

UC Davis

UC Davis Previously Published Works

Title

Flexibility of Timescales of Evidence Evaluation for Decision Making

Permalink

<https://escholarship.org/uc/item/8wj7s7xk>

Journal

Current Biology, 29(12)

ISSN

0960-9822

Authors

Ganupuru, Preetham
Goldring, Adam B
Harun, Rashed
[et al.](#)

Publication Date

2019-06-01

DOI

10.1016/j.cub.2019.05.037

Peer reviewed



Published in final edited form as:

Curr Biol. 2019 June 17; 29(12): 2091–2097.e4. doi:10.1016/j.cub.2019.05.037.

Flexibility of timescales of evidence evaluation for decision making

Preetham Ganupuru¹, Adam B. Goldring¹, Rashed Harun¹, Timothy D. Hanks²

Department of Neurology and Center for Neuroscience, University of California Davis, 1544 Newton Ct, Davis, CA, USA 95618

Summary

To understand the neural mechanisms that support decision making, it is critical to characterize the timescale of evidence evaluation. Recent work has shown that subjects can adaptively adjust the timescale of evidence evaluation across blocks of trials depending on context [1]. However, it's currently unknown if adjustments to evidence evaluation occur online during deliberations based on a single stream of evidence. To examine this question we employed a change detection task in which subjects report their level of confidence in judging whether or not there has been a change in a stochastic auditory stimulus. Using a combination of psychophysical reverse correlation analyses and single-trial behavioral modeling, we compared the time period over which sensory information has leverage on detection report choices versus confidence. We demonstrate that the length of this period differs on separate sets of trials based on what's being reported. Surprisingly, confidence judgments on trials with no detection report are influenced by evidence occurring earlier than the time period of influence for detection reports. Our findings call into question models of decision formation involving static parameters that yield a singular timescale of evidence evaluation and instead suggest that the brain represents and utilizes multiple timescales of evidence evaluation during deliberation.

Keywords

Decision making; audition; psychophysics; human; timescale; change detection; confidence; evidence evaluation; cognitive flexibility; behavioral modeling

Results and Discussion

The adaptive selection of behavior requires choosing appropriate actions based on available information. In many instances, adaptive behaviors are guided by detecting subtle signals in

Correspondence to: Timothy D. Hanks, thanks@ucdavis.edu.

Author Contributions

Conceptualization, T.D.H., P.G., R.H., A.B.G.; Methodology, T.D.H., P.G., R.H., A.B.G.; Software, T.D.H., P.G., R.H., A.B.G.; Validation, R.H.; Formal Analysis, T.D.H., P.G., R.H.; Investigation, T.D.H., P.G., R.H., A.B.G.; Resources, T.D.H.; Data Curation, T.D.H., P.G., R.H., A.B.G.; Writing -- Original Draft, T.D.H., P.G., R.H., A.B.G.; Writing -- Review and Editing, T.D.H., P.G., R.H., A.B.G.; Visualization, T.D.H., P.G., R.H., A.B.G.; Supervision, T.D.H., P.G., R.H., A.B.G.; Funding Acquisition, T.D.H.

¹Equal contribution

²Lead contact

Declaration of Interests

The authors declare no competing interests.

a dynamic environment. Previous studies have shown that humans, monkeys, and rodents are capable of quickly extracting information about the variability of changes in dynamic environments [2–5]; moreover, subjects can alter their timescale of evidence evaluation -- that is, the time period over which evidence has leverage over a decision -- based on the expected duration of signals in order to make judgments of when an actual signal occurs [1]. After individuals make a decision, the past evidence can be utilized for additional purposes, including the judgments of the degree of confidence that the selected option is correct [6–22]. To shed light on the flexibility of evidence evaluation, we examined if and how different timescales of evidence are utilized for change detection reports while subjects performed an auditory change detection task compared to confidence judgements in trials without detection reports.

Auditory change detection task

We trained subjects to perform an auditory change detection task in which they reported a change in the underlying rate in a sequence of auditory clicks generated by a stochastic Poisson process. Trials began when a subject placed their finger into a central port, which was followed by the onset of the auditory stimulus (Figure 1). The underlying rate was initially 50 Hz, and for 70% of trials, the rate increased at a random time and by a variable magnitude. The other 30% of trials ended without a change (catch trials). Subjects were required to remove their finger from the central port within 800 ms of change onset (hit) or withhold responding for catch trials (correct rejection; CR). There were two types of errors in this task: premature responses (false alarms; FA), which can occur in catch and non-catch trials, and failures to respond in time (misses). In all cases, subjects were then cued to report confidence using the two side ports. Confidence was assessed via a post-decision wager. Immediately following the confidence report, feedback was given via an auditory tone to indicate success or failure on that trial.

Task performance and confidence ratings

All subjects were able to perform the task with high hit rates for the easiest trials and diminishing hit rates for medium and high difficulty trials (Figure 2A, left; see Figure S1 for individual subject data), suggesting that they were attending to the stimuli. During catch trials, subjects displayed CR rates of ~85% or above, with a ~15% FA rate across both trial types (Figure 2A, right). For both hits and misses, confidence scaled with trial difficulty, with the largest changes in click rates evoking the highest level of confidence during hits and the lowest level of confidence during misses (Figure 2B, left). CR and FA trials by definition involved no change in the underlying generative click rate and evoked intermediate levels of confidence compared to hits and misses (Figure 2B, right).

Psychophysical reverse correlation

To examine the timescale of evidence evaluation, we first conducted psychophysical reverse correlation (RC) analyses [15,23–25]. RC traces were constructed by convolving click times with causal half-Gaussian filters ($\sigma = 0.05s$) and aligning the result to the end of the stimulus presentation. This allowed us to reconstruct the average stimulus that preceded a given trial outcome and confidence rating. We focused our initial analyses on FA trials and their associated confidence ratings. On these trials, responses were only affected by natural

fluctuations in the stochastic stimulus and not tied to a generative change in click rate as occurs for hit trials. To examine overall influence on choice independent of confidence, all FA RC traces were averaged together. Across subjects, FA choices were characterized by an average RC trace (hereafter referred to as the detection report kernel) showing a transient increase in click rate, which followed a time-course with a duration similar to the response window in which subjects were allowed to report an actual change in the generative click rate (Figure 3A; see Figure S2 for individual subject data). There was a sharp increase starting ~800 ms before the detection report that collapsed to baseline just before the report, indicative of sensorimotor delays limiting the influence of the time period just before the response. Comparing RCs from high and low confidence FA trials (Figure 3B; see Figure S2 for individual subject data), we found that evidence for confidence judgements was used during the period after which the detection report kernel returned to baseline (Figure S3), consistent with previous work using a different task design that showed that confidence is based on continued accumulation of evidence after the decision but before the confidence response [26].

We next asked whether there is an influence of evidence on confidence that extends earlier in time than its influence on detection reports as this would be indicative of evidence being evaluated at multiple timescales. Surprisingly, we found that the point in time when RC kernels deviated from baseline differs depending on what is reported. In particular, kernels began to deviate from baseline for CR confidence reports earlier in time than for detection reports. To quantify this for detection reports, we fit the ascending phase of the detection report kernel with a 2-piece linear function (see methods). Across subjects, the parameter estimate of the detection report kernel start point was ~0.74 seconds (95% CI: 0.73 to 0.75 s) preceding the detection report (Figure 3A, arrow).

In contrast, the period of influence of evidence for CR confidence reports extended considerably earlier in time. High confidence CRs were characterized by a lower average click rate preceding the end of the stimulus compared with low confidence CRs (Figure 3C; see Figure S2 for individual subject data). This difference gradually increased until the end of the trial. To estimate the point in time when the difference between high and low confidence CR reverse correlation deviated from baseline, we fit the difference (the “CR confidence difference kernel”) with a 2-piece linear function (see methods). Across subjects, the pooled parameter estimate for when the CR confidence difference kernel diverged (i.e. differed from 0 Hz) was ~1.76 seconds before the end of the trial (95% CI: 1.68 to 1.84 s), more than twice as early as the estimate of the start of the detection report kernels. This suggests that during the course of a trial, subjects have flexibility in the timescale of evidence evaluation, with different timescales more strongly linked to different types of reports.

Model-based analysis

While psychophysical RC analyses provide useful information for comparing the timescales of evidence evaluation, they are not a veridical representation of how evidence is temporally evaluated because they reflect the influence of a number of different components of evidence processing [27]. Therefore, we adopted a model-based approach to ask whether a single

timescale of evidence evaluation can explain the differing start points revealed by the RC analyses. In the model, evidence in the form of the auditory clicks was convolved with a half-Gaussian filter with its width (σ) as a free parameter (Figure 4A). The filter width determined the timescale of evidence evaluation, with a wider filter corresponding to a longer timescale. The output of the filter governed the average dynamics of a decision variable. Variability was included in the process by adding Gaussian noise to the decision variable at each timestep with the standard deviation of the Gaussian as a second free parameter in the model. A third free parameter set a decision bound that caused triggering of a detection report when the decision variable reached that value. To account for sensorimotor delays inherent to decision processes, an additional period of “non-decision time” described by a Gaussian distribution with mean set by a fourth free parameter was added to the bound-crossing time to determine the final response time.

We first used this model to capture the choice behavior exhibited by our subjects. In particular, we found the values for the four free parameters of the model that best fit the trial-by-trial choice responses made by our subjects (see methods). This yielded psychometric functions that closely approximated the behavioral data (Figure 4B). We then used model simulations with the best fit parameters to extract predictions for the detection report kernels and CR confidence difference kernels. These trials were classified as high or low confidence based on whether the final decision variable was higher or lower than a threshold set to 49 Hz. Comparing the predicted kernels to the actual data, we found that this single-timescale model predicted a start of the detection report kernel that was slightly earlier than the data (29.3 ± 6.4 ms earlier). Critically, this model predicted a start of the divergence of the CR confidence difference kernels that was substantially later than the experimental data (894.8 ± 12.6 ms later) (Figure 4C). We found similar results with a variety of filter shapes and non-decision time distributions (data not shown). We also tested whether trial-to-trial variability in decision bound could extend the divergence of CR confidence difference kernels closer to the experimental data, but we found no set of parameters capable of doing so (Figure S4). Finally, we extended the timescale of evidence evaluation in the model with a longer filter width to recapitulate the experimental divergence of the CR confidence difference kernels. In doing so, we found that the model predicted a start of the detection report kernel that was substantially earlier than the data (1082.6 ± 40.1 ms earlier). These analyses confirm our intuition that a neural mechanism with a single timescale of evidence evaluation cannot explain the result that the divergence of CR confidence difference kernels can extend more than twice as early relative to trial end than detection report kernels.

Our study demonstrates that, in perceptual decisions, the timescale of evaluation of past evidence can be flexibly adjusted for use in detection and separately for confidence judgments when no detection is reported. This was revealed using a change detection task in which perfect integration of evidence is suboptimal. With perfect integration, evidence accumulated early retains its influence on the decision for the full duration of the deliberation period, which is the optimal strategy during perceptual discrimination tasks based on the full stimulus. In contrast, evidence has a more transient influence in our change detection task [24]. By including an additional confidence report, this paradigm allowed us to test whether the limited temporal influence of evidence on decision formation was similar

for confidence judgments. We found that kernels for judgments of confidence could be influenced by evidence fluctuations earlier in time on trials without a detection report. Interestingly, information within different temporal epochs appears to be used for confidence judgments in a way that depends on how the trial ended. When the subject terminated the evidence stream by reporting a detection, as in the case of FAs, the confidence-influencing epoch began at approximately the same time as the detection report kernel. However, when the trial ended due to withholding a detection report, as in the case of CRs, the confidence-influencing epoch extended several hundred milliseconds earlier than the detection report kernel. While averaging over many trials prevents us from quantifying which epochs of time had influence on a trial-to-trial basis, the magnitude of the mean psychophysical kernel nonetheless relates to the magnitude and frequency of evidence evaluation during a given epoch across trials.

Model-based analyses confirmed that our results cannot be explained by a neural mechanism involving a single, fixed timescale of evidence evaluation. The differences in the evidence evaluation period that we found between conditions, which reveal flexibility in the process, provide insights for the requirements of any specific mechanism responsible for this discrepancy. In particular, our results suggest that the mechanism, whatever it may be, has the capacity to either adjust the timescale of evidence weighting during individual deliberative decisions or access multiple distinct timescales for different purposes. We discuss possible representational architectures of the brain that would enable this below.

It is unclear why subjects adjusted their timescales of evidence evaluation depending on how the trial terminated. It is optimal to utilize evidence over the last 800ms for detection and associated confidence because if there was a change, it could have only happened in this period in this task. Thus, it is suboptimal for subjects to base their confidence on earlier evidence, as was observed in CR trials. However, we suggest it may be optimal in the more general class of change detection decisions individuals encounter in real life. Typically, confidence in a choice based on a perceived change should be judged based on recent evidence that evoked the perceived change. In contrast, confidence that no change has occurred in a real-world situation often involves judgment based on a longer interval of time during which a change would have been possible. Our results are consistent with a strategy that would be appropriate for that more general type of situation.

Our results show that past evidence must be represented in a way that allows flexibility in the timescale of evidence evaluation. This extends the idea that adapting the timescale of evidence evaluation is necessary to optimize decision processes in changing environments [1,2,4,28]. In those studies, the dynamics of the environment dictate the optimal timescale, but only one timescale needs to be accessible for any given context or trial. In contrast, we find that multiple timescales of evidence are used within the same context. Thus, mechanisms are necessary to adjust the timescale of evaluation even as the evidence is being presented and used. Mid-deliberation adjustments of decision processes have been described extensively with regard to decision bounds. Collapsing decision bounds can furnish urgency onto the decision process [29–31], and changes of mind about decisions and confidence are best explained with altered post-commitment decision bounds [26]. In most of those cases, decisions were made in situations that involved near perfect integration of sensory evidence,

so there was no opportunity to look for changes in the timescale of evidence evaluation at earlier periods of deliberation. Here we show this timescale to be an important factor that can be adjusted online during deliberations based on a single stream of evidence.

Previous studies with yes-no detection tasks have suggested separate neural representations for stimulus present and stimulus absent choices [32,33]. While those tasks involved a delay between the stimulus epoch and the choice report (unlike our task), the separate neural representations they found could provide a substrate for distinct timescales of evidence evaluation. Thus, one possible neural architecture that could explain our results is having one population of neurons evaluating evidence over a shorter timescale for the detection decision and another population of neurons evaluating evidence in parallel over a longer timescale for confidence judgments on trials without a detection report. We suggest that neural mechanisms that allow flexibility in the timescale of evidence evaluation may be used in the service of multiple components of decision making behavior, rather than flexibility being a unique feature of decisions that are combined with confidence judgments.

It is tempting to speculate that recurrent network models of integration that require a fine level of tuning to avoid leaky dynamics [34], which is usually viewed as a shortcoming [35], may instead be a virtue by providing flexibility in the timescale of evidence evaluation. Under this idea, changes in the timescales of evidence evaluation could be introduced through small adjustments of tuning that would result in altered leakiness of integration. Alternatively, flexibility in the timescale of evidence evaluation could be implemented at an earlier stage of sensory processing, such as through gating of deliberation by stimulus salience [36,37]. In this alternative schema, sensory responses must exceed a salience threshold to be considered for the decision process, and alterations of the salience threshold would influence the evidence evaluation process. Neither of these mechanisms alone allows use of multiple timescales of evidence evaluation for the same stream of evidence. Memory traces that allow recall and re-processing of past evidence would be one mechanism to use multiple timescales of evidence evaluation for the same stream of evidence. Another related mechanism that would allow parallel access to multiple timescales for the same stream of evidence derives from theoretical work showing that memory traces may be encoded through neurons with heterogeneous dynamics that form a temporal basis set for previous events [38]. Selective readout of sets of neurons with differing timescales would allow flexible access that depends on task demands [39,40]. This would be readily achievable in networks that encode accumulated evidence with a diversity of timescales [41]. This mechanism could also allow adjustments of the timescale of evidence evaluation even after the evidence has been presented, which would enable meta-cognitive operations [11,42,43]. We therefore suggest that our paradigm provides a powerful approach to understand the neural mechanisms and representational architectures in the brain that could support meta-cognitive operations for decision making.

STAR Methods

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, Timothy Hanks (thanks@ucdavis.edu).

EXPERIMENTAL MODEL AND SUBJECT DETAILS

There were 7 subjects (2 female, 5 male) included in this study, all aged 18–34 and members of UC Davis. For the 7 subjects included for analysis, 3 subjects (S1, S2, S3) were knowledgeable about the task design and research motivations prior to data collection, while the remaining subjects were naive. Study procedures were approved by the UC Davis Institutional Review Board, and all subjects provided informed consent. Subjects were compensated with a \$10 Amazon gift card for each 1-hour experimental session completed, for a total of 6–11 sessions. Each subject received full payment, irrespective of task performance.

METHOD DETAILS

Apparatus—Control of the task was programmed in MATLAB (Mathworks, RRID: SCR_001622) and facilitated by Bpod (Sanworks, RRID: SCR_015943), which measures output of behavioral tasks in real time. Task stimuli were generated by the open source device Pulse Pal [44]. The stimulus-response apparatus consisted of 3 cone-shaped ports, each containing an infrared LED beam that can detect the insertion of a finger when the beam is obstructed. Each port can also be illuminated by an LED light, which signals to the subject that the port can be used during that stage of the trial. Sounds were played through headphones worn by the subject.

Change detection task—Subjects began each trial by inserting their index finger into the illuminated center port of the apparatus, which initiated a train of auditory clicks randomly generated by a Poisson process. The initial baseline frequency of this click train was 50 Hz, and the stimulus persisted at this frequency for a variable time period, during which the subject was to keep their finger in the port. In 70% of trials, the frequency of the stimulus increased with a magnitude of 10, 30, or 50 Hz at a random time sampled from a truncated exponential distribution (minimum 0.5 s, maximum 10 s, mean 4 s). This sampling produced an approximately flat hazard rate, such that the instantaneous probability of a change at the given moment did not increase or decrease as the trial progressed. When a change occurred, the subject was to respond by removing their finger from the port within 0.8 s of the change. The stimulus ended immediately upon finger removal. In the remaining 30% of trials (“catch” trials), no frequency increase occurred; in these trials, the subject was to maintain finger insertion for the full duration of the stimulus, which ended at a random time from 0.5 to 10.8 s. The same exponential distribution was used as for the change times in the non-catch trials plus the 0.8 s response window in order to match the distribution of catch trial durations to that of non-catch trials. Thus, the timing of trial termination provided no information about catch versus non-catch trial.

Finger removal occurring within the 0.8 s following a change was recorded as a “hit”. Failure to correctly respond in time to a change was recorded as a “miss”. Correctly responding to catch trials required the subject to maintain finger insertion until the stimulus ended, which was recorded as a “correct rejection” (CR). Whereas, if a subject removed their finger from the port while there was no change in the generative rate of clicks, either on a catch or non-catch trial, the response was recorded as a “false alarm” (FA).

After the auditory stimulus concluded, the peripheral ports of the apparatus illuminated, cuing the subject to report confidence in the decision. Subjects were given a choice between “low” or “high” confidence, which was reported by inserting a finger into either the left or right peripheral port, respectively. Subjects were instructed to report low confidence if the subject was “probably not successful” and high confidence if the subject was “probably successful” in the trial. Performance was tracked by a points system: Reporting high confidence on a correct decision awarded the subject with 2 points, while low confidence on a correct decision yielded only 1 point. Reporting high confidence on an incorrect decision cost the subject 3 points, while reporting low confidence on these trials cost the subject only 1 point. A running total of accumulated points in the experimental block was displayed on a monitor in front of the subject as a blue bar that changed size with the points total, which could not fall below 0 points. This points scheme encouraged subjects to report high confidence for trials in which the evidence especially favored their choices, because they were asymmetrically punished for erroneous high confidence reports. Subjects then received auditory feedback on their initial decisions, regardless of confidence report, indicating whether the response was correct. The center port then illuminated once again, allowing the subject to start a new trial.

If the subject removed their finger in response to a perceived change, a brief noise was played through the headphones to indicate the response preceded the end of the stimulus. This “haptic feedback” sound allowed subjects to determine whether they reacted to perceived changes in time so that they were registered as detection reports. Thus, subjects knew that trials with haptic feedback were either hits or FAs, because a detection report was registered, while trials without haptic feedback were either misses or CRs, because a detection report was not registered. This feedback allowed subjects to report confidence with full knowledge of the decision that had been registered. The feedback did not indicate the correctness of the decision.

Supplemental instruction and post-training criteria—Before subject data was used for analysis, subjects completed training sessions until reaching performance criteria. Subjects advanced past this training stage after completing a session in which they attained hits in 45% of non-catch trials, avoided FAs on at least 75% of all trials, and had fewer than 1 mean “haptic errors” (high confidence misses with confidence reports occurring within 0.5 s of stimulus end) per block. We established this criterion for identifying haptic errors because if subjects reported confidence this quickly, they would likely have failed to incorporate the haptic feedback sound, or lack thereof. Their high confidence reports would thus be informed only by recognition of the change and not success in responding to it. Any haptic errors that occurred during data collection were not excluded from our analyses, though post-training haptic errors were rare. During training, we occasionally provided subjects with supplemental instruction to allow them to better understand the haptic feedback if they accumulated excess haptic errors. Additionally, to better furnish analyses that required a large sample size of both confidence reports, we suggested to subjects who had low rates of high confidence judgments during training that they choose the “high confidence” option more often when certain of their decisions so that they may earn more

points; subjects still established their own criteria for which trials to assign high confidence, given those supplemental instructions.

Model-based analyses—To further test whether our experimental results could be explained with a mechanism involving a single timescale of evidence evaluation, we used a model-based approach. With maximum likelihood estimation methods, we fit the behavioral choice data with a model that had four free parameters (Figure 4A). The first stage of the model was to convolve the auditory clicks (sensory evidence) with a filter having a half-Gaussian functional form with a free parameter for its standard deviation and the filter defined out to 3 standard deviations. We note that we also used other functional forms including exponential filters, square wave filters, and trapezoid filters with similar conclusions (data not shown). In all cases, the result of the convolution stage delineated the evolution of a decision variable over time. At the second stage of the model, noise was added to the decision variable at each time step with the noise taken from a Gaussian distribution with a free parameter for its standard deviation. Thus, for any given stream of clicks, there was a distribution of possible decision variable values at each point in time. The model prescribes detection reports for any part of this distribution that reaches or exceeds a threshold level set by a decision bound, the third free parameter of the model. To account for attrition due to detection reports, the remaining probability distribution of the decision variable decreased by the probability of bound crossing at each time step. Finally, to account for non-decision sensory and motor processing that adds delays, an additional non-decision time (NDT) was added to the bound crossing time to yield the full reaction time. The non-decision time was taken from a Gaussian distribution with a free parameter for its mean and its standard deviation constrained to be one-fifth of the mean. It has been shown in other tasks that the shape of the non-decision time distribution can be quite variable across tasks/subjects and is not necessarily Gaussian [45]. Our model-based analysis results were robust to departures from the Gaussian non-decision time distribution that altered its skew (data not shown).

For any given set of parameters and sensory input, the model yields a probability distribution for the reaction times. Using this distribution of reaction times and the trial specifications, we calculated the probability of each trial outcome (hit, miss, FA, CR) for every trial performed by our subjects based on the stimulus that was presented. We used brute force grid search to find the values of the four free parameters that maximized the likelihood of the actual trial outcomes for every trial from the combined experimental data of all subjects. Thus, the exact timing of the auditory clicks for every trial was used for the parameter estimations.

To show the best fit behavior from the model, we applied the model with best fit parameters to both experimental stimuli and new stimuli generated in a fashion similar to the experimental stimuli to interpolate at intermediate stimulus strengths (Figure 4B). In particular, additional interpolated stimulus strengths were included at +15 Hz, +20 Hz, +25 Hz, +35 Hz, +40 Hz, +45 Hz to create a smooth psychometric curve. 8000 trials were generated for each of these new stimulus strengths, roughly the same number of trials of each experimental delta click rate. These interpolated values had no bearing on the fitting procedure itself. Psychometric behavioral rates for the model were computed by taking the

mean likelihood for a given trial type: for CR trials, the sum of CR likelihoods for all catch trials divided by the number of all catch trials; for FA trials, the sum of FA likelihoods for all trials divided by the number of all trials; for hits and misses, the sum of each trial type's likelihoods normalized by the number of non-FA trials (i.e. hits and misses) and divided by the number of non-catch trials.

Next, we used model simulations to generate predicted RCs (Figure 4C). This was done by simulating 17,383 trials (matching the total number of experimental trials) with the same stimulus parameters as used for the experiments and with 30% catch trial probability, also matched to the experiments. In these simulations, the model was applied to the trials as before, but with the trial type being classified depending on if and when the model predicted a response. Similar to the model fitting described above, the stimulus was convolved with the filter whose shape was defined by the fitted parameters to compute the mean decision variable. Instantaneous noise in the decision variable was drawn at each time interval from a Gaussian distribution with standard deviation specified by the fitted parameters, with NDT also drawn from a distribution with the fitted parameters in the event of bound crossing. FA RCs were generated similarly as the experimental RCs. FA trials were aligned to the time of response, and stimuli were convolved with a causal half-Gaussian filter. In the case of CR confidence RCs, trials were classified as high or low confidence by thresholding the final decision variable at 49 Hz, with high confidence CRs having a final decision variable lower than 49 Hz. Kernel start points were calculated as they were in analysis of the experimental data. In summary, the simulations used parameters fit to the behavioral choice data to make predictions for the RC analyses. In addition to generating predicted RCs with model fit parameters, we performed simulations using a wider half-Gaussian filter ($\sigma = 0.8$ s) such that the predicted CR confidence difference kernel approximated the experimental start of CR confidence difference kernel (1760.1 ± 73.6 ms vs. 1773.9 ± 43.5 ms). For this analysis, we decreased the bound from 103 Hz to 95.2 Hz in order to maintain FA rates at experimental levels.

We also sought to test the possibility of whether a result similar to the experimental RCs could be recovered by adjustments to the model that still involved a single timescale of evidence evaluation (Figure S4). This model recovery approach involved manipulating the trial-to-trial variability in bound height from 0–144 Hz² and adjusting the bound for each bound variability value to yield the same FA rate as observed experimentally (12.6%). Simulations with each of the 25 model variants produced their own RC kernel start points for FA choice reports and CR confidence reports, and the values of these kernel start points were compared to the experimental kernel start points to determine whether the experimental start points could be recapitulated.

QUANTIFICATION AND STATISTICAL ANALYSIS

Exclusion Criteria—Beyond the 7 subjects analyzed in this study, we excluded 2 subjects from post-training data collection for the inability to adequately detect changes in task stimuli at our criterion rate of 45% in any trial session and 2 subjects for failing to report high confidence for at least 10% of CRs of all trials, making the session data unviable for analysis.

Data analysis—Individual trials were classified as hits, misses, CRs, and FAs. Hit rates were calculated as proportion of hit trials out of trials in which a change occurred (non-catch, non-FA trials). FA rates were calculated as proportion of FA trials out of all trials, and CR rates were calculated as proportion of CR trials out of all catch trials (Figure 2A). Because there were two confidence ratings available (high and low), average confidence for each response type and stimulus condition (click rate) was calculated as proportion of high confidence reports for each response/stimulus combination (Figure 2B). Rates were calculated for both the combined data, which included every trial from each individual (Figure 2), and for each individual subject (Figure S1).

FA and CR RCs (Figure 3) were generated by smoothing click times with a causal half-Gaussian filter having a standard deviation of 0.05 s and sampling every 0.01 s. Note that trials differed in duration. Rather than discard trials with shorter durations, each time bin represents a mean over a different number of trials with shorter duration trials not contributing to earlier time points. For the confidence-based kernels (Figure 3B–C), data were first separated into sets of low and high confidence trials, and these individual data sets were each convolved with the half-Gaussian filter. The detection report kernel (Figure 3A) was created by convolving the click times of all FA trials, regardless of confidence. We calculated difference plots for FA RCs (Figure 3B, green line) by subtracting the mean low confidence kernel from the mean high confidence kernel at each time point. RCs included all trials of the associated trial type (e.g., FA RCs included every FA trial recorded).

The start point and slope of the detection report kernel's ascending phase were then quantified (Figure S2). For each subject's FA kernels, as well as FA kernels for the combined data, we fit the early phase of the FA kernel, from 3 s before stimulus end to the peak value of the kernel, with a 2-piece linear function with three free parameters using MATLAB's *fit* function. These parameters provided estimates for the average baseline click rate leading up to the FA, the average start time of the kernel, and the slope of the function from the start time to the peak of the kernel. The start time was the time point at which the function first diverged from the baseline.

Start point and slope of CR confidence difference kernels were estimated similarly (Figure S2), by fitting the last 5 s of the CR confidence difference kernel to a 2-piece linear function with two free parameters for start point and slope. The start point was the first time point that diverged from a 0 Hz differential click rate, and the slope was the slope of the line connecting this start point to the stimulus end.

To estimate kernel endpoint for FAs (Figure S3), we convolved the click times with a square-wave function encompassing the descending phase of the FA kernel. This allowed for a more precise estimate of the detection kernel endpoint that minimized the influence of clicks on later times in the kernel compared to using a causal half-Gaussian filter. Starting from 75% of the peak FA choice kernel, a 5 ms bin was moved by 1 ms toward the end of the stimulus, providing a 5 ms wide average click rate for each ms time point. This kernel was then fit to a 2-piece linear function as before, this time with two free parameters: slope and kernel endpoint. Because the subject's choice of whether to respond would no longer be influenced by the stimulus after the kernel endpoint, the mean click rate should return to the generative

click rate of 50 Hz. Therefore, the rate after the endpoint was set to a fixed constant of 50 Hz. The average excess click rate for high and low confidence trials was calculated by subtracting the mean click rate for high and low confidence trials, respectively, from 50 Hz.

DATA AND SOFTWARE AVAILABILITY

The data that support the findings of this study and the analysis code are available from the Lead Contact upon request.

ADDITIONAL RESOURCES

N/A

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

We thank Mark Goldman for helpful discussions and Shuqi Wang for helping to improve the efficiency of model-based analyses. This research was supported by a National Alliance for Research on Schizophrenia and Depression Young Investigator Award to T.D.H. and a Whitehall Foundation Grant.

References

- Ossmy O, Moran R, Pfeffer T, Tsetsos K, Usher M, and Donner TH (2013). The timescale of perceptual evidence integration can be adapted to the environment. *Curr. Biol* 23, 981–986. [PubMed: 23684972]
- Glaze CM, Kable JW, and Gold JI (2015). Normative evidence accumulation in unpredictable environments. *Elife* 4 Available at: 10.7554/eLife.08825.
- Sugrue LP, Corrado GS, and Newsome WT (2004). Matching behavior and the representation of value in the parietal cortex. *Science* 304, 1782–1787. [PubMed: 15205529]
- Piet AT, El Hady A, and Brody CD (2018). Rats adopt the optimal timescale for evidence integration in a dynamic environment. *Nat. Commun* 9, 4265. [PubMed: 30323280]
- Gold JI, and Stocker AA (2017). Visual Decision-Making in an Uncertain and Dynamic World. *Annu Rev Vis Sci* 3, 227–250. [PubMed: 28715956]
- Fetsch CR, Kiani R, and Shadlen MN (2014). Predicting the Accuracy of a Decision: A Neural Mechanism of Confidence. *Cold Spring Harb. Symp. Quant. Biol* 79, 185–197. [PubMed: 25922477]
- Hangya B, Sanders JI, and Kepecs A (2016). A Mathematical Framework for Statistical Decision Confidence. *Neural Comput.* 28, 1840–1858. [PubMed: 27391683]
- Sanders JI, Hangya B, and Kepecs A (2016). Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron* 90, 499–506. [PubMed: 27151640]
- Pouget A, Drugowitsch J, and Kepecs A (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nat. Neurosci* 19, 366–374. [PubMed: 26906503]
- Hanks TD, and Summerfield C (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron* 93, 15–31. [PubMed: 28056343]
- Yeung N, and Summerfield C (2012). Metacognition in human decision-making: confidence and error monitoring. *Philos. Trans. R. Soc. Lond. B Biol. Sci* 367, 1310–1321. [PubMed: 22492749]
- De Martino B, Fleming SM, Garrett N, and Dolan RJ (2012). Confidence in value-based choice. *Nat. Neurosci* 16, 105. [PubMed: 23222911]
- Insabato A, Pannunzi M, Rolls ET, and Deco G (2010). Confidence-related decision making. *J. Neurophysiol* 104, 539–547. [PubMed: 20393062]

14. Drugowitsch J, Moreno-Bote R, Churchland AK, Shadlen MN, and Pouget A (2012). The cost of accumulating evidence in perceptual decision making. *J. Neurosci* 32, 3612–3628. [PubMed: 22423085]
15. Zylberberg A, Barttfeld P, and Sigman M (2012). The construction of confidence in a perceptual decision. *Front. Integr. Neurosci* 6, 79. [PubMed: 23049504]
16. Fetsch CR, Kiani R, Newsome WT, and Shadlen MN (2014). Effects of cortical microstimulation on confidence in a perceptual decision. *Neuron* 83, 797–804. [PubMed: 25123306]
17. Kiani R, Corthell L, and Shadlen MN (2014). Choice certainty is informed by both evidence and decision time. *Neuron* 84, 1329–1342. [PubMed: 25521381]
18. Gherman S, and Philiastides MG (2015). Neural representations of confidence emerge from the process of decision formation during perceptual choices. *Neuroimage* 106, 134–143. [PubMed: 25463461]
19. Moran R, Teodorescu AR, and Usher M (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cogn. Psychol* 78, 99–147. [PubMed: 25868113]
20. Pleskac TJ, and Busemeyer JR (2010). Two-stage dynamic signal detection: a theory of choice, decision time, and confidence. *Psychol. Rev* 117, 864–901. [PubMed: 20658856]
21. Fleming SM, and Daw ND (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review* 124, 91–114. Available at: 10.1037/rev0000045. [PubMed: 28004960]
22. Peters MAK, Thesen T, Ko YD, Maniscalco B, Carlson C, Davidson M, Doyle W, Kuzniecky R, Devinsky O, Halgren E, et al. (2017). Perceptual confidence neglects decision-incongruent evidence in the brain. *Nat Hum Behav* 1 Available at: 10.1038/s41562-017-0139.
23. Churchland AK, and Kiani R (2016). Three challenges for connecting model to mechanism in decision-making. *Curr Opin Behav Sci* 11, 74–80. [PubMed: 27403450]
24. Johnson B, Verma R, Sun M, and Hanks TD (2017). Characterization of decision commitment rule alterations during an auditory change detection task. *J. Neurophysiol* 118, 2526–2536. [PubMed: 28794191]
25. Kiani R, Hanks TD, and Shadlen MN (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *J. Neurosci* 28, 3017–3029. [PubMed: 18354005]
26. van den Berg R, Anandalingam K, Zylberberg A, Kiani R, Shadlen MN, and Wolpert DM (2016). A common mechanism underlies changes of mind about decisions and confidence. *Elife* 5, e12192. [PubMed: 26829590]
27. Okazawa G, Sha L, Purcell BA, and Kiani R (2018). Psychophysical reverse correlation reflects both sensory and decision-making processes. *Nat. Commun* 9, 3479. [PubMed: 30154467]
28. Radillo AE, Veliz-Cuba A, Josi K, and Kilpatrick ZP (2017). Evidence Accumulation and Change Rate Inference in Dynamic Environments. *Neural Comput.* 29, 1561–1610. [PubMed: 28333591]
29. Hanks T, Kiani R, and Shadlen MN (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *Elife* 3, 1–17.
30. Purcell BA, and Kiani R (2016). Neural Mechanisms of Post-error Adjustments of Decision Policy in Parietal Cortex. *Neuron* 89, 658–671. [PubMed: 26804992]
31. Churchland AK, Kiani R, and Shadlen MN (2008). Decision-making with multiple alternatives. *Nat. Neurosci* 11, 693–702. [PubMed: 18488024]
32. Merten K, and Nieder A (2012). Active encoding of decisions about stimulus absence in primate prefrontal cortex neurons. *Proc. Natl. Acad. Sci. U. S. A* 109, 6289–6294. [PubMed: 22460793]
33. Deco G, Pérez-Sanagustín M, de Lafuente V, and Romo R (2007). Perceptual detection as a dynamical bistability phenomenon: a neurocomputational correlate of sensation. *Proc. Natl. Acad. Sci. U. S. A* 104, 20073–20077. [PubMed: 18077434]
34. Seung HS, Lee DD, Reis BY, and Tank DW (2000). Stability of the memory of eye position in a recurrent network of conductance-based model neurons. *Neuron* 26, 259–271. [PubMed: 10798409]
35. Goldman MS, Compte A, and Wang XJ (2009). Neural integrator models. *Encyclopedia of neuroscience* 6, 165–178.

36. Purcell BA, Schall JD, Logan GD, and Palmeri TJ (2012). From salience to saccades: multiple-alternative gated stochastic accumulator model of visual search. *J. Neurosci* 32, 3433–3446. [PubMed: 22399766]
37. Teichert T, Grinband J, and Ferrera V (2016). The importance of decision onset. *J. Neurophysiol* 115, 643–661. [PubMed: 26609111]
38. Goldman MS (2009). Memory without feedback in a neural network. *Neuron* 61, 621–634. [PubMed: 19249281]
39. Murray JD, Bernacchia A, Freedman DJ, Romo R, Wallis JD, Cai X, Padoa-Schioppa C, Pasternak T, Seo H, Lee D, et al. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nat. Neurosci* 17, 1661–1663. [PubMed: 25383900]
40. 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. [PubMed: 21317906]
41. Scott BB, Constantinople CM, Akrami A, Hanks TD, Brody CD, and Tank DW (2017). Frontoparietal Cortical Circuits Encode Accumulated Evidence with a Diversity of Timescales. *Neuron* 95, 385–398.e5. [PubMed: 28669543]
42. Boldt A, and Yeung N (2015). Shared Neural Markers of Decision Confidence and Error Detection. *Journal of Neuroscience* 35, 3478–3484. [PubMed: 25716847]
43. Kepecs A, and Mainen ZF (2012). A computational framework for the study of confidence in humans and animals. *Philos. Trans. R. Soc. Lond. B Biol. Sci* 367, 1322–1337. [PubMed: 22492750]
44. Sanders JI, and Kepecs A (2014). A low-cost programmable pulse generator for physiology and behavior. *Front. Neuroeng* 7, 43. [PubMed: 25566051]
45. Verdonck S, and Tuerlinckx F (2016). Factoring out nondecision time in choice reaction time data: Theory and implications. *Psychological Review* 123, 208–218. Available at: 10.1037/rev0000019. [PubMed: 26641558]




Figure 1. Auditory change detection task with confidence report showing the sequence of events for each trial:

A trial began when the center port was illuminated, cuing subjects to insert their finger. Once the finger was inserted, a stream of auditory clicks began to play. The clicks were generated by a Poisson process with a generative baseline click rate of 50 Hz. 30% of trials were catch trials with no change in the generative click rate. In the other 70% of trials, the generative click rate increased by 10, 30, or 50 Hz at a random point (red arrow in example). Subjects had to withdraw their finger within 800 ms of the change for the trial to qualify as a hit on change trials, or withhold a response for the trial to qualify as a CR on catch trials. At the end of the trial (after response or stimulus end), the two peripheral ports illuminated, cuing subjects to indicate confidence in their decision: engaging the left port reported low confidence while engaging the right reported high confidence. Immediately following the confidence report, feedback was given via an auditory tone to indicate success or failure on that trial.

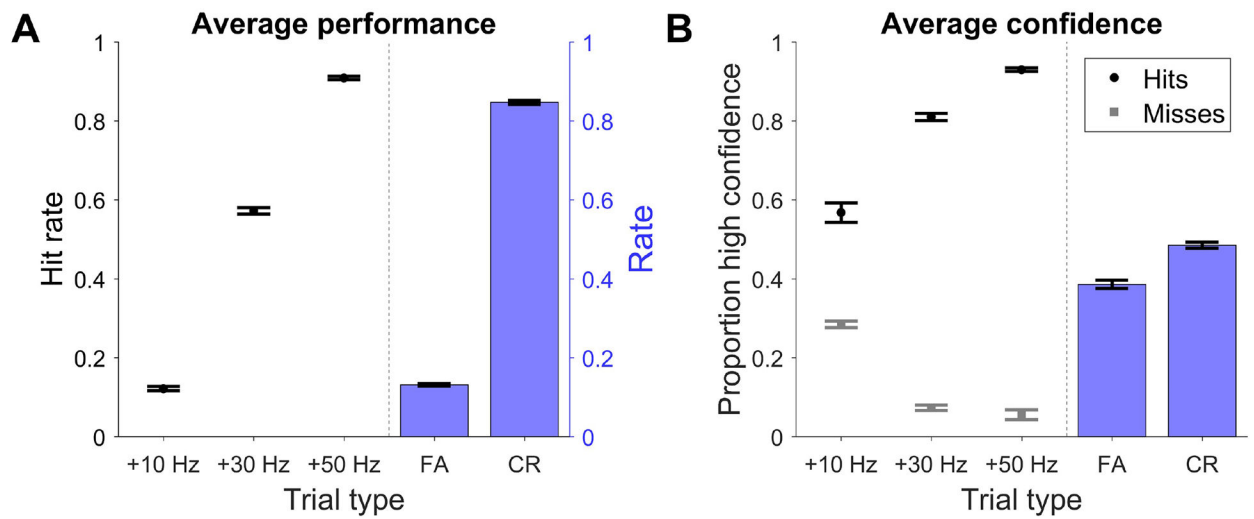


Figure 2. Task performance and confidence ratings:

A, Combined data from 7 subjects showing performance as a function of the change in click rate (left) and proportion of FA and CR trials (right). Hit rate was calculated excluding FA trials. FA rate was calculated from all trial types. CR rate was calculated from trials in which no change in generative click rate occurred (30% of trials). **B**, Proportion of high confidence hits (black) and misses (gray) as a function of change in click rate (left) and proportion of high confidence FA and CRs (right). Error bars indicate \pm SEM. See Figure S1.




Figure 3. FA and CR reverse correlations:

Combined data from 7 subjects was used to calculate average click rate over time for each outcome. RC traces were constructed by convolving click times preceding outcomes with causal half-Gaussian filters ($\sigma = 0.05$ s). **A**, Detection report kernel. RC trace (black line) is comprised of all FA trials, showing the average click rate preceding FAs. The start of the detection report kernel (arrow) was estimated by fitting the ascending phase of the kernel with a 2-piece linear function with 3 free parameters: the baseline click rate (left of arrow), the slope of the kernel's ascending phase (right of arrow), and the start of the ascending phase (arrow). Horizontal dotted line denotes the 50 Hz baseline generative click rate. **B**, Detection report confidence kernels. RC traces show the average click rate preceding high (red) and low (blue) confidence FAs, as well as the difference in click rate between the two confidence kernels (high - low; green). Same conventions as in **A**. The shaded portion of the graph near time 0 shows the temporal interval analyzed in Figure S3A. **C**, CR confidence kernels. As in **B**, but showing RC traces preceding CRs. The divergence point between the two confidence kernels (arrow) was estimated by fitting the difference between the two kernels (high - low; CR confidence difference kernel) with a 2-piece linear function with 2 free parameters (see Figure S2, column 4 for fits): the divergence point from a baseline difference of 0 Hz (arrow) and the slope from that point onward. For all kernels, shaded region shows \pm SEM. N = number of trials for each trace. See Figure S2 and S3.

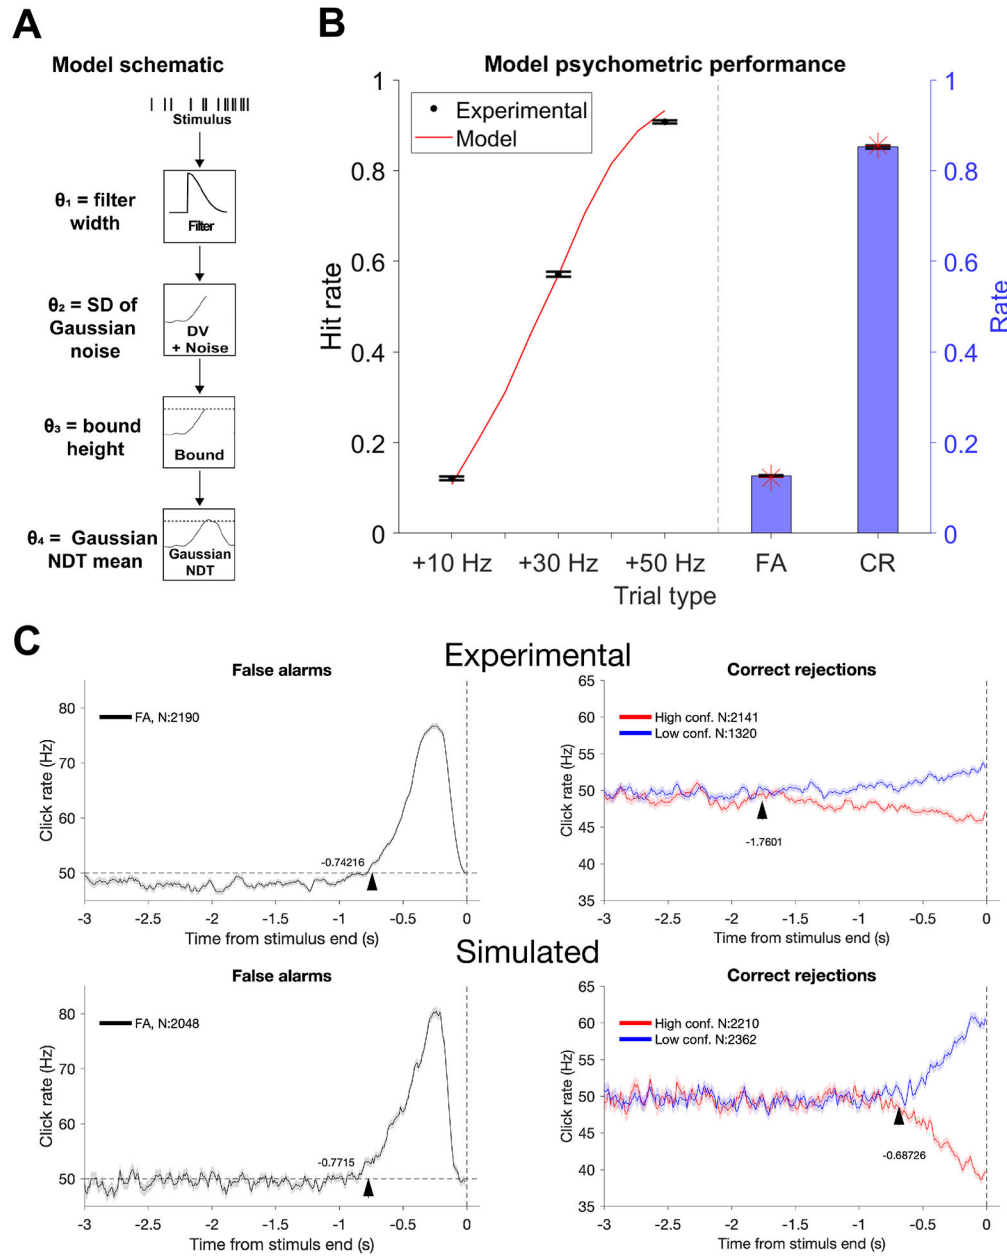


Figure 4. Model simulation of behavioral data:

A, Model schematic. The model included four free parameters: width of half-Gaussian filter (0.96 s), decision variable noise term standard deviation (13 Hz), decision bound (103 Hz), and mean Gaussian NDT (0.17 s). Best fit parameters are in parentheses. **B**, Model comparison with psychometric data based on best fit parameters. Error bars indicate \pm SEM (as in Fig. 2). The red line indicates model hit rate as a function of change in click rate, with additional interpolated delta click rates (+15 Hz, +20 Hz, +25 Hz, +35 Hz, +40 Hz, +45 Hz). Red stars indicate FA and CR rates for the model. **C**, Experimental and simulated RCs. See Figure S4.