Full Title: Perceptual change-of-mind decisions are sensitive to absolute evidence magnitude

William Turner¹, Daniel Feuerriegel¹, Milan Andrejević¹, Robert Hester¹, Stefan Bode^{1,2*}

1. Melbourne School of Psychological Sciences, The University of Melbourne, Australia

2. Department of Psychology, University of Cologne, Germany

* Corresponding Author

E-mail: sbode@unimelb.edu.au (SB)

Address: Melbourne School of Psychological Sciences, The University of Melbourne, Parkville,

Victoria, 3010, Australia

Abstract

To navigate the world safely, we often need to rapidly ‘change our mind’ about decisions. Current models assume that initial decisions and change-of-mind decisions draw upon common sources of sensory evidence. In two-choice scenarios, this evidence may be ‘relative’ or ‘absolute’. For example, when judging which of two objects is the brightest, the luminance difference and luminance ratio between the two objects are sources of ‘relative’ evidence, which are invariant across additive and multiplicative luminance changes. Conversely, the overall luminance of the two objects combined is a source of ‘absolute’ evidence, which necessarily varies across symmetric luminance manipulations. Previous studies have shown that initial decisions are sensitive to both relative and absolute evidence; however, it is unknown whether change-of-mind decisions are sensitive to absolute evidence. Here, we investigated this question across two experiments. In each experiment participants indicated which of two flickering greyscale squares was brightest. Following an initial decision, the stimuli remained on screen for a brief period and participants could change their response. To investigate the effect of absolute evidence, the overall luminance of the two squares was varied whilst either the luminance difference (Experiment 1) or luminance ratio (Experiment 2) was held constant. In both experiments we found that increases in absolute evidence led to faster, less accurate initial responses and slower changes of mind. Change-of-mind accuracy decreased when the luminance difference was held constant, but remained unchanged when the luminance ratio was fixed. We show that the three existing change-of-mind models cannot account for our findings. We then fit three alternative models, previously used to account for the effect of absolute evidence on one-off decisions, to the data. A leaky competing accumulator model best accounted for the changes in behaviour across absolute evidence conditions – suggesting an important role for input-dependent leak in explaining perceptual changes of mind.

Keywords: change-of-mind, decision-making models, evidence accumulation, absolute evidence

1. Introduction

Highly successful theoretical accounts of simple decision-making processes have arisen from the idea that decisions are reached via the accumulation of noisy evidence to a threshold level (Gold & Shadlen, 2007; Ratcliff, Smith, Brown, & McKoon, 2016; Smith & Ratcliff, 2004). In line with these accounts, neural activity in humans, monkeys, rodents, and other animals has been shown to display accumulation-like ramping patterns during decision making (Gold & Shadlen, 2007; Hanks et al., 2015; Hanks & Summerfield, 2017; O’Connell, Dockree, & Kelly, 2012). Moreover, evidence accumulation models have effectively captured both the choices people make, as well as the time taken to make them, across a wide range of experimental tasks and contexts (Bogacz, Brown, Moehlis, Holmes, & Cohen, 2006). These models typically include the assumption that the evidence accumulation process terminates once a decision threshold is crossed. This implies that, once formed, decisions are necessarily acted upon without alteration. However, this is at odds with a wealth of evidence suggesting that humans and animals are able to rapidly change their minds about decisions, even as they unfold (Albantakis, Branzi, Costa, & Deco, 2012; Burk, Ingram, Franklin, Shadlen, & Wolpert, 2014; Kaufman, Churchland, Ryu, & Shenoy, 2015; Kiani, Cueva, Reppas, & Newsome, 2014; Moher & Song, 2014; Resulaj, Kiani, Wolpert, & Shadlen, 2009; van den Berg et al., 2016).

To better understand changes of mind, and decision-making in general, it is important to consider the nature of the evidence being accumulated in the decision process. In simple choices between two alternatives, the decision maker may draw upon relative and/or absolute sources of evidence. ‘Relative evidence’ is information which is invariant across symmetric changes in the magnitude or intensity of the two choice options. For example, when deciding which of two objects is the brightest, the difference in luminance between the two objects constitutes relative evidence. This is because if the luminance of each object is increased by a constant amount (i.e. an additive stimulus manipulation) the difference in luminance between them will remain the same. Similarly, the luminance ratio (luminance of stimulus A/luminance of stimulus B) is also

relative evidence; if the luminance of each stimulus is multiplied by a constant factor (i.e. a multiplicative stimulus manipulation) the ratio of the luminance values will remain the same. ‘Absolute evidence’ on the other hand is information which necessarily varies with symmetric changes in stimulus magnitude. In the above example, the overall sum of the luminance values for each stimulus constitutes absolute evidence; if the luminance of each stimulus is increased by a fixed additive or multiplicative amount then their overall sum will also increase.

In previous research, the effects of variations in relative evidence on decision-making have been well characterised. In contrast, absolute evidence has often been overlooked as a potential source of decision-relevant information, perhaps because this information is task-irrelevant when making relative judgments (i.e. it tells the decision-maker nothing about which of the two objects is brighter). Recently however, a number of studies have shown that variations in absolute evidence do affect decision-making behaviour (Hunt et al., 2012; Polanía, Krajbich, Grueschow, & Ruff, 2014; Ratcliff, Voskuilen, & Teodorescu, 2018; Teodorescu, Moran, & Usher, 2016). In particular, these studies have demonstrated that people respond faster, and often less accurately, to stimuli containing high levels of absolute evidence (i.e. brighter pairs of squares). This occurs across a wide range of experimental tasks and contexts with similar findings also reported in monkeys (Pirrone, Habiba, Hayden, Stafford, & Marshall, 2018). Moreover, normative modelling has also shown that absolute evidence sensitivity is expected under optimal decision policies (e.g., when a speed up in response time helps maximise the reward rate across choices; Marshall, 2019; Steverson, Chung, Zimmermann, Louie, & Glimcher, 2019; Tajima, Drugowitsch, Patel, & Pouget, 2019; Tajima, Drugowitsch, & Pouget, 2016). As such, there is strong support for the notion that decision-making behaviour across a range of organisms is sensitive to absolute evidence. This raises the question of whether variations in absolute evidence affect the frequency and timing of subsequent change-of-mind decisions.

An influential theory developed by Rabbitt and colleagues posits that changes of mind occur because the decision process continues to unfold even after an initial decision is made

(Rabbitt & Vyas, 1981). According to this view, if enough late-arriving evidence is accumulated against an initial decision then a change of mind occurs. Recently, a number of computational models have been developed which incorporate this notion (Albantakis & Deco, 2011; Atiya, Rañó, Prasad, & Wong-Lin, 2019; Resulaj et al., 2009). The first is an extension of the diffusion model of decision-making (Ratcliff, 1978) in which ‘post-decisional’ evidence accumulation occurs (Resulaj et al., 2009). In this model, if enough late-arriving evidence is accumulated against an initial decision, such that a second decision threshold is crossed, then a change of mind occurs. The second model is a biophysically-plausible attractor network (Albantakis & Deco, 2011). This consists of a network of simulated neurons containing two outcome-selective pools. In this model, the decision-making process relates to a transition from a symmetric state, where both pools fire at approximately the same rate, to a decision state, where one pool fires at a higher rate than the other. Changes of mind occur when the firing rate of one pool crosses a threshold level, triggering an initial decision, but the alternative pool subsequently crosses this threshold and eventually predominates. The third model is a neural circuit model which encodes decision uncertainty (Atiya et al., 2019). In this model, changes of mind are driven in part by transient activity from a ‘decision-uncertainty monitoring module’ which is partially distinct from the core decision-making circuitry.

These three models differ in a number of important ways. However, at their core all models assume that changes of mind arise out of a continuation of the initial decision process. A corollary of this assumption is that initial decisions and change-of-mind decisions must be sensitive to common sources of sensory information. Given the findings showing that initial decisions are sensitive to absolute evidence, and the assumption that changes of mind arise out of the initial decision process, it follows that change-of-mind decisions should also be sensitive to variations in absolute evidence. However, this has yet to be tested. The primary aim of the current study was therefore to test this hypothesis. To foreshadow our results, we found that change-of-mind decisions, like the decisions which precede them, are indeed sensitive to

absolute evidence. Given this, our second aim was to investigate whether this sensitivity plays out in a manner which can be accounted for by existing models.

Considering the existing change of mind models, both the attractor network model (Albantakis & Deco, 2011) and the neural circuit model (Atiya et al., 2019) are inherently sensitive to absolute evidence. However, they make opposing predictions about the effect of absolute evidence magnitude on change of mind frequency. With increased levels of absolute evidence, the attractor network model predicts that changes of mind (following both correct and incorrect responses) will be more likely to occur (see Figure 7 in Albantakis & Deco, 2011). In contrast, the neural circuit model predicts that changes of mind (following both correct and incorrect responses) will be less likely (see section 4.2). Unlike these two models, the extended diffusion model is invariant to absolute evidence. To make this model sensitive to absolute evidence, auxiliary assumptions must be adopted (Ratcliff, Voskuilen, & Teodorescu, 2018). One such assumption is that the amount of noise within the decision process scales positively with the amount of absolute evidence (Ratcliff, Voskuilen, & Teodorescu, 2018). This assumption has been adopted in previous studies of human and animal decision making (Brunton, Botvinick, & Brody, 2013; Lu & Doshier, 2008; Teodorescu et al., 2016) and is in accord with the idea that neural firing is approximately Poisson distributed (Ratcliff, Voskuilen, & Teodorescu, 2018). Alternatively, one can assume that across-trial-variability in the average rate of evidence accumulation scales positively with absolute evidence magnitude (Ratcliff, Voskuilen, & Teodorescu, 2018). Finally, a fourth model, the leaky competing accumulator (LCA) model, has recently been used to give an alternative account of the effect of absolute evidence on perceptual decisions (Ratcliff, Voskuilen, & Teodorescu, 2018; Teodorescu et al., 2016). This model is similar to the attractor network model (Albantakis & Deco, 2011), however to our knowledge it has not been used to model change-of-mind behaviour (but see Evans, Dutilh, Wagenmakers, & Maas, 2019 for a recent application of the LCA to the related behavioural phenomenon of ‘double responding’).

In the current study, we first established whether any of the three existing change-of-mind models could predict the general pattern of change-of-mind results we observed. We then explored whether two variants of the extended diffusion model, which each incorporate one of the auxiliary assumptions outlined above, as well as a variant of the LCA model, which included a change of mind mechanism, could account for our observations.

1.1 The Current Study

To investigate whether change-of-mind decisions were sensitive to absolute evidence we ran two separate experiments employing the same dynamic luminance discrimination task. In this task participants had to rapidly indicate which of two flickering greyscale squares was on average the brightest by pressing one of two buttons on a response pad. Crucially, following an initial judgement the stimuli remained on screen for a fixed duration (1s), and participants were free to change their response. To investigate the effect of absolute evidence, the absolute luminance of the two squares was manipulated (low/high), whilst one source of relative evidence was held constant. In Experiment 1, the difference in luminance between the two stimuli was held constant across the low and high absolute evidence trials (i.e. an additive stimulus manipulation). In Experiment 2, the luminance ratio was held constant (i.e. a multiplicative stimulus manipulation). For both experiments, the main question of interest was whether the frequency and timing of changes of mind would vary across low and high absolute evidence trials.

2. Materials and Methods

2.1 Participants

In both Experiment 1 and Experiment 2, 30 right-handed participants each gave written informed consent. They were each remunerated \$15 AUD for their time. In Experiment 1, one dataset was excluded from all analyses due to an unusually high number of button presses per trial (~67% of trials contained ≥ 3 button presses). In Experiment 2, no participants were excluded. For Experiment 1, the final sample consisted of 29 participants aged 18-37 years ($M = 23.07$, $SD = 4.52$, 23 female). For Experiment 2, the final sample consisted of 30 participants

aged 19-39 years ($M = 24.73$, $SD = 5.13$, 17 female). The experimental procedures were approved by the University of Melbourne ethics committee (ID 1749951).

2.2 Materials

All stimuli were presented on a Sony Trinitron Multiscan G420 CRT Monitor (Resolution 1280 x 1024 pixels; Frame Rate 75 Hz). The monitor was gamma corrected using a ColorCAL MKII Colorimeter. Responses were recorded using a Tesoro Tizona Numpad (Polling Rate of 1000 Hz). The task was coded in MATLAB 2015b using functions from the Psychophysics Toolbox Version 3.0.14 (Brainard, 1997; Kleiner et al., 2007). Whilst performing the experiment participants were seated in a darkened room with their chin resting on a chinrest ~65 cm from the screen.

2.3 Stimuli and Procedures

In both experiments, participants were familiarised with the task requirements and stimuli in the task instructions, but did not undergo training prior to the main task. In each experiment, participants performed 1000 trials of a luminance discrimination task (depicted in Fig 1). On each trial they indicated which of two dynamic stimuli, which were flickering greyscale squares (70 x 70 pixels; $\sim 2.18 \times 2.18$ degrees of visual angle), was on average the brightest. The squares were presented side-by-side at equal distance from the centre, with 70 pixels separating them horizontally. In both experiments there were two stimulus conditions: a low absolute evidence condition and a high absolute evidence condition. For Experiment 1, we employed an additive stimulus manipulation. The mean RGB values for the target (the brighter stimulus) and the non-target (the darker stimulus) in the low condition were 112 and 100, respectively. The mean RGB values in the high condition were 200 and 188 respectively. For Experiment 2, we employed a multiplicative stimulus manipulation. The mean RGB values for the target and the non-target in the low condition were 116 and 100. The mean RGB values in the high condition were 203 and 175. Note that for Experiment 2, we increased the difference in luminance between the two squares slightly to increase initial response accuracy. In both experiments, on

each frame, independent greyscale values for the two stimuli were drawn from separate Gaussian distributions centered around their respective mean values. The standard deviation of the distributions was 25.5 and the distributions were truncated at 2 standard deviations from the mean. For discussion of our stimulus manipulations with respect to Weber's law – and the role that nonlinear perceptual processing plays in explaining task behaviour – see Section 4.5.

In both experiments, the low and high absolute evidence stimulus conditions were presented randomly interleaved within the blocks. Responses were given using the 1 (left response) and 3 (right response) keys on the numpad. Participants had 800 ms from stimulus onset to make an initial response. From the time of the initial response, the stimuli remained on screen for a fixed duration of 1 s. During this time, participants were able to change their mind and give a second response. Participants were told to be as accurate as possible in their initial responses but to change their mind whenever they felt that this was necessary. Following the end of each trial, feedback (“correct”, “error” or “too slow”) was presented for 300ms. This feedback was based on the last button that participants had pressed. A red fixation dot was presented for 500 ms before stimulus presentation. Self-paced breaks were provided every 100 trials.

Fig 1. Schematic of the trial structure. Each trial began with the presentation of a red fixation dot for 500 ms. The stimuli were then presented for up to 800 ms or until a button was first pressed. The luminance of each square was updated on each frame such that the two squares flickered slightly. From the time of the initial response the post-decision period (fixed duration 1 s) began. Feedback was then presented for 300 ms in the form of (“correct” or “error”). If participants failed to respond within 800 ms of the stimuli being presented, the post-decision period was skipped and “too slow” was presented for 300 ms.

The stimuli presented in the first half of each experiment were exactly replicated in the second half of each experiment. This was done so that we could conduct double-pass agreement analyses (Lu & Doshier, 2008). The logic behind such an analysis is that when individuals make perceptual decisions, there are two broad categories of noise which can influence their responses. These are: external noise, due to factors such as fluctuations in the stimulus evidence strength across time, and internal noise, due to factors such as variability in neuronal firing rates and fluctuations in attention or motivation over time. When physically identical stimuli are presented to participants multiple times, the limiting factor with respect to the consistency of their responses will be the level of internal noise (Green, 1964). Therefore, by examining response consistency across repeated presentations of the same stimuli it is possible to estimate the average level of internal noise for a participant. Moreover, it is possible to investigate whether there are differences in the consistency of responses to stimuli containing low and high levels of absolute evidence.

2.4 Statistical Analyses

Trials in which participants failed to respond, or in which they changed their mind more than once (i.e. 3 or more button presses per trial), were excluded. Trials in which the initial response time was less than 150 ms or in which the change of mind occurred less than 50 ms after the initial response were also excluded. All analyses were conducted using mixed-effects models in R (version 3.5) via the lme4 package (version 1.1; Bates, Mächler, Bolker, & Walker, 2015). All continuous predictor variables were centered and scaled. Likelihood ratio tests were performed to compare the goodness of fit of a full model, which contained the main effect or interaction of interest, to a null model which did not include the effect of interest. Alongside the outcome of each likelihood ratio test, we also estimated group mean differences between the absolute evidence conditions for choice proportions and RTs (computed using the effects package in R; Fox et al., 2016). Equations for all full models are reported below and the model outputs are presented in the supporting information (Tables A.1-A.5). The data from both

experiments was analysed separately using identical models for each analysis. Both datasets and all code are available at <https://osf.io/sr58p/>.

2.4.1 Random Effects Structure

In all models the intercept was allowed to vary among participants. Moreover, when possible (i.e. when the model still converged), a random intercept for stimuli nested within participants was also included. The logic behind these decisions was as follows. First, responses from a single participant are likely to be correlated. For example, some participants may be more prone to changing their mind than others. Additionally, responses to physically identical stimuli are also often correlated (Ratcliff, Voskuilen, & McKoon, 2018). For example, some stimuli may be more difficult to judge than others, due to random fluctuations in the noise added in each trial. By allowing the intercept to vary among participants and among stimuli, we could account for these sources of dependence in the data. As stimuli were not repeated across participants (noise was randomly generated for each participant, but was consistent across stimulus repetitions within participants), the random intercept for stimuli was nested within participants.

Where possible, random slopes by participant were also included for the predictors of theoretical interest. In the analysis of initial accuracy, the random slope for absolute evidence condition was omitted because the model including this parameter was degenerate (as indicated by a correlation of -1 between the random effects). Moreover, no random slopes were included in the analysis of change time due to convergence issues, which were likely due to the fact that only a small subset of trials (i.e. those containing a change of mind) are included in this analysis. In the response time and choice consistency models a single random slope for the absolute evidence condition variable was included. In the change-of-mind-frequency model a random slope was also included for initial response time and for the interaction between absolute evidence and initial accuracy. These were included for exploratory purposes after the data was plotted and the possibility of an interaction between initial accuracy and absolute evidence became apparent.

2.4.2 Regression model equations

The relationship between initial decision accuracy and absolute evidence magnitude was investigated using a generalized linear mixed-effects model (GLMM; binomial family) with a logit link function:

$$\text{Accuracy} \sim \text{Condition} + \text{RT}_i + (1 + \text{RT}_i | \text{Participant}) + (1 | \text{Participant:Stimulus})$$

In the above equation, Accuracy is a binary variable (0 = error, 1 = correct), Condition is a binary variable specifying absolute evidence magnitude (0 = low, 1 = high) and RT_i is a continuous variable specifying initial response time. When conducting the likelihood ratio test, this full model was compared to a null model which did not include the main effect of condition.

The relationship between initial response time and absolute evidence magnitude was investigated using a GLMM (Gamma family) with an identity link function as recommended by Lo and Andrews (2015):

$$\text{RT}_i \sim \text{Condition} + \text{Accuracy} + (1 + \text{Condition} | \text{Participant})$$

When conducting the likelihood ratio test, this full model was compared to a null model which did not include the main effect of condition but did include the random slope for condition.

The relationship between changes of mind and absolute evidence magnitude was investigated using GLMM (binomial family) with a logit link function:

$$\text{CoM} \sim \text{Condition} * \text{Accuracy} + \text{RT}_i + (1 + \text{RT}_i + \text{Condition} * \text{Accuracy} | \text{Participant}) + (1 | \text{Participant:Stimulus})$$

In the above equation, CoM is a binary variable (0 = no change, 1 = change of mind). For the change-of-mind analyses, a likelihood ratio test was first conducted between a model containing the main effect of condition and a null model which did not include this main effect but did include a random slope for condition. Subsequently, a likelihood ratio test was conducted between a model which included the main effect of condition and interaction between initial

accuracy and condition and a model which included only the main effect (but did include the random slope for the main effects and interaction).

The relationship between change time and absolute evidence magnitude was investigated using a GLMM (Gamma family) with an identity link function:

$$\text{Change Time} \sim \text{Condition} + \text{Accuracy} + \text{RT}_i + (1 | \text{Participant})$$

Finally, the relationship between choice consistency and absolute evidence magnitude was investigated using a GLMM (binomial family) with a logit link function:

$$\text{Consistency} \sim \text{Condition} + \text{Accuracy} + \text{RT}_i + (1 + \text{Condition} | \text{Participant})$$

In the above equation, Consistency is a binary variable (0 = different responses, 1 = same responses). For this analysis, likelihood ratio test was conducted between two models which included a random slope for the main effect of condition.

3. Results

3.1 Experiment 1 (Additive stimulus manipulation)

Mixed effects regression models were fit to response time and accuracy data to test for effects of absolute evidence on initial and change-of-mind decisions. These analyses revealed that participants made their initial decisions less accurately (an estimated 12.7% reduction in the probability of making a correct initial decision; likelihood ratio test, $\chi^2(1) = 303.91, p < 2.20 \times 10^{-16}$; Fig 2A) and faster (an estimated 14 ms decrease in response time; $\chi^2(1) = 10.96, p = 9.31 \times 10^{-4}$; Fig 2B) in high compared to low absolute evidence trials (see section 4.4.1 for a discussion on the role that perceptual nonlinearities play in explaining these, and the following, behavioural effects). Participants also changed their mind less often (an estimated 1.6% reduction in the overall probability of any change of mind occurring; $\chi^2(1) = 5.12, p = .024$) and more slowly (an estimated 36 ms increase in change-of-mind response time; $\chi^2(1) = 14.48, p = 1.42 \times 10^{-4}$; see Fig 2D) in high absolute evidence trials. Moreover, there was a significant interaction between initial response accuracy and absolute evidence condition, indicating that participants corrected fewer errors but spoiled more initially correct responses in high absolute evidence trials (an estimated

8.4% decrease in the probability of correcting an error, and an estimated 2% increase in the probability of spoiling an initially correct response; $\chi^2(1) = 24.77, p = 6.46 \times 10^{-7}$). This interaction indicates that participants made less accurate change-of-mind decisions in high absolute evidence trials. This pattern was evident across the course of the experiment (Fig A.1 in the supplementary materials shows the proportion of changes of mind across time for both experiments).

Fig 2. Experiment 1 behavioral results. A) Initial decision accuracy, B) initial response time, C) change of mind proportion, and D) change-of-mind response time (CoM RT; the latency of the change-of-mind response relative to the initial response) across low (yellow) and high (blue) absolute evidence conditions. Solid lines indicate correct initial responses and dashed lines indicate incorrect initial responses. Error bars indicate standard errors of the mean (SEM). The gray dots in sections A and C represent data from individual participants. Note that in section C the interaction between initial accuracy and absolute evidence condition is somewhat obscured as participants were less accurate in the high condition to begin with; the interaction can be seen more clearly when changes of mind are displayed as a proportion of the number of errors and correct responses separately (see Fig 4).

3.2 Experiment 2 (Multiplicative stimulus manipulation)

Participants responded less accurately (an estimated 3% reduction in the probability of making a correct initial decision; likelihood ratio test, $\chi^2(1) = 36.47, p = 1.55 \times 10^{-9}$; Fig 3A) and faster (an estimated 10 ms decrease in response time; $\chi^2(1) = 8.18, p = 0.004$; Fig 3B) in high compared to low absolute evidence trials. There was no significant difference in the proportion of changes of mind across absolute evidence conditions (an estimated 0.1% reduction in the probability of any change of mind occurring; $\chi^2(1) = 0.18, p = .67$), and no evidence of an interaction between initial response accuracy and absolute evidence condition (an estimated 1.4% decrease in the probability of correcting an error, and an estimated 0.1% decrease in the probability of spoiling an initially correct response; $\chi^2(1) = 0.04, p = 0.85$). However, changes of mind were significantly slower in the high absolute evidence trials (an estimated 42 ms increase in change-of-mind response time; $\chi^2(1) = 24.59, p = 7.12 \times 10^{-7}$; see Fig 3D). Change-of-mind latency here shows a different pattern to in experiment 1. In particular, in experiment 1 corrected errors were slower than spoilt responses. However, in experiment 2, corrected errors were more broadly distributed than spoilt responses. This difference across experiments is likely due to the fact that in experiment 2 the stimuli were slightly easier to discriminate (see section 2.3), so initial responses were more accurate.

Fig 3. Experiment 2 behavioral results. A) Initial decision accuracy, B) initial response time, C) change of mind proportion, and D) change-of-mind response time (i.e. the latency of the change-of-mind response relative to the initial response) across low (yellow) and high (blue) absolute evidence conditions. Solid lines indicated correct initial responses and dashed lines indicate incorrect initial responses. Error bars indicate standard errors of the mean (SEM). The gray dots in sections A and C represent data from individual participants.

3.3 Choice consistency analysis (Experiments 1 & 2)

For both experiments we also conducted a double-pass agreement analysis to investigate whether absolute evidence magnitude was related to the consistency of participants' responses across exact repetitions of the stimuli. To enable this, the stimuli in the first half of each experiment were exactly replicated in the second half of each experiment (see Methods). The purpose of these analyses was to examine the ratio of external (i.e. stimulus driven) to internal variability within the decision process – with the aim of better informing our understanding of the participants' decision process(es) and further constraining the computational models. To

perform this analysis, logistic mixed effects regression was used to predict whether participants would make the same or different responses across stimulus repetitions (coded as 1 or 0), and whether this was influenced by the absolute evidence condition.

The full model (which included the main effect of absolute evidence condition) fit the data significantly better than the null model for Experiment 1 (an estimated 5.1% reduction in the probability of repeating a response; $\chi^2(1) = 5.36, p = 0.021$) but not for Experiment 2 (an estimated 1.5% reduction in the probability of repeating a response; $\chi^2(1) = 2.19, p = 0.14$). This suggests that the additive stimulus manipulation has a larger effect on choice consistency than the multiplicative manipulation (in which evidence ratios are conserved). In section 4.5.5 below, we examine whether these changes in choice-consistency can be accounted for within a formal computational framework.

Fig 4. Changes of mind as a proportion of initial response type (i.e. correct and incorrect initial responses) for both experiments. The dashed bars represent the number of corrected errors as a proportion of the total number of initially incorrect responses made in each stimulus condition. The solid

bars display the number of spoilt responses (i.e. changes from a correct response to an incorrect response) as a proportion of the total number of initially correct responses. Error bars represent SEM. Gray dots represent data from individual participants.

3.4 Summary of results

The analyses above demonstrate that both initial decisions and subsequent change-of-mind decisions were affected by variations in absolute evidence magnitude. Across both experiments, we found that initial decisions were faster and less accurate on high absolute evidence trials. We also found that with an additive stimulus manipulation, change-of-mind decisions were less accurate. However, with a multiplicative stimulus manipulation the accuracy of change-of-mind decisions was unaffected. Finally, in direct contrast to the initial response time effects, we found that change-of-mind decisions were consistently slower on high absolute evidence trials across both experiments.

4. Computational modelling

Following the novel observation that change-of-mind decisions were sensitive to absolute evidence magnitude, we sought to account for this sensitivity within a formal modelling framework. Below, we first briefly demonstrate that all of the existing change-of-mind models ‘out of the box’ (i.e. with no additional modifications) cannot account for the current findings. Then, we examine whether three additional models (two modified DDMs and an extended LCA model) which have all been used to account for the effect of absolute evidence on one-off perceptual decisions, can account for the current observations.

Note, that for the current analyses we have restricted our focus to models which have previously been used to account for changes of mind or the effect of absolute evidence on one-off perceptual decisions. However, for discussions concerning the role that confidence (and associated models) might play in accounting for the current findings see sections 5.3 and 5.4.

4.1 The unmodified extended DDM (Resulaj et al., 2009)

As we mentioned above, the extended DDM in its original form is a purely relative model (i.e. only has access to evidence differences, not absolute values). As such, it cannot account for any of the effects of absolute evidence which we have observed. In Section 4.4. below, we examine whether variants of this model with additional modifications are able to capture the current findings.

4.2 Attractor network model (Albantakis & Deco, 2011)

Predictions for the attractor network model regarding the effect of absolute evidence on the frequency of changes of mind were derived in previous work by Albantakis and Deco (2011; see their Fig 7). In general, this model predicts that more changes of mind will occur (following both incorrect and correct initial responses) with increased absolute evidence. This is not consistent with the interaction between absolute evidence and initial response accuracy on the proportion of changes of mind which we observed in experiment 1. In particular, this model cannot explain the decrease in the proportion of corrected errors that occurs with higher absolute evidence. As such, this model does not provide a satisfactory account for the observed effect of absolute evidence on change-of-mind decisions (see section 5.7 for discussion of additional modifications which could be considered).

4.3 Neural circuit model (Atiya et al., 2019)

To derive predictions for the neural circuit model, we simulated this model across varying levels of absolute evidence and relative evidence strength, defined in the model as evidence quality (see Fig 5). Overall, we found that this model predicts fewer changes of mind (following both incorrect and correct initial responses) with higher absolute evidence. Again, this is not consistent with the interaction between absolute evidence and initial response accuracy which we observed in experiment 1. In particular, this cannot account for the increase in the number spoiled responses which we observed. As such, this model also does not provide a satisfactory account of the observed data (see section 5.7 for discussion of additional modifications which could be considered).

Fig 5. Simulated results for the neural circuit model. In these simulations the amount of absolute evidence was varied across three levels (low, medium, high) by changing the μ_0 parameter ($\mu_0 = 20$ for low, $\mu_0 = 30$ for medium, $\mu_0 = 40$ for high). All other parameter values were taken from Atiya et al. (2019) and were kept constant across simulations. The medium absolute evidence simulation is therefore a direct reproduction of the model simulations run in the original paper. In this figure the x-axis represents the level of relative evidence (in this case the evidence difference), with lower values indicate decreased relative evidence. Dashed lines indicate spoiled responses and solid lines indicate corrected errors. Simulations were run using the code provided at https://github.com/nidstigator/uncertainty_com_modelling. We simulated 8000 trials per evidence quality level.

4.4 Two modified DDMs (c.f. Ratcliff et al., 2018)

As mentioned above, two auxiliary assumptions have recently been proposed which allow absolute evidence sensitivity to be accounted for within the framework of the diffusion model (Ratcliff, Voskuilen, & Teodorescu, 2018). We therefore investigated whether two novel versions of the extended diffusion model, which each incorporate one of these assumptions, were able to account for the observed data. One model included the assumption that within-trial variability in the decision process differs across absolute evidence conditions. This will be referred to as the “sigma model” as assumptions were made regarding the sigma parameter, which specifies the degree of within-trial variability. The alternative model included the

assumption that across-trial variability in the decision process differs across absolute evidence conditions. This model will be referred to as the “eta model” as assumptions were made regarding the eta parameter, which specifies across-trial variability in the rate of evidence accumulation.

4.4.1 DDM model specifics

In both the sigma and eta models, the drift rate (i.e. the average rate of evidence accumulation) was allowed to vary between low and high absolute evidence trials. This was to account for the possibility of Weber-like scaling with our stimulus manipulation. For our task, Weber-like scaling (i.e. a compressive nonlinear transformation of perceptual evidence) would result in a smaller perceived difference in luminance between the two stimuli in the high absolute evidence condition, compared to the low absolute evidence condition (particularly in Experiment 1 where evidence ratios were not conserved). This is likely a key reason as to why participants made less accurate decisions in the high absolute evidence condition. In the DDM, the drift rate parameter represents the amount of relative evidence (i.e. the perceived difference in luminance between the two squares). We therefore let this parameter vary across the absolute evidence conditions to account for possible differences in the perceived amount of relative evidence across conditions. In section 4.4.3. we discuss what the estimated drift-rates for each model tell us about the relationship between objective stimulus values and perceived stimulus representations in our task.

In the sigma model, the degree of within-trial variability in the decision process was allowed to vary across low and high absolute evidence trials, whilst across-trial variability in drift rate was kept constant. In the eta model, the degree of across-trial variability in drift rate was allowed to vary whilst within-trial variability was kept constant.

In the sigma model, the degree of within-trial variability in the decision process was allowed to vary across low and high absolute evidence trials, whilst across-trial variability in drift

rate was kept constant. In the eta model, the degree of across-trial variability in drift rate was allowed to vary whilst within-trial variability was kept constant.

In previous modelling work, variation in the drift rate, eta, and sigma parameters across absolute evidence levels was tightly constrained. In particular, the parameter values were directly determined from the underlying stimulus luminance values (Ratcliff, Voskuilen, & Teodorescu, 2018; Teodorescu et al., 2016). In the current study, whilst we adopted the overarching assumption that these parameters varied across absolute evidence conditions, we did not constrain this variation to be a function of the underlying stimulus luminance values. In principle, this affords the models a greater (and potentially unreasonable) degree of flexibility. However, we believe that allowing these models to be maximally flexible helps rule out any concern that poor model fits are simply due to the specific nature of the constraints being put on the condition varying parameters.

4.4.2 DDM model fitting

Both the sigma and eta models were fit to initial response proportions, initial response time quantiles (0.1 0.3 0.5 0.7 0.9), change-of-mind proportions (proportion corrected errors and proportion spoilt responses) and change-of-mind latency quantiles simultaneously. This was carried out in MATLAB with custom code which implemented a discrete approximation of the extended diffusion model (6.667 ms timesteps, 1,000,000 trials per iteration). We adopted the simplifying assumption that the non-decision time was the same for initial and change-of-mind responses. In fitting the models, we collapsed across left and right responses. Hence, the starting point parameter was fixed to half the boundary separation parameter. Initially, we included across-trial-variability in starting point in both models; however, this resulted in a number of the parameter estimates converging to the limits of the parameter space. We therefore omitted this assumption from the final models.

Group averaged response proportions and vincentised response time quantiles for both initial responses and change-of-mind responses, were used to fit the models. We note that in

certain contexts it is more appropriate to consider individual level data (Liew, Howe, & Little, 2016). However, it has consistently been shown that for similar experimental designs, parameter estimates obtained from group-averaged data are closely matched to the average of parameters estimates obtained on the individual level (Ratcliff & McKoon, 2008; Ratcliff, Thapar, & McKoon, 2001, 2003, 2004). Moreover, given that changes of mind were relatively rare, we were concerned that individual level measures of change-of-mind timing and frequency would not yield precise and reliable model estimates. When fitting both initial and change-of-mind responses, the discrepancy between the data and model predictions was quantified as the root mean squared error between actual and simulated data. A simplex function (Nelder & Mead, 1965) was used to minimize this value. All code used to simulate and fit the models is available at <https://osf.io/sr58p/>.

4.4.3 DDM modelling results

Both models fit the initial responses well (Fig 6). In particular, both could recreate the qualitative pattern of responding across the stimulus conditions in each experiment (i.e. faster and less accurate responses in the high, compared to low, absolute evidence conditions).

Fig 6. Model fits for the initial responses. Plots A) and B) show the group-averaged data for initial responses in experiment 1, as well as the predictions from A) the sigma model and B) the eta model. Plots C) and D) show the group-averaged data for initial responses in experiment 2, as well as the predictions from the C) sigma model and D) eta model. In all plots, response proportions are plotted on the x-axis and response time quantiles are plotted on the y-axis. The hollow symbols denote the empirical data, and the solid symbols denote model predictions. Yellow data points are used to represent data from the low absolute evidence condition and blue data points are used to represent data from the high absolute evidence condition. Circular symbols denote correct responses and square symbols denote incorrect responses.

Interestingly, whilst the two models captured the pattern of responding across conditions relatively well, they both predicted that correct responses would be slightly faster than error responses, when in fact errors tended to be slightly faster than correct responses. We note that, with the same stimulus manipulation, this pattern of responding was also observed by Ratcliff et al. (2018), and that this feature was also not captured in their model fits. This may therefore be a general limitation of these models.

When considering the estimated drift rates, both models behave as if there is compressive nonlinear perceptual scaling within the decision process. For the sigma model, the

drift rates are negatively related to absolute evidence in Experiment 1, but are almost identical across the stimulus conditions in Experiment 2. This is consistent with logarithmic scaling of perceptual inputs, resulting in a smaller perceived difference in luminance for additive, but not multiplicative, stimulus manipulations. For the eta model, the drift rates are negatively related to absolute evidence in Experiment 1, indicating a compressive nonlinearity within the decision process. However, in Experiment 2 the drift rate is larger in the high absolute evidence condition, compared to the low condition, suggesting that the increase in the objective amount of relative evidence outweighs the impact of the underlying compressive nonlinearity.

Table 1. Parameter estimates for the sigma model and eta model for both experiments.

Experiment 1												
	a	T_{er}	st	η_{low}	η_{high}	v_{low}	v_{high}	σ_{low}	σ_{high}	a_{CoM}	t_{CoM}	RMSE
Sigma	0.080	0.349	0.281	0.300	0.300	0.141	0.052	0.1	0.118	0.119	0.721	0.3428
Eta	0.069	0.349	0.279	0.014	0.308	0.095	0.064	0.1	0.1	0.104	0.720	0.3178
Experiment 2												
	a	T_{er}	st	η_{low}	η_{high}	v_{low}	v_{high}	σ_{low}	σ_{high}	a_{CoM}	t_{CoM}	RMSE
Sigma	0.075	0.403	0.262	0.387	0.387	0.352	0.365	0.1	0.121	0.126	0.735	0.2821
Eta	0.072	0.411	0.287	0.469	0.666	0.424	0.491	0.1	0.1	0.114	0.754	0.2252

The parameters were estimated by simultaneously fitting the initial responses and change of mind responses. a is the boundary separation parameter, T_{er} is the non-decision time parameter (s), st is the parameter specifying the range of across-trial-variability in non-decision time (assumed to be uniform), η_{low} and η_{high} are the across-trial variability in drift rate parameters for low and high absolute evidence trials (assumed to be normally distributed), v_{low} and v_{high} are the drift rates for low and high absolute evidence trials, σ_{low} and σ_{high} are the within trial noise parameters for low and high absolute evidence trials (σ_{low} was fixed to 1 for the sigma model), a_{CoM} is the distance of the change of mind threshold from the initial decision bound, t_{CoM} is the additional processing time parameter which specifies how much additional evidence is processed, and RMSE is the root mean squared deviation of the simulated and actual data.

4.4.4 DDM fit for change-of-mind responses

The sigma model fit the patterns of change-of-mind responses poorly. This model could not capture the qualitative changes in either the frequency or timing of changes of mind across absolute evidence conditions (see Fig 7). In particular, across both experiments this model predicted that changes of mind would be faster and more frequent in high absolute evidence trials. The eta model in comparison performed slightly better. This model was able to predict, at least qualitatively, the unexpected interaction between initial response accuracy and absolute evidence which we observed in Experiment 1 (see Fig 7E). However, it underestimated the proportion of corrected errors in Experiment 2, and, like the sigma model, it also tended to incorrectly predict that changes of mind would be faster rather than slower on higher absolute evidence trials. As such, neither model provided a comprehensive account of the effects of absolute evidence magnitude on change-of-mind responses.

Fig 7. Change-of-mind responses and model predictions. Panels A) and B) show the group averaged data for change-of-mind responses in Experiment 1. Panels C) and D) show the predictions from the sigma model, whilst Panels E) and F) show the predictions from the eta model. Panels G) and H) show the group averaged data for change-of-mind responses in Experiment 2. Panels I) and J) show the predictions from the sigma model, whilst K) and L) show the predictions from the eta model. In all plots, yellow denotes data from the low absolute evidence condition and blue denotes data from the high absolute evidence condition. Dashed lines indicate corrected error trials (i.e. changes away from an

initially incorrect response), whilst solid lines indicate spoilt responses (i.e. changes away from an initially correct response). Error bars indicate the SEM.

4.5 An extended Leaky Competing Accumulator (LCA) model

The final model we considered was the Leaky Competing Accumulator (LCA) model (Usher & McClelland, 2001). This model has been shown to account for the effect of absolute evidence on one-off perceptual decisions, so is important to consider when searching for an account of the current findings (Teodorescu et al., 2016). In this model, like in the attractor network and neural circuit models, two competing accumulators encode the evidence for each decision alternative. When the activity of one accumulator crosses a threshold level, an initial decision is made. In our extension of the LCA model, we assumed that the decision process then continues to unfold, and, if the activity of the initially unsuccessful accumulator then crossed the decision threshold, and predominated (i.e. was higher than the activity of the initially winning accumulator), then a change of mind occurs (see Evans, Dutilh, Wagenmakers, & Maas, 2019 for related work).

4.5.1 LCA model specifics

In fitting the LCA model, the starting point of each accumulator on each trial was determined by a uniformly distributed random value with a mean of 0.1 and a range of Sz. Increases in absolute evidence were assumed to lead to increases in the shared input to each accumulator (I). The activity of each accumulator was half-wave rectified (i.e. limited to a minimum 0). As in the DDM analysis above, we allowed the drift rate (i.e. the additional input to the accumulator associated with the correct response; v) to vary across the absolute evidence conditions. This was to account for the possibility of a nonlinear perceptual transformation within the decision process (see section 4.1.1). After preliminary fits, we found it was necessary to assume that the amount of leakage (k) in the decision process differs across the stimulus conditions, to capture the slowing of change-of-mind responses with increased absolute evidence. Leakage—a key feature of the LCA model—refers to the dissipation of information

over time from the decision process. This reflects the exponential decay in neural firing that has been observed in neurophysiological experiments (Usher & McClelland, 2001). Finally, the amount of lateral inhibition (β) between the decision accumulators was fixed across stimulus conditions.

Unlike in the DDMs, we did not need to assume the presence of a second decision threshold for changes of mind, nor did we need to assume that there was a time-limit on the processing of post-decisional evidence. Instead, we could simply assume that if the activity of the initially unsuccessful accumulator crossed the initial decision threshold at any point in the post-decision period, and predominated (i.e. was higher than the activity of the initially winning accumulator), then a change of mind would occur.

4.5.2 LCA model fitting

The LCA model was fit to group-average initial response proportions, initial response time quantiles, change-of-mind proportions and change-of-mind latency quantiles simultaneously. This was carried out in MATLAB with custom code available at <https://osf.io/sr58p/> (13.33 ms timesteps, 1,000,000 trials per iteration). As in the DDM analysis above, we adopted the simplifying assumption that the non-decision time was the same for initial and change-of-mind responses. The discrepancy between the data and model predictions was again quantified as the root mean squared error between actual and simulated data and a simplex function (Nelder & Mead, 1965) was used to minimize this value.

4.5.3 LCA modelling results

The LCA model captured the initial response patterns across the stimulus conditions (Fig 8). In particular, it was able to predict the speed up in response times and decrease in accuracy which we observed with increases in absolute evidence magnitude.

Fig 8. Initial responses and LCA model predictions. Panels A) and B) show the group averaged data for initial responses in Experiments 1 and 2 as well as the predictions from the LCA model. In all plots, response proportions are plotted on the x-axis and response time quantiles are plotted on the y-axis. The hollow symbols denote the empirical data, and the solid symbols denote model predictions. Yellow data points are used to represent data from the low absolute evidence condition and blue data points are used to represent data from the high absolute evidence condition. Circular symbols denote correct responses and square symbols denote incorrect responses.

Considering the estimated drift-rate parameters for the LCA model, it is clear that this model again behaves as if there is a compressive nonlinearity within the decision process. This is because the drift rates are negatively associated with absolute evidence in Experiment 1 but are practically identical in Experiment 2 – consistent with roughly logarithmic scaling of sensory inputs.

Table 2. Parameter estimates for the LCA model for both experiments.

Experiment 1											
B	T_{er}	st	I_{low}	I_{high}	V_{low}	V_{high}	k_{low}	k_{high}	β	Sz	RMSE
0.391	0.318	0.268	0.019	0.035	0.019	0.007	0.014	0.035	0.900	0.161	0.259
Experiment 2											
B	T_{er}	st	I_{low}	I_{high}	V_{low}	V_{high}	k_{low}	k_{high}	β	Sz	RMSE
0.437	0.350	0.218	0.020	0.037	0.036	0.035	0.039	0.071	0.702	0.171	0.46

The parameters were estimated by simultaneously fitting the initial responses and change of mind responses. B is the decision threshold parameter, T_{er} is the non-decision time parameter (s), st is the parameter specifying the range of across-trial-variability in non-decision time (assumed to be uniform), I_{low} and I_{high} specify the share input to each decision accumulator parameters in the low and high absolute evidence trials, V_{low} and V_{high} are the drift rates for low and high absolute evidence trials (i.e. the additional

input into the correct accumulator), k_{low} and k_{high} are the leak parameters for the low and high absolute evidence trials, β is lateral inhibition parameter, $\mathcal{S}_x$ starting point variability parameter (specifying the range of a uniform distribution around a mean value of 0.1), and RMSE is the root mean squared deviation of the simulated and actual data.

4.5.4 LCA fit for the change-of-mind responses

For both experiments, the LCA model was able to capture the main qualitative changes in the timing and frequency of changes of mind across the stimulus conditions (see Fig 9). Notably, this model was able to predict the slowing of changes of mind with increases in absolute evidence – a behavioural feature which neither of the DDM’s could fully capture. Nevertheless, there were still some features of the data that the LCA model could not capture. In particular, this model struggled to fully capture the change-of-mind response times in Experiment 2. Whilst it could predict the general slowing in change-of-mind response times with higher absolute evidence, it incorrectly predicted that for spoilt responses (i.e. changes away from a correct initial response) the effect of absolute evidence on change-of-mind speed grows with time (Fig 9. H). Similarly, the model was also unable to capture the crossover of the change time distributions (i.e. the broader distribution of response times for corrected errors).

Fig 9. Change-of-mind responses and LCA model predictions. Panels A) and B) show the group averaged data for change-of-mind responses in Experiment 1. Panels C) and D) show the predictions of the LCA model for change-of-mind proportions and change-of-mind response time respectively. Panels E) and F) show the group averaged data for change-of-mind responses in Experiment 2. Panels G) and H) show the LCA model. In all plots, blue denotes data from high absolute evidence trials, yellow denotes data from low absolute evidence trials, dashed lines indicate trials in which the initially response was incorrect and solid lines denote trials in which the initial response was correct (spoilt responses). Error bars indicate SEM.

4.5.5 LCA predictions for choice consistency

Since the LCA model was best able to capture the effects of absolute evidence on the participants' behaviour, we examined whether this model could also account for the changes in choice-consistency across the absolute evidence conditions. To this end, we decomposed the within-trial variability in the model into two components. One component we termed 'internal variability', which accounts for variability in the decision process which differs across stimulus repetitions (e.g., fluctuations in attention or neural firing). The other component we termed 'external variability', which accounts for stimulus-driven variability (i.e. the random flicker in the two squares). Critically, this variability component is assumed to be identical across stimulus

repetitions. By varying the ratio of these two variability components, whilst keeping combined variability (σ) fixed to 0.1 (as in the fitting procedure), we found it was possible to account for the choice-consistency patterns in each experiment (Fig 10). For Experiment 1, the decrease in choice consistency could be captured by a ratio of external to internal variability of ~ 0.4 . For Experiment 2, the results could be captured with a ratio of ~ 0.75 . The higher ratio of external to internal variability in Experiment 2, compared to Experiment 1, may be due to the fact there was more relative evidence in Experiment 2 (i.e. the stimuli were more discriminable), leading participants to place more weight on stimulus fluctuations. Alternatively, Poisson-like encoding of relative evidence strength could also explain the increase in stimulus-driven variability (as stronger evidence would lead to more variable encoding).

Fig 10. LCA model predicts choice consistency. These plots show the group averaged data for proportion of initial correct initial responses (y-axis) and the proportion of responses that were repeated when participants were presented with an exact stimulus repetition (x-axis). The dots with the error bars (indicating SEM) denote the actual data, whilst the dots joined by the black lines represent the model predictions.

5. Discussion

In this study, we report that the timing and accuracy of perceptual change-of-mind decisions are affected by variations in absolute evidence magnitude. We show that the observed pattern of effects cannot be accounted for by existing change-of-mind models, nor by two modified DDMS which previously have been used to account for the effect of absolute evidence on one-off perceptual decisions. Out of the models we examined, the best account of the behavioural findings is given by an extended LCA model in which leak is positively associated with absolute evidence magnitude. This suggests that input-dependent leak, and the dynamics of lateral inhibition, are important factors in accounting for perceptual changes of mind.

5.1 How plausible is the extended LCA model?

Given that the LCA model provided the best account of the current data, it is worth examining the core assumptions of this model in greater detail. To account for the effects of absolute evidence, three main assumptions needed to be made: First, increases in absolute evidence lead to greater mutual input to the decision accumulators. Second, increases in absolute evidence lead to decreases in drift rate (i.e. decreased relative evidence) – particularly with additive stimulus manipulations where evidence ratios are not conserved. Finally, it was also necessary to assume that leak was positively associated with absolute evidence magnitude. The first two assumptions are relatively straightforward – the second being consistent with a Weber-like compressive nonlinearity in the decision process (see section 5.7). However, the third assumption was somewhat arbitrary and deserves further consideration.

One way of further testing the plausibility of the extended LCA model would be to examine the patterns of neural activity which occur when manipulating absolute evidence. With increases in absolute evidence, the LCA model predicts that the average activity of decision-selective neural pools will decrease (see Fig A.2 in the supplementary materials). This is the case even when the drift rate between stimulus conditions is identical (i.e. when the amount of ‘perceived’ relative evidence is matched). Interestingly, recordings from neurons in area MT during transparent dot motion (i.e. the presentation of dot stimuli which are moving in opposing

directions) support this prediction, with increases in bi-directional motion (i.e. increases in absolute evidence) leading to decreases in the firing rate of motion-selective neurons (see Snowden, Treue, Erickson, & Andersen, 1991 Fig 12). However, whether this holds for manipulations of other forms of absolute evidence (e.g., luminance) remains to be seen.

5.2 Using neural variability to distinguish between competing models

Recordings of neural activity could also be used to further arbitrate between the modelling frameworks considered in the current paper. In particular, measures of firing-rate variability would help to distinguish whether the effects of absolute evidence are best understood as resulting from the effects of input-dependent noise or the dynamics of leak and lateral inhibition. Both the sigma and eta models (which include input-dependent noise sources) predict that the variability of the decision process will increase in conditions of high absolute evidence. In contrast, the LCA and attractor network models predict that with higher absolute evidence, there will be less variation in firing rate across trials (see Albantakis & Deco, 2011 and Fig A.3 in the supplementary materials). Given this, if future studies examined firing rate variability across conditions of high and low absolute evidence, this would provide a strong test of whether the behavioral effects of absolute evidence manipulations are best explained by input-dependent noise or the dynamics of leak and lateral inhibition.

5.3 Do change-of-mind mechanisms account for evidence variability?

Both the eta and sigma models rely on the assumption that input-dependent noise, varying either within or across trials, underlies the effects of absolute evidence on initial responses. If this assumption is correct, the fact that we observed either a decrease (main effect in Experiment 1) or no difference (Experiment 2) in the proportion of changes of mind with increases in absolute evidence may point to an adaptive change-of-mind mechanism which attempts to avoid costly vacillation. Recently, it has been proposed that an explicit representation of evidence reliability could be encoded in the decision process (Yeung & Summerfield, 2012). If

this is true, then such a representation could plausibly be drawn upon to flexibly adjust the change of mind threshold within the course of a single trial. When evidence is noisy the threshold for changing one's mind could be set higher than when evidence is reliable, so as to avoid unnecessary changes of mind. Such a mechanism would make it possible to simultaneously capture the speed-up in initial response time (due to the effect of input dependent noise) and the slow-down in changes of mind (as with a higher threshold, more evidence, and thus time, is required to overrule a decision).

5.4 Does the change of mind threshold depend on initial confidence?

Plausibly, the position of the change-of-mind threshold may also depend on initial decision confidence. For high confidence decisions the change-of-mind threshold may be set further away from the initial decision threshold, than for low confidence decisions. This would result in more contradictory evidence being required to overrule high confidence decisions. Previously, it has been shown that initial response time is often negatively associated with decision confidence, whereby confidence is greater for fast decisions (Kiani, Corthell, & Shadlen, 2014). In the current study, participants' initial response times were faster in high absolute evidence trials, compared to low absolute evidence trials. It is therefore possible that participants had an inflated sense of confidence in their initial decisions on high absolute evidence trials, despite being objectively less accurate. Consequently, participants may have set higher change-of-mind thresholds. This offers an alternative explanation as to why changes of mind were slower on high absolute evidence trials, as more time would be needed to accumulate the additional evidence. However, this would also predict fewer changes of mind (of both types) with higher absolute evidence, which is not consistent with our observations. As such, a dynamic change-of-mind threshold alone cannot account for the current data. However, future work could consider whether a dynamic threshold, in concert with other mechanisms, might capture the current observations.

5.5 Can a metacognitive bias towards decision-congruent evidence explain the results?

Recently, a number of studies have demonstrated that humans overweight decision-congruent information when rating their confidence in a previous perceptual decision (Koizumi, Maniscalco, & Lau, 2015; Maniscalco, Peters, & Lau, 2016; Peters et al., 2017; Zylberberg, Barttfeld, & Sigman, 2012). For example, when asked to judge how confident they are that they correctly chose the brighter of two squares, participants will tend to ignore information which provides evidence against their choice (i.e. the brightness of the unchosen square), and instead focus on information which is decision-congruent (i.e. the brightness of the chosen square; Zylberberg et al., 2012). Under the assumption that one's confidence in their initial decision affects the position of their change-of-mind threshold, a bias towards decision-congruent information also offers an explanation for the change of mind latency effects which we observed. In particular, on high absolute evidence trials the chosen square will be brighter, leading participants to be more confident. If, as a consequence of this increase in confidence, they then set a higher threshold for changing their mind, those responses will slow down. However, as we noted above, this view cannot explain the interaction between initial response accuracy and absolute evidence (i.e. the increase in the number of spoilt responses). If the change-of-mind threshold is higher, then the number of spoilt responses should decrease. Finally, it cannot explain why the proportion of changes of mind was unaffected by a multiplicative stimulus manipulation. As such, a bias towards-decision congruent information alone cannot account for the current findings.

5.6 Are changes of mind driven by a second order process?

The current results do not rule out the possibility that changes of mind arise from a second order process (i.e. a process which is, at least partially, distinct from the initial decision process). Indeed, the fact that increases in absolute evidence had opposing effects on the timing of initial decisions and change-of-mind decisions may suggest a dissociation between the processes which underlie these two responses (Fleming & Daw, 2017). From the modelling results, it is clear that input-dependent noise can explain the speed-up in initial response times

across conditions. However, models that incorporate input-dependent noise also tend to predict faster, rather than slower, change-of-mind latencies. Given this, frameworks built on partial dissociations between the initial decision process and the change-of-mind process, where the change-of-mind process does not share all the dynamics of the initial decision process (e.g. does not inherit input-dependent noise), may be better suited to accounting for the different response time effects. Nevertheless, the fact that the extended LCA model was able to predict the simultaneous speeding and slowing of initial and change of mind response demonstrates that it is possible to explain the opposing response time effects within a single decision process.

5.7 Weber's law

It is worth considering the current results with respect to Weber's law. According to Weber's law, the just-noticeable difference between two stimuli is inversely proportional to the overall intensity of the two stimuli (i.e. to absolute evidence magnitude). For the current study, this means that the perceived difference in luminance between the stimuli in the high absolute evidence condition will have been diminished compared to the perceived difference between the stimuli in the low condition (at least in Experiment 1 where evidence ratios were not conserved). Indeed, the reason we allowed the drift rates (representing relative evidence strength) to vary across stimulus conditions in all fitted models was to account for this very possibility. Considering the fitted drift rates for all models, the parameter values indeed suggest the presence of a compressive nonlinearity within the decision process, which is roughly consistent with logarithmic scaling of perceptual inputs. Critically however, the effect of a compressive nonlinearity alone cannot fully account for our findings. This is because, whilst a compressive nonlinearity causes initial decisions to be less accurate in trials with higher absolute evidence, it also causes initial responses to be slower, not faster as we reliably observed. Given this, we conclude that our absolute evidence manipulations are having an effect over and above the effect of a diminished perceptual difference between the stimuli in high absolute evidence trials.

So what is this additional effect? For the LCA model we assumed that increases in absolute evidence also led to greater shared input to the decision accumulators as well as greater leak – allowing us to capture the simultaneous speeding and slowing of initial and change-of-mind responses. For the DDMs we assumed that increases in absolute evidence lead to greater variability, allowing us to capture the speeding of initial responses. However, this also caused the models to incorrectly predict a speeding of change-of-mind responses.

Given that the difference in overall luminance between the high and low absolute evidence conditions was quite large, one could alternatively argue that participants may have adopted condition-dependent strategies, for example in the setting of their initial decision threshold, and that this may explain the changes in behaviour we observed, over and above those driven by nonlinear transformation of perceptual inputs. For example, if participants had adopted a lower decision threshold in response to brighter stimuli, this could explain the decrease in accuracy and response time we observed. However, our results for initial responses are in line with Teodorescu et al. (2015) and Ratcliff et al. (2019), who used more closely overlapping stimulus distributions, which could not have been easily discriminated. Moreover, to our knowledge there is no evidence that it is possible to make sub-second, reactive decision threshold adjustments (as would be required in our task). Finally, even if it is possible to implement such a strategy, participants would have been incentivised against doing so by the feedback telling them they were already less accurate at judging between brighter stimuli. Hence, we argue that our results are not best explained by condition-dependent strategy use. Given that the LCA model was best able to account for the current findings, the effect of absolute evidence magnitude on behaviour is best understood in terms of the combined effects of nonlinear perceptual scaling, increased mutual input to the decision accumulators, and input-dependent leak.

For our simulation of the Neural Circuit model, it is important to note that the amount of relative evidence was assumed to be constant across stimulus conditions. This was in keeping

with the simulations conducted by Albantakis and Deco (2011) for the predictions of the Attractor network model across changes in absolute evidence magnitude. Further modification and fitting of the Attractor Network and Neural Circuit models was beyond the scope of this paper, due to their sheer complexity. However, future theoretical work could examine whether these models, or other related models (e.g., Pais et al., 2013), with additional modifications (i.e. nonlinear scaling of sensory inputs) can better account for the current behaviour. Given the similarities between the LCA and these models, it is possible that with the same set of additional assumptions as those of the extended LCA (i.e. changes in drift and leak across conditions), they may offer a similar account of the current findings.

5.8 Limitations

The findings of our study should be interpreted with the following limitations in mind. First, behavioural responses were recorded using button presses rather than by tracking continuous movement trajectories, as has been done in a number of previous studies investigating changes of mind (Burk et al., 2014; Moher & Song, 2014; Resulaj et al., 2009; van den Berg et al., 2016). Tracking movement trajectories has the advantage that changes of mind can be more directly observed, for example in a change of direction or slowing of the movement. However, by recording button presses, we were afforded with a unique opportunity to characterise the onset times of change-of-mind responses. This has not been done in previous change-of-mind studies as responses were made unimanually, making it more difficult to define the point at which the change of mind began (Albantakis et al., 2012; Burk et al., 2014; Moher & Song, 2014; Resulaj et al., 2009; van den Berg et al., 2016). Overall, the results of the current study suggest that accounting for the latencies of changes-of-mind decisions constitutes a critical test of computational models, which has been overlooked in previous work.

Another potential limitation of this study is that we imposed time limits for initial responses (800ms) and changes of mind (1s). The limit for initial responses was imposed as a means of generating errors, and consequently changes of mind (Resulaj et al., 2009). Not

implementing any deadline at all would have encouraged the use of a very liberal decision criterion, which would make changes of mind unnecessary. However, as a result of the deadline the response time distributions will have been censored (i.e. the tails of the distributions will have been cut off). It is also possible that participants may have adopted a hybrid decision strategy involving an accumulation to bound mechanism plus a fast guessing process, which was triggered in the case of long decision times (e.g. Noorbaloochi, Sharon, & McClelland, 2015). This would have provided a means of circumventing the deadline to avoid a high number of missed responses, and offers an explanation as to why errors were faster than correct responses. Future theoretical work may therefore consider exploring whether novel models, based on hybrid decision processes which include a random guessing mechanism, are better able to capture the observed data.

Finally, because participants did not undergo training prior to each experiment, their behavioural performance was not completely stationary across blocks (Fig A.1 in the supplementary materials shows the proportion of changes of mind across time in both experiments). This non-stationarity is important to consider, particularly when interpreting the results of double-pass analyses – which often rely on the assumption of stationarity. However, since the critical comparison for our double-pass analysis was between the two (interleaved) stimulus conditions, non-stationary behaviour will have equally affected the choice-consistency estimates for each stimulus condition. As such, the condition-wise differences we observe cannot be driven by learning-related changes in behaviour across the experiments.

5.9 Conclusion

To conclude, in the current study we have shown that perceptual change-of-mind decisions are sensitive to variations in absolute evidence. We found that changes of mind are consistently slower and often less accurate in conditions of high absolute evidence. We have shown that this pattern of effects is best accounted for by an extended LCA model in which leak is positively associated with absolute evidence magnitude.

900 **6. Acknowledgements**

901 We thank Prof. Philip Smith for helpful comments on an early draft of this manuscript.

902 **7. Funding**

903 This work was supported by an Australian Research Council (ARC) Discovery Project

904 Grant [DP160103353] to SB and RH and an Australian Government Research Training Program

905 (RTP) Scholarship to WT.

References

- Albantakis, L., Branzi, F. M., Costa, A., & Deco, G. (2012). A multiple-choice task with changes of mind. *PLoS ONE*, 7(8), e43131. <https://doi.org/10.1371/journal.pone.0043131>
- Albantakis, L., & Deco, G. (2011). Changes of mind in an attractor network of decision-making. *PLoS Computational Biology*. <https://doi.org/10.1371/journal.pcbi.1002086>
- Atiya, N., Rañó, I., Prasad, G., & Wong-Lin, K. (2019). A neural circuit model of decision uncertainty and change-of-mind. *Nature Communications*, 10, 1–12. <https://doi.org/10.1038/s41467-019-10316-8>
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4 Douglas. *Journal of Statistical Software*, 67(1), 1-48. <https://doi.org/10.18637/jss.v067.i01>
- 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>
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. <https://doi.org/10.1163/156856897X00357>
- Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. *Science*, 340(6128), 95–99. <https://doi.org/10.1126/science.1233912>
- Burk, D., Ingram, J. N., Franklin, D. W., Shadlen, M. N., & Wolpert, D. M. (2014). Motor effort alters changes of mind in sensorimotor decision making. *PLoS ONE*, 9(3), e92681. <https://doi.org/10.1371/journal.pone.0092681>

929 Evans, N. J., Dutilh, G., Wagenmakers, E., & Maas, L. J. Van Der. (2019). Double Responding :
930 A New Constraint for Models of Speeded Decision Making. *PsyArXiv*.

931 Fleming, S. M., & Daw, N. D. (2017). Self-Evaluation of Decision-Making: A General Bayesian
932 Framework for Metacognitive Computation. *Psychological Research*, 124(1), 91–114.
933 <https://doi.org/10.1037/rev0000045>

934 Fox, J., Weisberg, S., Price, B., Friendly, M., Firth, D., & Taylor, S. (2016). Package ‘ effects ’
935 Retrieved from <https://cran.r-project.org/web/packages/effects/effects.pdf>

936 Gold, J. I., & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of*
937 *Neuroscience*, 30(1), 535–574. <https://doi.org/10.1146/annurev.neuro.29.051605.113038>

938 Green, D. M. (1964). Consistency of auditory detection judgments. *Psychological Review*, 71(5),
939 392–407. <https://doi.org/10.1037/h0044520>

940 Hanks, T., Kopec, C. D., Brunton, B. W., Duan, C. A., Erlich, J. C., & Brody, C. D. (2015).
941 Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature*,
942 520(7546), 220–223. <https://doi.org/10.1038/nature14066>.

943 Hanks, T., & Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and
944 Humans. *Neuron*, 93(1), 15–31. <https://doi.org/10.1016/j.neuron.2016.12.003>

945 Hunt, L. T., Kolling, N., Soltani, A., Woolrich, M. W., Rushworth, M. F. S., & Behrens, T. E. J.
946 (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature*
947 *Neuroscience*, 15(3), 470–476. <https://doi.org/10.1038/nn.3017>

948 Kaufman, M. T., Churchland, M. M., Ryu, S. I., & Shenoy, K. V. (2015). Vacillation, indecision
949 and hesitation in moment-by-moment decoding of monkey motor cortex. *ELife*, 4, e04677.
950 <https://doi.org/10.7554/eLife.04677>

951 Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence
952 and decision time. *Neuron*. Retrieved from
953 <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC4271191/pdf/nihms647759.pdf>

954 Kiani, R., Cueva, C. J., Reppas, J. B., & Newsome, W. T. (2014). Dynamics of neural population
955 responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology*,
956 24(13), 1542–1547. <https://doi.org/10.1016/j.cub.2014.05.049>

957 Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What’s new
958 in Psychtoolbox-3. *Perception*, 36(14), 1–16. <https://doi.org/10.1068/v070821>

959 Koizumi, A., Maniscalco, B., & Lau, H. (2015). Does perceptual confidence facilitate cognitive
960 control? *Atten Percept Psychophys*, (March), 1295–1306. [https://doi.org/10.3758/s13414-015-](https://doi.org/10.3758/s13414-015-0843-3)
961 0843-3

962 Liew, S. X., Howe, P. D. L., & Little, D. R. (2016). The appropriacy of averaging in the study of
963 context effects. *Psychonomic Bulletin and Review*, 23(5), 1639–1646.
964 <https://doi.org/10.3758/s13423-016-1032-7>

965 Lu, Z. L., & Doshier, B. A. (2008). Characterizing Observers Using External Noise and Observer
966 Models: Assessing Internal Representations With External Noise. *Psychological Review*, 115(1),
967 44–82. <https://doi.org/10.1037/0033-295X.115.1.44>

968 Lüdecke, D. (2018). sjPlot - Data Visualization for Statistics in Social Science.
969 <https://doi.org/10.5281/ZENODO.2400856>

970 Maniscalco, B., Peters, M. A. K., & Lau, H. (2016). Heuristic use of perceptual evidence leads to
971 dissociation between performance and metacognitive sensitivity. *Atten Percept Psychophys*,
972 78(3), 923–937. <https://doi.org/10.3758/s13414-016-1059-x>.Heuristic

973 Marshall, J. A. R. (2019). Comment on ‘optimal policy for multi-alternative decisions.’ *BioRxiv*.

974 Moher, J., & Song, J.-H. (2014). Perceptual decision processes flexibly adapt to avoid change-of-
975 mind motor costs. *Journal of Vision*, 14(8), 1–13. <https://doi.org/10.1167/14.8.1>

976 Nelder, J. A., & Mead, R. (1965). A simplex method for function minimization. *The Computer*
977 *Journal*, 7(4), 308–313. <https://doi.org/10.1093/comjnl/7.4.308>

978 Noorbaloochi, S., Sharon, D., & McClelland, J. L. (2015). Payoff information biases a fast guess
979 process in perceptual decision making under deadline pressure: Evidence from behavior,
980 evoked potentials, and quantitative model comparison. *Journal of Neuroscience*, 35(31), 10989–
981 11011. <https://doi.org/10.1523/JNEUROSCI.0017-15.2015>

982 O’Connell, R. G., Dockree, P. M., & Kelly, S. P. (2012). A supramodal accumulation-to-bound
983 signal that determines perceptual decisions in humans. *Nature Neuroscience*, 15(12), 1729–
984 1735. <https://doi.org/10.1038/nn.3248>

985 Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., & Marshall, J. A. R. (2013). A
986 mechanism for value-sensitive decision-making. *PloS One*, 8(9).
987 <https://doi.org/10.1371/journal.pone.0073216>

988 Peters, M., Thesen, T., Ko, Y., Maniscalco, B., Carlson, C., Davidson, M., ... Lau, H. (2017).
989 Perceptual confidence neglects decision-incongruent evidence in the brain. *Nature Human*
990 *Behaviour*, 1(7). <https://doi.org/10.1038/s41562-017-0139>

991 Pirrone, A., Habiba, A., Hayden, B. Y., Stafford, T., & Marshall, J. A. R. (2018). Evidence for the
992 speed-value trade-off: human and monkey decision making is magnitude sensitive. *Decision*,
993 5(2), 129–142. <https://doi.org/10.1037/dec0000075>.Evidence

994 Polanía, R., Krajbich, I., Grueschow, M., & Ruff, C. C. (2014). Neural Oscillations and
995 Synchronization Differentially Support Evidence Accumulation in Perceptual and Value-
996 Based Decision Making. *Neuron*, 82(3), 709–720.

997 <https://doi.org/10.1016/j.neuron.2014.03.014>

998 Rabbitt, P. M., & Vyas, S. M. (1981). Processing a display even after you make a response to it.
999 how perceptual errors can be corrected. *The Quarterly Journal of Experimental Psychology Section*
1000 *A*, 33(3), 223–239. <https://doi.org/10.1080/14640748108400790>

1001 Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108.
1002 <https://doi.org/10.1037/0033-295X.85.2.59>

1003 Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-
1004 Choice Decision Tasks. *Neural Computation*, 20(4), 873–922.
1005 <https://doi.org/10.1126/scisignal.2001449>.

1006 Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion Decision Model:
1007 Current Issues and History. *Trends in Cognitive Sciences*, 20(4), 260–281.
1008 <https://doi.org/10.1016/j.tics.2016.01.007>.

1009 Ratcliff, R., Thapar, A., & McKoon, G. (2001). The effects of aging on reaction time in a signal
1010 detection task. *Psychology and Aging*, 16(2), 323–341. Retrieved from
1011 <https://www.ncbi.nlm.nih.gov/pubmed/11405319>

1012 Ratcliff, R., Thapar, A., & McKoon, G. (2003). A diffusion model analysis of the effects of aging
1013 on brightness discrimination. *Perception and Psychophysics*, 65(4), 523–535.
1014 <https://doi.org/10.3758/BF03194580>

1015 Ratcliff, R., Thapar, A., & McKoon, G. (2004). A diffusion model analysis of the effects of aging
1016 on recognition memory. *Journal of Memory and Language*, 50(4), 408–424.
1017 <https://doi.org/10.1016/j.jml.2003.11.002>

1018 Ratcliff, R., Voskuilen, C., & McKoon, G. (2018). Internal and external sources of variability in
1019 perceptual decision-making. *Psychological Review*, 125(1), 33–46.

1020 <https://doi.org/10.1037/rev0000080>

1021 Ratcliff, R., Voskuilen, C., & Teodorescu, A. (2018). Modeling 2-Alternative Forced-Choice
 1022 Tasks: Accounting for both Magnitude and Difference Effects. *Cognitive Psychology*, *103*, 1–
 1023 22. <https://doi.org/10.1016/j.cogpsych.2018.02.002>

1024 Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-
 1025 making. *Nature*, *461*(7261), 263–266. <https://doi.org/10.1038/nature08275>

1026 Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in*
 1027 *Neurosciences*, *27*(3). <https://doi.org/10.1016/j.tins.2004.01.006>

1028 Snowden, R. J., Treue, S., Erickson, R. G., & Andersen, R. a. (1991). The response of area MT
 1029 and V1 neurons to transparent motion. *Journal of Neuroscience*, *11*(9), 2768–2785. Retrieved
 1030 from <http://www.ncbi.nlm.nih.gov/pubmed/1880548>

1031 Steverson, K., Chung, H., Zimmermann, J., Louie, K., & Glimcher, P. (2019). Sensitivity of
 1032 reaction time to the magnitude of rewards reveals the cost-structure of time. *Scientific*
 1033 *Reports*, 1–14. <https://doi.org/10.1038/s41598-019-56392-0>

1034 Tajima, S., Drugowitsch, J., Patel, N., & Pouget, A. (2019). Optimal policy for multi-alternative
 1035 decisions. *Nature Neuroscience*, 1–27.

1036 Tajima, S., Drugowitsch, J., & Pouget, A. (2016). Optimal policy for value-based decision-
 1037 making. *Nature Communications*, 1–12. <https://doi.org/10.1038/ncomms12400>

1038 Teodorescu, A. R., Moran, R., & Usher, M. (2016). Absolutely relative or relatively absolute:
 1039 violations of value invariance in human decision making. *Psychonomic Bulletin and Review*,
 1040 *23*(1), 22–38. <https://doi.org/10.3758/s13423-015-0858-8>

1041 Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing
 1042 accumulator model. *Psychological Review*, *108*(3), 500.

1043 van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D.
 1044 M. (2016). A common mechanism underlies changes of mind about decisions and
 1045 confidence. *ELife*, 5, e12192. <https://doi.org/10.7554/eLife.12192>
 1046 Yeung, N., & Summerfield, C. (2012). Metacognition in human decision-making: confidence and
 1047 error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 367(1594),
 1048 1310–1321. <https://doi.org/10.1098/rstb.2011.0416>
 1049 Zylberberg, A., Barttfeld, P., & Sigman, M. (2012). The construction of confidence in a
 1050 perceptual decision. *Frontiers in Integrative Neuroscience*, 6(September), 1–10.
 1051 <https://doi.org/10.3389/fnint.2012.00079>
 1052

1053 **Table A.1. Parameter estimates for the accuracy and change-of-mind (CoM) regression**
1054 **models (Experiment 1).**

	Accuracy			CoM		
<i>Predictors</i>	<i>Odds Ratios</i>	<i>CI</i>	<i>p</i>	<i>Odds Ratios</i>	<i>CI</i>	<i>p</i>
(Intercept)	2.38	2.15 – 2.63	<0.001	0.22	0.12 – 0.38	<0.001
Condition(1)	0.57	0.54 – 0.61	<0.001	0.48	0.35 – 0.65	<0.001
Initial RT	1.11	1.04 – 1.17	0.001	0.97	0.85 – 1.11	0.648
Accuracy(1)				0.09	0.05 – 0.15	<0.001
Condition:Accuracy				4.32	2.85 – 6.55	<0.001

1055 Note: This table was made using the tab_model function in the sjPlot R package (Lüdecke,
1056 2018).

Table A.2. Parameter estimates for the initial response time and change time regression models (Experiment 1).

<i>Predictors</i>	Initial RT			Change Time		
	<i>Estimates</i>	<i>CI</i>	<i>p</i>	<i>Estimates</i>	<i>CI</i>	<i>p</i>
(Intercept)	0.47	0.44 – 0.50	<0.001	0.46	0.41 – 0.51	<0.001
Condition(1)	-0.01	-0.02 – -0.01	<0.001	0.04	0.02 – 0.05	<0.001
Accuracy(1)	0.01	0.01 – 0.01	<0.001	-0.05	-0.07 – -0.03	<0.001
Initial RT				-0.02	-0.09 – 0.05	0.584

Note: This table was made using the `tab_model` function in the `sjPlot` R package (Lüdtke, 2018).

Table A.3. Parameter estimates for the accuracy and change-of-mind (CoM) regression models (Experiment 2).

<i>Predictors</i>	Accuracy			CoM		
	<i>Odds Ratios</i>	<i>CI</i>	<i>p</i>	<i>Odds Ratios</i>	<i>CI</i>	<i>p</i>
(Intercept)	5.01	4.22 – 5.95	<0.001	0.53	0.31 – 0.90	0.018
Condition(1)	0.82	0.77 – 0.87	<0.001	0.94	0.71 – 1.24	0.665
Initial RT	0.93	0.85 – 1.02	0.123	1.00	0.90 – 1.12	0.958
Accuracy(1)				0.02	0.02 – 0.04	<0.001
Condition:Accuracy				0.96	0.61 – 1.51	0.850

Note: This table was made using the `tab_model` function in the `sjPlot` R package (Lüdtke, 2018)

Table A.4. Parameter estimates for the initial response time and change time regression models (Experiment 2).

<i>Predictors</i>	Initial RT			Change Time		
	<i>Estimates</i>	<i>CI</i>	<i>p</i>	<i>Estimates</i>	<i>CI</i>	<i>p</i>
(Intercept)	0.50	0.49 – 0.52	<0.001	0.42	0.37 – 0.47	<0.001
Condition(1)	-0.010	-0.02 – 0.00	0.002	0.04	-0.03 – 0.06	<0.001
Accuracy(1)	-0.00	-0.00 – 0.00	0.936	0.01	-0.01 – 0.03	0.339
Initial RT				-0.03	-0.10 – 0.05	0.455

1067 **Table A.5. Parameter estimates for the choice consistency models.**

<i>Predictors</i>	Consistency (Exp. 1)			Consistency (Exp. 2)		
	<i>Odds</i>	<i>CI</i>	<i>p</i>	<i>Odds</i>	<i>CI</i>	<i>p</i>
	<i>Ratios</i>			<i>Ratios</i>		
(Intercept)	1.04	0.86 – 1.26	0.683	0.88	0.72 – 1.07	0.201
Condition(1)	0.81	0.68 – 0.96	0.015	0.91	0.81 – 1.03	0.135
Accuracy(1)	2.20	2.03 – 2.38	<0.001	6.92	6.29 – 7.61	<0.001
Initial RT	0.89	0.85 – 0.93	<0.001	0.83	0.80 – 0.87	<0.001

1068 Note: This table was made using the `tab_model` function in the `sjPlot` R package (Lüdtke, 2018)

Fig A.1. Changes of mind as a proportion of correct and error responses, across the course of each experiment. The proportion of changes of mind cross 5 stages of each experiment (i.e. for neighbouring pairs of runs) is plotted separately for correct and error responses. Low absolute evidence trials are shown in orange and high absolute evidence trials are shown in blue. Dashed histograms indicate trials in which the initial response was an error ('corrected errors'), solid histograms indicate trials in which the initial response was correct ('spoil correct'). Interestingly, there is a general trend towards both types of changes of mind becoming more common as the each experiment progresses. The relative pattern of changes of mind between the absolute evidence conditions remains constant across these 5 stages of the task, suggesting that the effect of absolute evidence has little to do with learning.

Fig A.2. Average activity in the winning accumulator of the LCA model. We simulated 300,000 trials of the LCA model using the parameter estimates from each experiment. We then plotted the average activity in the winning accumulator on correct trials across time. Blue lines denote activity on high absolute evidence trials and yellow lines denote activity on low absolute evidence trials.

Fig A.3. Variance of the activity in the winning accumulator of the LCA model. We simulated 300,000 trials of the LCA model using the parameter estimates from each experiment. We then plotted the variance of the activity in the winning accumulator on correct trials across time. Blue lines denote activity on high absolute evidence trials and yellow lines denote activity on low absolute evidence trials.