Report

Current Biology

Decision Making through Integration of Sensory Evidence at Prolonged Timescales

Highlights

- A new task quantifies the computational properties of prolonged deliberation
- Humans accurately integrate sensory information over protracted durations
- Memory noise is minimal, and integration time constants exceed tens of seconds
- Results put constraints on plausible models of integration in biophysical networks

Authors

Michael L. Waskom, Roozbeh Kiani

Correspondence

mwaskom@nyu.edu (M.L.W.), roozbeh@nyu.edu (R.K.)

In Brief

Waskom and Kiani quantitatively evaluate human decision-making behavior using a new psychophysical task that extends deliberation times to previously untested durations. They show that sensory-guided choice reflects linear integration of evidence with minimal loss of information over time, violating predictions from many biophysical models.

Waskom & Kiani, 2018, Current Biology *28*, 1–7
December 3, 2018
https://doi.org/10.1016/j.cub.2018.10.021

CellPress

Please cite this article in press as: Waskom and Kiani, Decision Making through Integration of Sensory Evidence at Prolonged Timescales, Current Biology (2018), https://doi.org/10.1016/j.cub.2018.10.021

Current Biology
Report

CellPress

Decision Making through Integration of Sensory Evidence at Prolonged Timescales

Michael L. Waskom[1,*] and Roozbeh Kiani[1,2,3,4,*]
[1]Center for Neural Science, New York University, 4 Washington Pl, New York, NY 10003, USA
[2]Neuroscience Institute, NYU Langone Medical Center, 550 First Avenue, New York, NY 10016, USA
[3]Department of Psychology, New York University, 4 Washington Pl, New York, NY 10003, USA
[4]Lead Contact
*Correspondence: mwaskom@nyu.edu (M.L.W.), roozbeh@nyu.edu (R.K.)
https://doi.org/10.1016/j.cub.2018.10.021

SUMMARY

When multiple pieces of information bear on a decision, the best approach is to combine the evidence provided by each one. Evidence integration models formalize the computations underlying this process [1–3], explain human perceptual discrimination behavior [4–9], and correspond to neuronal responses elicited by discrimination tasks [10–14]. These findings suggest that evidence integration is key to understanding the neural basis of decision making [15–18]. But while evidence integration has most often been studied with simple tasks that limit deliberation to relatively brief periods, many natural decisions unfold over much longer durations. Neural network models imply acute limitations on the timescale of evidence integration [19–23], and it is currently unknown whether existing computational insights can generalize beyond rapid judgments. Here, we introduce a new psychophysical task and report model-based analyses of human behavior that demonstrate evidence integration at long timescales. Our task requires probabilistic inference using brief samples of visual evidence that are separated in time by long and unpredictable gaps. We show through several quantitative assays how decision making can approximate a normative integration process that extends over tens of seconds without accruing significant memory leak or noise. These results support the generalization of evidence integration models to a broader class of behaviors while posing new challenges for models of how these computations are implemented in biological networks.

RESULTS

The normative basis of the evidence integration framework and its ability to explain perceptual discrimination behavior suggest that its principles are generally relevant to understanding decision making. But at present, potential limitations on the practicable timescale of integration are a critical obstacle to this generalization. While perceptual discrimination tasks afford tight experimental control and embody many important aspects of decision making, they rarely demand integration over durations that are characteristic of complex natural behaviors. Humans can deliberate about natural decisions for many seconds—or even much longer—and often consider multiple discrete pieces of information before committing to a choice. Making such decisions through evidence integration would require maintaining a high-fidelity representation over time and updating it in response to new information. Biophysical models show how this can be achieved in neural networks over relatively short timescales, but they imply acute limitations on prolonged deliberation [20, 23].

We developed a new psychophysical paradigm to quantify temporal limitations on evidence integration (Figure 1). Our paradigm builds on the success of established tasks by providing evidence to the subject in the form of simple visual stimuli that can be parametrically controlled with high precision and whose basic encoding and representation in sensory cortex are well understood. Instead of displaying these stimuli in a continuous stream, we presented multiple brief exposures with variable intensity that were separated in time by long and unpredictable gaps. This simple manipulation produced a long and variable timescale of deliberation, extending it on many trials to tens of seconds (mean trial duration, 10.1 ± 5.6 s; range, 2.2–34 s). We leveraged our experimental control over the evidence and the variability in its strength, quantity, and timing to identify the computations underlying decision-making behavior.

We trained five human subjects on this task until they reached an accuracy criterion (> 76% correct, achieved over 2–4 sessions) and then collected multiple sessions of data for analysis. We framed this analysis with three questions. First, did the subjects combine information from multiple samples before reaching a commitment? Second, did they use the graded quantity of evidence provided by each sample? Third, was their performance limited by loss of information over time?

To answer these questions, we evaluated several different computational models that enable quantification of the decision-making process (Figure S1). This quantitative approach is necessary because behavior can qualitatively resemble integration while lacking its key characteristics. We focus here on four specific models: one implements the normative policy and serves as a baseline for comparison, while the others

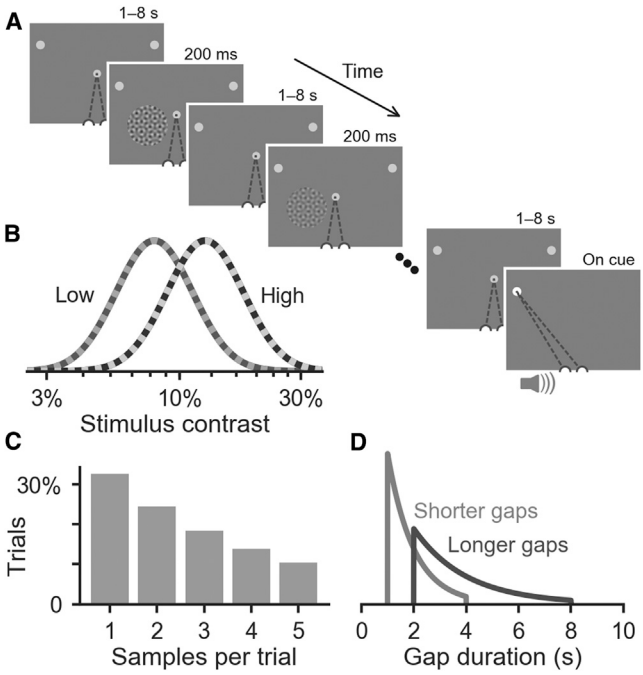

Figure 1. Experimental Design

(A) Subjects viewed brief samples of a contrast pattern while maintaining central fixation. They were cued at the end of the trial to report their decision by making a saccade to one of two targets and received feedback about the accuracy of their choice.

(B) Each sample had a different contrast, randomly drawn from one of two overlapping Gaussian distributions in log contrast space. Each trial was generated using samples from the same distribution; the subject's task was to infer which one.

(C) 1–5 samples were shown before cuing a response, determined by drawing from a truncated geometric distribution.

(D) Each sample was followed by a gap lasting either 1–4 s (shorter gap sessions) or 2–8 s (longer gap sessions), determined by drawing from one of two truncated exponential distributions.

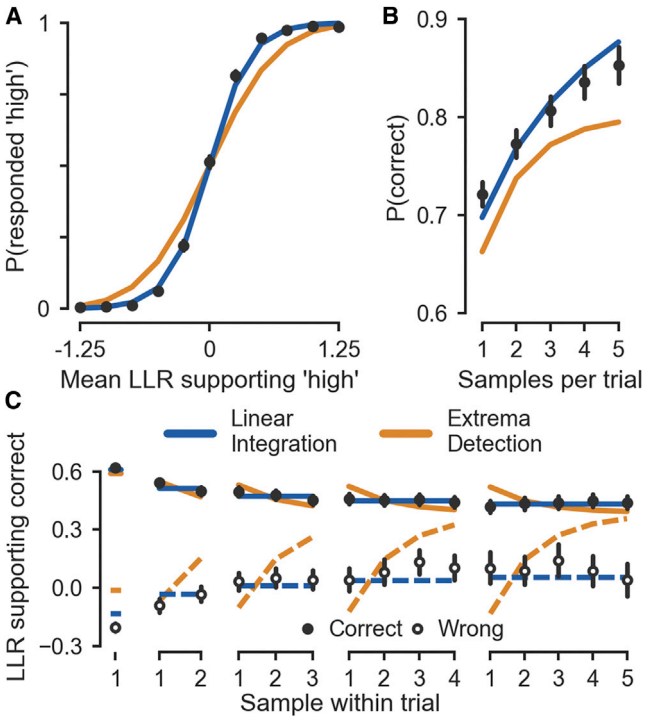

Figure 2. Integration of Evidence across Samples

(A) Sample mean psychometric function (mPMF), showing the relationship between mean strength of evidence supporting a choice of "high" and the probability of making that choice.

(B) Sample count psychometric function (cPMF), showing the relationship between the number of samples in a trial and the probability of making a correct choice.

(C) Reverse correlation functions (RCFs) shown separately for trials with different sample counts.

In all panels, black points and error bars show means and bootstrap 95% CIs, and blue and gold lines show analytic functions from the best-fitting linear integration and extrema detection models, respectively. All panels show aggregate data and model fits; see Figure S1 for model predictions and Figure S2 for individual data and fits.

establish the presence of integration and quantify its limitations. The models span a space of mechanisms within the general sequential sampling framework. They operate by using each sample, denoted x, to update a decision variable, denoted V. The evidence is quantified in terms of the log-likelihood ratio (LLR) that the sample was generated from the high-contrast distribution, and the sign of the decision variable at the end of the trial determines the choice.

Integration of Evidence across Samples

Optimal performance in the task can be achieved by summing the evidence afforded by each sample in units of LLR with constant weighting across time. This computation can be formalized in a "linear integration" model defined by the following update equation:

$$V_i = V_{i-1} + x_i + \xi_\eta \quad (1)$$

where V_i is the decision variable after observing sample i, and $V_0 = 0$. Choice variability arises because the internal representation of each sample is corrupted by noise, represented by ξ_η (Figure S1A). This noise is stimulus-dependent, but it may arise from sources beyond sensory encoding [4, 24, 25]. Assuming Gaussian noise, the only free parameter is σ_η, the standard deviation of ξ_η in units of LLR. The decision variable updates are indexed by the ordinal sample number, as this model does not use information about the duration of the gaps. Linear integration is optimal in that it is limited only by variability in stimulus generation and noisy encoding of each sample; no other information is distorted or lost during deliberation [1–3].

Using this model, we derived analytic expressions for three assays of decision-making behavior (aggregate performance: Figure 2; individual performance: Figure S2; blue lines). Details of the mathematical derivations are provided in the STAR Methods section. The first behavioral assay shows how choice depends on the mean strength of the samples ("sample mean psychometric function" or "mPMF"; Equation 9; Figures 2A and S2A). It demonstrates high behavioral sensitivity to the evidence. The second assay shows how accuracy depends on the number of samples in a trial ("sample count psychometric function" or "cPMF"; Equation 10; Figures 2B and S2B). It demonstrates that performance benefits with more samples. The third assay

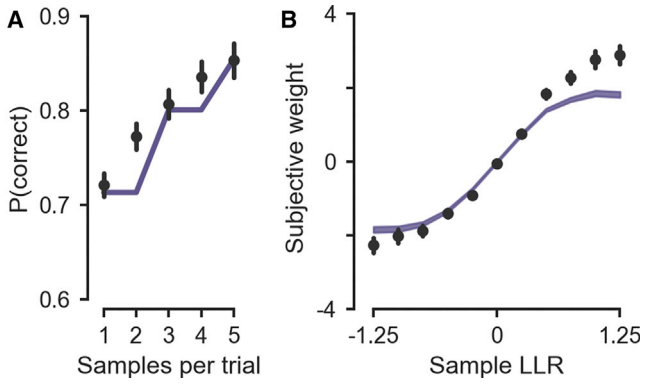

Figure 3. Integration of Graded Stimulus Evidence

(A) Data and model cPMFs, showing how the data deviate from a qualitative signature of binarized evidence transformation (counting model). Black points and error bars show means and bootstrap 95% CIs; purple line shows analytic prediction of the best-fitting counting model.

(B) Estimated subjective weighting of samples with different evidence values. Black points and lines show logistic regression coefficients and 95% CIs for the behavioral data; purple band shows logistic regression coefficients and 95% CIs for simulated data from the counting model with parameters that best fit the choices.

Both panels show aggregate data and model fits; see Figure S1 for model predictions and Figure S3 for individual data and fits.

relates choice accuracy to stochastic fluctuations in information across samples ("reverse correlation function" or "RCF"; Equation 11; Figures 2C and S2C). It demonstrates approximately constant weighting of the evidence over time. Each effect was replicated in individual subjects (Figure S2), suggesting that integration of evidence can explain both average behavior and individual idiosyncrasies.

Based on these three assays, the choice data appear generally consistent with integration. But minor deviations from the linear integration model (e.g., a shallower cPMF) imply that it is not a full account of the decision-making process. While several different factors could produce these deviations (see Discussion), a primary concern is whether the data can be explained by a process that lacks integration altogether. For example, a network with a relatively short integration time constant but the ability to maintain a discrete choice state in working memory could perform our task by sequentially sampling before deciding based on a single stimulus. We evaluate this approach with a model that compares each sample individually against a detection threshold, ignoring samples prior to and following a commitment. The process can be termed "extrema detection" [26] and leads to the following update equation:

$$V_i = \begin{cases} +1, & \text{if } V_{i-1} = 0 \text{ and } x_i + \xi_\eta > +\theta_x \\ -1, & \text{if } V_{i-1} = 0 \text{ and } x_i + \xi_\eta < -\theta_x \\ V_{i-1}, & \text{if } V_{i-1} \neq 0 \text{ or } -\theta_x < x_i + \xi_\eta < +\theta_x. \end{cases} \quad (2)$$

As in the linear integration model, ξ_η represents Gaussian noise with standard deviation σ_η; the other free parameter, θ_x, represents the threshold.

Despite using fundamentally different computations, extrema detection can superficially mimic evidence integration: both the strength and number of samples will influence choice (Figure S1B). But fitting the model to the choice data shows that it cannot explain the behavior. The linear integration model had a higher cross validated log-likelihood (aggregate ΔLL_{CV}: 788.7; individual ΔLL_{CV}: 95.2–288.4), and extrema detection could not account for performance in any of the behavioral assays (Figures 2 and S2; gold lines). Therefore, the observed behavioral performance could not have been obtained using only individual samples, implying that evidence was integrated across the gaps.

Integration of Graded Stimulus Evidence

The linear integration model operates on an untransformed representation of the stimulus evidence and stores the decision variable in an analog format. However, the long gaps in our experiment might encourage alternate strategies that would permit discretized storage. For example, the subject could determine whether each sample appears higher or lower than average, maintaining only the outcome of these decisions over time. Doing so would cause the decision variable to reflect a discretized count, implementing a linear integration process over binarized representations of the evidence. This "counting" model can be defined with the following update equation, using the same terms as in Equation 1:

$$V_i = V_{i-1} + \text{sign}\left(x_i + \xi_\eta\right). \quad (3)$$

The counting model is structurally similar to the linear integration model (compare Equations 1 and 3) and makes broadly similar predictions (Figure S1C). Its likelihood was only moderately lower than linear integration (aggregate ΔLL_{CV}: 71.5; individual ΔLL_{CV}: −0.5–27.6). Despite this similarity, the counting model exhibits a distinctive qualitative signature that is not present in the data. Representing the decision variable with a discrete count leads to a tie on many trials with an even number of samples, requiring a random guess to generate a choice. As a result, the model's expected accuracy does not improve relative to trials with the next smallest odd number of samples, a prediction that is independent of the noise parameter (Figure S1C; Equation 37). In contrast, the data exhibit a clear improvement in accuracy with additional samples (Figures 3A and S3B; Equation 6; aggregate model: $\beta_1 = 0.25 \pm 0.04$; $p < 10^{-8}$; individual models: 5/5 subjects $p < 0.05$).

Potential loss of information from a binarized representation can also be evaluated by estimating the subjective weight of evidence afforded by samples with different objective strengths (Equation 5). In the counting model, discretization curtails the influence of strong samples. By comparing subjective weights estimated from choice data to weights obtained by simulating choices from the best-fitting counting model (Figures 3B and S3C), it is apparent that strong samples contributed more to the decision than would be expected for counting. This analysis does suggest moderate compression of the evidence from strong samples, although this would not be inconsistent with linear integration in general (see STAR Methods). With respect to the decision-making process itself, we conclude that the data are best explained by integration of graded evidence across multiple samples.

Minimal Influence of Memory Leak or Noise

The normative model in Equation 1 does not require any information about the timing of the sample presentations or the duration

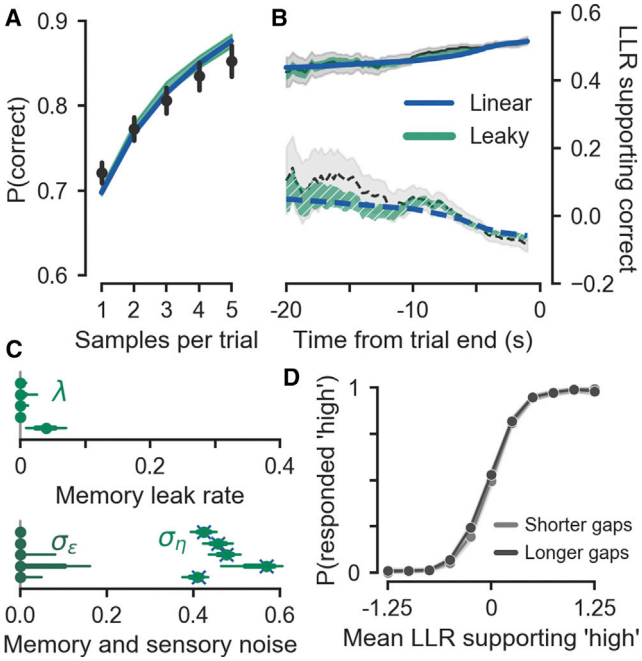

Figure 4. Minimal Influence of Memory Leak or Noise

(A) Data and model cPMFs. Black points and error bars show means and bootstrap 95% CIs for the behavioral data, blue line shows analytic function for the best-fitting linear integration model, and green band shows simulated performance for the best-fitting leaky integration model.

(B) Data and model RCFs, aligned to the end of the trial. Element colors are as in (A); solid lines show functions conditioned on correct choices and dashed lines show functions conditioned on incorrect choices. Note the apparent nonlinearity in the linear integration model RCF, an artifact of combing trials with different sample counts (see STAR Methods).

(C) Integration model parameter fits. Green elements show parameters for the leaky integration model: points show maximum likelihood estimates; thick and thin error bars show bootstrap 68% and 95% CIs, respectively. Blue crosses show σ_η estimated with the linear integration model.

(D) Data mPMF plotted separately for trials with shorter and longer gaps between samples.

All panels show aggregate data and model fits; see Figure S1 for model predictions and Figure S4 for individual data and fits.

of the gaps between them. We expected, however, that the prolonged deliberation timescale would make two limitations on working memory apparent [20, 23]. First, information from earlier samples would exert less influence because it would "leak" from memory. Second, additional behavioral variability would be caused by integration of diffusive noise in the absence of sensory input. To quantify the influence of these limitations, we extended Equation 1 into the temporal domain and explicitly modeled them (Figure S1D). This "leaky integration" model can be defined by the following update equation:

$$V(t) = V(t - \Delta t) - \Delta t \lambda V(t - \Delta t) + x(t) + I(t)\xi_\eta + \xi_\varepsilon. \quad (4)$$

In this equation, λ represents the memory leak rate (the inverse of the integration time constant τ) in units of s^{-1}. $x(t)$ represents the strength of evidence at time t and is 0 during the gaps. $I(t)$ is an indicator variable that governs the influence of stimulus-dependent noise, $\xi_\eta \sim \mathcal{N}(0, \sigma_\eta)$, and $\xi_\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon)$ represents memory noise, which accumulates throughout the gaps.

Surprisingly, fitting the leaky integration model to choice behavior revealed minimal influence of these factors. After optimizing the free parameters to maximize the likelihood of the choices, we simulated data for comparison to the behavioral dataset and fits of the normative linear integration model. The best-fitting leaky integration model could not be distinguished from linear integration in terms of its simulated cPMF (Figures 4A and S4A) or RCF (Figures 4B and S4B), and a quantitative comparison only marginally favored the full model (aggregate ΔLL_{CV}: −1.4; aggregate likelihood ratio test: χ^2_2 = 10.1; p = 0.006; individual likelihood ratio tests p < 0.05 in only 1/5 subject). The best-fitting leak rates and memory noise magnitudes were both close to 0 (Figure 4C), implying an integration time constant that was longer than 25 s for all subjects and effectively infinite for 4 out of 5 subjects.

Minimal loss of information can also be seen using a model-free approach. Subjects performed the task in two different conditions where the gaps were sampled from distributions with different minimum, maximum, and mean durations (Figure 1D). Despite large differences in gap durations, performance did not appear to substantially vary (Figures 4D and S4C). Statistically, behavioral sensitivity to the strength of evidence was similar (Equation 7; aggregate model: β_3 = 0.29 ± 0.18; p = 0.11; individual models: p > 0.05 in 4/5 subjects and p = 0.03 in S4). In fact, behavior in one of the conditions could be used to predict behavior in the other. These results further emphasize the functional invariance of the decision-making process to the timescale of deliberation, implying the existence of a highly flexible and efficient system for combining the evidence from each sample to make decisions at a level approaching normative performance.

DISCUSSION

Evidence integration models reflect a synthesis of computational insights spanning cognitive and biological levels of analysis [8, 27]. They have been tested in a variety of experimental contexts, including perceptual discrimination of continuous [10, 11, 28] or pulsed [4, 12, 29, 30] stimuli and probabilistic inferences similar to those in our experiment [13, 24, 31, 32]. However, because previous tasks had relatively short timescales, there has been persistent concern about the scope of behaviors to which these insights apply.

Behavioral performance and neural responses in existing tasks can be explained by network models that overcome biophysically limited time constants of individual neurons through emergent attractor dynamics [3, 19, 21, 33]. These networks form an important bridge between cognitive theory and biological data. However, existing models have specific failure modes that would be apparent at prolonged timescales [3, 23]. Rapid discrimination tasks are often modeled by networks with bistable point attractor dynamics [21, 34]. In the prolonged absence of input, however, a decision variable computed this way would either decay back to its starting point or converge onto a terminal decision state [35]. Alternatively, maintenance and manipulation of analog information can be achieved using line attractor dynamics [36, 37], even without persistent input. But it is widely thought that line attractor networks require fine tuning to avoid memory leak and that even small amounts of memory noise

would accumulate to dominate the representation over prolonged timescales [19, 20, 23, 38].

None of these failure modes emerged when we pushed the decision-making process an order of magnitude beyond the timescale of conventional tasks. This suggests either that the intrinsic timescale of biological networks is much longer than currently thought [22] or that network models require augmentation. Possible accommodations could involve additional mechanisms that add robustness to leak and noise [39, 40] or different, functionally feedforward network architectures [41]. More radically, prolonged integration could recruit fundamentally distinct neural mechanisms to perform functionally similar computations. Might evidence integration engage a long-term memory system? As this question comes into focus, we note that prolonged yet controlled deliberation would make the integration process amenable to study with functional MRI, complementing previous investigations of gradual perceptual judgments [42, 43].

Integration without leak is the normative policy in our task, but forgetting can be adaptive in nonstationary environments. Previous work has suggested that humans can adopt a leak rate that matches environmental statistics when performing signal [44] or change-point [45] detection. Our results show that it is also possible to eliminate leak in situations where it is optimal to do so. By ruling out fundamental limitations on integration, our findings may also prompt a re-interpretation of recency bias, which is sometimes observed in similar tasks [24, 46, 47]. Recency bias is often attributed to memory leak, but absence of leak at prolonged timescales suggests that these effects, when present, might be better explained by decision strategy [48].

While behavior in our task was approximately normative, small but systematic deviations from the model fits imply that linear integration cannot fully explain the data. Figure 2 suggests that suboptimal performance arises from modest over-weighting of early samples. This would be consistent with the presence of a decision bound that can terminate integration prior to the response cue [5, 8, 11, 27, 49]. We do not quantitatively evaluate bounded integration because, despite the presence of these qualitative signatures, our experimental design makes the height of the bound difficult to estimate. The asymmetry between the RCF for correct and error trials also suggests that a bound, if present, enacts only provisional commitments that can be revised by further evidence [50], analagous to a confirmation bias [32, 51]. Variants of our task that accentuate the influence of these factors may afford computational modeling of their origin and provide insight into suboptimal decision-making in naturalistic contexts.

Our results quantify the influence of memory leak and noise on prolonged integration, demonstrating that neither pose a fundamental challenge for naturalistic decisions. While surprising given the biophysical properties of neurons, these findings create an opportunity to generalize the powerful evidence integration framework beyond simple perceptual judgments. Many open questions remain about the computational and neural basis of natural decision-making. Answering them will require formal models that can make connections across levels of analysis to explain how biological systems generate complex behaviors. Evidence integration is a promising candidate for these efforts.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- CONTACT FOR REAGENT AND RESOURCE SHARING
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
- QUANTIFICATION AND STATISTICAL ANALYSIS
  - Data analysis
  - General computational modeling approach
  - Linear integration model
  - Extrema detection model
  - Counting model
  - Leaky integration model
- DATA AND SOFTWARE AVAILABILITY

SUPPLEMENTAL INFORMATION

Supplemental Information includes four figures and can be found with this article online at https://doi.org/10.1016/j.cub.2018.10.021.

ACKNOWLEDGMENTS

R.K. is supported by the NIMH (R01-MH109180), the Simons Collaboration on the Global Brain (542997), and a Pew Scholarship in the Biomedical Sciences. M.L.W. is supported by the Simons Foundation as a Junior Fellow in the Simons Society of Fellows (527794).

AUTHOR CONTRIBUTIONS

M.L.W and R.K. designed the experiment, collected the data, developed the computational models, analyzed the data, and wrote the paper.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: August 4, 2018
Revised: September 19, 2018
Accepted: October 8, 2018
Published: November 21, 2018

REFERENCES

1. Neyman, J., and Pearson, E.S. (1933). On the Problem of the Most Efficient Tests of Statistical Hypotheses. Phil. Trans. R. Soc. A *231*, 289–337.
2. Wald, A., and Wolfowitz, J. (1948). Optimum Character of the Sequential Probability Ratio Test. Ann. Math. Stat. *19*, 326–339.
3. Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765.
4. Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and humans can optimally accumulate evidence for decision-making. Science *340*, 95–98.
5. Link, S.W. (1992). The Wave Theory of Difference and Similarity (Hillsdale, NJ: Erlbaum).
6. Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. *5*, 376–404.
7. Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. *20*, 873–922.

8. Smith, P.L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. Trends Neurosci. *27*, 161–168.
9. Vickers, D. (1979). Decision Processes in Visual Perception (New York: Academic Press).
10. Mazurek, M.E., Roitman, J.D., Ditterich, J., and Shadlen, M.N. (2003). A role for neural integrators in perceptual decision making. Cereb. Cortex *13*, 1257–1269.
11. Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. J. Neurosci. *28*, 3017–3029.
12. Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223.
13. Kira, S., Yang, T., and Shadlen, M.N. (2015). A neural implementation of Wald's sequential probability ratio test. Neuron *85*, 861–873.
14. Scott, B.B., Constantinople, C.M., Akrami, A., Hanks, T.D., Brody, C.D., and Tank, D.W. (2017). Fronto-parietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. Neuron *95*, 385–398.e5.
15. Platt, M.L. (2002). Neural correlates of decisions. Curr. Opin. Neurobiol. *12*, 141–148.
16. Glimcher, P.W. (2003). The neurobiology of visual-saccadic decision making. Annu. Rev. Neurosci. *26*, 133–179.
17. Heekeren, H.R., Marrett, S., and Ungerleider, L.G. (2008). The neural systems that mediate human perceptual decision making. Nat. Rev. Neurosci. *9*, 467–479.
18. Shadlen, M.N., and Kiani, R. (2013). Decision making as a window on cognition. Neuron *80*, 791–806.
19. Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. Psychol. Rev. *108*, 550–592.
20. Brody, C.D., Romo, R., and Kepecs, A. (2003). Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. Curr. Opin. Neurobiol. *13*, 204–211.
21. Wang, X.-J. (2008). Decision making in recurrent neuronal circuits. Neuron *60*, 215–234.
22. Murray, J.D., Bernacchia, A., Freedman, D.J., Romo, R., Wallis, J.D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., and Wang, X.J. (2014). A hierarchy of intrinsic timescales across primate cortex. Nat. Neurosci. *17*, 1661–1663.
23. Chaudhuri, R., and Fiete, I. (2016). Computational principles of memory. Nat. Neurosci. *19*, 394–403.
24. Drugowitsch, J., Wyart, V., Devauchelle, A.-D., and Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. Neuron *92*, 1398–1411.
25. Li, V., Herce Castañón, S., Solomon, J.A., Vandormael, H., and Summerfield, C. (2017). Robust averaging protects decisions from noise in neural computations. PLoS Comput. Biol. *13*, e1005723.
26. Stine, G., Zylberberg, A., and Shadlen, M. (2018). Disentangling evidence integration from memoryless strategies in perceptual decision making (Denver, CO, USA: Cosyne Abstracts).
27. Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.
28. Huk, A.C., and Shadlen, M.N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. J. Neurosci. *25*, 10420–10436.
29. Scott, B.B., Constantinople, C.M., Erlich, J.C., Tank, D.W., and Brody, C.D. (2015). Sources of noise during accumulation of evidence in unrestrained and voluntarily head-restrained rats. eLife *4*, e11308.
30. Pinto, L., Koay, S.A., Engelhard, B., Yoon, A.M., Deverett, B., Thiberge, S.Y., Witten, I.B., Tank, D.W., and Brody, C.D. (2018). An Accumulation-of-Evidence Task Using Visual Pulses for Mice Navigating in Virtual Reality. Front. Behav. Neurosci. *12*, 36.
31. de Lange, F.P., Jensen, O., and Dehaene, S. (2010). Accumulation of evidence during sequential decision making: the importance of top-down factors. J. Neurosci. *30*, 731–738.
32. Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., de Gardelle, V., Herce Castañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. Neuron *81*, 1429–1441.
33. Lim, S., and Goldman, M.S. (2013). Balanced cortical microcircuitry for maintaining information in working memory. Nat. Neurosci. *16*, 1306–1314.
34. Wang, X.-J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. Neuron *36*, 955–968.
35. Wong, K.-F., and Wang, X.-J. (2006). A recurrent network mechanism of time integration in perceptual decisions. J. Neurosci. *26*, 1314–1328.
36. Seung, H.S., Lee, D.D., Reis, B.Y., and Tank, D.W. (2000). Stability of the memory of eye position in a recurrent network of conductance-based model neurons. Neuron *26*, 259–271.
37. Mante, V., Sussillo, D., Shenoy, K.V., and Newsome, W.T. (2013). Context-dependent computation by recurrent dynamics in prefrontal cortex. Nature *503*, 78–84.
38. Burak, Y., and Fiete, I.R. (2012). Fundamental limits on persistent activity in networks of noisy neurons. Proc. Natl. Acad. Sci. USA *109*, 17645–17650.
39. Koulakov, A.A., Raghavachari, S., Kepecs, A., and Lisman, J.E. (2002). Model for a robust neural integrator. Nat. Neurosci. *5*, 775–782.
40. Goldman, M.S., Levine, J.H., Major, G., Tank, D.W., and Seung, H.S. (2003). Robust persistent neural activity in a model integrator with multiple hysteretic dendrites per neuron. Cereb. Cortex *13*, 1185–1195.
41. Goldman, M.S. (2009). Memory without feedback in a neural network. Neuron *61*, 621–634.
42. Ploran, E.J., Nelson, S.M., Velanova, K., Donaldson, D.I., Petersen, S.E., and Wheeler, M.E. (2007). Evidence accumulation and the moment of recognition: dissociating perceptual recognition processes using fMRI. J. Neurosci. *27*, 11912–11924.
43. Krueger, P.M., van Vugt, M.K., Simen, P., Nystrom, L., Holmes, P., and Cohen, J.D. (2017). Evidence accumulation detected in BOLD signal using slow perceptual decision making. J. Neurosci. Methods *281*, 21–32.
44. Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. Curr. Biol. *23*, 981–986.
45. Glaze, C.M., Kable, J.W., and Gold, J.I. (2015). Normative evidence accumulation in unpredictable environments. eLife *4*, 4.
46. Hubert-Wallander, B., and Boynton, G.M. (2015). Not all summary statistics are made equal: Evidence from extracting summaries across time. J. Vis. *15*, 5.
47. Tsetsos, K., Moran, R., Moreland, J., Chater, N., Usher, M., and Summerfield, C. (2016). Economic irrationality is optimal during noisy decision making. Proc. Natl. Acad. Sci. USA *113*, 3102–3107.
48. Summerfield, C., and Tsetsos, K. (2015). Do humans make good decisions? Trends Cogn. Sci. *19*, 27–34.
49. Tsetsos, K., Gao, J., McClelland, J.L., and Usher, M. (2012). Using Time-Varying Evidence to Test Models of Decision Dynamics: Bounded Diffusion vs. the Leaky Competing Accumulator Model. Front. Neurosci. *6*, 79.
50. Resulaj, A., Kiani, R., Wolpert, D.M., and Shadlen, M.N. (2009). Changes of mind in decision-making. Nature *461*, 263–266.
51. Talluri, B.C., Urai, A.E., Tsetsos, K., Usher, M., and Donner, T.H. (2018). Confirmation bias through selective overweighting of choice-consistent evidence. Curr. Biol. *28*, 3128–3135.e8, Epub ahead of print.
52. Peirce, J.W. (2009). Generating stimuli for neuroscience using PsychoPy. Front. Neuroinform. *2*, 10.

53. Jones, E., Oliphant, T., and Peterson, P. (2001). SciPy: Open source scientific tools for Python. Available at: http://www.scipy.org/.

54. Okazawa, G., Sha, L., Purcell, B.A., and Kiani, R. (2018). Psychophysical reverse correlation reflects both sensory and decision-making processes. Nat. Commun. *9*, 3479.

55. Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

56. Cumming, G., and Finch, S. (2005). Inference by eye: confidence intervals and how to read pictures of data. Am. Psychol. *60*, 170–180.

57. de Gardelle, V., and Summerfield, C. (2011). Robust averaging during perceptual judgment. Proc. Natl. Acad. Sci. USA *108*, 13341–13346.

58. Boynton, G.M., Demb, J.B., Glover, G.H., and Heeger, D.J. (1999). Neuronal basis of contrast discrimination. Vision Res. *39*, 257–269.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Software and Algorithms		
PsychoPy	[52]	http://www.psychopy.org/
SciPy	[53]	https://www.scipy.org/
Custom code	This paper	https://github.com/kianilab/Waskom_CurrBiol_2018

CONTACT FOR REAGENT AND RESOURCE SHARING

Further information and requests for resources and reagents should be directed to and will be fulfilled by the Lead Contact, Roozbeh Kiani (roozbeh@nyu.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

Five human subjects, (four male and one female; ages 19–40) participated in the experiment. All had normal or corrected-to-normal vision. One of the subjects was author RK; the others were naive to the purposes of the experiment. All experimental procedures were approved by the Institutional Review Board at New York University, and the subjects provided informed written consent before participating.

METHOD DETAILS

The subjects were seated in an adjustable chair in a semi-dark room with chin and forehead supported before a cathode ray tube display monitor (21 in Sony GDM-5402; 1600 × 1200 screen resolution; 75 Hz refresh rate; 8 bit color; 52 cm viewing distance). The display was calibrated with a photometer and gamma corrected to have a linear response. The mean luminance of the display during the experiment was 35 cd/m^2. Viewing was binocular. Gaze position was monitored at 1 kHz using a high-speed infrared camera (SR-Research; Ottowa, Canada). Stimulus presentation was controlled using PsychoPy [52].

The experimental task is diagrammed in Figure 1. Each trial was initiated when the subject looked at a central fixation point (0.3° diameter). 500 ms after fixation, two response targets (0.5° diameter) appeared 5° above and 10° to the left or right of fixation, and a small protuberance (0.1° length) extended from the fixation point to cue the location at which the visual stimulus would be presented on that trial.

The visual stimulus was a circular (6° diameter) contrast patch constructed as the average of 8 sine wave gratings at evenly spaced orientations. Each grating had a spacial frequency of 2 cycles per degree, and the phases of the gratings were independently randomized on each presentation. The stimulus edges (20% of the radius) were blended with the background using a raised-cosine mask. The stimulus was centered 2° below and 5.6° either to the left or right of fixation and was always compatible with the cue. Within a given trial, the stimulus would appear in the same location, but the location varied randomly between trials. Each stimulus presentation had a duration of 200 ms.

There were 1–5 stimulus presentations (“samples”) per trial, drawn from a truncated geometric distribution with p = 0.25. This distribution has a nearly flat hazard function, meaning that the probability of the trial ending is approximately constant after each sample. This makes it impossible for the subject to adopt a strategy of preferentially allocating attention to the samples immediately preceding the response. We note that this design feature makes us less likely to observe a recency bias, which might be incorrectly attributed to memory leak.

The contrast of the stimulus varied across each presentation. Stimulus contrast, C, was defined in units of Michelson contrast and was determined by setting each constituent grating to a contrast of $\sqrt{8}C$ before averaging. This procedure generated contrast patches with root-mean-square (RMS) contrast that was equivalent to single gratings drawn at the specified value of C.

The experimental task required probabilistic inference on the basis of stimulus contrast. Each trial was defined, with equal probability, as either a “low-contrast” or “high-contrast” trial. This determined which of two overlapping distributions would be used to generate the contrast for each stimulus presentation in that trial. The distributions were defined as Gaussians in $\log_{10}$ contrast space with means of -1.1 or -0.9 and standard deviations of 0.15. The subject’s task was to infer, using the perceived contrast of each stimulus presentation, whether the trial was generated using the low-contrast or high-contrast distribution.

Each stimulus presentation was separated by a gap with a duration drawn from a truncated exponential distribution. The data reported here were collected across two different experimental conditions: “shorter” gaps (1–4 s) and “longer” gaps (2–8 s). These conditions appeared in separate experimental sessions, and subjects performed the shorter gap sessions first. The gaps prior to the first and following the last stimulus presentations were determined using the same distribution as the inter-stimulus gaps.

At the end of the trial, the fixation point disappeared, which cued the subject to report their decision by making a saccadic response to one of the two response targets. After choosing one of the targets, the subject received auditory and visual feedback about the accuracy of their response. Accuracy was defined in terms of whether the subject chose the target corresponding to the distribution that was actually used to generate the stimuli for that trial. Response target assignments were stable across the experiment. The subject was allowed to blink during the trial but was otherwise required to maintain fixation prior to responding. Trials were separated by an inter-trial interval of at least 3 s.

Each subject was trained on the task for 2–4 sessions prior to collection of the data reported here, and they were required to reach a criterion level of performance (session-wise accuracy > 76%) before continuing. They additionally performed two practice sessions with longer gaps prior to collection of longer-gap data for analysis. No subject failed to reach the accuracy criterion. Experimental instructions emphasized accuracy, and subjects were explicitly informed that the best way to maximize performance would be to make their decision on the basis of what they perceived as the average contrast of all the stimuli presented in each trial. Subjects typically performed 6 blocks (7 min duration) in each session and received feedback about their accuracy after each block.

Not including training, each subject contributed 18 separate 1-hour sessions to the data reported here (6 sessions with shorter gaps and 12 sessions with longer gaps). In total, our analyses involved 14,869 trials (6,500 with shorter gaps and 8,369 with longer gaps). Control analyses showed that behavior was stable over the course of the experiment; therefore, our main results are unlikely to depend on the amount of experience subjects had with the task.

QUANTIFICATION AND STATISTICAL ANALYSIS

Data analysis

For our primary analysis of the different computational models, we characterized the data using three behavioral assays: one relating variability in behavioral responses to average stimulus strength (the sample mean psychometric function; mPMF), one relating choice accuracy to the number of evidence samples (the sample count psychometric function; cPMF), and one relating choice accuracy to stochastic variability in the stimulus over time (the reverse correlation function; RCF).

The mPMF illustrates how behavioral responses depended on the average strength of evidence in each trial. The evidence value corresponding to each stimulus contrast was defined using the log-likelihood ratio of the two generating distributions (positive evidence supports "high" and negative supports "low.") Our experimental design produces a large amount of trial-to-trial variability in the strength of available evidence. To characterize the effect of this variability, we computed the proportion of "high" choices within evenly spaced bins of trials defined by the mean evidence across samples (bin width: 0.25, except for the lowest and highest bins which were unbounded on one side).

The cPMF illustrates how accuracy depended on the number of stimulus samples that were presented on each trial. To characterize this relationship, we computed the proportion of trials with the same number of samples where the subject chose the correct target (the target corresponding to the distribution that generated the observed evidence).

The RCF provides insight into the dynamics of the decision-making process by estimating the relative leverage on choice of evidence presented at different points in time. This approach exploits the random variability in the strength of evidence afforded by each stimulus presentation. To estimate this function, we grouped trials based on response accuracy and then computed the mean strength of evidence supporting the correct target across time. Note that because we are conditioning on accuracy and not choice, the resulting function should not be considered an estimate of the sensory kernel. Instead, we are using the logic of reverse correlation to define a function that can be compared between data and model. This function represents a conditional expectation as a function of time. For some analyses, we averaged over ordinal sample positions (Figures 2C and S2C). For other analyses (Figures 4B and S4B), we used the actual timing of stimulus events by computing a running average over all stimuli occurring within a bin of width 4 s at evenly spaced time points (bin width: 0.25 s). Under linear integration, the height of the RCF depends on the square root of the sample count [54]. This causes an artifact in plots that show an RCF computed from all trials (Figures 4B and S4B). Therefore, nonlinear effects in those plots do not necessarily reflect nonlinear weighting of evidence. This can also be seen in an ordinal analysis by comparing Figures 2 and S1.

To estimate the influence of individual samples with different objective weights of evidence on the subjects' choices, we used an approach described in previous work [55]. This analysis estimates the subjective weight of evidence afforded by samples falling into evenly spaced bins (bin spacing: 0.25) using the following logistic regression model:

$$\text{logit}[P(\text{response} = \text{high})] = \sum_{k \in \text{bins}} \beta_k N_k, \quad (5)$$

Where N_k is the number of samples appearing in each trial with objective strength falling in bin k. We binned stimulus strengths because the objective values were continuous. The regression coefficient corresponding to each bin, β_k, can be interpreted as an estimate of the subjective weight of samples falling within that bin.

Figures that show these functions for behavioral data have 95% bootstrap confidence intervals [56]. Confidence intervals were computed by resampling trials with replacement and estimating the statistic of interest across 10,000 iterations. In figures showing aggregate data, bootstrapping was performed at the level of trials and was irrespective of subjects. The error bars correspond to the 2.5 and 97.5 percentiles of the resulting distribution.

We additionally performed several statistical analyses of choice behavior using logistic regression. To determine whether choice accuracy improved consistently with additional samples, we tested whether accuracy differed for trials with odd and even sample counts using the following function:

$$\text{logit}[P(\text{correct})] = \beta_0 + \beta_1 P + \beta_2 M. \quad (6)$$

In this equation, P is an indicator variable specifying the parity of the sample count on each trial (odd versus even), and M is an indicator variable specifying whether the sample count on each trial was 1–2 or 3–4. This analysis excluded trials with 5 samples because corresponding trials with 6 samples did not occur in the experiment. The null hypothesis (a prediction of the counting model) is that sample parity would have no effect ($H_0 : \beta_1 = 0$).

To evaluate differences in behavioral performance between the longer and shorter gap conditions, we used a logistic regression model to test for an interaction between the influence of the mean evidence strength and the gap condition:

$$\text{logit}[P(\text{response} = \text{high})] = \beta_0 + \beta_1 X + \beta_2 I + \beta_3 IX. \quad (7)$$

In this equation, X is the mean strength of evidence across samples in a trial, and I is an indicator variable specifying the gap duration condition. The null hypothesis was that the influence of stimulus evidence on choice did not vary across gap duration conditions ($H_0 : \beta_3 = 0$).

General computational modeling approach

We identified the computational properties of the decision-making process by fitting and evaluating several quantitative models. The models spanned a space of computations within the general sequential sampling framework, sharing a common structure but differing in ways that provided leverage on our three main questions about the influence of temporal prolongnation and discontinuity.

Our general approach was to fit the free parameters of each model using the sequence of stimuli and the subject's response on each trial and then to evaluate the optimized model performance using the behavioral assays described above. Where feasible, we derived analytic expressions for model performance on the behavioral assays. When analytic solutions were too complex or not possible to derive, we relied on Monte Carlo simulations. In this section, we provide general expressions for the model-fitting procedure and the predictions about the behavioral assays. In subsequent sections, we derive the specific equations used for each model and assay.

Model fitting was performed in the maximum likelihood framework. For each model, we found the parameter set θ that maximized the log-likelihood of observing the set of responses R given the set of stimulus sequences S across all trials. Let x_i represent the evidence afforded by the ith stimulus presentation on trial j, and let r_j represent the subject's response ("low" or "high") on that trial. In practice, we need to derive only the probability of one of the responses because $P(r_j = \text{low}) = 1 - P(r_j = \text{high})$. Assuming independence between trials, the log-likelihood of the data under each model is

$$\ln\mathcal{L}(R \mid S; \theta) = \sum_{j=1}^{m} \ln[P(r_j \mid x_1 \ldots x_n; \theta)]. \quad (8)$$

This function was numerically optimized for each model using the Nelder-Mead algorithm as implemented in scipy [53].

To obtain confidence intervals on the optimized log-likelihood and parameter values, we bootstrapped the maximum likelihood estimation procedure. On each of 1000 iterations, we randomly resampled trials (with replacement) before re-estimating the model parameters. The resampling procedure maintained the proportion of trials with each number of total samples. To obtain confidence intervals on the model parameter values, we used percentiles of the bootstrap distribution (i.e., for 95% CIs, we used the 2.5 and 97.5 percentile of the distribution).

We quantitatively compared the fit of the normative linear integration model to each of the other three candidates in terms of differences in the cross-validated log-likelihood ($\Delta\text{LL}_{\text{CV}}$) of the choice data after optimizing the parameters. The $\Delta\text{LL}_{\text{CV}}$ measure was defined so that positive scores indicated higher likelihoods for the linear integration model. To obtain cross-validated log-likelihood scores, we performed k-fold cross-validation by splitting the behavioral datasets by experimental session (k = 18). Aggregate $\Delta\text{LL}_{\text{CV}}$ scores were calculated by summing the log-likelihoods from the model corresponding to each individual subject before taking the difference. Because the linear integration model is nested within the leaky integration model (they are equivalent if λ and σ_ε both equal 0), we also used a likelihood ratio test to compare these models.

We further evaluated model performance in terms of how well model expressions for the three behavioral assays described in the previous section corresponded to the behavior of the subjects. Model expressions for the behavioral assays were derived using only the optimized model parameters and the generating statistics of the task, not the specific stimulus sequences on each trial.

In the following equations, μ_x and σ_x represent, respectively, the mean and standard deviation of the evidence distribution for the "high" condition in units of LLR. Note that the distributions for the two conditions are symmetric around LLR = 0, so the generating mean for the "low" condition is $-\mu_x$. Sometimes we will work only with the positive distribution corresponding to μ_x, and at other times we will work with the set $D = \{-\mu_x, +\mu_x\}$. Finally, let N represent the distribution of sample counts across trials, and let n represent a specific count.

To obtain a model mPMF, we compute separate predictions for each generating distribution and sample count, and we then combine them into the full prediction using a weighted sum. Therefore, the mPMF can be expressed, in general, as

$$P(r = \text{high} \mid \bar{x}, D, N, \theta) = \sum_{n \in N, \mu_x \in D} P(N = n) P(D = \mu_x) P(r = \text{high} \mid \bar{x}, n, \mu_x, \sigma_x, \theta). \tag{9}$$

The model cPMF can be expressed, in general, as

$$P(\text{correct} \mid n, \mu_x, \sigma_x, \theta). \tag{10}$$

To generate a model RCF, we take the same approach as with the mPMF: we compute separate predictions for trials with different sample counts and then combine them into a full prediction. While the two psychometric functions represent probabilities, the reverse correlation represents a conditional expectation of stimulus evidence given response accuracy, *C*. Because we are conditioning on accuracy, and not choice, the RCF is defined in terms of evidence supporting the correct alternative, which we denote as x_i', where $x_i' = x_i$ for the "high" distribution and $x_i' = -x_i$ for the "low" distribution. Therefore, the RCF can be expressed, in general, as

$$E\left(x_i' \mid C, N, \mu_x, \sigma_x, \theta\right) = \sum_{n \in N} P(N = n) E\left(x_i' \mid C, n, \mu_x, \sigma_x, \theta\right), \tag{11}$$

where *N* is the set of sample counts across trials, and *C* defines whether the RCF is for a correct or error response.

To draw a time-resolved RCF, we took a weighted average of these values across ordinal sample positions where the weights corresponded to the probability that a sample at time *t* was the *i*th sample in that trial. Probabilities were estimated empirically from the behavioral dataset.

In the following sections, we derive specific expressions for the forms of these functions under the assumptions of different models. Except where noted, the model predictions were computed using numerical integration algorithms as implemented in scipy.

Linear integration model

In the linear integration model, each stimulus affords a quantity of evidence toward the decision, which is noisily encoded, and the perceived evidence afforded by different samples are summed. The summed evidence represents a decision variable, *V*, which is defined in Equation 1. The response is determined by the sign of the decision variable at the end of the trial:

$$r_j = \begin{cases} \text{high} & \text{if } V_n > 0 \\ \text{low} & \text{if } V_n < 0. \end{cases} \tag{12}$$

It is important to note that, in this model, linearity applies to the integration process. Our main question concerns a general class of models with this property, including those in which each sample of evidence undergoes a nonlinear transformation during encoding. In this section, we define a model with untransformed evidence. We also consider a special case of nonlinear transformation when we evaluate the counting model. An intermediate subclass of models, which non-linearly compress evidence from strong samples, has also been proposed to explain behavior in rapid perceptual averaging tasks [25, 57]. Some features of our data (Figure 3B) appear consistent with this proposal. Similarly, complexities in contrast perception [58] could lead to a nonlinear transformation of evidence during encoding. Further understanding of these issues will be necessary for developing a complete account of decision-making. We do not pursue them here because they are orthogonal to our main question, which concerns the properties and limitations of the integration process itself.

Assuming that the noise is independent for each sample, the decision variable *V* can be represented by a Gaussian distribution with a mean at the sum of the evidence, $\sum x_i$, and a standard deviation that grows with the square root of the number of samples, *n*. Integrating over the possible values of accumulated noise provides the probability of choosing each alternative (*cf.* Equation 8):

$$\begin{aligned} P(r_j = \text{high} \mid x_1 \ldots x_n, \sigma_\eta) &= P(V_n > 0 \mid x_1 \ldots x_n, \sigma_\eta) \\ &= \int_0^\infty \mathcal{N}\left(\sum x_i, \sqrt{n}\,\sigma_\eta\right) \end{aligned} \tag{13}$$

where $\int_a^b \mathcal{N}(\mu, \sigma)$ corresponds to the integral of a one-dimensional Gaussian distribution with mean μ and standard deviation σ taken with respect to the one-dimensional variable whose probability is defined by the distribution.

Equation 13 shows that the choice probability given a particular sequence of samples depends only on the observation noise, number of samples, and mean evidence ($\bar{x} = \sum x_i / n$). This establishes the mPMF as a key assay for comparison of behavior with the linear integration model (*cf.* Equation 9):

$$\begin{aligned} P(r = \text{high} \mid \bar{x}, n, \sigma_\eta) &= P(V_n > 0 \mid \bar{x}, n, \sigma_\eta) \\ &= \int_0^\infty \mathcal{N}\left(\bar{x}, \frac{\sigma_\eta}{\sqrt{n}}\right). \end{aligned} \tag{14}$$

Because the cPMF is defined based on accuracy, it depends on both the stimulus distribution and the stimulus-dependent noise. It can be computed from the expected distribution of decision variables. The mean of the decision variable distribution scales linearly with the number of samples, and its width is determined by summing the variance attributable to the random sampling of stimuli with the variance of the accumulated noise (*cf.* Equation 10):

$$P(\text{correct} \mid n, \mu_x, \sigma_x, \sigma_\eta) = \int_0^\infty \mathcal{N}\left(n\mu_x, \sqrt{n\sigma_x^2 + n\sigma_\eta^2}\right). \tag{15}$$

Under linear integration, the RCF will be flat across time for a given total number of samples. The height of the RCF conditional on a given sample count can be found by integrating over possible values of evidence and computing the conditional probability of accuracy on a trial featuring a sample with that value (*cf.* Equation 11):

$$E(x_i' \mid C, n, \mu_x, \sigma_x, \sigma_\eta) = \frac{1}{P(C \mid n, \mu_x, \sigma_x, \sigma_\eta)} \int_{-\infty}^{\infty} x_i' \, P(x_i' \mid \mu_x, \sigma_x) P(C \mid x_i', n, \mu_x, \sigma_x, \sigma_\eta) dx_i'. \tag{16}$$

The first term in the right hand side of the equation is a normalization factor, which can be obtained from Equation 15. The second term is a weighted integral of x_i', where the weights are defined by both the likelihood of observing sample x_i' and the likelihood that a trial including that sample will lead to response C. The computation of this term depends on whether the trial has more than one sample. Single-sample predictions require marginalization only over noise:

$$P(\text{correct} \mid x_i', n = 1, \mu_x, \sigma_x, \sigma_\eta) = \int_0^\infty \mathcal{N}(x_i', \sigma_\eta). \tag{17}$$

With multiple samples, however, it is necessary to marginalize over the contributions of noise on sample i, y_i, and the joint influence of evidence and noise on other samples:

$$\begin{aligned} P(C \mid x_i', n > 1, \mu_x, \sigma_x, \sigma_\eta) &= \int_{-\infty}^{\infty} P(y_i \mid \sigma_\eta) P(C \mid y_i, \sigma_\eta, \mu_x, \sigma_x) dy_i \\ P(\text{correct} \mid x_i', n > 1, \mu_x, \sigma_x, \sigma_\eta) &= \int_{-\infty}^{\infty} P(y_i \mid \sigma_\eta) \left[\int_{-y_i - x_i'}^{\infty} \mathcal{N}\left[(n-1)\mu_x, \sqrt{(n-1)\left(\sigma_x^2 + \sigma_\eta^2\right)}\right] \right] dy_i. \end{aligned} \tag{18}$$

Extrema detection model

In the extrema detection model, the noisy perception of evidence afforded by each sample is compared against a threshold; if the threshold is exceeded, the process terminates in a commitment to the corresponding decision. Otherwise, the sample is discarded and the process continues. If the trial ends without a commitment, the process generates a random choice. The update equation that produces a decision variable is defined in Equation 2. The response is determined either by a commitment during the trial or by drawing a random response g_j from $G \sim \text{Bernoulli}(.5)$:

$$r_j = \begin{cases} \text{high} & \text{if } V_n = +1 \text{ or} \left(V_n = 0 \text{ and } g_j = 1\right) \\ \text{low} & \text{if } V_n = -1 \text{ or} \left(V_n = 0 \text{ and } g_j = 0\right). \end{cases} \tag{19}$$

Unlike with linear integration, predicting the probability of choosing each alternative under the assumption of extrema detection requires knowledge about the specific sequence of stimuli. Intuitively, the prediction of the "high" choice is based on the probability that the noisy perception of a given sample, $\hat{x}_i = x_i + \xi_\eta$, will exceed the positive threshold weighted by the probability that the process has not yet terminated before sample i (*cf.* Equation 8):

$$\begin{aligned} P(r_j = \text{high} \mid x_1 \ldots x_n, \sigma_\eta, \theta_x) &= \sum_{i=1}^{n} \left[P(+\theta_x < \hat{x}_i \mid \sigma_\eta) \prod_{k=1}^{i-1} P(-\theta_x < \hat{x}_k < +\theta_x \mid \sigma_\eta) \right] \\ &\quad + \frac{1}{2} \prod_{i=1}^{n} P(-\theta_x < \hat{x}_i < +\theta_x \mid \sigma_\eta) \\ &= \sum_{i=1}^{n} \left[\int_{+\theta_x}^{\infty} \mathcal{N}(x_i, \sigma_\eta) \prod_{k=1}^{i-1} \int_{-\theta_x}^{+\theta_x} \mathcal{N}(x_k, \sigma_\eta) \right] \\ &\quad + \frac{1}{2} \prod_{i=1}^{n} \int_{-\theta_x}^{+\theta_x} \mathcal{N}(x_i, \sigma_\eta). \end{aligned} \tag{20}$$

To predict the choice probability given a specific mean evidence value, $\bar{x}$, it is necessary to marginalize over all possible stimulus sequences that could have produced that mean. Therefore, the basic approach is to compute (*cf.* Equation 9):

$$P(r_j = \text{high} \mid \bar{x}, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x) = \frac{1}{Z} \int \dots \int_R P(x_1 \dots x_n \mid \mu_x, \sigma_x)\, I(x_1 \dots x_n, \bar{x})\, H(x_1 \dots x_n, \sigma_\eta, \theta_x)\, dx_1 \cdots dx_n. \tag{21}$$

Working backward, the function H gives the probability that the stimulus sequence $x_1 \dots x_n$ will lead to a "high" response and can be computed as in Equation 20. The function I is an indicator function that selects only those sequences with mean $\bar{x}$. Non-zero response probabilities are further weighted by the probability of observing that sequence of stimuli, which is computed using the parameters of the evidence-generating distribution. Finally, the integrated values are divided by a normalizing constant:

$$Z = \int \dots \int_R P(x_1 \dots x_n \mid \mu_x, \sigma_x)\, I(x_1 \dots x_n, \bar{x})\, dx_1 \cdots dx_n. \tag{22}$$

In practice, the discontinuity imposed by I requires us to approximate the integrals by sampling on a dense grid with bin width ω, so I becomes

$$I(x_1 \dots x_n, \bar{x}, \omega) = \begin{cases} 1 \text{ if } \left| \bar{x} - \frac{1}{n} \sum_{i=1}^{n} x_i \right| < \frac{\omega}{2} \\ 0 \text{ otherwise} \end{cases} \tag{23}$$

For predicting accuracy as a function of the number of samples, we take a similar approach as in Equation 20 but use the generating statistics rather than the specific stimulus sequences (*cf.* Equation 10):

$$P(\text{correct} \mid n, \mu_x, \sigma_x, \sigma_\eta, \theta_x) = \sum_{i=1}^{n} \int_{+\theta_x}^{\infty} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \left(\int_{-\theta_x}^{+\theta_x} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \right)^{i-1} + \frac{1}{2} \left(\int_{-\theta_x}^{+\theta_x} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \right)^{n}. \tag{24}$$

Predicting the reverse correlation function depends on the additional observation that the conditional distribution of evidence depends on whether or not the extrema detection process has terminated at that point. Let T_i represent whether the process is still considering evidence at the ith sample, with $T_i = 0$ representing an active sampling process and $T_i = 1$ representing a terminated process. Once the process has terminated, the samples do not bear on choice; their conditional probability given the choice is determined by the generating distribution for the stimulus, with mean μ_x. Further, because we are conditioning on accuracy, the active process must be sampling from a distribution that is truncated at $-\theta_x$. Therefore, the contributions of samples from these two distributions on sample i are weighted by the probability that the process has terminated prior to that sample (*cf.* Equation 11):

$$E\left(x_i' \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x\right) = P(T_i = 1 \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)\, \mu_x + P(T_i = 0 \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)\, E\left(x_i' \mid T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x\right). \tag{25}$$

The conditional distribution for active processes can be further broken down by whether the noisily perceived evidence from a given sample, $\hat{x}_i' = x_i' + \xi_\eta$, is above, below, or between the thresholds:

$$\begin{aligned} E\left(x_i' \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x\right) \quad & = P(T_i = 1 \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)\, \mu_x + P(T_i = 0 \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x) \\ & \Big[P\left(+\theta_x < \hat{x}_i' \,\middle|\, T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) E\left(x_i' \mid +\theta_x < \hat{x}_i', \mu_x, \sigma_x, \sigma_\eta\right) \\ & + P\left(-\theta_x < \hat{x}_i' < +\theta_x \mid T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) E\left(x_i' \mid -\theta_x < \hat{x}_i' < +\theta_x, \mu_x, \sigma_x, \sigma_\eta\right) \\ & + P\left(\hat{x}_i' < -\theta_x \mid T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) E\left(x_i' \mid \hat{x}_i' < -\theta_x, \mu_x, \sigma_x, \sigma_\eta\right) \Big]. \end{aligned} \tag{26}$$

The conditional expectations in Equation 26 require marginalizing over possible values of noise:

$$\begin{aligned} E\left(x_i' \mid +\theta_x < \hat{x}_i', \mu_x, \sigma_x, \sigma_\eta\right) \quad & = \int_{-\infty}^{\infty} x_i'\, P(x_i') \left[\int_{+\theta_x - x_i'}^{\infty} P(\xi_\eta)\, d\xi_\eta \right] dx_i' \\ E\left(x_i' \mid -\theta_x < \hat{x}_i' < +\theta_x, \mu_x, \sigma_x, \sigma_\eta\right) \quad & = \int_{-\infty}^{\infty} x_i'\, P(x_i') \left[\int_{-\theta_x - x_i'}^{+\theta_x - x_i'} P(\xi_\eta)\, d\xi_\eta \right] dx_i' \\ E\left(x_i' \mid \hat{x}_i' < -\theta_x, \mu_x, \sigma_x, \sigma_\eta\right) \quad & = \int_{-\infty}^{\infty} x_i'\, P(x_i') \left[\int_{-\infty}^{-\theta_x - x_i'} P(\xi_\eta)\, d\xi_\eta \right] dx_i'. \end{aligned} \tag{27}$$

The conditional probabilities of observing evidence at different positions with respect to the threshold in Equation 26 can be found using Bayes rule:

$$\begin{aligned}
P\left(+\theta_x < \hat{x}_i^{'} \,\middle|\, T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) &= \frac{P\left(C \mid +\theta_x < \hat{x}_i^{'}, T_i = 0\right) P\left(+\theta_x < \hat{x}_i^{'} \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right)}{P(C \mid T_i = 0, n, \mu_x, \sigma_x, \sigma_\eta)} \\
P\left(-\theta_x < \hat{x}_i^{'} < +\theta_x \,\middle|\, T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) &= \frac{P\left(C \mid -\theta_x < \hat{x}_i^{'} < +\theta_x, T_i = 0\right) P\left(-\theta_x < \hat{x}_i^{'} < +\theta_x \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right)}{P(C \mid T_i = 0, n, \mu_x, \sigma_x, \sigma_\eta)} \\
P\left(\hat{x}_i^{'} < -\theta_x \,\middle|\, T_i = 0, C, n, \mu_x, \sigma_x, \sigma_\eta\right) &= \frac{P\left(C \mid \hat{x}_i^{'} < -\theta_x, T_i = 0\right) P\left(\hat{x}_i^{'} < -\theta_x \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right)}{P(C \mid T_i = 0, n, \mu_x, \sigma_x, \sigma_\eta)}.
\end{aligned} \tag{28}$$

The conditional probability of accuracy given an observation between the thresholds in Equation 28 can be found using Equation 24 but substituting $n - i + 1$ for n. The conditional probabilities given observations above or below the threshold in Equation 28 are given by the definition of the extrema detection process:

$$\begin{aligned}
P\left(\text{correct} \mid +\theta < \hat{x}_i^{'}, T_i = 0\right) &= 1 \\
P\left(\text{wrong} \mid +\theta < \hat{x}_i^{'}, T_i = 0\right) &= 0 \\
P\left(\text{correct} \mid \hat{x}_i^{'} < -\theta, T_i = 0\right) &= 0 \\
P\left(\text{wrong} \mid \hat{x}_i^{'} < -\theta, T_i = 0\right) &= 1.
\end{aligned} \tag{29}$$

The second term in each numerator of Equation 28 is the marginal probability of observing evidence in each segment, which depends on the parameters of the distribution that generates evidence and also on the noise:

$$\begin{aligned}
P\left(+\theta_x < \hat{x}_i^{'} \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right) &= \int_{+\theta_x}^{\infty} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \\
P\left(-\theta_x < \hat{x}_i^{'} < +\theta_x \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right) &= \int_{-\theta_x}^{+\theta_x} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \\
P\left(\hat{x}_i^{'} < -\theta_x \,\middle|\, \mu_x, \sigma_x, \sigma_\eta\right) &= \int_{-\infty}^{-\theta_x} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right).
\end{aligned} \tag{30}$$

Each denominator in Equation 28 is the probability of response accuracy conditional on the process not having terminated at sample i, which again can be found using Equation 24 with $n - i + 1$ substituted for n.

So far, we have derived all the terms necessary for the calculation of the RCF (Equation 26) except for the probability of having or not having terminated before sample i. A process that has not terminated prior to sample i has observed only intermediate values of noisy evidence up to that point, and the probability of this occurring can be computed using Equation 28:

$$\begin{aligned}
P(T_i = 0 \mid C, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x) &= \frac{P(C \mid T_i = 0, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)}{P(C \mid n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)} P\left(-\theta_x < \hat{x}_i^{'} < +\theta_x \mid C, \mu_x, \sigma_x, \sigma_\eta\right)^{i-1} \\
P(T_i = 1 \mid \mathit{correct}, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x) &= \sum_{k=1}^{i} \int_{+\theta_x}^{\infty} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \left(\int_{-\theta_x}^{+\theta_x} \mathcal{N}\left(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2}\right) \right)^{k-1},
\end{aligned} \tag{31}$$

where $P(C \mid n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)$ is derived from Equation 24. $P(C \mid T_i = 0, n, \mu_x, \sigma_x, \sigma_\eta, \theta_x)$ is derived from the same equation but by replacing n with $n - i + 1$.

Counting model

In the counting model, each sample undergoes a transformation to a positive or negative increment, and a discrete count is maintained across samples. The process is defined in Equation 3. The response is determined by the sign of the decision variable at the end of the trial. Because V is maintained as a discrete count, it is possible that trials with an even number of samples will end in a tie ($V_n = 0$). In this case, as with the extrema detection model, the response is generated using a random guess g_j from $G \sim \text{Bernoulli}(.5)$:

$$r_j = \begin{cases} \text{high} & \text{if } V_n > 0 \text{ or} \left(V_n = 0 \text{ and } g_j = 1\right) \\ \text{low} & \text{if } V_n < 0 \text{ or} \left(V_n = 0 \text{ and } g_j = 0\right). \end{cases} \tag{32}$$

Therefore, the probability of a "high" response equals the sum of two probabilities: the probability of positive increments for more than half of samples or a positive increment in exactly half and a positive random guess. If we define V_n^+ as representing the number of positive increments from n samples, the choice probability becomes (*cf.* Equation 8)

$$P(r_j = \text{high} \mid x_1 \ldots x_n, \sigma_\eta) = P(V_n > 0 \mid x_1 \ldots x_n, \sigma_\eta) + \frac{1}{2} P(V_n = 0 \mid x_1 \ldots x_n, \sigma_\eta)$$
$$= \sum_{k=\left\lceil \frac{n+1}{2} \right\rceil}^{n} P(V_n^* = k \mid x_1 \ldots x_n, \sigma_\eta) + \frac{1}{2} P\left(V_n^* = \frac{n}{2} \middle| x_1 \ldots x_n, \sigma_\eta\right), \tag{33}$$

where $\lceil . \rceil$ is the ceiling function. There is a unique probability of each sample producing a positive increment in the accumulator. For sample x_i the probability of an increment is

$$P(c_i = 1 \mid x_i, \sigma_\eta) = \int_0^\infty \mathcal{N}(x_i, \sigma_\eta). \tag{34}$$

Further, observe that

$$P(V_n^* = k) = P(V_{n-1}^* = k) P(c_i = 0 \mid x_i, \sigma_\eta) + P(V_{n-1}^* = k - 1) P(c_i = 1 \mid x_i, \sigma_\eta). \tag{35}$$

We can recursively apply this equation to compute the probabilities in Equation 33.

Predicting choice probability from a specific mean evidence value (the mPMF) follows a similar approach to the one taken for the extrema detection model using Equation 21. The only difference is the $H(.)$ function; for the counting model, we can use Equation 33 to compute the choice probability for a given sequence of stimuli.

To generate a prediction for the cPMF under the counting model, we also take a similar approach as in Equation 33. But it is not necessary to consider each idiosyncratic sequence. Instead, the probability that the count of increments supporting the correct response takes any particular value can be calculated based on a Binomial random variable with success probability $p = \int_0^\infty \mathcal{N}(\mu_x, \sqrt{\sigma_x^2 + \sigma_\eta^2})$. The count psychometric function can therefore be expressed as (*cf.* Equation 10)

$$P(\text{correct} \mid n, p) = \sum_{k=\left\lceil \frac{(n+1)}{2} \right\rceil}^{n} \binom{n}{k} p^k (1-p)^{n-k} + \frac{1}{2} \begin{cases} \binom{n}{\frac{n}{2}} p^{\frac{n}{2}} (1-p)^{\frac{n}{2}} & n \text{ even} \\ 0 & n \text{ odd} \end{cases} \tag{36}$$

where $\binom{n}{k}$ is the binomial coefficient function (*n* choose *k*). The first term on the right hand side of the equation gives the probability of success due to a count that favors the right choice, while the second term gives the probability of success due to guessing correctly after a tie.

The counting model makes a qualitative prediction that accuracy on trials with an even number of samples will not exceed accuracy on trials with the next smallest odd number. This result follows from Equation 36, but the logic behind it can also be seen by considering the possible sequence of events on each trial. If the probability that a sample will support the correct choice is p, then the probability of responding correctly on a trial with two samples can be written as

$$\begin{aligned} P(\text{correct} \mid n = 2) &= p^2 + \frac{1}{2} p(1-p) + \frac{1}{2} (1-p) p \\ &= p \\ &= P(\text{correct} \mid n = 1). \end{aligned} \tag{37}$$

This shows the equivalence between accuracy on one-sample and two-sample trials. The logic can be extended to pairs of trial types with larger numbers of samples.

Similar to the case with linear integration, the RCF for the counting model will be flat for a given total number of samples. It can be computed in general using Equation 16 with a change in the conditional probabilities and normalization constant. The normalization constant can be found using Equation 36. As with linear integration, the conditional probability of response accuracy given evidence on sample *i* is different for single and multiple sample trials. In fact, the two models make identical predictions for single sample trials (Equation 17). For multi-sample trials, the conditional probability depends on whether sample *i* will produce a positive or negative increment in the count:

$$\begin{aligned} P(C \mid x_i', n, \mu_x, \sigma_x, \sigma_\eta) &= P\left(C \mid \hat{x}_i' > 0, n, \mu_x, \sigma_x, \sigma_\eta\right) P\left(\hat{x}_i' > 0 \mid x_i', \sigma_\eta\right) \\ &+ P\left(C \mid \hat{x}_i' < 0, n, \mu_x, \sigma_x, \sigma_\eta\right) P\left(\hat{x}_i' < 0 \mid x_i', \sigma_\eta\right) \\ &= P\left(C \mid \hat{x}_i' > 0, n, \mu_x, \sigma_x, \sigma_\eta\right) \int_0^\infty \mathcal{N}(x_i', \sigma_\eta) \\ &+ P\left(C \mid \hat{x}_i' < 0, n, \mu_x, \sigma_x, \sigma_\eta\right) \int_{-\infty}^0 \mathcal{N}(x_i', \sigma_\eta) \, . \end{aligned} \tag{38}$$

The conditional probabilities use the Binomial distribution following a similar logic to what we used when computing the cPMF (Equation 36). If $\hat{x}_i > 0$, getting a correct response requires the remaining $n-1$ samples to produce at least $n/2$ positive increments or $n/2-1$ positive increments and a lucky guess. If $\hat{x}_i < 0$, the remaining $n-1$ samples must produce at least $(n+1)/2$ positive increments or $n/2$ positive increments and a lucky guess.

Leaky integration model

The leaky integration model extends the linear integration model to account for two additional factors that could influence the decision variable as a function of time. The process is defined in Equation 4. As in linear integration, the response is determined by the sign of the decision variable at the end of the trial.

To fit the model, we use the fact that as $\Delta t \to 0$, the change in the decision variable during the gaps can be described with an Ornstein-Uhlenbeck (OU) process. If sample i at time t_i is followed by gap duration $\mathfrak{T}_g$, the OU process transforms the Gaussian distribution of the decision variable after sample i, $\mathcal{N}(\mu(t_i), \sigma(t_i))$, to another Gaussian distribution with the following mean and variance at the end of the gap:

$$\begin{aligned} \mu\left(t_i + \mathfrak{T}_g\right) &= e^{-\lambda \mathfrak{T}_g} \mu(t_i) \\ \sigma^2\left(t_i + \mathfrak{T}_g\right) &= e^{-2\lambda \mathfrak{T}_g} \sigma^2(t_i) + \left(1 - e^{-2\lambda \mathfrak{T}_g}\right) \frac{\sigma_\varepsilon^2}{2\lambda}. \end{aligned} \tag{39}$$

At the end of the trial, time T, the choice likelihood is given by integrating the positive density of the resulting distribution following all samples and gaps (*cf.* Equation 8):

$$P(r_j = \text{high} \mid x_i \ldots x_n; \sigma_\eta, \sigma_\varepsilon, \lambda) = \int_0^\infty N(\mu(T), \sigma(T)). \tag{40}$$

Analytic predictions for the behavioral assays could be obtained by extending the equations of the linear integration model for an OU process. However, the resulting equations would be complex because they would need to marginalize over all gap durations in the experiment. We avoid this complexity by using Monte Carlo simulation of model performance for the trials in the behavioral datasets (5 simulations per trial; Δt = 100 ms) and comparing the simulated model behavior to that of our subjects.

This modeling framework also allowed us to evaluate whether time-independent leak could explain the data, which we accomplished by setting the gap between each sample to a uniform duration and re-fitting the model. This produced results that were entirely consistent with those using time-dependent leak that are reported in the main results.

Finally, we confirmed that our conclusions about the magnitude of memory leak would be robust to transformations of the evidence that introduce a compressive nonlinearity, as Figure 3B suggests. Model simulations showed that the leaky integration model, as defined here, would only over-estimate the leak rate when fit to data generated by a process that includes both significant leak and a compressive nonlinearity. This is expected because compression and leak both disproportionately reduce the influence of strong samples. Therefore, to the extent that compression exists in the data, it is only causing us to under-estimate the integration time constant that humans can achieve.

DATA AND SOFTWARE AVAILABILITY

Experimental data and code will be made available at https://github.com/kianilab/Waskom_CurrBiol_2018.

Current Biology, Volume *28*

Supplemental Information

Decision Making through Integration of Sensory Evidence at Prolonged Timescales

Michael L. Waskom and Roozbeh Kiani

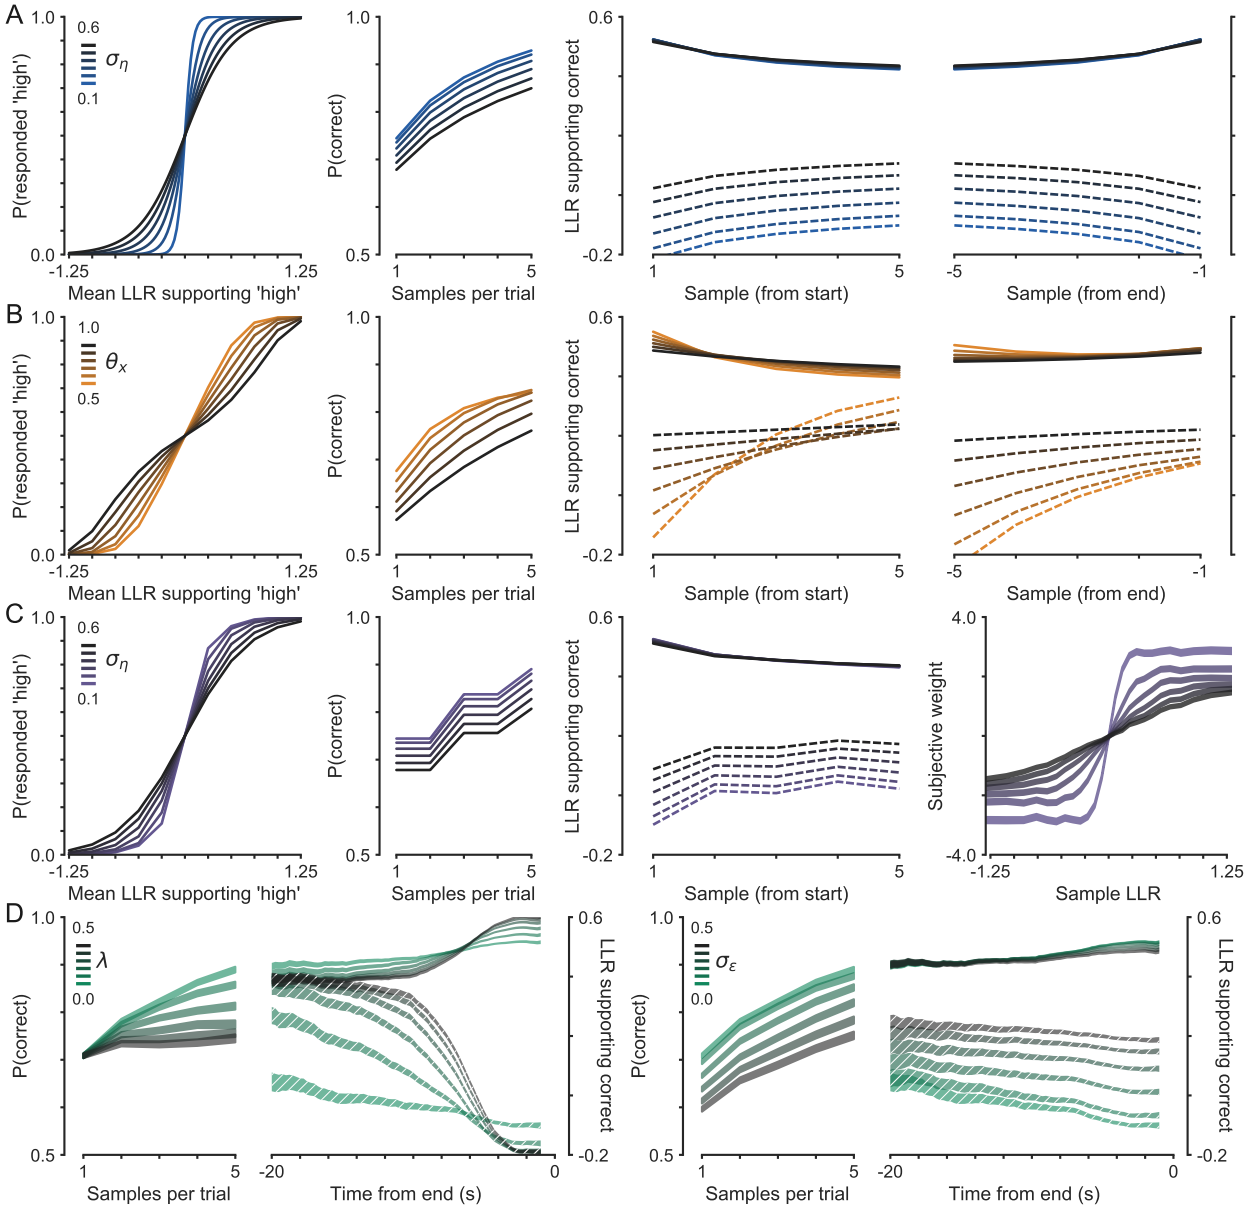

Figure S1: Computational model predictions for different behavioral assays Related to Figures 2–4. (A) Predictions of the Linear Integration model when varying the magnitude of noise, σ_η. Unlike in Figures 2 and S2, RCF predictions are shown averaged over ordinal sample positions, aligned to either the start or end of the trial. Note how the variable distribution of sample counts produces an apparent nonlinearity for Linear Integration even though samples are weighted equally across time. (B) Predictions of the Extrema Detection model when varying the height of the detection threshold, θ_x, while holding noise constant (σ_η = 0.2). (C) Predictions of the Counting model when varying the magnitude of noise, σ_η. (D) Simulated performance for the Leaky Integration model (N = 100,000 trials). The left panels show simulations varying the memory leak rate, λ, while holding the noise terms constant (σ_η = 0.4 and σ_ϵ = 0). The right panels show simulations varying the memory noise magnitude, σ_ϵ while holding stimulus-dependent noise and memory leak constant (σ_η = 0.4 and λ = 0). The widths of the bands show 95% confidence intervals.

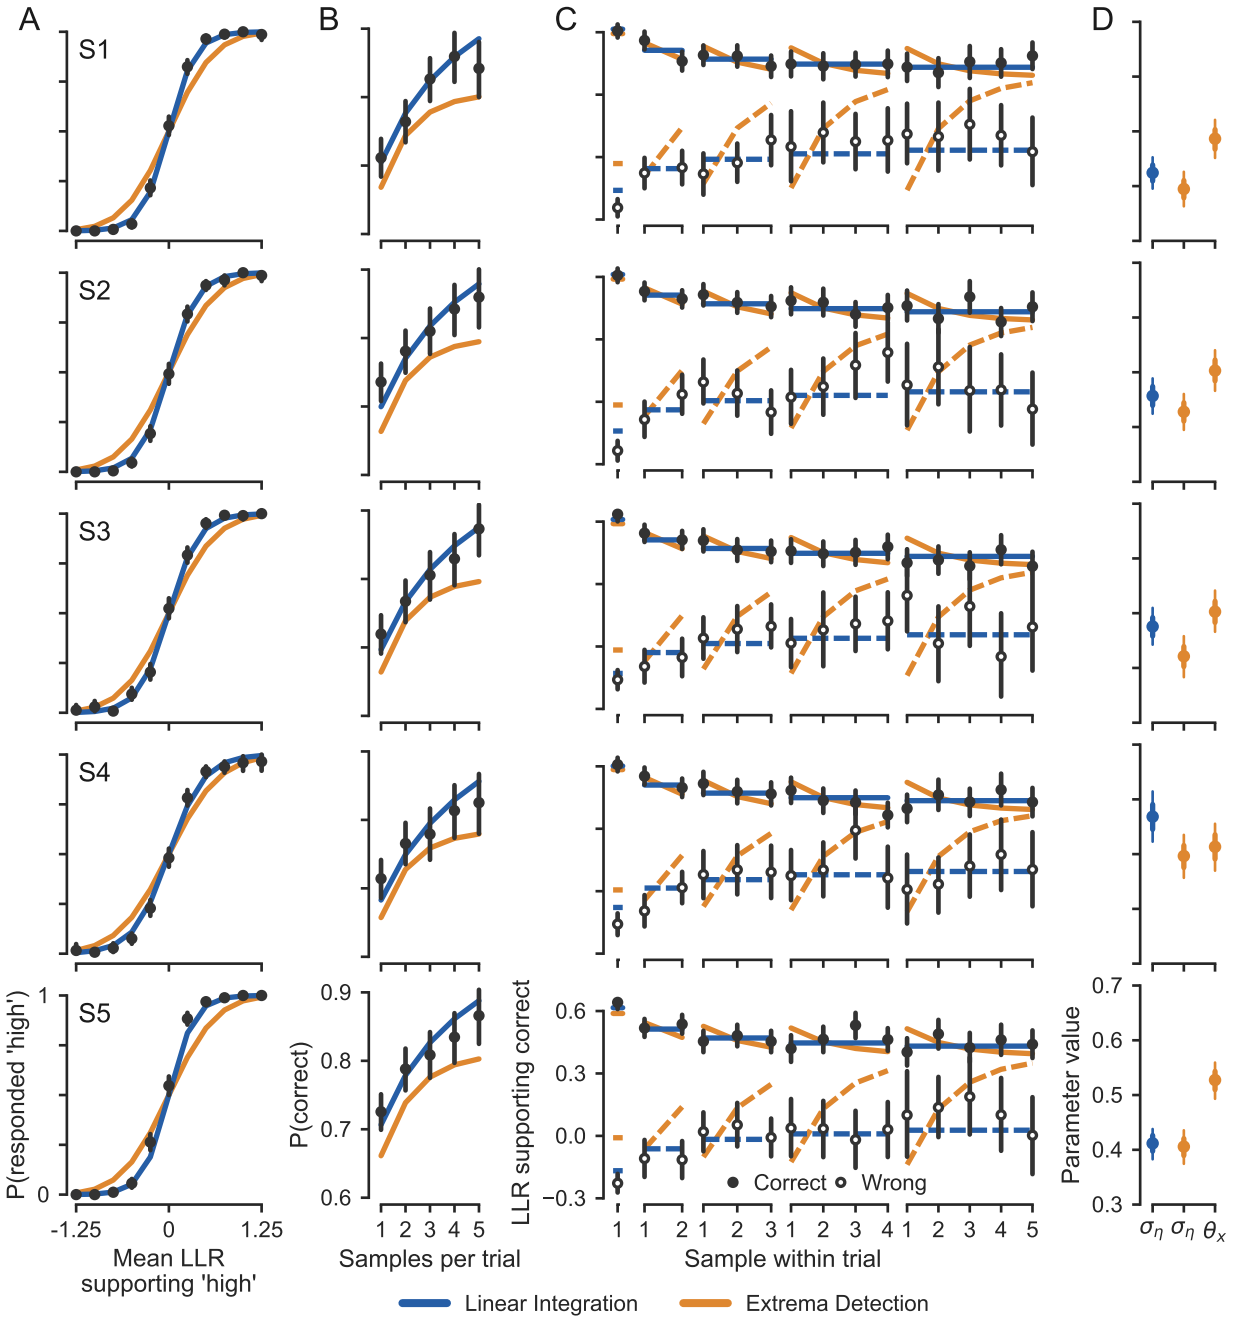

Figure S2: Integration of evidence across samples. Related to Figure 2. Each row shows data and model fits for individual subjects. (A-C) Individual data and model fits for the three main behavioral assays, as in Figure 2A-C. (D) Estimated model parameters (points) and bootstrap confidence intervals (thick and thin error bars show 68% and 95% CIs, respectively).

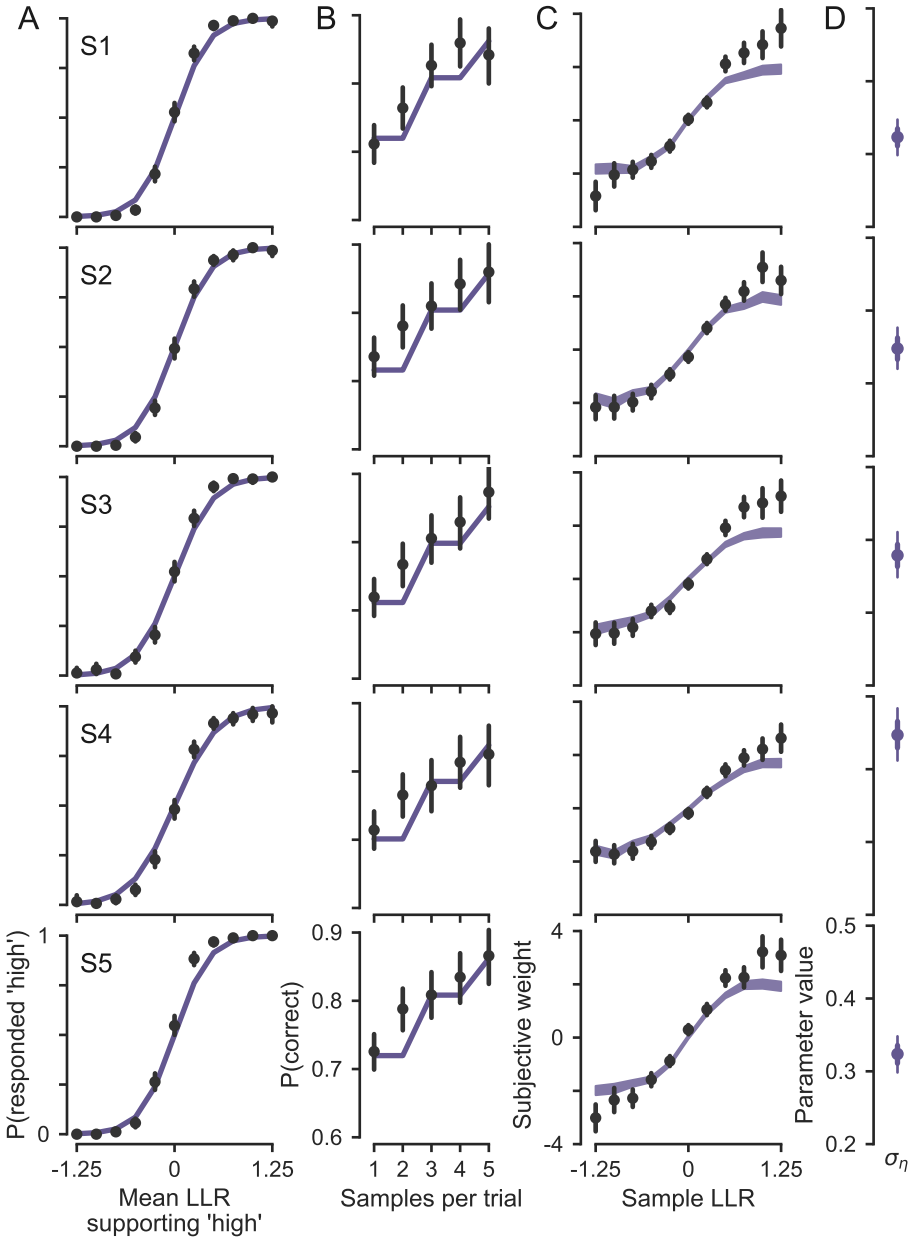

Figure S3: Integration of graded stimulus evidence. Related to Figure 3. Each row shows data and model fits for individual subjects. (A) Individual data and model mPMFs. (B-C) Individual data and model fits corresponding to Figure 3. (D) Estimated model parameters (points) and bootstrap confidence intervals (thick and thin error bars show 68% and 95% CIs, respectively).

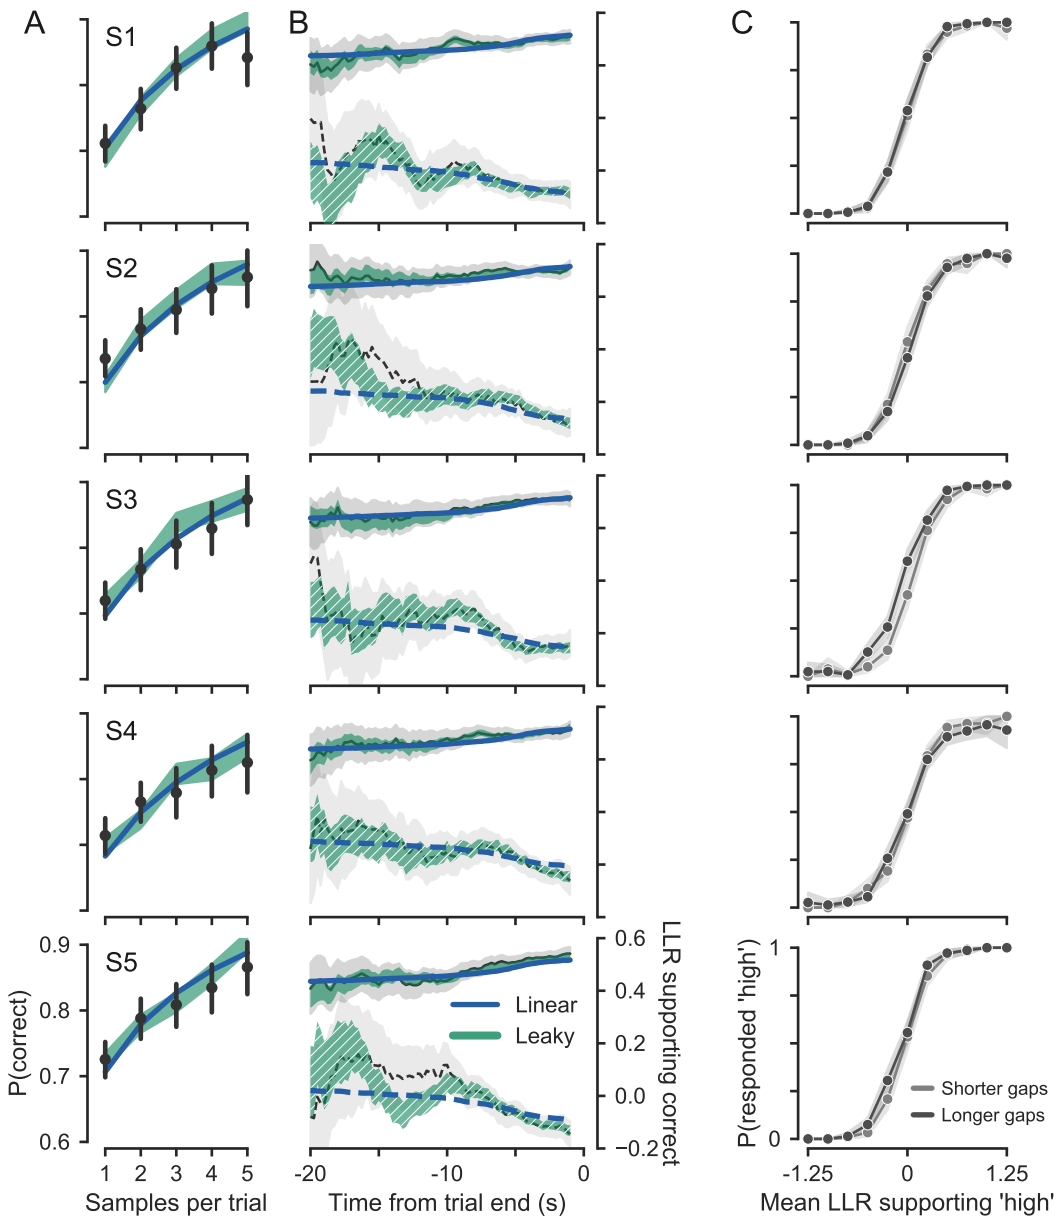

Figure S4: Minimal influence of memory leak or noise. Related to Figure 4. Each row shows data and model fits for individual subjects. (A-C) Individual data and model fits corresponding to Figure 4.