

nlm.nih.gov/sites/sra) archives as E-MEXP-3798 and SRP017718, respectively. We thank A. Vodala, J. Menet, K. Abruzzi, and N. Francis for advice with techniques. Supported by a Wellcome Trust Senior Research Fellowship in Basic Biomedical Sciences, the Gatsby Charitable Foundation, Oxford Martin School and grants MH069883 and MH081982 from NIH (S.W.), NIH grants

NS044232 and NS045713 and the Ellison Medical Foundation (M.R.), and NIH grant R01HD049116 (Z.W. and W.T.).

Supplementary Materials

www.sciencemag.org/cgi/content/full/340/6128/91/DC1
Materials and Methods

Figs. S1 to S6
Tables S1 to S10
References (45–49)

24 October 2012; accepted 1 February 2013
10.1126/science.1231965

Rats and Humans Can Optimally Accumulate Evidence for Decision-Making

Bingni W. Brunton,^{1,2*} Matthew M. Botvinick,^{1,3} Carlos D. Brody^{1,2,4†}

The gradual and noisy accumulation of evidence is a fundamental component of decision-making, with noise playing a key role as the source of variability and errors. However, the origins of this noise have never been determined. We developed decision-making tasks in which sensory evidence is delivered in randomly timed pulses, and analyzed the resulting data with models that use the richly detailed information of each trial's pulse timing to distinguish between different decision-making mechanisms. This analysis allowed measurement of the magnitude of noise in the accumulator's memory, separately from noise associated with incoming sensory evidence. In our tasks, the accumulator's memory was noiseless, for both rats and humans. In contrast, the addition of new sensory evidence was the primary source of variability. We suggest our task and modeling approach as a powerful method for revealing internal properties of decision-making processes.

Decisions in real life often need to be made based on noisy or unreliable evidence. Accumulating evidence from a set of noisy observations made over time makes it possible to average over different noise samples, thus improving estimates of the underlying signal. This principle is the basis for the influential class

of “drift-diffusion” models (1–5), which have been broadly applied to explain a variety of phenomena in biology (6–8). Accumulation involves both maintaining a memory of evidence accrued so far and adding new evidence to the memory. Yet no test to date has distinguished between noise associated with each of these two components.

We developed tasks in which subjects (humans and rats) were concurrently presented with two trains of pulses, one train representing “left”-labeled pulses and the other, “right”-labeled pulses. At the end of each trial, the subjects had to report which of the two trains had the greater total number of pulses. The timing of pulses was random and varied widely, both within and across individual trials (9, 10). We reasoned that the precisely known pulse timing would enable detailed modeling of the subjects' choices on each individual trial, whereas its variability would allow exploration of the stimulus space and would thus provide statistical power.

In an auditory version of the task, performed by three humans and 19 rats, left pulse trains were clicks presented on a speaker to the left of the subject, and right pulse trains were clicks presented on a speaker to the right of the sub-

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

*Present address: Department of Biology and Department of Applied Mathematics, University of Washington, Seattle, WA 98195, USA.

†Corresponding author. E-mail: brody@princeton.edu

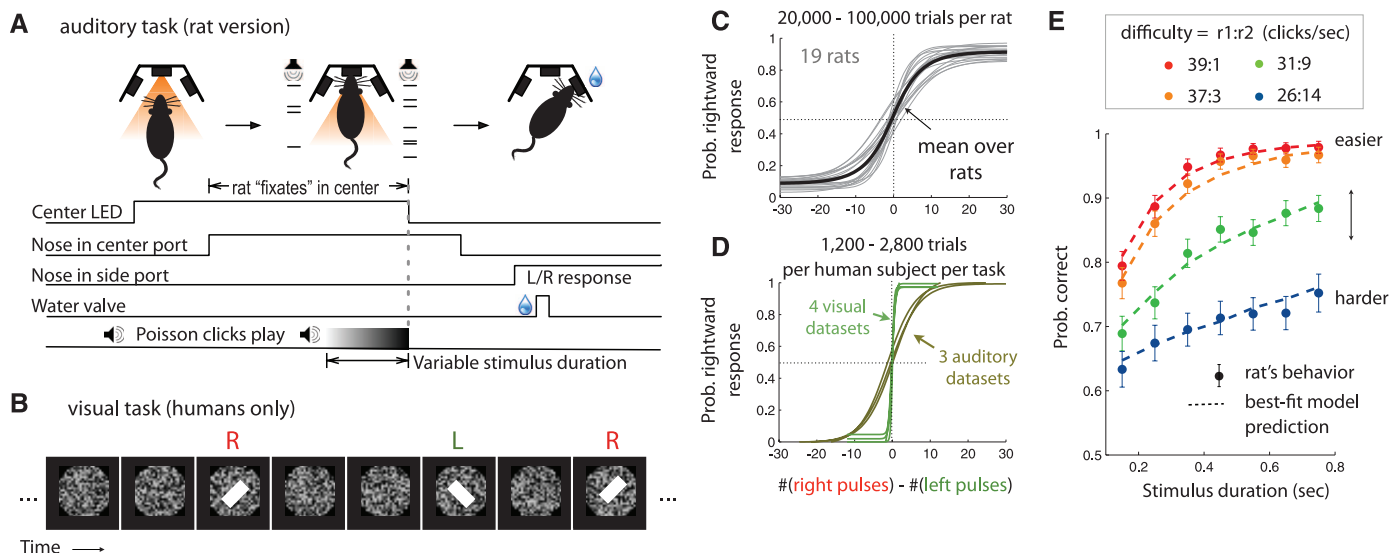


Fig. 1. Psychophysical tasks and summary of behavior. (A) Sequence of events in each trial of the rat auditory task. After light onset from a light-emitting diode (LED) in a center port, trained rats placed their nose into the port and “fixated” their nose there for a fixed amount of time until the light was turned off (1 to 2 s). Trains of randomly timed clicks were played concurrently from left and right free-field speakers during the last portion of the fixation time. After nose fixation and sounds ended, the rat made a choice, poking in the left or the right port to indicate which side played more clicks. Humans performed an analogous version of the task on a computer while

wearing headphones. (B) Schematic diagram of a stimulus in the visual pulses version of the task, performed by humans on a computer. (C) Psychometric curves (fits to a four-parameter logistic function for each subject; see methods) for rat subjects. (D) Psychometric curves, as in (C), for human subjects. (E) Chronometric curves for an example rat. Difficulty is labeled by the ratio of click rates played on the two sides. For each difficulty, performance improves with longer stimulus durations. Dashed lines show the best-fit model predictions for this rat, as described in the text. The vertical axis shows mean accuracy and 95% confidence interval (CI).

ject (Fig. 1A; free-field speakers for rats, headphones for humans). In a visual version of the task, performed by four humans, left pulses were flashed white bars, tilted anticlockwise from the vertical, and right pulses were flashed white bars tilted clockwise (3) (Fig. 1B). On each trial, the stimulus was presented for a duration controlled by the experimenter. The sum of the two pulse rates was kept fixed within each task, and discrimination difficulty was controlled on each trial by the ratio of the two rates (Fig. 1, C to E).

To examine a large variety of possible mechanisms consistent with the performance improvement at longer stimulus durations seen in Fig. 1E [(11, 12); see also fig. S4], we expanded on the drift-diffusion framework and implemented a flexible model (Box 1) in which different regimes of model parameter values represent widely different mechanisms (three examples in Fig. 2A), with mixtures of mechanisms represented by intermediate parameter values.

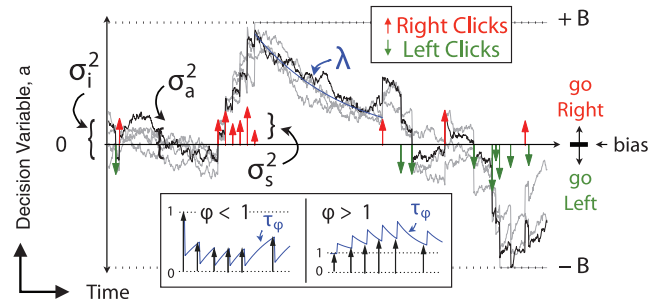
Given a trial's specific pulse times and a set of parameter values, the model produces the probability of observing a left versus a right response on that trial. Methods to compute the gradient of this probability with respect to model parameters (see the supplementary materials) were critical for efficiently finding the parameter values that gave the maximum likelihood of observing the complete set of a subject's responses. Numerical tests always found only one maximum (fig. S6), suggesting that we always found the global maximum. Consistent with this observation, a mathematically related model has been proven to have a concave log likelihood (6), suggesting that our model may also be provably concave and have a single maximum.

Figure 2, B to D, shows the likelihood landscape around best-fit (i.e., maximum-likelihood) parameters, given the data of a representative rat subject. Confidence intervals are given by the parameter width of the maximum (blue contours). Figure 2B shows λ ($= 1/\tau$, the memory time constant), which represents accumulator memory leak (if $\lambda < 0$) or instability (if $\lambda > 0$), and B , the height of the decision-commitment evidence bounds. λ was statistically indistinguishable from zero. That is, the decision dynamics were neither leaky nor unstable, suggesting that sensory evidence from throughout the entire stimulus period was given equal weight. The best-fit B was large enough that it produced model fits indistinguishable from those produced by $B = \infty$. Across subjects (Fig. 2, E and H), species, sensory modalities, and task parameters, the accumulator's memory time constant $|\tau| = 1/|\lambda|$ was long ($|\lambda| = 0.91 \pm 0.15 \text{ s}^{-1}$ mean \pm SE across rats; $|\lambda| = 0.23 \pm 0.071 \text{ s}^{-1}$ across humans), in the sense that $|\tau|$ was comparable to or greater than the longest stimulus duration used [1 s for rats, 4 s for humans (13)]. The best-fit values of B and λ were thus in the gradual accumulation regime (Fig. 2A, top).

In our tasks, noise in the sensory evidence (σ_s^2) adds total variance proportional to the sum

of the amplitude of the clicks, whereas the memory diffusion noise (σ_a^2) adds total variance proportional to the stimulus duration. This separability allows us to isolate the magnitude of the diffusion noise σ_a^2 that gives the drift-diffusion model its name (2, 5). To our surprise, in 13 out of 19 rats and in all three humans performing the auditory task and all four humans performing the visual task, the value that best fit the data was the ideal $\sigma_a^2 = 0$ (Fig. 2C for an example subject, Fig. 2, F and I, for all subjects). Consistent with the easily

distinguishable right versus left pulses used with humans, the best-fitting values of sensory evidence noise σ_s^2 for humans were substantially lower than those for the rats (Fig. 2I). Again in this much lower σ_s^2 regime, the best-fitting memory diffusion noise σ_a^2 was zero. The dominant source of variability was thus noise in the evidence associated with each incoming pulse ($\sigma_s^2 = 1.90 \pm 0.28 \text{ pulses}^2$ per incoming pulse for rats, $0.50 \pm 0.059 \text{ pulses}^2$ in the human auditory task, and $0.24 \pm 0.10 \text{ pulses}^2$ in the human visual task.)



Box 1.

At each time point, the accumulator memory a (black trace) represents an estimate of the right versus left evidence accrued so far. At stimulus end, the model decides right if $a > \text{bias}$ and left otherwise, where bias is a free parameter. The light gray traces indicate alternate runs with different instantiations of model noise.

Right \uparrow (left \downarrow) pulses change the value of a by positive (negative) impulses of magnitude C .

σ_i^2 parameterizes noise in the the initial value of a .

σ_a^2 is a diffusion constant, parameterizing noise in a .

σ_s^2 parameterizes noise when adding the evidence from a right or left pulse: For each click, variance σ_s^2 is scaled by the amplitude of C and then added to the evidence contributed by the click.

λ parameterizes consistent drift in the memory a . In the “leaky” or forgetful case ($\lambda < 0$, illustrated), drift is toward $a = 0$, and later pulses affect the decision more than earlier pulses. In the “unstable” or impulsive case ($\lambda > 0$), drift is away from $a = 0$, and earlier pulses affect the decision more than later pulses. The memory's time constant $\tau = 1/\lambda$.

B is the height of the sticky decision bounds and parameterizes the amount of evidence necessary to commit to a decision.

ϕ and τ_ϕ parameterize sensory adaptation by defining the dynamics of C . Immediately after a click, the magnitude C is multiplied by ϕ . C then recovers toward an unadapted value of 1 with time constant τ_ϕ . Facilitation is thus represented by $\phi > 1$, whereas depression is represented by $\phi < 1$ (inset).

These properties are implemented by the following equations

if $|a| \geq B$ then $da/dt = 0$; else

$$da = \sigma_a dW + (\delta_{t,t_r} \cdot \eta_R \cdot C - \delta_{t,t_l} \cdot \eta_L \cdot C) dt + \lambda a dt \tag{1}$$

where $\delta_{t,t_{r,l}}$ are delta functions at the times of the pulses; η are i.i.d. Gaussian variables drawn from $N(1, \sigma_c)$; and dW is a white-noise Wiener process. The initial condition $a(t = 0)$ is drawn from the Gaussian $N(0, \sigma_i)$.

Adaptation dynamics are given by

$$\frac{dC}{dt} = \frac{1-C}{\tau_\phi} + (\phi - 1)C(\delta_{t,t_r} + \delta_{t,t_l}) \tag{2}$$

In addition, a lapse rate parameterizes the fraction of trials on which a random response is made. Ideal performance ($a = \text{\#right clicks} - \text{\#left clicks}$) would be achieved by

$$\lambda = 0, B = \infty, \sigma_a^2 = \sigma_s^2 = \sigma_i^2 = 0, \phi = 1, \text{bias} = 0 \tag{3}$$

This variability could be introduced by sensory uncertainty in the left-versus-right classification of each individual pulse, or by noise in the process of adding new sensory evidence to the accumulator memory.

The pulsatile nature of our task made it straightforward to parameterize sensory adaptation [(14); Eq. 2]. We found strong, quickly-recovering depression for rats (Fig. 2, D, G, and J; adaptation magnitude $\varphi = 0.17 \pm 0.021$, recovery time constant $\tau_\varphi = 0.080 \pm 0.024$ s, mean \pm SE across rats). This is consistent with the depression observed in auditory cortex neural responses to click train stimuli (15, 16). Stimuli in the human tasks were constructed with a minimum inter-pulse interval (30 ms in the auditory task, 150 ms in the visual task), and this greatly reduced the adaptation effects as compared to those in rats (Fig. 2, G and J). Across all adaptation regimes, we found long accumulator time constants and zero memory diffusion noise.

If our subjects' behavior depends on a process that cannot be approximated by the model [such as collapsing bounds (17), variability in attention (18), or other possibilities not yet formalized in a model], our interpretation of the best-fit values may be problematic. We therefore tested the model-derived conclusions. To assess the memory time constant τ , we calculated the "psychophysical reverse correlation" (17, 19–21), which estimates the extent to which click rates at each point in time influence left and right decisions. This analysis indicated that all periods of the trial have similar influence on the decision (approximately constant separation between the two traces in Fig. 3A), which is consistent with the long τ found in the model-based analysis. To assess our estimates of the single-pulse noise σ_s^2 and the starting variability σ_a^2 , we fit the model to data from trials with multiple clicks on each of the two sides, and used those best-fit parameters to predict performance on trials in

which only one single-side pulse happened to be presented (for which performance is dominated by σ_s^2 and σ_a^2). The prediction was accurate, even on an individual subject-by-subject basis (Fig. 3B). To assess the memory diffusion noise σ_a^2 , we controlled for sensory evidence by dividing trials into groups, so that all trials within a group had the same number of right clicks and the same number of left clicks; assuming large $|\tau|$, performance within each group is then dominated by σ_a^2 and the click depression parameters φ and τ_φ . Large σ_a^2 would predict decreasing within-group performance at longer stimulus durations. The data showed the opposite trend and was precisely predicted by the best-fit model, where $\sigma_a^2 = 0$ and clicks are depressing (Fig. 3C and fig. S13). The tests of Fig. 3 thus provided model-independent confirmation of the model-fit parameter values.

Using highly variable yet precisely known stimuli, together with a trial-by-trial model that

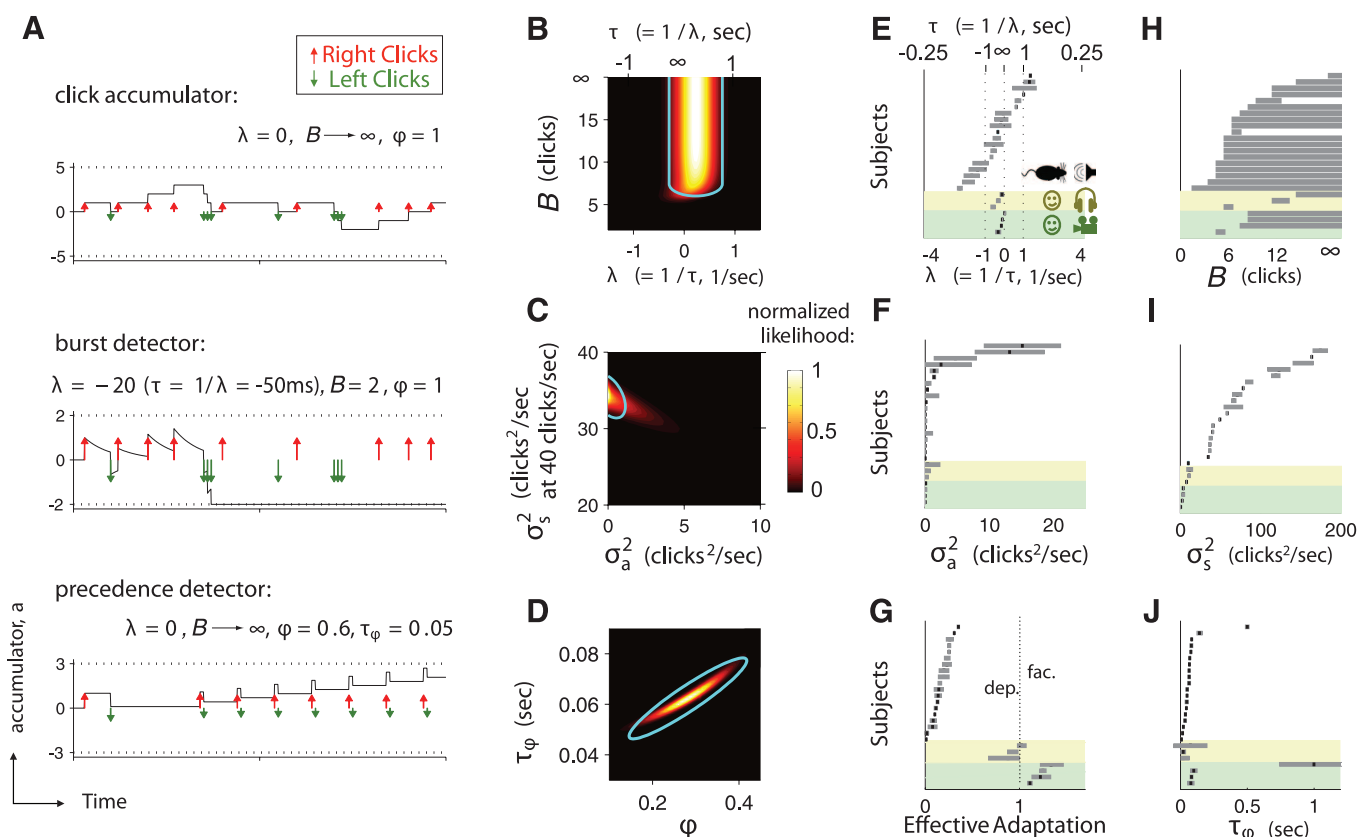


Fig. 2. The model can fit a variety of mechanisms, but the data consistently fit to a pulse-accumulating mechanism with zero noise in the accumulator's memory. (A) Three examples of the mechanisms that the model can represent. Top: The ideal, a pulse accumulator that weights all pulses equally. Middle: A burst detector. If three or more pulses from the same side arrive within 50 ms, the sticky bounds are reached, meaning that a commitment to orient to that side is made. Bottom: A precedence detector. If pulses from one side tend to arrive shortly before pulses from the other side, the adaptation minimizes the second side's pulses, and the decision tends toward the preceding side. (B to D) Parameter-likelihood landscapes that result from the data of one example rat. Panels are two-dimensional slices, cut through the full nine-dimensional parameter space, around the best-fit values. The blue curves represent CIs [2 SD of the multidimensional normal distribution fit to the

likelihood landscape (28)]. The best-fit parameter values found, which are $\lambda \approx 0$, $B \gg 1$, $\sigma_a^2 \approx 0$, $\sigma_s^2 \gg 0$, and $\varphi \ll 1$, correspond to pulse accumulation [(A), top] with a perfect memory but imperfect processing of sensory inputs. (E to J) Summaries of best-fit parameters over all subjects and tasks. Black ticks are best-fit values; gray bars span the CIs. Each panel has been divided by task (yellow highlight for human auditory task, green highlight for human visual task) and then sorted independently in order of parameter value. (E) All subjects, in all tasks, had long accumulator memories. (H) Most subjects were best fit with large bounds ($B \rightarrow \infty$). (F) Thirteen of 19 rats and all humans in both auditory and visual tasks were best fit with $\sigma_a^2 = 0$. (I) A wide range of values, all large compared to σ_a^2 , were found for σ_s^2 . (G) and (J) All rats showed strong, rapidly recovering depression ($\varphi < 1$, mean $\tau_\varphi = 0.040$ s). Humans showed weak depression in the auditory task and weak facilitation in the visual task.

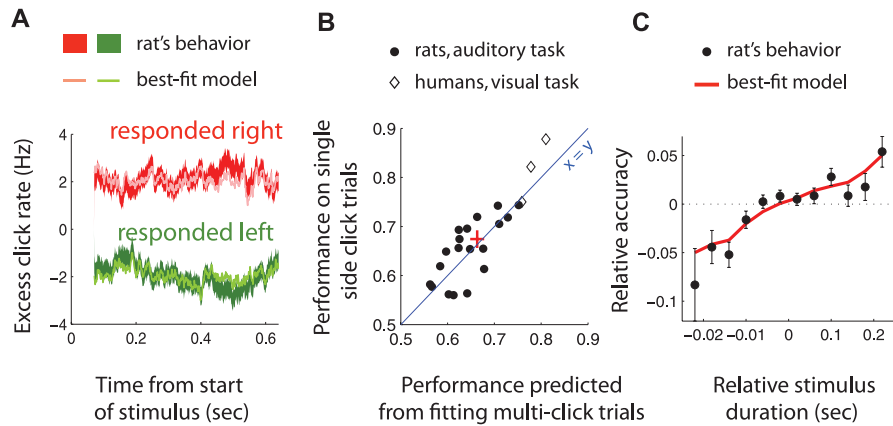


Fig. 3. Model-independent analyses support model-fitting results. (A) Long τ : psychophysical reverse correlation for an example rat (longest quarter of trials only). For each time point in each trial, we computed the excess pulse rate difference (right pulses/s – left pulses/s, relative to the value expected given the random processes used to generate the trial) and then obtained an average for trials resulting in a right (red) and an average for trials resulting in a left (green) decision. The separation between the two indicates how strongly clicks from each time point influenced the final decision. Thick solid lines were obtained from the rat's responses; the thickness of the line represents the standard error. Narrow shaded lines were predicted by the best-fit model. (B) Accurate estimates of σ_s^2 and σ_i^2 ; actual performance (fraction correct) on short-duration trials in which one side had one pulse and the other side had two pulses (i.e., trials for which performance was dominated by σ_s^2 and σ_i^2), versus performance predicted by fitting the model to trials with multiple pulses on each of the two sides. Too few trials of the short type to perform this analysis were presented for one of the human visual data sets and all human auditory data sets. Solid circles are individual rats, and open diamonds are individual humans. The red cross shows the mean and standard error across subjects. The accurate estimates of σ_s^2 and σ_i^2 suggested by the good predictions also suggest that σ_a^2 was estimated accurately, because the sum $\sigma_s^2 + \sigma_i^2 + \sigma_a^2$ is tightly constrained by the data (fig. S12). (C) Assessing σ_a^2 , ϕ , and τ_ϕ : Trials were divided into groups, with sensory evidence and sensory noise controlled by keeping the total number of right clicks and the total number of left clicks fixed within each group. Performance within each group was then dominated by σ_a^2 , ϕ , and τ_ϕ . (C) shows the performance of an example rat, averaged across trial groups and relative to the overall mean of each group, as a function of stimulus duration. The red line is the prediction from the best-fit model, with $\sigma_a^2 = 0$.

uses the full information about each trial's richly detailed stimulus (22), is a powerful approach for precisely quantifying multiple properties of decision-making processes. The approach provided strong evidence that rats can indeed gradually accumulate evidence for decision-making (23–26), thus establishing that this important cognitive phenomenon can be studied in a widely available animal model that is amenable to a rapidly growing arsenal of molecular tools.

With its capacity to provide moment-by-moment estimates of the temporal evolution of the accumulator, the approach will combine particularly well with neurobiological measurements. The model used for analysis can be readily expanded to consider and quantify further decision-making parameters, and the approach is easily generalized to different species, sensory modalities, and types of decision-making, including value-based decision-making (8, 27).

References and Notes

1. R. Ratcliff, *Psychol. Rev.* **85**, 59 (1978).
2. S. W. Link, *The Wave Theory of Difference and Similarity* (Lawrence Erlbaum Associates, Mahwah, NJ, 1992).
3. M. Usher, J. L. McClelland, *Psychol. Rev.* **108**, 550 (2001).
4. R. Bogacz, E. Brown, J. Moehlis, P. Holmes, J. D. Cohen, *Psychol. Rev.* **113**, 700 (2006).
5. R. Ratcliff, G. McKoon, *Neural Comput.* **20**, 873 (2008).
6. L. Paninski, J. W. Pillow, E. P. Simoncelli, *Neural Comput.* **16**, 2533 (2004).
7. J. I. Gold, M. N. Shadlen, *Annu. Rev. Neurosci.* **30**, 535 (2007).
8. M. Milosavljevic, J. Malmaud, A. Huth, C. Koch, A. Rangel, *Judgm. Decis. Mak.* **5**, 437 (2010).
9. B. W. Brunton, C. D. Brody, paper presented at the Annual Meeting of the Society for Neuroscience, Chicago, IL, 18 October 2009.
10. J. I. Sanders, A. Kepecs, *J. Neurophysiol.* **108**, 3416 (2012).
11. A. B. Watson, *Vision Res.* **19**, 515 (1979).
12. J. Ditterich, *Neural Netw.* **19**, 981 (2006).
13. A. C. Huk, M. N. Shadlen, *J. Neurosci.* **25**, 10420 (2005).
14. B. Wark, B. N. Lundstrom, A. Fairhall, *Curr. Opin. Neurobiol.* **17**, 423 (2007).
15. D. Bendor, X. Wang, *Nat. Neurosci.* **10**, 763 (2007).
16. M. Wehr, A. M. Zador, *Neuron* **47**, 437 (2005).
17. R. Kiani, T. D. Hanks, M. N. Shadlen, *J. Neurosci.* **28**, 3017 (2008).
18. I. Krajbich, C. Armel, A. Rangel, *Nat. Neurosci.* **13**, 1292 (2010).
19. P. Neri, A. J. Parker, C. Blakemore, *Nature* **401**, 695 (1999).
20. H. Nienborg, B. G. Cumming, *Nature* **459**, 89 (2009).
21. D. Raposo, J. P. Sheppard, P. R. Schrater, A. K. Churchland, *J. Neurosci.* **32**, 3726 (2012).
22. J. W. Pillow *et al.*, *Nature* **454**, 995 (2008).
23. N. Uchida, Z. F. Mainen, *Nat. Neurosci.* **6**, 1224 (2003).
24. N. M. Abraham *et al.*, *Neuron* **44**, 865 (2004).
25. D. Rinberg, A. Koulakov, A. Gelperin, *Neuron* **51**, 351 (2006).
26. H. Zariwala, Z. F. Mainen, thesis, Stony Brook University, Stony Brook, NY (2007).
27. I. Krajbich, A. Rangel, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 13852 (2011).
28. D. MacKay, *Information Theory, Inference, and Learning Algorithms* (Cambridge Univ. Press, Cambridge, 2003).

Acknowledgments: We thank Y. Niv for suggesting trial-by-trial analysis; D. Buonomano for suggesting we examine sensory adaptation; and A. Akrami, J. Erlich, C. Kopec, J. Kubanek, T. Hanks, B. Scott, M. Shadlen, D. Tank, and M. Yartsev for comments on the manuscript.

Supplementary Materials

www.sciencemag.org/cgi/content/full/340/6128/95/DC1
 Methods
 Supplementary Text
 Modeling Methods

11 December 2012; accepted 1 February 2013
 10.1126/science.1233912