Multiple timescales of sensory-evidence accumulation across the dorsal cortex

Lucas Pinto[1,§], David W. Tank[1,2,#] , Carlos D. Brody[1,2,3,#,*]

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

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

[#] these senior authors contributed equally to this work

Correspondence:
David W. Tank, dwtank@princeton.edu
Carlos D. Brody, brody@princeton.edu

*** Lead contact:** brody@princeton.edu

Highlights

- Mice navigated in virtual reality while accumulating evidence over seconds
- We briefly inactivated dorsal cortical regions at different points in the trials
- Inactivation of different regions resulted in accumulation deficits on distinct timescales
- Dorsal cortical regions displayed a hierarchy of activity timescales

Key words
Decision making, evidence accumulation, intrinsic timescales, cortex, virtual reality, optogenetic inactivation, widefield Ca^{2+} imaging, mice

SUMMARY

Decisions requiring the gradual accrual of sensory evidence appear to recruit widespread cortical areas. However, the nature of the contributions of different regions remains unclear. Here we trained mice to accumulate evidence over seconds while navigating in virtual reality, and optogenetically silenced the activity of many cortical areas during different brief trial epochs. We found that the inactivation of different areas primarily affected the evidence-accumulation computation per se, rather than other decision-related processes. Specifically, we observed selective changes in the weighting of evidence over time, such that frontal inactivations led to deficits on longer timescales than posterior cortical ones. Likewise, large-scale cortical Ca^{2+} activity during task performance displayed different temporal integration windows matching the effects of inactivation. Our findings suggest that distributed cortical areas accumulate evidence following their hierarchy of intrinsic timescales.

INTRODUCTION

The decisions we make in our daily lives often unfold over time as we deliberate between competing choices. In the laboratory, 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). Decisions under uncertainty can be potentially solved in the brain through a chain of largely feedforward computations progressing from stimulus processing to accumulation of evidence to post-accumulation categorization and decision, with different brain areas implementing different steps of the process (Brincat et al., 2018; Brody and Hanks, 2016; Erlich et al., 2015; Gold and Shadlen, 2007; Hanks et al., 2015; Yartsev et al., 2018). On the other hand, neural correlates of decisions relying on evidence accumulation have been found in a number of cortical and subcortical structures, in both primates and rodents (Ding and Gold, 2010; 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., 2019; 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).

Thus, the findings reviewed above point towards the existence of widespread cortical and subcortical mechanisms supporting decision making under uncertainty. However, the exact nature of the contribution of different brain areas is largely unknown. The striatum appears to be a central hub for evidence accumulation (Yartsev et al., 2018), but it remains unclear whether widespread cortical dynamics reflect evidence accrual, other decision processes, or both. Further, if these dynamics do reflect accumulation of sensory evidence, we do not understand the logic underlying the observed widespread cortical recruitment. One hitherto unexplored possibility is that it exploits an intrinsic hierarchy of integration timescales, whereby frontal areas integrate input over longer time windows than sensory areas (Chaudhuri et al., 2015; Gao et al., 2020; Hasson et al., 2008; Ito et al., 2020; Murray et al., 2014; Runyan et al., 2017; Spitmaan et al., 2020). Here, to distinguish between these possibilities, we trained

mice to accumulate evidence over seconds towards navigational decisions, and used brief inactivations 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 of this hierarchy.

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)(Figs. 1A, B)(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 (Fig. 1C, 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 regression model that used the net sensory evidence (Δ towers, or #R – #L) from each of four equally-spaced bins from the cue region to predict the choice the mice made. In other words, we inferred the weight of sensory evidence from different positions in the maze on the final decision. While individual mice showed different evidence-weighting profiles, fitting the model on aggregate data yielded a flat evidence-weighting curve (Fig. 1D, n = 100,787 trials), indicating that on average the mice weighted evidence equally from throughout the maze (Pinto et al., 2018).

Our previous results have shown that cortical contributions to the performance of this task are widespread (Pinto et al., 2019), but our whole-trial inactivations did not allow us to tease apart the nature of the contributions of different areas. Here, we addressed this by asking how different dorsal cortical regions contribute to the weighting of sensory evidence in order to make a perceptual decision. To do this we cleared the intact skull of mice expressing Channelrhodopsin-2 (ChR2) in inhibitory interneurons (VGAT-ChR2-EYFP, n = 20), and used a scanning laser system to bilaterally silence different cortical regions, by activating inhibitory cells (Figs. 1A, B) (Guo et al., 2014; Pinto et al., 2019). We targeted 7 different areas – primary visual cortex (V1), medial secondary visual cortex (mV2, roughly corresponding to area AM), posterior parietal cortex (PPC), retrosplenial cortex (RSC), the posteromedial portion of the premotor cortex (mM2), the anterior portion of the premotor cortex (aM2), and the primary motor cortex (M1) – as well as two combinations of these individual ares, namely posterior cortex (V1, mV2, PPC and RSC) and frontal cortex (mM2, aM2 and M1). Cortical silencing occured in one of six 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). We tested all the 54 possible area-epoch combinations (Table S1). This large number of experimental conditions allowed us to assess how the inactivation of different areas affects the use of current or past sensory evidence towards a final decision.

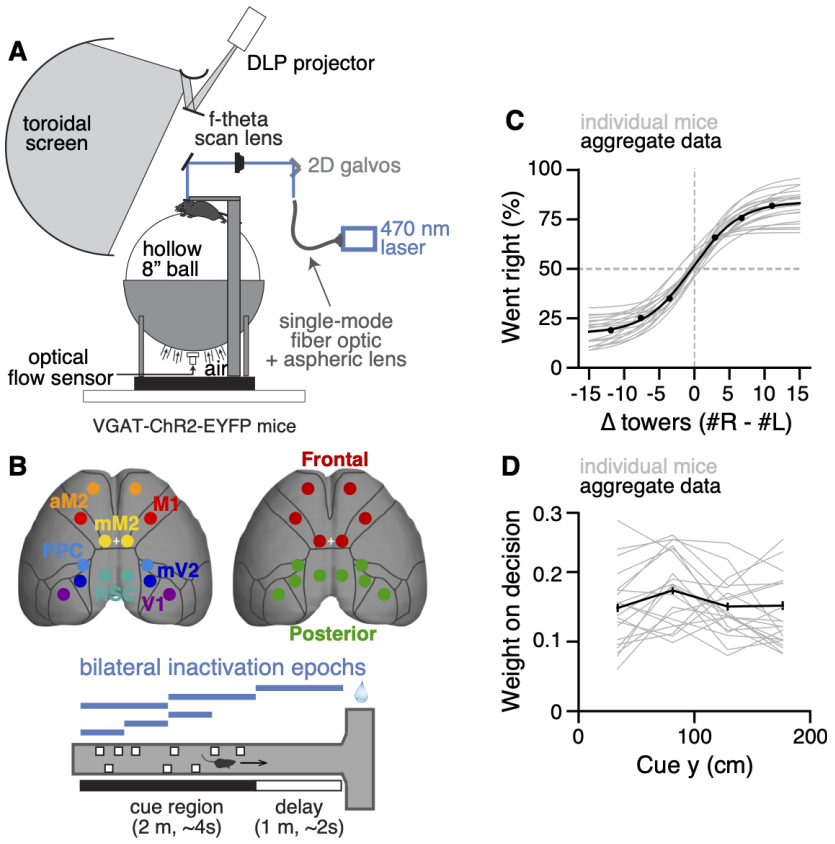

Figure 1. Temporally specific inactivation of multiple dorsal cortical regions during performance of a VR-based evidence-accumulation task.
(A) Schematics of the experimental set-up.
(B) Experimental design. We bilaterally inactivated 7 dorsal cortical areas, alone or in combination, while mice performed the accumulating-towers task. Bilateral inactivations happened during one of six regions in the maze spanning different parts of the cue region or delay. We thus tested a total of 54 area-epoch combinations.
(C) Psychometric functions for control trials, showing the probability of right choice as a function of the strength of right sensory evidence, Δ towers (#R – #L). Thin gray lines: best-fitting psychometric functions for each individual mouse (n = 20). Black circles: aggregate data (n = 100,787 trials), black line: fit to aggregate data, error bars: binomial confidence intervals.
(D) Logistic regression curves for the weight of sensory evidence from four equally-spaced bins on the final decision, from control trials. Thin gray lines: individual animals, thick black line: aggregate data, error bars: ± SD from 200 bootstrapping iterations.
See also Figures S1, S2 and Table S1.

Compatible with our previous whole-trial inactivation experiments (Pinto et al., 2019), we found that the inactivation of all tested cortical areas significantly affected behavioral performance, though to varying degrees (Figs. S1, S2). Furthermore, we observed a variety of effect profiles across regions and inactivation epochs, as assessed by the difference between the evidence-weighting curves separately calculated for 'laser off' and 'laser on' trials (Fig. S3). Different effects were observed even comparing regions that were in close physical proximity (e.g. V1 and mV2). Additionally, all tested areas had significant effects in at least a subset of conditions (Fig. S3, $p<0.05$, bootstrapping).

Most changes in the evidence-weighting curves happened for evidence concomitant to or preceding laser onset, indicating that the manipulations primarily affected the processing and/or memory of the evidence, i.e. the accumulation process itself (Fig. S3). To quantify this, for each cortical area we aligned the control-subtracted evidence-weighting curves from different inactivation epochs by the position of laser onset, and focused on the changes in weights of evidence occurring up to 100 cm in the past (~2s, Fig. 2A). While some variability in these laser-onset-triggered curves suggests that the effects of inactivation depend somewhat on the exact inactivation epoch, the aligned curves from different epochs were fairly consistent (Fig. 2A, gray lines), providing a concise summary of the multiple experimental conditions. Interestingly, the effects varied systematically according to the inactivated region. For example, mV2 inactivation led to a significant drop in the weight of evidence occurring while the laser was on ($p = 0.004$, one-sided paired t test), a trend towards affecting the memory of evidence occurring 50 cm in the past (~1 s, $p = 0.045$, not significant after false discovery rate correction), and no discernible effect on evidence occurring 100 cm in the past (~2 s, $p = 0.41$). Conversely, aM2 inactivation led to significant decreases in weighting all evidence between 100 cm in the past and the time of laser onset ($p < 0.05$, one-sided paired t test), with no differences in magnitude between position bins ($F_{2,10} = 0.27$, $p = 0.77$, one-way ANOVA with repeated measures). This lack of modulation of effect size across y position bins was also true for the two other frontal areas, mM2 and M1 (Fig. 2A, $p > 0.88$). Thus, subsets of cortical areas resembled each other in terms of the effects of their inactivation. Indeed, they could be optimally grouped into three clusters using spectral clustering (Fig. 2B). Cluster 1 contained all frontal areas in our dataset: M1, mM2 and aM2. On average, this cluster resembled the effects described for aM2 above. In other words, the inactivation of frontal areas tended to equally and significantly affect weights for evidence occurring up to 100 cm in the past, suggesting that these areas accumulate evidence at fairly long timescales ($p < 0.001$ for all position bins, one-sided paired t test). Cluster 2 contained mV2 and PPC, and on average showed monotonically decreasing effects of inactivation on the weight of evidence as it gets more distal from laser onset ($p < 0.02$ for 0 and 50 cm, $p = 0.47$ for 100 cm). Thus, compared to the frontal area cluster, these posterior areas contributed to evidence accumulation on shorter timescales. Finally, cluster 3 contained V1 and RSC, whose inactivation led to non-monotonic changes in evidence weighting, affecting current and long-past evidence ($p < 0.001$), but not evidence occurring in between ($p = 0.07$). This is potentially compatible with findings that multiple timescales of processing can be present within the same cortical regions (Bernacchia et al., 2011; 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.

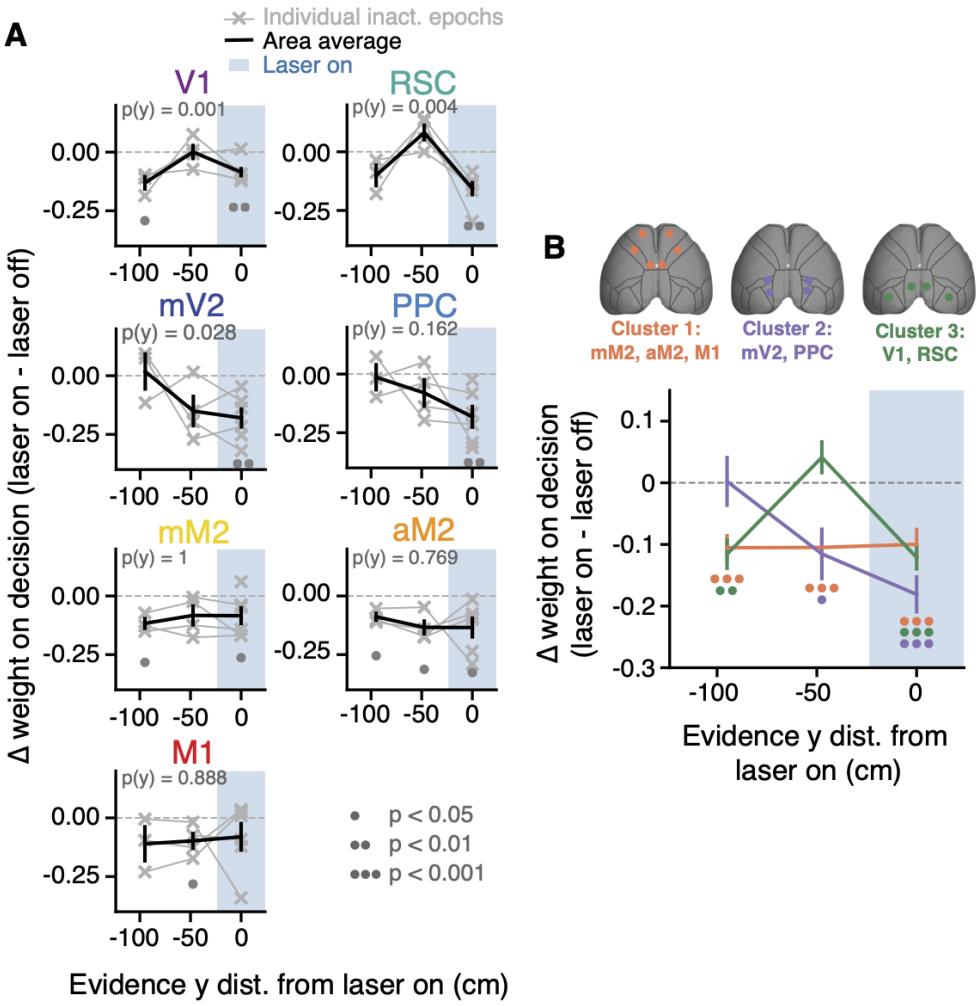

Figure 2. Inactivating different cortical areas leads to evidence-accumulation deficits on distinct timescales.
(A) Laser-onset-aligned changes in evidence-weighting curves for each bilaterally targeted area (laser on - laser off). Thin gray lines, individual inactivation epochs (n = 6 for y = 0, 4 for y = 50 and 3 for y = 100). Thick black lines, average across conditions. Error bars, ± 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.
(B) Average laser-onset-aligned changes in evidence-weighting curves for each custer (caption on top). Error bars, ± 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).
See also Figures S3 – S5.

Next, we performed a similar analysis to assess changes in evidence weights after laser offset. Confirming our initial impression, the inactivation of most areas did not impact evidence weighting prospectively (Fig. S4). Interestingly, however, mV2 and mM2 inactivation led to

moderately but significantly decreased use of evidence occurring 100 cm in the future ($p < 0.05$, one-sided paired t test), perhaps suggesting an additional role for these areas also in post-accumulation decision processes (Hanks et al., 2015). Thus, our inactivation data suggest that the widespread cortical involvement in this task is largely related to the accumulation of sensory evidence, and that different cortical areas accumulate on distinct timescales. We next wondered whether evidence information from different areas is linearly combined, at least from a behavioral standpoint. To do this, we compared the effects of simultaneously inactivating all frontal or posterior areas to that expected by a linear combination of the effects of inactivating areas individually (i.e. their average). Neither posterior nor frontal areas significantly deviated from the linear prediction (Fig. S5, $p > 0.05$, two-way ANOVA with repeated measures, factors y position and inactivation type). This suggests that signals from the different dorsal cortical areas could be combined by downstream regions in a near-linear fashion. Candidate regions include the medial prefrontal cortex, or subcortical structures such as the striatum and the cerebellum, which have been shown to be causally involved in evidence accumulation (Deverett et al., 2019; Yartsev et al., 2018). Other subcortical candidates are midbrain regions shown to have a high incidence of choice signals in a contrast discrimination task (Steinmetz et al., 2019).

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 (Fig. 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)(Fig. S6). 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 the held-out data (Fig. 3B, C, $F_{model(6,2,12)} = 1994.85$, two-way ANOVA with repeated measures; the full model had an accuracy of ~0.95).

Motivated by our inactivation findings, we focused our analysis on the auto-regressive coefficients of the model. We observed that across animals the rate of decay of these coefficients over lags slowed systematically from visual to premotor areas, with intermediate values for M1, PPC and RSC (Fig. 3D). To quantify this, we fitted exponential decay functions to the auto-regressive coefficients averaged across hemispheres (Fig. 3D, E), and extracted

decay time constants (τ, Fig. 3F). Compatible with our observations, τ differed significantly across cortical areas ($F_{6,30}$ = 3.84, p = 0.006, one-way ANOVA with repeated measures), being larger for frontal than posterior areas, in particular PPC and mV2. Note that, while it is possible that these coefficients capture auto-correlations introduced by intrinsic GCaMP6f dynamics, there is no reason to believe that this affects our conclusions, as indicator dynamics should be similar across regions. Thus, during the evidence-accumulation task, cortical regions display increasing intrinsic timescales going from visual to frontal areas. This is consistent with previous reports for spontaneous activity and other behavioral tasks (Chaudhuri et al., 2015; Hasson et al., 2008; Murray et al., 2014; Runyan et al., 2017). Moreover, it is in overall agreement with our inactivation findings (Fig. 2), and suggests that the different intrinsic timescales across the cortex support evidence integration over time windows of different durations.

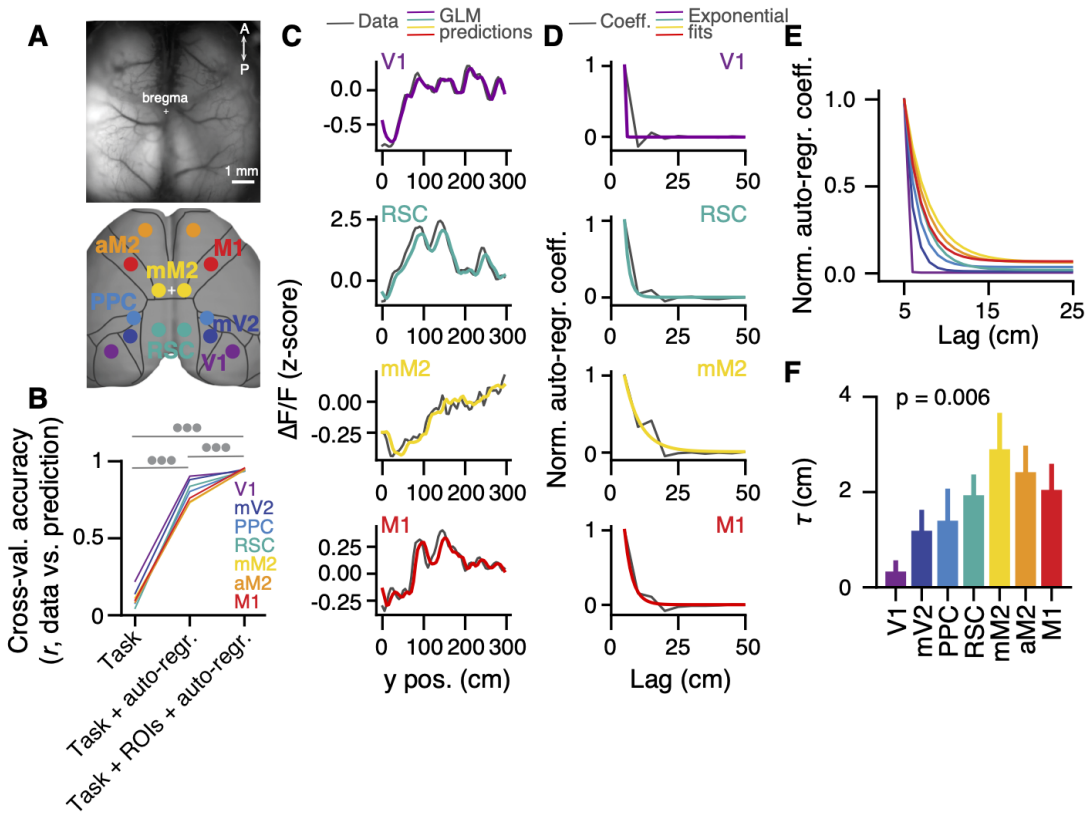

Figure 3. A hierarchy of activity timescales during evidence accumulation.
(A) Top: example widefield imaging field of view showing GCaMP6f fluorescence across the dorsal cortex. Bottom: approximate correspondence between the field of view and ROIs defined from the Allen Brain Atlas, ccv3.
(B) Cross-validated accuracies of different versions of the GLM. Each line corresponds to the average for each ROI across hemispheres and mice (n = 6). Error bars are omitted for clarity. Circles indicate statistical significance (p < 0.001, Tukey's post-hoc test).
(C) Example of actual ΔF/F (gray) and GLM predictions (colored lines) for 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 spatial lags for an example imaging session and four example ROIs. Gray, coefficient values. Colored lines, best-fitting exponential decay functions. Note that this model was fitted in spatially downsampled activity for symmetry with the inactivation data, such that time constants are in units of cm. Given average running speeds, each cm corresponds to ~ 20 ms.

(E) Exponential decay functions for all seven cortical areas, fitted to the average across mice (n = 6).
(F) Time constants extracted from the exponential decay fits, for each area. Error bars, ± SEM across mice (n = 6). P-value is from a one-way ANOVA with repeated measures with ROIs as factors.
See also Figure S6.

DISCUSSION

Taken together, our results suggest that distributed cortical areas contribute to sensory-evidence accrual on different timescales. Specifically, brief sub-trial inactivations during performance of a decision-making task requiring seconds-long evidence accumulation resulted in distinct deficits in the weighting of sensory evidence from different points in the stimulus stream. This was such that, on average, the inactivation of frontal cortical areas resulted in decreased use of evidence occurring further in the past from laser onset compared to a subset of posterior regions (Fig. 2). Compatible with this, using an encoding model of large-scale cortical dynamics, we found that activity timescales vary systematically across the cortex in a way that mirrors the inactivation results (Fig. 3).

Our results add causal evidence to a growing body of literature that has revealed that the cortex of rodents and primates appears to be organized in a hierarchy of temporal processing windows across regions (Chaudhuri et al., 2015; Gao et al., 2020; Hasson et al., 2008; Ito et al., 2020; Murray et al., 2014; Runyan et al., 2017; Spitmaan et al., 2020). Further, they 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.

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.).

AUTHOR CONTRIBUTIONS

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

METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Experimental Models: Organisms/Strains		
Mouse: B6.Cg-Tg(Slc32a1-COP4*H134R/EYFP)8Gfng/J	The Jackson Laboratory	JAX: 014548
Mouse: IgS6$^{tm93.1(tetO-GCaMP6f)Hze}$ Tg(Camk2a-tTA)1Mmay/J	The Jackson Laboratory	JAX: 024108
Mouse: B6.129S2-Emx1$^{tm1(cre)Krj}$/J	The Jackson Laboratory	JAX: 005628
Software and Algorithms		
Matlab 2015b, 2016b, 2017b	Mathworks	https://www.mathworks.com/products/matlab.html
ViRMEn	Aronov and Tank, 2014	https://pni.princeton.edu/pni-software-tools/virmen
NI DAQmx 9.5.1	National Instruments	https://www.ni.com/en-us/support/downloads/drivers/download.ni-daqmx.html
HCImage	Hamamatsu	https://hcimage.com
Scanning laser control software	Pinto et al., 2019	https://github.com/BrainCOGS/laserGalvoControl.git
Widefield data analysis software	Pinto et al., 2019	https://github.com/BrainCOGS/widefieldImaging.git
Inactivation analysis	This paper	https://github.com/BrainCOGS/PintoEtAl2020_subtrial_inact.git
Python 3.7	Python	https://www.python.org
Numpy 1.18.1	Walt et al., 2011	https://www.numpy.org
Scipy 1.4.1	Virtanen et al., 2020	https://www.scipy.org

Deepdish 0.3.4	University of Chicago	https://github.com/uchicago-cs/deepdish
Statsmodels 0.11.0	Skipper, 2010	https://www.statsmodels.org
Matplotlib 3.1.3	Hunter, 2007	https://www.matplotlib.org
Pandas 1.0.1	McKinney, 2010	https://www.pandas.pydata.org
Pingouin 0.3.8	Vallat, 2018	https://pingouin-stats.org
Mat7.3	Simon Kern	https://github.com/skjerns/mat7.3
Other		
Virtual reality setup tools	Pinto et al., 2018	https://github.com/sakoay/AccumTowersTools

LEAD CONTACT AND MATERIALS AVAILABILITY

This study did not generate new unique reagents. Further information requests may be directed to the Lead Contact, Dr. Carlos D. Brody (brody@princeton.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

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) and diluted clear metabond (Parkell) were applied evenly to the dorsal surface of the skull, and polished after curing using a dental polishing kit (Pearson dental). The headplate was attached to the cleared skull using metabond, and a layer of transparent nail polish (Electron Microscopy Sciences) was applied and allowed to cure for 10 – 15 min. The procedure was done under isoflurane anesthesia (2.5% for induction, 1.5% for maintenance). The animals received two doses of

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

The analysis reported in Fig. 3 (widefield Ca^{2+} imaging) is from data collected in the context of a previous study (Pinto et al., 2019), although the analysis is novel. The data was from 6 male and female mice from triple transgenic crosses expressing GCaMP6f under the CaMKIIα promoter from the following two lines: Ai93-D;CaMKIIα-tTA [IgS6$^{tm93.1(tetO\text{-}GCaMP6f)Hze}$ Tg(Camk2a-tTA)1Mmay/J, Jackson Laboratories, stock # 024108] and Emx1-IRES-Cre [B6.129S2-Emx1$^{tm1(cre)Krj}$/J, Jackson Laboratories, stock # 005628]. These animals also underwent the surgical procedure described above.

METHODS DETAILS

Virtual reality apparatus

The mice were trained in a virtual reality (VR) environment (Fig. 1A) described in detail elsewhere (Pinto et al., 2018). Briefly, they sat on an 8-inch hollow Styrofoam® ball that was suspended by compressed air at ~60 p.s.i, after passing through a laminar flow nozzle to reduce noise (600.326.5K.BC, Lechler). They were head-fixed such that their snouts were aligned to the ball equator and at a height such that they could run comfortable without hunching, while still being able to touch the ball with their full paw pads (corresponding to a headplate-to-bal height of ~1 inch for a 25-g animal). Ball movements were measured using optical flow sensors (ADNS-3080 APM2.6) and transformed into virtual world displacements using custom code running on Arduino Due (https://github.com/sakoay/AccumTowersTools/tree/master/OpticalSensorPackage). The ball sat on a custom 3D-printed cup that contained both the air outlet and the movement sensor. The VR environment was projected onto a custom-built toroidal Styrofoam® screen using a DLP projector (Optoma HD141X) at a refresh rate of 120 Hz and a pixel resolution of 1024 x 768. The screen spanned ~270˚ of azimuth and ~80˚ of elevation in the mouse's visual field. The whole set-up was enclosed in a custom-built sound-attenuating chamber. The VR environment was programmed and controlled using ViRMEn (Aronov and Tank, 2014) (https://pni.princeton.edu/pni-software-tools/virmen), running on Matlab on a PC.

Behavioral task

We trained the mice in the accumulating-towers task (Pinto et al., 2018). The mice ran down a virtual T-maze that was 3.3 m in length (y), 5 cm in height and a nominal 10 cm in width (x, though they were restricted to the central 1 cm). The length of the maze consisted of a 30-cm start region to which they were teleported at the start of each trial, followed by a 200-cm cue region and a 100-cm delay region. The cue and the delay region had the same wallpaper designed to provide optical flow. During the cue region, the mice 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 x 11 x 5 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) was directed to 2-D galvanometers using a 125-μm single-mode optic fiber optic (Thorlabs) and reached the cortical surface after passing through an f-theta scanning lens (LINOS). 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 ~ 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 μs) 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 Ca^{2+} imaging

Details on the experimental setup and data preprocessing can be found elsewhere (Pinto et al., 2019). Briefly, we used a tandem-lens macroscope (1x – 0.63x planapo, Leica M series) 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), with an image size of 512 x 512 pixels (pixel size of ~17 μm). Images were acquired with HCImage (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) 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 ΔF/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.

QUANTIFICATION AND STATISTICAL 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 (Mathworks), 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 Table S1). 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.

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 in the number of right and left towers (#R – #L, or Δ). Δ was binned in increments of 5 between

-15 and 15, and its value defined as the average Δ weighted by the number of trials. We fitted the psychometric curves using a 4-parameter sigmoid:

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

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

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

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

Laser-triggered analysis. For each area, we aligned the evidence-weighting curves by the position of laser onset, which was defined as y = 0, and used y position bins going up to 100 cm in the past. Thus, each inactivation condition contributed up to 3 bins of data, depending on the position of laser onset. For inactivations lasting more than one position bin (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 (Fig. S4), 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 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 and evidence weights on decision. To analyze the laser-onset triggered curves (Fig. 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 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 (Fig. S5), 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 x 3 (areas x 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 inactivated simultaneously from this analysis. We then performed spectral clustering into *k* clusters on that matrix. We tested *k* = 2 – 5, and chose the value of *k* that maximized clustering quality as measured by the Calinski-Harabasz index. Given the small number of areas per cluster, we generated error estimates for each y position bin by concatenating the individual inactivation conditions (epochs) for all areas of the cluster.

Generalized linear model (GLM) of widefield data

We fitted Ca^{2+} activity averaged over each anatomically defined ROI with a generalized linear model (GLM)(Pinto and Dan, 2015; Pinto et al., 2019; Scott et al., 2017). Activity was spatially downsampled in 5-cm steps so as to avoid the confound of different occupancies across different spatial positions, as follows. For each trial and y position in the maze, we sampled ΔF/F from a single frame corresponding to the first point in time that the mouse visited that position, limited to $0 \leq y \leq 300$ cm (i.e. trial start, outcome and inter-trial periods were not included). Thus, all the trials had the same length. The spatially downsampled activity was then z-scored across all trials. ΔF/F of each area was modeled as a linear combination of different predictors at different spatial lags. In addition to the previously used task-event predictors (Pinto et al., 2019), we added coupling terms i.e. the zero-lag activity of the other simultaneously imaged ROIs (Pillow et al., 2008; Runyan et al., 2017), as well as auto-regressive terms to capture activity auto-correlations that were independent of task events (Spitmaan et al., 2020). Finally, we added a term to penalize the L2 norm of the coefficients, i.e. we performed ridge regression. The full model was thus defined as:

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

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

$$A = \sum_{i=5}^{100} \beta_i^{autoregr} \Delta F/F(y-i)$$

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

$$T = \sum_{i=0}^{100} \beta_i^{tR} E_{y-i}^{tR} + \sum_{i=0}^{100} \beta_i^{tL} E_{y-i}^{tL} + \sum_{i=0}^{50} \beta_i^{\Delta} E_{y-i}^{\Delta} + \sum_{i=-30}^{30} \beta_i^{d\theta/dt} E_{y-i}^{d\theta/dt} + \sum_{i=-30}^{30} \beta_i^{sp} E_{y-i}^{sp} + \beta^{\theta}\theta(y) + \beta^{y} y + \beta^{ch} ch + \beta^{pch} pch + \beta^{prw} prw$$

In the above equations, β_i^x is the encoding weight for predictor *x* at spatial lag *i* (in steps of 5 cm), where *x* is either a task event or the activity of the ROI at a previous spatial location, and $\beta_j^{coupling}$ is the weight for the zero-lag activity for simultaneously imaged ROI *j* (we had a total of 16 ROIs across the two hemispheres). In the task term, E_y^x is a delta function indicating the occurrence of event *x* at position *y*. Specifically, *tR* indicates the occurrence of a right tower, *tL* of a left tower, *Δ* = cumulative #R – #L towers, *dθ/dt* is virtual view angle velocity, *sp* is running speed, *θ* is view angle (no lags), *y* is spatial position in the maze stem (no lags), and *ch*, *pch* and *prw* are constant offsets for a given trial, indicating upcoming choice, previous choice (+1 for right and –1 for left) and previous reward (1 for reward and –1 otherwise), respectively.

Cross-validation. The model was fitted using 3-fold cross-validation. For each of 20 values of the penalty term *λ*, we trained the model using ⅔ of the trials (both correct and wrong choices), and tested it on the remaining ⅓ of trials. We picked the value of *λ* that maximized 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 ΔF/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 Fig. 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 5 and 100 cm in the past, normalized to the coefficient at 5 cm (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 $P_i \leq (\alpha i)/n$, 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.

REFERENCES

Aronov, D., and Tank, D.W. (2014). Engagement of Neural Circuits Underlying 2D Spatial Navigation in a Rodent Virtual Reality System. Neuron *84*, 442–456.

Benjamini, Y., and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. J. R. Stat. Soc. Series B Stat. Methodol. *57*, 289–300.

Bernacchia, A., Seo, H., Lee, D., and Wang, X.-J. (2011). A reservoir of time constants for memory traces in cortical neurons. Nat. Neurosci. *14*, 366–372.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765.

Brincat, S.L., Siegel, M., von Nicolai, C., and Miller, E.K. (2018). Gradual progression from sensory to task-related processing in cerebral cortex. Proc. Natl. Acad. Sci. U. S. A. *115*, E7202-E7211.

Brody, C.D., and Hanks, T.D. (2016). Neural underpinnings of the evidence accumulator. Curr. Opin. Neurobiol. *37*, 149–157.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.

Carandini, M., and Churchland, A.K. (2013). Probing perceptual decisions in rodents. Nat. Neurosci. *16*, 824–831.

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

Deverett, B., Kislin, M., Tank, D.W., and -H. Wang, S.S. (2019). Cerebellar disruption impairs working memory during evidence accumulation. Nat. Comm. *10*, 3128.

Ding, L., and Gold, J.I. (2010). Caudate encodes multiple computations for perceptual decisions. J. Neurosci. *30*, 15747–15759.

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

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*, e05457.

Fetsch, C.R., Odean, N.N., Jeurissen, D., El-Shamayleh, Y., Horwitz, G.D., and Shadlen, M.N. (2018). Focal optogenetic suppression in macaque area MT biases direction discrimination and decision confidence, but only transiently. Elife *7*, e36523.

Fulcher, B.D., Murray, J.D., Zerbi, V., and Wang, X.-J. (2019). Multimodal gradients across mouse cortex. Proc. Natl. Acad. Sci. U. S. A. *116*, 4689–4695.

Ganupuru, P., Goldring, A.B., Harun, R., and Hanks, T.D. (2019). Flexibility of Timescales of Evidence Evaluation for Decision Making. Curr. Biol. *29*, 2091–2097.e4.

Gao, R., van den Brink, R.L., Pfeffer, T., and Voytek, B. Neuronal timescales are functionally dynamic and shaped by cortical microarchitecture.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Guo, Z.V., Li, N., Huber, D., Ophir, E., Gutnisky, D., Ting, J.T., Feng, G., and Svoboda, K. (2014). Flow of Cortical Activity Underlying a Tactile Decision in Mice. Neuron *81*, 179–194.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223.

Hasson, U., Yang, E., Vallines, I., Heeger, D.J., and Rubin, N. (2008). A hierarchy of temporal receptive windows in human cortex. J. Neurosci. *28*, 2539–2550.

Horwitz, G.D., and Newsome, W.T. (1999). Separate signals for target selection and movement specification in the superior colliculus. Science *284*, 1158–1161.

Hunter, J.D. (2007). Matplotlib: A 2D Graphics Environment. Comput. Sci. Eng. *9*, 90–95.

Ito, T., Hearne, L.J., and Cole, M.W. (2020). A cortical hierarchy of localized and distributed processes revealed via dissociation of task activations, connectivity changes, and intrinsic timescales. NeuroImage *221*, 117141.

Jeneson, A., and Squire, L.R. (2012). Working memory, long-term memory, and medial temporal lobe function. Learn. Mem. *19*, 15–25.

Joshua, M., and Lisberger, S.G. (2015). A tale of two species: Neural integration in zebrafish

and monkeys. Neuroscience *296*, 80–91.

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285–288.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185.

Koay, S.A., Thiberge, S.Y., Brody, C., and Tank, D.W. (2020). Amplitude modulation of cortical sensory responses in pulsatile evidence accumulation. Elife *9*, e60628.

Krueger, P.M., van Vugt, M.K., Simen, P., Nystrom, L., Holmes, P., and Cohen, J.D. (2017). Evidence accumulation detected in BOLD signal using slow perceptual decision making. J. Neurosci. Methods *281*, 21–32.

McKinney, W. (2010). Data structures for scientific computing in Python. In: Proceedings of the 9th Python in Science Conference, *445*.

Miri, A., Daie, K., Arrenberg, A.B., Baier, H., Aksay, E., and Tank, D.W. (2011). Spatial gradients and multidimensional dynamics in a neural integrator circuit. Nat. Neurosci. *14*, 1150–1159.

Morcos, A.S., and Harvey, C.D. (2016). History-dependent variability in population dynamics during evidence accumulation in cortex. Nat. Neurosci. *19*, 1672–1681.

Murphy, P.R., Wilming, N., Hernandez-Bocanegra, D.C., Ortega, G.P., and Donner, T.H. (2020). Adaptive circuit dynamics across human cortex during evidence accumulation in changing environments. bioRxiv, https://doi.org/10.1101/2020.01.29.924795.

Murray, J.D., Bernacchia, A., Freedman, D.J., Romo, R., Wallis, J.D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., et al. (2014). A hierarchy of intrinsic timescales across primate cortex. Nat. Neurosci. *17*, 1661–1663.

Newsome, W.T., Britten, K.H., and Movshon, J.A. (1989). Neuronal correlates of a perceptual decision. Nature *341*, 52–54.

Odoemene, O., Pisupati, S., Nguyen, H., and Churchland, A.K. (2018). Visual evidence accumulation guides decision-making in unrestrained mice. J. Neurosci. *38*, 10143–10155.

Orsolic, I., Rio, M., Mrsic-Flogel, T.D., and Znamenskiy, P. (2019). Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making. bioRxiv, https://doi.org/10.1101/552026.

Pillow, J.W., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J., and Simoncelli, E.P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. Nature *454*, 995–999.

Pinto, L., and Dan, Y. (2015). Cell-type-specific activity in prefrontal cortex during goal-directed behavior. Neuron *87*, 437–450.

Pinto, L., Koay, S.A., Engelhard, B., Yoon, A.M., Deverett, B., Thiberge, S.Y., Witten, I.B., Tank, D.W., and Brody, C.D. (2018). An accumulation-of-evidence task using visual pulses for mice

navigating in virtual reality. Front. Behav. Neurosci. *12*, 36.

Pinto, L., Rajan, K., DePasquale, B., Thiberge, S.Y., Tank, D.W., and Brody, C.D. (2019). Task-dependent changes in the large-scale dynamics and necessity of cortical regions. Neuron *104*, 810–824.e9.

Runyan, C.A., Piasini, E., Panzeri, S., and Harvey, C.D. (2017). Distinct timescales of population coding across cortex. Nature *548*, 92–96.

Scott, B.B., Constantinople, C.M., Akrami, A., Hanks, T.D., Brody, C.D., and Tank, D.W. (2017). Fronto-parietal cortical circuits encode accumulated evidence with a diversity of timescales. Neuron *95*, 385–398.e5.

Shadlen, M.N., and Newsome, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. *86*, 1916–1936.

Skipper, S. (2010). Statsmodels: econometric and statistical modeling with Python. In: Proceedings of the 9th Python in Science Conference, *445*.

Spitmaan, M., Seo, H., Lee, D., and Soltani, A. (2020). Multiple timescales of neural dynamics and integration of task-relevant signals across cortex. Proc. Natl. Acad. Sci. U. S. A. *117*, 22522-22531.

Steinmetz, N.A., Zatka-Haas, P., Carandini, M., and Harris, K.D. (2019). Distributed coding of choice, action and engagement across the mouse brain. Nature *576*, 266–273.

Stine, G.M., Zylberberg, A., Ditterich, J., and Shadlen, M.N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. Elife *9*.

Sun, P., and Landy, M.S. (2016). A two-stage process model of sensory discrimination: an alternative to drift-diffusion. J. Neurosci. *36*, 11259–11274.

Tsetsos, K., Gao, J., McClelland, J.L., and Usher, M. (2012). Using time-varying evidence to test models of decision dynamics: bounded diffusion vs. the leaky competing accumulator model. Front. Neurosci. *6*, 79.

Vallat, R. (2018). Pingouin: statistics in Python. J. Open Source Software *3*, 1026.

Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., et al. (2020). SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods *17*, 261–272.

Walt, S. van der, van der Walt, S., Chris Colbert, S., and Varoquaux, G. (2011). The NumPy array: a structure for efficient numerical computation. Comput. Sci. Engin. *13*, 22–30.

Wang, X.-J. (2020). Macroscopic gradients of synaptic excitation and inhibition in the neocortex. Nat. Rev. Neurosci. *21*, 169–178.

Waskom, M.L., and Kiani, R. (2018). Decision making through integration of sensory evidence at prolonged timescales. Curr. Biol. *28*, 3850–3856.e9.

Wasmuht, D.F., Spaak, E., Buschman, T.J., Miller, E.K., and Stokes, M.G. (2018). Intrinsic

neuronal dynamics predict distinct functional roles during working memory. Nat. Comm. *9*, 3499.

Wilming, N., Murphy, P.R., Meyniel, F., and Donner, T.H. (2020). Large-scale dynamics of perceptual decision information across human cortex. Nat. Comm. *11*, 5109

Yartsev, M.M., Hanks, T.D., Yoon, A.M., and Brody, C.D. (2018). Causal contribution and dynamical encoding in the striatum during evidence accumulation. Elife *7*, e35929.

Yeshurun, Y., Nguyen, M., and Hasson, U. (2017). Amplification of local changes along the timescale processing hierarchy. Proc. Natl. Acad. Sci. U. S. A. *114*, 9475–9480.

Zatka-Haas, P., Steinmetz, N.A., Carandini, M., and Harris, K.D. (2020). A perceptual decision requires sensory but not action coding in mouse cortex. bioRxiv, https://doi.org/10.1101/501627.

Supplementary Table and Figures

Region	Epoch	Num. control trials	Num. laser trials	Num. mice	Num. sessions	Total trial count
V1	cue, 1st quarter (0 – 50 cm)	21,741	692	7	189	22,433
V1	cue, 2nd quarter (50 – 100 cm)	21,646	686	7	189	22,332
V1	cue, 1st half (0 – 100 cm)	19,753	540	7	158	20,293
V1	cue, 3rd quarter (100 – 150 cm)	21,580	680	7	185	22,260
V1	cue, 2nd 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, 1st quarter (0 – 50 cm)	18,372	589	7	159	18,961
mV2	cue, 2nd quarter (50 – 100 cm)	18,495	577	7	162	19,072
mV2	cue, 1st half (0 – 100 cm)	14,446	367	3	121	14,813
mV2	cue, 3rd quarter (100 – 150 cm)	18,429	594	7	160	19,023
mV2	cue, 2nd 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, 1st quarter (0 – 50 cm)	15,929	403	6	133	16,332
PPC	cue, 2nd quarter (50 – 100 cm)	15,993	429	6	134	16,422
PPC	cue, 1st half (0 – 100 cm)	9,006	615	9	63	9,621
PPC	cue, 3rd quarter (100 – 150 cm)	15,950	419	6	134	16,369
PPC	cue, 2nd 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, 1st quarter (0 – 50 cm)	15,923	422	6	134	16,345
RSC	cue, 2nd quarter (50 – 100 cm)	16,064	418	6	136	16,482
RSC	cue, 1st half (0 – 100 cm)	11,736	925	10	86	12,661
RSC	cue, 3rd quarter (100 – 150 cm)	15,952	432	6	134	16,384
RSC	cue, 2nd 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, 1st quarter (0 – 50 cm)	16,147	455	4	126	16,602
Posterior	cue, 2nd quarter (50 – 100 cm)	16,287	482	4	128	16,769
Posterior	cue, 1st half (0 – 100 cm)	14,439	355	3	119	14,794
Posterior	cue, 3rd quarter (100 – 150 cm)	16,252	474	4	127	16,726
Posterior	cue, 2nd 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, 1st quarter (0 – 50 cm)	21,787	703	7	191	22,490
mM2	cue, 2nd quarter (50 – 100 cm)	21,615	692	7	187	22,307
mM2	cue, 1st half (0 – 100 cm)	19,972	588	7	161	20,560
mM2	cue, 3rd quarter (100 – 150 cm)	21,766	674	7	190	22,440
mM2	cue, 2nd 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, 1st quarter (0 – 50 cm)	16,025	426	6	135	16,451
aM2	cue, 2nd quarter (50 – 100 cm)	16,023	420	6	135	16,443

aM2	cue, 1st half (0 – 100 cm)	20,338	576	7	166	20,914
aM2	cue, 3rd quarter (100 – 150 cm)	16,046	425	6	135	16,471
aM2	cue, 2nd 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, 1st quarter (0 – 50 cm)	13,894	369	6	117	14,263
M1	cue, 2nd quarter (50 – 100 cm)	13,938	382	6	119	14,320
M1	cue, 1st half (0 – 100 cm)	14,806	387	3	125	15,193
M1	cue, 3rd quarter (100 – 150 cm)	14,009	364	6	121	14,373
M1	cue, 2nd 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, 1st quarter (0 – 50 cm)	16,255	475	4	126	16,730
Frontal	cue, 2nd quarter (50 – 100 cm)	16,178	485	4	128	16,663
Frontal	cue, 1st half (0 – 100 cm)	14,410	369	3	120	14,779
Frontal	cue, 3rd quarter (100 – 150 cm)	16,065	483	4	123	16,548
Frontal	cue, 2nd 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

Table S1. Numbers of mice, sessions and trials for each of the 54 experimental conditions, related to Figure 1.
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.

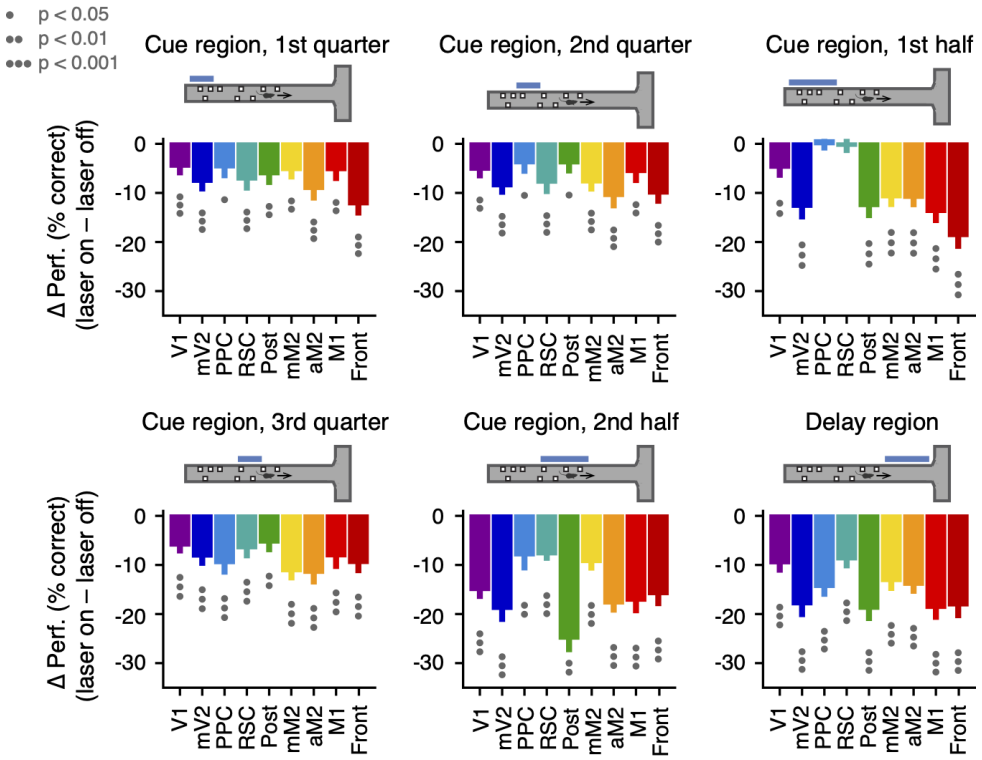

Fig. S1. Effects of subtrial inactivations on overall performance during all 54 area-epoch combinations, related to Figure 1.
Each panel shows inactivation-induced change in overall % correct performance for each inactivation epoch. Error bars: S.D. across 10,000 bootstrapping iterations. Circles indicate significance according to the captions on the leftmost panel.

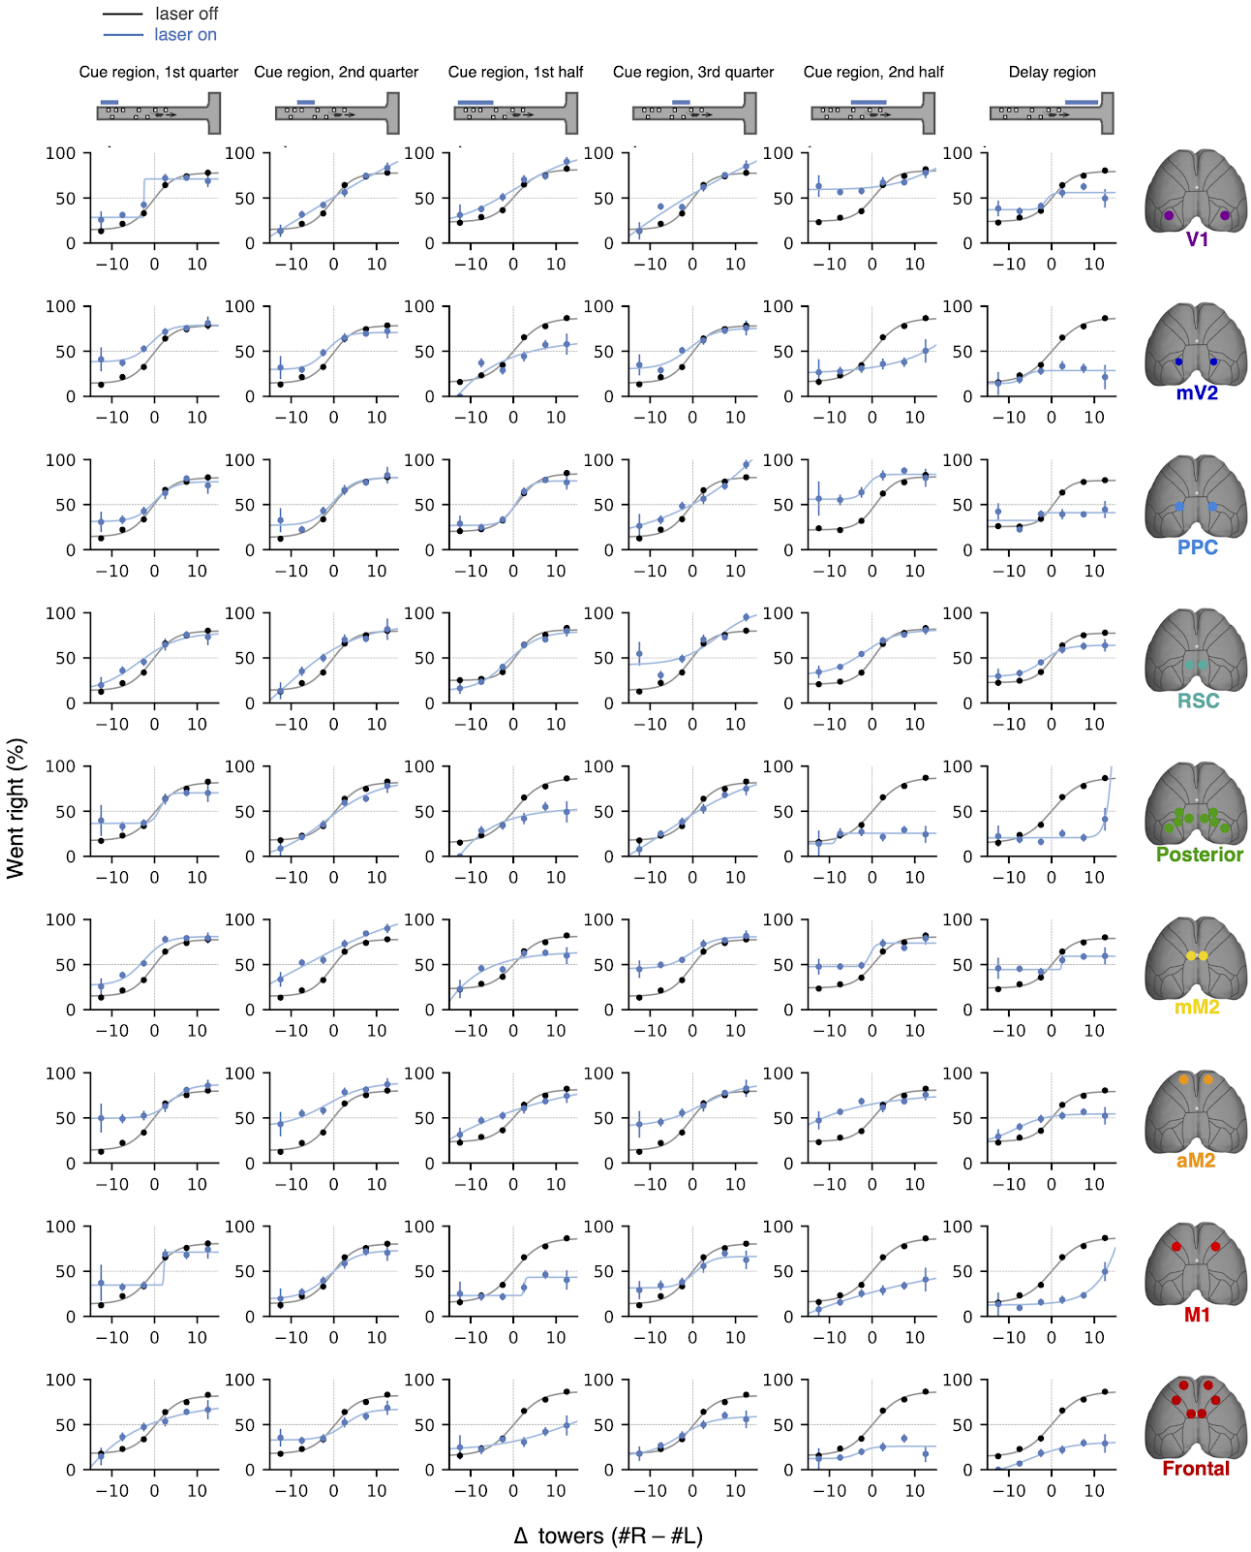

759 **Fig. S2. Effects of subtrial inactivations on psychometric functions during all 54 area-epoch**
760 **combinations, related to Figure 1.**
761 Black lines show control trials and blue lines show inactivation trials. Error bars, binomial confidence
762 intervals. Lines are best-fitting psychometric functions.

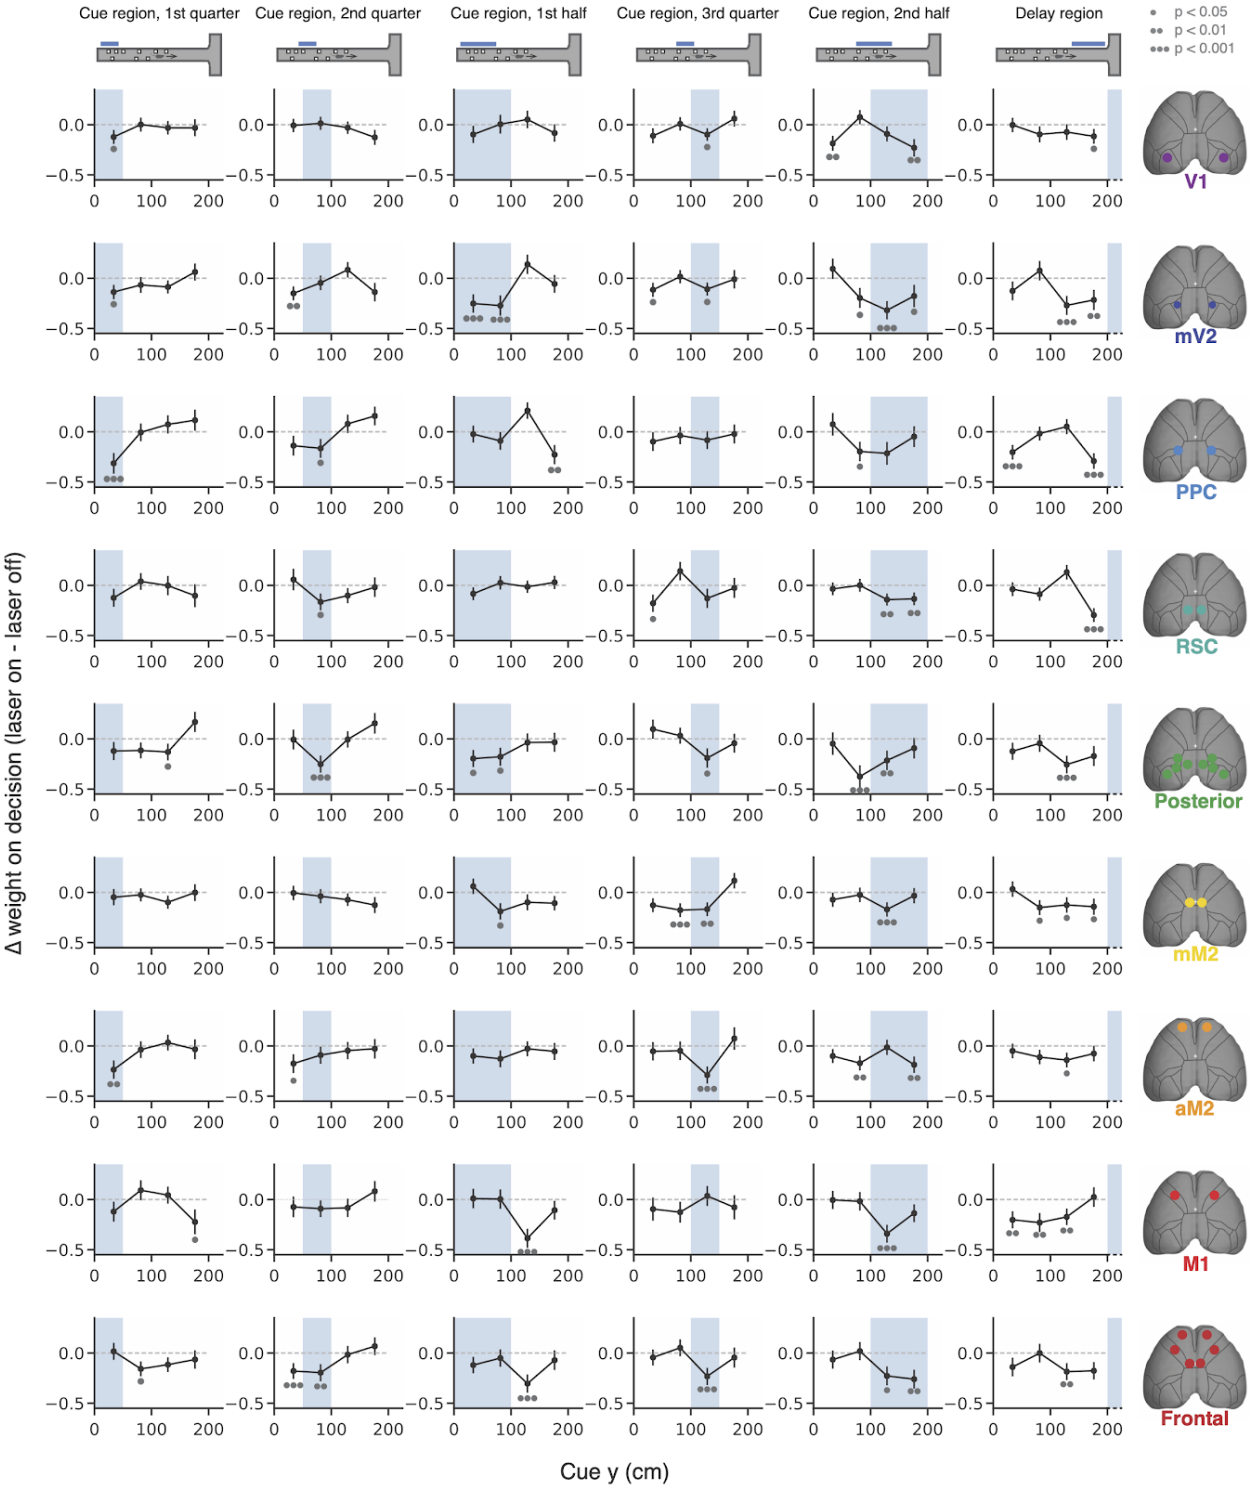

Fig. S3. Effects of subtrial inactivations on evidence-weighting functions during all 54 area-epoch combinations, related to Figure 2.
Black lines show the inactivation-induced change in evidence weights (laser on – laser off), and shaded areas indicate inactivation epoch. Error bars, S.D. across 10,000 bootstrapping iterations. Gray circles indicate statistical significance according to the caption on top.

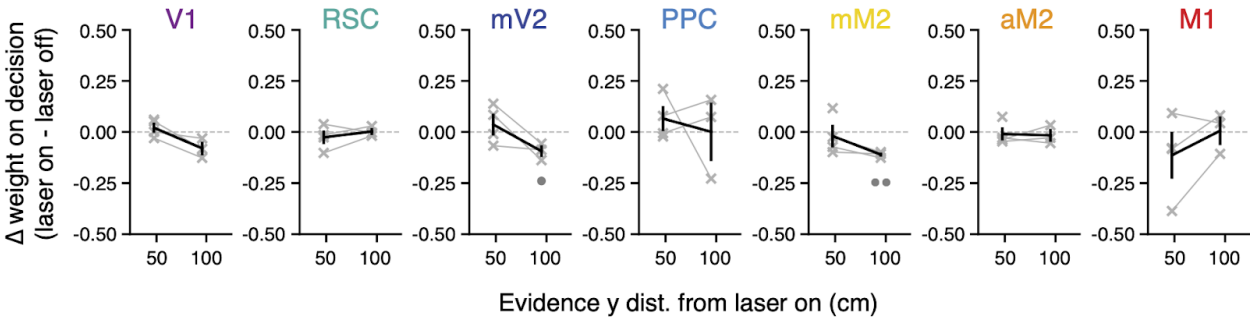

Fig. S4. Little effect on evidence weighting after laser offset, related to Figure 2.
Laser-offset-aligned changes in evidence-weighting curves for each bilaterally targeted area (laser on - laser off). Thin gray lines, individual inactivation epochs (n = 4 for y = 50, 3 for y = 100). Thick black lines, average across conditions. Error bars, ± SEM across experimental conditions. Circles below the lines indicate statistical significance (one circle: $p < 0.05$, two circles: $p < 0.01$; one-sided paired t test vs. zero, corrected for multiple comparisons).

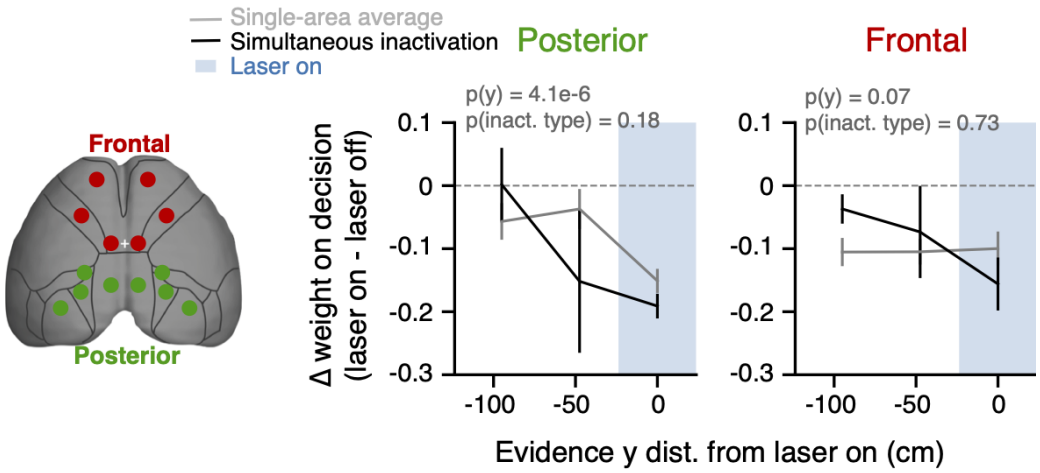

Fig. S5. Comparison between simultaneous multi-area inactivation and average across the corresponding individually inactivated areas, related to Figure 2.
Average laser-onset-aligned changes in evidence-weighting curves for each set of experimental conditions (caption on top). Error bars, ± SEM across experimental conditions. Shaded areas, laser on periods. P-values on top are from a two-way ANOVA with repeated measures, factors inactivation type (simultaneous vs. individual) and y position. There were no differences between simultaneous and individual inactivation of either frontal or posterior areas. On the other hand, effects depended on the distance of evidence from laser onset for posterior but not frontal areas.

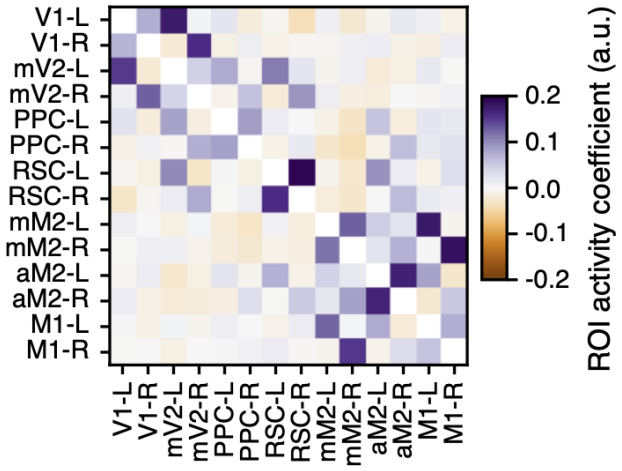

Fig. S6. GLM coupling coefficients for ROI activity predictors, related to Figure 3. For each ROI (rows), shown are the average coefficients (n=6 mice) for the predictors consisting of zero-lag activity of the other simultaneously imaged ROIs. Note that the diagonal elements are not defined since zero-lag self-activity predictors were not in the model. Data from the somatosensory cortex ROI has been omitted for symmetry with the inactivation data. L and R indicate left and right hemispheres, respectively.