Wasmuht et al., 2018), and/or
 302 that, other decision-making processes beyond evidence accumulation are also affected by our
 303 inactivations. For instance, both mV2 and mM2 appeared to contribute to post-accrual
 304 decision processes (Figure 2–figure supplement 2). Nevertheless, our results point to accrual

timescale hierarchies being a significant factor explaining the large-scale functional organization of cortical dynamics during evidence-based decisions.

Our findings further suggest the possibility that the logic of widespread recruitment of cortical regions in complex, time-extended decisions may in part rely on intrinsic temporal integration properties of local cortical circuits, rather than specific evidence-accumulation mechanisms. For instance, it is possible that simple perceptual decisions primarily engage only the relevant sensory areas because they can be made on the fast intrinsic timescales displayed by these regions (Zatka-Haas et al., 2020). Along the same lines, it is conceivable that discrepancies in the literature regarding the effects of perturbing different cortical areas during evidence accumulation stem in part from differences in the timescales of the various tasks (Erlich et al., 2015; Fetsch et al., 2018; Hanks et al., 2015; Katz et al., 2016; Pinto et al., 2019).

An important remaining question is whether evidence from the different time windows is accumulated in parallel or as a feedforward computation going from areas with short to those with long integration time constants. The parallel scheme would be compatible with recent psychophysical findings in humans reporting confidence of their evidence-based decisions (Ganupuru et al., 2019). Conversely, a feedforward transformation would be in agreement with human fMRI findings during language processing (Yeshurun et al., 2017), and with a previously published model whereby successive (feedforward) convolution operations lead to progressively longer-lasting responses to sensory evidence (Scott et al., 2017). Interestingly, the oculomotor integrator of both fish and monkeys appears to be organized as largely feedforward chains of integration leading to systematically increasing time constants (Joshua and Lisberger, 2015; Miri et al., 2011), perhaps suggesting that this architecture is universal to neural integrators.

Much work remains before obtaining a complete circuit understanding of gradually evolving decisions. Our findings highlight the fact that, much like in memory systems (Jeneson and Squire, 2012), the timescale of decision processes is an important feature governing their underlying neural mechanisms, a notion which should be incorporated into both experimental and theoretical accounts of decision making.

Materials and Methods

Animals and surgery. All procedures were approved by the Institutional Animal Care and Use Committee at Princeton University and were performed in accordance with the Guide for the Care and Use of Laboratory Animals (National Research Council, 2011). We used both male and female VGAT-ChR2-EYFP mice aged 2 – 16 months [B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J, Jackson Laboratories, stock # 014548, n = 28]. Part of the inactivation data from some of these animals was collected in the context of previous work (Pinto et al., 2019), but the analyses reported here are completely novel. The mice underwent sterile surgery to implant a custom titanium headplate and optically clear their intact skulls, following a procedure described in detail elsewhere (Pinto et al., 2019). Briefly, after exposing the skull and removing the periosteum, successive layers of cyanoacrylate glue (krazy glue, Elmers, Columbus, OH) and diluted clear metabond (Parkell, Brentwood, NY) were applied evenly to the dorsal surface of the skull, and polished after curing using a dental polishing kit (Pearson dental, Sylmar, CA). The headplate was attached to the cleared skull using metabond, and a layer of transparent nail polish (Electron Microscopy Sciences, Hatfield, PA) was applied and allowed to cure for 10 – 15 min. The procedure was done under isoflurane

349 anesthesia (2.5% for induction, 1.5% for maintenance). The animals received two doses of
 350 meloxicam for analgesia (1 mg/kg I.P or S.C.), given at the time of surgery and 24 h later, as
 351 well as peri-operative I.P. injections of body-temperature saline to maintain hydration. Body
 352 temperature was maintained constant using a homeothermic control system (Harvard
 353 Apparatus, Holliston, MA). The mice were allowed to recover for at least 5 days before starting
 354 behavioral training. After recovery they were restricted to 1 – 2 mL of water per day and
 355 extensively handled for another 5 days, or until they no longer showed signs of stress. We
 356 started behavioral training after their weights were stable and they accepted handling. During
 357 training, the full allotted fluid volume was typically delivered within the behavioral session, but
 358 supplemented if necessary. The mice were weighed and monitored daily for signs of
 359 dehydration. If these were present or their body mass fell below 80% of the initial value, they
 360 received supplemental water until recovering. They were group housed throughout the
 361 experiment, and had daily access to an enriched environment (Pinto et al., 2018). The animals
 362 were trained 5 – 7 days/week.

363 The analysis reported in Figure 3 (widefield Ca²⁺ imaging) is from data collected in the
 364 context of a previous study (Pinto et al., 2019), although the analysis is novel. The data was
 365 from 6 male and female mice from triple transgenic crosses expressing GCaMP6f under the
 366 CaMKII α promoter from the following two lines: Ai93-D;CaMKII α -tTA [IgS6^{tm93.1(tetO-GCaMP6f)Hze}
 367 Tg(Camk2a-tTA)1Mmay/J, Jackson Laboratories, stock # 024108] and Emx1-IRES-Cre
 368 [B6.129S2-Emx1^{tm1(cre)Krl}/J, Jackson Laboratories, stock # 005628]. These animals also
 369 underwent the surgical procedure described above.

370 **Virtual reality apparatus.** The mice were trained in a virtual reality (VR) environment (Figure 1A)
 371 described in detail elsewhere (Pinto et al., 2018). Briefly, they sat on an 8-inch hollow

372 Styrofoam® ball that was suspended by compressed air at ~60 p.s.i, after passing through a
 373 laminar flow nozzle to reduce noise (600.326.5K.BC, Lechler, St. Charles, IL). They were
 374 head-fixed such that their snouts were aligned to the ball equator and at a height such that
 375 they could run comfortable without hunching, while still being able to touch the ball with their
 376 full paw pads (corresponding to a headplate-to-bal height of ~1 inch for a 25-g animal). Ball
 377 movements were measured using optical flow sensors (ADNS-3080 APM2.6) and transformed
 378 into virtual world displacements using custom code running on Arduino Due
 379 (<https://github.com/sakoay/AccumTowersTools/tree/master/OpticalSensorPackage>). The ball
 380 sat on a custom 3D-printed cup that contained both the air outlet and the movement sensor.
 381 The VR environment was projected onto a custom-built toroidal Styrofoam® screen using a
 382 DLP projector (Optoma HD141X, Fremont, CA) at a refresh rate of 120 Hz and a pixel resolution
 383 of 1024 x 768. The screen spanned ~270° of azimuth and ~80° of elevation in the mouse's
 384 visual field. The whole set-up was enclosed in a custom-built sound-attenuating chamber. The
 385 VR environment was programmed and controlled using ViRMEn (Aronov and Tank, 2014)
 386 (<https://pni.princeton.edu/pni-software-tools/virmen>), running on Matlab (Mathworks, Natick,
 387 MA) on a PC.

388 **Behavioral task.** We trained the mice in the accumulating-towers task (Pinto et al., 2018). The
 389 mice ran down a virtual T-maze that was 3.3 m in length (y), 5 cm in height and a nominal 10
 390 cm in width (x, though they were restricted to the central 1 cm). The length of the maze
 391 consisted of a 30-cm start region to which they were teleported at the start of each trial,
 392 followed by a 200-cm cue region and a 100-cm delay region. The cue and the delay region had
 393 the same wallpaper designed to provide optical flow. During the cue region, the mice
 394 encountered tall white objects (2 x 6 cm, width x height), or towers, that appeared at random

locations in each trial at a Poisson rate of 7.7 m^{-1} and 2.3 m^{-1} on the rewarded on non-rewarded side, respectively (or 8.0 and 1.6 m^{-1} in some sessions), with a 12-cm refractory period and an overall density of 5 m^{-1} . The towers appeared when the mice were 10 cm away from their drawn locations, and disappeared 200 ms later (roughly corresponding to the time over which the tower sweeps across the visual field given average running speeds). After the maze stem the mice turned into one of the two arms ($10.5 \times 11 \times 5 \text{ cm}$, length x width x height), and received a reward if they turned to the arm corresponding to the highest tower count (4–8 μL of 10% v/v sweet condensed milk). This was followed by a 3-s inter-trial interval, consisting of 1 s of a frozen frame of the VR environment and 2 s of a black screen. An erroneous turn resulted in a loud sound and a 12-s timeout.

Each daily behavioral session (~1 h, ~200 – 250 trials) started with warm-up trials of a visually guided task in the same maze, in which towers appeared only on the rewarded side and additionally a 30-cm tall visual guide visible from the start of the trial was placed in the arm corresponding to the reward location. The animals progressed to the main task when they achieved at least 85% correct trials over a running window of 10 trials in the warm-up task. During the accumulating-towers task, performance was evaluated over a 40-trial running window, both to assess side biases and correct them using an algorithm described elsewhere (Pinto et al., 2018), and to trigger a transition into a 10-trial block of easy trials if performance fell below 55% correct. These blocks consisted of towers only on the rewarded side, and were introduced to increase motivation but were not included in the analyses. No optogenetic inactivation was performed during either warm-up or easy-block trials. In the widefield imaging experiments, the behavioral sessions contained several visually guided (warm up) blocks (Pinto et al., 2019). These were excluded from the present analyses.

Laser-scanning optogenetic inactivation. We used a scanning laser setup described in detail elsewhere (Pinto et al., 2019). Briefly, a 473-nm laser beam (OBIS, Coherent, Santa Clara, CA) was directed to 2-D galvanometers using a 125- μ m single-mode optic fiber optic (Thorlabs, Newton, NJ) and reached the cortical surface after passing through an f-theta scanning lens (LINOS, Waltham, MA). We used a 40-Hz square wave with an 80% duty cycle and a power of 6 mW measured at the level of the skull. This corresponds to an inactivation spread of $\sim 1.5 - 2$ mm (Pinto et al., 2019). While this may introduce confounds regarding ascribing exact functions to specific cortical areas, we have previously shown that the effects of whole-trial inactivations at much lower powers (corresponding to smaller spatial spreads) are consistent with those obtained at 6 mW. To minimize post-inactivation rebounds, the last 100 ms of the laser pulse consisted of a linear ramp-down of power (Guo et al., 2014; Pinto et al., 2019). We performed inactivations during the following trial epochs: 1st, 2nd or 3rd quarter of the cue region (0 – 50 cm, 50 – 100 cm or 100 – 150 cm, respectively), 1st or 2nd half of the cue region (0 – 100 cm or 100 – 200 cm, respectively), or delay region (200 – 300 cm). Thus, the epochs were defined according to the animals' y position in the maze. Because of this, the onset time of the power ramp-down was calculated in each trial based on the current speed and the expected time at which the mouse would reach the laser offset location. The system was controlled using custom-written code in Matlab running on a PC, which sent command analog voltages to the laser and galvanometers through NI DAQ cards. This PC received instructions for laser onset, offset and galvanometer position from the ViRMEn PC through digital lines.

We targeted a total of 9 area combinations, either consisting of homotopic bilateral pairs or multiple bilateral locations. The galvanometers alternated between locations at 200 Hz (20-mm travel time: ~ 250 μ s) and, in the case of more than 2 locations, the sequence of visited

locations was chosen to minimize travel distance. The inactivated locations were defined based on stereotaxic coordinates using bregma as reference, as follows:

- Primary visual cortex (V1): -3.5 AP, 3 ML
- Medial secondary visual cortex (mV2, ~ area AM): -2.5 AP, 2.5 ML
- Posterior parietal cortex (PPC): -2 AP, 1.75 ML
- Retrosplenial cortex (RSC): -2.5 AP, 0.5 ML
- Posteromedial portion of the premotor cortex (mM2): 0.0 AP, 0.5 ML
- Anterior portion of the premotor cortex (aM2): +3 AP, 1 ML
- Primary motor cortex (M1): +1 AP, 2 ML
- Posterior cortex: V1, mV2, PPC and RSC
- Frontal cortex: mM2, aM2 and M1

To ensure consistency in bregma location across behavioral sessions, the experimenter set bregma on a reference image and for each session the current image of the mouse's skull was registered to this reference using rigid transformations. Different sessions contained different combinations of areas and inactivation epochs, resulting in partially overlapping mice and sessions for each condition. The probability of inactivation trials therefore varied across sessions, ranging from a total of 0.15 – 0.35 across conditions, and from 0.02 – 0.15 per condition. In our experience, capping the probability at ~0.35 is important to maintain motivation throughout the behavioral session.

Widefield Ca²⁺ imaging. Details on the experimental setup and data preprocessing can be found elsewhere (Pinto et al., 2019). Briefly, we used a tandem-lens microscope (1x – 0.63x planapo, Leica M series, Wetzlar, Germany) with alternating 410-nm and 470-nm LED epifluorescence illumination for isosbestic hemodynamic correction, and collected 525-nm

emission at 20 Hz, using an sCMOS (OrcaFlash4.0, Hamamatsu, Hamamatsu City, Japan), with an image size of 512 x 512 pixels (pixel size of ~17 μm). Images were acquired with HCLImage (Hamamatsu) running on a PC, and synchronized to the behavior using a data acquisition-triggering TTL pulse from another PC running ViRMEn, which in turn received analog frame exposure voltage traces acquired through a DAQ card (National Instruments, Austin, TX) and saved in the behavioral log file. The image stacks were motion-corrected by applying the x-y shift that maximized the correlation between successive frames, and then were spatially binned to a 128 x 128 pixel image (~68 x 68 μm). The fluorescence values from pixels belonging to different anatomical ROIs were averaged into a single trace, separately for 410-nm (F_v) and 470-nm excitation (F_b). After applying a heuristic correction to F_v (Pinto et al., 2019), we calculated fractional fluorescence changes as $R = F/F_0$, where F_0 for each excitation wavelength was calculated as the mode of all F values over a 30-s sliding window with single-frame steps. The final $\Delta F/F$ was calculated using a divisive correction, $\Delta F/F = R_b / R_v - 1$. ROIs were defined based on the Allen Brain Mouse Atlas (ccv3). We first performed retinotopic mapping to define visual areas, and used the obtained maps to find, for each mouse, the optimal affine transformation to the Allen framework.

Data analysis. All analyses of the behavioral effects of cortical inactivations were performed in Python 3.7. Generalized linear model (GLM) fitting of widefield data was performed in Matlab, and the results were analyzed in Python.

Behavioral data selection. Because of the warm-up and easy-block trials, the sessions are naturally organized into a block structure, such that the duration of each block of the accumulating-towers task is of at least 40 trials (see above). We selected all trials from blocks

in which the control (laser off) performance was at least 60% correct, collapsed over all levels of sensory evidence. After block selection, we excluded trials in which the animals failed to reach the end of the maze, or in which the total traveled distance exceeded the nominal maze length by more than 10% (Pinto et al., 2018, 2019). Additionally, because we were interested in assessing the effects of inactivation on accumulation timescales, we excluded animals that failed to use evidence from all quarters of the cue region of the maze to make their decisions in control trials. To do this, we fitted the logistic regression model (see below) separately for each animal, bootstrapping by sampling trials with replacement 200 times. We then computed the significance at each y position bin as the fraction of trials in which the model coefficient was equal to or greater than zero. Mice with any coefficients not significantly different than zero after false discovery rate correction (see below) were excluded from further analyses. These selection criteria yielded a total of 855 optogenetic inactivation sessions from 20 mice (average ~43/mouse), corresponding to 100,787 control (laser off) trials, and 27,606 inactivation trials (average ~511/condition, see Figure 1–table supplement 1). Twenty-five sessions from six mice were selected for widefield imaging data analysis.

Analysis of behavioral data. *Overall performance.* We calculated overall performance as the percentage of trials in which the mice turned to side with the highest tower counts, separately for control and inactivation trials.

Running speed. Speed was calculated for each inactivation segment using the total x-y displacement. We compared laser-induced changes in speed to control trials from the same maze segment.

Psychometric curves. We computed psychometric curves separately for control and inactivation trials by plotting the percentage of right-choice trials as a function of the difference

509 in the number of right and left towers (#R – #L, or Δ). Δ was binned in increments of 5 between
 510 -15 and 15, and its value defined as the average Δ weighted by the number of trials. We fitted
 511 the psychometric curves using a 4-parameter sigmoid:

$$p_R = b + \frac{a}{1 + \exp(-(\Delta - \Delta_0)/\lambda)}$$

512 *Evidence-weighting curves.* To assess how mice weighted sensory evidence from
 513 different segments of the cue region, we performed a logistic regression analysis in which the
 514 probability of a right choice was predicted from a logistic function of the weighted sum of the
 515 net amount of sensory evidence from each of 4 equally-spaced segments (10 – 200 cm, since
 516 no towers can occur before $y = 10$):

$$p_R = \frac{1}{1 + \exp(-(\beta_0 + \sum_{i=1}^4 \beta_i \Delta_i))}$$

517 where $\Delta = \# \text{ right} - \# \text{ left}$ towers calculated separately for each segment. These weighting
 518 functions were calculated separately for 'laser on' and 'laser off' trials. To quantify the
 519 laser-induced changes in evidence weighting, we simply subtracted the 'laser on' from the
 520 'laser off' curves, such that negative values indicate smaller evidence weights in the 'laser on'
 521 condition. Bin sizes were chosen to match the resolution of our inactivation epochs.

522 *Laser-triggered analysis.* For each area, we aligned the evidence-weighting curves by
 523 the position of laser onset, which was defined as $y = 0$, and used y position bins going up to
 524 100 cm in the past. Thus, each inactivation condition contributed up to 3 bins of data,
 525 depending on the position of laser onset. For inactivations lasting more than one position bin
 526 (i.e. 100 cm), we used only the first bin during the inactivation as the $y = 0$ datapoint. Given the

laser onset positions in our experiments, $y = 0$ had six data points, $y = -50$ had four, and $y = -100$ had three. For the laser-offset triggered analysis (Figure 2–figure supplement 2), we aligned the evidence-weighting curves by the first bin following laser offset. Given the laser offset positions in our data, we analyzed $y = 50$ ($n = 4$ conditions) and $y = 100$ ($n = 3$).

Statistics of inactivation effects. Error estimates and statistics for general performance, running speed and logistic regression weights were generated by bootstrapping this procedure 10,000 times, where in each iteration we sampled trials with replacement. P-values were calculated as the fraction of bootstrapping iterations in which the control-subtracted inactivation value was above zero. In other words, we performed a one-sided test of the hypothesis that inactivation decreases performance, speed and evidence weights on decision. To analyze the laser-onset triggered curves (Figure 2), we used a one-way ANOVA with repeated measures with the y position bin as a factor to establish the significance of the difference in the effects across bins. To account for the different number of datapoints per spatial bin, we implemented this as a mixed model with experimental conditions as the random effect. To assess whether laser effects were significant for each bin in the laser-onset (or offset)-aligned curves, we performed a one-sided t test against zero, with inactivation epochs as data points. Finally, to compare the effects of simultaneous and individual area inactivations (Figure 2–figure supplement 3), we performed a two-way ANOVA with repeated measures with factors y position bin and inactivation type, using the individual inactivation epochs as data points (in the case of individual inactivations, epochs from different areas were concatenated).

Clustering of evidence-weighting curves. We generated a 7×3 (areas $\times$ laser-onset-triggered inactivation bins) matrix containing the average laser-subtracted evidence-weighting curves, aligned by laser onset, for each individually targeted area. Thus, we excluded the experimental conditions in which frontal or posterior cortical areas were

551 inactivated simultaneously from this analysis. We then performed spectral clustering into k
 552 clusters on that matrix. We tested $k = 2 - 5$, and chose the value of k that maximized clustering
 553 quality as measured by the Calinski-Harabasz index. Given the small number of areas per
 554 cluster, we generated error estimates for each y position bin by concatenating the individual
 555 inactivation conditions (epochs) for all areas of the cluster.

556 **Generalized linear model (GLM) of widefield data.** We fitted Ca^{2+} activity averaged over each
 557 anatomically defined ROI with a generalized linear model (GLM)(Pinto et al., 2019; Pinto and
 558 Dan, 2015; Scott et al., 2017). For each trial and y position in the maze, we extracted $\Delta F/F$
 559 (with native 10-Hz sampling frequency) limited to $0 \leq y \leq 300$ cm (i.e. trial start, outcome and
 560 inter-trial periods were not included). Activity was then z-scored across all trials. $\Delta F/F$ of each
 561 area was modeled as a linear combination of different predictors at different time lags. In
 562 addition to the previously used task-event predictors (Pinto et al., 2019), we added coupling
 563 terms i.e. the zero-lag activity of the other simultaneously imaged ROIs (Pillow et al., 2008;
 564 Runyan et al., 2017), as well as auto-regressive terms to capture activity auto-correlations that
 565 were independent of task events (Spitmaan et al., 2020). Finally, we added a term to penalize
 566 the L2 norm of the coefficients, i.e. we performed ridge regression. The full model was thus
 567 defined as:

$$\Delta F/F(t) = \beta_0 + A + C + T + \lambda ||\vec{B}||$$

568 where β_0 is an offset term, λ is the penalty term and $||\vec{B}||$ is the L2 norm of the weight vector.
 569 Additionally, A , C and T are the auto-regressive, coupling and task terms, respectively:

$$A = \sum_{i=0.1}^2 \beta_i^{autoregr} \Delta F / F(t-i)$$

$$C = \sum_{j=1}^{15} \beta_j^{coupling} \Delta F / F_j(t)$$

$$T = \sum_{i=0}^2 \beta_i^{tR} E_{t-i}^{tR} + \sum_{i=0}^2 \beta_i^{tL} E_{t-i}^{tL} + \sum_{i=0}^2 \beta_i^{\Delta} E_{t-i}^{\Delta} + \sum_{i=-0.3}^{0.3} \beta_i^{\theta} E_{t-i}^{\theta} + \sum_{i=-0.3}^{0.3} \beta_i^{d\theta/dt} E_{t-i}^{d\theta/dt} + \sum_{i=-0.3}^{0.3} \beta_i^{sp} E_{t-i}^{sp} + \beta^y y + \beta^{ch} ch + \beta^{pch} pch + \beta^{prw} prw$$

570 In the above equations, β_i^x is the encoding weight for predictor x at time lag i (in steps of 0.1
571 s), where x is either a task event or the activity of the ROI at a previous time point, and $\beta_j^{coupling}$
572 is the weight for the zero-lag activity for simultaneously imaged ROI j (we had a total of 16 ROIs
573 across the two hemispheres). In the task term, E_{t-i}^x is a delta function indicating the
574 occurrence of event x at time $t-i$. Specifically, tR indicates the occurrence of a right tower, tL of
575 a left tower, Δ = cumulative #R – #L towers, θ is view angle, $d\theta/dt$ is virtual view angle velocity,
576 sp is running speed, y is spatial position in the maze stem (no lags), and ch , pch and prw are
577 constant offsets for a given trial, indicating upcoming choice, previous choice (+1 for right and
578 –1 for left) and previous reward (1 for reward and –1 otherwise), respectively.

579 *Cross-validation.* The model was fitted using 3-fold cross-validation. For each of 20
580 values of the penalty term λ , we trained the model using $\frac{2}{3}$ of the trials (both correct and wrong
581 choices), and tested it on the remaining $\frac{1}{3}$ of trials. We picked the value of λ that maximized
582 accuracy, and used median accuracy and weight values across all 10 x 3 runs for that λ . Model

accuracy was defined as the linear correlation coefficient between actual $\Delta F/F$ and that predicted by the model in the test set.

Model comparison. We tested three versions of the GLM, one with just the task term T , another one adding the auto-regressive term A , and the other with the coupling term C in addition to A and T . All versions were fitted using exactly the same cross-validation data partitioning to allow for direct comparison. We averaged cross-validated predictions over hemispheres and sessions for each mouse, performing the comparison with mouse-level data. Statistical significance of the differences between the accuracy of different models was computed using a two-way ANOVA with repeated measures with factors ROI and model type, and individual model comparisons were made using Tukey's post-hoc test. Coefficient analysis in Figure 3 is from the full model, which had the highest performance.

Quantification of timescales from the GLM. To quantify the timescales from the fitted auto-regressive coefficients, for each behavioral session we fitted an exponential decay function to the coefficients between 0.1 and 2 s in the past, normalized to the coefficient at 0.1 s (first bin):

$$B + A\exp(-x/\tau)$$

where B is the offset term, A controls the amplitude of the curve, x is the vector of normalized coefficients and τ is the decay time constant. Fits were performed using the non-linear least squares algorithm. The extracted time constants (τ) were first averaged over hemispheres and sessions for each mouse, and statistics were performed on mouse averages. Significance of

the differences in the time constants across regions was assessed by performing a one-way ANOVA with repeated measures, with cortical regions as the factor.

False discovery rate correction. We corrected for multiple comparisons using a previously described method for false discovery rate (FDR) correction (Benjamini and Hochberg, 1995; Guo et al., 2014; Pinto et al., 2019). Briefly, p-values were ranked in ascending order, and the i th ranked p-value, P_i , was deemed significant if it satisfied $P_i \leq (\alpha)/n$, where n is the number of comparisons and α is the significance level. In our case, $\alpha = 0.05$ because we defined all tests as one sided.

Data and code availability

Data analysis code and source code for figures is available at https://github.com/BrainCOGS/PintoEtAl2020_subtrial_inact.git. Behavioral data from inactivation experiments and GLM summary data will be deposited on a public repository upon peer-reviewed publication of this manuscript.

Acknowledgements

We thank Sue Ann Koay and Kanaka Rajan for discussions, Abigail Russo and E. Mika Diamanti for comments on the manuscript, and Samantha Stein and Scott Baptista for technical assistance. This work was supported by the NIH grants U01NS090541, U19NS104648, F32NS101871 (L.P.) and K99MH120047 (L.P.).

620 **Competing interests**

621 The authors declare no competing interests.

622 **Author contributions**

623 L.P. performed the experiments and analyzed the data; L.P. wrote the manuscript with input
624 from C.D.B. and D.W.T.; L.P., C.D.B. and D.W.T. conceived the project.

625 **References**

- 626 Aronov D, Tank DW. 2014. Engagement of Neural Circuits Underlying 2D Spatial Navigation in a
627 Rodent Virtual Reality System. *Neuron* **84**:442–456.
- 628 Benjamini Y, Hochberg Y. 1995. Controlling the false discovery rate: a practical and powerful
629 approach to multiple testing. *J R Stat Soc Series B Stat Methodol.* **57**:289–300
- 630 Bernacchia A, Seo H, Lee D, Wang X-J. 2011. A reservoir of time constants for memory traces
631 in cortical neurons. *Nat Neurosci* **14**:366–372.
- 632 Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006. The physics of optimal decision
633 making: A formal analysis of models of performance in two-alternative forced-choice tasks.
634 *Psychol Rev* **113**:700–765.
- 635 Brincat SL, Siegel M, von Nicolai C, Miller EK. 2018. Gradual progression from sensory to
636 task-related processing in cerebral cortex. *Proc. Natl. Acad. Sci. U. S. A.* **115**:E7202-E7211
- 637 Brody CD, Hanks TD. 2016. Neural underpinnings of the evidence accumulator. *Curr Opin*
638 *Neurobiol* **37**:149–157.
- 639 Brunton BW, Botvinick MM, Brody CD. 2013. Rats and Humans Can Optimally Accumulate
640 Evidence for Decision-Making. *Science* **340**:95–98.

641 Carandini M, Churchland AK. 2013. Probing perceptual decisions in rodents. *Nat Neurosci*
642 **16**:824–831.

643 Cavanagh SE, Hunt LT, Kennerley SW. 2020. A Diversity of Intrinsic Timescales Underlie Neural
644 Computations. *Front Neural Circuits* **14**:615626.

645 Chaudhuri R, Knoblauch K, Gariel M-A, Kennedy H, Wang X-J. 2015. A Large-Scale Circuit
646 Mechanism for Hierarchical Dynamical Processing in the Primate Cortex. *Neuron*
647 **21**:419–431.

648 Devereett B, Kislin M, Tank DW, -H. Wang SS. 2019. Cerebellar disruption impairs working
649 memory during evidence accumulation. **10**:3128.

650 Ding L, Gold JI. 2010. Caudate encodes multiple computations for perceptual decisions. *J*
651 *Neurosci* **30**:15747–15759.

652 Duarte R, Seeholzer A, Zilles K, Morrison A. 2017. Synaptic patterning and the timescales of
653 cortical dynamics. *Curr Opin Neurobiol* **43**:156–165.

654 Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD. 2015. Distinct effects of prefrontal and
655 parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* **4**:e05457.

656 Fetsch CR, Odean NN, Jeurissen D, El-Shamayleh Y, Horwitz GD, Shadlen MN. 2018. Focal
657 optogenetic suppression in macaque area MT biases direction discrimination and decision
658 confidence, but only transiently. *Elife* **7**:e36523.

659 Fulcher BD, Murray JD, Zerbi V, Wang X-J. 2019. Multimodal gradients across mouse cortex.
660 *Proc Natl Acad Sci U S A* **116**:4689–4695.

661 Ganupuru P, Goldring AB, Harun R, Hanks TD. 2019. Flexibility of Timescales of Evidence
662 Evaluation for Decision Making. *Curr Biol* **29**:2091–2097.e4.

663 Gao R, van den Brink RL, Pfeffer T, Voytek B. 2020. Neuronal timescales are functionally
664 dynamic and shaped by cortical microarchitecture. **9**:61277.

665 Gold JI, Shadlen MN. 2007. The neural basis of decision making. *Annu Rev Neurosci*
666 **30**:535–574.

667 Guo ZV, Li N, Huber D, Ophir E, Gutnisky D, Ting JT, Feng G, Svoboda K. 2014. Flow of
668 Cortical Activity Underlying a Tactile Decision in Mice. *Neuron* **81**:179–194.

669 Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD. 2015. Distinct relationships
670 of parietal and prefrontal cortices to evidence accumulation. *Nature* **520**:220–223.

671 Hasson U, Yang E, Vallines I, Heeger DJ, Rubin N. 2008. A hierarchy of temporal receptive
672 windows in human cortex. *J Neurosci* **28**:2539–2550.

673 Horwitz GD, Newsome WT. 1999. Separate signals for target selection and movement
674 specification in the superior colliculus. *Science* **284**:1158–1161.

675 Huth AG, Nishimoto S, Vu AT, Gallant JL. 2012. A continuous semantic space describes the
676 representation of thousands of object and action categories across the human brain.
677 *Neuron* **76**:1210–1224.

678 Ito T, Hearne LJ, Cole MW. 2020. A cortical hierarchy of localized and distributed processes
679 revealed via dissociation of task activations, connectivity changes, and intrinsic timescales.
680 *NeuroImage* **221**:117141.

681 Jeneson A, Squire LR. 2012. Working memory, long-term memory, and medial temporal lobe
682 function. *Learn Mem* **19**:15–25.

683 Joshua M, Lisberger SG. 2015. A tale of two species: Neural integration in zebrafish and
684 monkeys. *Neuroscience* **296**:80–91.

685 Katz LN, Yates JL, Pillow JW, Huk AC. 2016. Dissociated functional significance of
686 decision-related activity in the primate dorsal stream. *Nature* **535**:285–288.

687 Kiebel SJ, Daunizeau J, Friston KJ. 2008. A hierarchy of time-scales and the brain. *PLoS*
688 *Comput Biol* **4**:e1000209.

689 Kim JN, Shadlen MN. 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex
690 of the macaque. *Nat Neurosci* **2**:176–185.

691 Koay SA, Thiberge SY, Brody C, Tank DW. 2020. Amplitude modulation of cortical sensory
692 responses in pulsatile evidence accumulation. *Elife* **9**:e60628.

693 Krueger PM, van Vugt MK, Simen P, Nystrom L, Holmes P, Cohen JD. 2017. Evidence
694 accumulation detected in BOLD signal using slow perceptual decision making. *J Neurosci*
695 *Methods* **281**:21–32.

696 Miri A, Daie K, Arrenberg AB, Baier H, Aksay E, Tank DW. 2011. Spatial gradients and
697 multidimensional dynamics in a neural integrator circuit. *Nat Neurosci* **14**:1150–1159.

698 Morcos AS, Harvey CD. 2016. History-dependent variability in population dynamics during
699 evidence accumulation in cortex. *Nat Neurosci.* **19**:1672–1681

700 Murphy PR, Wilming N, Hernandez-Bocanegra DC. 2020. Normative Circuit Dynamics Across
701 Human Cortex During Evidence Accumulation in Changing Environments. *BioRxiv*
702 <https://doi.org/10.1101/2020.01.29.924795>.

703 Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C,
704 Pasternak T, Seo H, Lee D, Wang X-J. 2014. A hierarchy of intrinsic timescales across
705 primate cortex. *Nat Neurosci* **17**:1661–1663.

706 Newsome WT, Britten KH, Movshon JA. 1989. Neuronal correlates of a perceptual decision.
707 *Nature* **341**:52–54.

708 Odoemene O, Pisupati S, Nguyen H, Churchland AK. 2018. Visual Evidence Accumulation
709 Guides Decision-Making in Unrestrained Mice. *J Neurosci* **38**:10143–10155.

710 Orsolic I, Rio M, Masic-Flogel TD, Znamenskiy P. 2021. Mesoscale cortical dynamics reflect the
711 interaction of sensory evidence and temporal expectation during perceptual
712 decision-making. *Neuron* **in press**.

713 Pillow JW, Shlens J, Paninski L, Sher A, Litke AM, Chichilnisky EJ, Simoncelli EP. 2008.
714 Spatio-temporal correlations and visual signalling in a complete neuronal population.
715 *Nature* **454**:995–999.

716 Pinto L, Dan Y. 2015. Cell-Type-Specific Activity in Prefrontal Cortex during Goal-Directed
717 Behavior. *Neuron* **87**:437–450.

718 Pinto L, Koay SA, Engelhard B, Yoon AM, Deverett B, Thiberge SY, Witten IB, Tank DW, Brody
719 CD. 2018. An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in
720 Virtual Reality. *Front Behav Neurosci* **12**:36.

721 Pinto L, Rajan K, DePasquale B, Thiberge SY, Tank DW, Brody CD. 2019. Task-Dependent
722 Changes in the Large-Scale Dynamics and Necessity of Cortical Regions. *Neuron*
723 **104**:810–824.e9.

724 Runyan CA, Piasini E, Panzeri S, Harvey CD. 2017. Distinct timescales of population coding
725 across cortex. *Nature* **548**:92–96.

726 Scott BB, Constantinople CM, Akrami A, Hanks TD, Brody CD, Tank DW. 2017. Fronto-parietal
727 Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron*
728 **95**:385–398.e5.

729 Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex
730 (area LIP) of the rhesus monkey. *J Neurophysiol* **86**:1916–1936.

731 Soltani A, Murray JD, Seo H, Lee D. 2021. Timescales of cognition in the brain. *Curr Opin*
732 *Behav Sci* **41**:30–37.

733 Spitmaan M, Seo H, Lee D, Soltani A. 2020. Multiple timescales of neural dynamics and
734 integration of task-relevant signals across cortex. *Proc Natl Acad Sci USA*
735 **117**:22522–22531.

736 Steinmetz NA, Zatka-Haas P, Carandini M, Harris KD. 2019. Distributed coding of choice,

737 action and engagement across the mouse brain. *Nature* **576**:266–273.

738 Stine GM, Zylberberg A, Ditterich J, Shadlen MN. 2020. Differentiating between integration and

739 non-integration strategies in perceptual decision making. *Elife* **9**:e55365.

740 Sun P, Landy MS. 2016. A Two-Stage Process Model of Sensory Discrimination: An Alternative

741 to Drift-Diffusion. *J Neurosci* **36**:11259–11274.

742 Tsetsos K, Gao J, McClelland JL, Usher M. 2012. Using Time-Varying Evidence to Test Models

743 of Decision Dynamics: Bounded Diffusion vs. the Leaky Competing Accumulator Model.

744 *Front Neurosci* **6**:79.

745 Wang X-J. 2020. Macroscopic gradients of synaptic excitation and inhibition in the neocortex.

746 *Nat Rev Neurosci* **21**:169–178.

747 Waskom ML, Kiani R. 2018. Decision Making through Integration of Sensory Evidence at

748 Prolonged Timescales. *Curr Biol* **28**:3850–3856.e9.

749 Wasmuht DF, Spaak E, Buschman TJ, Miller EK, Stokes MG. 2018. Intrinsic neuronal dynamics

750 predict distinct functional roles during working memory. *Nat Commun* **9**:3499.

751 Wilming N, Murphy PR, Meyniel F, Donner TH. 2020. Large-scale Dynamics of Perceptual

752 Decision Information across Human Cortex. *Nat Commun* **11**:5109.

753 Yartsev MM, Hanks TD, Yoon AM, Brody CD. 2018. Causal contribution and dynamical

754 encoding in the striatum during evidence accumulation. *Elife* **7**:e35929

755 Yeshurun Y, Nguyen M, Hasson U. 2017. Amplification of local changes along the timescale

756 processing hierarchy. *Proc Natl Acad Sci USA* **114**:9475–9480.

757 Zatka-Haas P, Steinmetz NA, Carandini M, Harris KD. 2020. A perceptual decision requires

758 sensory but not action coding in mouse cortex. *BioRxiv* doi:10.1101/501627

759 Table and Figure Supplements

Region	Epoch	Num. control trials	Num. laser trials	Num. mice	Num. sessions	Total trial count
V1	cue, 1 st quarter (0 – 50 cm)	21,741	692	7	189	22,433
V1	cue, 2 nd quarter (50 – 100 cm)	21,646	686	7	189	22,332
V1	cue, 1 st half (0 – 100 cm)	19,753	540	7	158	20,293
V1	cue, 3 rd quarter (100 – 150 cm)	21,580	680	7	185	22,260
V1	cue, 2 nd half (100 – 200 cm)	20,492	585	7	163	21,077
V1	delay (200 - 300 cm)	20,013	596	8	162	20,609
mV2	cue, 1 st quarter (0 – 50 cm)	18,372	589	7	159	18,961
mV2	cue, 2 nd quarter (50 – 100 cm)	18,495	577	7	162	19,072
mV2	cue, 1 st half (0 – 100 cm)	14,446	367	3	121	14,813
mV2	cue, 3 rd quarter (100 – 150 cm)	18,429	594	7	160	19,023
mV2	cue, 2 nd half (100 – 200 cm)	14,449	353	3	119	14,802
mV2	delay (200 - 300 cm)	15,058	366	5	124	15,424
PPC	cue, 1 st quarter (0 – 50 cm)	15,929	403	6	133	16,332
PPC	cue, 2 nd quarter (50 – 100 cm)	15,993	429	6	134	16,422
PPC	cue, 1 st half (0 – 100 cm)	9,006	615	9	63	9,621
PPC	cue, 3 rd quarter (100 – 150 cm)	15,950	419	6	134	16,369
PPC	cue, 2 nd half (100 – 200 cm)	7,505	266	5	52	7,771
PPC	delay (200 - 300 cm)	8,947	616	6	67	9,563
RSC	cue, 1 st quarter (0 – 50 cm)	15,923	422	6	134	16,345
RSC	cue, 2 nd quarter (50 – 100 cm)	16,064	418	6	136	16,482
RSC	cue, 1 st half (0 – 100 cm)	11,736	925	10	86	12,661
RSC	cue, 3 rd quarter (100 – 150 cm)	15,952	432	6	134	16,384
RSC	cue, 2 nd half (100 – 200 cm)	13,116	1,018	8	94	14,134
RSC	delay (200 - 300 cm)	9,533	729	8	70	10,262
Posterior	cue, 1 st quarter (0 – 50 cm)	16,147	455	4	126	16,602
Posterior	cue, 2 nd quarter (50 – 100 cm)	16,287	482	4	128	16,769
Posterior	cue, 1 st half (0 – 100 cm)	14,439	355	3	119	14,794
Posterior	cue, 3 rd quarter (100 – 150 cm)	16,252	474	4	127	16,726
Posterior	cue, 2 nd half (100 – 200 cm)	14,038	312	3	118	14,350
Posterior	delay (200 - 300 cm)	14,117	357	3	115	14,474
mM2	cue, 1 st quarter (0 – 50 cm)	21,787	703	7	191	22,490
mM2	cue, 2 nd quarter (50 – 100 cm)	21,615	692	7	187	22,307
mM2	cue, 1 st half (0 – 100 cm)	19,972	588	7	161	20,560
mM2	cue, 3 rd quarter (100 – 150 cm)	21,766	674	7	190	22,440
mM2	cue, 2 nd half (100 – 200 cm)	21,015	634	8	167	21,649
mM2	delay (200 - 300 cm)	19,755	605	7	159	20,360
aM2	cue, 1 st quarter (0 – 50 cm)	16,025	426	6	135	16,451

aM2	cue, 2 nd quarter (50 – 100 cm)	16,023	420	6	135	16,443
aM2	cue, 1 st half (0 – 100 cm)	20,338	576	7	166	20,914
aM2	cue, 3 rd quarter (100 – 150 cm)	16,046	425	6	135	16,471
aM2	cue, 2 nd half (100 – 200 cm)	21,210	658	8	169	21,868
aM2	delay (200 - 300 cm)	19,737	616	7	158	20,353
M1	cue, 1 st quarter (0 – 50 cm)	13,894	369	6	117	14,263
M1	cue, 2 nd quarter (50 – 100 cm)	13,938	382	6	119	14,320
M1	cue, 1 st half (0 – 100 cm)	14,806	387	3	125	15,193
M1	cue, 3 rd quarter (100 – 150 cm)	14,009	364	6	121	14,373
M1	cue, 2 nd half (100 – 200 cm)	14,774	395	3	122	15,169
M1	delay (200 - 300 cm)	14,721	386	3	123	15,107
Frontal	cue, 1 st quarter (0 – 50 cm)	16,255	475	4	126	16,730
Frontal	cue, 2 nd quarter (50 – 100 cm)	16,178	485	4	128	16,663
Frontal	cue, 1 st half (0 – 100 cm)	14,410	369	3	120	14,779
Frontal	cue, 3 rd quarter (100 – 150 cm)	16,065	483	4	123	16,548
Frontal	cue, 2 nd half (100 – 200 cm)	14,392	377	3	121	14,769
Frontal	delay (200 - 300 cm)	14,511	365	3	118	14,876
Total unique count		100,787	27,606	20	855	128,393

Figure 1–table supplement 1. Numbers of mice, sessions and trials for each of the 54 experimental conditions.

Last line shows the number of unique mice and trials across all experiments, as conditions were partially overlapping for a given mouse and behavioral session.

764 **Figure 1-figure supplement 1. Effects of subtrial inactivations on psychometric functions during all**
765 **54 area-epoch combinations.**

766 Black lines show control trials and blue lines show inactivation trials, for data combined across mice.

767 Error bars, binomial confidence intervals. Lines are best-fitting psychometric functions.

768 **Figure 1-figure supplement 2. Effects of subtrial inactivations on evidence-weighting functions**
 769 **during all 54 area-epoch combinations.**

770 Black lines show the inactivation-induced change in evidence weights (laser on – laser off), and shaded
 771 areas indicate inactivation epoch. Data were combined across mice. Error bars, S.D. across 10,000
 772 bootstrapping iterations. Gray circles indicate statistical significance according to the caption on top.

Figure 2-figure supplement 1. Inactivation of cortical areas does not decrease running speeds.

Effects of subtrial inactivations on running speed during all 54 area-epoch combinations. Each panel shows inactivation-induced change in speed during the laser-on epoch or equivalent maze regions in control trials, for data combined across mice. Error bars: S.D. across 10,000 bootstrapping iterations. There were no significant decreases, as assessed by a one-sided bootstrapping test (see Materials and Methods).

779 **Figure 2-figure supplement 2. Little effect on evidence weighting after laser offset.**
780 Laser-offset-aligned changes in evidence-weighting curves for each bilaterally targeted area (laser on -
781 laser off). Thin gray lines, individual inactivation epochs ($n = 4$ for $y = 50$, 3 for $y = 100$). Thick black lines,
782 average across conditions. Error bars, $\pm$ SEM across experimental conditions, for data combined across
783 mice. Circles below the lines indicate statistical significance (one circle: $p < 0.05$, two circles: $p < 0.01$;
784 one-sided paired t test vs. zero, corrected for multiple comparisons).

785 **Figure 2-figure supplement 3. Comparison between simultaneous multi-area inactivation and**
786 **average across the corresponding individually inactivated areas.**
787 Average laser-onset-aligned changes in evidence-weighting curves for each set of experimental
788 conditions (caption on top), for data combined across mice. Error bars, $\pm$ SEM across experimental
789 conditions. Shaded areas, laser on periods. P-values on top are from a two-way ANOVA with repeated
790 measures, factors inactivation type (simultaneous vs. individual) and y position. There were no
791 differences between simultaneous and individual inactivation of either frontal or posterior areas. On the
792 other hand, effects depended on the distance of evidence from laser onset for posterior but not frontal
793 areas.

794 **Figure 3—figure supplement 1. GLM coupling coefficients for ROI activity predictors.**

795 For each ROI (rows), shown are the average coefficients (n=6 mice) for the predictors consisting of
 796 zero-lag activity of the other simultaneously imaged ROIs. Note that the diagonal elements are not
 797 defined since zero-lag self-activity predictors were not in the model. Data from the somatosensory
 798 cortex ROI has been omitted for symmetry with the inactivation data. L and R indicate left and right
 799 hemispheres, respectively.