Dynamic expressions of confidence within an evidence accumulation framework

Kobe Desender[1,2], Tobias H. Donner[1], & Tom Verguts[2]

1. **Section Computational Cognitive Neuroscience, Department of Neurophysiology and Pathophysiology, University Medical Center Hamburg-Eppendorf, Germany**
2. **Department of Experimental Psychology, Ghent University, Belgium**

Corresponding author:

Dr. Kobe Desender
Department of Neurophysiology and Pathophysiology
University Medical Center Hamburg-Eppendorf
Martinistraße 52, 20251 Hamburg
Germany
E-mail: Kobe.Desender@gmail.com

Short title: Dynamic expressions of confidence

Abstract

Human observers can reliably report their confidence in the choices they make. An influential framework conceptualizes decision confidence as the probability of a decision being correct, given the choice made and the evidence on which it was based. This framework accounts for three diagnostic signatures of human confidence reports, including an opposite dependence of confidence on evidence strength for correct and error trials. However, the framework does not account for the temporal evolution of these signatures, because it only describes the transformation of a static evidence representation into choice and the associated confidence. Here, we combine this framework with another influential framework: the temporal accumulation of evidence towards decision bounds. We propose that confidence at any point in time reflects the probability of being correct, given the choice and accumulated evidence *up until that point*. This model predicts a systematic time-dependence of all diagnostic signatures of decision confidence, most critically: an increase of the opposite dependence of confidence on evidence strength and choice correctness with time. We tested, and confirmed, these predictions in human subjects performing a random dot motion discrimination task, in which confidence judgments were queried at different points in time. We conclude that confidence reports track the temporal evolution of the probability of being correct.

Author summary

Humans are able to report a sense of confidence in the accuracy of a choice. An influential framework states that confidence reflects the probability that a choice is correct. We combined human experimenting with computational modelling and extended this notion in the time domain, thus to formally describe the dynamics of confidence. Both human data and our model show that the sense of confidence depends on the point in time, at which it is queried. We conclude that human confidence reports reflect the dynamics of the probability of a choice being correct.

Introduction

Human observers can reliable judge the confidence about their choices. They often report high confidence for correct choices and low confidence for errors. Accurate internal representations of confidence are useful for the adaptive control of future behaviour [1–3]. An influential framework posits that internal representations of decision confidence, and agents' overt reports thereof, reflect the probability of being correct, given the choice made and given the evidence on which it was based [4–6]. In this framework, both choice and confidence are directly based on the same underlying computations. One advantage of this approach is that it predicts three qualitative signatures of confidence [5]: (i) an interaction between evidence strength and choice accuracy, whereby confidence increases with evidence strength for correct choices, but decreases for incorrect choices; (ii): confidence predicts a monotonic increase in accuracy; (iii): a steeper psychometric performance for high versus low confidence trials. These three signatures have been observed in neural data [5], several implicit behavioural measures of confidence [4,5,7,8], and explicit confidence reports of human observers [4,9]. While this framework is highly influential, an important limitation is that it is static: a fixed quantity of evidence determines both the choice and associated confidence. Therefore, this framework does not account for the dynamics of decision-making, the associated trade-off between speed and accuracy, and their effect on confidence reports.

Another influential framework, bounded accumulation, holds that perceptual decisions are based on the temporal accumulation of noisy sensory evidence towards decision bounds [10,11]. In two-choice tasks, a decision maker accumulates evidence for each option, and the option for which the integrated evidence first crosses a decision threshold is selected, indicating commitment to choice [11]. The efficiency (i.e., signal-to-noise ratio) of the accumulation process is governed by the so-called drift rate.

Here, we extend the framework of statistical confidence into the time domain, by connecting it to the framework of evidence accumulation towards decision bounds, and show that this dynamic framework of statistical decision confidence accounts well for human behavior in a classic perceptual choice task widely used in the study of decision-making.

Results

Dynamic statistical decision confidence

We propose that confidence at any point in time reflects the probability of being correct, given the choice made and the evidence accumulated *up until that point*. We first unpack and solidify this idea through simulations of the drift diffusion model (DDM), a popular evidence accumulation model (Fig. 1A). The key insight, supported by recent data, is that the evidence accumulation does not necessarily terminate at the time of bound crossing: evidence can continue to accumulate following the response [12,13]. Therefore, confidence reports may differ, depending on whether they are probed around the time of the response [14,15] or only later in time, after additional post-decision processing [16–18]. Even so, in both cases, confidence reflects the probability of being correct, given the choice and accumulated evidence up until that point. The heat map in Fig. 1A reflects the probability of being correct given evidence (Y-axis) and time (X-axis), conditional on the choice made. Note that the heat map is flipped vertically when the lower boundary is reached instead. Thus, confidence in our model reflects the probability of being correct, given choice, evidence and time. Most importantly, with respect to Signature 1 (an interaction between evidence strength and choice accuracy), our model simulations show that confidence increases for both corrects and errors when confidence is quantified at the time the bound is reached (Figure 1B), whereas the interaction between evidence strength and choice accuracy emerges when confidence is queried later in time (Figure 1C-D).

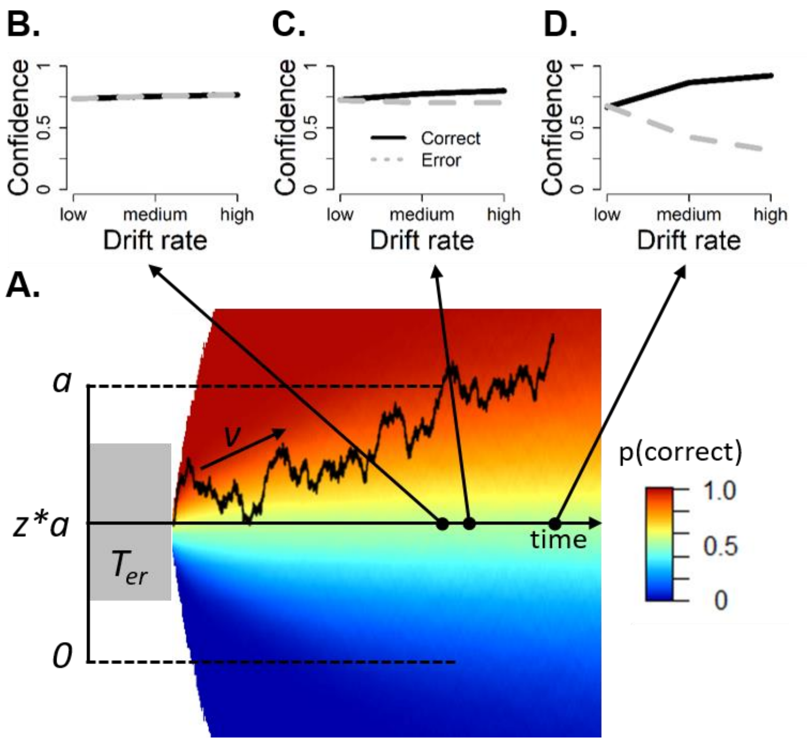

Figure 1. Quantifying decision confidence within an evidence accumulation framework. A. *Noisy sensory evidence is accumulated over time, until the decision variable reaches one of two bounds (a or 0), corresponding to a left or right choice, respectively. After the decision variable reaches a bound, evidence continues to accumulate. The heat map shows the probability of being correct given time and evidence, conditional on the (left) choice made. Confidence is quantified as the probability of the choice being correct, given elapsed time and the integrated evidence (i.e., represented by the color of the heat map). Confidence can be queried at different points in time.* ***B-D****. Model predictions about signature 1, an interaction between evidence strength and accuracy, depending on when in time confidence is quantified.*

We next tested the model predictions in the behavior of human participants during the widely used random dot motion discrimination task, in which we prompted confidence ratings at different latencies. We first show that behavioral performance was well explained by the drift diffusion model. Second, we tested and confirmed dynamic predictions about these three diagnostic statistical signatures of confidence. Third, using a manipulation of evidence volatility, we shed light on the stopping rule used for post-decisional processing.

Explaining behavior through bounded evidence accumulation

Twenty-six human participants viewed random dot motion stimuli and decided as fast and accurately as possible whether a subset of dots was coherently moving towards the left or the right side of the screen [10]. The difficulty of these decisions was manipulated by varying the proportion of coherently moving dots. Five different levels of coherence were used, ranging from 0 up to .4, all of which were randomly intermixed within a block. We also manipulated the volatility of motion coherence over the course of a single trial. Specifically, on each frame, the input coherence was either sampled from a Gaussian distribution with *SD* = 0 (low volatility), or from a Gaussian distribution with *SD* = .256 (high volatility) around the generative coherence. In the high volatility condition, additional noise is thus introduced in the decision process, which previous work has shown to speed up RTs and increase confidence [14]. Depending on the block that participants were in, responses were collected in a different way (see Figure 2). In the *immediate condition* participants jointly indicated their choice (left or right) and their level of confidence (guess correct, probably correct or certainly correct) via a single response. In the *delayed condition*, participants first indicated their choice (left or right), and then after a 1s blank screen or 1s of continued motion (same coherence, volatility and motion direction as the initial stimulus) they indicated the level of confidence in their choice on a six-point scale.

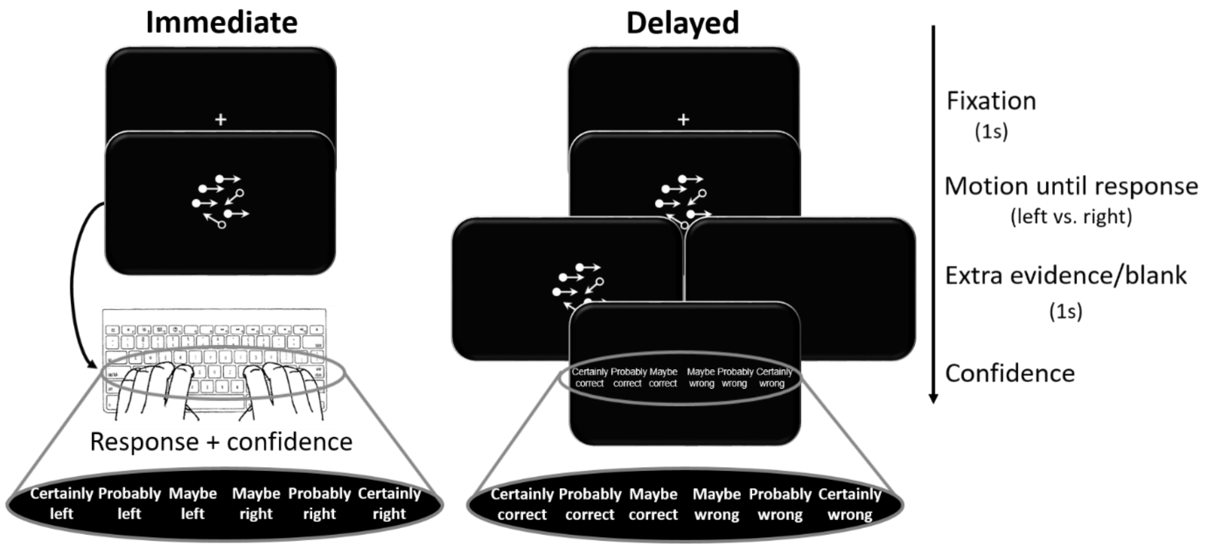

Figure 2. Experimental task. *Sequence of events in the experimental task. Participants decided, as fast and accurately as possible, whether the majority of dots were moving left or right. In the immediate condition, they did so by jointly indicating their choice (left or right) and confidence (sure correct, probably correct or guess correct) in a single response. In the delayed condition, participants first indicated their choice with their thumbs (left or right), and after a 1s blank or 1s of continued motion, they were prompted to indicate the degree of confidence in their decision using a six-point confidence scale (ranging from certainly correct to certainly wrong).*

To unravel how coherence and volatility affected latent cognitive variables in the decision process, we fitted choices and reaction times using a hierarchical version of the drift diffusion framework [19]. Because the effects of coherence and volatility were not modulated by the timing of confidence reports (immediate vs delayed) for both RTs, $F < 1$, *Bayes Factor (BF)* = .008, and accuracy, $F < 1$, $BF = .01$, the RT and accuracy data were combined. First, as typically observed in random dot motion tasks, drift rates increased monotonically with coherence level (see Figure 3A), with significant differences in drift rate between all coherence levels (averaged across volatility levels), *p*s < .001. Estimated drift rates did not depend on the level of evidence volatility, *p*s > .119. Second, as we predicted [14], our manipulation of within-trial evidence volatility was captured by the within-trial drift variability parameter σ (see Figure 3B; Methods). When averaged over different coherences, estimated within-trial variability was higher for high compared to low volatility, $p = .014$ (pair-wise comparisons within each coherence value: 0% coherence: p = .091; 5% coherence: p = .049; 10% coherence: p = .259; 20% coherence: p = .106; 40% coherence: p = .457).

These model fits captured key qualitative patterns evident in the behavioral data (Figure 3C-D). Accuracy increased with the level of coherence (data: $F(4,22) = 267.48$, $p < .001$; model: $F(4,22) = 619.57$, $p < .001$), whereas evidence volatility and the interaction between both variables left accuracy unaffected (data: Fs < 1; model: ps > .213). Reaction times decreased with increasing coherence levels (data: $F(4,22) = 30.68$, $p < .001$; model: $F(4,22) = 52.25$, $p < .001$), and were shorter with high compared to low volatility (data: $F(1,25) = 9.10$, $p = .006$; model: $F(1,25) = 17.91$, $p < .001$), an effect that was mostly pronounced at low coherence levels (data: $F(4,22) = 13.21$, $p < .001$; model: $F(4,22) = 15.53$, $p < .001$).

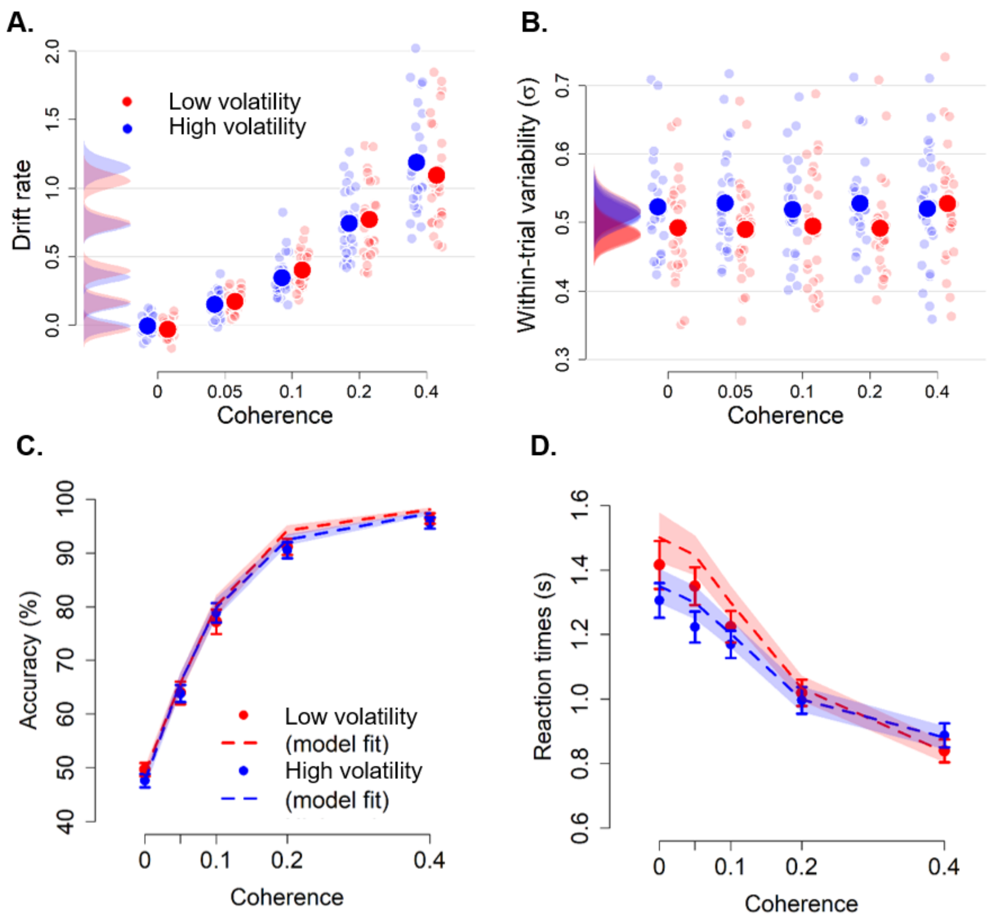

Figure 3. Model fits and task performance. A. *Drift rate scales monotonically with the proportion of coherently moving dots, but did not differ for high and low volatility conditions.* ***B.*** *Within-trial variability (σ) selectively varied as a function of evidence volatility, whereas it was unaffected by motion coherence. Large dots: group averages; small dots: individual participants. Distributions show the group posteriors. Statistical significance is reflected in overlap between posterior distributions over parameter estimates (Materials and Methods).* ***C-D.*** *Accuracy (C) and RTs (D) as a function of coherence and evidence volatility, separately for the empirical data (points and bars) and model fits (lines and shades). Shades and error bars reflect SEM of model and data, respectively.*

Post-decision accumulation explains dynamic signatures of statistical confidence

Next, we used our model fits to obtain qualitative and quantitative predictions of confidence reports about the three dynamic signatures of confidence. In order to create a heat map reflecting the probability of being correct, we simulated a large number of trials and calculated average accuracy for each combination of time and evidence. Confidence predictions were quantified by reading out the values from this heat map (reflecting the probability of being correct) for each combination of evidence, time, and choice.

Signature 1: interaction between evidence strength and choice accuracy. The first diagnostic signature of statistical confidence established previously [5] is an increase of confidence with evidence strength for correct trials, but a decrease for error trials. In the immediate condition, confidence increased with coherence level, $F(4,44.81) = 15.62$, $p < .001$. Crucially, there was also the predicted interaction between coherence level and choice accuracy, $F(4,1990.70) = 14.09$, $p < .001$. Confidence increased with evidence strength for correct trials (linear contrast: $p < .001$), but there was no significant effect for error trials (linear contrast: $p = .070$; see Figure 4A). In contrast, as visualized in Figure 1B, the model predicts that when confidence is quantified at the time when the decision boundary is reached, confidence scales with coherence, $F(4,25) = 14.14$, $p < .001$, but there is no interaction between coherence and choice accuracy, $F(4,125.02) = 1.03$, $p = .39$.

The above mismatch can easily be remedied by assuming that choice and confidence cannot be simultaneously computed, or accessed for report – for example due to the psychological 'refractory period' [20,21]. Indeed, when confidence was calculated with a small temporal delay (100 ms, Figure 1C; see Methods), the model did predict the interaction between coherence and choice accuracy, $F(4,200) = 84.05$, $p < .001$. As in the behavioral data, the model with the small temporal delay predicted increasing confidence with coherence for correct trials (linear contrast: $p < .001$), but not for error trials (linear contrast: $p = .541$; Figure 4A). In the remainder, we will continue with predictions from the model with temporal delay.

In both delayed conditions, confidence scaled with coherence level (blank condition: $F(4,51.8) = 5.49$, $p < .001$; extra evidence condition: $F(4,4571.1) = 4.75$, $p < .001$). In both conditions, there was also an interaction between coherence and choice accuracy (blank condition: $F(4,3625.6) = 53.38$, $p < .001$; extra evidence condition: $F(4,4568.7) = 71.45$, $p < .001$). Within the correct trials, confidence increased with coherence levels (blank and extra evidence conditions, linear contrasts: $p < .001$. Instead, within the error trials, confidence decreased as a function of coherence (blank and extra evidence conditions, linear contrasts: $p = .001$). This interaction was captured by a model which terminated post-decision accumulation after a fixed amount of time (cf. Figure 1D; Materials and Methods). This model also showed the scaling of confidence with coherence ($F(4,69.79) = 39.9$, $p < .001$), as well as the interaction with choice accuracy ($F(4,225) = 1634.3$, $p < .001$). Similar to the human data, confidence increased with coherence for correct trials (linear contrast: $p < .001$) and decreased for error trials (linear contrast: $p < .001$; figure 4B). Finally, there was a three-way interaction between coherence, choice accuracy and interrogation condition (data: $F(8,13466.5) = 18.22$, $p < .001$, model: $F(4,475) = 161.54$, $< .001$).

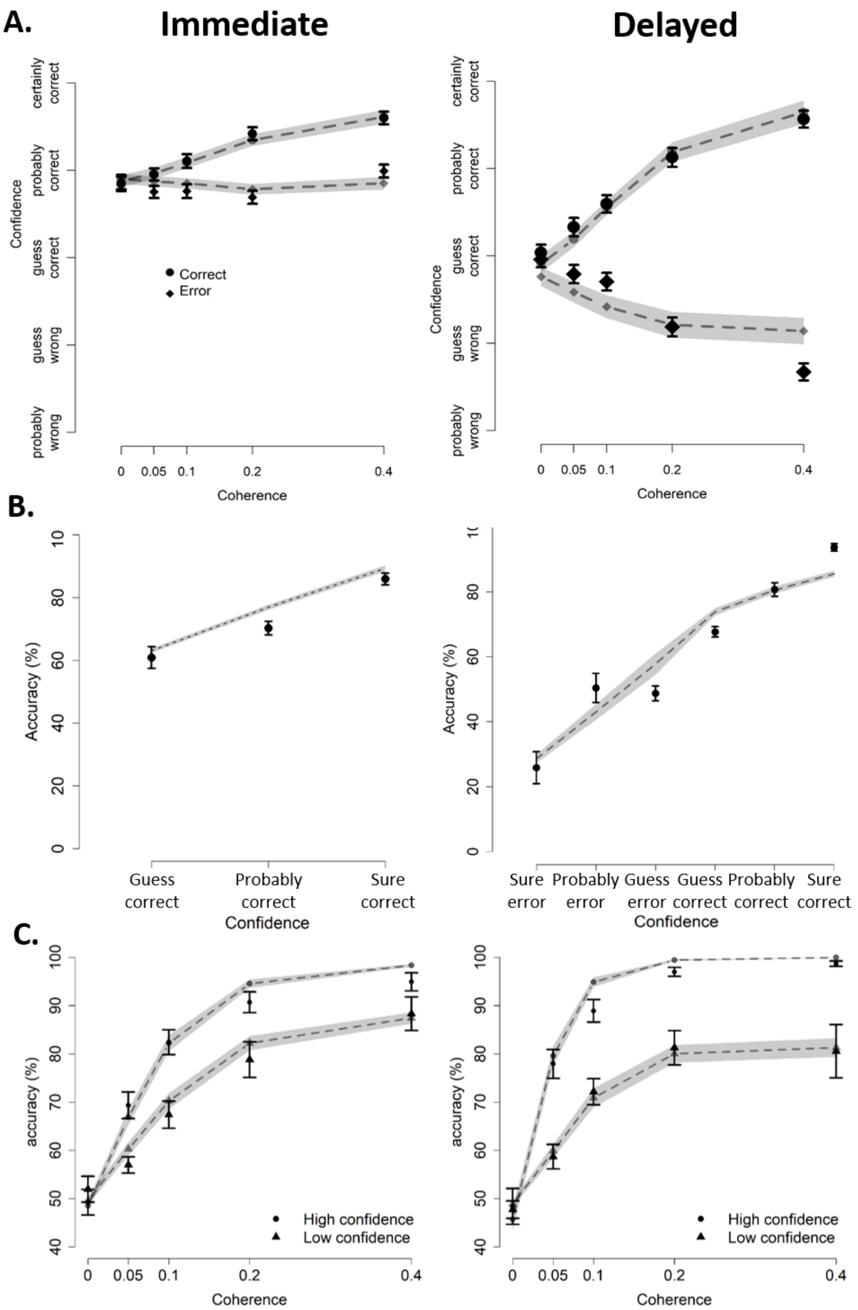

215

216 ***Figure 4. Three dynamic signatures of statistical confidence. A. Signature 1: an interaction***
217 ***between evidence strength and choice accuracy.*** *When confidence is quantified shortly after the*
218 *decision bound has been reached ("immediate"), both model and data show an interaction between*
219 *evidence strength and choice accuracy in the immediate condition. The same pattern was observed*
220 *for the delayed condition, although the interaction effect was clearly much stronger here.* ***B. Signature***
221 ***2: monotonically increasing accuracy as a function of confidence.*** *Both model and data show a*
222 *monotonic scaling of accuracy depending on the level of confidence.* ***C. Signature 3: Steeper***
223 ***psychometric performance for high versus low confidence.*** *Both model and data show a steeper*

psychometric performance for trials judged with high versus low confidence. Notes: data for the delayed conditions are averaged over blank and extra evidence conditions. All plots show empirical data (black points and bars) and model predictions (grey lines and shades). Shades and error bars reflect SEM of model and data, respectively.

Signature 2: Monotonically increasing accuracy as a function of confidence. The second signature of statistical confidence is that it monotonically predicts choice accuracy. Indeed, an approximately linear relation between confidence and mean accuracy was observed in the data for both the immediate condition, $b = .13$, $t(29.92) = 12.82$, $p < .001$, the delayed blank, $b = .12$, $t(27.29) = 16.12$, $p < .001$, and the delayed extra evidence condition, $b = .13$, $t(23.33) = 15.45$, $p < .001$. This pattern was also captured by the model in the immediate condition, $b = .12$, $t(26.4) = 11.76$, $p < .001$, and in the delayed condition, $b = .12$, $t(26) = 25.05$, $p < .001$ (see Figure 4B). Note that these slopes did not differ depending on the moment in time when confidence was queried (data: $X^2 = 2.03$, $p = .363$; model: $X^2 = 3.76$, $p = .152$).

Signature 3: Steeper psychometric performance for high versus low confidence. The third signature of statistical confidence is that the relation between accuracy and evidence strength should be steeper for trials judged with high versus low confidence. The model predicts that this difference should be larger for the delayed compared to the immediate condition (Figure 4C). To test this prediction, confidence reports were divided into high or low using a split-median, separately per participant. As expected, the interaction between coherence and confidence in predicting accuracy was observed both in the immediate condition (data: $X^2(4) = 30.9$, $p < .001$; model: $X^2(4) = 2212.4$, $p < .001$), and in the delayed condition (data: delayed blank: $X^2(4) = 84.15$, $p < .001$, extra evidence: $X^2(4) = 56.64$, $p < .001$; model: $X^2(4) = 9018.7$, $p < .001$; see Figure 4C). Finally, there was a significant three-way interaction between coherence, confidence and interrogation condition (data: $X^2(8) = 24.51$, $p = .002$; model: $X^2(4) = 228.90$, $p < .001$).

Evidence volatility dissociates time-based and evidence-based stopping criteria

If decision confidence is 'read out' and reported after additional post-decision processing, a stopping rule has to be implemented that determines when confidence is evaluated. In the previous simulations, following previous research a so-called 'time-based stopping rule' was implemented [16,17]: confidence was extracted after a fixed latency following initial bound crossing. An alternative implementation, however, is that the stopping rule for confidence reports is also based on accumulated evidence, just like the stopping rule for the first-order decision process [18]. According to this 'evidence-based stopping rule', after reaching the initial choice threshold, agents impose a second threshold and a delayed confidence report is given when this second threshold is reached. Because the statistical signatures discussed before do not arbitrate between the two delayed confidence stopping criteria (see Supplementary Materials), we next turn towards our manipulation of evidence volatility. Previous work has shown that an evidence-based model can explain the volatility effect on confidence for immediate confidence judgments [14]. We reasoned that the same manipulation could be used to disentangle a time-based versus an evidence-based stopping rule for delayed confidence judgments.

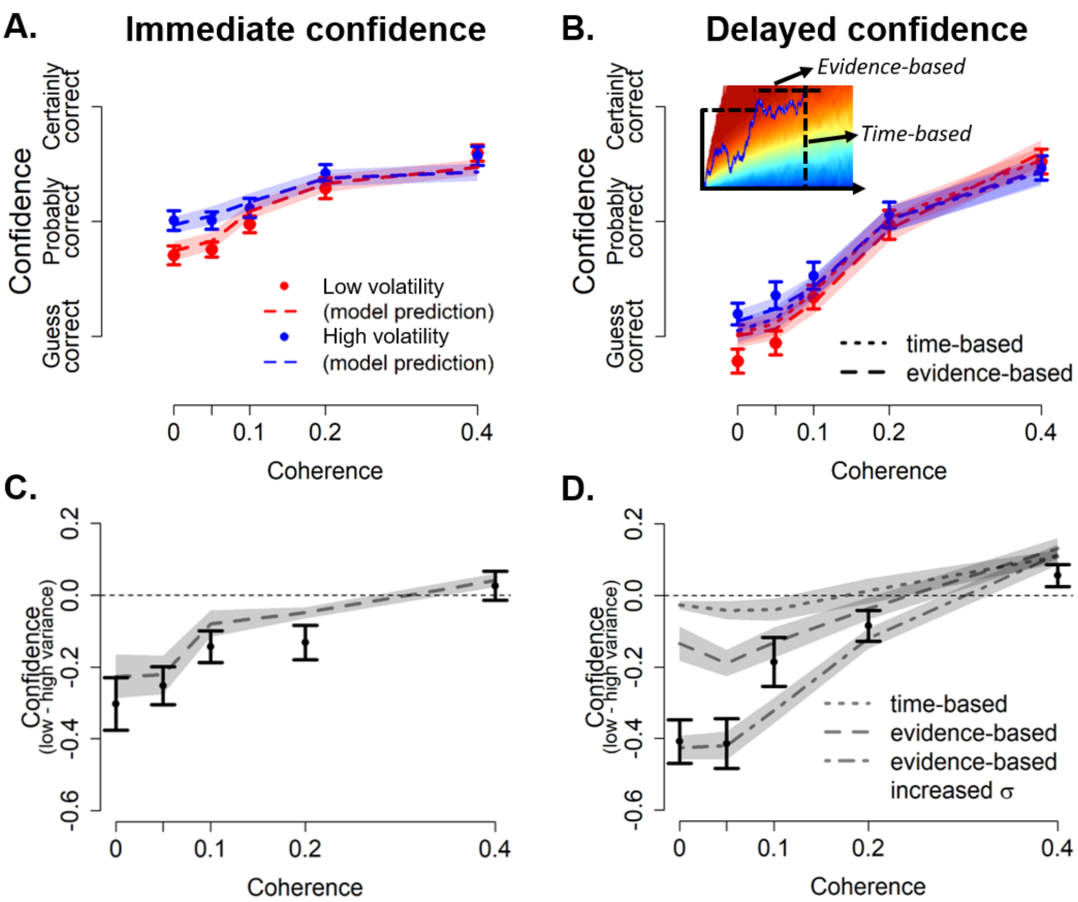

Figure 5. Within-trial evidence volatility arbitrates between an evidence-based and a time-based stopping rule. *Immediate confidence (A and C) and delayed confidence (B and D) as a function of coherence and evidence volatility, separately for the empirical data (points and bars) and model predictions (lines and shades). A and B show average confidence, C and D show differences between low and high evidence volatility. The inset on the top right shows two potential stopping criteria for post-decision processing: post-decision accumulation can stop after a fixed period of time (i.e., a vertical time-based rule) or when a fixed amount of evidence is reached (i.e., a horizontal evidence-based rule). Notes: shades and error bars reflect SEM of model and data, respectively.*

For *immediate confidence reports*, model predictions closely capture the pattern seen in human confidence ratings (see Figure 5A). Confidence monotonically increased with coherence levels (data: $F(4,22) = 27.47$, $p < .001$; model: $F(4,22) = 27.68$, $p < .001$), and was higher with high evidence volatility (data: $F(1,25) = 41.19$, $p < .001$; model: $F(1,25) = 9.90$, p

= .004). Similar to RTs, the effect of evidence volatility on confidence was most pronounced with low coherence values (data: $F(4,22) = 4.46$, $p = .008$; model: $F(4,22) = 30.79$, $p < .001$). To easily interpret this effect, Figure 5C shows differences between the low and high volatility condition. As can be seen, for both model and data, confidence was increased with high evidence volatility, particularly with low coherence values.

For *delayed confidence* reports, the data favored the evidence-based stopping rule over the time-based stopping rule (see Figure 5B and 6D). The data and both models showed a monotonic increase of confidence with coherence levels (data extra evidence: $F(4,22) = 46.67$, $p < .001$; data blank: $F(4,22) = 33.38$, $p < .001$; time-based model: $F(4,22) = 60.83$, $p < .001$; evidence-based model: $F(4,22) = 46.03$, $p < .001$), and an interaction between coherence and volatility (data extra evidence: $F(4,22) = 10.39$, $p < .001$; data blank: $F(4,22) = 8.42$, $p < .001$; time-based model: $F(4,22) = 11.94$, $p < .001$; evidence-based model: $F(4,22) = 23.50$, $p < .001$). However, evidence volatility affected confidence in the data and the model with the evidence-based stopping rule (extra evidence: $F(1,25) = 23.78$, $p < .001$; blank: $F(1,25) = 28.69$, $p < .001$; evidence-based rule, $F(1,25) = 8.96$, $p = .006$), but not with the time-based stopping rule, $F < 1$. Finally, in the human data, delayed confidence reports were similar irrespective of whether post-decision evidence or a blank screen was presented following the choice (data not shown). This was further confirmed by an analysis including post-decision evidence (extra evidence or blank), which did not show a three-way interaction, $F < 1$, $BF = .037$.

Figure 5D suggests that the effect of volatility on confidence for the lowest coherence values is even stronger than predicted by the model with the evidence-based stopping rule. This is most likely because the sigma parameter, which captures evidence volatility, was estimated based on choices and RTs only (i.e., not based on confidence). Therefore, our predictions about immediate and delayed confidence are entirely constrained by the decision process itself. Some evidence hints at the possibility that post-decision accumulation is different from pre-decision accumulation [16]. In the current context, it could therefore be that post-decision processing from memory amplifies noise in the sampling process. Indeed, when simulating the model with an evidence-based stopping rule using a slightly increased sigma value in the high volatility condition (σ = .575), it captures the pattern in the data even more tightly (see Figure 5D). This finding is in line with the possibility that post-decision accumulation is not fully determined by the pre-decision choice process.

Discussion

And influential framework conceptualizes the sense of confidence in a decision as the probability of a choice being correct. Although this formalization is principled and fruitful, it has remained unclear whether and how it can account for dynamic expressions of confidence. To close this gap, we have formalized confidence within an evidence accumulation framework as the probability of being correct, given the accumulated evidence up until that point. We tested model predictions concerning three diagnostic signatures of statistical confidence, most notably an interaction between evidence strength and choice accuracy, both for immediate and delayed confidence reports. There was a close correspondence between model and human data for all three signatures, showing that these signature of statistical confidence depend on the time at which confidence is queried.

Dynamic signatures of statistical confidence

Statistical models have conceptualized confidence as the probability of being correct [4,5,22]. Intuitively, when option A has a high (vs low) probability of being the correct answer, the model will give response A with high (vs low) confidence. One advantage of such a formalization is that it predicts the 3 qualitative signatures of confidence [4]. A limitation of such an account is that this framework is inherently static, and therefore does not take time into account. To resolve this, we relied instead on a dynamic evidence accumulation framework to probe these different signatures across time. We are not the first to account for confidence within an evidence accumulation framework [14,17,18,23,24]. Previous work has conceptualized immediate confidence as the probability of being correct given evidence and elapsed time [14,15,23]. Choices are formed when evidence reaches a fixed decision threshold, and both choice and confidence are quantified when this threshold is reached. This model is similar to ours, but it did not consider post-decision accumulation. As shown in our simulations, such a model does not predict an interaction between evidence strength and choice accuracy, a prediction at odds with many existing datasets. By quantifying confidence across time, our model can account for these discrepancies. Specifically, our model was able to explain signature 1, an interaction between evidence strength and choice accuracy, in the immediate condition, as seen in behavioral data, by assuming that immediate confidence is quantified with a small temporal delay from the choice, suggesting a brief refractory period [20,21]. Thus, an important novel insight of the current work is that some form of post-decision evidence accumulation is necessary, even to explain immediate confidence reports.

Previous modelling work has unraveled boundary conditions of this first diagnostic signature, the interaction between evidence strength and choice accuracy. Model simulations have shown that this interaction disappears if stimuli are only probabilistically related to choices [25], and if the statistical model has knowledge about evidence strength on the single-trial level [26]. Remarkably, however, no previous work has unraveled the role of time in this signature. The current work overcomes this limitation, by incorporating the notion of confidence reflecting the probability of being correct within a dynamic evidence accumulation framework. Our model simulations show that at the time of the boundary crossing, confidence increases with evidence strength for both corrects and errors, whereas the interaction effect only emerges with time. Crucially, this pattern was also observed in the empirical data. This has important consequences for studies relying on this signature to identify brain regions coding for decision confidence [5,27].

Post-decision processing terminates using an evidence-based stopping rule

Post-decision evidence accumulation has been proposed as a mechanism explaining confidence [17,18] and biases in confidence judgments [28]. It remains unclear, however, which stopping rule terminates this process of post-decision accumulation. Our data favored an evidence-based stopping rule (i.e., the sampling process terminates when a certain level of evidence has been reached), while it was incompatible with a time-based stopping rule (i.e., sampling terminates after a certain time has elapsed). Only the evidence-based rule could explain increased confidence with high evidence volatility. Intuitively, high evidence volatility increases (immediate) confidence because the injection of noise in the decision process speeds up RTs [14], and faster RTs are associated with higher confidence. The model with an evidence-based stopping rule for delayed confidence judgments similarly predicts higher confidence with high evidence volatility, because the noise again pushes the decision variable towards a certain level of evidence (i.e., a second bound). This effect does not appear with a time-based stopping rule, however, because the noise only affects the evidence (i.e., how fast is a certain level of evidence reached), but not the timing of post-decision accumulation itself. Therefore, using a time-based stopping rule the effects of evidence volatility are averaged out, and no differences in confidence are predicted. In sum, a second important insight of the current work is that human participants also use an evidence-based stopping rule in delayed confidence judgments.

Sources of post-decisional evidence accumulation

The hypothesis that confidence is affected by post-decisional evidence accumulation has evoked a strong interest in neural signatures of post-decisional processing [9,16,29]. For example, recent neuroimaging work has linked this process of post-decision evidence accumulation to a specific neural signal in the EEG [29], that is sensitive to fine-grained levels of decision confidence [1,30]. One question that has been largely overlooked so far, is what kind of information determines post-decisional evidence accumulation. For example, external information could drive post-decisional evidence accumulation [9]. Alternatively, internal sources, such as additional evidence from the sensory buffer [12] or resampling from memory [31], could determine such accumulation. To contrast these two possibilities, the current work featured conditions with and without additional external evidence during the post-decisional period. Interestingly, confidence judgments were highly similar between these two conditions. This demonstrates that, at least in our current experimental design, participant benefit exclusively from internal resampling of the earlier evidence, whereas continued external sampling has no measurable influence. This does not imply that post-decisional evidence will never play a role in confidence. For example, in a recent study that de-correlated the strength of pre-decisional and post-decisional evidence (i.e., so that sometimes post-decision evidence was highly informative when pre-decision evidence was not), external post-decisional evidence did have a reliable effect on confidence [9]. Presumably, the correlational structure of post- versus pre-decision evidence determines whether sampling continues or not.

Conclusion

The current work quantified confidence within an evidence accumulation framework as the probability of being correct given the accumulated evidence up until that point. Both model and data showed that three key signatures of statistical confidence depend on the point in time when confidence is queried. Finally, post-decision confidence reports were best explained by an evidence-based stopping rule.

Acknowledgments

The authors want to thank Peter R Murphy for advice on the modeling and Niklas Wilming, Cristian Buc Calderon and Konstantinos Tsetsos for fruitful discussions. This research was supported by an FWO [PEGASUS]² Marie Skłodowska-Curie fellowship (12T9717N, to K.D.), DFG grants DO 1240/2-1 and DO 1240/3-1 (to. T.H.D.), and a Research Council-Flanders grant G010419N (to K.D and T.V.).

Methods

Participants

Thirty participants (two men; age: *M* = 18.5, *SD* = .78, range 18 – 21) took part in return for course credit. All participants reported normal or corrected-to-normal vision and were naïve with respect to the hypothesis. All but four participants were right handed. Four participants were excluded because their performance was not different from chance level in the immediate condition (as assessed by a binomial test). Participants provided written informed consent before participation. All data have been made publicly available via the Open Science Framework and can be accessed at osf.io/83x7c. Non-overlapping analyses of these data have been published elsewhere [3].

Stimuli and apparatus

Stimuli were presented in white on a black background on a 20-inch LCD monitor with a 75 Hz refresh rate, using Psychtoolbox3 [32] for MATLAB (The MathWorks, Natick, MA). Random moving white dots were drawn in a circular aperture centered on the fixation point. The current experiment was based on code provided by Kiani and colleagues [33]. Parameter details can be found there.

Procedure

Participants completed a random dot motion task in which they additionally rated the confidence in their response. Each experimental trial started with a fixation dot for 750ms followed by random dot motion that lasted until a response was made, with a maximum of 3 seconds. On each trial, the proportion of coherently moving dots was either 0, .05, .1, .2 or .4. In each block, there was an equal number of leftward and rightward movement. In the low evidence volatility condition, this proportion was the same on every timeframe within a trial.

In the high evidence volatility condition, the proportion of coherently moving dots was on each timeframe sampled from a Gaussian distribution with mean equal to the generative distribution of that trial and a standard deviation of .256. There were three different interrogation conditions. In the *immediate* condition, participants jointly indicated their response and their level of confidence. The numerical keys ‘1’, ‘2’, ‘3’, ‘8’, ‘9’, and ‘0’ on top of the keyboard mapped onto ‘sure left’, ‘probably left’, ‘guess left’, ‘guess right’, ‘probably right’, and ‘sure right’, respectively. In the *delayed blank* condition, participants indicated their response (left or right) by pressing ‘c’ or ‘n’ with the thumbs of their right and left hand, respectively. Then, a blank screen was presented for 1s, after which the following six confidence options were presented on the screen: ‘sure correct’, ‘probably correct’, ‘guess correct’, ‘guess error’, ‘probably error’, ‘sure error’ (reversed order for half of the participants). Participants had unlimited time to indicate their level of confidence by pressing one of the corresponding numerical keys (i.e., ‘1’, ‘2’, ‘3’, ‘8’, ‘9’, and ‘0’) on top of the keyboard. The *delayed extra evidence* condition was similar to the delayed blank condition, except that now 1s of continued random motion was presented during the 1s interval between the response and the confidence judgment. The continued motion had the same direction, the same motion coherence and the same level of evidence volatility as the pre-decisional motion.

The entire experiment comprised twelve blocks of sixty trials each, including three practice blocks. In the first practice block, participants only indicated the direction of the dots (i.e., no confidence), and each trial stopped after a response was given. Only coherence levels of .2 and .4 were presented. When participants made an error, the message ‘Error’ was shown on the screen for 750ms. This block was repeated until mean accuracy exceeded 75%. The second practice block was similar, except that now the full range of coherence levels was used. This block was repeated until mean accuracy exceeded 60%. Block three served as a last practice block, and was identical to the main experiment. No more feedback was presented from this block on. Each participant then performed three blocks of each interrogation condition, with the specific order depending on a Latin square. Before the start of block seven and block ten (i.e., start of a new interrogation condition), participants performed eight practice trials with .4 coherence using the procedure of the subsequent block, to get familiarized with the response keys. These eight trials were repeated until accuracy exceeded 75%. After each block, participants received feedback about their performance in that block, including mean response time on correct trials, mean accuracy, and the absolute value of the correlation between accuracy and confidence. Participants were motivated to maximize these three values.

Data analysis

Behavioral data and model predictions were analyzed using mixed regression modeling. This method allows analyzing data at the single-trial level. We fitted random intercepts for each participant; error variance caused by between-subject differences was accounted for by adding random slopes to the model. The latter was done only when this significantly increased the model fit. RTs and confidence were analyzed using linear mixed models, for which F statistics are reported and the degrees of freedom were estimated by Satterthwaite's approximation [34]. Accuracy was analyzed using logistic linear mixed models, for which X^2 statistics are reported. Model fitting was done in R (R Development Core Team, 2008) using the lme4 package [35].

Drift diffusion modelling

Fitting. Drift diffusion model parameters were estimated using hierarchical Bayesian estimation within the HDDM toolbox [19]. The HDDM uses Markov-chain Monte Carlo (MCMC) sampling, which generates full posterior distributions over parameter estimates, quantifying not only the most likely parameter value but also uncertainty associated with each estimate. Due to the hierarchical nature of the HDDM, estimates for individual subjects are constrained by group-level prior distributions. In practice, this results in more stable estimates for individual subjects. For each model, we drew 100.000 samples from the posterior distribution. The first ten percent of these samples were discarded as burn-in and every second sample was discarded for thinning, reducing autocorrelation in the chains. Group level chains were visually inspected to ensure convergence, i.e. ruling out sudden jumps in the posterior and ruling out autocorrelation. Additionally, all models were fitted three times, in order to compute the Gelman-Rubin R hat statistics (comparing within-chain and between-chain variance). We checked and confirmed that all group-level parameters had an R hat between 0.98-1.02, showing convergence between these three instantiations of the same model. Because individual parameter estimates are constrained by group-level priors, frequentist statistics cannot be used because data are not independent. The probability that a condition differs from another can be computed by calculating the overlap in posterior distributions.

When fitting the data (choices and reaction times), both drift rate (v) and decision bound (a) were allowed to vary as a function of coherence and evidence volatility, whereas non-decision time (ter) was fixed across conditions. According to our hypothesis, the effect of evidence volatility will be expressed in the within-trial variability parameter (σ). When fitting the DDM this parameter is fixed (i.e., to .1 in the Ratcliff Diffusion model or to 1 in the currently used HDDM). Because σ is a scaling factor, after fitting the model, we next scaled drift rate, decision bound and within-trial variability for each condition so that decision bound was equal to 1. Thus, this is approach allows estimating within-trial variability. Note that under this approach, an implicit assumption is that the decision bound does not differ between the different conditions.

Simulations. Using the estimates obtained from the HDDM fit, predictions were generated using a random walk approximation of the diffusion process [36]. This method simulates a random walk process that starts at z*a (here, *z* was an unbiased starting point of .5) and stops once the integrated evidence crosses 0 or a. At each time interval *t*, a displacement Δ occurs with probability *p* and a displacement -Δ with probability *1-p*. Both quantities are given in equation (1).

$$p = \frac{1}{2}\left(1 + \frac{\mu\sqrt{\tau}}{\sigma}\right)$$
$$\Delta = \sigma\sqrt{\tau} \tag{1}$$

Drift rate is given by μ, and within-trial variability is given by σ. In all simulations τ was set to 1e-4. In order to construct the heat map representing the probability of being correct shown in Figure 1, 300.000 random walks without absorbing bounds were generated, with drift rates sampled from a uniform distribution between zero and one. This assured sufficient data points across the relevant part of the heat map. Subsequently, the average accuracy was calculated for each (response time, evidence) combination, based on all trials that had a data point for that (response time, evidence) combination. Smoothing was achieved by aggregating over evidence windows of .01 and τ windows of 2.

To generate model fits for choices and RTs and model predictions for confidence, we used the parameters obtained by the HDDM fit. For each combination of coherence levels, within-trial evidence volatility and interrogation condition, we simulated 5000 trials per participant. Both immediate and delayed confidence predictions were obtained by reading out the probability of being correct from the heat map given RT and evidence, conditional on the response given. Model predictions about confidence were then converted from a continuous scale to a categorical scaling by dividing them into three (immediate condition) or six (delayed condition) equal-sized bins. For the immediate condition, confidence predictions were obtained without any post-decision accumulation. In the adapted version confidence was quantified with a small temporal delay of .1s; other (small) values led to very similar results. For delayed confidence predictions with a time-based stopping rule, after reaching the decision bound, the random walk process continued for one second (i.e., the duration of the ITI) plus the average response speed of confidence judgments in that condition minus the non-decision time of that condition. For the evidence-based stopping rule, after the evidence crosses *a*, an evidence-based stopping rule (i.e., a horizontal boundary) was placed at *a* + *a*.125* and 0 (or similarly at –a*.125 and *a* if evidence initially crossed 0), and confidence was quantified at the time when the continued evidence accumulation crossed this second-order threshold. To project model confidence onto the same scale as human confidence, we used a linear transformation.

References

1. Desender, K., Murphy, P., Boldt, A., Verguts, T. & Yeung, N. A Postdecisional Neural Marker of Confidence Predicts Information-Seeking in Decision-Making. *J. Neurosci.* **39**, 3309–3319 (2019).
2. van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N. & Wolpert, D. M. Confidence Is the Bridge between Multi-stage Decisions. *Curr. Biol.* **26**, 3157–3168 (2016).
3. Desender, K., Boldt, A., Verguts, T. & Donner, T. H. Confidence predicts speed-accuracy tradeoff for subsequent decisions. *Elife* **e43499**, 1:25 (2019).
4. Sanders, J. I., Hangya, B. & Kepecs, A. Signatures of a Statistical Computation in the Human Sense of Confidence. *Neuron* **90**, 499–506 (2016).
5. Kepecs, A., Uchida, N., Zariwala, H. a & Mainen, Z. F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227–31 (2008).
6. Pouget, A., Drugowitsch, J. & Kepecs, A. Confidence and certainty: distinct probabilistic quantities for different goals. *Nat. Neurosci.* **19**, 366–374 (2016).
7. Urai, A. E., Braun, A. & Donner, T. H. Pupil-linked arousal is driven by decision uncertainty and alters serial choice bias. *Nat. Commun.* **8**, 1–11 (2017).
8. Braun, A., Urai, A. E. & Donner, T. H. Adaptive history biases result from confidence-weighted accumulation of past choices. *J. Neurosci.* **38**, 2418–2429 (2018).
9. Fleming, S. M., van der Putten, E. J. & Daw, N. D. Neural mediators of changes of mind about perceptual decisions. *Nat. Neurosci.* (2018). doi:10.1038/s41593-018-0104-6
10. Gold, J. I. & Shadlen, M. N. The neural basis of decision making. *Annu. Rev. Neurosci.* **30**, 535–561 (2007).
11. Ratcliff, R. & McKoon, G. The Diffusion Decision Model : Theory and Data for Two-Choice Decision Tasks. *Neural Comput.* **20**, 873–922 (2008).
12. Resulaj, A., Kiani, R., Wolpert, D. M. & Shadlen, M. N. Changes of mind in decision-making. *Nature* **461**, 263–266 (2009).
13. Calderon, C. B., Verguts, T. & Gevers, W. Losing the boundary: Cognition biases action well after action selection. *J. Exp. Psychol. Gen.* **144**, 737–743 (2015).
14. Zylberberg, A., Fetsch, C. R. & Shadlen, M. N. The influence of evidence volatility on choice , reaction time and confidence in a perceptual decision. *Elife* **5**, 1–31 (2016).
15. Kiani, R. & Shadlen, M. N. Representation of confidence associated with a decision by neurons in the parietal cortex. *Science (80-.).* **324**, 759–764 (2009).
16. Yu, S., Pleskac, T. J. & Zeigenfuse, M. D. Dynamics of Postdecisional Processing of Confidence. *J. Exp. Psychol. Gen.* **144**, 489–510 (2015).

17. Pleskac, T. J. & Busemeyer, J. R. Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychol. Rev.* **117**, 864–901 (2010).
18. Moran, R., Teodorescu, A. R. & Usher, M. Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cogn. Psychol.* **78**, 99–147 (2015).
19. Wiecki, T. V., Sofer, I. & Frank, M. J. HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front. Neuroinform.* **7**, 1–10 (2013).
20. Pashler, H. Dual-task interference in simple tasks: Data and theory. *Psychol. Bull.* **116**, 220–244 (1994).
21. Marti, S., King, J.-R. & Dehaene, S. Time-Resolved Decoding of Two Processing Chains during Dual-Task Interference. *Neuron* 1–11 (2015). doi:10.1016/j.neuron.2015.10.040
22. Maniscalco, B. & Lau, H. A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Conscious. Cogn.* **21**, 422–30 (2012).
23. Kiani, R., Corthell, L. & Shadlen, M. N. Choice Certainty Is Informed by Both Evidence and Decision Time. *Neuron* **84**, 1329–1342 (2014).
24. Ratcliff, R. & Starns, J. J. Modeling confidence judgments, response times, and multiple choices in decision making: recognition memory and motion discrimination. *Psychol. Rev.* **120**, 697–719 (2013).
25. Adler, W. T. & Ma, W. J. Limitations of Proposed Signatures of Bayesian Confidence. *Neural Comput.* **30**, 1–28 (2018).
26. Rausch, M. & Zehetleitner, M. The folded X-pattern is not necessarily a statistical signature of decision confidence. *PLoS Comput. Biol.* **15**, e1007456 (2019).
27. Kepecs, A. & Mainen, Z. F. A computational framework for the study of confidence in humans and animals. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **367**, 1322–37 (2012).
28. Navajas, J., Bahrami, B. & Latham, P. E. Post-decisional accounts of biases in confidence. *Curr. Opin. Behav. Sci.* **11**, 55–60 (2016).
29. Murphy, P. R., Robertson, I. H., Harty, S. & O'Connell, R. G. Neural evidence accumulation persists after choice to inform metacognitive judgments. *Elife* **4**, 1–23 (2015).
30. Boldt, A. & Yeung, N. Shared Neural Markers of Decision Confidence and Error Detection. *J. Neurosci.* **35**, 3478–3484 (2015).
31. Vlassova, A. & Pearson, J. Look Before You Leap: Sensory Memory Improves Decision Making. *Psychol. Sci.* **24**, 1635–1643 (2013).
32. Brainard, D. H. The Psychophysics Toolbox. *Spat. Vis.* **10**, 433–6 (1997).
33. Kiani, R., Churchland, A. K. & Shadlen, M. N. Integration of Direction Cues Is Invariant

to the Temporal Gap between Them. *J. Neurosci.* **33**, 16483–16489 (2013).

34. Kuznetsova, A., Brockhoff, P. B. & Christensen, R. H. B. lmerTest: Test in Linear Mixed Effect Models. R package version 2.0-20. http://CRAN.R-project.org/package=lmerTest. (2014).

35. Bates, D. M., Maechler, M., Bolker, B. & Walker, S. Fitting Linear Mixed-Effects Models Using lme4. *J. Stat. Softw.* **67**, 1–48 (2015).

36. Tuerlinckx, F., Maris, E., Ratcliff, R. & De Boeck, P. A comparison of four methods for simulating the diffusion process. *Behav. Res. Methods, Instruments, Comput.* **33**, 443–456 (2001).

Supplementary Materials

Statistical confidence signatures with an evidence-based stopping rule.

Model predictions about the statistical confidence signatures for the delayed condition were quantified using a time-based stopping rule. Here, we report that these predictions were highly similar when using an evidence-based stopping rule instead. First, this model also predicted that confidence scales with coherence, $F(4,225) = 84.43$, $p < .001$, as well as the interaction between coherence and choice accuracy, $F(4,225) = 232.31$, $p < .001$, reflecting increasing confidence with coherence levels for correct trials (linear contrast: $p < .001$) and decreasing for error trials (linear contrast: $p < .001$). Second, this model also predicted a monotonic positive relation between confidence and mean accuracy, $b = .06$, $t(129) = 20.42$, $p < .001$.