

1 **Neural population dynamics underlying evidence accumulation in multiple rat**
2 **brain regions**

3

4 Brian DePasquale^{1,4}, Carlos D. Brody^{1,2,*}, Jonathan W. Pillow^{1,3,*}

5

6 ¹Princeton Neuroscience Institute,

7 ²Howard Hughes Medical Institute,

8 ³Department of Psychology,

9 Princeton University, Princeton NJ, USA

10

11 ⁴ Present address: Department of Biomedical Engineering,

12 Boston University, Boston MA, 02215

13

14 *These senior authors contributed equally to this work

15

16 **Correspondence:** B. DePasquale (depasquale@princeton.edu), C. Brody (brody@princeton.edu),
17 and J. Pillow (pillow@princeton.edu)

18

19 **Keywords:** latent variable models, decision making, evidence accumulation

20

21 **Abstract**

22

23 Accumulating evidence to make decisions is a core cognitive function. Previous studies have tended
24 to estimate accumulation using either neural or behavioral data alone. Here we develop a unified
25 framework for modeling stimulus-driven behavior and multi-neuron activity simultaneously. We
26 applied our method to choices and neural recordings from three rat brain regions — the posterior
27 parietal cortex (PPC), the frontal orienting fields (FOF), and the anterior-dorsal striatum (ADS) —
28 while subjects performed a pulse-based accumulation task. Each region was best described by a
29 distinct accumulation model, which all differed from the model that best described the animal's
30 choices. FOF activity was consistent with an accumulator where early evidence was favored while
31 the ADS reflected near perfect accumulation. Neural responses within an accumulation framework
32 unveiled a distinct association between each brain region and choice. Choices were better predicted
33 from all regions using a comprehensive, accumulation-based framework and different brain regions
34 were found to differentially reflect choice-related accumulation signals: FOF and ADS both reflected
35 choice but ADS showed more instances of decision vacillation. Previous studies relating neural data
36 to behaviorally-inferred accumulation dynamics have implicitly assumed that individual brain
37 regions reflect the whole-animal level accumulator. Our results suggest that different brain regions
38 represent accumulated evidence in dramatically different ways and that accumulation at the whole-
39 animal level may be constructed from a variety of neural-level accumulators.

40

41

42 Introduction

43

44 Accumulation of evidence is a critical process underlying decision-making in complex environments
45 where relevant information is distributed across time. Choice data from evidence accumulation
46 tasks (e.g., Brunton et al., 2013; Raposo et al., 2012; Sanders and Kepecs, 2012) have allowed for the
47 development of sophisticated models of animals' accumulation strategies (e.g., Bogacz et al., 2006;
48 Brunton et al., 2013; Genkin et al., 2021; Gold and Shadlen, 2007; Ratcliff et al., 2016; Ratcliff and
49 McKoon, 2008; Shinn et al., 2020; Wiecki et al., 2013). In parallel, neural correlates of accumulated
50 evidence have been found in a wide variety of brain regions (e.g., Brody and Hanks, 2016;
51 Churchland et al., 2011; Ding and Gold, 2010; Erlich et al., 2011; Gold and Shadlen, 2007; Hanks et
52 al., 2015; Huk and Shadlen, 2005; Kim and Shadlen, 1999; Mante et al., 2013; Ratcliff et al., 2003;
53 Roitman and Shadlen, 2002; Shadlen and Newsome, 2001; Yartsev et al., 2018) and methods have
54 been developed to describe the statistical relationship between neural activity and accumulated
55 evidence (e.g., Aoi et al., 2020; Beck et al., 2008; Churchland et al., 2011; Hanks et al., 2015; Latimer
56 et al., 2015; Latimer and Freedman, 2021; Park et al., 2014; Zoltowski et al., 2020, 2019).

57

58 Obtaining a comprehensive account of how stimulus-influenced accumulated evidence underlies
59 neural activity and subject choice remains an open problem. For example, few analysis methods
60 which use precise spike timing information take into account the timing of stimulus information or
61 use choice data directly (e.g., Latimer et al., 2015). Likewise few methods that use the precise
62 timing of stimulus information to infer accumulated evidence use neural responses directly (e.g.,
63 Hanks et al., 2015). To address this gap, we developed a framework for inferring probabilistic
64 evidence accumulation models jointly from choice data, neural activity, and precisely controlled
65 stimuli.

66

67 A complete understanding of decision-making necessitates models that can comprehensively
68 combine stimuli, neural activity, and behavior. The evidence accumulation process inferred from
69 behavioral data alone need not correspond to the accumulation process that best matches data
70 from a single brain region; behavior is the result of interactions between multiple brain regions. For
71 example, two brain regions, one favoring accumulation of early evidence (e.g., an unstable
72 accumulator) and the other favoring accumulation of late evidence (e.g., a leaky accumulator) could
73 together support stable behavior-level accumulation. By fitting accumulator models to neural data
74 from multiple brain regions and to subject choice data, we gained the opportunity to probe for the
75 first time whether different brain regions reflect the same, or different, accumulation processes and
76 how those individual processes correspond to the animal's overall behavior.

77

78 We applied our model to choices and neural responses from three brain regions known to be
79 involved in evidence accumulation while animals perform a pulse-based evidence accumulation
80 task. A single variable representing accumulated evidence, shared across neurons within a brain
81 region, accurately accounted for both neural and choice data. We identified distinct signatures of
82 accumulation reflected in each brain region, all of which differed from the accumulation model that
83 best described behavior, supporting the idea that whole-organism accumulation likely results from
84 multiple accumulation processes. Prior analysis of these data found that the anterodorsal striatum
85 (ADS) represented accumulated evidence in a graded manner (Yartsev et al., 2018) while the
86 frontal orienting fields (FOF) represented choice more categorically (Hanks et al., 2015). Our
87 analysis confirms the ADS as a veracious representation of accumulated evidence while offering a
88 more nuanced view of the FOF: the accumulation model that best described FOF activity was
89 dynamically unstable, producing neural responses that looked like a categorical representation of
90 choice but that were in fact unstable accumulators sensitive to early stimulus information.
91 Additionally, we analyzed recordings from the posterior parietal cortex (PPC), a brain region long

92 studied in connection to evidence accumulation (Hanks et al., 2015; Roitman and Shadlen, 2002;
93 Shadlen and Newsome, 2001), where we identified neural correlates of graded evidence
94 accumulation, albeit more weakly than in the ADS.

95
96 Incorporating neural activity into accumulation models reduced the uncertainty in the moment-by-
97 moment value of accumulated evidence when compared to models fit only to animal choices. This
98 reduction in uncertainty led to a more refined picture of the moment-by-moment value of
99 accumulated evidence, which made the model more informative about what choice the animal
100 intended to make. Our model allowed us to implement a novel analysis to examine how subject
101 provisional choice changed during individual trials, commonly referred to as ‘changes of mind’
102 (Boyd-Meredith et al., 2022; Kiani et al., 2014; Peixoto et al., 2021), that revealed extensive choice
103 vacillation reflected in ADS activity and greater choice certainty reflected in FOF activity.

104
105 Broadly, our framework offers a unified, mechanistic, and probabilistic description of the moment-
106 by-moment accumulation process that underlies decision-making. Our flexible framework offers a
107 computationally efficient method for identifying a key normative decision-making model using
108 multiple types of data, and can easily accommodate simultaneous recordings from many neurons or
109 recordings performed sequentially over many days. It provides a platform for quantitatively
110 characterizing choice-related information in neural responses and can be used to understand how
111 different brain regions implement an accumulation strategy.

112 113 **Results**

114
115 We analyzed behavioral and neural data from rats trained to perform a perceptual decision-making
116 task (Brunton et al., 2013). Rats listened to two simultaneous series of randomly timed auditory
117 clicks, one from a speaker on the left and one from a speaker on the right. After the end of the click
118 train, the rat had to orient to the side with a greater number of clicks to receive a reward (**Figure**
119 **1A**).

120
121 We analyzed behavioral choice data and electrophysiological neural recordings from 11 rats. In
122 total, we analyzed 37,179 behavioral choices and 141 neurons from three brain areas—the
123 posterior parietal cortex (PPC), the frontal orienting fields (FOF), and the anterior-dorsal striatum
124 (ADS). Prior electrophysiological and lesions studies have shown that these brain regions play a key
125 role in evidence accumulation (Ding and Gold, 2013, 2010; Erlich et al., 2015, 2011; Gold and
126 Shadlen, 2007, 2000; Hanks et al., 2015; Huk and Shadlen, 2005; Kim and Shadlen, 1999; Mante et
127 al., 2013; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001; Yartsev et al., 2018).

128
129 Data were collected after the animals were well-trained and exhibiting a high-level of performance
130 (Brunton et al., 2013; Hanks et al., 2015; Yartsev et al., 2018); these data were collected as part of
131 two earlier studies and have been previously analyzed (Hanks et al., 2015; Yartsev et al., 2018).
132 Data were subject to a selection criterion for inclusion in our study. We selected neurons with
133 significant tuning for choice during the stimulus period (two-sample t-test, $p < 0.01$) because choice
134 tuning is a prerequisite for reflecting accumulation-like signals. Information about the data is
135 summarized in **Table 1**. Once tuning significance was determined, our dataset consisted of 68
136 neurons from FOF, with 7,382 behavioral choices recorded from five rats over 46 behavioral
137 sessions; 25 neurons from PPC, with 9,037 behavioral choices from three rats over 24 sessions; and
138 48 neurons from ADS, with 10,760 behavioral choices from three rats over 27 behavioral sessions.

139 140 **A latent variable model of behavioral choice and neural activity**

141
 142 One of the most common normative models of the internal mental processes that underlie evidence
 143 accumulation is the drift-diffusion to bound model (DDM; **Figure 1B**; Bogacz et al., 2006; Brunton
 144 et al., 2013; Gold and Shadlen, 2007; Ratcliff and McKoon, 2008). While previous work has tended
 145 to fit this model (either explicitly or implicitly) using either choice data (e.g., Brunton et al., 2013;
 146 Chandrasekaran and Hawkins, 2019; Gold and Shadlen, 2007; Ratcliff et al., 2016; Shinn et al., 2020;
 147 Wiecki et al., 2013; Zylberberg et al., 2016) or neural response data (e.g., Bollimunta et al., 2012;
 148 Brody and Hanks, 2016; Churchland et al., 2011; Ditterich, 2006; Genkin et al., 2021; Hanks et al.,
 149 2015; Howard et al., 2018; Latimer et al., 2015; Zoltowski et al., 2020, 2019), here we seek to jointly
 150 model the relationship between accumulated evidence, choices, and neural activity.

151
 152 The essence of our model is to describe a DDM based accumulation process driven by sensory
 153 stimuli following (Brunton et al., 2013) and relate the latent accumulation process to both neural
 154 responses and the rat's choice. Previous results have shown that this model is sufficiently flexible to
 155 accommodate the various behavioral strategies rats exhibit while performing this task (Brunton et
 156 al., 2013). The resulting model has a single latent variable, denoted $a(t)$, that evolves in time and
 157 represents the current, inner mental representation of the evidence in support of a left or right
 158 choice at each moment in time. This latent variable is shared by the neurons within a region (except
 159 where explicitly noted), so that each neuron's time-varying firing rate is a function of $a(t)$ on each
 160 trial. The key distinction of our approach is that the accumulator variable $a(t)$ drives both choices
 161 and neural activity, as described below.

162
 163 Formally, the temporal evolution of the latent evidence $a(t)$ is governed by:

$$164 \quad da = \lambda a dt + \Delta(t) dt + \sigma_a dW + \sigma_s \Sigma(t) \eta dt, \quad (\text{Equation 1})$$

165
 166 where da is the amount $a(t)$ changes in a time dt . λ is a leak parameter. $\Delta(t)$ and $\Sigma(t)$ indicate the
 167 difference and sum, respectively, in the number of right and left sensory clicks at time t , after the
 168 magnitude of the clicks has been adapted based on recent stimulus history (see parameters
 169 governing adaptation below, and Methods for additional detail). $\sigma_a dW$ is a diffusive Gaussian noise
 170 process (or Weiner process) with scaling σ_a . $\sigma_s \Sigma(t) \eta$ is additive Gaussian noise induced by each
 171 click input, where $\sigma_s \Sigma(t)$ is the standard deviation of the click noise and η is a Gaussian random
 172 variable with a mean of zero and standard deviation 1.

173
 174 If $a(t)$ becomes greater in magnitude than a symmetric boundary with magnitude B (**Figure 1B**,
 175 dotted lines), then $da = 0$, and accumulation ceases for the remainder of the trial. To illustrate, the
 176 blue trajectory in **Figure 1B** crosses the boundary B roughly one-third of the way through the trial,
 177 and thus remains constant thereafter.

178
 179 The four terms of **Equation 1** each account for specific ways $a(t)$ might reflect accumulated
 180 evidence. The first two terms are designed to account for deterministic (non-random) dynamics
 181 exhibited by $a(t)$. The first term specifies how recent values of $a(t)$ influence future values and is
 182 governed by λ that determines the timescale of this effect. Positive values of λ correspond to
 183 unstable dynamics so that $a(t)$ grows exponentially. In this setting, early clicks have greater
 184 influence on $a(t)$ than recent clicks, because their impact grows with time. By contrast, negative
 185 values of λ correspond to leaky dynamics. In this setting, early clicks have a weaker influence on
 186 $a(t)$ than recent clicks because the impact of early clicks decays with time. When λ equals zero, the
 187 sensory clicks are perfectly integrated. Previous results have shown that rats exhibit a range of
 188 accumulation strategies spanning these values of λ (Brunton et al., 2013). The second term, $\Delta(t) dt$,
 189 specifies how the click stimulus is incorporated into $a(t)$. Because the task requires reporting
 190

191 whether there was a greater number of left or right clicks, only the total click difference is required
192 to correctly perform it.

193
194 To account for stochasticity in the accumulation dynamics, the model also contains two forms of
195 noise in $a(t)$. The first noise term, $\sigma_a dW$, corresponds to diffusive noise that corrupts $a(t)$
196 continuously in time. The final term, $\sigma_s \Sigma(t) \eta dt$, introduces noise into $a(t)$ that is proportional to the
197 total number of clicks that occur at a given moment. The sum of clicks $\Sigma(t)$ is included so that the
198 magnitude of the noise increases depending on the number of sensory clicks experienced at time t .
199 **Figure 1B** illustrates the effects of these two noise terms: although the sensory inputs and leak are
200 identical for both blue and black trajectories of $a(t)$, differences in noise lead the two trajectories to
201 diverge so that one hits the boundary $+B$ while the other remains sub-threshold and continues to
202 integrate the sensory stimulus.

203
204 To model animal choices, we assume that the accumulation variable $a(t)$ directly governs the
205 animal's choice on each trial. Specifically, we describe the probability of a rightward choice as
206 depending on $a(T)$, the accumulated evidence at the end of the stimulus period T , using a step
207 function with 'lapses'. With probability γ the animal picks one of the two sides without considering
208 the stimulus, referred to as a 'lapse'. With probability $(1 - \gamma)$ the animal does not lapse, and makes
209 a rightward choice if $a(T) > c$ and a leftward choice if $a(T) < c$, where c denotes the choice criterion.
210 This model can be expressed as

$$P(d = R) = \gamma/2 + (1 - \gamma)H(a_T - c) \quad (\text{Equation 2})$$

211
212
213
214 where $d \in \{L, R\}$ is the decision variable and $H(\cdot)$ is the Heaviside step function. As described
215 above, when $a(t)$ crosses the decision bound B a choice commitment is made, either to the left or
216 the right, and no further evidence accumulation occurs. Previous work has found that
217 parameterizing choice this way creates a model that is sufficiently flexible to describe animals'
218 choice (Brunton et al., 2013) while remaining as simple as possible.

219
220 To model spike train data, we describe the time-varying firing rate of each neuron as a soft-rectified
221 linear function of the same accumulated evidence variable $a(t)$:

$$f_{\theta_n}(a(t)) = \text{softplus}(\theta_n a(t) + \theta_n^0(t)), \quad (\text{Equation 3})$$

222
223
224
225 where n indexes neurons, the softplus function (**Figure 1B**) is given by $\text{softplus}(x) = \log(1 + \exp(x))$,
226 and θ_n denotes the slope of the linear relationship between $a(t)$ and neuron n 's firing rate. The
227 slope parameter, θ_n , is fit separately for each neuron. A time-varying offset, $\theta_n^0(t)$, is included to
228 capture time-varying changes in firing rate that do not depend on $a(t)$ (see Methods). The spikes of
229 each neuron are modeled as a Poisson process with a time-dependent conditional intensity function
230 $f_{\theta_n}(a(t))$. The softplus function (smooth rectified linear function) was used to ensure the expected
231 firing rate was positive, and was selected because it is the simplest function to achieve this goal, and
232 also based on prior success in similar studies (e.g., Latimer et al. 2015).

233
234 We refer to the set of all parameters that govern $a(t)$, and its relationship to the neural activity and
235 choice data as $\theta = \{\sigma_i, B, \lambda, \sigma_a, \sigma_s, \phi, \tau_\phi, \theta_{1:N}, c, \gamma\}$, where σ_i is the variance of $a(t)$ at the start of the
236 trial, and ϕ and τ_ϕ determine how the magnitude of each click is adapted based on the timing of
237 recent clicks (see Methods). We fit θ separately for each brain region using maximum likelihood
238 (see Methods). Maximizing the likelihood of the data requires computing the temporal evolution of
239 the probability distribution of $a(t)$ over the duration of a single trial, for all trials, and computing the

240 probability of the observed spikes and choices under this distribution. The dynamics of this
241 probability distribution can be expressed using the Fokker-Planck equation, and previous work has
242 developed methods for numerically solving it (Brunton et al., 2013; see Methods). We refer to the
243 value of θ that maximizes the likelihood of the data as $\hat{\theta}$. We verified that our method was able to
244 recover the parameters that generated synthetic physiologically-relevant spiking and choices data (
245 **Figure 1 — figure supplement 1**), and that parameter recovery was robust across a range of
246 parameter values (Figure 1 — figure supplement 2).

247

248 **Shared accumulator model captures neural responses and choices**

249

250 We fit the model separately to data from each brain region. To verify model fits were consistent
251 with data, we compared the peri-stimulus time histograms (PSTHs; **Figure 2A & B**) and
252 psychometric curves (**Figure 2C**) of the empirical data to synthetic data simulated from the fitted
253 model for each brain region. The PSTH of most neurons showed a characteristic choice preference
254 that increased over time, consistent with accumulation. The model was able to capture this (**Figure**
255 **2A**). The model provided an accurate account of mean responses in all three brain areas (**Figure**
256 **2B**), with a median R^2 of 0.91, 0.68, and 0.87 for the FOF, PPC, and ADS respectively (**Figure 2B**,
257 colored lines). **Figure 2C** shows a comparison between true psychometric curves and the
258 psychometric curve of the fitted model, confirming that the model also accounted for
259 psychophysical choice behavior. (R^2 : 0.99 - FOF; 0.99 - PPC; ADS - 0.97; see Methods for details).
260 These analyses confirm that a shared accumulator model for each brain region is sufficient to
261 capture the animals' choice sensitivity to the stimulus and strength of accumulated evidence
262 reflected in each neuron's response.

263

264 **Different regions reflect different accumulator models, which all differ from model** 265 **describing behavior**

266

267 The primary motivation of our study was to learn accumulator models that incorporate precise
268 stimulus timing information and describe the animal's choices and temporally structured neural
269 activity. Previous efforts only modeled choices using stimulus-timing information (Brunton et al.,
270 2013) or modeled neural activity without choices for tasks without detailed stimulus-timing
271 information (Latimer et al., 2015; Zoltowski et al., 2019). We refer to our model that describes both
272 neural activity and choices as the 'joint neural-behavioral model' or the 'joint model'. We compared
273 the joint neural-behavioral model to a model where only the stimulus is used to model the animal's
274 choice (i.e., neural activity is not used). To fit such a 'choice-only' accumulator model we fit the
275 same latent variable model using only choice data (see Methods).

276

277 **Figure 3A** shows the maximum likelihood parameters for the joint and choice-only accumulator
278 models for each brain region. Neural data was not used for the choice model so brain region
279 designates the cohort of animals from which the choice data was taken. We stress that because of
280 this, each fitted choice model uses different behavioral choice data, and thus the fitted parameters
281 vary from fitted model to fitted model. Both fitted models exhibited strong adaptation ($\phi \ll 1$)
282 consistent with prior work fitting choice accumulator models (Brunton et al., 2013). This indicates
283 that a stimulus pulse that occurs in rapid succession following other pulses has a smaller effect on
284 $a(t)$ than an isolated pulse. Each model was impacted by different forms of noise: choice models
285 exhibited small diffusive noise ($\sigma_a \approx 0$) and large stimulus noise ($\sigma_s \gg 1$), consistent with earlier
286 findings, while joint models exhibited large diffusive noise ($\sigma_a > 0$) and large initial variability in
287 $a(t)$ ($\sigma_i \gg 0$). The effect of these different parameters can be seen in **Figure 3B**: choice models
288 have smaller initial variance and more variability when clicks arrive, while joint accumulator
289 models have larger initial variance and diffusive noise. Large initial variance in the joint model

290 likely reflects variability in neural responses prior to stimulus onset (Churchland et al., 2010).
291 Strong accumulation noise in the joint model was also found when the negative binomial
292 distribution, a more flexible observation model, was used, suggesting that this finding was not
293 sensitive to the Poisson observation model (**Figure 3 — figure supplement 1**). Differences in
294 diffusive noise between the joint and choice-only models suggest that accumulation dynamics
295 underlying neural activity is impacted by noise that is resolved at the level of a behavioral
296 accumulator model.

297
298 We also compared the best-fit parameters across the three, separately fit, brain regions (**Figure**
299 **3A**). We focus on one of the most salient differences — the leak or instability parameter λ . Although
300 there was no significant difference in the value of λ across the cohorts of animals in the choice-only
301 model, we found substantial differences across brain regions in the joint model fits (**Figure 3A**).
302 The PPC and ADS data were best fit by leaky accumulator models ($\lambda < 0$). Surprisingly however, the
303 FOF data was best described by a model with unstable accumulation dynamics ($\lambda > 0$) meaning
304 that the model's accumulator (and thus firing rates) are more strongly affected by early stimulus
305 clicks. The stronger weighting of earlier clicks was compounded further by the low accumulation
306 bound of the model that best described FOF data. Such a low bound, in conjunction with unstable
307 accumulation, causes $a(t)$ to stop evolving early in the trial (**Figure 3B**). This results in a
308 phenomenon known as 'primacy encoding', in which early

309
310 clicks more strongly impact the animal's choice while later clicks are ignored. We confirmed this
311 finding in the FOF using a generalized linear model (GLM; see Methods & **Figure 3 — figure**
312 **supplement 2**). This result is consistent with previous work suggesting that the FOF has a
313 categorical representation of $a(t)$ (Hanks et al., 2015). We expand on these findings in light of other
314 studies of the FOF in the Discussion. Collectively, these results indicate that all three brain regions
315 were best described by accumulator models that differed in their best fitting parameters (and thus
316 exhibit dramatically different accumulation dynamics) and that each region's data was likewise best
317 described by a model that differed from that which best described accumulation at the level of the
318 animal's choice.

319 320 **ADS is better described by multiple, independent accumulators**

321
322 Our model describes the spiking activity of a population of simultaneously recorded neurons as
323 relying on a single shared latent variable. To assess whether this is indeed the best description of
324 the data, we compared it to an 'independent noise accumulator model' where each neuron is driven
325 by an accumulator with its own independent noise (**Figure 4A**; Methods). It is worth emphasizing
326 that the independent noise model is identical to the shared noise model in the way it is
327 parameterized (i.e. number and form of the model parameters) but only differs in the structure of
328 the latent accumulation noise. If trial-to-trial spiking covariation is produced by temporal
329 covariation in the accumulator due to noise, the independent noise model (which does not share
330 this covariation) should not account for the data as well, suggesting that correlations in the data can
331 be attributed to correlated diffusive noise reflected in the shared model. We fit the parameters of
332 the independent noise model using the same optimization method but with a different log-
333 likelihood function (see Methods). Because the independent noise model contained multiple
334 accumulators (one for each neuron), the animal's choice was modeled differently than for the
335 shared noise model (see Methods). We focused on the FOF and ADS datasets because they
336 contained a sufficient number of simultaneously recorded neurons to make this comparison (**Table**
337 **1**). The maximum likelihood parameters for the two models for both regions were similar (**Figure 4**
338 **— figure supplement 1**), except for the initial accumulator variance parameter which differed
339 significantly.

340
341 We used 5-fold cross-validation to determine which model better described each data set.
342 Comparing cross-validated log-likelihood, we found that the independent noise model provided a
343 better description of choices and neural activity from ADS, while the shared noise model provided a
344 slightly better description of FOF data (**Figure 4B**). This finding supports the conclusion that neural
345 responses within the ADS reflect independent accumulation processes, while neurons in the FOF
346 reflect a single latent accumulator. Although ADS datasets with 4 or more neurons provided the
347 primary contribution to these results (**Figure 4 — figure supplement 2 A**), when the number of
348 neurons in ADS datasets were subsampled to match the maximum number of neurons in FOF
349 sessions (3 neurons), the ADS recordings still favored an independent noise accumulator model
350 (**Figure 4 — figure supplement 2 B**). We fit the shared noise and independent noise model to
351 neural data only (excluding choice data) and found consistent results (**Figure 4 — figure**
352 **supplement 2 D**) suggesting this difference is not due to contributions from the animal's choice,
353 which was modeled differently in each model (see above).

354
355 To further examine this result, we computed the 'shuffle corrected' cross-correlation function
356 (Methods; Perkel et al., 1967; Smith and Kohn, 2008) for all pairs of simultaneously recorded
357 neurons to examine spiking covariation in the empirical data and synthetic data from the fit models
358 (**Figure 4C & D**). To shuffle-correct, we took the raw cross-correlation and subtracted the cross-
359 correlation of the PSTHs of two neurons (for left and for right trials separately). This provides a
360 measure of the neurons' correlation beyond what is to be expected from the PSTHs (i.e., **Figure 2A**).

361
362 Synthetic data of both models captured trends in the shuffle-corrected cross-correlation function at
363 slower time scales but failed to capture fluctuations on short time scales. Across all pairs of
364 simultaneously recorded neurons (70 pairs in total), we found that the shared and independent
365 noise accumulator models provided approximately equally accurate fits to the shuffle corrected
366 cross-correlations (mean r of 0.55 for shared model and 0.57 for independent noise model for FOF;
367 0.63 for shared model and 0.60 for independent noise model for ADS). This shows that both models
368 capture correlations in trial-to-trial neural responses beyond those accounted for by the PSTH.
369 These correlations likely arise from trial-to-trial differences in the exact sequence of clicks, which
370 are not reflected in the PSTH for left- or right-choice trials. Although FOF weakly favored a shared-
371 noise model and ADS favored an independent-noise model (**Figure 4B**) the comparable ability for
372 each model to capture the shuffle-corrected cross-correlation function for each region suggests that
373 these correlations are primarily stimulus-induced and not a manifestation of non-stimulus induced
374 (i.e., 'noise') correlations, which are weak if present at all. Although these results suggest that each
375 model fits the data equally well, the results of **Figure 4B** suggest that the independent noise model
376 may be accounting for intricate features of the ADS data not reflected in the shuffle-corrected cross-
377 correlation function.

378
379 To validate that neural responses in the ADS weakly covary, as suggested by an independent noise
380 model, we computed a measure of response dimensionality known as the participation ratio
381 (Litwin-Kumar et al., 2017). The participation ratio is computed using the eigenvalues of the
382 covariance matrix of firing rates (Methods). If all firing rates are independent the eigenvalues will
383 all be equal and the participation ratio will equal the number of neurons. If the firing rates are
384 correlated such that some eigenvalues are small (or perhaps zero) the participation ratio will
385 reflect this and the dimensionality of the data will be less than the number of neurons. Consistent
386 with our modeling results, we found that responses in ADS had higher dimensionality than in FOF
387 (i.e., ADS exhibited less firing rate covariation) and that ADS sessions with greater dimensionality
388 were those that favored the independent noise model (**Figure 4 — figure supplement 2 C**).

389

390 **Neural data provides more information about accumulated evidence than choice**

391
392 Next we examined how neural data affected inferences about accumulated evidence. We computed
393 the posterior distribution over the accumulator variable $a(t)$ for the joint model, given choice data
394 only, or given neural and choice data. The posterior distribution combines information from
395 multiple sources — stimulus, choice, and neural activity — to offer a concise window into the
396 animal's internal state of evidence accumulation. **Figure 5A** shows the posterior distribution for
397 three example trials (one for each brain region) when only choice data was included and when both
398 choice and neural data were included. The choice data posterior was broad; a large set of $a(t)$
399 trajectories were all consistent with the animal's choice. However, when we considered both choice
400 and neural spiking activity, we obtained a substantially narrower distribution over $a(t)$, meaning
401 including neural data in the joint model offers greater confidence in the precise value of
402 accumulated evidence at each moment within a trial.

403 To quantify this difference, we computed the standard deviation of the two posteriors (**Figure 5B**).
404 For all brain regions, the median posterior standard deviation given neural data and choice was
405 substantially smaller than when conditioning only on choice (**Figure 5B**; median difference FOF:
406 0.46; PPC: 0.72; ADS: 2.23). This reduction in the posterior width increased with the number of
407 neurons (**Figure 5C**). The increased certainty about $a(t)$ provided by neural activity makes intuitive
408 sense: temporally specific spiking activity (e.g., in the middle of a trial) allows one to infer that $a(t)$
409 has increased in favor of a choice, whereas choice information can only offer certainty about the
410 range of $a(t)$ at the end of the trial.

411 **Joint neural-behavioral model improves choice decoding**

412
413 We designed our joint model with the expectation that combining choice data, neural responses,
414 and stimulus information within an accumulation framework would lead to greater insight into
415 decision-making than models that lacked these features. We tested this expectation by comparing
416 choice decoding accuracy of the joint model on single trials to models that used stimulus
417 information and only choice data or only neural data (see Methods). We found that choices could be
418 predicted more accurately under the joint model, which took into account the stimulus, neural
419 activity, and choices, than under the choice model, which used stimulus information and choices
420 alone. We quantified this improvement in test log-likelihood and percent correct (**Figure 6A**). The
421 joint model had higher test log-likelihood for choice data and choice prediction accuracy for all
422 three brain regions, with the joint model of FOF data showing an almost 50% improvement in test
423 log likelihood and a 6% increase in prediction accuracy. The posterior mean of the joint model and
424 the posterior mean of the choice model is shown in **Figure 6B** for three example trials. In all
425 examples, the joint model correctly predicted the choice the animal made (indicated by the arrow),
426 whereas the choice-only model failed because its prediction was based on the stimulus. This
427 increased performance derives from the choice-informative spiking information contained in the
428 posterior (**Figure 5**) that the choice model lacks.

429
430
431 If neural activity is highly correlated with the motor report (for example, activity from motor
432 neurons controlling orientation), we would expect the neural activity to be a good predictor of the
433 animal's choice. In such a case, a model that predicted choice without the framework of the DDM
434 accumulator but using neural activity, would have high accuracy. We compared our accumulator-
435 based joint model to a logistic regression model (i.e., Bernoulli generalized linear model, GLM)
436 which used the final accumulated click difference and the trial-summed spike count for each neuron
437 as regressors (Methods). Decoding under the joint accumulator model significantly outperformed
438 logistic regression (**Figure 6A, GLM**). The performance of the GLM did not depend strongly on the
439 time window considered: decoding of choice using spikes from the last 50 ms (**Figure 6A, GLM 50**

440 ms), 100 ms, 150 ms, 200 ms and 250 ms before a decision all performed similarly (**Figure 6 —**
441 **figure supplement 1**). This shows that the joint accumulation framework and the fine timescale
442 dynamics of the joint model captures features of the spike trains that are useful for predicting the
443 animal's choice, above and beyond the information carried by spike counts in particular time
444 windows before the choice.

445 **Putative changes of mind are common in ADS, rare in FOF**

446 The previous analysis illustrated how the joint accumulation framework, combined with temporally
447 precise neural responses, can accurately predict animal choices. Numerous studies have shown that
448 subjects making decisions based on noisy stimuli will vacillate before reporting a decision
449 (Kaufman et al., 2015; Kiani et al., 2014; Resulaj et al., 2009). Switches of a subject's provisional
450 decision have been referred to as 'changes of mind' (Boyd-Meredith et al., 2022; Peixoto et al.,
451 2021). We used our joint accumulator model to identify putative changes of mind from our neural
452 recordings, to examine how decision commitment is manifested in different brain regions. We
453 examined the temporal dynamics of the joint model posterior, conditioned on neural activity only,
454 to find putative changes of mind: moments when posterior mean crossed from one side of the
455 decision threshold to the other. We required that the conditioned posterior mean remained on one
456 side of the decision threshold for at least 50 ms before and after the crossing and achieved an
457 absolute magnitude greater than 2 during that 100 ms window (see Methods).
458

459 **Figure 6C** shows three example putative change of mind trials. We also plot the posterior mean of
460 the choice model (black) and the cumulative click difference (gray) for comparison. In all three
461 examples, the joint model posterior mean crossed the decision threshold, ending on the side
462 corresponding to the animal's choice. Sign changes in the cumulative click difference were rare, as
463 were putative change of mind events under the choice-only model, both of which could only be
464 caused by the stimulus (**Figure 6D**). In contrast, putative change of mind events were observed
465 frequently under the joint model for all three brain regions (**Figure 6D**). This shows that putative
466 change of mind events reflect information about the accumulator carried in neural activity. Putative
467 change of mind events were observed least frequently in the FOF and most frequently in the ADS
468 (**Figure 6D**); compounded by our initial finding, that different brain regions are best fit by different
469 accumulator models (**Figure 3**), these results further support the view that the decision making
470 dynamics in each brain region are fundamentally and consequentially different.
471

472 The animal's performance improved on putative change of mind event trials (fraction correct: FOF:
473 0.88 vs. 0.74; PPC: 0.87 vs. 0.74; ADS: 0.85 vs. 0.76; **Figure 6 — figure supplement 2 A**) and the
474 choice prediction of the joint model was also more accurate (fraction correct: FOF: 0.92 vs. 0.80;
475 PPC: 0.88 vs. 0.77; ADS: 0.88 vs. 0.78; **Figure 6 — figure supplement 2 B**), suggesting that the
476 decision making dynamics that give rise to these events primarily correct incorrect decision-
477 making dynamics early within a trial. Initial variability in the accumulation dynamics, as reflected in
478 neural responses, was found to be greater in both PPC and ADS (**Figure 3A**), regions for which
479 putative changes of mind were more likely (**Figure 6D**), consistent with this assumption.
480 Furthermore, putative change of mind events were more likely to occur at later moments in the
481 trial, usually not long before the stimulus ended (**Figure 6E**), consistent with the assumption that
482 they generally correct incorrect early-trial dynamics. To more firmly connect putative change of
483 mind events to the animal's behavior, we performed linear regression to compare the time of the
484 event relative to the end of the stimulus to the response latency (**Figure 6F**). We found a
485 statistically significant effect for the PPC and the ADS (PPC: $p < 0.003$; ADS: $p < 0.0008$; two-sided t -
486 test), which both showed a slower response time when a change of mind event occurred closer to
487 the end of the stimulus. These results illustrate the potential of our framework for uncovering
488

490 putative covert changes of mind within neural activity, and demonstrate the varying way in which
491 decision-making dynamics — both prior to stimulus onset and during the stimulus period — differ
492 in different brain regions.

493

494 **Discussion**

495

496 We developed a probabilistic latent process model to simultaneously describe neural activity and
497 choices during an evidence accumulation decision-making task. We fit the model to data from three
498 brain regions and found that the dynamics of accumulation that best fit choices and neural data
499 from each brain region differed significantly across brain regions, and from the accumulation model
500 that best described the animal's choices. We found that including neural activity in the model
501 provided rich, moment-by-moment information about the animal's choice. The inferred
502 accumulation model could be used to examine estimates of the animal's moment-by-moment
503 provisional choice, and by doing so, we found differing choice-related dynamics in each brain
504 region, dynamics that meaningfully related to other measures of behavior such as reaction time.
505 Collectively, our results argue for the existence of very different accumulation dynamics in different
506 brain regions, dynamics which each differ greatly from the dynamics giving rise to behavior. An
507 exciting future application of our modeling framework is to model multiple, independent
508 accumulators in several brain regions which collectively give rise to the animal's behavior. Such a
509 model would provide incredible insight into how the brain collectively gives rise to behavioral
510 choices.

511

512 There has been substantial work relating neural activity to evidence accumulation. The logic
513 underlying this work (e.g., Churchland et al., 2011; Gold and Shadlen, 2007; Hanks et al., 2015;
514 Mante et al., 2013; Ratcliff et al., 2003; Yartsev et al., 2018) is that behavior is well approximated by
515 gradual evidence accumulation (Ratcliff and McKoon, 2008). Numerous studies have probed
516 whether neurons in any given brain are involved in encoding or computing a correlate of this
517 behavior-level evidence accumulation. A rarely emphasized assumption is that the accumulation
518 process, at the level of individual brain regions, will be similar to the accumulation process at the
519 level of the organism's behavior. This assumption need not be correct. As in the example mentioned
520 in the Introduction, two brain regions, one representing a leaky accumulator from which recent
521 evidence is best decoded, and another representing an unstable accumulator from which the
522 earliest evidence is best decoded, could combine to generate behavior that is well-described by
523 stable evidence accumulation, in which evidence from throughout behavioral trials is weighted
524 approximately equally. One should not conclude that neural activity best explained by a leaky or by
525 an unstable accumulator is unrelated to behavior that is best explained by stable accumulation.
526 Other properties, in addition to leakiness/instability, may also differ across contributing brain
527 regions. Developing a formal approach to fit the parameters of evidence accumulation models from
528 neural data as well as from choices provided us with the opportunity to probe this assumption. Our
529 results suggest that it is *not* correct. Elucidating the neural basis of evidence accumulation for
530 decision-making may require understanding how brain regions with neural activity that appears
531 driven by accumulators with potentially very different properties combine, and perhaps
532 counterbalance each other, so as to produce the organism's behavior.

533

534 Our approach extends and complements existing approaches that construct formal mathematical
535 models of decision making which combine both behavioral data and neural data. These models
536 leverage both neural and behavioral observations to jointly infer decision making parameters, as
537 we've done here (see Turner et al., 2019 for a comprehensive overview). However, the majority of
538 these approaches have tended to emerge from the field of cognitive neuroscience, and as such, have
539 predominantly focused on models for application to neural data acquired by other methods, such as

540 EEG, fMRI, etc. (e.g., Turner et al. 2015; but also see Frank et al., 2015). Our approach adds to these
541 efforts by offering a method that can combine fine timescale single unit recordings with behavioral
542 measurements specifically during pulse-based evidence accumulation tasks, thereby offering a
543 moment-by-moment picture into the latent dynamics that underlies cognition. Continued
544 development of joint models such as our and existing approaches in the field of cognitive
545 neuroscience are critical to quantitatively understand the latent processes underlying cognition.
546

547 One of our most surprising discoveries was that neural data from the FOF was best modeled by an
548 accumulator consistent with a ‘primacy’ strategy in which early stimulus clicks have an out-size
549 impact on neural activity and choice compared to later clicks. Coupled with the low accumulation
550 bound of the model fit to the FOF, our analysis suggests a model of FOF accumulation where a
551 subject prematurely commits to a decision based on early sensory evidence. Previous analysis of
552 these data did not find that FOF activity was described by an unstable accumulator because the
553 accumulator model was not learned from neural activity, only choices (Hanks et al., 2015). This
554 prior analysis identified an alternative interpretation of FOF activity: FOF activity exhibited a step-
555 like encoding of accumulated evidence that was unbounded, consistent with the FOF encoding a
556 categorical representation of choice (Hanks et al., 2015). At a strategic level, this interpretation is
557 consistent with the model of FOF activity we identified. Noting that, in this task, the stimulus will
558 rarely cause the accumulator to switch sign (**Fig 6D**), a step-like encoding of an unbounded
559 accumulator that does not switch sign will appear very much like a bounded accumulator: for
560 either model, the accumulator will quickly jump to its largest value and remain there. Additional
561 experiments and modeling are required to differentiate these two models.
562

563 A primacy encoding model of the FOF is supported by our change of mind analysis. Putative change
564 of mind events identified from neural activity occurred less frequently in the FOF than other
565 regions ($p < 4.5694e-82$ FOF vs. PPC; $p < 3.4585e-323$ FOF vs. ADS; Fisher’s exact test) consistent
566 with an early-commitment strategy in the FOF. A recent study of the FOF during an accumulation
567 task in which evidence dynamically changed throughout a trial found that FOF activity reflected
568 evidence across stimulus-induced ‘overt’ changes of mind, and that these events were common in
569 the FOF (Boyd-Meredith et al., 2022). It’s important to note that we likewise found that FOF reflects
570 evidence across changes of mind, but we identified rarely-occurring non-stimulus-induced ‘covert’
571 changes of mind during a task in which the evidence was static, and thus our results do not conflict
572 with those findings.
573

574 A primacy encoding model of the FOF is also both supported by and offers context to prior FOF
575 inactivation studies (Erlich et al., 2015). Behavioral modeling of choices in conjunction with
576 bilateral pharmacological inactivation found that FOF inactivation led to leakier accumulation when
577 producing choices (Erlich et al., 2015). Leakier accumulation at the level of choice also implies that
578 later stimulus information disproportionately impacts choice, precisely the impact predicted if an
579 early stimulus favoring brain region, such as the FOF, was silenced. A more complete model relating
580 accumulation dynamics in multiple brain regions to choice-related accumulation dynamics at the
581 level of behavior would aid in understanding how silencing individual brain regions, with their
582 region specific accumulation dynamics, impacts accumulation at the level of behavior.
583

584 Our novel change of mind analysis identified both the ADS and PPC as regions that showed frequent
585 instances of choice vacillation during this task. Prior studies in related tasks found that neural
586 responses in one of these regions, the PPC, (or its primate homolog), reflect information related to
587 already experienced trials (Akrami et al., 2018; Purcell and Kiani, 2016), consistent with our
588 interpretation of prestimulus neural responses being suboptimally tuned for the upcoming trial and
589 thus requiring mid-trial correction. Given the large initial accumulator variance of ADS and the

590 presence of frequent putative change of mind events in this region, activity in ADS seems poised to
591 also reflect these types of trial-history dependent responses as well. Future experiments and
592 analysis are required to determine this.

593
594 Previous studies that fit this model to only choices developed specific interpretations of the
595 accumulation strategy used by animals (Brunton et al., 2013). One difference between choice
596 accumulator models and joint neural-behavior models is the differential impact of accumulator
597 noise versus stimulus noise. Choice-only models have typically indicated that stimulus noise is the
598 primary cause of systematic behavioral uncertainty (Brunton et al., 2013), whereas our joint
599 models suggest that this impact is weaker than diffusion noise. One interpretation of this difference
600 is that at the level of a single neural population, diffusive noise plays a stronger role in producing
601 uncertainty in $a(t)$ than stimulus noise, whereas at the level of the entire brain's encoding of
602 accumulated evidence, this diffusive noise 'averages out' and residual stimulus noise remains.
603 Understanding how multiple brain regions work together to produce a model of accumulated
604 evidence at the level of behavior is an important future direction of this work.

605
606 Several extensions of our framework are readily apparent. Increasing the number of recorded
607 neurons led to an improved estimate of $a(t)$. As the density of neural recordings increases (Luo et
608 al., 2020), the explanatory power of our model will increase. Although we have extended the
609 evidence accumulation model to include neural responses and choice, we could extend it further to
610 describe additional physiological or behavioral variables (e.g., from annotated video data, pupil-
611 dilation measurements, response time, etc.). Including these additional behavioral measures would
612 further inform the inferred accumulator model, providing a clearer window into the internal factors
613 governing choices. Although we considered a specific evidence accumulation model due to its
614 normative interpretation, our framework can readily accept modifications and extensions of its
615 dynamical equations (e.g., Genkin et al., 2021). More sophisticated (e.g., nonlinear) dynamics of
616 accumulated evidence or more refined models of accumulation noise are two examples. Our
617 framework can also accommodate more elaborate and/or appropriate relationships between
618 accumulated evidence and neural responses, as we briefly explored by considering the negative
619 binomial distribution (**Figure 3 — figure supplement 1**). Changing this relationship would open
620 the door to using this approach with other types of data, such as imaging data. Although our
621 framework was developed with the specific application to a pulsed-based accumulation task in
622 mind, it is not confined to this. Our framework can be adapted to any task where noisy temporal
623 accumulation of evidence is thought to play a role, and for which neural recordings and behavioral
624 choices reflect this process (International Brain Laboratory et al., 2021). Finally, while a major
625 motivation of our approach was to develop a framework for identifying a specific normative and
626 mechanistic accumulation model, its rigidity makes it difficult to capture varying features present in
627 the data. Extending the model to include additional latent processes alongside a rigid accumulation
628 model (Zoltowski et al., 2020) would enable the model to simultaneously account for currently
629 unexplained variance in the data while preserving the model's ability to account for variance with
630 an accumulation model. Doing so may offer a clearer picture of the evidence accumulation process
631 by sweeping away unrelated variance with a more flexible, but less interpretable, latent process
632 model.

633 634 **Acknowledgments**

635
636 This work was supported by Simons Collaboration on the Global Brain (SCGB AWD543027 and
637 AWD542593), and NIH-NINDS BRAIN Initiative Award (5U19NS104648-02). We thank Michael
638 Yartsev, Tim Hanks, and Charles Kopec for providing the neural and behavioral data analyzed here.
639 We thank members of the Brody and Pillow labs for comments on this work, especially Thomas

640 Luo, Tim Kim, Diksha Gupta, and Tyler Boyd-Meredith. B.D. would like to thank Carol Mason, Paul
 641 DiMaggio, Betsy Levy Paluck, and Nathan Paluck for their support while this study was conducted.

642

643 **Methods**

644

645 **Latent variable model**

646

647 We model accumulated evidence as a one-dimensional drift diffusion model (DDM) with a
 648 symmetric absorbing boundary (Brunton et al., 2013). On a single behavioral trial, the evolution of
 649 the accumulated evidence, $a(t)$, is governed by

650

$$651 da = \lambda a dt + \sigma_a dW + \sigma_s dt (\eta' \delta_{t,t_R} C_R(t) - \eta' \delta_{t,t_L} C_L(t)). \quad (\text{Equation 4})$$

652

653 λ is the inverse of the drift time constant. $\sigma_a dW$ is a Wiener process with scaling σ_a . $\sigma_s \eta'$ are
 654 Gaussian variables with variance σ_s^2 and mean 1. δ_{t,t_L} and δ_{t,t_R} are the timing of left and right
 655 pulses respectively, and $C_L(t)$ and $C_R(t)$ are the magnitude that each left or right click, respectively,
 656 has at time t . The impact of each click is modulated by sensory adaptation, based on the following
 657 equation:

658

$$659 \frac{dC_\alpha}{dt} = \frac{1-C_\alpha}{\tau_\phi} + (\phi - 1)(C_\alpha \delta_{t,t_\alpha}), \quad (\text{Equation 5})$$

660

661 where $\alpha = \{L, R\}$. We define the difference of the adapted click magnitude at time t as $\Delta(t) =$
 662 $\delta_{t,t_R} C_R(t) - \delta_{t,t_L} C_L(t)$ and the sum of the adapted click magnitude at time t as $\Sigma(t) =$
 663 $\delta_{t,t_R} C_R(t) + \delta_{t,t_L} C_L(t)$. By doing so, we can express Equation 4 as,

664

$$665 da = \lambda a dt + \Delta(t) dt + \sigma_a dW + \sigma_s \Sigma(t) \eta dt, \quad (\text{Equation 6})$$

666

667 where η is a standard Normal. An absorbing boundary, B , if crossed, prevents $a(t)$ from evolving
 668 further (i.e. $da = 0$ if $a(t) > B$). The initial state of $a(t)$ is distributed normally with mean of 0 and
 669 variance of σ_i^2 . We refer to all parameters that govern the dynamics of $a(t)$ as

$$670 \theta_\alpha = \{\sigma_i, \lambda, B, \sigma_a, \sigma_s, \phi, \tau_\phi\}.$$

671

672 **Computing the distribution of the latent state**

673

674 The temporal dynamics of the probability distribution of $a(t)$, $P(a(t))$, can be expressed as a Fokker-
 675 Planck equation,

676

$$677 \frac{\partial P(a(t))}{\partial t} = \frac{\sigma_a^2 + \sigma_s^2 \Sigma(t)}{2} \frac{\partial^2 P}{\partial a^2} - \frac{\partial((\lambda a + \Delta(t))P)}{\partial a}. \quad (\text{Equation 7})$$

678

679 We numerically compute the solution to Equation 7 by dividing $P(a(t))$ into a set of n discrete
 680 spatial bins, and determine how mass moves after a discrete temporal interval, Δt . The discrete
 681 time dynamics of $P(a_t)$ are Markov, and obey the following equation,

682

$$683 P(a_t) = M(\theta_\alpha, \delta_t) P(a_{t-1}), \quad (\text{Equation 8})$$

684

685 where δ_t is the collection of left and right clicks at time t . The transition matrix $M(\theta_\alpha, \delta_t)$ is
 686 determined using methods established in Brunton et al., 2013. Briefly, for each spatial bin, the
 687 deterministic effect of the dynamics on the probability mass is computed, and this is convolved with

688 a discrete approximation to a Gaussian distribution with the appropriate variance and a finer
689 spatial resolution than the initial spatial resolution described above, to determine the various
690 locations of that probability mass at the next time bin. Because the location of each bin of mass after
691 the Gaussian convolution is not likely to correspond to the spatial grid defined for $P(a_t)$, the mass is
692 'settled' into appropriate bins based on the distance of each bit of mass and the nearest two bins.
693 Mass located in the first and last bin, corresponding to mass that has been captured by the
694 boundary, cannot change locations, and the entries of $M(\theta_a, \delta_t)$ that determines how the mass in
695 these bins moves, reflects this. $n = 53$ and $\Delta t = 10$ ms for all results presented here.

696

697 **Relating $a(t)$ to spikes and choices**

698

699 On a single behavioral trial, the observed spike count of the n^{th} neuron at time t , $y_{n,t}$, is a Poisson
700 random variable,

701

$$702 P(y_{n,t}|a_t, \theta_n) = (f_{\theta_n}(a_t))\Delta t^{y_{n,t}} \exp(-f_{\theta_n}(a_t)\Delta t), \quad \text{(Equation 9)}$$

703

704 where θ_n defines the expected firing rate function f for the n^{th} neuron. We choose f_{θ_n} to be a
705 softplus function, i.e., $\text{softplus}(x) = \log(1 + \exp(x))$. Each neuron has their own parameter θ_n that
706 relates f_{θ_n} to a_t . $\theta_y = \{\theta_1, \theta_2, \dots, \theta_N\}$ is the collection of all neural parameters for the population of N
707 neurons.

708

709 We define $f_{\theta_n}(a_t)$ as

710

$$711 f_{\theta_n}(a_t) = \text{softplus}(\theta_n a_t + \theta^0_{n,t}), \quad \text{(Equation 10)}$$

712

713 where $\theta^0_{n,t}$ accounts for the time-varying trial-average (i.e., invariant to $a(t)$) firing rate of the n^{th}
714 neuron. $\theta^0_{n,t}$ is learned prior to fitting the full model, i.e., before learning θ_a and θ_y . We
715 approximate $\theta^0_{n,t}$ with a set of six Gaussian radial basis functions

716

$$717 \theta^0_{n,t} = \sum_i^6 w_{i,n}^{RBF} N(\mu_i, \sigma^2_{RBF}).$$

(Equation 11)

718

719 The mean of the functions, μ_i are spaced uniformly from time 0 to the maximum trial length for
720 each respective neuron. The variance of the functions, σ^2_{RBF} , is equal to the distance between the
721 function means. We learn $w_{i,n}^{RBF}$ by assuming that $y_{n,t}$ is distributed Poisson with an intensity
722 function $\theta^0_{n,t}$ and maximize the likelihood. In other words, for the n^{th} neuron we define the
723 likelihood of the observed spikes for a trial of duration T , y_n , assuming a time-varying intensity
724 function $\theta^0_{n,t}$

725

$$727 P(y_n|\theta^0_n) = \prod_{t=1}^T (\theta^0_{n,t}\Delta t)^{y_{n,t}} \exp(-\theta^0_{n,t}\Delta t),$$

(Equation 12)

728

729 and maximize this likelihood across K trials with respect to the parameters $w_{i,n}^{RBF}$.

730

731 Although both $\theta_n a(t)$ and $\theta^0_n(t)$ vary in time to define each neuron's expected firing rate, they are
732 uniquely identifiable, because $\theta_n a(t)$ varies from trial to trial depending on the stimulus while
733 $\theta^0_n(t)$ does not. We verified through numerical experimentation and parameter recovery using
734 synthetic data that each process can be identified.

735

736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783

On a single behavioral trial, with a probability $1 - \gamma$ the subject's choice, d , is a deterministic function of $a(t)$ at the end of the trial (time T), (Brunton et al., 2013); with probability γ the choice is made without considering $a(t)$. γ captures "lapses" in the subject's performance. For choices that depend on $a(t)$, if $a(T)$ is greater than a cutoff value c , $d = 1$, otherwise $d = 0$. Thus, the probability of the choice, given $a(t)$ and θ_d can be written as,

$$P(d|a_T, \theta_d) = \left(\frac{\gamma}{2} + (1 - \gamma)H(a_T - c)\right)^d \left(\frac{\gamma}{2} + (1 - \gamma)(1 - H(a_T - c))\right)^{1-d},$$

(Equation 13)

where $H(\cdot)$ is the Heaviside function. We refer to the parameters relating $a(t)$ to the likelihood of a subject's choice as $\theta_d = \{c, \gamma\}$.

Relative binning of clicks and spikes

A minor but key implementation detail concerns defining the start and end times of the temporal bin edges that are used to bin the click inputs and the spikes trains. Through numerical experimentation, we identified that our numerical procedure produces a systematic error in estimating the model parameters when the temporal bins for the clicks are aligned with the temporal bins for the spikes. To circumvent this issue, we offset the bins for the spikes by $\Delta t/2$, so that the bin edges for spikes at time t surround the forward bin edge of the clicks by $\pm \Delta t/2$. This procedure is similar to the central difference formulation of a finite difference approximation to a differential equation.

Inferring model parameters with maximum likelihood

We refer to the set of all parameters for models fit to spikes and choices as $\theta = \{\theta_a, \theta_y, \theta_d\}$. Given the Markov dynamics described above, the relationship between $a(t)$ and the observed data, and the model parameters, we can write out the likelihood of the spike train data \mathbf{Y} from N neurons for T time bins, the behavioral choice d , and the latent variable \mathbf{a} for T time bins as

$$P(a, Y, d|\theta) = P(a_0|\theta_a) \prod_{t=1}^T P(a_t|a_{t-1}, \theta_a, \delta_t) \prod_{n=1}^N P(y_{n,t}|a_t, \theta_n) P(d|a_T, \theta_d)$$

(Equation 14)

We compute the likelihood of the data by integrating over \mathbf{a}

$$P(Y, d|\theta) = \sum_{\mathbf{a}} P(a, Y, d|\theta).$$

(Equation 15)

Because of the way in which we compute $P(a_t|a_{t-1}, \theta_a, \delta_t)$ (see above) computing the log-likelihood of the data can be done with a single forward pass over the data using the 'forward-backward' algorithm method for Hidden Markov models (Bishop, 2006). We maximize the sum over K behavioral trials of the logarithm of this quantity with respect to θ via gradient ascent. To compute the gradient of $\sum^K \log P(Y_k, d_k|\theta)$ with respect to θ we use a standard automatic differentiation package (Revels et al., 2016). We refer to the set of parameters that maximizes the likelihood as $\hat{\theta}$.

784 We note that all K trials for many of the models we fit were not recorded on the same behavioral
785 session, and therefore, all N neurons are not recorded for every trial. For example, neurons 1-3
786 might be recorded on trials 1-500, while neurons 4-6 might be recorded on trials 501-1000.
787 Although our notation does not reflect this in order to keep the notation simple, only neurons
788 recorded on a trial contribute to the likelihood on that trial.

789

790 **Bounded optimization**

791

792 Several model parameters are only defined within a restricted domain; for example, all variances
793 parameters, such as σ_a^2 , are only defined on the positive real axis. Alternatively, other parameters,
794 although defined on a more expansive domain, have values that correspond to models that are not
795 very likely; for example, although B is defined on the positive real axis, values much greater than 40
796 are not likely to be exhibited in the data, given the specifics of the stimulus, where greater than 40
797 clicks were rare. For these reasons, we define the following domain over which parameter
798 optimization was performed:

799

- 800 • $1e^{-3} \leq \sigma_a^2 \leq 100$
- 801 • $8 \leq B \leq 40$
- 802 • $-5 \leq \lambda \leq 5$
- 803 • $1e^{-3} \leq \sigma_a^2 \leq 400$
- 804 • $1e^{-3} \leq \sigma_s^2 \leq 10$
- 805 • $1e^{-3} \leq \phi \leq 1.2$
- 806 • $5e^{-3} \leq \tau_\phi \leq 1$
- 807 • $-10 \leq c \leq 10$
- 808 • $0 \leq \gamma \leq 1$
- 809 • $-10 \leq \theta_n \leq 10 \forall n$

810

811 The occurrence of parameters hitting the bound can be seen in Figure 3 & Figure 3 — figure
812 supplement 4. The most common boundary hitting situation was a variance parameter ($\sigma_i, \sigma_a, \sigma_s$)
813 hitting the lower boundary of zero, which means that the model did not support noise of that kind
814 in the model fit. σ_i and σ_a were found to do this for the choice only model, consistent with the
815 results of Brunton et al. The other bound that was frequently hit was the upper bound for the
816 accumulation bound parameter B , a result also consistent with the results of Brunton et al. The log-
817 likelihood surface as B grows very large becomes very flat, because it becomes increasingly unlikely
818 that probability mass $P(a(t))$ crosses the boundary. Thus, the model fits do not change appreciably
819 if this optimization boundary is relaxed.

820

821 **Confidence intervals for maximum likelihood parameters**

822

823 To compute confidence bounds of estimated parameters (as in **Figure 3** and **Figure 1 — figure**
824 **supplement 1**, Figure 4 — figure supplement 1, Figure 3 — figure supplement 3, Figure 3 — figure
825 supplement 4), we use the Laplace approximation to the log-likelihood. Using automatic
826 differentiation, we compute the Hessian (the matrix of second derivatives) of the log-likelihood at
827 the maximum likelihood parameters. The diagonal entries of the Hessian's inverse quantify the
828 sharpness of the curvature of the log-likelihood surface, and therefore the uncertainty of the
829 estimate of each parameter. We define the confidence bound as +/- two times the square root of
830 each diagonal entry; approximating the log-likelihood surface as Gaussian, this describes the range
831 of parameters that would fall within approximately 95% of the log-likelihood volume.

832

833 For some sets of maximum likelihood parameters, further consideration was necessary. In cases
834 where confidence bounds extend beyond an optimization bound that corresponds to a strict
835 boundary on the domain of a parameter (e.g., variance parameters being strictly positive), we
836 truncate these intervals at the bound. In some cases, we found that Hessian was not positive semi-
837 definite, a necessary condition to invert it. This most often occurred when a maximum likelihood
838 parameter encroached upon a strict parameter boundary (e.g., variance parameters being strictly
839 positive). We dealt with these scenarios in two ways. In some cases, numerical line search along any
840 Hessian eigenvector with negative eigenvalue confirmed the convexity of the log-likelihood was
841 local whereas more globally the log-likelihood was concave. In light of this, we numerically
842 computed the global concavity of the log-likelihood with a numerical line search and approximated
843 this curve with a quadratic function. We replaced the negative eigenvalue of the Hessian with two
844 times the coefficient of this quadratic approximation (the multiplier two is used because the
845 Hessian is two times the second-order approximation of the log-likelihood via Taylor series
846 approximation, where the second-order term contains a 1/2 prefactor). In other cases, computing
847 the Hessian in a transformed space (e.g., log space) where troublesome parameters were free to
848 take on any value, rectified the non-concavity (Yartsev et al., 2018). After computing confidence
849 intervals in the transformed space, we mapped these values back into the standard space by the
850 inverse transform.

851 852 **Data selection**

853
854 Details regarding behavioral data collection and neural recordings and spike sorting can be found
855 in (Hanks et al., 2015) and (Yartsev et al., 2018). To select which neurons were used, a firing rate
856 for each neuron was computed by summing spikes over the duration of the stimulus period and
857 dividing this by the length of the stimulus period. A two-sided t-test was applied, comparing the
858 firing rate distribution on trials when the animal chose left and when the animal chose right.
859 Neurons with a p -value less than 0.01 were included for analysis.

860 861 **Data grouping**

862
863 We grouped together rats that had neural recordings from the same brain region (five FOF rats,
864 three PPC rats, three ADS rats; see **Table 1** for information about the data) to improve our
865 estimation of the model parameters for each region. For the PPC and ADS recordings, the majority
866 of recorded neurons came from a single rat (**Table 1**). Although individual FOF rats had enough
867 neurons to support fitting each rat alone, the maximum likelihood parameters for FOF rats fit
868 individually were qualitatively similar (**Figure 3 — figure supplement 3**).

869 870 **Response latency**

871
872 Previous analyses have identified a response latency between the stimulus and the neural
873 responses, and that this latency can be different in different brain regions (Hanks et al., 2015). To
874 account for this, we shifted the time of the neural responses relative to the clicks based on these
875 prior results. FOF and ADS responses had a latency of 60 ms, while PPC responses had a latency of
876 120 ms.

877 878 **Specifics of data selection for each analyses**

879
880 Our reports of the maximum likelihood parameters for each model are for models fit to the entire
881 dataset. Each model was also fit using cross-validation (i.e., training on a subset of the data, while
882 reserving data for testing) but the maximum likelihood parameters did not qualitatively change

883 from those identified using the entire dataset, and the log-likelihood computed on test data using
884 parameters identified with training data did not differ appreciably from the log-likelihood
885 computed on those same trials using parameters identified with the entire dataset (Figure 3 —
886 figure supplement 4). This consistency is likely due to the modest number of model parameters.
887

888 When we compute various quantities related to the data, such as peri-stimulus time histograms
889 (PSTHs), cross-correlation functions, and psychometric functions, we likewise use the entire
890 dataset. We did not find that we could accurately estimate the PSTH when only using a small subset
891 of the data (i.e., test data) due to the fact that our task lacks repeated stimulus conditions.

892 Additionally, when we simulate data from a fit model (e.g., **Figure 2A**), we used the maximum
893 likelihood parameters derived from model fits to the entire dataset, and used the stimuli of the
894 entire dataset to generate these data. Again, because the maximum likelihood parameters did not
895 qualitatively change when the model was fit to a subset of the data, we found it easier to focus our
896 analyses on a single model. The above statements apply to analyses in the following figures: **Figure**
897 **2, Figure 3, Figure 4C & D, Figure 5, Figure 6 B-F, and Figure 3 — figure supplement 2, Figure 4**
898 **— figure supplement 1, Figure 6 — figure supplement 2, Figure 3 — figure supplement 3.**
899

900 When comparing performance across models, cross-validation is necessary, and we did so in those
901 cases (e.g. **Figure 4B, Figure 6A, and Figure 4 — figure supplement 2, Figure 6 — figure**
902 **supplement 1, Figure 3 — figure supplement 4 B**). In these cases, we performed five-fold cross-
903 validation by dividing the dataset into a training set that consisted of 80% of the data and a test set
904 that consisted of 20% of the data. We fit each model using the training data of each fold, and
905 computed the test log-likelihood using the test data and the parameters derived from the training
906 data. Test performance was averaged across the five folds. Again, we stress that the test
907 performance on cross-validated data did not appreciably differ from that computed using a model
908 trained to the entire dataset (Figure 3 — figure supplement 4). We note, however, that even in
909 cases when we performed cross-validation, we still computed an approximation to each neuron's
910 trial-averaged firing rates, $\theta_{n,t}^0$, using all available data, prior to fitting the full model.
911

912 Because most of our models were fit simultaneously to data from multiple experimental sessions
913 (in which different neurons are recorded), to perform cross-validation, we randomly divided trials
914 within each session into a train and test set, and trained and tested the model collectively on those
915 groups of trials. Testing the model in this way will determine parameter robustness across all
916 sessions (for model parameters that are shared across all sessions) and individual parameter
917 robustness within a session (for parameters that are specific to an individual session). This
918 procedure also worked for the ‘independent noise model’, for which model parameters were shared
919 across all sessions, but individual neuron parameters were session specific.
920

921 **Other fit models**

923 **Independent noise accumulator models**

924
925 We refer to the set of all parameters for the model with independent accumulator noise per neuron
926 as θ_{ind} . The likelihood of the spike train data from the n^{th} neuron Y_n for T time bins is
927

$$928 P(Y_n|\theta_{ind}) = \sum_{a_n} P(a_{0,n}|\theta_a) \prod_{t=1}^T P(a_{n,t}|a_{n,t-1}, \theta_a, \delta_t) P(y_{n,t}|a_{n,t}, \theta_n). \\ 929 \text{(Equation 16)} \\ 930$$

931 The joint likelihood for the spike train data from all neurons is the product of the likelihood for each
 932 neuron: $P(Y|\theta_{ind}) = \prod_{n=1}^N P(Y_n|\theta_{ind})$. Our primary interest in this analysis was capturing the
 933 neural responses, so we considered a simple model of choice for this model: on each trial, choice is
 934 determined by randomly selecting one of the accumulators. The likelihood of the choice d under
 935 such a model is the average of the the n accumulators at time T :
 936

$$937 \quad P(d|\theta_{ind}) = \frac{1}{N} \sum_{n=1}^N P(d|a_{n,T}\theta_d). \quad \text{(Equation 17)}$$

938
 939 The full likelihood is the product of these terms: $P(Y, d|\theta_{ind}) = P(d|\theta_{ind})P(Y|\theta_{ind})$.
 940

941 **Choice-only model**

942
 943 We refer to the set of all parameters for the model fit to choices only as $\theta_d = \{\theta_a, \theta_d\}$. The
 944 likelihood of the behavioral choice d is
 945

$$946 \quad P(d|\theta_d) = \sum_a P(a_0|\theta_a) \prod_{t=1}^T P(a_t|a_{t-1}, \theta_a, \delta_t) P(d|a_T, \theta_d)$$

(Equation 18)

948 949 **Bernoulli GLM**

950
 951 To benchmark our method's ability to predict the animal's choice, we considered a basic logistic
 952 regression model (i.e., Bernoulli GLM) that included stimulus information and neural activity (e.g.
 953 **Figure 6A** and **Figure 6 — figure supplement 1**). For each trial, we computed the total number of
 954 spikes each neuron produced during the specified temporal window and the final cumulative click
 955 difference, and used them as regressors in a standard Bernoulli generalized linear model to predict
 956 the animal's choice. A constant bias was also included, as well as a single lapse parameter that
 957 scaled the minimum and maximum values of the logistic inverse link function. Cross-validation was
 958 performed on this model as described above.
 959

960 **Null choice model**

961
 962 In **Figure 6A**, we assess how well each of our fitted models can predict choice. We compare all
 963 models against a baseline model where each choice is a Bernoulli random variable with probability
 964 of making a right choice equal to the empirical fraction of choices made to the right.
 965

966 **Null joint model**

967
 968 To compare the improvement of the joint model in absolute terms (i.e., when not comparing two
 969 fitted models) we compute a null model of the spiking activity and choices (**Figure 3 — figure
 970 supplement 4 B**). The null likelihood of the choice data is as described above. The null likelihood of
 971 the spike train data assumes that the time-varying expected firing rate of each neuron is equal to its
 972 estimated time-varying trial-average firing rate, i.e., $f_{\theta_n}(t) = \theta_{n,t}^0$.
 973

974 The improved performance (i.e. cross-validated log likelihood) of our joint model over the null
 975 model shown in Figure 3 — figure supplement 4 further confirms that $\theta_n a(t)$ and $\theta_n^0(t)$ are
 976 uniquely identifiable, and that they are not redundant (i.e. the joint model is not
 977 overparameterized).
 978

979 **Poisson GLM**

980

981 To validate the maximum likelihood parameters derived from the joint model, we fit a variant of a
 982 Poisson GLM to the spiking responses (**Figure 3 — figure supplement 2**). As a regressor, we used
 983 the adapted, exponentially filtered click inputs,

984

$$985 \quad da = \lambda a dt + dt \Delta(t), \quad \text{(Equation 19)}$$

986

987 where $\Delta(t)$ is defined as above. The expected firing rate of each neuron is defined as in the full
 988 model, by Equation 10. For the bounded Poisson GLM model, the dynamics of $a(t)$ follow Equation
 989 19, except that if $a(t)$ crosses B , $a(t)$ stops evolving (i.e. $da = 0$ if $a(t) > B$). The parameters λ , B , ϕ , τ_ϕ ,
 990 and θ_y that maximize the likelihood of the spike data were learned using gradient ascent. The null
 991 model described in Figure 3 — figure supplement 2 is the null joint model, described above.

992

993 **Negative binomial**

994

995 In **Figure 3 — figure supplement 1**, we compare a Poisson observation model to a negative
 996 binomial model. To do this, we model the spikes as

997

$$998 \quad P(y_{n,t} | a_t, \theta_n) = NB(\theta_n^{NB}, \frac{\theta_n^{NB}}{f_{\theta_n}(a_t) \Delta t + \theta_n^{NB}}) \quad \text{(Equation 20)}$$

999

1000 where $NB(\cdot, \cdot)$ is the negative binomial distribution, and θ_n^{NB} controls the variance of the
 1001 distribution for each neuron and can take values between 0 and positive infinity. When θ_n^{NB}
 1002 becomes large the negative binomial distribution approaches the Poisson distribution. θ_n^{NB} was fit
 1003 for each neuron using gradient ascent, as described above.

1004

1005 **Quantifying model fit**

1006

1007 **Computing PSTHs and cross-correlation functions on empirical data**

1008

1009 We computed a ‘single-trial’ firing rate for each neuron by convolving its binned spikes with a
 1010 Gaussian kernel of standard deviation 50 ms. We call this single-trial rate $r_{t,k,n}$ for the n^{th} neuron
 1011 on the k^{th} trial at time t . We divide all the trials into two equally-sized groups based on the
 1012 cumulative click difference at the end of the trial and average $r_{t,k,n}$ based on these groupings.
 1013 Because trials are not of equal duration, at time t we use whichever trials have data at that time. We
 1014 refer to this average as $\bar{r}_{c,n,t}$ where the index c runs from 1 to 2.

1015

1016 We used the empirical binned spikes counts to compute cross-correlation functions. Raw cross-
 1017 correlation functions were normalized by the (across time) mean firing rates of the two neurons
 1018 being used so they provided a measure of excess spike rate. The equation for the raw cross-
 1019 correlation function was,

1020

$$1021 \quad R_{m,n}(\tau) = \frac{1}{m_m} \left(\frac{1}{K} \sum_k \frac{1}{N_k(\tau)} \sum_t \frac{y_{n,k,t} y_{m,k,t-\tau}}{\Delta t \Delta t} \right) - m_n, \quad \text{(Equation 21)}$$

1022

1023 where t is over all bins for the k^{th} trial, $y_{n,k,t}$ and $y_{m,k,t-\tau}$ are the binned spike train of neuron n
 1024 and m at time t and $t - \tau$ respectively, and $N_k(\tau)$ is the number of bins such that both $y_{n,k,t}$ and
 1025 $y_{m,k,t-\tau}$ are valid. m_n and m_m are the mean firing rates of the n^{th} and m^{th} neuron respectively,
 1026 computed by taking the average spike count across all times.

1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073

To compute the shuffled corrected cross-correlation, we computed the cross-correlation of the expected firing rate of each neuron provided by the PSTH, i.e. $\bar{r}_{n,c,t}$,

$$R^{PSTH}_{m,n}(\tau) = \frac{1}{m_m} \left(\frac{1}{C} \sum_c \frac{1}{N_c(\tau)} \sum_t \bar{r}_{n,c,t} \bar{r}_{m,c,t-\tau} \right) - m_n, \quad \text{(Equation 22)}$$

where $C=2$ is the number of conditions used to define the PSTH, $N_c(\tau)$ is defined similarly as above, and m_n and m_m are as defined above. The shuffle corrected cross correlation is the raw cross correlation minus the cross-correlation of the expected firing rate: $R_{m,n}(\tau) - R^{PSTH}_{m,n}(\tau)$.

Computing PSTHs and cross-correlation functions on synthetic data

We generated synthetic data from a model by using the maximum likelihood parameters to generate the expected firing rate of each neuron on each trial, i.e. $f_{t,k,n}$. We averaged this expected rate for each neuron on each trial over 20 different realizations of the latent noise to reduce variation due to the latent process. We then grouped and averaged these average expected rates, as described above, to generate a synthetic PSTH, which we denote by $\bar{f}_{n,c,t}$, as used in **Figure 2** and **Figure 1 — figure supplement 1**.

We used the synthetic expected firing rate, $f_{t,k,n}$, to compute cross-correlation function for synthetic data,

$$R^{syn}_{m,n}(\tau) = \frac{1}{m_m} \left(\frac{1}{K} \sum_k \frac{1}{N_k(\tau)} \sum_t f_{n,k,t} f_{m,k,t-\tau} \right) - m_n, \quad \text{(Equation 23)}$$

where K , $N_k(\tau)$, m_n and m_m are as defined above. The shuffle corrected cross correlation function of synthetic data is the raw cross correlation function minus the cross correlation function of the expected synthetic firing rate provided by the synthetic PSTH, $\bar{f}_{n,c,t}$.

Goodness-of-fit metrics

To compare empirical and synthetic PSTHs, we computed the coefficient of determination. Because fewer and fewer trials were included in computing the PSTH at large time values (because trials of great length were rare) we included PSTH values 200 ms before the stimulus onset up until 500 ms after stimulus onset in this calculation. Based on the definitions of the empirical and synthetic PSTHs, the coefficient of determination is defined as:

$$R^2_n = 1 - \frac{\sum_c \sum_t (\bar{r}_{n,c,t} - f_{n,c,t})^2}{\sum_c \sum_t (\bar{r}_{n,c,t} - \langle \bar{r}_{n,c,t} \rangle_{ct})^2}, \quad \text{(Equation 24)}$$

where $\langle \bar{r}_{n,c,t} \rangle_{ct}$ is the mean of $\bar{r}_{n,c,t}$ over trial groupings and times. Pearson correlation (r) was used to compare empirical and synthetic cross-correlation functions. When computing r we considered values of τ between -800 and 800 ms.

Psychometric functions

1074 We used a Bernoulli GLM (i.e. logistic regression) to compute psychometric functions for empirical
1075 and synthetic data. We generated synthetic data from a model by using the maximum likelihood
1076 parameters to generate the probability of a choice, and sampled the choice from a Bernoulli
1077 distribution. For the Bernoulli GLM, for each trial, we computed the final click difference and used it
1078 as a regressor to predict the animal's choice. A constant bias was also included, as well as a single
1079 lapse parameter that scaled the minimum and maximum values of the logistic inverse link function.
1080 R^2 values comparing empirical and synthetic psychometric functions were defined as above, but
1081 using the psychometric functions whose domain was from the minimum final cumulative click
1082 difference to the maximum final cumulative click difference.

1083

1084 **Choice decoding**

1085

1086 We used two metrics to determine how well choice could be decoded from various models: choice
1087 prediction accuracy and test log likelihood. Test likelihood was reported in bits per trial, i.e.

1088

$$1089 \Delta LL = \frac{LL_{model} - LL_{null}}{\log_2(K)} \quad (\text{Equation 25})$$

1090

1091 where K is the number of trials in the test set and LL_{null} is the appropriate null model, as described
1092 above, or a second model with which to test against. Five-fold cross validation was performed, as
1093 described above. Accuracy was determined, depending on the model, by computing the probability
1094 that the model predicted a right choice, given all available data (i.e., inputs and spikes in a model
1095 that includes spikes). If the model had a greater than 0.5 probability of choosing right, we
1096 considered that a prediction of a rightward choice. Accuracy is the fraction of correct choice
1097 predictions.

1098

1099 **Identifying putative changes of mind**

1100

1101 Based on a recent study (Peixoto et al., 2021) we defined putative changes in mind in the following
1102 way. For each model and each trial, we computed the posterior distribution of $a(t)$ given all
1103 available data except for the choice. In the case of the choice only model, this means using only the
1104 stimulus, and is equivalent to the forward pass of the model. In the case of the joint model, this is
1105 equivalent to the posterior distribution of $a(t)$ given the spikes on that trial. We computed the
1106 expected value of the posterior distribution for each trial and identified moments when it crossed
1107 the decision threshold as determined for each model (i.e., the c parameter of the choice likelihood).
1108 We required that the expected value remain on one side of the threshold for 50 ms, remain on the
1109 other side following the crossing for 50 ms, and achieve an absolute magnitude greater or equal to
1110 2 at some point during that 100 ms window.

1111

1112 To relate putative change of mind events to the animal's behavior we performed linear regression
1113 between the time of the event relative to the end of the stimulus (i.e., how close to a decision the
1114 event occurred) and a measure of the animal's reaction time. In this task, the animal is required to
1115 fixate in the center poke for the duration of the stimulus, so it does not exhibit a true reaction time
1116 in the standard sense of the term. However, following the end of the stimulus, it does take the
1117 animal time to withdraw from the center port to make its choice (see **Figure 1A**, bottom, upper two
1118 lines). We refer to the difference between the end of the stimulus and when the animal withdrew
1119 from the center port as the animal's reaction time, which we used in our analysis.

1120

1121 **Estimating dimension**

1122

1123 To estimate the effective dimension of groups of simultaneously recorded neurons, we computed
1124 the ‘participation ratio’ (Litwin-Kumar et al., 2017). Single-trial firing rates were computed by
1125 convolving the spike trains with a Gaussian kernel (std=50 ms), and the covariance matrix of these
1126 rates was computed. The participation ratio is

1127
1128
$$\frac{(\sum_n^N \lambda_n)^2}{\sum_n^N (\lambda_n)^2}$$

1129 (Equation 26)

1130
1131 where λ are the eigenvalues of the covariance matrix. If the firing rates are independent, the
1132 eigenvalues will all be equal and the participation ratio will equal the number of neurons. If the
1133 firing rates are correlated such that some eigenvalues are small (or perhaps even zero) the
1134 participation ratio will be less than the number of neurons.

1135 **Code availability**

1136
1137 All code was written in the Julia programming language. The core codebase for fitting the models
1138 described in this manuscript can be found here: <https://github.com/Brody-Lab/PulseInputDDM>.
1139 Code and data for performing the analyses described in this manuscript can be found here:
1140 <https://github.com/Brody-Lab/DePasquale-eLife-2023>.

1141 **Figure Legends**

1142
1143 **Figure 1: Accumulating evidence task and latent variable model.** (A) Rats performed a pulsed-
1144 based evidence accumulation task. A central LED illuminates, indicating that the rat can begin a trial
1145 by poking its nose in a central port. After a delay of random duration, an auditory stimulus of
1146 variable duration is delivered—a series of brief auditory pulses played from a left and a right
1147 speaker. Upon cessation of the stimulus, the rat must orient to the direction of the greater number
1148 of pulses to receive a water reward. (B) The model relates the click-based sensory stimulus to two
1149 types of observations—the animal’s choice and neural activity observed during the task. The latent
1150 variable model is a bounded accumulator. Left and right clicks (green and red arrows, respectively)
1151 push the variable to one side or the other; if the accumulator variable reaches the bound B (dotted
1152 line) accumulation ceases. Seven parameters govern the dynamics of $a(t)$ (see main text). Two
1153 different hypothetical trajectories of $a(t)$ are illustrated (black and blue) for the same click
1154 stimulus; the two trajectories differ due to the diffusive and stimulus noise in the model. $a(t)$ relates
1155 to the animal’s choice by a Heaviside step function and to neural activity by way of a softplus
1156 nonlinearity and a Poisson distribution. $a(t)$ is common for all simultaneously recorded neurons
1157 and each neuron has its own parameters that determine its tuning curve.

1158
1159
1160
1161 **Figure 1 — figure supplement 1: Recovering the parameters of synthetic data.** Synthetic data
1162 was generated with parameters $\sigma_i = 5$, $B=15$, $\lambda = -0.5$, $\sigma_a = 100$, $\sigma_s = 20$, $\phi = 0.4$, $\tau_\phi = 0.02$.
1163 Two synthetic ‘sessions’ were generated, with 400 trials and 3 neurons each. Softplus gain
1164 parameters were randomly generated between -2 and 2. $c = 1$, $\gamma = 0.05$. (A) PSTHs for two example
1165 neurons for synthetic data and simulated data after modeling fitting. (B) Psychometric curves for
1166 synthetic data and simulated data after modeling fitting. (C) Optimization was initialized at a
1167 random set of parameters (‘initial’). Maximum likelihood parameters (‘final’) converged to within
1168 two standard deviations (error bars computed by Laplace approximation) of the parameters used

1169 to generate the data (dotted lines). θ_{ij} refers to the neuron parameters for the j th neuron from the
1170 i th session.

1171
1172 Figure 1 — figure supplement 2: Recovering the parameters of synthetic data for multiple datasets.
1173 Four synthetic datasets (red, cyan, green, blue) were generated as in Figure 1 — figure supplement
1174 1 (two sessions per dataset, with three neurons in each session). Dotted lines in each panel indicate
1175 the generative parameters. Optimization was initialized at a random set of parameters ('init.').
1176 Maximum likelihood parameters ('final') almost always converge to within two standard deviations
1177 (error bars computed by Laplace approximation) of the parameters used to generate the data. θ_{ij}
1178 refers to the neuron parameters for the j th neuron from the i th session.

1179
1180 **Figure 2: A shared accumulator model captures neural response and choice for each brain**
1181 **region.** (A) Peri-stimulus time histograms (PSTHs) of three example neurons for each brain region
1182 (each row; FOF: red/green, PPC: blue/orange, ADS: purple/yellow). Spike trains were binned,
1183 filtered with a Gaussian kernel (std = 50 ms), grouped based on the strength of evidence, and
1184 averaged. Transparent shaded regions are +/- 1 standard error of the mean for the empirical data
1185 for each grouping. Colored curves are the mean of synthetic data simulated from the model with the
1186 parameters that maximize the likelihood of the data, grouped in a similar fashion. The black curve
1187 shows the trial-averaged firing rate, for all evidence strengths. Gray vertical lines indicate the
1188 average delay between the stimulus and the response for each brain region (see Methods). (B)
1189 Coefficient of determination (R^2) between empirical PSTH and synthetic data PSTH, for each neuron
1190 in each brain region. The data are plotted as a function of average firing rate. The median across the
1191 population is shown as a line. Points indicated with a 'star' refer to the data plotted in (A). (C)
1192 Probability of making a rightward choice as a function of cumulative difference in the number of
1193 clicks (psychometric curves) for empirical data (black lines) and data simulated from the model
1194 with the best fitting parameters (colored curves; FOF: red, PPC: blue, ADS: purple). Each curve is
1195 the curve of best fit, as computed by logistic regression.

1196
1197 **Figure 3: Data from different regions is best fit by different accumulator models.** (A)
1198 Maximum likelihood parameters that govern $a(t)$ for the joint neural-behavioral model and the
1199 choice-only model. Error bars, computed by the Laplace approximation (Methods), are +/- 2
1200 standard deviations. Parameters are σ_i : initial variance, B: accumulation bound, λ : drift, σ_a :
1201 accumulation noise variance, σ_s : click noise variance, ϕ : adaptation strength, τ_ϕ : adaptation
1202 timescale. (B) 10 example trajectories with different noise instantiations for one trial for the choice
1203 model (top) and the joint model (middle) model for each brain region, and cumulative sum of the
1204 click stimulus for each trial (bottom). The dotted black lines (top and middle) indicate the
1205 accumulation boundary value for each model.

1206
1207 Figure 3 — figure supplement 1: **Model comparison using Poisson or negative binomial**
1208 **observation model.** (A) Parameters are σ_i : initial variance, B: accumulation bound, λ : drift, σ_a :
1209 accumulation noise variance, σ_s : click noise variance, ϕ : adaptation strength, τ_ϕ : adaptation
1210 timescale. Each point is a data fold (1 of 5). Maximum likelihood parameters were similar for the
1211 two observation models. Cross-validated log-likelihood was statistically indistinguishable (FOF:
1212 $p=0.99$; PPC:0.93; ADS:0.98) and the average difference in cross-validated log-likelihood was small
1213 (FOF: 1.11e-5; PPC: -0.036; ADS: -0.035). (B) Histogram of the negative binomial dispersion
1214 parameter (r) across all neurons for each region. For large values of r , as seen here, the negative
1215 binomial approaches the Poisson distribution.

1216

1217 **Figure 3 — figure supplement 2: GLM analysis of individual sessions.** (A) Poisson GLM with a
1218 softplus nonlinearity was fit with exponentially filtered clicks as the regressors (see Methods),
1219 using the same data as in **Figure 2** and **Figure 3**. Each dot is the maximum likelihood drift (λ)
1220 parameter for a session. Sessions are ordered (from left to right) based on the fraction of the
1221 cumulative sum (across all sessions for a brain region) of the change in log likelihood (LL) over the
1222 null model (see Methods for null model). For example, the leftmost dot for each brain region is the
1223 session with the largest change in LL. Dots on the right were from sessions with the smallest change
1224 in LL over the null. The colored lines are the cumulative mean of λ weighted by that session's
1225 normalized change in LL. Dots on the far right have little change in LL and thus contribute to this
1226 mean only weakly. (B) GLM as in (A) but fit with a boundary, such that if the filtered clicks crossed a
1227 boundary B , the value of the regressors remained equal to B henceforth in a trial (see Methods). In
1228 each plot, the dashed colored lines are the values of λ from the full model fit (as in **Figure 2** and
1229 **Figure 3**).

1230

1231 **Figure 3 — figure supplement 3: Maximum likelihood parameters of joint model for each**
1232 **FOF rat individually.** Error bars, computed by the Laplace approximation (Methods), are ± 2
1233 standard deviations. Parameters are σ_i : initial variance, B : accumulation bound, λ : drift, σ_a :
1234 accumulation noise variance, σ_s : click noise variance, ϕ : adaptation strength, τ_ϕ : adaptation
1235 timescale.

1236

1237 **Figure 3 — figure supplement 4: Comparison of maximum likelihood parameters for three**
1238 **models: joint (neural/choice) model, choice-only model, and independent noise joint model,**
1239 **when fit to all data, or using cross-validation data.** (A) Circles with error bars are for models fit
1240 to all data. Error bars for models fit to full data computed by the Laplace approximation (Methods)
1241 are ± 2 standard deviations. 'Diamond' marks are models (5 for each model type) fit to cross-
1242 validation data (5-fold). Parameters are σ_i : initial variance, B : accumulation bound, λ : drift, σ_a :
1243 accumulation noise variance, σ_s : click noise variance, ϕ : adaptation strength, τ_ϕ : adaptation
1244 timescale. (B) Test log likelihood for models fit to all data (i.e., using trials reserved as testing trials
1245 when cross-validation is done) plotted against test log likelihood for cross validation models, for
1246 each model type (joint, choice, joint (ind.)), for all three brain regions.

1247

1248 **Figure 4: ADS is better described by independent accumulators.** (A) For the shared noise
1249 accumulator model (top), a set of parameters defines the dynamics of a single accumulator, which
1250 drives the spiking activity of the entire population. In the independent noise accumulator model, a
1251 set of parameters defines the dynamics of an ensemble of independent accumulator models, which
1252 each individually determine the spiking of a single neuron. (B) Difference in test log-likelihood
1253 (bits/trial) for the shared noise versus independent noise accumulator models. (C) Empirical (red)
1254 and synthetic (shared: black; independent; gray) shuffle-corrected cross-correlation function for
1255 three simultaneously recorded neurons from the FOF. Corresponding PSTHs are shown below for
1256 reference. (D) Same as (C) for three (of five) simultaneously recorded neurons from the ADS.

1257

1258 **Figure 4 — figure supplement 1: Maximum likelihood parameters for the joint**
1259 **(neural/choice, i.e., shared noise) model and independent ('ind.') noise joint model.** Error
1260 bars, computed by the Laplace approximation (Methods), are ± 2 standard deviations. Parameters
1261 are σ_i : initial variance, B : accumulation bound, λ : drift, σ_a : accumulation noise variance, σ_s : click
1262 noise variance, ϕ : adaptation strength, τ_ϕ : adaptation timescale.

1263
1264 **Figure 4 — figure supplement 2: ΔLL between the shared-noise and independent-noise**
1265 **accumulator model.** (A) Difference in log likelihood for each session for FOF and ADS data plotted
1266 as a function of the number of neurons in each session. (B) When the number of neurons in each
1267 session for the ADS dataset was subsampled to match the maximum number of neurons in a FOF
1268 session (3 neurons) the ADS was still favored by an independent noise accumulator model (purple,
1269 no fill; averaged across 2 subsample permutations of the ADS recordings). (C) Same as (A) but
1270 plotted as a function of dimension, as computed by the participation ratio (see Methods). Sessions
1271 in the ADS with higher dimension favored the independent noise accumulator model, leading to the
1272 net effect seen in **Figure 4B**. (D) The difference in log likelihood was similar when the choice data
1273 was omitted from both models.

1274
1275 **Figure 5: Neural data provides more information about accumulated evidence on single**
1276 **trials than choice alone.** (A) Posterior distribution of $a(t)$ under the joint model (excluding
1277 captured mass at the boundary) given only the choice (top row) and given spike times and choice
1278 (bottom row), for a single example trial. Columns show example trials for different brain regions.
1279 (B) Histogram of joint model posterior standard deviations given choice data (*black*) or both neural
1280 and choice data (*colors*) for all three brain regions. (C) Difference in choice-conditioned joint
1281 posterior standard deviation and neural- and choice-conditioned joint posterior standard deviation
1282 as a function of the number of simultaneously recorded neurons. Each point is the difference in the
1283 average posterior standard deviation for a session. Negative values indicate that the neural- and
1284 choice-conditioned posterior had smaller average standard deviation than the choice-conditioned
1285 posterior.

1286
1287 **Figure 6: Joint neural-behavioral model improves choice decoding.** (A) Choice-prediction
1288 accuracy, quantified with log-likelihood (left) and percent correct (right) on test choice data for 4
1289 models: joint neural-behavioral model, choice-only model, and two logistic regression models
1290 (Methods). Values greater than zero indicate that the model can predict choices better than a
1291 baseline model that only knows the marginal probability of a rightward choice. (B) Posterior mean
1292 of $a(t)$ conditioned on the neural activity for the joint model (colors), the distribution of $a(t)$ for the
1293 choice only model (black), and the cumulative click difference (gray) for three example trials (one
1294 for each brain region). ‘animal’s choice’ arrow indicates the choice (left or right) the animal made
1295 on that trial. (C) Putative change of mind events, where the posterior mean of the joint model
1296 crossed the decision threshold. The corresponding distribution of $a(t)$ for the choice only model
1297 (black) and the cumulative click difference (gray) for the same trial are shown for comparison.
1298 ‘animal’s choice’ arrow indicates the choice (left or right) the animal made on that trial. (D) Fraction
1299 of trials that contain at least one putative change of mind event for the cumulative click difference,
1300 the choice model, and the joint model, for each brain region. (E) Fraction of trials for which a
1301 putative change of mind event occurs at the specified time relative to the end of the stimulus for the
1302 joint model (color) and the cumulative click difference (black) for each brain region. (F) Choice
1303 response latency as a function of timing of putative change of mind events relative to stimulus
1304 offset for each brain region. Bar plots show the 25-75 percentiles of the choice response latency for
1305 putative change of mind events occurring at similar times. The colored lines indicate the line of best
1306 fit for each brain region computed by linear regression.

1307
1308 **Figure 6 — figure supplement 1: GLM choice decoding (as in Figure 6A) using spikes in different**
1309 **time windows relative to stimulus offset.**

1310

1311 **Figure 6 — figure supplement 2: Accuracy on putative change of mind event trials and non-**
 1312 **event trials.** (A) Accuracy of the rat for data from each brain region for putative change of mind
 1313 event trials and trials that lacked events ('no event'). (B) Same as (A) but for accuracy of the joint
 1314 model for each brain region.
 1315
 1316

Rat	Region	Sessions	Neurons	Trials	Sessions with greater than 1 neuron	Max. # of simultaneously recorded neurons
B068	FOF	11	13	5859	2	2
T034	FOF	9	10	4138	1	2
T036	FOF	8	12	3026	4	2
T063	FOF	17	32	4002	9	3
T030	FOF	1	1	357	0	1
T035	PPC	15	16	5919	1	2
T011	PPC	7	7	2235	0	1
B053	PPC	2	2	883	0	1
T080	ADS	5	6	1731	1	2
T103	ADS	19	38	8332	9	5
E021	ADS	3	4	697	1	2

1317
 1318 **Table 1:** Number of neurons, sessions and trials for each rat.
 1319

1320 **References**

1321 Akrami A, Kopec CD, Diamond ME, Brody CD. 2018. Posterior parietal cortex represents sensory
 1322 history and mediates its effects on behaviour. *Nature* **554**:368–372.
 1323 Aoi MC, Mante V, Pillow JW. 2020. Prefrontal cortex exhibits multidimensional dynamic encoding
 1324 during decision-making. *Nat Neurosci* **23**:1410–1420.
 1325 Beck JM, Ma WJ, Kiani R, Hanks T, Churchland AK, Roitman J, Shadlen MN, Latham PE, Pouget A.
 1326 2008. Probabilistic population codes for Bayesian decision making. *Neuron* **60**:1142–1152.
 1327 Bishop CM. 2006. Pattern Recognition and Machine Learning. Springer-Verlag, 0387310738.
 1328 Bogacz R, Brown E, Moehlis J, Holmes P, Cohen JD. 2006. The physics of optimal decision making: a
 1329 formal analysis of models of performance in two-alternative forced-choice tasks. *Psychol Rev*
 1330 **113**:700–765.
 1331 Bollimunta A, Totten D, Ditterich J. 2012. Neural dynamics of choice: single-trial analysis of
 1332 decision-related activity in parietal cortex. *J Neurosci* **32**:12684–12701.

1333 Boyd-Meredith JT, Piet AT, Dennis EJ, El Hady A, Brody CD. 2022. Stable choice coding in rat frontal
1334 orienting fields across model-predicted changes of mind. *Nat Commun* **13**:3235.

1335 Brody CD, Hanks TD. 2016. Neural underpinnings of the evidence accumulator. *Curr Opin Neurobiol*
1336 **37**:149–157.

1337 Brunton BW, Botvinick MM, Brody CD. 2013. Rats and humans can optimally accumulate evidence
1338 for decision-making. *Science* **340**:95–98.

1339 Chandrasekaran C, Hawkins GE. 2019. ChaRTr: An R toolbox for modeling choices and response
1340 times in decision-making tasks. *J Neurosci Methods* **328**:108432.

1341 Churchland AK, Kiani R, Chaudhuri R, Wang X-J, Pouget A, Shadlen MN. 2011. Variance as a
1342 signature of neural computations during decision making. *Neuron* **69**:818–831.

1343 Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark
1344 AM, Hosseini P, Scott BB, Bradley DC, Smith MA, Kohn A, Movshon JA, Armstrong KM, Moore T,
1345 Chang SW, Snyder LH, Lisberger SG, Priebe NJ, Finn IM, Ferster D, Ryu SI, Santhanam G, Sahani
1346 M, Shenoy KV. 2010. Stimulus onset quenches neural variability: a widespread cortical
1347 phenomenon. *Nat Neurosci* **13**:369–378.

1348 Ding L, Gold JI. 2013. The basal ganglia’s contributions to perceptual decision making. *Neuron* **79**.
1349 doi:10.1016/j.neuron.2013.07.042

1350 Ding L, Gold JI. 2010. Caudate encodes multiple computations for perceptual decisions. *J Neurosci*
1351 **30**:15747–15759.

1352 Ditterich J. 2006. Stochastic models of decisions about motion direction: behavior and physiology.
1353 *Neural Netw* **19**:981–1012.

1354 Erlich JC, Bialek M, Brody CD. 2011. A cortical substrate for memory-guided orienting in the rat.
1355 *Neuron* **72**:330–343.

1356 Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD. 2015. Distinct effects of prefrontal and
1357 parietal cortex inactivations on an accumulation of evidence task in the rat. *Elife* **4**.
1358 doi:10.7554/eLife.05457

1359 Frank MJ, Gagne C, Nyhus E, Masters S, Wiecki TV, Cavanagh JF, Badre D. 2015. fMRI and EEG
1360 predictors of dynamic decision parameters during human reinforcement learning. *J Neurosci*.
1361 2015 Jan 14;35(2):485-94. doi: 10.1523/JNEUROSCI.2036-14.2015. PMID: 25589744; PMCID:
1362 PMC4293405.

1363 Genkin M, Hughes O, Engel TA. 2021. Learning non-stationary Langevin dynamics from stochastic
1364 observations of latent trajectories. *Nat Commun* **12**:5986.

1365 Gold JI, Shadlen MN. 2007. The neural basis of decision making. *Annu Rev Neurosci* **30**:535–574.

1366 Gold JI, Shadlen MN. 2000. Representation of a perceptual decision in developing oculomotor
1367 commands. *Nature* **404**:390–394.

1368 Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD. 2015. Distinct relationships of
1369 parietal and prefrontal cortices to evidence accumulation. *Nature* **520**:220–223.

1370 Howard MW, Luzardo A, Tiganj Z. 2018. Evidence accumulation in a Laplace domain decision space.
1371 *Comput Brain Behav* **1**:237–251.

1372 Huk AC, Shadlen MN. 2005. Neural activity in macaque parietal cortex reflects temporal integration
1373 of visual motion signals during perceptual decision making. *J Neurosci* **25**:10420–10436.

1374 International Brain Laboratory, Aguillon-Rodriguez V, Angelaki D, Bayer H, Bonacchi N, Carandini
1375 M, Cazzettes F, Chappuis G, Churchland AK, Dan Y, Dewitt E, Faulkner M, Forrest H, Haetzl L,
1376 Häusser M, Hofer SB, Hu F, Khanal A, Krasniak C, Laranjeira I, Mainen ZF, Meijer G, Miska NJ,
1377 Mrsic-Flogel TD, Murakami M, Noel J-P, Pan-Vazquez A, Rossant C, Sanders J, Socha K, Terry R,
1378 Urai AE, Vergara H, Wells M, Wilson CJ, Witten IB, Wool LE, Zador AM. 2021. Standardized and
1379 reproducible measurement of decision-making in mice. *Elife* **10**. doi:10.7554/eLife.63711

1380 Kaufman MT, Churchland MM, Ryu SI, Shenoy KV. 2015. Vacillation, indecision and hesitation in
1381 moment-by-moment decoding of monkey motor cortex. *Elife* **4**:e04677.

1382 Kiani R, Cueva CJ, Reppas JB, Newsome WT. 2014. Dynamics of neural population responses in

1383 prefrontal cortex indicate changes of mind on single trials. *Curr Biol* **24**:1542–1547.

1384 Kim JN, Shadlen MN. 1999. Neural correlates of a decision in the dorsolateral prefrontal cortex of
1385 the macaque. *Nat Neurosci* **2**:176–185.

1386 Latimer KW, Freedman DJ. 2021. Low-dimensional encoding of decisions in parietal cortex reflects
1387 long-term training history. *bioRxiv*. doi:10.1101/2021.10.07.463576

1388 Latimer KW, Yates JL, Meister MLR, Huk AC, Pillow JW. 2015. NEURONAL MODELING. Single-trial
1389 spike trains in parietal cortex reveal discrete steps during decision-making. *Science* **349**:184–
1390 187.

1391 Litwin-Kumar A, Harris KD, Axel R, Sompolinsky H, Abbott LF. 2017. Optimal Degrees of Synaptic
1392 Connectivity. *Neuron* **93**:1153–1164.e7.

1393 Luo TZ, Bondy AG, Gupta D, Elliott VA, Kopec CD, Brody CD. 2020. An approach for long-term, multi-
1394 probe Neuropixels recordings in unrestrained rats. *Elife* **9**. doi:10.7554/eLife.59716

1395 Mante V, Sussillo D, Shenoy KV, Newsome WT. 2013. Context-dependent computation by recurrent
1396 dynamics in prefrontal cortex. *Nature* **503**:78–84.

1397 Park IM, Meister MLR, Huk AC, Pillow JW. 2014. Encoding and decoding in parietal cortex during
1398 sensorimotor decision-making. *Nat Neurosci* **17**:1395–1403.

1399 Peixoto D, Verheijn JR, Kiani R, Kao JC, Nuyujukian P, Chandrasekaran C, Brown J, Fong S, Ryu SI,
1400 Shenoy KV, Newsome WT. 2021. Decoding and perturbing decision states in real time. *Nature*
1401 **591**:604–609.

1402 Perkel DH, Gerstein GL, Moore GP. 1967. Neuronal spike trains and stochastic point processes. II.
1403 Simultaneous spike trains. *Biophys J* **7**:419–440.

1404 Purcell BA, Kiani R. 2016. Neural Mechanisms of Post-error Adjustments of Decision Policy in
1405 Parietal Cortex. *Neuron* **89**:658–671.

1406 Raposo D, Sheppard JP, Schrater PR, Churchland AK. 2012. Multisensory decision-making in rats
1407 and humans. *J Neurosci* **32**:3726–3735.

1408 Ratcliff R, Cherian A, Segraves M. 2003. A comparison of macaque behavior and superior colliculus
1409 neuronal activity to predictions from models of two-choice decisions. *J Neurophysiol* **90**:1392–
1410 1407.

1411 Ratcliff R, McKoon G. 2008. The diffusion decision model: theory and data for two-choice decision
1412 tasks. *Neural Comput* **20**:873–922.

1413 Ratcliff R, Smith PL, Brown SD, McKoon G. 2016. Diffusion Decision Model: Current Issues and
1414 History. *Trends Cogn Sci* **20**:260–281.

1415 Resulaj A, Kiani R, Wolpert DM, Shadlen MN. 2009. Changes of mind in decision-making. *Nature*
1416 **461**:263–266.

1417 Revels J, Lubin M, Papamarko T. 2016. Forward-Mode Automatic Differentiation in Julia. arXiv,
1418 <http://arxiv.org/abs/1607.07892>.

1419 Roitman JD, Shadlen MN. 2002. Response of neurons in the lateral intraparietal area during a
1420 combined visual discrimination reaction time task. *J Neurosci* **22**:9475–9489.

1421 Sanders JI, Kepecs A. 2012. Choice ball: a response interface for two-choice psychometric
1422 discrimination in head-fixed mice. *J Neurophysiol* **108**:3416–3423.

1423 Shadlen MN, Newsome WT. 2001. Neural basis of a perceptual decision in the parietal cortex (area
1424 LIP) of the rhesus monkey. *J Neurophysiol* **86**:1916–1936.

1425 Shinn M, Lam NH, Murray JD. 2020. A flexible framework for simulating and fitting generalized
1426 drift-diffusion models. *Elife* **9**. doi:10.7554/eLife.56938



1427 Smith MA, Kohn A. 2008. Spatial and temporal scales of neuronal correlation in primary visual
1428 cortex. *J Neurosci* **28**:12591–12603.



1429 Turner BM, van Maanen L, Forstmann BU. Informing cognitive abstractions through
1430 neuroimaging: the neural drift diffusion model. *Psychol Rev.* 2015 Apr;122(2):312-336. doi:
1431 10.1037/a0038894. PMID: 25844875.

1432 Turner BM, van Maanen L, Forstmann BU. Informing cognitive abstractions through

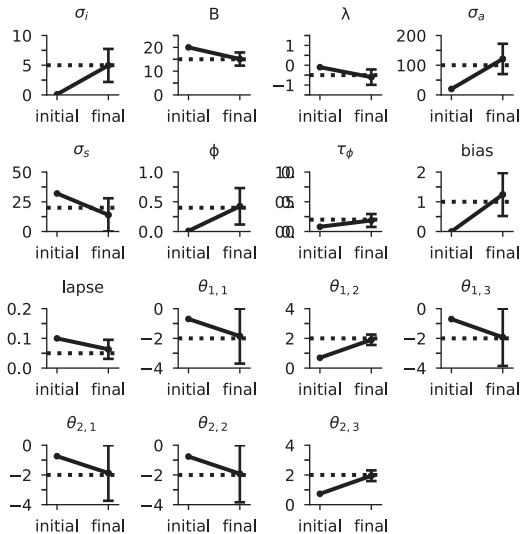
1433 neuroimaging: the neural drift diffusion model. *Psychol Rev.* 2015 Apr;122(2):312-336. doi:
1434 10.1037/a0038894. PMID: 25844875.
1435 Wiecki TV, Sofer I, Frank MJ. 2013. HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion
1436 Model in Python. *Front Neuroinform* 7:14.
1437 Yartsev MM, Hanks TD, Yoon AM, Brody CD. 2018. Causal contribution and dynamical encoding in
1438 the striatum during evidence accumulation. *Elife* 7. doi:10.7554/eLife.34929
1439 Zoltowski DM, Latimer KW, Yates JL, Huk AC, Pillow JW. 2019. Discrete Stepping and Nonlinear
1440 Ramping Dynamics Underlie Spiking Responses of LIP Neurons during Decision-Making.
1441 *Neuron* 102:1249–1258.e10.
1442 Zoltowski D, Pillow J, Linderman S. 2020. A general recurrent state space framework for modeling
1443 neural dynamics during decision-making In: Iii HD, Singh A, editors. Proceedings of the 37th
1444 International Conference on Machine Learning, Proceedings of Machine Learning Research.
1445 PMLR. pp. 11680–11691.
1446 Zylberberg A, Fetsch CR, Shadlen MN. 2016. The influence of evidence volatility on choice, reaction
1447 time and confidence in a perceptual decision. *Elife* 5. doi:10.7554/eLife.17688

1448

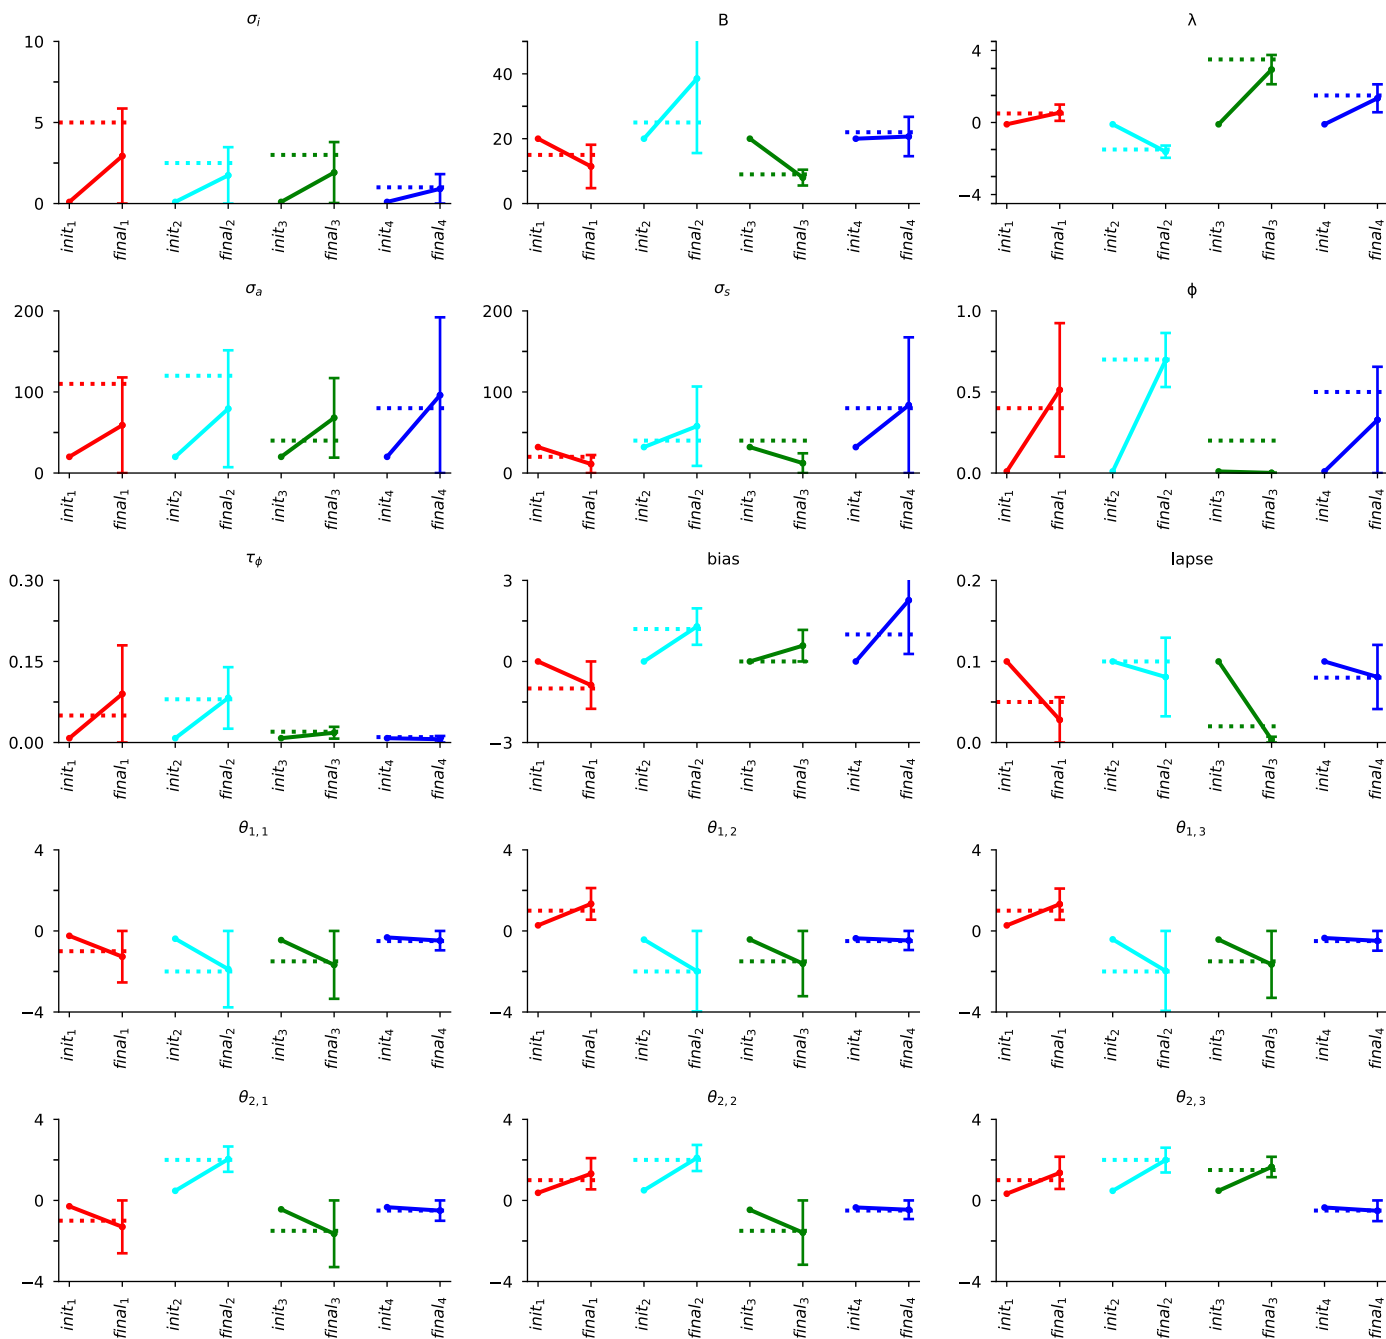


A example recovered PSTHs**B** example recovered psychometric curve**C**


example recovered parameters




example recovered parameters




A



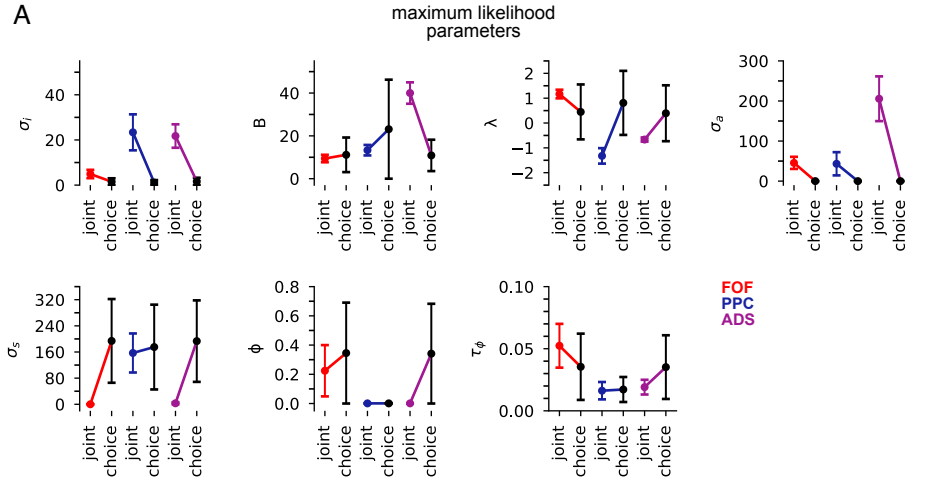
B



C



A



B




example latent trajectories

choice model







FOF

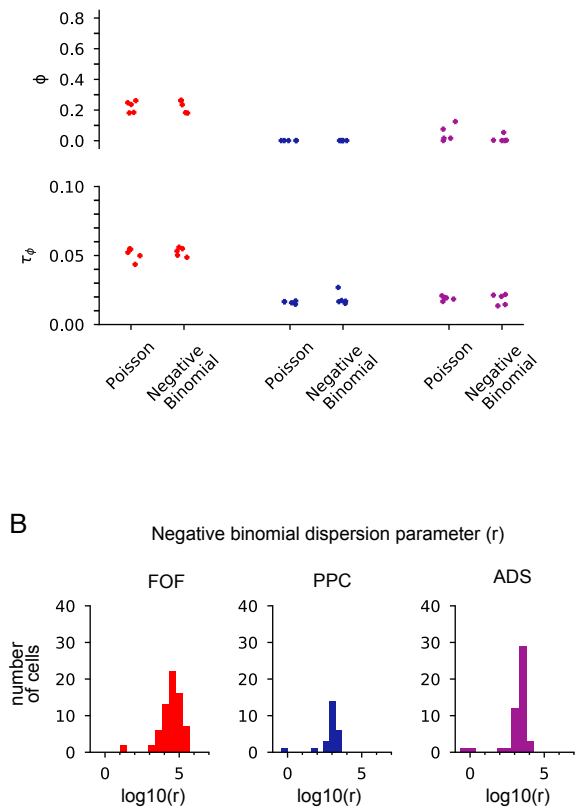

PPC



ADS









joint model

 Δ clicks














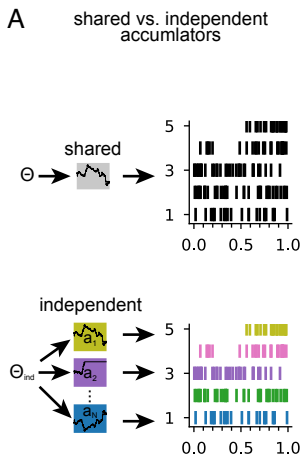
maximum likelihood parameters





A




B

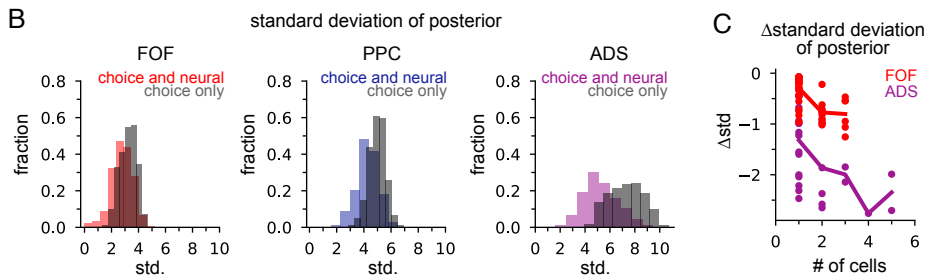
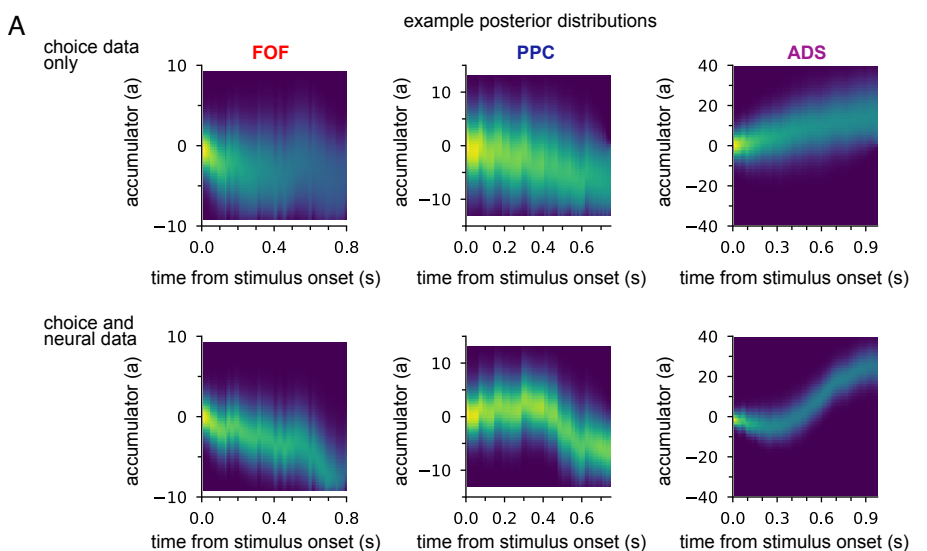










A**B****C****D**



maximum likelihood parameters








choice decoding



Aanimal
accuracy**B**joint model
accuracy