

The Timescale of Perceptual Evidence Integration Can Be Adapted to the Environment

Ori Ossmy,^{1,7} Rani Moran,^{1,7} Thomas Pfeffer,^{2,3,4} Konstantinos Tsetsos,⁵ Marius Usher,^{1,8,*} and Tobias H. Donner^{2,3,6,8,*}

¹School of Psychology and Sagol School of Neuroscience, Tel-Aviv University, 69978 Ramat-Aviv, Israel

²Department of Psychology, University of Amsterdam, Weesperplein 4, 1018 XA Amsterdam, the Netherlands

³Cognitive Science Center Amsterdam, University of Amsterdam, 1018 WS Amsterdam, the Netherlands

⁴Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg- Eppendorf, 20246 Hamburg, Germany

⁵Department of Experimental Psychology, University of Oxford, 9 South Parks Road, Oxford OX1 3UD, UK

⁶Bernstein Center for Computational Neuroscience, Charité-Universitätsmedizin, Haus 6, Philippstraße 13, 10115 Berlin, Germany

Summary

A key computation underlying perceptual decisions is the temporal integration of “evidence” in favor of different states of the world. Studies from psychology and neuroscience have shown that observers integrate multiple samples of noisy perceptual evidence over time toward a decision [1–11]. An influential model posits perfect evidence integration (i.e., without forgetting), enabling optimal decisions based on stationary evidence [2, 3, 12]. However, in real-life environments, the perceptual evidence typically changes continuously. We used a computational model to show that, under such conditions, performance can be improved by means of leaky (forgetful) integration, if the integration timescale is adapted toward the predominant signal duration. We then tested whether human observers employ such an adaptive integration process. Observers had to detect visual luminance “signals” of variable strength, duration, and onset latency, embedded within longer streams of noise. Different sessions entailed predominantly short or long signals. The rate of performance improvement as a function of signal duration indicated that observers indeed changed their integration timescale with the predominant signal duration, in accordance with the adaptive integration account. Our findings establish that leaky integration of perceptual evidence is flexible and that cognitive control mechanisms can exploit this flexibility for optimizing the decision process.

Results

Numerous studies of perceptual decision-making in psychology and neuroscience have shown that human and animal decision-makers integrate multiple samples of noisy “evidence”

about the state of the outside world over time [1–11]. An influential model, the drift diffusion model, posits a perfect integration (i.e., without forgetting) of evidence toward a critical level, henceforth termed “decision bound.” Crossing the decision bound triggers the response and hence determines reaction time [2, 3, 12–14]. When the evidence is stationary (i.e., its mean does not vary over time), this model produces the fastest decisions for a fixed error rate [1, 2, 15]. In most real-life perceptual decisions, however, we face changing environments that yield changes in evidence over time (Figure 1A). Consider a radar operator who has to decide whether the trace displayed on the monitor corresponds to a missile, a passenger plane, or just “noise”: the operator has to search for a weak signal that emerges from a continuous stream of noise, at an unknown time, and needs to respond to it as soon as she detects the signal. Henceforth, we will refer to this situation as signal detection under nonstationary evidence and temporal uncertainty.

Here, we used a computational model of the decision process to show that, in this situation, perfect integration is sub-optimal, because it results in an excessive level of “false alarms” due to integration of presignal noise. Instead, we show that leaky (i.e., forgetful) integration, which limits the integration of presignal noise, is more suitable, provided that the decision-maker can adapt the integration time constant to the typical signal duration (e.g., the typical duration of the signal emitted by a missile). We simulated a leaky integrator model that detected signals of varying duration in protracted streams of noise using different integration time constants (Figure 1B; see also Figure S1 and Supplemental Experimental Procedures). The time constant corresponded to the time it took the integrator’s response to a sudden signal increase (step function) to reach $1 - (1/e)$ of its final, asymptotic value. The model’s detection threshold (the inverse of sensitivity, not to be confused with the decision bound shown in Figure S1) decreased as a function of signal duration in an approximately linear fashion in log-log coordinates [16]. This linear decrease of the logarithm of the threshold (i.e., the increase of sensitivity) with the logarithm of signal duration is a hallmark of temporal integration of sensory evidence [1, 4, 17, 18]. Its slope provides an index of the integration time constant (Figure S1C). For large time constants (approaching perfect integration), the slope was close to -1 . For small time constants (leak approaching 1), the decrease was shallower, governed by “probability summation” of correct detection events, rather than temporal integration [19]. Intermediate time constants yielded slopes between about -0.3 and -1 , like the ones shown in Figure 1B. As expected, the long time constant was better suited than the short time constant (i.e., yielded a lower threshold) for detecting the longest signals (compare red and blue lines in Figure 1B). However, crucially, for the shortest signals, the short time constant (blue line) was advantageous over the long time constant (i.e., lower threshold than for red line). Consequently, the threshold versus duration functions produced by the two different time constants intersected.

If human observers are able to adapt to the statistics of signal durations to enhance their performance, then they should, likewise, shift their integration timescale toward the

⁷Co-first authors

⁸Co-senior authors

*Correspondence: marius@post.tau.ac.il (M.U.), t.h.donner@uva.nl (T.H.D.)



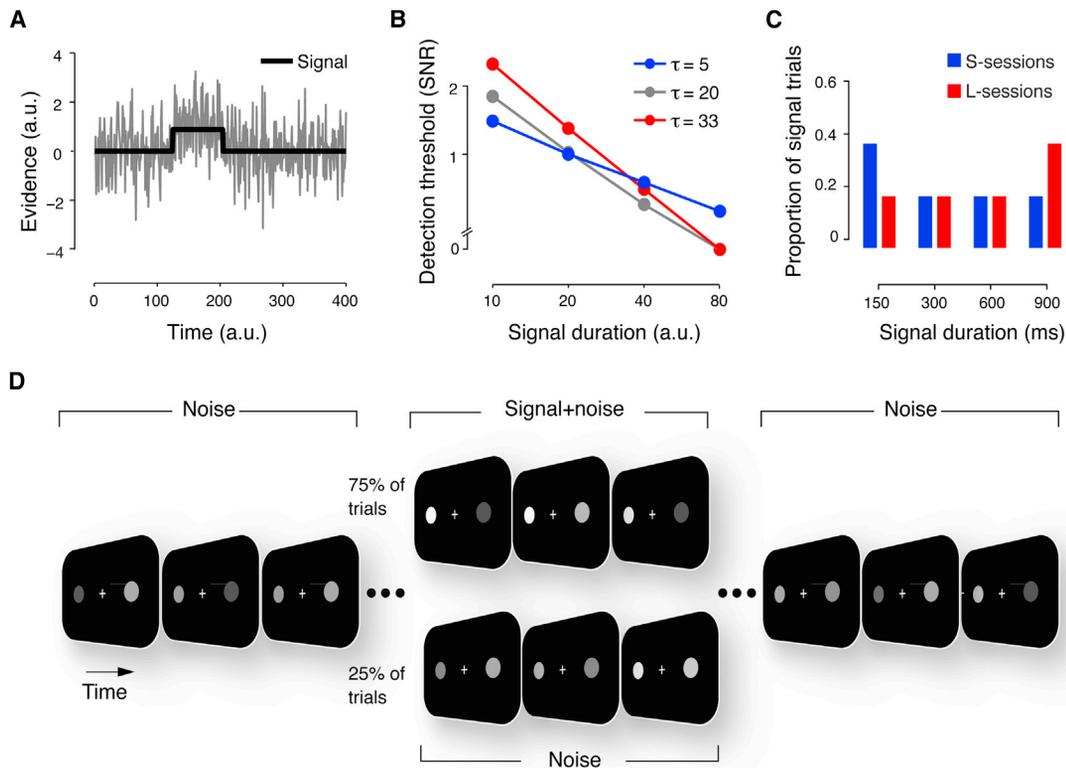


Figure 1. Model Predictions and Behavioral Detection Task

(A) Example time course of evidence for a detection decision with temporal uncertainty. A “signal” (an increment in the mean level, black line) is superimposed onto a continuous stream of random fluctuations (“noise”) at a random time, resulting in the noisy evidence (gray line). This time course describes any type of evidence that a decision-maker might integrate (e.g., the luminance of a visual stimulus or the value of a stock). (B) Detection thresholds as a function of signal duration for a leaky integrator model employing different time constants τ . Thresholds correspond to the signal strength yielding a hit rate of 80%. The decision bound was adjusted for each τ to maintain a fixed false alarm rate of 20%. The linear slope on log-log axes provides an estimate of τ . See also Figures S1–S3. (C) Histograms of the signal durations on signal + noise trials during S- and L-sessions of the psychophysical detection task. Observers were informed about these contingencies at the start of each session. (D) Schematic of the time course of the stimulus during an example signal + noise trial. The two discs fluctuated in luminance around a common mean level. During the signal interval, the mean luminance level of one of the discs increased (variable onset latency, magnitude, and duration). The signal is exaggerated for illustration. In the actual experiments, signal strengths were selected to span observers’ psychophysical detection threshold.

typical signal duration. This should be evident in the rate of observers’ performance (measured in terms of hit rate and detection threshold, respectively) improvement as a function of signal duration (Figure 1B): one should find a stronger improvement when long signals predominate compared to when short signals predominate. Thus, the hit rate versus duration functions (and threshold versus duration functions, respectively) measured under different predominant signal durations should intersect.

To test this prediction, we asked 12 human observers to detect visual “signals” embedded in a longer “noise” stream (Figures 1C and 1D; Experimental Procedures). On each “signal + noise” trial, the signal was an increment in mean luminance level of one of two fluctuating discs, which occurred at different latencies within the longer noise stream and varied (across trials) in duration and intensity (Figure 1D). Critically, we systematically manipulated the typical signal duration (and, thereby, presumably the observers’ expectation of signal duration) by presenting either predominantly the shortest or the longest signal durations within each of a number of different experimental sessions (henceforth referred to as “S-” or “L-sessions,” respectively; Figure 1C). Observers were informed about the predominant signal duration at the start of each session. We hypothesized that they would

integrate the difference between the two input streams on the left and right, respectively, and respond whenever this accumulated difference surpassed one of two symmetric decision bounds (a positive and negative one for left and right, respectively). More importantly, we further hypothesized that observers would employ a longer integration timescale in the L- than in the S-sessions.

This is what we found (Figures 2 and 3). All observers performed the task with low rate of false alarms (mean across observers: 14%; range: 8%–24%) and other errors (Table S1). Across the group, there was no significant difference in false alarms between L and S sessions ($t_{11} = 0.94$; $p = 0.36$). Hit rates increased monotonically with signal duration, indicating temporal integration of the signal. Importantly, for the shortest signals (150 ms), hit rates were significantly higher in the S- than in the L-sessions, while the opposite was the case for the longest signals (900 ms; Figure 2). Accordingly, in signal trials, there was a highly significant interaction between the session type (L versus S) and signal duration (two-way repeated-measures ANOVA with factors signal duration and session type; $F_{3,33} = 17.91$; $p < 10^{-3}$). This interaction was not evident in the reaction times (RTs). Although RTs increased with signal duration (main effect of signal duration: $F_{3,33} = 6.81$; $p = 0.02$; effect of session type: $F_{3,33} = 2.32$;

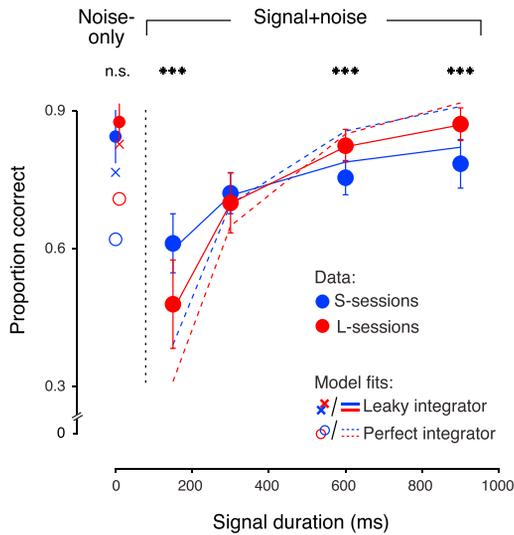


Figure 2. Interacting Effects of Expectation and Signal Duration on Hit Rate Reveal Adaptation of Integration Time Constant

The group average ($n = 12$) of proportion of correct choices for the noise-only trials and signal + noise trials are shown separately for the different signal durations and for L-sessions (red) and S-sessions (blue). For noise-only trials, numbers correspond to the correct rejection rate (no response). For signal + noise trials, numbers correspond to hit rates, collapsed across the five different signal strengths. Solid lines, average prediction of best-fitting leaky integrator model; dashed lines, average prediction of best-fitting drift diffusion model; error bars represent 95% confidence intervals; $***p < 10^{-3}$ (paired t test). See also Figure S2 and Tables S1–S3.

$p = 0.15$), there was no statistically significant interaction between signal duration and session type ($F_{3,33} = 0.9$; $p = 0.45$). Therefore, we focused our subsequent analyses on the accuracy data.

The behavioral results in Figure 2 indicate that observers indeed changed their integration time constant, depending on which signal duration predominated in a given session. We used two complementary approaches to quantify this effect and link it to our theoretical predictions. First, we compared alternative computational models of the decision process in their ability to account for the observers' behavioral data (i.e., the false alarm rates, and the hit rates as a function of signal strength and duration; Figures 2 and S2; Tables S2 and S3). One class of models incorporated our theoretical predictions: a leaky integrator, the timescale of which was free to vary with the session type (see Supplemental Experimental Procedures). This model provided a reasonable account of the psychophysical data (solid lines in Figure 2; see Figure S2A and Table S2 for individual observers). The integration time constants estimated by this model were consistently longer for L- than for S-sessions (group average: S: 80 ms; L: 490 ms; Wilcoxon signed-rank test: 10^{-3}).

The variant of the leaky integrator model shown in Figure 3 had different parameters for the internal “noise” in both session types and for an exponent describing the nonlinear relationship between physical signal intensity and neuronal input to the integrator at the decision stage. A simpler variant of this model, without these additional parameters, provided qualitatively identical results (Figure S2B).

We also fitted a perfect integrator (drift diffusion) model to the behavioral data (Supplemental Experimental Procedures). As expected, this model (dashed lines in Figure 2) consistently

provided a worse performance level (low correct rejection rates for noise-only trials) and a worse account of the data (underestimation of hit rate at short signals and overestimation of hit rate at long signals) than the leaky integrator model with variable time constant (see Figure S2A and Table S3 for individual observers). Quantitative model comparison was consistently in favor of the leaky integrator model (Table 1). The difference in the Bayesian information criterion (BIC, which takes into account both the goodness of fit and the number of model parameters) values, ranged between 82 and 766 across observers, providing strong support for the leaky integrator. In sum, fitting alternative models of the decision process strengthens the case for leaky integration with adaptive timescale.

In the second, model-independent approach, we estimated observers' psychophysical detection thresholds for all signal durations (see Figure 3A for an example observer; Supplemental Experimental Procedures). For all observers individually (Figure S3A), as well as for the group (Figure 3B), the thresholds were approximately linear in duration (in log-log coordinates). Just as for the model in Figure 1, the slopes of observers' empirical threshold versus duration functions depended on the expected signal duration (Figures 3B and 3C and Figure S3A). Slopes were significantly larger (i.e., more negative) in the L- than S-sessions in nine of the 12 individual observers ($p < 0.05$; one-sided permutation test). Further, the difference in slopes was highly significant when tested for the group (Wilcoxon signed-rank test; $p < 10^{-3}$; Figure 3C). We obtained qualitatively similar results when using an alternative, more constrained, approach for estimating the slopes of the threshold versus duration functions (see Supplemental Experimental Procedures and Figures S3B–S3D). Taken together, our results provide strong and complementary support for our hypothesis that observers changed their integration timescale with the predominant signal duration, in line with a decision process employing an adaptive timescale.

Discussion

The temporal integration of pieces of evidence supporting different choice options is a fundamental computational process underlying all decisions. This process operates during perceptual decisions like the one studied here [1–11], as well as during valued-based, economic decisions (e.g., a stock-buyer choosing stock options by integrating their fluctuating values) [14, 17, 20]. Here, we show that under conditions of nonstationary evidence with temporal uncertainty, perfect evidence integration is suboptimal and is not the computation employed by human decision-makers. To this end, we used an experimental protocol, which, albeit still far from real-life decisions, extends the standard laboratory tasks used for probing evidence integration in an important way. Our results demonstrate that human observers can boost their decision performance by flexibly changing the timescale of a leaky integration process according to changes in the expected signal duration.

Our estimates of integration timescales showed a consistent separation between long and short expected signals. But these timescales also varied substantially across observers and tended to be shorter than the actual duration of the typical signal, especially during the S-sessions. One possible explanation for the latter is that, in the S-sessions, some observers may have based their decision, in addition to the integrated signal, on transient responses at signal on- and offsets (see

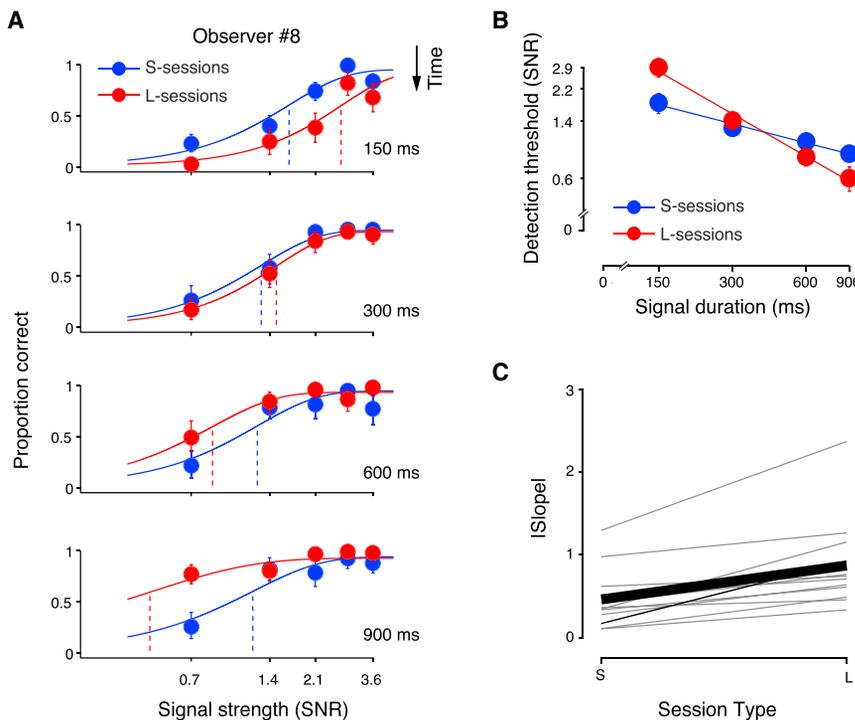


Figure 3. Change of Psychophysical Thresholds with Signal Duration Reveals Adaptation of Integration Time Constant

(A) Psychometric functions of one example observer, separately for the four signal durations and two session types. Psychophysical detection thresholds (expressed in units of signal-to-noise ratio; SNR) decreased with duration but more strongly in L- than in S-sessions. Solid lines, best-fitting cumulative Weibull functions. Vertical dashed lines indicate location of threshold parameter on the x axis. Error bars represent 95% bootstrap confidence intervals. (B) Group average thresholds, plotted on log-log scales as a function of signal duration. Error bars represent SEM. (C) Individual (thin lines) and group average (thick black line), best-fitting regression slopes for S- versus L-sessions. The absolute value of the slopes is shown (all fitted individual slopes were negative; see Figure S3A). The thin black line corresponds to the observer shown in (A). See also Figure S3.

legend of Figure S2A for details). In general, a substantial fraction of the interindividual variability in the performance of our 12 observers may reflect differences between strategies (or integration capacities) among individuals. Future studies should examine more complex models of the decision process to capture such individual differences.

Our results have a number of implications for understanding the mechanisms of decision-making. First, they provide strong support for models of perceptual choice that are based on leaky, rather than perfect, integration of perceptual evidence. The notion of leaky integration is inspired by fundamental principles of neural computation [4, 18], and it is consistent with neurophysiological data suggesting a reservoir of time constants in the cerebral cortex [21, 22]. Leaky integration is also consistent with a common finding in psychophysical experiments, which manipulated the stimulus duration and hence available decision time, showing that observers integrate sensory evidence only across limited periods of time ([8, 10, 11, 23, 24], but see [8] for an alternative interpretation involving perfect integration). While information leak has so far been regarded as a limitation in perceptual decision-making, our results reveal that it is, in fact, advantageous for real-life decisions: the integration timescale provides a critical degree of freedom for adapting the decision process to the environmental contingencies. Similar conclusions have been reached for the temporal integration of reward (across trials) in a dynamic foraging task, in which monkeys used a time constant that was closely matched to the statistics of the environment [25]. However, in this study, the optimal timescale was not systematically manipulated.

Our results also shed new light on the question of how “top-down” cognitive control mechanisms shape decision computations [26]. Since the emergence of signal detection theory in perceptual psychophysics [16], perceptual decision-making has been viewed as a two-component process: (1) the encoding and integration of sensory evidence and (2) the criterion

level (decision bound), against which the integrated evidence is compared to reach a decision. In this view, decision-makers exert control over the decision process only by adapting the decision bound [1–3, 27, 28], while the evidence integration operates automatically. By contrast, our results reveal that adaptive control mechanisms can directly shape the evidence integration computation.

Further, the results imply that the brain is remarkably flexible in selecting an integration timescale suitable for the environmental context at hand. Timescales of several hundreds of milliseconds, like the ones observed here, are probably an emergent property of dynamic network interactions rather than a fixed property of individual neurons [4, 18, 24]. Physiological evidence suggests that such network interactions during decision-making span multiple regions that are widely distributed across the brain [29–31]. The adaptive changes in integration timescales found in the present study may, therefore, reflect the flexible adjustment of such large-scale decision networks, which might be achieved by neuromodulation [32–34].

Finally, our study sets the stage for future neurophysiological studies of the biophysical circuit mechanisms underlying perceptual evidence integration. Such studies could exploit our experimental protocol to experimentally manipulate integration timescale, while characterizing neurophysiological signatures of the evidence integration mechanism [22, 29].

Experimental Procedures

Below, we describe the psychophysical experiments and provide a brief summary of the psychometric fitting procedures. All data analysis and modeling procedures are described in detail in the [Supplemental Experimental Procedures](#).

Observers

Twelve healthy observers (eight females, age range: 20–37 years, normal or corrected-to-normal vision) participated in this study after informed consent. The observers were psychology students at TAU, who were naive to the purpose of the study, did the experiment for credit, and were, in addition, paid for their participation in proportion to task performance (20–40 NIS per hr). The study was approved by the local ethics committee.

Table 1. Comparison between Leaky and Perfect Integrator Models

Observer	Perfect Integrator	Leaky Integrator	Leaky Better?
1	468.06	298.32	TRUE
2	596.55	263.14	TRUE
3	527.07	300.60	TRUE
4	505.74	278.92	TRUE
5	436.49	244.50	TRUE
6	483.94	313.87	TRUE
7	1,144.75	379.02	TRUE
8	737.00	358.21	TRUE
9	1,016.56	494.46	TRUE
10	346.15	263.81	TRUE
11	470.75	332.18	TRUE
12	624.11	445.98	TRUE

Values correspond to Bayesian information criterion (BIC). BIC differences perfect integrator – leaky integrator of >6 are considered strong support for the leaky integrator model.

Stimuli

Stimuli were presented on a linearized CRT screen (frame rate: 100 Hz). On each trial, two discs were presented in the right and left visual hemifields on a black background, for a total duration of 5 s (Figure 1C). Discs were presented at an eccentricity of 5° and subtended about 2.85° of visual angle. We controlled the level of (external) noise entering the decision process by randomly and independently changing each disc's luminance level at each monitor refresh (10 ms). These random luminance fluctuations were drawn from a truncated Gaussian noise (SD after truncation = 0.11), and they were added to each disc's mean luminance level. The truncation was used to prevent the luminance variable to exceed the (0,1) boundaries (see below). Throughout each trial, the mean luminance level had one of two values. Twenty-five percent of the trials ("noise-only" trials) contained only fluctuations around the same baseline level (0.40, on a 0–1 scale). This was constructed by using RGB values with equal components. In the remaining 75% ("signal + noise" trials), both discs fluctuated around the same baseline level for most of the trial, and a luminance increment was added to the baseline level of one of the discs at various different latencies within the noise stream (uniform distribution of signal onsets; range: 0.6–3.5 s). This increment was the signal that observers had to detect. The signal's location (left or right), strength (0.08, 0.16, 0.24, 0.32, or 0.40), and duration (150, 300, 600, or 900 ms) were randomly selected on each signal + noise trial.

Procedure

Observers' task was to maintain fixation throughout the whole trial and to respond whenever they judged that one of the discs had increased in brightness (the signal; Figure 1D). They had to indicate the location of the signal by pressing a left or a right button within a predetermined response window (from signal onset to 600 ms after signal offset). Responses to signals were classified as correct ("hits") if they were made within the response window and with the correct button. There were four possible types of errors: "false alarms" (response on noise-only trials or before the response window on signal + noise trials), "slow responses" (response on signal + noise trials after the response window), "misses" (no response whatsoever on signal + noise trials), or "mislocalization" (response on signal + noise trials within response window but with incorrect button). At the end of "mislocalization" trials, a sound feedback was provided. At the end of the other error trials, a text on the screen informed observers about the type of error. After the completion of each trial, observers pressed a key to continue to the next trial. In the absence of any key press, the next trial started automatically after 2.5 s. Every 100 trials, observers were allowed to pause.

After a short practice session, each observer performed several experimental sessions of 500 trials each (duration: ~1 hr). All sessions consisted of 125 noise-only trials and 375 signal + noise trials. To manipulate the observers' expectation of signal duration, we introduced a predominance of a factor of two (rate of occurrence: 0.4 versus 0.2) of either the longest or the shortest signal duration over the other three durations (Figure 1C). In one type of session ("S-sessions"), there were 150 trials (5 signal strengths × 30 trials) of the shortest duration (150 ms). In the other type of session ("S-sessions"), there were 150 trials of the longest duration (900 ms). In all sessions, there were 75 trials (5 signal strengths × 15 trials) of the remaining three durations (see Supplemental Experimental Procedures for rationale behind these distributions of signal durations.)

The two types (S/L) were alternated from session to session, with the order counterbalanced across observers. Observers received an explicit instruction about each session's type before the start of testing. Nine observers performed four sessions. Three observers performed six sessions.

Data Analysis and Model Fits

We computed observers' psychophysical detection performance as the percentage of hits on signal + noise trials, separately for each signal duration and intensity, and as the percentage of no responses ("correct rejections") on noise-only trials (Figure 2). We fitted two different classes of computational models of the decision process (leaky integrator and drift diffusion) to the complete performance data of each individual observer (including noise-only trials). To estimate observers' detection thresholds, we fitted a cumulative Weibull function to the hit rates as a function of signal strength and extracted the threshold parameter of the best fits (Figure 3). See Supplemental Experimental Procedures for details of the data analysis and model fitting procedures.

Supplemental Information

Supplemental Information includes Supplemental Experimental Procedures, three figures, and three tables and can be found with this article online at <http://dx.doi.org/10.1016/j.cub.2013.04.039>.

Acknowledgments

We thank Michael Herrmann for helpful discussions on the computational model. M.U. is funded by the Israeli Science Foundation (grant: 743/12) and by the German Israeli Foundation (grant, 1130-158.4/2010).

Received: November 22, 2012

Revised: April 11, 2013

Accepted: April 15, 2013

Published: May 16, 2013

References

- 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.
- Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.* 30, 535–574.
- Smith, P.L., and Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends Neurosci.* 27, 161–168.
- Usher, M., and McClelland, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychol. Rev.* 108, 550–592.
- 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.
- Donner, T.H., Siegel, M., Fries, P., and Engel, A.K. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision making. *Curr. Biol.* 19, 1581–1585.
- Gold, J.I., and Shadlen, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature* 404, 390–394.
- 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.
- van Ravenzwaaij, D., van der Maas, H.L., and Wagenmakers, E.J. (2012). Optimal decision making in neural inhibition models. *Psychol. Rev.* 119, 201–215.
- 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.
- Gao, J., Tortell, R., and McClelland, J.L. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS ONE* 6, e16749.
- Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Comput.* 20, 873–922.
- Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J. Neurosci.* 22, 9475–9489.

14. Krajbich, I., and Rangel, A. (2011). Multialternative drift-diffusion model predicts the relationship between visual fixations and choice in value-based decisions. *Proc. Natl. Acad. Sci. USA* *108*, 13852–13857.
15. Wald, A., and Wolfowitz, J. (1948). Optimum characteristic of sequential probability ratio test. *Ann. Math. Stat.* *19*, 326–339.
16. Green, D.M., and Swets, J.A. (1966). *Signal Detection Theory and Psychophysics* (New York: Wiley).
17. Busemeyer, J.R., and Townsend, J.T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* *100*, 432–459.
18. Wang, X.J. (2002). Probabilistic decision making by slow reverberation in cortical circuits. *Neuron* *36*, 955–968.
19. Watson, A.B. (1979). Probability summation over time. *Vision Res.* *19*, 515–522.
20. Tsetsos, K., Chater, N., and Usher, M. (2012). Saliency driven value integration explains decision biases and preference reversal. *Proc. Natl. Acad. Sci. USA* *109*, 9659–9664.
21. Bernacchia, A., Seo, H., Lee, D., and Wang, X.J. (2011). A reservoir of time constants for memory traces in cortical neurons. *Nat. Neurosci.* *14*, 366–372.
22. Honey, C.J., Thesen, T., Donner, T.H., Silbert, L.J., Carlson, C.E., Devinsky, O., Doyle, W.K., Rubin, N., Heeger, D.J., and Hasson, U. (2012). Slow cortical dynamics and the accumulation of information over long timescales. *Neuron* *76*, 423–434.
23. Burr, D.C., and Santoro, L. (2001). Temporal integration of optic flow, measured by contrast and coherence thresholds. *Vision Res.* *41*, 1891–1899.
24. Uchida, N., Kepecs, A., and Mainen, Z.F. (2006). Seeing at a glance, smelling in a whiff: rapid forms of perceptual decision making. *Nat. Rev. Neurosci.* *7*, 485–491.
25. Sugrue, L.P., Corrado, G.S., and Newsome, W.T. (2004). Matching behavior and the representation of value in the parietal cortex. *Science* *304*, 1782–1787.
26. Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. *Annu. Rev. Neurosci.* *24*, 167–202.
27. Botvinick, M.M., Braver, T.S., Barch, D.M., Carter, C.S., and Cohen, J.D. (2001). Conflict monitoring and cognitive control. *Psychol. Rev.* *108*, 624–652.
28. Lo, C.C., and Wang, X.J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nat. Neurosci.* *9*, 956–963.
29. Donner, T.H., Siegel, M., Oostenveld, R., Fries, P., Bauer, M., and Engel, A.K. (2007). Population activity in the human dorsal pathway predicts the accuracy of visual motion detection. *J. Neurophysiol.* *98*, 345–359.
30. Pesaran, B., Nelson, M.J., and Andersen, R.A. (2008). Free choice activates a decision circuit between frontal and parietal cortex. *Nature* *453*, 406–409.
31. Siegel, M., Donner, T.H., and Engel, A.K. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nat. Rev. Neurosci.* *13*, 121–134.
32. Aston-Jones, G., and Cohen, J.D. (2005). An integrative theory of locus coeruleus-norepinephrine function: adaptive gain and optimal performance. *Annu. Rev. Neurosci.* *28*, 403–450.
33. Eckhoff, P., Wong-Lin, K.F., and Holmes, P. (2009). Optimality and robustness of a biophysical decision-making model under norepinephrine modulation. *J. Neurosci.* *29*, 4301–4311.
34. Usher, M., and Davelaar, E.J. (2002). Neuromodulation of decision and response selection. *Neural Netw.* *15*, 635–645.