Modelling Speed-Accuracy Tradeoffs in the Stopping Rule for Confidence Judgments

Stef Herregods[1], Pierre Le Denmat[1], & Kobe Desender[1]

1. Brain and Cognition, KU Leuven, Belgium

Corresponding author:

Prof. dr. Kobe Desender

Brain and Cognition

KU Leuven

Tiensestraat 102, 3000 Leuven

Belgium

Abstract (232 words)

Making a decision and reporting your confidence in the accuracy of that decision are thought to reflect a similar mechanism: the accumulation of evidence. Previous research has shown that choices and reaction times are well accounted for by a computational model assuming noisy accumulation of evidence until crossing a decision boundary (e.g., the drift diffusion model). Decision confidence can be derived from the amount of evidence following post-decision evidence accumulation. Currently, the stopping rule for post-decision evidence accumulation is underspecified. Inspired by recent neurophysiological evidence, we introduce additional confidence boundaries that determine the termination of post-decision evidence accumulation. If this conjecture is correct, it implies that confidence judgments should be subject to the same strategic considerations as the choice itself, i.e. a tradeoff between speed and accuracy. To test this prediction, we instructed participants to make fast or accurate decisions, and to give fast or carefully considered confidence judgments. Results show that our evidence accumulation model with additional confidence boundaries successfully captured the speed-accuracy tradeoffs seen in both decisions and confidence judgments. Most importantly, instructing participants to make fast versus accurate decisions influenced the decision boundaries, whereas instructing participants to make fast versus careful confidence judgments influenced the confidence boundaries. Our data show that the stopping rule for confidence judgments can be well understood within the context of evidence accumulation models, and that the computation of decision confidence is under strategic control.

Keywords: Confidence, decision-making, drift diffusion model, computational modeling

Introduction

Human decision making is accompanied by a sense of confidence. Humans often report high confidence when they make correct decisions and low confidence when they make incorrect decisions (Fleming et al., 2010). Understanding the computational underpinnings of decision confidence is of high importance, given that humans use decision confidence to adapt subsequent behavior (Desender et al., 2018, 2019; Folke et al., 2016). In recent work, identifying the computational underpinnings of decision confidence has been identified as an important common goal for the field of metacognition (Rahnev et al., 2022). Given that decision confidence reflects an evaluation of the accuracy of a decision, computational accounts of decision confidence usually depart from decision making models and aim to explain the computation of confidence within these models.

In many decision making scenarios, human observers face the challenging task to make accurate decisions based on noisy evidence. Many theories of decision making assume that people solve this challenge by accumulating multiple pieces of evidence. Accumulation-to-bound models specifically propose that evidence is accumulated sequentially until the accumulated evidence reaches a predefined decision boundary. Once the decision boundary is reached, the model makes a choice (for review, see Gold & Shadlen, 2007). Within the drift diffusion model (DDM), evidence accumulates towards one of two opposite decision boundaries, with the additional assumption that evidence for both choice options is perfectly anti-correlated (Ratcliff & McKoon, 2008). In its most basic implementation, the DDM explains the dynamics of decision making using only three main parameters: a drift rate, reflecting the strength of the evidence accumulation process, a decision boundary, reflecting the degree of evidence required before a decision is made, and non-decision time, capturing non-decision related components. This simple tenet has proven to be a powerful framework that can account for a realm of behavioral and neurophysiological data. For example, accumulation-to-bound signals such as described by the DDM have been observed in human (Donner et al., 2009; O'Connell et al., 2012) and primate (Gold & Shadlen, 2007) neurophysiology. Most prominently, the DDM can explain the tradeoff between speed and accuracy that characterizes all forms of speeded decision making (Bogacz et al., 2006; Bogacz, Wagenmakers, et al., 2010). When participants are instructed to make speeded versus accurate decisions, the DDM explains these data by changing the height of the decision boundary (althoug the selectivity of this effect has been debated; Rafiei & Rahnev, 2021). Decreasing the decision boundary effectively lowers the required level of evidence before reaching it, promoting fast responses at the expense of accuracy. Given that participants are able to change the decision boundary based on instructions (amongst many other manipulations), it is believed that the height of the decision boundary is under voluntary strategic control (Balci et al., 2011; Bogacz, Hu, et al., 2010).

Given the success of the DDM in explaining decision making, several attempts have been made to explain decision confidence within these models. Capitalizing on the notion that the sense of confidence seems to arise *after* a decision has been made, Pleskac and Busemeyer (2010) proposed that the process of evidence accumulation does not terminate once a choice boundary has been crossed, but rather there is continued accumulation of (post-decision) evidence, which further informs decision confidence. If additional post-decision evidence confirms the initial decision, the model will produce a high confidence response. If additional post-decision evidence contradicts the initial decision, the model produces low confidence, or even changes its mind about the initial decision (Resulaj et al., 2009; Van Den Berg et al., 2016). Given that post-decision evidence is most likely to contradict initial decisions when these were incorrect, this account can explain why confidence is usually higher for correct than for incorrect decisions (Moran et al., 2015; Pleskac & Busemeyer, 2010), and why confidence better tracks accuracy when participants take more time to report confidence (Yu et al., 2015).

Previous modeling work mostly used this model to jointly explain choices, reaction times and decision confidence. Strikingly, much less attention has been devoted towards the speed with which confidence reports are provided. This is remarkable, given that confidence RTs are highly informative about the underlying computations (Moran et al., 2015). As a consequence, most of the abovementioned models use a too simplistic stopping rule for the post-decision evidence accumulation process. For example, typical DDM models with post-decision processing include an additional parameter which controls the duration of the post-decision processing time (i.e. the time between the choice and the confidence report; Hellmann et al., 2021; Pleskac & Busemeyer, 2010; Yu et al., 2015). Thus, in these models the stopping rule for confidence judgments is to stop accumulating post-decision evidence once a certain amount of time has passed. However, such a static implementation seems incompatible with the considerable variation in confidence RTs that is usually observed in empirical data. Indeed, under a strict interpretation, this account predicts that confidence judgments will always be provided after a fixed latency. Contrary to this, confidence RTs show the same right-skewed distributions as decision RTs. This critique also applies to a more recent proposal which quantified confidence as the maximal evidence accumulated by a leaky evidence accumulation process (Pereira et al., 2021, 2022). Although such an account explains confidence in detection tasks very well, it does not make any prediction regarding the stopping rule for confidence judgments (see also Balsdon et al., 2020).

Recent neurophysiological work suggests that the stopping rule for confidence judgments is very similar to the stopping rule for decisions (for recent review, see Desender et al., 2021). For example, Murphy and colleagues (2015) showed that both choices and error detection judgments were associated with a similar accumulation-to-bound signature over parietal electrodes in human EEG recordings. In line with this observation, Moran and colleagues (2015) established a list of empirical patterns involving confidence RTs and showed that a model with continuous post-decision accumulation until reaching a slowly collapsing confidence boundary was able to account for all empirical patterns, where other models failed to explain some of them. The work from Moran et al. (2015) provides initial evidence for (collapsing) confidence boundaries as the stopping rule for confidence judgments. Importantly, as discussed previously, there is extensive evidence that choice boundaries are under strategic control. Consequently, if a similar accumulation-to-bound mechanism underlies the stopping rule for confidence judgments, it is predicted that the termination of post-decision evidence accumulation should be similarly under strategic control. Remarkably, although there are numerous studies that have investigated speed-accuracy tradeoffs in choice formation (for review, see Bogacz, Wagenmakers, et al., 2010), to our knowledge it has yet to be investigated whether similar tradeoffs can be observed in confidence formation, and if so whether these are best accounted for by changes in the confidence boundary controlling post-decision evidence termination. Therefore, in the current work we modeled the stopping rule for confidence judgments as an accumulation-to-bound mechanism (see Figure 1 for a visual description of the model and hypotheses). In two experiments (one with a binary and one with a 6-choice confidence report), we then investigated the following hypotheses: 1) the stopping rule for confidence judgments is well described by an accumulation-to-bound mechanism similar to that for the primary decision, 2) participants can selectively modulate the height of the choice boundary and the height of the confidence boundary when instructed to do so via speed-accuracy tradeoff instructions.

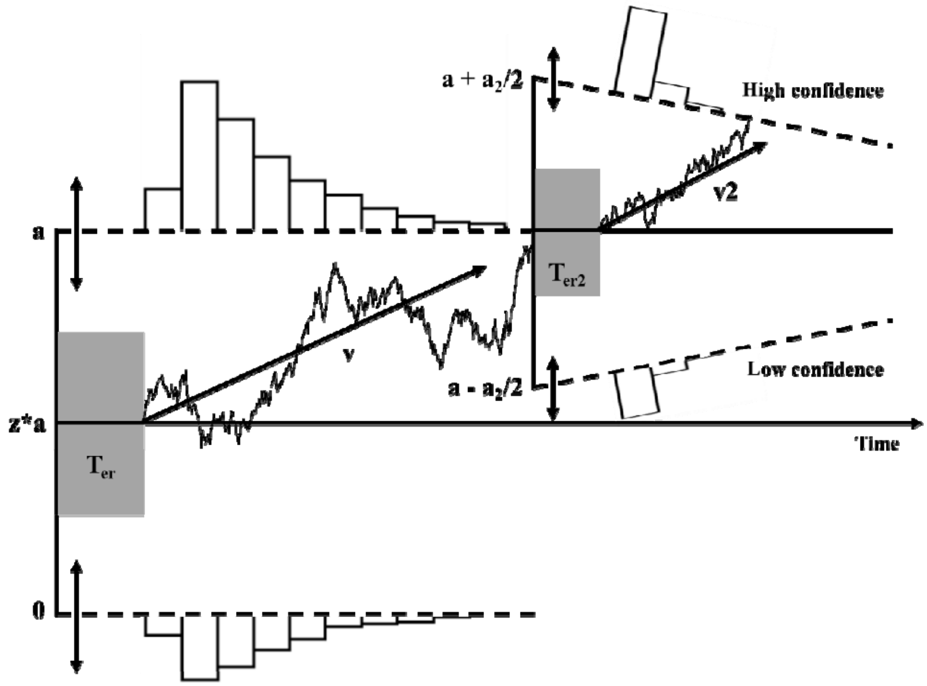

Figure 1. Drift Diffusion Model (DDM) with additional Confidence Boundaries. *In the classical DDM, evidence is assumed to accumulate (departing from starting point z*a) until it reaches one of two opposing boundaries (a or 0). The strength of the accumulation process is captured by the drift rate (v). The decision boundaries are thought to be under strategic control, and thus can be strategically increased or decreased (indicated by the arrows). To explain decision confidence, the model continues to accumulate post-decision evidence, the strength of which is controlled by the post-decision drift rate (*v_2*). Note that the ratio between v and* v_2 *has been coined v-ratio (Desender et al., 2022). The post-decision accumulation process continues until it reaches one of two opposing confidence boundaries (*a_2 *or -*a_2*). The reported level of confidence depends on the confidence boundary that was reached. Similar to the choice boundaries, the height of these confidence boundaries are thought to be under strategic control (indicated by the arrows). Non-decision related components are captured by* T_{er} *and* T_{er2} *Note that confidence boundaries are allowed to slowly collapse over time (controlled by an urgency parameter u), to account for possible speed pressure on confidence formation.*

Experiment 1

Methods and Materials

Preregistration and Code

All hypotheses, sample sizes, exclusion criteria for participants, analyzed variables, the experimental design and planned analyses were preregistered on the Open Science Framework (OSF) registries (https://doi.org/10.17605/OSF.IO/Z2UCM), unless specified as exploratory. Additionally, all code and data are made publicly available on GitHub (https://github.com/StefHerregods/ConfidenceBounds).

Participants

We decided *a priori* to test a minimum of 40 viable participants, in line with previous speed-accuracy trade-off research (Desender et al., 2022). Participant recruitment continued until this sample size was met after applying exclusion criteria. In total, 51 participants took part in Experiment 1 in return for course credit. From the total dataset, one participant gave the same confidence rating in more than 95% of the trials and 10 participants required too many training trials or did not complete the experiment in time. Data from these participants was excluded from further analyses. The final dataset comprised 40 participants (36 female), with a mean age of 18.0 (*SD* = 0.6, range = 17-19). All participants had normal or corrected-to-normal vision, and signed informed consent before their participation. The experiment was approved by the local ethics committee.

Stimuli and Apparatus

The experiment was programmed using Python v3.6.6 and PsychoPy (Peirce et al., 2019). Participants completed the experiment on 24-inch LCD screens using an AZERTY-keyboard, with blue stickers indicating buttons used for confidence judgments and red stickers indicating decision-making buttons.

Procedure

Each experimental trial started with the display of a white fixation cross on a black background for 1s (see Figure 2). Instructions regarding the speed-accuracy regime were shown above and below the fixation cross for decision-making and confidence judgments, respectively. Depending on the block, the instructions were to either ‘Make fast decisions’ or ‘Make accurate decisions’, and to ‘Give fast confidence ratings’ or ‘Think carefully about your confidence ratings’, for choices and confidence reports, respectively. For convenience, we will refer to both type of instructions as choice SAT and confidence SAT, respectively. Next, a dynamic random dot motion stimulus was presented until participants gave a response. If participants did not provide a response within 5s, the message “Too slow, please respond faster” was shown on the screen. Motion coherence was controlled by the proportion of dots consistently moving towards the left versus right side of the screen. During the main experiment, three levels of coherence were used (.1, .2 and .4). Participants were instructed to press the ‘c’ or ‘n’ key with the thumbs of their left and right hand, to indicate whether they thought dots were moving towards the left or the right, respectively. If participants responded within 5s, they were subsequently asked about their level of confidence. The text ‘How confident are you that you made the correct choice?’ appeared on top of the screen, and participants pressed the ‘e’ or the ‘u’ key with their index fingers, mapped to high and low confidence, respectively (mapping counterbalanced across participants). Confidence judgments were transformed to numeric values, with ‘low confidence’ as zero and ‘high confidence’ as one.

The experiment started with three practice blocks (24 trials each). In block 1 participants only made random dot motion decisions with a coherence of .5 for all trials. During this block they received immediate feedback about choice accuracy. Participants repeated block 1 until achieving average accuracy of 85% or more. Block 2 was identical except that the same three coherence levels as in the main phase were used (.1, .2 and .4). Participants repeated block 2 until achieving average accuracy of 60% or more. In block 3, participants no longer received trial-by-trial feedback but instead were asked about their level of confidence after each trial. Afterwards, participants took part in twelve blocks of 60 trials each. In each block there was a similar number of coherent left and right dot motion trials, and an equal occurrence of the three coherence levels. Finally, each block had specific instructions about the speed-accuracy regime for decision-making and confidence judgments. These instructions appeared both before each block and at the start of each trial (i.e. during the fixation cross). Speed-accuracy regime instructions were constant within a block, but switched after each block. Each combination of instructions appeared three times, and the order of appearance was counterbalanced across participants using a Latin square. After each block, participants received feedback about their average accuracy, average reaction time and average confidence reaction time of the preceding block.

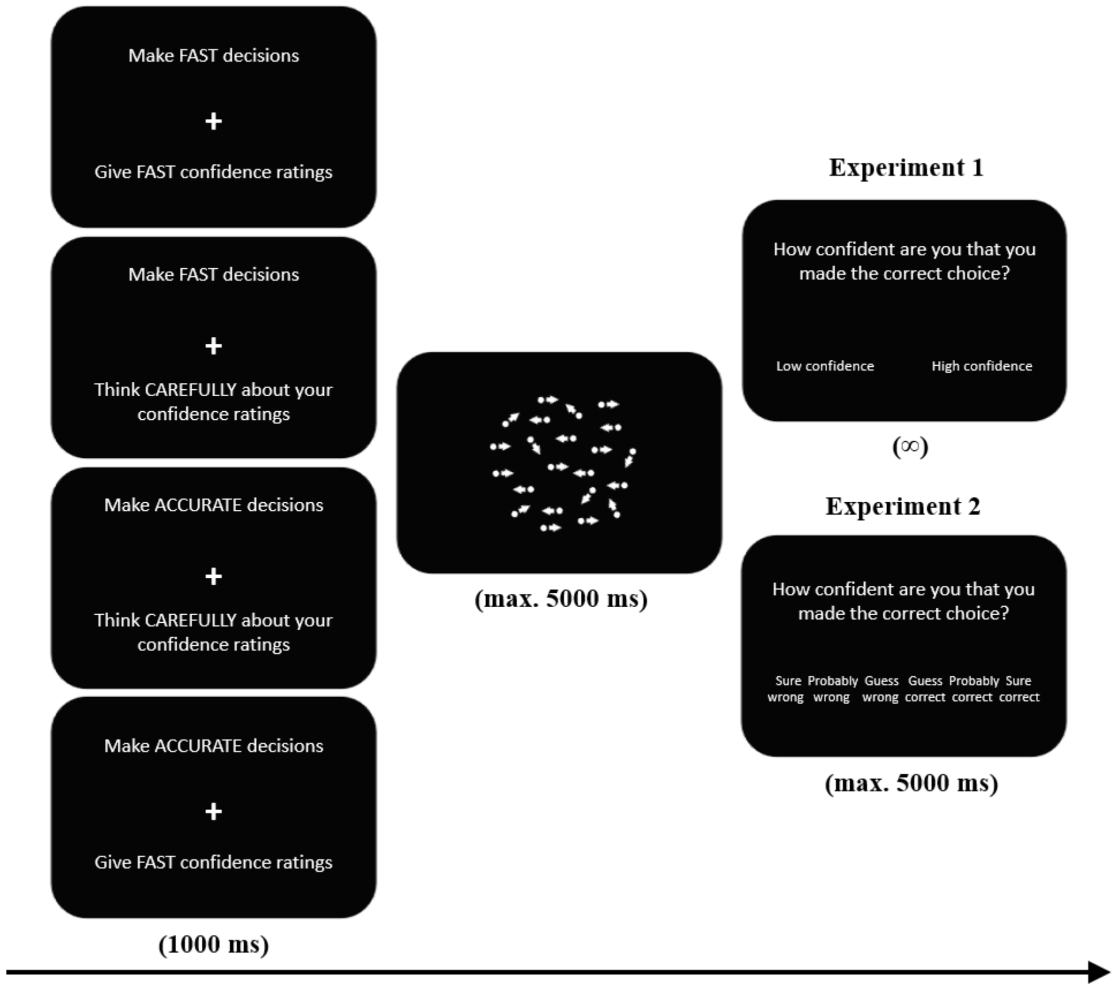

Figure 2. Example of an experimental trial. *During presentation of the fixation cross participants received specific instructions regarding the speed-accuracy regime for choices (above fixation) and confidence (below fixation). These instructions were constant within a block, but switched each block. Next, participants made binary choices about random dot motion, and afterwards indicated their level of confidence on a two-point scale (Experiment 1), or a six-point scale (Experiment 2).*

Statistical analyses

Reaction times on correct trials, accuracy, confidence judgments on correct trials and confidence RT's on correct trials were analyzed using mixed effects models. All models included at least a random intercept per participant, and all manipulations (choice SAT, confidence SAT and coherence) and their interactions as fixed effects, unless otherwise specified. These models were then extended with random slopes in order of biggest increase in BIC, until the addition of random slopes led to a non-significant increase in likelihood or until the random effects structure was too complex to be supported by the data (leading to an unstable fit). We used the lmer and glmer functions of the lme4 package (Bates et al., 2015) to fit the linear and generalized linear mixed models, respectively, in R (R Core Team, 2021). The calculation of *p* values is based on chi-square estimations using the Wald test from the car-package (Fox & Weinberg, 2019). Due to violations of the assumptions of normally distributed residuals and homoscedasticity, all RT's and confidence RT's were log transformed and mean-centered. Finally, the influence of the speed-accuracy manipulations on estimated model parameters was examined using a repeated measures ANOVA's and follow-up paired t-tests, as implemented in the rstatix package (Kassambara, 2021).

Model Specification

We simulated noisy evidence accumulation using a random walk approximation of the drift diffusion process (Tuerlinckx et al., 2001). A random walk process started at $z*a$, with z being an unbiased starting point of .5, and continued to accumulate until the accumulated evidence reaches 0 or a (reflecting the height of the decision boundaries). At each time step τ the accumulated evidence was updated with Δ, with the update rule shown in equation (1):

$$\Delta = v * \tau + \sigma * \sqrt{\tau} * N(0,1) \tag{1}$$

with v reflecting the drift rate, N reflecting the standard normal distribution, τ reflecting precision, which was set to .001 in all simulations, and σ reflecting within-trial noise which was fixed to 1. Choice and RT are quantified at the moment of boundary crossing. An additional time *ter* is added to predicted RTs to capture non-decision-related processes. After the accumulated evidence reached 0 or *a*, evidence continued to accumulate at each time step τ with displacement Δp, with the post-decision update rule shown in equation (2):

$$\Delta p = v_2 * \tau + \sigma * \sqrt{\tau} * N(0,1) \tag{2}$$

with v_2 reflecting the drift rate governing post-decisional processing. Allowing dissociations between drift rate and post-decisional drift rate is necessary to account for differences in metacognitive accuracy (Desender et al., 2022). Post-decisional accumulation continued until the lower or upper confidence boundary was reached. The height of the confidence boundaries is given by equation (3):

$$\begin{aligned} &if(choice=a)\ confidence\ boundary = a \pm a_2 \pm u * t2 \\ &if(choice=0)\ confidence\ boundary = 0 \pm a_2 \pm u * t2 \end{aligned} \tag{3}$$

with a_2 reflecting the height of the confidence boundaries, u reflecting the amount of linear urgency, and $t2$ reflecting post-decision time. The ± sign indicates that this value should be added or subtracted depending on whether it reflects the upper or lower boundary. Finally, an additional time ter_2 was added to predicted confidence RTs to capture non-confidence related processes (e.g. pressing a confidence button). In contrast to *ter*, which is by definition always positive, we also allowed ter_2 to take negative values, to account for the possibility that post-decision evidence accumulation already starts before an overt response has been made (e.g., during the motor execution of the first response). An overview of all parameters can be found in table 1.

Table 1. Parameters of the extended drift diffusion model.

Parameter	Meaning	Description
v	Drift rate	Average rate of evidence accumulation
σ	Drift coefficient	Noise in the accumulation process, fixed to .1
a	Boundary	Determines the amount of evidence required before making a choice
ter, ter_2	Non-decision time	Non-decision related components (e.g. motor execution)
z	Starting point	Determines the starting point of the accumulation process, fixed to .5.
v_2	Post-decision drift rate	Average rate of post-decision evidence accumulation
a_2	Confidence boundary	Determines the amount of evidence requires before making a confidence judgment
u	Urgency	Evidence-independent constant subtracted from the confidence boundary each time step (i.e. collapsing boundary)

Parameter Estimation and Model Fit

We estimated best fitting parameters by minimizing an error function based on quantile optimization of the RT and confidence RT distributions. Quantiles were computed in observed and simulated data for (i) decision RT quantiles, separately for correct and error trials, (ii) confidence RT quantiles, separately for correct and error trials, and (iii) confidence RT quantiles, separately for high and low confidence ratings. The resulting error function is shown in equation (4):

$$\text{RSS} = 2 * \sum(oRT_{i,q} - sRT_{i,q})^2 + \sum(oRTconf_{i,q} - sRTconf_{i,q})^2 + \sum(oRTconf_{j,q} - sRTconf_{j,q})^2 \quad (4)$$

with *oRT* and *sRT* referring to observed and simulated RT proportions, and *oRTconf* and *sRTconf* referring to observed and simulated confidence RT proportions, across multiple quantiles (q) (.1, .3, .5, .7, and .9), for correct- and error-trials (i), and high- and low confidence trials (j) separately. We minimized the abovementioned error function using differential evolution optimization as operationalized in the DEoptim package, and by setting the amount of iterations to 1000 (Mullen et al., 2011). Model fitting was done separately per participant. To test model fits, we simulated choices, RTs, confidence and confidence RTs from the estimated parameters.

Parameter Recovery

Before estimating parameters based on empirical data, we performed parameter recovery to ensure that the extended DDM accurately recovers known parameters. For this end, we simulated data for N = 40 synthetic participants, once with N_{trials} = 10.000 (i.e., simulating an ideal scenario) and once with N_{trials} = 180 (i.e. the number of trials per cell in our design). Parameters were randomly sampled from a uniform distribution with the minimum and maximum values chosen such that they were in line with the empirically observed fits; a = [.5, 3], a_2 = [.5, 5], u = [0, 3], v = [0, 3], ter = [0, 1], ter_2 = [-.5, .5], v_2 = [0, 7]. Subsequently, we performed linear regression predicting the true parameter value by the estimated value. Inspection of the regression results revealed that with 10.000 trials all slopes were significant (ps < .002) and close to 1, reflecting excellent recovery. The only exception was the urgency parameter, *u*, which yielded a slope of .19. Given the low recovery of this parameter, estimates of *u* should be interpreted with caution. Results did not drastically change when the parameter recovery was repeated with 180 trials. Full results of the parameter recovery for Experiment 1 can be found in table S1.

Results

Behavioral Analysis

Trials with RTs below .2s were excluded from the dataset (00.40%) (Moran et al., 2015). In addition, confidence RTs slower than 5s were excluded (00.10%; note that choice RTs slower than 5s were excluded by design). Next, we report a set of analyses testing how RTs, confidence RTs, accuracy and confidence judgments were influenced by motion coherence (3 levels: .1, .2 and .4), choice SAT (2 levels: fast vs accurate) and confidence SAT (2 levels: fast vs careful).

For reaction times on correct trials (shown in Figure 3A), as expected we found a significant effect of choice SAT instructions, $\chi^2(1) = 68.87$, $p < .001$, but not of confidence SAT instructions, $\chi^2(1) = 0.56$, $p = .455$. Choice RTs were shorter when participants were instructed to respond fast ($M = 0.92$s) versus accurate ($M = 1.38$s). Also the main effect of motion coherence was significant, $\chi^2(2) = 687.25$, $p < .001$, reflecting shorter RTs with increasing motion coherence. Additionally, we found a significant interaction between the choice SAT and confidence SAT, $\chi^2(1) = 15.35$, $p < .001$, reflecting that the choice SAT effect was more expressed when participants were instructed to provide accurate versus careful confidence ratings. There was also a significant interaction between choice SAT and coherence, $\chi^2(1) = 43.24$, $p < .001$, reflecting that the choice SAT effect was slightly larger for low coherence trials. All other effects were not significant, *p*s > .525. For accuracy, we likewise found a significant effect of the choice SAT, $\chi^2(1) = 8.25$, $p = .004$, and coherence, $\chi^2(2) = 165.58$, $p < .001$, but not of the confidence SAT, $\chi^2(1) = 0.63$, $p = .429$. As shown in Figure 3B, participants responded more correct when instructed to be accurate ($M = 72.28\%$) compared to when instructed to be fast ($M = 70.78\%$), and accuracy increased with motion coherence. All other effects were not significant, *p*s > .071.

For confidence RTs on correct trials, we found significant effects of the confidence SAT instructions, $\chi^2(1) = 77.06$, $p < .001$, and coherence, $\chi^2(2) = 14.29$, $p = .001$. As expected, choice SAT instructions did not influence confidence RTs, $\chi^2(1) = 0.42$, $p = .518$. As can be seen in Figure 3C, confidence RTs were faster when participants were instructed to make fast ($M = 0.31$s) vs careful ($M = 0.73$s) confidence judgments. Additionally, we found a significant interaction between choice SAT and confidence SAT, $\chi^2(1) = 6.28$, $p = .012$, reflecting a small spill-over from choice SAT into confidence RTs (mostly visible in the "accurate" condition). All other effects were not significant, *p*s > .464. Finally, for confidence judgments (see Figure 3D) we observed a significant main effect of coherence, $\chi^2(2) = 120.71$, $p < .001$, reflecting that confidence increased with the proportion of motion coherence. There were no significant main effects of choice SAT, $\chi^2(1) = 1.23$, $p = .267$, nor confidence SAT, $\chi^2(1) = 0.27$, $p = .606$. There was only a small but significant interaction between choice SAT and confidence SAT, $\chi^2(1) = 4.89$, $p = .027$, reflecting that participants more often reported high confidence for fast ($M = .64$) than for accurate ($M = .61$) choices in the fast confidence condition, whereas there were was no such difference in the careful confidence condition (.62 vs .62, respectively). Finally, there was an interaction between choice SAT and coherence, $\chi^2(2) = 9.97$, $p = .007$, reflecting that the relation between confidence and coherence was slightly stronger in the accurate compared to the fast choice condition. All other effects were not significant, *p*s > .160.

Given that there was no effect of the SAT manipulations on average confidence, we additionally examined whether there was a difference in confidence resolution (i.e. the relation between confidence and accuracy). To do so, we computed type II ROC separately for each condition (ignoring coherence). Note that these analyses were not pre-registered. A 2-way ANOVA on these estimates showed a main effect of confidence SAT, $F(1,39) = 15.42$, $p < .001$, but not from choice SAT, $p = .599$, nor was there an interaction, $p = .491$. As can be seen in Figure 4A, the relation between confidence and accuracy (expressed in AUC units) was higher when participants were instructed to make deliberate versus fast confidence ratings. Thus, although confidence did not strongly change on average, there was clear evidence in our data that confidence SAT did influence confidence resolution.

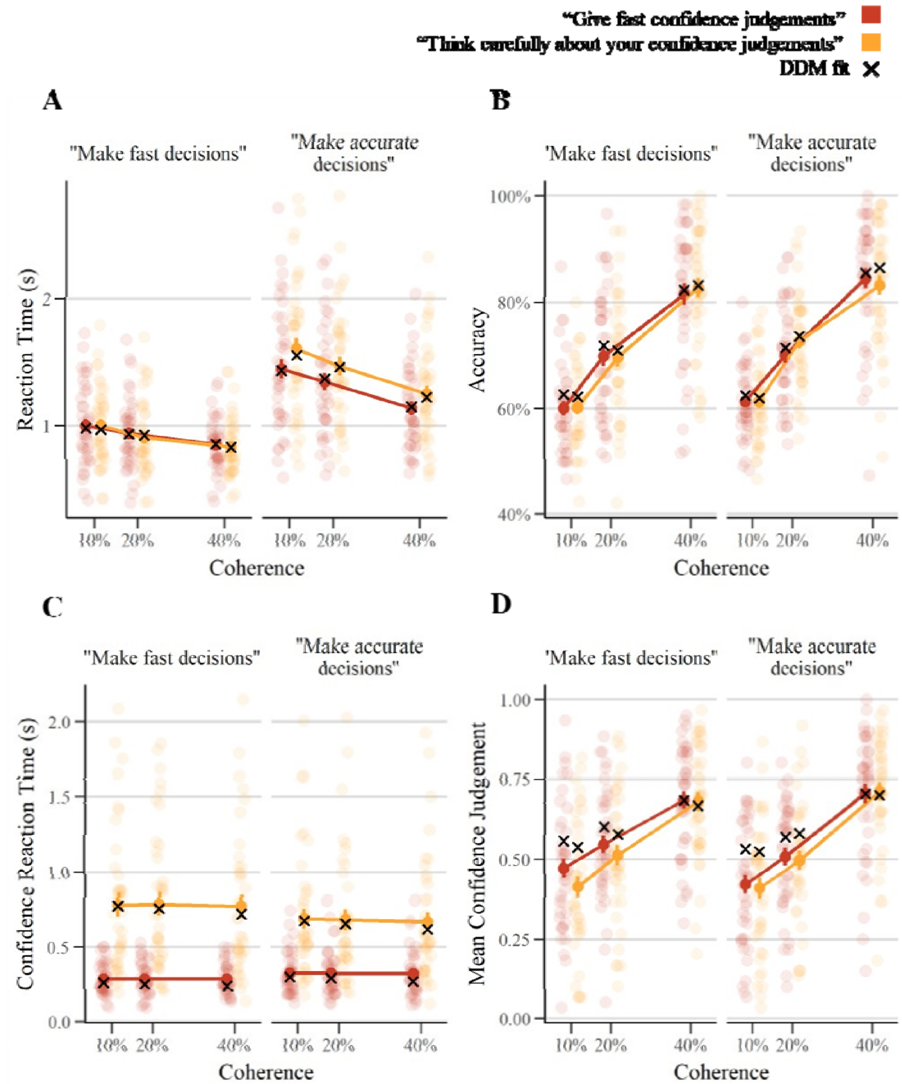

Figure 3. The influence of choice SAT and confidence SAT on reaction times (A), accuracy (B), confidence RTs (C) and confidence (D) for Experiment 1. *As expected, when participants were instructed to make fast versus accurate choices this led to fast versus slow choice RTs (A) and to a lesser extent to less and more accurate choices (B), respectively. When participants were instructed to make fast vs deliberate confidence judgments, this led to fast versus slow confidence RTs (C). The effects on confidence judgments (D) were pronounced. Note: error bars reflect SEM, transparent dots reflect means of individual participants, black crosses reflect extended DDM fits.*

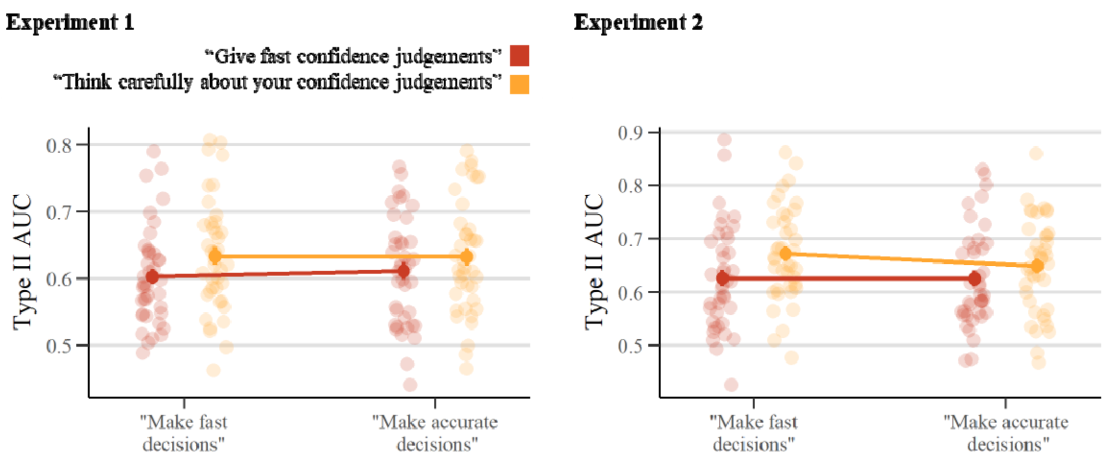

***Figure 4**. Confidence resolution in Experiment 1 and Experiment 2, expressed as Type II AUC*. *Although confidence SAT instructions did not have a clear effect on average confidence, we did observe a clear effect on confidence resolution, which was not the case for choice SATs. Same conventions as in Figure 3.*

Modeling speed-accuracy tradeoffs in choices and confidence

Our modeling framework departed from the classical drift diffusion model (DDM), a popular evidence accumulation model that accounts well for choices and reaction times in perceptual decisions (Ratcliff & McKoon, 2008). To also account for confidence within the DDM, we allow the evidence to accumulate after it has reached a threshold (post-decision evidence accumulation; Pleskac & Busemeyer, 2010). Critically, post-decision evidence accumulation continues until the evidence reaches a second boundary. We will refer to this second boundary as the confidence boundary to dissociate it from the (first) decision boundary. Given that Experiment 1 only has two levels of confidence (high vs low), confidence here fully coincides with the boundary that was reached (i.e., high versus low confidence when reaching the upper vs lower confidence boundary; see Figure 1). In addition to a confidence boundary, we also allowed the height of the confidence boundaries to collapse over time, an effect often referred to as urgency since it accounts for possible time-pressure effects.

Model fit

Having confirmed that our speed accuracy tradeoff instruction had the desired effect for both choices and for confidence judgments, we turned towards computational modeling of our data. Before looking at the estimated parameters, we first ensured that our model provided a good account of the decision process underlying the data, by examining whether it successfully captures both choices, RTs, confidence judgments and confidence RTs at the same time. In Figure 3, model predictions for RTs, choices, confidence and confidence RTs are plotted on top of the observed data. As can be seen, our model captured the trends in the data very well. This is further confirmed by analyzing model predictions (generated using the same number of trials as in the empirical data) in the same way as previously done with empirical data, which provided highly similar results. Most importantly, choice RTs and accuracy were both modulated by choice SAT instructions (RTs: $\chi^2(1) = 63.52$, $p < .001$; accuracy: $\chi^2(1) = 5.30$, $p = .021$), but not by confidence SAT instructions (RTs: $p = .262$; accuracy: $p = 388$). Reversely, confidence RTs and confidence were modulated by confidence SAT instructions (confidence RTs; $\chi^2(1) = 71.75$, $p < .001$; confidence: $\chi^2(1) = 5.15$, $p = .023$), but not by choice SAT instructions (confidence RTs: $p = .998$; confidence: $p = .708$). Note that the main effect of confidence SAT instructions on confidence was not significant in the empirical data, reflecting a subtle but qualitative difference. We also note that the model slightly overestimates confidence for low coherence trials. The full results for the analysis of model predictions can be found in the Supplementary Materials, table S3. Finally, because the DDM aims to explain entire RT distributions, and not just summary statistics, we also inspected similarities between observed and simulated RT and confidence

RT distributions. As can be seen in Figure 5, the model captured the RT and confidence RT distributions very well across the different SAT manipulations.

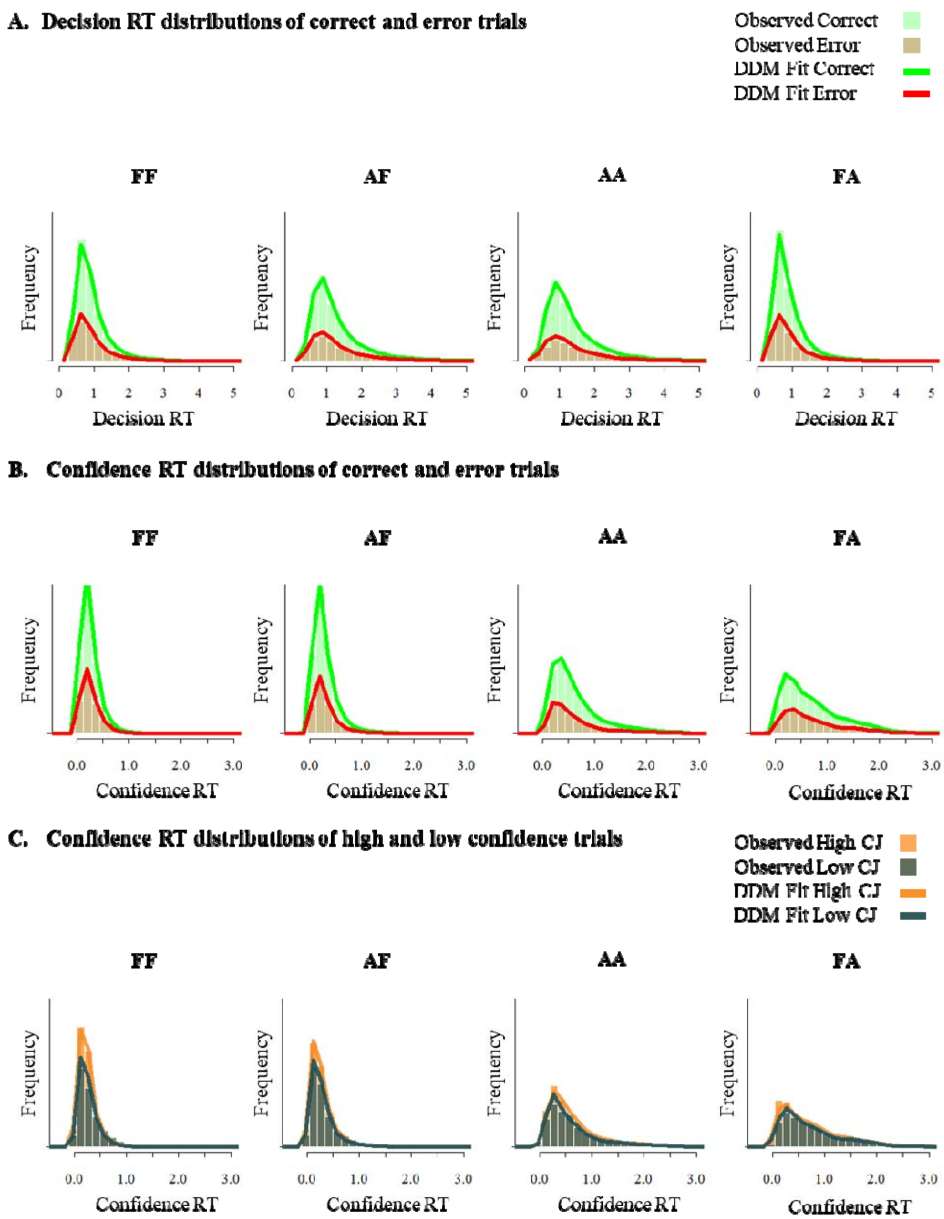

Figure 5. Distributions of reaction times (A) and confidence RTs (B-C) for data and model fit for Experiment 1. *Inspection of the model fit reveals that an extended DDM accurately captures the distributions in reaction times and confidence RTs seen in the data, across all four SAT manipulations. Note: AA, AF, FA and FF refer to SAT instructions to be*

accurate/cautious (A) or fast (F) with the first index referring to the decision and the second index referring to confidence.

Model parameters

Having established that our model provides a good fit to the experimental data, we next turn towards the actual parameters. We hypothesized that SAT instructions for choices selectively affected choice boundaries, leaving confidence boundaries unaffected. Likewise, we expected SAT instructions about confidence to selectively affect confidence boundaries, leaving choice boundaries unaffected. These observations would support our hypothesis that indeed the stopping rule for both choices and choice confidence are under strategic control.

First, we used a repeated measures ANOVA to examine the influence of choice SAT (fast vs accurate), confidence SAT (fast vs careful), and their interaction on estimated decision boundaries. As expected, we found a strong and significant effect of the choice SAT, $F(1, 39) = 75.10$, $p < .001$, $\eta_p^2 = .66$, but not of confidence SAT, $F(1, 39) = 0.78$, $p = .384$, $\eta_p^2 = .02$, nor was there an interaction, $F(1, 39) = 2.82$, $p = .101$, $\eta_p^2 = .07$. As can be seen in Figure 6A, when participants were asked to make fast decisions, the separation between both choice boundaries was smaller ($M = 1.46$) compared to when they were asked to make accurate decisions ($M = 1.96$).

Second, the same analysis on the estimated confidence boundaries revealed, as expected, a significant main effect of confidence SAT, $F(1, 39) = 36.14$, $p < .001$, $\eta_p^2 = .48$, but not of choice SAT, $F(1, 39) = 3.29$, $p = .078$, $\eta_p^2 = .08$. There was, however, a significant interaction between both types of instruction, $F(1, 39) = 9.20$, $p = .004$, $\eta_p^2 = .19$. Paired t-tests showed that indeed confidence boundaries were higher when confidence SAT instructions were to make careful as opposed to fast confidence judgments, and this effect was slightly stronger when choice SAT focused on speed, $t(39) = -6.45$, $p < .001$, then when it focused on accuracy, $t(39) = -4.00$, $p < .001$, although it was highly significant in both cases (see Figure 6B). Choice SAT instructions did not influence confidence boundaries when confidence SAT instructed people to respond fast, $t(39) = -0.45$, $p = .658$, but when confidence SAT instructed people to respond carefully, choice SAT seemed to influence estimated confidence boundaries, $t(39) = 2.83$, $p = .007$.

Third, we looked at the influence of SAT instructions on estimated urgency parameters. There were no significant effect of the choice SAT, $F(1, 39) = 1.88$, $p = .179$, $\eta_p^2 = .05$, nor of the confidence SAT, $F(1, 39) = 0.002$, $p = .967$, $\eta_p^2 < 0.001$, nor was there an interaction, $F(1, 39) = 0.05$, $p = .824$, $\eta_p^2 = 0.001$. Thus, it seems that confidence SAT instructions are implemented by changing the height of the confidence decision boundaries, while leaving urgency constant. In Figure 6C, two outliers with an urgency value higher than 5 can be noticed. Removing these outliers from the analysis did not alter any of the conclusions.

Finally, as a sanity check we confirmed that estimated drift rates scaled with motion coherence using a repeated measures ANOVA, $F(1.18, 45.98) = 108.304$, $p < .001$, $\eta_p^2 = 74$. The other parameters were not allowed to vary, their mean estimates can be found in Table 2.

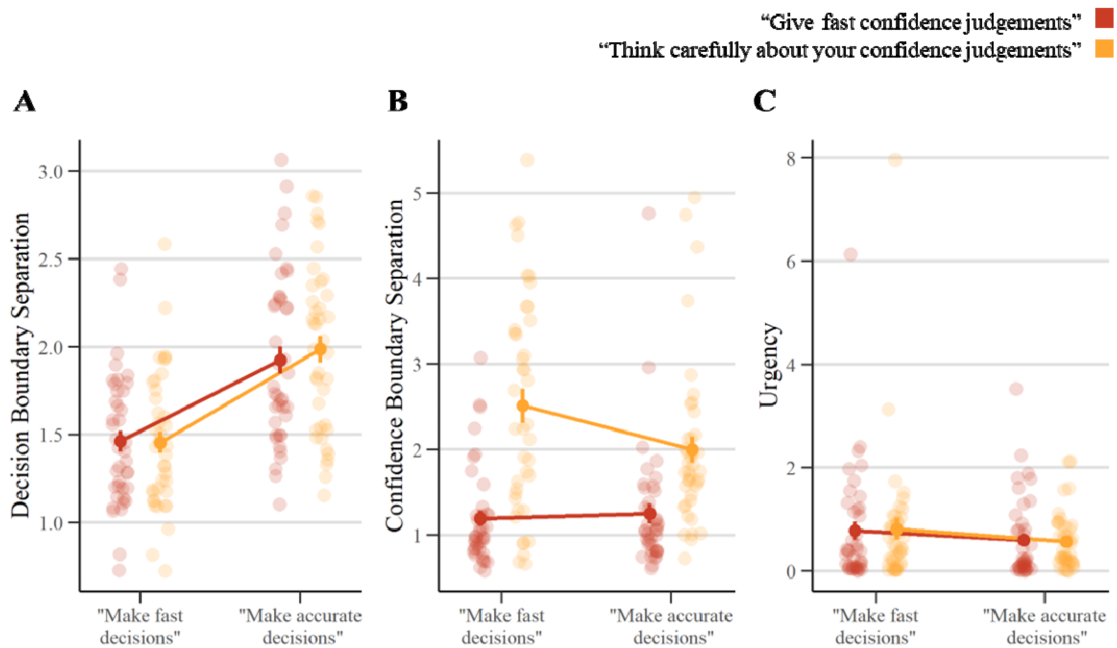

Figure 6. Influence of choice SAT and confidence SAT on decision boundaries and confidence boundaries in Experiment 1. *Instructing participants to make fast vs accuracy choices influences estimated decision boundaries (A), but rarely influences confidence boundaries or urgency (B-C). Instructing participants to provide fast vs careful confidence ratings influences estimated confidence boundaries (B), but does not affect decision bounds and urgency (A,C). Same conventions as in Figure 4.*

***Table 2. Mean (SD) Estimates for the Extended Drift Diffusion Model**. Note, AA, AF, FA and FF refer to SAT instructions to be accurate/cautious (A) or fast (F) with the first index referring to the decision and the second index referring to confidence.*

Parameter	AA	AF	FA	FF
Experiment 1				
a	1.99 (0.48)	1.92 (0.49)	1.46 (0.39)	1.46 (0.38)
v_1 (coherence = 0.1)	0.26 (0.15)	0.27 (0.15)	0.34 (0.26)	0.36 (0.27)
v_2 (coherence = 0.2)	0.56 (0.31)	0.51 (0.27)	0.65 (0.39)	0.68 (0.41)
v_3 (coherence = 0.4)	1.12 (0.64)	1.17 (0.68)	1.31 (0.68)	1.27 (0.70)
ter	0.51 (0.25)	0.44 (0.17)	0.39 (0.13)	0.41 (0.14)
ter_2	0.08 (0.20)	0.04 (0.08)	0.05 (0.30)	0.05 (0.09)
v-ratio (= v_2 / v)	1.32 (1.73)	2.02 (2.19)	0.97 (1.27)	1.78 (1.76)
a_2	2.00 (0.97)	1.25 (0.73)	2.51 (1.28)	1.20 (0.57)
u	0.57 (0.55)	0.59 (0.79)	0.82 (1.31)	0.78 (1.13)
Experiment 2				
a	2.11 (0.46)	1.81 (0.41)	1.55 (0.39)	1.41 (0.30)
v_1 (coherence = 0.1)	0.30 (0.19)	0.36 (0.23)	0.36 (0.24)	0.34 (0.22)
v_2 (coherence = 0.2)	0.60 (0.39)	0.58 (0.31)	0.70 (0.38)	0.70 (0.44)
v_3 (coherence = 0.4)	1.18 (0.68)	1.08 (0.69)	1.33 (0.81)	1.32 (0.83)
ter	0.52 (0.18)	0.50 (0.18)	0.42 (0.12)	0.43 (0.10)
ter_2	-0.53 (0.61)	-0.33 (0.42)	-0.76 (0.94)	-0.23 (0.38)
v-ratio (= v_2 / v)	1.05 (1.28)	1.66 (2.40)	1.25 (1.91)	1.40 (1.85)
a_2*_upper*	3.76 (2.91)	3.30 (3.01)	4.22 (3.17)	2.63 (2.16)
a_2*_lower*	-5.47 (3.90)	-5.44 (3.46)	-5.87 (3.35)	-4.58 (3.58)

u_upper	2.12 (1.65)	3.95 (2.77)	1.96 (1.64)	3.22 (2.74)
u_lower	2.69 (2.74)	4.15 (3.30)	2.67 (2.74)	3.36 (3.26)

Interim Summary

In Experiment 1, participants were instructed to make fast or accurate decisions and to make fast or careful confidence judgments, depending on the block they were in. At the behavioral level, we observed that participants were indeed able to selectively speed up choices or confidence judgments when instructed to do so. More importantly, model fits using an extended DDM with additional confidence boundaries revealed that the mechanism underlying such behavior was a change in the decision boundary for choices, and the confidence boundary for confidence. Thus, these findings show that the stopping rule for confidence judgments, just like the choice boundary for choices, is under voluntary strategic control. One limitation of Experiment 1 is that participants were only allowed to give binary confidence ratings (high or low). This design choice made for an easy modeling approach, because it allows to directly map high and low confidence onto the upper and lower confidence boundary, respectively. It is well known, however, that humans can provide more fine-grained estimates of their performance. Thus, this begs the question whether our extended DDM can also account for tasks with more fine-grained confidence scales. For this end, in Experiment 2 we replicated Experiment 1, but now using a more fine-grained 6-choice confidence scale.

Experiment 2

Methods and Materials

Preregistration and Code

The preregistration of this experiment can be found on OSF registries (https://doi.org/10.17605/OSF.IO/VYH4K), all code and data can be found on GitHub (https://github.com/StefHerregods/ConfidenceBounds).

Participants

A total of 54 participants participated in Experiment 2. Requirements and recruitment was identical to Experiment 1, with the additional criterium of not having participated in Experiment 1. Data of six participants were removed for not having an accuracy above chance level as assessed by a binomial test, and four participants for requiring more than seven training blocks. Finally, four participants did not finish the experiment in time. The final sample comprised 40 participants (33 female), with a mean age of 18.5 (*SD* = 1.3, range = 17 - 24).

Stimuli and Apparatus

Experiment 2 used the same apparatus and stimuli as in Experiment 1.

Procedure

The experiment was identical to Experiment 1, except for the following two exceptions: First, instead of a binary confidence rating, participants could choose between six options; ‘Sure wrong’, ‘Probably wrong’, ‘Guess wrong’, ‘Guess correct’, ‘Probably correct’ and ‘Sure correct’, using the ‘1’, ‘2’, ‘3’, ‘8’, ‘9’ and ‘0’ keys on top of the keyboard. These six options were mapped onto a 1-6 confidence scale (counterbalanced between participants). Second, a time-limit of 5s was imposed on indicating confidence judgments, equal to the time-limit during decision-making. If a participant did not respond within this limit, they were instructed to respond faster in future trials with the following text: ‘Too slow… Please respond faster’.

Model Specification and fit

The modeling strategy was identical to Experiment 1 except for the following: we estimated separate parameters for upper and lower confidence boundary separation, and separate parameters for urgency of the upper and the lower boundary. This was done to allow the model to account for the negative relationship between confidence and confidence RTs (discussed in more detail below). In order for the model to be able to produce six levels of confidence, we changed the implementation such that confidence no longer corresponded to the confidence boundary that was reached (as in Experiment 1). Instead, confidence depended on the level of accumulated evidence at the time of reaching the confidence boundary. We evenly divided the space in between the two confidence boundaries into six categories, and the model produced a level of confidence between 1 and 6 depending on the state of the accumulated evidence.

Parameter Recovery

Because the extended DDM used in Experiment 2 differed in important aspects from Experiment 1, we repeated the parameter recovery exercise. The following minimum and maximum values were used: a = [.5, 3], a_{2_upper} = [.1, 15], a_{2_lower} = [.1, 15], u_{upper} = [0, 15], u_{lower} = [0, 15], v = [0, 3], ter = [0, 1], ter_2 = [-.2, .2], v_2 = [0, 5]. Inspection of the regression analyses results revealed that with 10.000 trials all slopes were significant (*p*s < .001), but the estimates were less close to 1 compared to Experiment 1. Given these results, particularly the estimate of v_2 should be interpreted with caution. Results did not change drastically when the parameter recovery was repeated with only 180 trials. Full results of the parameter recovery for Experiment 2 can be found in table S2.

Results

Behavioral Analysis: Mixed Effects Modelling

Data were analyzed in the same way as described in Experiment 1. Similar to Experiment 1, trials with a decision time of less than 0.2 s were excluded (0.30%). A mixed effects model on decision RTs on correct trials showed a significant effect of choice SAT, $\chi^2(1) = 57.91$, $p < .001$, and coherence, $\chi^2(2) = 956.89$, $p < .001$. Unexpectedly, there also was a significant effect of confidence SAT, $\chi^2(1) = 34.02$, $p < .001$. Additionally, we found significant interactions between the choice SAT and confidence SAT, $\chi^2(1) = 8.72$, $p = .003$, between coherence and choice SAT, $\chi^2(2) = 21.80$, $p < .001$, and between coherence and confidence SAT, $\chi^2(2) = 12.10$, $p = .002$. The three-way interaction between choice SAT, confidence SAT and coherence was not significant, $\chi^2(2) = 0.65$, $p = .722$. As can be seen in Figure 7A, choice RTs were shorter when participants were instructed to respond fast (M = 0.93s) versus accurate (M = 1.34s), however the effect was not as selective as in Experiment 1, because choice RTs were also shorter when participants were instructed to provide fast (M = 1.06s) versus careful confidence ratings (M = 1.20s). The same analysis on accuracy likewise showed significant main effects of choice SAT, $\chi^2(1) = 6.05$, $p = .014$, confidence SAT, $\chi^2(1) = 4.76$, $p = .029$, and coherence $\chi^2(2) = 1202.14$, $p < .001$ (see Figure 7B). Accuracy was lower when participants were instructed to make fast (M = 73%) compared to accurate choices (M = 75%), and likewise when participants were instructed to make fast (M = 74%) versus careful confidence ratings (M = 75%). All other effects were not significant, *p*s > .257.

The same analysis on confidence RTs on correct trials, showed significant main effects of confidence SAT, $\chi^2(1) = 85.62$, $p < .001$, and coherence, $\chi^2(2) = 71.12$, $p < .001$. Unexpectedly, there was also a significant main effect of choice SAT, $\chi^2(1) = 9.64$, $p = .002$. Finally, the interaction between the confidence SAT and coherence was significant, $\chi^2(2) = 6.26$, $p = .044$. All other effects were not significant, *p*s > .161. As can be seen in Figure 7C, although confidence SAT clearly affected confidence RTs in the expected way, the effect was not as selective as in Experiment 1. Confidence RTs were shorter when participants were instructed to make fast (*M* = .39s) versus careful (*M* = .76s) confidence ratings, and counterintuitively confidence RTs were slightly longer when participants were instructed to make fast (*M* = .62s) versus accurate (*M* = .53s) decisions.

Finally, the same analysis was carried out on confidence for correct trials. Note that for this analysis, the three-way interaction and the interaction between the choice SAT and confidence SAT were excluded because they caused variance inflation factors higher than 10. In the final model, there was a significant main effect of coherence, $\chi^2(2) = 100.06$, $p < .001$, and the confidence SAT, $\chi^2(1) = 4.36$, $p = .037$, but not of the choice SAT, $\chi^2(1) = 0.84$, $p = .359$. As can be seen in Figure 7D, variations in confidence were mostly driven by coherence, but confidence was also slightly lower when participants were instructed to make fast (*M* = 4.54) versus careful (*M* = 4.84) confidence judgments. Finally, we found a significant interaction between the confidence SAT and coherence, $\chi^2(2) = 22.10$, $p < .001$, reflecting that the confidence SAT was more pronounced on low coherence trials. The interaction between the choice SAT and coherence was found to be not significant, $\chi^2(2) = 1.57$, $p = .455$.

Similar to Experiment 1, in a non-preregistered analysis we additionally looked at confidence resolution by calculating type II AUC separately for each condition. Again, a 2-way ANOVA showed a main effect of confidence SAT, $F(1,39) = 14.49$, $p < .001$, but not from choice SAT, $p = .066$, nor was there an interaction, $p = .125$. As can be seen in Figure 4B, the relation between confidence and accuracy (expressed in AUC units) was higher when participants were instructed to make careful vs fast confidence ratings

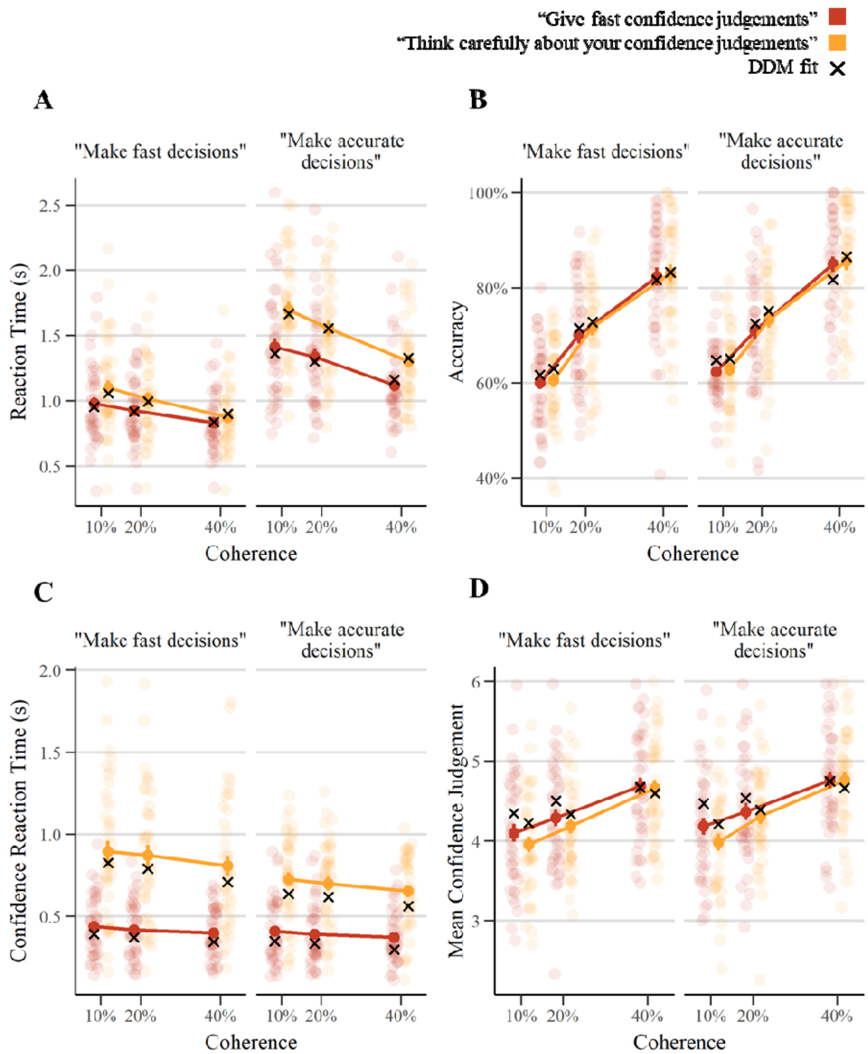

Figure 7. The influence of choice SAT and confidence SAT on reaction times (A), accuracy (B), confidence RTs (C) and confidence (D) for Experiment 2. *As expected, when participants were instructed to make fast versus accurate choices this led to fast versus slow choice RTs (A) and to a lesser extent to less and more accurate choices (B), respectively. When participants were instructed to make fast vs deliberate confidence judgments, this led to fast versus slow confidence RTs (C), with less pronounced effects on confidence judgments. Note: same conventions as in Figure 3.*

Extended Drift Diffusion Model Fits

Model fit

Similar to Experiment 1, we again first inspected whether our model captured the patterns in the data well. In Figure 7, it can be appreciated that model predictions capture the trends seen in the behavioral data. This is further confirmed by analyzing model predictions (generated using the same number of trials as in the empirical data) in the same way as previously done with empirical data, which provided highly similar results. Most importantly, choice RTs and accuracy were both modulated by choice SAT instructions (RTs: $\chi^2(1)$ = 61.15, $p < .001$; accuracy: $\chi^2(1) = 7.95$, $p = .005$), as well as by confidence SAT instructions (RTs: $\chi^2(1) = 23.62$, $p < .001$; accuracy: $\chi^2(1) = 13.61$, $p = .005$), similar to what was found in the behavioral data. Confidence RTs and confidence were modulated by confidence SAT instructions (confidence RTs; $\chi^2(1) = 65.47$, $p < .001$; confidence: $\chi^2(2) = 4.88$, $p = .027$), and choice SAT instructions influenced confidence RTs, $\chi^2(1) = 11.92$, $p < .001$, but not confidence, $p = .929$. Note that, similar to the fits in Experiment 1, the model slightly overestimates confidence for low coherence trials. The full results for the analysis of model predictions can be found in the Supplementary Materials, Table S4. Finally, in Figure 8 it can be seen that the model captures RT and confidence RT distributions very well across the different SAT manipulations. Because Experiment 2 features six confidence levels, we additionally investigated the relation between the reported level of confidence and confidence RT. As can be seen in Figure 8D, participants tended to be faster when reporting high then low confidence, a pattern that was captured well by the computational model. Note that a model with a single parameter controlling the height of both the upper and the lower confidence boundary (cf. the model used in Experiment 1) could not capture this pattern.

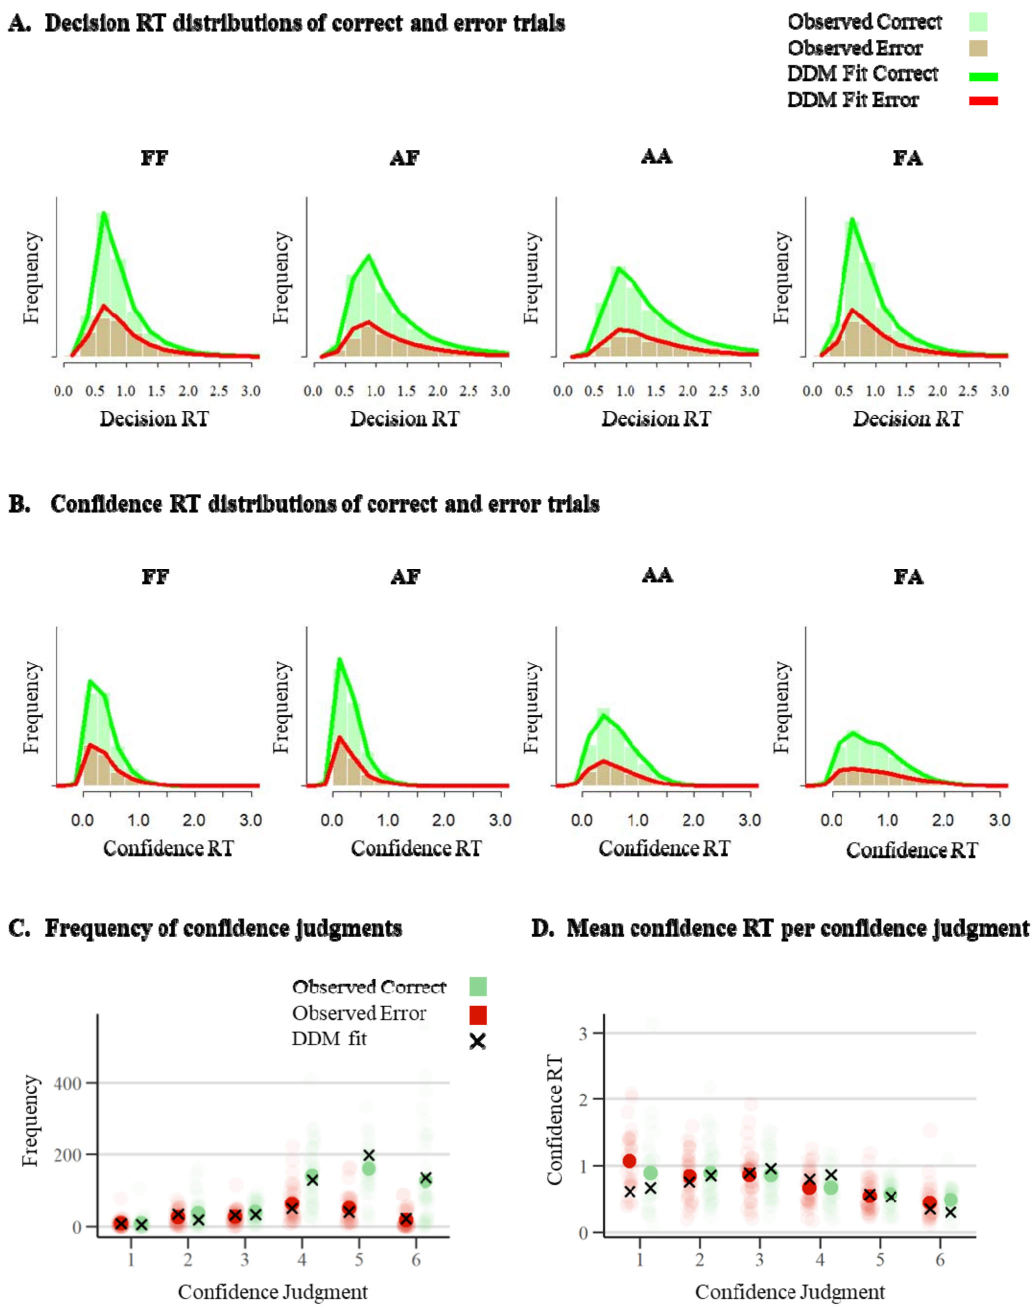

Figure 8. Distributions of reaction times (A), confidence RTs (B) and confidence (C-D) for data and model fit for Experiment 2. *Inspection of the model fit reveals that our model accurately captures the distributions in reaction times, confidence RTs and confidence seen in the data, across all four SAT manipulations.*

Model parameters

Replicating the findings of Experiment 1, we again observed that estimated decision boundaries were affected by choice SAT, $F(1, 39) = 56.27$, $p < .001$, $\eta_p^2 = .66$. However, we also found a significant effect of the confidence SAT, $F(1, 39) = 40.46$, $p < .001$, $\eta_p^2 = .02$, and an interaction between both, $F(1, 39) = 8.28$, $p < .001$, $\eta_p^2 = .07$. Follow-up paired *t*-tests showed that the effect of the decision RT instructions on decision boundary separation was significant both when confidence SAT was to be accurate, $t(39) = -7.14$, $p < .001$, as well as when confidence SAT was to be fast, $t(39) = -6.69$, $p < .001$. As expected, choice boundaries were modulated by choice SAT instructions (Figure 9), although the effect also seemed to scale, to a lesser extent, with confidence SAT.

Second, we analyzed the confidence boundaries (see Figure 9B-C). Notice that, different from Experiment 1, both confidence boundaries were allowed to vary independently and thus are analyzed separately. First, analysis of the upper confidence boundary revealed a significant effect of the confidence SAT, $F(1, 39) = 12.80$, $p < .001$, $\eta_p^2 = .25$, but not of the choice SAT, $F(1, 39) = 0.08$, $p = .779$, $\eta_p^2 = .002$, nor was there an interaction, $F(1, 39) = 3.01$, $p = .086$, $\eta_p^2 = .07$. In line with our hypothesis, participants increased the upper confidence boundary when instructed to make careful confidence judgments compared to when instructed to make fast confidence judgments. Interestingly, the lower confidence boundary seemed unaffected by any of our manipulations: neither confidence SAT, $F(1, 39) = 2.69$, $p = .109$, $\eta_p^2 = .06$, nor choice SAT, $F(1, 39) = 0.33$, $p = .570$, $\eta_p^2 = .008$, nor the interaction between both, $F(1, 39) = 2.96$, $p = .093$, $\eta_p^2 = .07$, was significant.

Finally, we analyzed the urgency parameters controlling the slope of the confidence boundaries. For the upper confidence boundary we observed that a strong difference in urgency depending on confidence SAT, $F(1, 39) = 22.75$, $p < .001$, $\eta_p^2 = .37$, and to a lesser extent also on choice SAT, $F(1, 39) = 5.68$, $p = .022$, $\eta_p^2 = .13$. There was no significant interaction, $F(1, 39) = 1.51$, $p = .227$, $\eta_p^2 = .04$. As can be seen in Figure 9D, instructions to make fast decisions and instructions to give fast confidence ratings led to increased urgency on the upper confidence boundary. Finally, for the lower confidence boundary we found that urgency was significantly affected by confidence SAT, $F(1, 39) = 6.37$, $p = .016$ $\eta_p^2 = .14$, but not by choice SAT, $F(1, 39) = 0.71$, $p = .405$, $\eta_p^2 = .02$, nor was there an interaction, $F(1, 39) = 0.94$, $p = .339$, $\eta_p^2 = .02$. When participants were instructed to make fast confidence judgments, this led to increased urgency for the lower confidence boundary (Figure 9E). All parameter estimates can be found in Table 2.

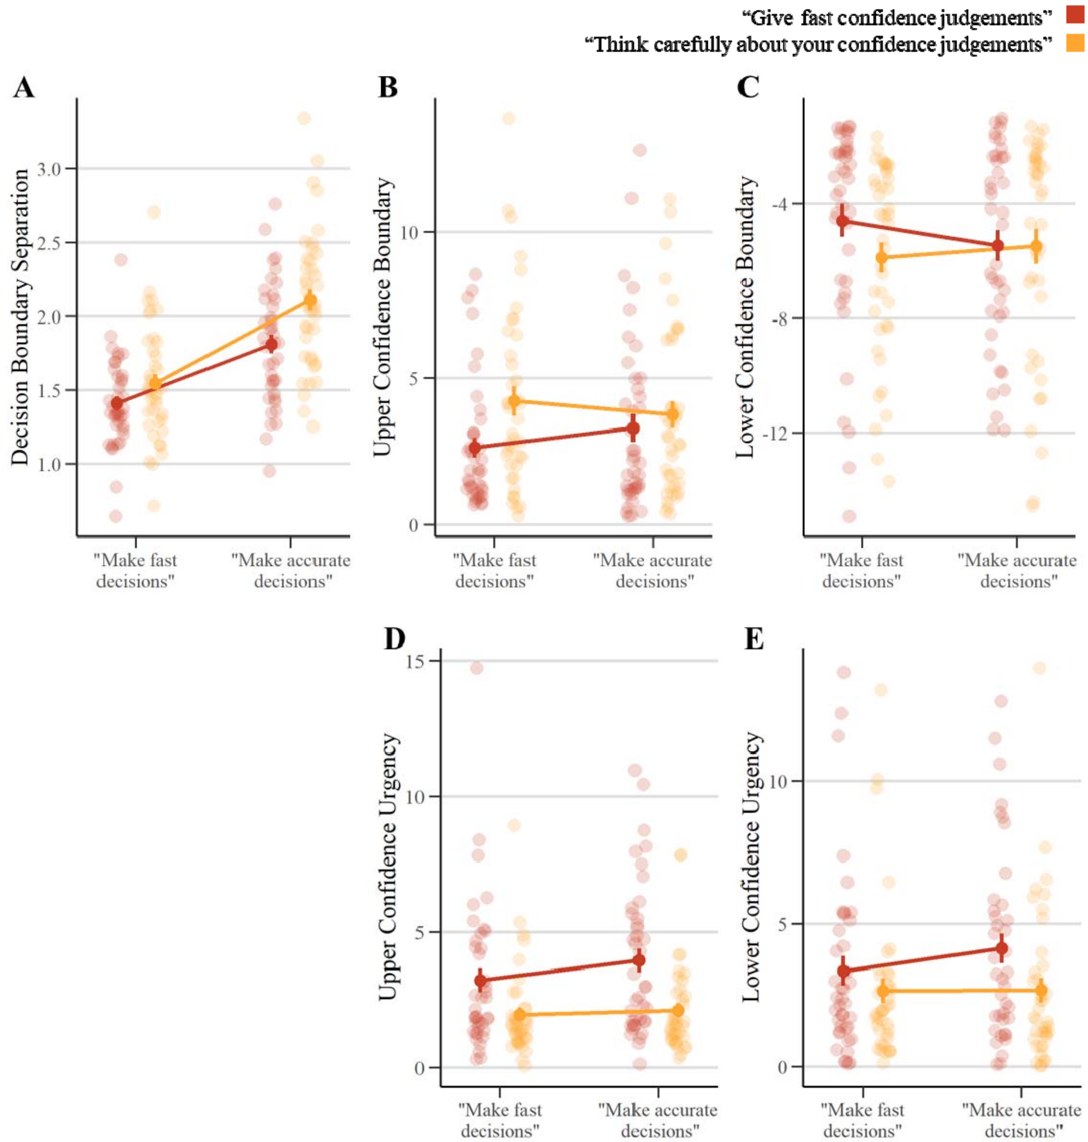

Figure 9. Influence of choice and confidence SAT on decision boundaries and confidence boundaries in Experiment 2. *Instructing participants to make fast vs accuracy choices influences estimated decision boundaries (A), but rarely influences upper or lower confidence boundaries and urgency (B-E). Instructing participants to provide fast vs careful confidence ratings influences the estimated lower but not upper confidence boundary (B-C) and urgency (D-E). Same conventions as in Figure 4.*

Discussion

The human ability to estimate and report the level of confidence in their decisions has been the central topic of much recent investigations (Rahnev et al., 2022). Despite a large number of studies examining how confidence is computed, how people decide *when* to provide a confidence rating has been unresolved. This is remarkable, because the timing of confidence judgments can be highly diagnostic about the computations underlying decision confidence (Moran et al., 2015). In the current work, we propose to model the stopping rule for confidence judgments using an accumulation-to-bound mechanism similar to the one underlying decisions. Given that decision boundaries are believed to be under strategic control, this account predicts that confidence boundaries should also be under strategic control. We investigated this prediction by providing participants with different instructions regarding the tradeoff between speed and accuracy, both for decisions and for confidence judgments. In two experiments, we found that participants made faster and less accurate decisions when instructed to favor speed over accuracy, and that they made faster confidence judgments when instructed to favor speed over careful deliberation. Although the effects on average confidence were subtle or even absent, in both experiments the relation between confidence and accuracy (cf. confidence resolution) was clearly stronger when participants were more cautious in their confidence ratings. When modeling these data with an extension of the DDM with additional confidence boundaries for post-decision processing, results were as expected: SAT instructions about the decision influenced decision boundaries, while SAT instructions about confidence influenced confidence boundaries. Our findings have important consequences for the field of decision confidence, as they shed light on the importance of considering the dynamics of confidence RTs when investigating the computations underlying decision confidence.

The stopping rule for confidence is under strategic control

Previous work investigating the dynamics of decision confidence has mostly focused on explaining variations in decision confidence, with less focus on the speed with which confidence reports are given. The most common approach is to simply have a free parameter that controls the duration of post-decision evidence accumulation (Hellmann et al., 2021; Pleskac & Busemeyer, 2010; Yu et al., 2015). Such an implementation, however, predicts that confidence judgments will always be provided at the same post-decision latency. This prediction is at odds with the observation that confidence RTs show a right-skewed distribution that is also characteristic of decision RTs. Confidence boundaries, on the other hand, provide a plausible mechanism for the stopping rule of post-decision evidence accumulation. A notable exception to this critique is a study by Moran and colleagues (2015) who proposed a single confidence boundary that collapses slowly over time, with the level of confidence being determined by the height of the boundary at the time of crossing. Their model has three free parameters that control the termination of post-decision processing: i) a parameter that controls the initial height of the boundary, similar to a_2 in our model, ii) a parameter controlling the collapse rate, similar to u in our model, and iii) a parameter controlling the time before the first collapse, for which there is no equivalent in our model. The proposal from Moran and colleagues further differs from ours because it does not consider a lower confidence boundary; the model provides a confidence rating of .5 if the collapsing confidence boundary has not been reached before it collapsed to .5. Contrastingly, in our implementation the model features both an upper and a lower confidence boundary, which can be mapped onto high versus low confidence (Experiment 1), but critically can also account for changes-of-mind by further dividing the area in between the two confidence boundaries (Experiment 2). Most importantly, although Moran and colleagues also considered speed-accuracy tradeoff in the decision, they did not investigate whether a similar tradeoff exists for confidence judgments and whether this can be accounted for within their model.

Speed-accuracy tradeoffs can be implemented in accumulation-to-bound models via two different mechanisms: changing the overall height of the boundary or changing the collapse rate of the boundary over time. Previous work in decision making has unraveled that instructions regarding the tradeoff between speed and accuracy for the decision tend to modulate the height of the decision boundary, while not affecting urgency (Katsimpokis et al., 2020). Reversely, when providing participants with a response deadline (e.g. respond within 1s) data are best accounted for by a slowly collapsing decision boundary (Katsimpokis et al., 2020; Murphy et al., 2016). In theory, the same two mechanisms could be used to implement speed-accuracy tradeoffs for confidence judgments. However, in our experiments where confidence SAT was modulated by means of instructions, given the evidence cited above, it was expected that participants would change the height of the confidence boundary, while leaving the urgency constant. In both our experiments, there was clear evidence that participants changed the height of the confidence boundary in response to SAT instructions. Results were more mixed concerning urgency. In Experiment 1, we did not observe any difference between conditions in terms of urgency, suggesting that participants selectively changed the height of the confidence boundaries but not the slope. In Experiment 2, however, we found that in response to instructions requiring fast confidence responses, participants also increased the level of urgency for both the upper and the lower confidence boundary. Given that both experiments were identical to each other except for the number of confidence options, this suggest that complexity of the design (i.e. arbitrating between high and low confidence versus arbitrating between six fine-grained levels of confidence) is an important factor determining the specificity with which these manipulations have an effect. As already noted, however, parameter recovery for urgency was rather low, suggesting these values should be interpreted with great caution. Nevertheless, in addition to further unravelling the role of design complexity, future work might also investigate the influence of providing a hard deadline for confidence judgments (e.g. you have to provide a confidence rating within 1s) on the confidence boundary and associated urgency signal.

Characteristics of post-decision processing

If confidence can be understood as an accumulation-to-bound signal, it follows that the reported level of confidence should depend on the height of the confidence boundary. Similar to how decreasing the decision boundary induces faster RTs and less accurate responses, it follows that decreasing the confidence boundaries should induce faster confidence RTs and lower confidence. Contrary to this, we did not observe a clear influence of confidence SAT on average confidence despite a clear difference in the height of the confidence boundary. As can be seen in Table 2, the extended DDM explained these data by assuming that decreasing the confidence boundaries was associated with increased post-decision drift rates. Future work might examine whether this prediction holds in post-decision centro-parietal EEG signals, which are thought to reflect the post-decision accumulation-to-bound signal (Desender et al., 2021). Although we did not find an effect on average confidence, there was a clear effect on confidence resolution: the relation between confidence and accuracy was much stronger when participants increased the confidence boundary. This finding could be anticipated, because increasing the confidence boundaries effectively requires collecting more post-decision evidence before reporting confidence, i.e. making a more informed confidence judgment. This finding adds to a number of reports showing that measures of metacognitive accuracy critically depend on the timing of confidence reports (Rosenbaum et al., 2022; Yu et al., 2015).

Inspection of the estimated model parameters in Table 2 reveals an interesting difference in magnitude between the non-decision component associated with the decision, *Ter*, and that associated with the confidence report, Ter_2. In line with the literature, values of *Ter* are in the range of .4s - .5s on average, suggesting that this is the time participants spend on processes unrelated to the actual decision (e.g. stimulus processing, motor components). These estimates are by definition positive. Contrary to this, values of Ter_2 are very low for Experiment 1 and even negative for Experiment 2. Although negative values of Ter_2 might seem counterintuitive at first, they suggest that "post-decision" processing already initiates prior to the execution of the decision motor response (e.g., Verdonck et al., 2020). In line with this observation, there is some work that has suggested that pre-choice or peri-choice neural signals contribute to the computation of decision confidence (Feuerriegel et al., 2022; Gherman & Philiastides, 2015; Murphy et al., 2015).

The computations underlying decision confidence

Humans differ in the extent to which they can accurately judge the accuracy of their decisions via confidence judgments (i.e. metacognitive accuracy). Such variability is not without consequences, as metacognitive accuracy has been associated with political extremism (Rollwage et al., 2018), anxiety and depression (Rouault et al., 2018). Although this variability is widely accepted, there is much debate regarding the best way to quantify, so-called, metacognitive accuracy (Fleming & Lau, 2014). In recent years, there has been an increase in the number of studies investigating the computations underlying decision confidence, and accompanying proposals of novel ways to quantify metacognitive accuracy (Dayan, 2022; Desender et al., 2022; Guggenmos, 2022; Mamassian & Gardelle, 2021; Maniscalco & Lau, 2012). To our knowledge, none of these proposals takes the dynamics of confidence RTs into account. This is of critical importance, though, given our demonstration that different strategies in the reporting of decision confidence (i.e., fast versus carefully) have a consistent influence on metacognitive accuracy: the relation between confidence and accuracy is stronger when participants are more cautious in their reporting of confidence. Thus, measures that do not take these dynamics into account risk to confound variability in the caution with which confidence judgments are provided with variability in genuine metacognitive accuracy. This cautionary tale bears close resemblance to a previous study in which we showed that static measures of metacognition confound caution with metacognition because they do not take the height of the decision boundary into account (Desender et al., 2022). Although this might sound very similar to the conclusion of the current work (i.e. the height of the boundary influences confidence) the underlying mechanism is very different. In our previous work, the reasoning is that impulsive decisions made with a low decision boundary lead to many premature errors that are easy to detect. Although detecting premature errors is obviously an act of metacognition, the fact that these errors are easier to detect should be taken into account when quantifying metacognition. In the current work, the reasoning is that (all) errors are easier to detect when the confidence boundaries are increased, because more evidence is accumulated to inform about the level of confidence (see also Yu et al., 2015).

Conclusion

We demonstrated that the stopping rule for confidence judgments is well described by a set of confidence boundaries that terminate post-decision processing. Importantly, just like with decision boundaries, these confidence boundaries are under strategic control, and can be increased or decreased by instructing participants to make very careful or very fast confidence judgments, respectively. When prompted to be more careful about their confidence judgments, individuals tend to be slower but metacognitively more accurate when reporting their confidence. Taken together, these results highlight the importance of taking into account the dynamics of confidence computation to unravel its underlying mechanisms.

Acknowledgments

This research was supported by a starting grant from the KU Leuven (STG/20/006) and the Francqui Foundation (PXF-D8830), and a CELSA grant from the KU Leuven (CELSA/21/010).

References

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: Reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics*, *73*(2), 640–657. https://doi.org/10.3758/s13414-010-0049-7.Acquisition

Balsdon, T., Wyart, V., & Mamassian, P. (2020). Confidence controls perceptual evidence accumulation. *Nature Communications*, *11*(1). https://doi.org/10.1038/s41467-020-15561-w

Bates, D. M., Maechler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, *67*(1), 1–48.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. https://doi.org/10.1037/0033-295X.113.4.700

Bogacz, R., Hu, P. T., Holmes, P., & Cohen, J. D. (2010). Do humans produce the speed-accuracy trade-off that maximizes reward rate? *Quarterly Journal of Experimental Psychology*, *63*(5), 863–891. https://doi.org/10.3174/ajnr.A1256.Functional

Bogacz, R., Wagenmakers, E.-J., Forstmann, B. U., & Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. *Trends in Neurosciences*, *33*(1), 10–16. https://doi.org/10.1016/j.tins.2009.09.002

Dayan, P. (2022). *Metacognitive Information Theory*. 1–14.

Desender, K., Boldt, A., & Yeung, N. (2018). Subjective confidence predicts information seeking in decision making. *Psychological Science*, *29*(5), 761–778.

Desender, K., Murphy, P., Boldt, A., Verguts, T., & Yeung, N. (2019). A Postdecisional Neural Marker of Confidence Predicts Information-Seeking in Decision-Making. *The Journal of Neuroscience*, *39*(17), 3309–3319.

Desender, K., Ridderinkhof, K. R., & Murphy, P. R. (2021). Understanding neural signals of post- decisional performance monitoring□: An integrative review. *ELife*, 1–21.

Desender, K., Vermeylen, L., & Verguts, T. (2022). Dynamic influences on static measures of metacognition. *Nature Communications*, *2022*, 1–30. https://doi.org/10.1038/s41467-022-31727-0

Donner, T. H., Siegel, M., Fries, P., & Engel, A. K. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Current Biology*, *19*(18), 1581–1585. https://doi.org/10.1016/j.cub.2009.07.066

Feuerriegel, D., Murphy, M., Konski, A., Mepani, V., Sun, J., Hester, R., & Bode, S. (2022). Electrophysiological correlates of confidence differ across correct and erroneous perceptual decisions. *NeuroImage*, *259*(November 2021), 119447. https://doi.org/10.1016/j.neuroimage.2022.119447

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*(July), 1–9. https://doi.org/10.3389/fnhum.2014.00443

Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. (2010). Relating introspective accuracy to individual differences in brain structure. *Science*, *329*(5998), 1541–1543. https://doi.org/10.1126/science.1191883

Folke, T., Jacobsen, C., Fleming, S. M., & De Martino, B. (2016). Explicit representation of confidence informs future value-based decisions. *Nature Human Behaviour*, *1*(0002), 1–8. https://doi.org/10.1038/s41562-016-0002

Fox, J., & Weinberg, S. (2019). *An R Companion to Applied Regression* (Third Edit). Sage.

Gherman, S., & Philiastides, M. G. (2015). Neural representations of confidence emerge from the process of decision formation during perceptual choices. *NeuroImage*, *106*, 134–143. https://doi.org/10.1016/j.neuroimage.2014.11.036

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–561. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Guggenmos, M. (2022). Reverse engineering of metacognition. *ELife*, *11*, 1–29. https://doi.org/10.7554/eLife.75420

Hellmann, S., Zehetleitner, M., & Rausch, M. (2021). Simultaneous modeling of choice, confidence and response time in visual perception. *BioxRxiv*, 1–41.

Katsimpokis, D., Hawkins, G. E., & van Maanen, L. (2020). Not all Speed-Accuracy Trade-Off Manipulations Have the Same Psychological Effect. *Computational Brain & Behavior*, *3*(3), 252–268. https://doi.org/10.1007/s42113-020-00074-y

Mamassian, P., & Gardelle, V. De. (2021). *Modeling Perceptual Confidence and the Confidence Forced-Choice Paradigm*.

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, *21*(1), 422–430. https://doi.org/10.1016/j.concog.2011.09.021

Moran, R., Teodorescu, A. R., & Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cognitive Psychology*, *78*, 99–147. https://doi.org/10.1016/j.cogpsych.2015.01.002

Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, *7*(May), 13526. https://doi.org/10.1038/ncomms13526

Murphy, P. R., Robertson, I. H., Harty, S., & O'Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *ELife*, *4*(DECEMBER2015), 1–23. https://doi.org/10.7554/eLife.11946

O'Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. *Nature Neuroscience*, *15*(12), 1729–1735. https://doi.org/10.1038/nn.3248

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pereira, M., Megevand, P., Tan, M. X., Chang, W., Wang, S., Rezai, A., Seeck, M., Corniola, M., Momjian, S., Bernasconi, F., Blanke, O., & Faivre, N. (2021). Evidence accumulation relates to perceptual consciousness and monitoring. *Nature Communications*, *12*(1), 1–11. https://doi.org/10.1038/s41467-021-23540-y

Pereira, M., Perrin, D., & Faivre, N. (2022). A leaky evidence accumulation process for perceptual experience. *Trends in Cognitive Sciences*, *26*(6), 451–461. https://doi.org/10.1016/j.tics.2022.03.003

Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, *117*(3), 864–901. https://doi.org/10.1037/a0022399

Rafiei, F., & Rahnev, D. (2021). Qualitative speed-accuracy tradeoff effects that cannot be explained by the diffusion model under the selective influence assumption. *Scientific Reports*, *11*(1), 1–19. https://doi.org/10.1038/s41598-020-79765-2

Rahnev, D., Balsdon, T., Charles, L., de Gardelle, V., Denison, R., Desender, K., Faivre, N., Filevich, E., Fleming, S. M., Jehee, J., Lau, H., Lee, A. L. F., Locke, S. M., Mamassian, P., Odegaard, B., Peters, M., Reyes, G., Rouault, M., Sackur, J., … Zylberberg, A. (2022). Consensus Goals in the Field of Visual Metacognition. *Perspectives on Psychological Science*, 174569162210756. https://doi.org/10.1177/17456916221075615

Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model□: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*, 873–922.

Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, *461*(7261), 263–266. https://doi.org/10.1038/nature08275

Rollwage, M., Dolan, R. J., & Fleming, S. M. (2018). Metacognitive Failure as a Feature of Those Holding Radical Beliefs. *Current Biology*, *28*(24), 4014-4021.e8. https://doi.org/10.1016/j.cub.2018.10.053

Rosenbaum, D., Glickman, M., Fleming, S. M., & Usher, M. (2022). The Cognition/Metacognition Trade-Off. *Psychological Science*, *33*(4), 613–628. https://doi.org/10.1177/09567976211043428

Rouault, M., Seow, T., Gillan, C. M., & Fleming, S. M. (2018). Psychiatric Symptom Dimensions Are Associated With Dissociable Shifts in Metacognition but Not Task Performance. *Biological Psychiatry*, *84*(6), 443–451. https://doi.org/10.1016/j.biopsych.2017.12.017

Tuerlinckx, F., Maris, E., Ratcliff, R., & De Boeck, P. (2001). A comparison of four methods for simulating the diffusion process. *Behavior Research Methods, Instruments, and Computers*, *33*(4), 443–456. https://doi.org/10.3758/BF03195402

Van Den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert,

D. M. (2016). A common mechanism underlies changes of mind about decisions and confidence. *ELife*, *5:e12192*, 1–21. https://doi.org/10.7554/eLife.12192

Verdonck, S., Loossens, T., & Philiastides, M. G. (2020). The Leaky Integrating Threshold and Its Impact on Evidence Accumulation Models of Choice Response Time (RT). *Psychological Review*, *July*. https://doi.org/10.1037/rev0000258

Yu, S., Pleskac, T. J., & Zeigenfuse, M. D. (2015). Dynamics of Postdecisional Processing of Confidence. *Journal of Experimental Psychology: General*, *144*(2), 489–510.

Supplementary Materials

Table S1. Parameter recovery for the extended DDM used in Experiment 1.

					95% CI		
			N_{trials} = 10.000				
Model	Predictor	Estimate	SE	t	LL	UL	p
a	β_0	0.03	0.06	0.52	-0.09	0.14	.606
	$\beta 1$	0.99	0.03	30.96	0.92	1.05	< .001
a_2	β_0	0.40	0.31	1.32	-0.22	1.02	.196
	$\beta 1$	0.78	0.09	8.89	0.60	0.96	< .001
u	β_0	1.30	0.16	7.92	0.97	1.63	< .001
	$\beta 1$	0.19	0.06	3.40	0.08	0.30	.002
ter	β_0	0.01	0.00	1.82	-0.00	0.02	.077
	$\beta 1$	0.98	0.01	116.55	0.96	1.00	< .001
ter_2	β_0	0.01	0.01	0.90	-0.01	0.02	.373
	$\beta 1$	0.96	0.02	57.86	0.93	1.00	< .001
v	β_0	-0.00	0.04	-0.08	-0.09	0.08	.941
	$\beta 1$	1.02	0.02	79.41	0.97	1.06	< .001
v_2	β_0	0.35	0.30	1.16	-0.26	0.96	.255
	$\beta 1$	0.84	0.08	10.28	0.67	1.00	< .001
			N_{trials} = 180				
					95% CI		
Model	Predictor	Estimate	SE	t	LL	UL	p
a	β_0	0.15	0.11	1.35	-0.07	0.37	.187
	$\beta 1$	0.91	0.06	14.93	0.79	1.04	< .001

a_2	β_0	1.22	0.41	2.95	0.38	2.05	.005
	$\beta 1$	0.56	0.12	4.52	0.31	0.81	< .001
u	β_0	1.36	0.17	7.82	1.01	1.71	< .001
	$\beta 1$	0.13	0.05	2.70	0.03	0.23	.010
ter	β_0	0.00	0.01	0.07	-0.02	0.02	.948
	$\beta 1$	1.00	0.02	57.54	0.96	1.03	< .001
ter_2	β_0	-0.00	0.01	-0.13	-0.02	0.02	.895
	$\beta 1$	1.02	0.03	37.91	0.97	1.10	< .001
v	β_0	0.12	0.08	1.49	-0.04	0.29	.144
	$\beta 1$	0.94	0.04	22.46	0.86	1.03	< .001
v_2	β_0	1.14	0.35	3.27	0.43	1.85	.002
	$\beta 1$	0.59	0.09	6.62	0.41	0.77	< .001

943 *Table S2. Parameter recovery for the extended DDM used in Experiment 2.*

N_{trials} = 10.000							
					95% CI		
Model	Predictor	Estimate	SE	*t*	LL	UL	*p*
a	β_0	0.22	0.10	2.30	0.03	0.41	.027
	$\beta 1$	0.86	0.05	17.77	0.76	0.96	< .001
a_2_upper	β_0	1.31	1.31	1.00	-1.34	3.97	.323
	$\beta 1$	0.72	0.14	5.31	0.45	0.99	< .001
a_2_lower	β_0	2.32	1.08	2.16	0.14	4.50	.037
	$\beta 1$	0.69	0.12	5.95	0.46	0.93	< .001
u_upper	β_0	1.59	0.69	2.29	0.18	2.99	.028
	$\beta 1$	0.82	0.08	9.74	0.65	1.00	< .001
u_lower	β_0	2.74	0.99	2.77	0.74	4.74	.009
	$\beta 1$	0.65	0.12	5.52	0.41	0.89	< .001
ter	β_0	0.01	0.01	1.39	-0.01	0.03	.171
	$\beta 1$	0.98	0.02	56.85	0.95	1.02	< .001
ter_2	β_0	-0.04	0.05	-0.85	-0.13	0.05	.402
	$\beta 1$	0.78	0.08	10.21	0.62	0.93	< .001
v	β_0	0.07	0.09	0.87	-0.10	0.25	.393
	$\beta 1$	0.98	0.05	18.64	0.87	1.09	< .001
v_2	β_0	1.07	0.32	3.31	0.41	1.72	.002
	$\beta 1$	0.51	0.08	6.10	0.34	0.68	< .001

N_{trials} = 180							
					95% CI		
Model	Predictor	Estimate	SE	*t*	LL	UL	*p*

a	β_0	0.22	0.11	2.21	0.01	0.44	.040
	$\beta 1$	0.86	0.05	15.92	0.75	0.97	< .001
a_2_upper	β_0	3.02	1.16	2.60	0.67	5.38	.013
	$\beta 1$	0.57	0.12	4.57	0.32	0.82	< .001
a_2_lower	β_0	2.35	1.16	2.03	0.01	4.69	.049
	$\beta 1$	0.74	0.14	5.46	0.47	1.02	< .001
u_upper	β_0	1.17	0.55	2.13	0.06	2.28	.039
	$\beta 1$	0.86	0.07	13.24	0.73	0.99	< .001
u_lower	β_0	0.90	1.05	0.86	-1.22	3.03	.396
	$\beta 1$	0.93	0.14	6.87	0.66	1.21	< .001
ter	β_0	0.02	0.02	1.25	-0.01	0.05	0.22
	$\beta 1$	0.97	0.03	34.15	0.91	1.03	< .001
ter_2	β_0	-0.08	0.04	-1.95	-0.16	0.00	0.058
	$\beta 1$	0.77	0.07	10.64	0.62	0.92	< .001
v	β_0	0.01	0.11	0.07	-0.21	0.23	.942
	$\beta 1$	1.00	0.06	15.45	0.87	1.13	< .001
v_2	β_0	0.99	0.30	3.32	0.39	1.59	.002
	$\beta 1$	0.55	0.08	7.00	0.39	0.71	< .001

Table S3. Full results table of the analyses on the extended DDM model predictions for Experiment 1.

Experiment 1

Dependent variable = log(RT)	**Chi-square**	**df**	***p* value**
Choice SAT	63.52	1	< .001
Confidence SAT	1.26	1	.262
Coherence	563.29	2	< .001
Choice SAT x confidence SAT	9.59	1	.002
Choice SAT x coherence	20.11	2	< .001
Confidence SAT x coherence	3.49	2	.175
Choice SAT x confidence SAT x coherence	0.95	2	.622
Dependent variable = log(confidence RT)			
Choice SAT	0.00	1	.998
Confidence SAT	71.75	1	< .001
Coherence	140.67	2	< .001
Choice SAT x confidence SAT	5.39	1	.020
Choice SAT x coherence	4.28	2	.117
Confidence SAT x coherence	2.89	2	.236
Choice SAT x confidence SAT x coherence	1.95	2	.377
Dependent variable = accuracy			
Choice SAT	5.30	1	.021
Confidence SAT	0.75	1	.388

Coherence	85.68	2	< .001
Choice SAT x confidence SAT	0.64	1	.424
Choice SAT x coherence	14.00	2	< .001
Confidence SAT x coherence	2.30	2	.317
Choice SAT x confidence SAT x coherence	3.61	2	.164
Dependent variable = confidence			
Choice SAT	0.14	1	.708
Confidence SAT	5.15	1	.023
Coherence	92.41	2	< .001
Choice SAT x confidence SAT	1.29	1	.257
Choice SAT x coherence	3.11	2	.211
Confidence SAT x coherence	0.83	2	.660
Choice SAT x confidence SAT x coherence	2.57	2	.277

Table S4. Full results table of the analyses on the extended DDM model predictions for Experiment 2. Note that two interaction effects were omitted from the confidence model in the behavioral data (due to inflated VIF values), and these were likewise omitted here in order to be able to directly compare both fits.

Experiment 2

Dependent variable = log(RT)	**Chi-square**	**df**	***p* value**
Choice SAT	61.15	1	< .001
Confidence SAT	23.62	1	< .001
Coherence	636.23	2	< .001
Choice SAT x confidence SAT	6.89	1	.009
Choice SAT x coherence	20.50	2	< .001
Confidence SAT x coherence	8.64	2	.013
Choice SAT x confidence SAT x coherence	0.63	2	.731
Dependent variable = log(confidence RT)			
Choice SAT	11.92	1	< .001
Confidence SAT	65.47	1	< .001
Coherence	160.92	2	< .001
Choice SAT x confidence SAT	1.34	1	.247
Choice SAT x coherence	1.83	2	.400
Confidence SAT x coherence	2.50	2	.287
Choice SAT x confidence SAT x coherence	1.36	2	.507
Dependent variable = accuracy			

Choice SAT	7.95	1	.005
Confidence SAT	13.61	1	< .001
Coherence	1001.67	2	< .001
Choice SAT x confidence SAT	5.43	1	.020
Choice SAT x coherence	4.55	2	.103
Confidence SAT x coherence	4.21	2	.122
Choice SAT x confidence SAT x coherence	0.58	2	.746
Dependent variable = confidence			
Choice SAT	0.01	1	.929
Confidence SAT	4.88	2	.027
Coherence	39.25	1	< .001
Choice SAT x confidence SAT	-	-	-
Choice SAT x coherence	1.44	2	.486
Confidence SAT x coherence	10.20	2	.006
Choice SAT x confidence SAT x coherence	-	-	-