Explicit representation of confidence informs future value-based decisions

Tomas Folke[1], Catrine Jacobsen[2], Stephen M. Fleming[3] and Benedetto De Martino[4, 1*]

1. Department of Psychology, University of Cambridge Downing Street, Cambridge CB2 3EB, United Kingdom
2. University of Copenhagen, Institute of Food and Resource Economics, Copenhagen, Denmark
3. University College London, Wellcome Trust Centre for Neuroimaging, 12 Queen Square, WC1N 3BG London, United Kingdom
4. Institute of Cognitive Neuroscience, 17-19 Queen Square, WC1N 3AR, London, United Kingdom

* Correspondence should be addressed to benedettodemartino@gmail.com

Abstract

Humans can reflect on previous decisions and report variable levels of confidence. But why maintain an explicit representation of confidence for choices that have already been made and therefore cannot be undone? Here we show that an explicit representation of confidence is harnessed for subsequent changes of mind. Specifically, when confidence is low, participants are more likely to change their minds when the same choice is presented again, an effect that is most pronounced in participants with greater fidelity in their confidence reports. Furthermore, we show that choices reported with high confidence follow a more consistent pattern (fewer transitivity violations). Finally, by tracking participants' eye movements we demonstrate that lower-level gaze dynamics can track uncertainty but do not directly impact changes of mind. Taken together, these results suggest that an explicit and accurate representation of confidence has a positive impact on the quality of future value-based decisions.

Introduction

As we navigate through life we are constantly faced with choices that require us to assign and compare the values of different options or actions. Some of these value-based choices seem relatively straightforward ('What should I eat for lunch?') and others less so ('Which job offer should I take?'). No matter how simple or complex these choices are, they are often accompanied by a sense of confidence in having made the right choice. Recent work has shown that it is possible to behaviourally and computationally dissociate a value estimate ('*How much do I like something?*') from internal fluctuations in confidence ('*How sure am I?')*. For example, at a behavioural level it has been shown that confidence shares only a limited amount of variance with value, and instead reflects an assessment of choice accuracy[1]. This relation between value and confidence is neatly accounted for computationally by assuming that confidence emerges from the dynamics of noisy accumulators in an evidence-accumulation framework [1, 2, 3, 4]. More recently, Lebreton and colleagues have shown that confidence may be an inherent property of value estimation, sharing a quadratic relationship with a linear rating of value[5] (see also the work of Barron and colleagues[6]). But what is the function of confidence? Why maintain an explicit representation of confidence when a choice has already been made and therefore cannot be undone?

According to one view, confidence can be thought of as a by-product of a stochastic accumulation process implemented in the ventromedial prefrontal cortex (vmPFC) during value comparison. Previous work indicates the brain constructs an explicit representation of confidence that underpins verbal reports [7, 8]. A range of studies suggests that the rostrolateral prefrontal cortex represents confidence in both value-based and perceptual decisions [1, 9, 10, 11]. Explicit representations of confidence allow individuals to communicate the strength of their beliefs to others, facilitating group decisions [12, 13], but may play little role in one's own decision process.

An alternative view is that explicit representations of confidence are critical for guiding one's own future behaviour [14]. Work in perceptual decision-making has revealed commonalities between mechanisms supporting confidence construction and error-monitoring [15, 16], suggesting changes of mind may be informed by confidence [4]. However, it is unknown whether confidence is harnessed over a longer timescale to guide future choices. Here we aim to test the hypothesis that an explicit (and well-tuned) representation of confidence in a recent choice can guide a decisions maker's choice when faced by the same (or a similar) decision again. To test this hypothesis, we presented participants with the same set of choices more than once during the course of two experiments and tested which factors were associated with a change of mind. We then investigated how confidence related to the degree of internal consistency in their patterns of choice. Choice consistency can be quantified by measuring the degree of transitivity across choices. Here we introduced a novel method for tagging choices as conforming to or violating transitivity. Using this method we were able to show that explicit representations of confidence are associated with more consistent patterns of choice as a consequence of changes of mind. Finally, we directly contrasted the effect of explicit confidence reports with lower-level markers of uncertainty that we gathered using eye tracking, revealing that changes of mind were specifically associated with explicit reports of confidence.

Results

We collected data in two experiments in which hungry participants made choices between food items (which they could consume later) while their eye movements were monitored. In the first experiment the twenty-eight participants included in the study were shown high-definition pictures of *two* snacks and were asked to choose the preferred one (Figure 1 A). In a second experiment twenty-four participants chose their preferred snack among *three* snacks available in each trial (Figure 1 D). After making each choice, participants reported their degree of confidence in having made the 'correct' choice, which in this design equates to choosing the higher valued item. The value for individual items was elicited using a standard incentive compatible BDM-method[17]. The experimental procedure we used was adapted (with modifications) from a task we developed previously[1] (see methods for more details).

Relation between confidence and choice

In line with a wealth of previous research [18, 19, 20, 21, 22] we found that the difference in value between the two items (constructed from values elicited through an incentive-compatible bidding procedure, BDM) was a reliable predictor of participants' choices in both experiments (hierarchical logistic regression; *Experiment 1:* $z = 11.48$, $p < .0001$, Figure 1C and F; *Experiment 2:* $z = 6.66$, $p < .0001$, Figure 1B and E). Note that in the three-choice design (Experiment 2) DV was calculated as the difference between the value of the reference item and the average of the two other available options (following Krajbich and Rangel[23]). In the supplemental materials (S1) we additionally report the result of a multinomial logistic regression model in which the value of each option was inputted independently and therefore does not require a priori specification of DV. This analysis yielded the same pattern of results. In both studies we also identified a significant negative interaction between the summed value of all options (SV) and value difference (DV) (*Experiment 1:* z=-3.08, p<.005; *Experiment 2:* z=-2.84, p<.005), indicating that DV had a stronger influence on choice when item values were low, compared to when items were high in value (Figure 1 C and F). To our knowledge this effect has not been reported before but is consistent with the Weber–Fechner law in sensory perception in which the resolution of precepts diminishes for stimuli of greater magnitude. The effect is also compatible with the notion of normalization [24, 25, 26]. Confidence, unlike DV, was not in itself a predictor of choice (right or left item) but instead correlated with choice accuracy, with a steeper slope relating DV to choice when confidence was high, as found previously[1] (Fig. 1B, E; *Experiment 1:* z=7.43, p<.0001; *Experiment 2:* z=5.82, p<.0001).

Using eye tracking we measured the dynamics of eye movements between items during the choice, both the total amount of time participants spent looking at each item and how frequently gaze shifted back and forth between items (see supplementary materials, section S2). Replicating previous studies[23, 27] we found that the difference in dwell time (DDT) was a robust predictor of choice in both two-option and three-option experiments (*Experiment 1:* z=4.95, p<.0001;*Experiment 2:* z=9.81, p<.0001; Figure 1 C and F).

For a full list of fitted models and their respective BIC scores see the supplementary materials (section S3).

Figure 1

Factors that contribute to confidence

We next investigated which variables contributed to subjective confidence during value-based choice. Our previous work shows an interrelationship between absolute difference in value (|DV|), response time (RT) and confidence (i.e. participants are more confident when |DV| is high and their choice are faster)[1]. These findings are in line with the conceptual relation between confidence, strength of evidence (indexed by |DV| in the value-based framework) and decision time[3, 28]. We observed this same relation in the current study. In both experiments we found that |DV| was a significant predictor of confidence (*Experiment 1:* t=13.43 p<.0001; *Experiment 2:* t=7.46, p<.0001). We also found that RT was a negative predictor of confidence (*Experiment 1:* t=-10.01, p<.0001; *Experiment 2:* t=-7.53, p<.0001). Additionally, we found that summed value positively predicted confidence, meaning that participants tended to be more confident when the options were all high in value (*Experiment 1:* t=3.50, p<.005; *Experiment 2:* t=4.80, p<.0001). This finding indicates that overall value might boost confidence, despite paradoxically making choices *less* accurate. More broadly these findings highlight how evidence and confidence, though related, play partially independent roles in the decision making process. Note that all of the predictors analysed in this section were entered into the same hierarchical linear regression; therefore all the effects hold when controlling for the other variables reported.

We also hypothesized that lower-level features of information sampling may reflect an individual's explicit confidence reports. To test this idea, we constructed a novel measure that captured uncertainty in information-sampling behaviour. This new measure, which we label "gaze shift frequency" (GSF), indexes how frequently gaze shifted back and forth among the options presented on the screen. This measure is independent of difference in dwell-time (Experiment 1 *r*=-.02, Experiment 2 *r*=.04): for a constant allocation of time between the options (e.g. 3 seconds for the left-hand option and 5 seconds for the right-hand option) one may shift fixation only once (switching from left to right after 3 seconds have elapsed, for example; low gaze shift frequency) or shift many times between the two options (high gaze shift frequency). We found that GSF was a robust negative predictor of confidence in both experiments (*Experiment 1:* t=-3.67, p<.005; *Experiment 2:* t=-8.94, p<.0001) see Figures 2 A and B. In other words, in trials in which participants shifted their gaze more often between the available options their confidence was lower, even after accounting for changes in |DV| and RT. The four-way relationship between |DV|, RT, GSF and confidence is plotted in Figures 2 C and D. Correlation tables can be found in the supplementary materials (S4).

Figure 2

Confidence predicts change of mind

In both experiments participants saw the same exact choice sets on more than one occasion. In Experiment 1 each pair was presented twice; in Experiment 2 each triad was presented three times (counterbalancing for different spatial locations). This design allowed us to determine factors affecting changes of mind when the same choice is encountered again. Note that the way we define change of mind in this study is different from how it is often defined in perceptual decision-making, as a reversal in an ongoing motor plan due to additional processing of sensory information[4, 15, 29, 30]. The hypothesis we sought to test was that an explicit report of confidence in an initial choice at time t would influence behaviour when the same decision was presented again at a future time t_{future}. In a hierarchical logistic regression, lower confidence at time t was indeed associated with increased changes of mind at time t_{future} in both experiments (*Experiment 1:* z=-6.70, p<.0001; *Experiment 2:* z=-5.71, p<.0001). The effect of confidence in predicting change of mind remained robust after controlling for several other factors that might correlate with the stability of a choice such as |DV| and RT. Because |DV| correlated positively with confidence (see the previous section and S4) we checked the covariance matrices and Variance Inflation Factors (VIFs) to ensure that these correlations did not influence the interpretation of our findings. Both the covariances and the VIFs were below standard thresholds, allowing straightforward interpretation of coefficients (see S5). Furthermore, to rule out the possibility that the effect we observed was driven by the presence of fast errors that were later corrected by the participant, we reanalysed the data excluding all trials that were faster than each participant's mean response time. This analysis produced comparable results (See S6). Notably GSF, itself a correlate of confidence, did not predict change of mind when included in the regression analysis (Fig. 3 A and B, *coefficients in blue)*, even when excluding reported confidence from the regression analysis (see supplementary materials, section S7a). Together, these results suggest that a low-level (and possibly implicit) representation of uncertainty indexed by GSF is insufficient to trigger a future change of mind. On the contrary, individuals may use an explicit representation of uncertainty (expressed through confidence) to reverse their initial decision when the same (or a similar) choice is presented again.

We next harnessed individual differences in metacognition to provide a more stringent test of this hypothesis. We reasoned that the impact of confidence on changes of mind would be more prominent in participants who have enhanced metacognitive skills, i.e. those whose explicit confidence ratings more accurately track the level of uncertainty underlying their decision process. In order to test this hypothesis we calculated an individual index of metacognitive sensitivity by computing the difference in slope between psychometric functions fitted to high and low confidence trials[1, 31, 32]. We then ran a logistic regression to predict changes of mind at time t_{future} using confidence measured at time t. In line with our initial hypothesis, we were able to show that the impact of confidence on changes of mind (here the negative coefficient of confidence predicting change of mind) is stronger in those subjects with greater metacognitive accuracy ($r = -0.35$, $p=0.01$) (Figure 3 C).

Figure 3

Link between confidence and choice transitivity

In the analyses presented above we established a link between an explicit representation of confidence and future changes of mind. However, these analyses are agnostic to the quality of the decisions that emerge as a consequence of changes of mind. Not all choices are born equal; some are more consistent than others, which is formally captured by the notion of transitivity. A transitive ranking is characterized by the following structure: if an option A is preferred over option B and option B is preferred over option C, then it follows that A should be preferred over C (i.e. A≻B and B≻C then A≻C). Transitivity is a normative prescription in utility theory[33]; however, failures of transitivity are commonly observed in human choices and represent a prominent violation of economic rationality and, more generally, of logical consistency[34, 35]. In order to test the relation between confidence and transitivity we found the (idiosyncratic) preference ranking of items that led to the lowest number of transitivity violations for each subject. Finding an optimal ranking of choice sets with more than a handful of items is extremely complex; however, a number of efficient algorithms that approximate a numerical solution have been developed for pairwise comparisons. In our study we used the Minimum Violations Ranking (MVR) algorithm[36] that minimizes the number of inconsistencies in the ranking of items conditional on each participant's choices. This method is conceptually similar to other methods based on revealed preferences such as Afriat's efficiency index[37, 38]. The MVR algorithm provides an optimal ranking of items for each participant so that we could tag choices violating this ranking, hereafter labelled transitivity violations (TV). Because most of these methods are not suited for ternary choice the analyses presented in this section were performed only on data collected for the experiment using binary choice (Experiment 1). An alternative way to assess choice quality is to compute the choice ranking using BDM and test whether participants chose the item with the highest ranking. This method gives qualitatively similar results to those reported below (see S8).

After ordering the participants' choices according to the MVR algorithm, 4.5% of all decisions were classified as transitivity violations. We then split the dataset into trials in which participants reported high confidence and trials in which they reported low confidence (median split). A dramatic reduction in transitivity violations was observed in high confidence trials (16% of transitivity violations) in comparison to low confidence trials (84% of transitivity violations). (Figure 4 A) While these results are consistent with previous evidence provided in this paper and elsewhere[1], note that we did not rely on BDM value estimates (collected post-choice), instead relying only on subjects' choices to generate the optimal ranking. In other words, the link between confidence and the quality of a value-based decision is robust to the method used to elicit preference. In order to statistically quantify the relation between confidence and transitivity violations on a trial-by-trial basis (while accounting for other factors that may result in violations of transitivity) we constructed a set of hierarchical logistic regression models. We found that absolute difference in value (|DV|) was a robust negative predictor of TV ($z=-6.59$, $p<0.0001$; Figure 4 B) such that participants were more likely to violate transitivity when items were closer in value. Critically, this same model showed that even when |DV| was accounted for, confidence was a negative predictor of transitivity violations ($z=-6.75$, $p<.0001$). In other words, participants were less confident during those trials in which they went against their best-fitting preference order. Finally, both response time ($z=2.55$, $p=.01$) and summed value ($z=2.55$, $p=.01$) positively predicted transitivity violations, such that trials in which the value of both options was higher and/or in trials in which their responses were slower, participants' choices were more likely to result in transitivity violations. Similar to the change of mind analysis, eye tracking variables did not reliably predict transitivity violations (GSF=-1.74, $p=.08$; |DDT| $z=-0.47$, $p=.64$) (Figure 5B). Note that this was still true when reported confidence was excluded from the regression analysis (see supplementary materials, section S7b).

Finally, we examined whether intersubject variability in metacognitive ability affected transitivity violations. We reasoned that if a well-calibrated, explicit representation of uncertainty plays a role in guiding future decisions, participants with greater metacognitive ability would show a decrease in the number of transitivity violations when the same option was presented a second time. In line with this hypothesis we observed that greater metacognitive ability was associated with a marked reduction in transitivity violations between the first and second presentation of the same choice (beta=0.85, SE=0.42, $z(26)=2.03$, $p<.05$; Figure 4 C). We also confirmed that this effect was not due to a relationship between metacognition and choice instability: the total number of transitivity violations was unrelated to metacognitive accuracy (beta=-1.83, SE=1.61, z(26)=-1.14, p=0.25). Together these analyses show that a more accurate explicit representation of confidence is associated with more optimal choices when participants are given the opportunity to change their minds.

Figure 4

Discussion

What is the advantage of explicitly representing one's confidence in value-based decision-making? Most experimental setups elicit confidence after a decision has been made and cannot be changed. Our hypothesis was that an explicit representation of confidence might serve an important role in decision-making by signalling the need to explore different alternatives when the same (or a similar) choice is presented again.

Value-based decisions are often perceptually unambiguous (i.e. a banana is noticeably different from an apple) and most of the uncertainty is contingent on a number of internal processes such as memories or homeostatic states that are often difficult to manipulate experimentally. For example, a choice between two food items might be affected both by uncertainty about the tastes of the items and by uncertainty about one's own level of hunger. In order to take advantage of this information, a decision-maker should be able to correctly monitor uncertainty that arises from the different constitutive computations. A wealth of work has shown that humans can introspect on their choice process and report their level of confidence, an ability that has been associated with the psychological concept of metacognition. However, the functions of these explicit representations of confidence (as opposed to implicit markers of uncertainty such as decision time) have remained unclear. Furthermore, individuals show wide variations in how accurately they can track and report fluctuations in uncertainty (i.e. metacognitive accuracy).

In two independent experiments we showed that confidence reports (elicited after a value-based decision) reliably predicted a change of mind when the same choice was presented again. This effect is robust after controlling for other factors associated with the difficulty of a decision, such as difference in value and reaction time. Furthermore, intersubject variability in metacognitive accuracy modulated the degree to which confidence predicted change of mind: confidence was a stronger predictor of change of mind in participants with better metacognitive abilities. Critically, and in contrast to our findings on explicit confidence reports, a lower-level marker of uncertainty (GSF) did not predict subsequent changes of mind, suggesting that an explicit representation of uncertainty expressed through confidence is important for guiding future choices. Instead, we suggest that gaze shift frequency can be considered an ingredient that agents use to construct a subjective sense of certainty, together with decision time and strength of evidence (cf. 3). An alternative interpretation of our results is that gaze shift frequency does not contribute directly to subjective confidence but reflects an agent's attempt to gather more information to adaptively reduce uncertainty (a situation in which confidence would be low and reaction time slow). Future work is required to distinguish between these two hypotheses. A further methodological appeal of GSF as a trial-by-trial measure of uncertainty is that it can be easily gathered in animals. Recent years have seen a resurgence of interest in studying uncertainty and confidence using animal models[39]. This promising line of work relies heavily on the development of experimental paradigms (such as opt-out or post decision wagering) to measure the fluctuation in uncertainty during a decision process. GSF (which can be measured in rodents by tracking head movements) may prove a useful tool to monitor, on a trial-by-trial basis, internal fluctuations in uncertainty and its relation to the neural encoding of decision time and strength of evidence.

Tracking the level of decision uncertainty is helpful in guiding behaviour in a number of contexts; for example, in guiding learning [40], in deciding whether to explore a new alternative or stick with the current one[41, 42], or in evaluating an alternative course of action[18]. At the neural level, the rostrolateral prefrontal (RLPFC) cortex and frontopolar cortex have been shown to play key roles in tracking trial-by-trial evolution of uncertainty[43, 44, 45] and modulating uncertainty-driven behaviours[18, 41, 42, 46, 47, 48]. At the same time, the RLPFC and frontal pole have also been shown (using a number of different methods) to play a key role in enabling metacognitive abilities[1, 10, 11, 14, 32]. It is therefore possible that these two processes are linked anatomically and computationally: individuals whose prefrontal cortex more closely tracks the trial-by-trial evolution of uncertainty might also have more accurate explicit representations of confidence. In turn, superior metacognitive abilities might confer the advantage of knowing how uncertain one's choice was and therefore guide future behavioural strategies, such as uncertainty-driven exploration[42] or changes of mind. Since we did not collect neural measures in this study we cannot test this hypothesis directly, but our findings provide a foundation for future studies of the neurobiology of changes of mind.

Another question we sought to address was: are changes of mind associated with more optimal decisions? In value-based decisions, the difference between correct decisions is often murky since value is a subjective construct. However, when people make a series of value-based choices across a set of options, their pattern of decisions is characterized by a variable degree of internal consistency. In experiment 1 we used a recently developed algorithm to find an optimal ranking of items that produced the lowest number of transitivity violations for each individual. In this way we identified when participants' decisions were inconsistent with their overall (idiosyncratic) pattern of decisions. Violations of transitivity are a paradigmatic example of irrationality in economic choice since they are easy to exploit. For example, when individual preferences are not transitive, it is possible to construct a choice set in which each decision appears fair on its own, but, when combined together, guarantees a loss (a phenomenon known as a Dutch book or arbitrage in finance)[49]. Here we showed that choices made with high confidence are overall more transitive and therefore more optimal according to the normative prescriptions of utility theory. Noticeably, this effect is robust after controlling for the absolute difference in value and reaction time. This finding suggests that individuals can monitor and report that a given decision was noisier and therefore more likely to result in a decision inconsistent with their overall preference patterns, establishing confidence as a correlate of choice accuracy without relying on the BDM procedure to derive independent estimates of subjective utility. This result also resonates with the well-established finding in perceptual decision making that people are able to detect and signal errors as soon as they respond[16, 50] and with the proposal that confidence can facilitate cognitive control[51]. Here, we suggest that a similar process might operate in value-based decisions, in which errors can be thought of as choices that are at odds with one's overall preferences. Consistent with this proposal, we found that individuals who have a more accurate representation of confidence (greater metacognitive ability) were more likely to move towards a more internally consistent decision-making pattern over time. Our work sheds light on the reasons for an explicit representation of confidence in human decision-making. It explores value-based choices (aka economic choices) by borrowing methods and concepts from perceptual decision-making[52]. Similar to perceptual decision-making, we found that the same 'strength of evidence' in value (i.e. |DV|) is accompanied by a variable level of uncertainty that is represented explicitly as confidence. We suggest these representations play a functional role not only in allowing confidence to be shared with others, but also in guiding our own future choices. Taken together our results show that an explicit and accurate representation of confidence can have a positive impact on the quality of future value-based decisions.

Methods

Experimental Procedures

Experiment 1: Participants were required to make binary choices between 16 common snack items. Participants were asked to choose between each combination of the items (N = 120) twice, counterbalanced across the left-right spatial configurations (total number of choices = 240). After each choice, participants indicated their confidence in their decision on a continuous rating scale. Neither choices nor confidence ratings were time constrained. Trial order was randomized with the only constraint being that the same pair was never repeated in subsequent trials. Participants' eye movements were recorded throughout this task.

At the end of the experiment, one choice from this phase was played out and the subject had the opportunity to buy the chosen item by means of an auction administered according to the Becker-DeGroot-Marschak (BDM) procedure: The experimenter randomly extracted a price from a uniform distribution (£0 to £3)—the 'market price' of that item. If the participant's bidding price (willingness-to-pay) was below the market price, no transaction occurred. The computer-generated value was drawn to a precision greater than 2 decimals to avoid the possibility of a tie but was rounded to pennies in the event of a transaction. If the subject's bidding price was above the market price, the participant bought the snack item at the market price[17]. At the end of the experiment, participants had to remain in the lab for an additional hour. During this hour, the only food they were allowed to eat was the item purchased in the auction, if any. At the end of the waiting period participants were debriefed and thanked for their participation. Participants were paid £25 for their time, deducting the cost of the food item, if they bought any. Both tasks were programmed using MATLAB 8.0 (MathWorks) running the Psychophysics toolbox (http://psychtoolbox.org) as well as the Eyelink toolbox extensions[53, 54]. The procedure of this experiment was approved by the UCL Research Ethics Committee (Project ID: 3736/004).

Experiment 2: Participants gave their willingness to pay for 72 common snack food items on a scale ranging from £0-£3, in a BDM procedure[17] similar to the one in experiment 1. Next they completed a choice task where, in each trial, they had to pick their favourite item out of three options. The triplets presented in the choice task were tailored for each participant from their willingness-to-pay ratings. The items were divided into high-value and low-value sets by a median split. The 36 high-value items were randomly combined into 12 high-value triplets; this procedure was mirrored to generate 12 low-value triplets. The high-value and low-value items were then mixed to generate medium value triplets, with 12 triplets consisting of two high-value items and one low-value item, and 12 triplets with the reverse ratio. This resulted in 48 unique triplets, with counterbalanced spatial configurations (total trials =144), split into three blocks. Each triplet was shown once in each block, the presentation order inside blocks was randomized with the constraint that the triplet that ended one block was never shown first in the next block.

In the subsequent choice task, the triplets were presented inside 3 squares in an equidistant 2x2 grid (one randomly-determined position on the grid was left empty). We used a gaze-contingent paradigm in which the items were only visible when the participant fixated inside one of the squares, so that the participant could only see one item at a time. They had unlimited time to make up their mind and could make as many fixations as they wished. After each choice, participants indicated their confidence in their decision on a visual analogue rating scale without any time constraints. Participants' eye movements were recorded throughout the choice task. Both the choice task and the willingness to pay procedure were programmed in Experiment Builder version 1.10.1640, SR-Research, Ontario.

Following the choice task, an auction based on the BDM-ratings took place (see experiment 1). After the auction, participants had to remain in the lab for an additional hour as in experiment 1. At the end of the waiting period participants were debriefed and thanked for their participation. Participants were paid £15 for their time, deducting the cost of the food item, if they bought any. The procedure of this experiment was approved by the University of Cambridge Psychology Research Ethics Committee (Application number: Pre2014.113).

Exclusion Criteria

Because the aim of the experiment was to explore the relationship between confidence and value, it was essential that we had enough measurement sensitivity in both the confidence scale and in the value scale (the BDM ratings), and that participants' choices reflected their stated preferences. We therefore excluded participants if any of the following criteria were met:

1. Participants used less than 25% of the BDM Scale.
2. Participants gave exactly the same BDM rating for more than 25% of the items.
3. Participants used less than 25% of the confidence scale.
4. Participants gave exactly the same confidence rating for more than 25% of their choices.
5. Participant choices did not correspond to their BDM ratings (When predicting choices from differences in value, the DV coefficient deviated more than 2 SD from the experimentwise mean).

Participants

Experiment 1: 30 participants took part in the study. One participant did not complete the task and one participant was excluded because the BDM estimates were poor predictors of his choice (failed criterion 5). Thus 28 participants were included in the analysis (13 females, age: 19-73). All participants were required to fast for four hours prior to taking part in the experiment. Blood glucose levels were taken to test their adherence to this criterion (mean glucose level = 83.57mg/dl, sd = 10.90mg/dl; by comparison, the mean fasting blood glucose levels for adults is 86.4mg/dl[55]). All participants gave informed consent prior to participating in this experiment.

Experiment 2: 30 participants completed the study. Of these 30, three were excluded due to a limited range in their BDM ratings (failed criterion 2). An additional three participants were excluded for a limited range in their use of the confidence scale (failed criterion 4). 24 participants were included in the main analyses (17 females, age: 21-38). All participants were required to fast for four hours prior to doing the experiment. All participants gave informed consent prior to participating in this experiment.

Sample size was determined a-priori. A power estimation was based on previously published work that used a similar experimental setup[26]. We implemented a fixed sample stopping rule set a-priori (N=30). Statistical inferences were conducted only after all data were collected. If a participant did not fulfil one of the exclusion criteria (decided before data collection) would have been excluded from the analysis without replacement.

Eye Trackers

For experiment 1, eye gaze was sampled at 250 Hz with a head-mounted SR Research Eyelink II eye-tracker (SR-Research, Ontario). For experiment 2, eye movements were recorded at 1000Hz with an EyeLink 1000 Plus eye-tracker (SR-Research, Ontario).

Preparation of the Eye-tracking Data

Experiment 1: Areas of Interest (AI) were defined by splitting the screen in half, creating two equal sized areas. Fixations in the left AI were assumed to be directed towards the left snack item, and vice versa. We constructed two variables from the eye tracking data: the difference in dwell time between the two AIs (DDT), and gaze shift frequency (GSF). DDT was calculated by subtracting the total dwell time on the left side from the total dwell time on the right side. GSF was calculated as the number of times participants shifted their gaze from one AI to the other during each trial.

Experiment 2: AIs were pre-defined by the 3 squares that participants had to fixate to view the items (given the gaze-contingent design). We derived two variables from the eye tracking data: the total dwell time in each AI for a given trial, and GSF. Following experiment 1, GSF measured the number of fixations in one AI immediately followed by a fixation in another AI. To ensure that participants paid attention, we excluded trials where participants had not fixated on every option available at least once. 13 trials out of 3457 were excluded from the analysis for this reason.

Hierarchical Models

All hierarchical analyses reported in the results section were conducted using the *lme4* package (version 1.1-7[56]) in R. For the linear models degrees of freedom and p-values were obtained using the Kenward-Roger approximation, as implemented in the *pbkrtest* package[57]. For the choice models (Figures 1 C and F) we ran two hierarchical logistic regressions: In Experiment 1 we predicted the log odds ratio of picking the right-hand option on a given trial; for Experiment 2 we predicted the log odds ratio of picking the reference item. The reference item was determined as the first item encountered according to reading order in Latin languages (i.e. the upper left item for the trials when an item was presented in that position and the upper right item for the remaining trials). Fixed effect confidence intervals were estimated by multiplying the standard errors by 1.96[58]. Because these confidence intervals are estimates that do not take the covariance between parameters into account[59] they should not be interpreted too strictly, but rather serve to give the reader a sense of the precision of the fixed effect coefficients. Note that all predictors reported are z-scored on the participant level, and that all models allow for random slopes at the participant level. For completeness we report coefficients from the full model, while noting that this model is not in every case the most parsimonious. For a comprehensive list of models tested and a formal model comparison using BIC scores see supplementary materials, section S3.

Note that the regression models for confidence in experiment 1 had issues converging. We addressed these issues by square root transforming the |DV| predictor. Notably, for the individual difference analyses investigating change of mind and transitivity we did not implement hierarchical models, but unpooled (individual-level) models. The rationale behind this choice was that for both analyses we were interested in studying between-subjects variation (Figure 3 C and Figure 4 C) that could be potentially affected by the shrinkage of parameters towards the group mean that is characteristic of hierarchical models[60].

Data availability & Code availability

The data and the code for the analyses presented in this article can be found at the BDM Lab GitHub page: https://github.com/BDMLab

The data can also be found on figshare: https://dx.doi.org/10.6084/m9.figshare.3756144.v2

References

1. De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. Confidence in value-based choice. *Nature Neuroscience* **16,** 105–110 (2013) http://doi.org/10.1038/nn.3279
2. Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. Neural correlates, computation and behavioural impact of decision confidence. *Nature* **455**, 227-231 (2008) doi:10.1038/nature07200
3. Kiani, R., Corthell, L., & Shadlen, M. N. Choice Certainty Is Informed by Both Evidence and Decision Time. *Neuron* **84**, 1329–1342 (2014) http://doi.org/10.1016/j.neuron.2014.12.015
4. Van den Berg, R., Anandalingam, K., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. A common mechanism underlies changes of mind about decisions and confidence. *Elife* **5,** e12192 (2016) http://dx.doi.org/10.7554/eLife.12192
5. Lebreton, M., Abitbol, R., Daunizeau, J., & Pessiglione, M. Automatic integration of confidence in

the brain valuation signal. *Nature Neuroscience* **18,** 1159–1167 (2015) http://doi.org/10.1038/nn.4064

6. Barron, H. C., Garvert, M. M., & Behrens, T. E. Reassessing VMPFC: full of confidence?. *Nature neuroscience* **18,** 1064-1066 (2015)

7. Meyniel, F., Sigman, M., & Mainen, Z. F. Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron* **88,** 78–92 (2015) http://doi.org/10.1016/j.neuron.2015.09.039

8. Pouget, A., Drugowitsch, J., & Kepecs, A. Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience* **19,** 366–374 (2016) http://doi.org/10.1038/nn.4240

9. Fleming, S. M., Huijgen, J., & Dolan, R. J. Prefrontal Contributions to Metacognition in Perceptual Decision Making. *Journal of Neuroscience* **32,** 6117–6125 (2012) http://doi.org/10.1523/JNEUROSCI.6489-11.2012

10. Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. Relating Introspective Accuracy to Individual Differences in Brain Structure. *Science* **329** 1541–1543 (2010) http://doi.org/10.1126/science.1191883

11. Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., & Lau, H. Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognitive Neuroscience* **1**, 165–175 (2010) http://doi.org/10.1080/17588921003632529

12. Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. What failure in collective decision-making tells us about metacognition. *Philosophical Transactions of the Royal Society B: Biological Sciences* **367,** 1350–1365 (2012) http://doi.org/10.1098/rstb.2011.0420

13. Bang, D., Fusaroli, R., Tylén, K., Olsen, K., Latham, P. E., Lau, J. Y. F., et al. Does interaction matter? Testing whether a confidence heuristic can replace interaction in collective decision-making. *Consciousness and Cognition* **26**, 13–23 (2014) http://doi.org/10.1016/j.concog.2014.02.002

14. Lau, H., & Rosenthal, D. Empirical support for higher-order theories of conscious awareness. *Trends in cognitive sciences* **15**, 365-373 (2011)

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

16. Yeung, N., & Summerfield, C. Metacognition in human decision-making: confidence and error monitoring. *Philosophical Transactions of the Royal Society B: Biological Sciences* **367**, 1310–1321 (2012) http://doi.org/10.1098/rstb.2011.0416

17. Becker, G. M., DeGroot, M. H., & Marschak, J. Measuring utility by a single-response sequential method. *Behavioral Science* **9,** 226–232 (1964)

18. Boorman, E. D., Behrens, T. E. J., Woolrich, M. W., & Rushworth, M. F. S. How Green Is the Grass on the Other Side? Frontopolar Cortex and the Evidence in Favor of Alternative Courses of Action. *Neuron* **62,** 733–743 (2009) http://doi.org/10.1016/j.neuron.2009.05.014

19. FitzGerald, T. H. B., Seymour, B., & Dolan, R. J. The Role of Human Orbitofrontal Cortex in Value Comparison for Incommensurable Objects. *Journal of Neuroscience* **29,** 8388–8395 (2009) http://doi.org/10.1523/JNEUROSCI.0717-09.2009

20. Lebreton, M., Jorge, S., Michel, V., Thirion, B., & Pessiglione, M. An Automatic Valuation System in the Human Brain: Evidence from Functional Neuroimaging. *Neuron* **64,** 431–439 (2009) http://doi.org/10.1016/j.neuron.2009.09.040

21. Levy, I., Lazzaro, S. C., Rutledge, R. B., & Glimcher, P. W. Choice from Non-Choice: Predicting Consumer Preferences from Blood Oxygenation Level-Dependent Signals Obtained during Passive Viewing. *Journal of Neuroscience* **31,** 118–125 (2011) http://doi.org/10.1523/JNEUROSCI.3214-10.2011

22. Plassmann, H., O'Doherty, J. P., & Rangel, A. Appetitive and Aversive Goal Values Are Encoded in the Medial Orbitofrontal Cortex at the Time of Decision Making. *Journal of Neuroscience* **30,** 10799–10808 (2010) http://doi.org/10.1523/JNEUROSCI.0788-10.2010

23. Krajbich, I., & Rangel, A. Multialternative drift-diffusion model predicts the relationship between

visual fixations and choice in value-based decisions. *Proceedings of the National Academy of Sciences of the United States of America* **108,** 13852–13857 (2011) http://doi.org/10.1073/pnas.1101328108

24. Carandini, M., & Heeger, D. J. Normalization as a canonical neural computation. *Nature Reviews Neuroscience* **13,** 51-62 (2011) http://doi.org/10.1038/nrn3136

25. Louie, K., & Khaw, M. W. Normalization is a general neural mechanism for context-dependent decision making. *Proceedings of the National Academy of Sciences of the United States of America* **110,** 6139-6144 (2013) http://doi.org/10.1073/pnas.1217854110/-/DCSupplemental

26. Soltani, A., De Martino, B., & Camerer, C. A Range-Normalization Model of Context-Dependent Choice: A New Model and Evidence. *PLoS Computational Biology* **8,** e1002607 (2012) http://doi.org/10.1371/journal.pcbi.1002607.t001

27. Krajbich, I., Armel, C., & Rangel, A. Visual fixations and the computation and comparison of value in simple choice. *Nature Neuroscience* **13,** 1292–1298 (2010)

28. Kiani, R., & Shadlen, M. N. Representation of Confidence Associated with a Decision by Neurons in the Parietal Cortex. *Science* **324,** 759–764 (2009) http://doi.org/10.1126/science.1169405

29. Bronfman, Z. Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., & Usher, M. Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences* **282,** 20150228 (2015) http://doi.org/10.1098/rspb.2015.0228

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

31. de Gardelle, V., & Mamassian, P. Does Confidence Use a Common Currency Across Two Visual Tasks? *Psychological Science* **25,** 1286–1288 (2014) http://doi.org/10.1177/0956797614528956

32. Fleming, S. M., & Lau, H. C. How to measure metacognition. *Frontiers in Human*

Neuroscience, **8,** 1-9 (2014) http://dx.doi.org/10.3389/fnhum.2014.00443

33. Neumann, Von, J., Morgenstern, O., Rubinstein, A., & Kuhn, H. W. Theory of games and economic behavior. (Princeton Univ Press, 2007)
34. Camerer, C. F., & Ho, T.-H. Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty* **8,** 167–196 (1994) http://doi.org/10.1007/BF01065371
35. Loomes, G., Starmer, C., & Sugden, R. Observing violations of transitivity by experimental methods. *Econometrica: Journal of the Econometric Society* **59,** 425-439 (1991) http://doi.org/10.2307/2938263
36. Pedings, K. E., Langville, A. N., & Yamamoto, Y. A minimum violations ranking method. *Optimization and Engineering*, **13,** 349–370 (2011) http://doi.org/10.1007/s11081-011-9135-5
37. Afriat, S. N. Efficiency estimation of production functions. *International Economic Review*, 568-598 (1972)
38. Varian, H. R. Goodness-of-fit in optimizing models. *Journal of Econometrics* **46,** 125-140 (1990)
39. Kepecs, A., & Mainen, Z. F. A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences* **367,** 1322–1337 (2012) http://doi.org/10.1098/rstb.2012.0037
40. Meyniel, F., Schlunegger, D., & Dehaene, S. The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLoS Computational Biology* **11,** e1004305 (2015) http://doi.org/10.1371/journal.pcbi.1004305
41. Daw, N. D., O'Doherty, J. P., Dayan, P., Seymour, B., & Dolan, R. J. Cortical substrates for exploratory decisions in humans. *Nature* **441** 876–879 (2006) http://doi.org/10.1038/nature04766
42. Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. Rostrolateral prefrontal cortex and individual differences in uncertainty-driven exploration. *Neuron* **73,** 595-607 (2012)

http://doi.org/10.1016/j.neuron.2011.12.025

43. Koechlin, E., & Hyafil, A. Anterior Prefrontal Function and the Limits of Human Decision-Making. *Science* **318,** 594–598 (2007) http://doi.org/10.1126/science.1142995

44. Rushworth, M. F. S., & Behrens, T. E. J. Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience* **11,** 389–397 (2008) http://doi.org/10.1038/nn2066

45. Yoshida, W., & Ishii, S. Resolution of Uncertainty in Prefrontal Cortex. *Neuron* **50,** 781–789 (2006) http://doi.org/10.1016/j.neuron.2006.05.006

46. Donoso, M., Collins, A., & Koechlin, E. Foundations of human reasoning in the prefrontal cortex. *Science* **344,**1481-1486 (2014) http://doi.org/10.1126/science.1252254

47. Lee, S. W., Shimojo, S., & O'Doherty, J. P. Neural Computations Underlying Arbitration between Model-Based and Model-free Learning. *Neuron* ***81,*** 687–699 (2014) http://doi.org/10.1016/j.neuron.2013.11.028

48. Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. The Neural Representation of Unexpected Uncertainty during Value-Based Decision Making. *Neuron* **79,** 191–201 (2013) http://doi.org/10.1016/j.neuron.2013.04.037

49. Hájek, A. Dutch book arguments. *The Oxford Handbook of Rational and Social Choice*, 173–195. (Oxford University Press, 2008)

50. Boldt, A., & Yeung, N. Shared Neural Markers of Decision Confidence and Error Detection. *The Journal of Neuroscience* ***35,*** 3478–3484 (2015) http://doi.org/10.1523/JNEUROSCI.0797-14.2015

51. Koizumi, A., Maniscalco, B., & Lau, H. Does perceptual confidence facilitate cognitive control?. *Attention, Perception, & Psychophysics* **77,** 1295-1306 (2015) http://doi.org/10.3758/s13414-015-0843-3

52. Summerfield, C. Building bridges between perceptual and economic decision-making: neural and computational mechanisms. *Frontiers in neuroscience* **6,** 70 (2012) http://doi.org/10.3389/fnins.2012.00070/abstract

53. Brainard, D. H. The psychophysics toolbox. *Spatial Vision* **10,** 433–436 (1997)

54. Cornelissen, F. W., Peters, E. M., & Palmer, J. The Eyelink Toolbox: eye tracking with MATLAB and the Psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers* **34,** 613–617 (2002)

55. Rapp, K., et al. Fasting blood glucose and cancer risk in a cohort of more than 140,000 adults in Austria. *Diabetologia* **49,** 945-952 (2006)

56. Bates, D., Maechler, M., Bolker, B., & Walker, S. lme4: Linear mixed-effects models using Eigen and S4. *R Package Version*, **1** (2014)

57. Halekoh, U., & Højsgaard, S. A kenward-roger approximation and parametric bootstrap methods for tests in linear mixed models–the R package pbkrtest. *Journal of Statistical Software* **59,** 1–30 (2014)

58. Gelman, A., & Hill, J. Data analysis using regression and multilevel/hierarchical models. (Cambridge University Press, 2006)

59. Bolker, B. How trustworthy are the confidence intervals for lmer objects through the effects package? http://stats.stackexchange.com/questions/117641/how-trustworthy-are-the-confidence-intervals-for-lmer-objects-through-effects-pa (2014, accessed 10th December 2015)

60. Gelman, A., & Pardoe, I. Bayesian measures of explained variance and pooling in multilevel (hierarchical) models. *Technometrics* **48,** 241-251 (2006)

Acknowledgements

This work has been supported by the Wellcome Trust and Royal Society (Henry Dale Fellowship to BDM – 102612/Z/13/Z) and the Economics and Social Research Council (PhD Scholarship for TF). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. We would like to specially thank Yoshitsugu Yamamoto for sharing the methods he developed to rank choice in Experiment 1 and Chris Street and Sebastian Bobadilla Suarez for help in collecting the data and pre-processing the eye-tracking raw data used in Experiment 1.

Author Contribution Statement

BDM, CJ and SF *designed the first experiment reported in this paper. The data for the first experiment was collected by CJ. The second experiment was designed by TF and BDM. The data for the second experiment was collected by TF, the data from both experiments was analysed by TF. The article was written by BDM and TF. All the authors revised the manuscript*

Competing Interests Statement

The authors declare no competing interests.

Figure legends

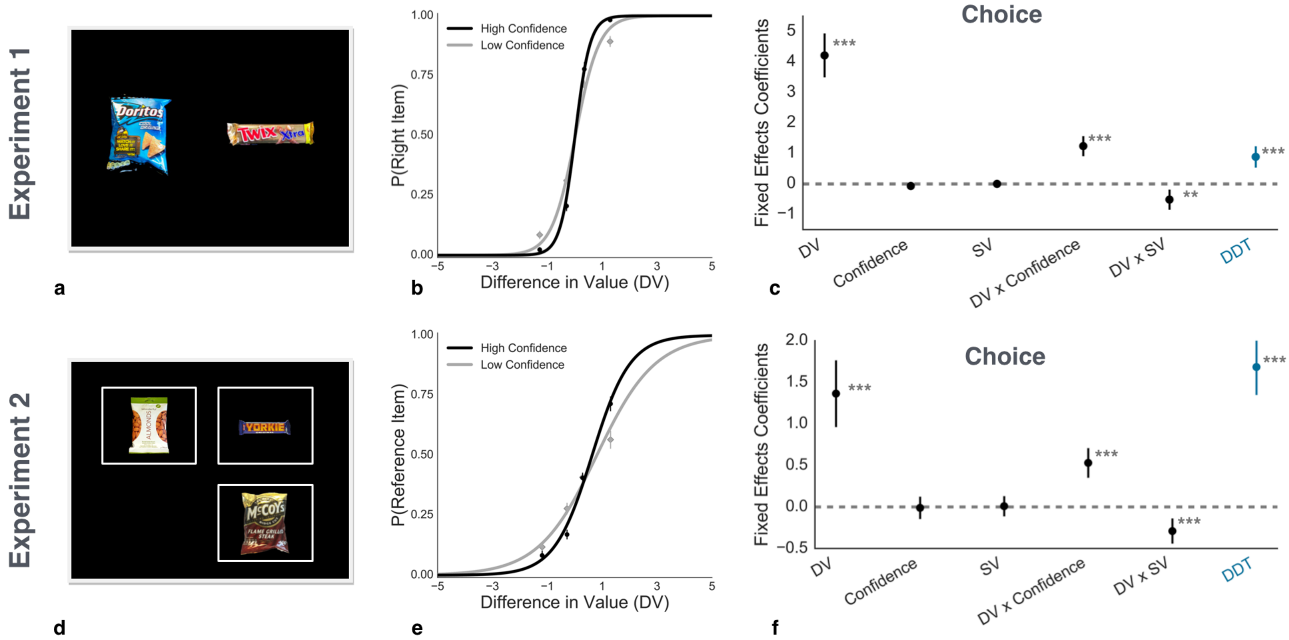

Figure 1: ***Relation between confidence and choice***

Eye-tracking tasks: ***(a)*** *In Experiment 1, participants were presented with two snack items and were then required to choose one item to consume at the end of the experiment.* ***(d)*** *In Experiment 2, participants chose between three options, and the presentation of the stimuli was contingent on which box participants looked at. In both experiments, participants indicated their confidence that they had made a correct decision on a visual analogue scale after each choice they made.* ***(b)*** *Probability of choosing the item on the right as a function of the difference in value between the options, data from.* ***(e)*** *Probability of choosing the reference item (see methods), as a function of the value difference between the reference item and the mean value of the alternatives. Black line, high confidence trials, grey line, low confidence trials (as determined by a median split). Each graph shows the z-scored data pooled across participants. Points represent quartiles of DV. Error bars show standard errors.* ***(c and f)*** *Fixed effects coefficients from hierarchical logistic regression* *models predicting choice (DV= difference in value; SV= summed value; DDT = difference in dwell time, DV x Confidence= Interaction of difference in value and confidence; DV x SV= Interaction of difference in value and summed value). The graph for experiment 1 (C) shows the coefficients predicting the probability of choosing the right-hand option; the graph for experiment 2 (F) shows the coefficients predicting the probability of choosing the reference option (see Methods). Error bars show 95% CIs. The sample size for experiment 1 was 28 participant (each completing 240 trials), the sample size for experiment 2 was 24 participants (each completing 144 trials) *** = p < .001; ** = p < .01; * = p < .05 (two-sided).*

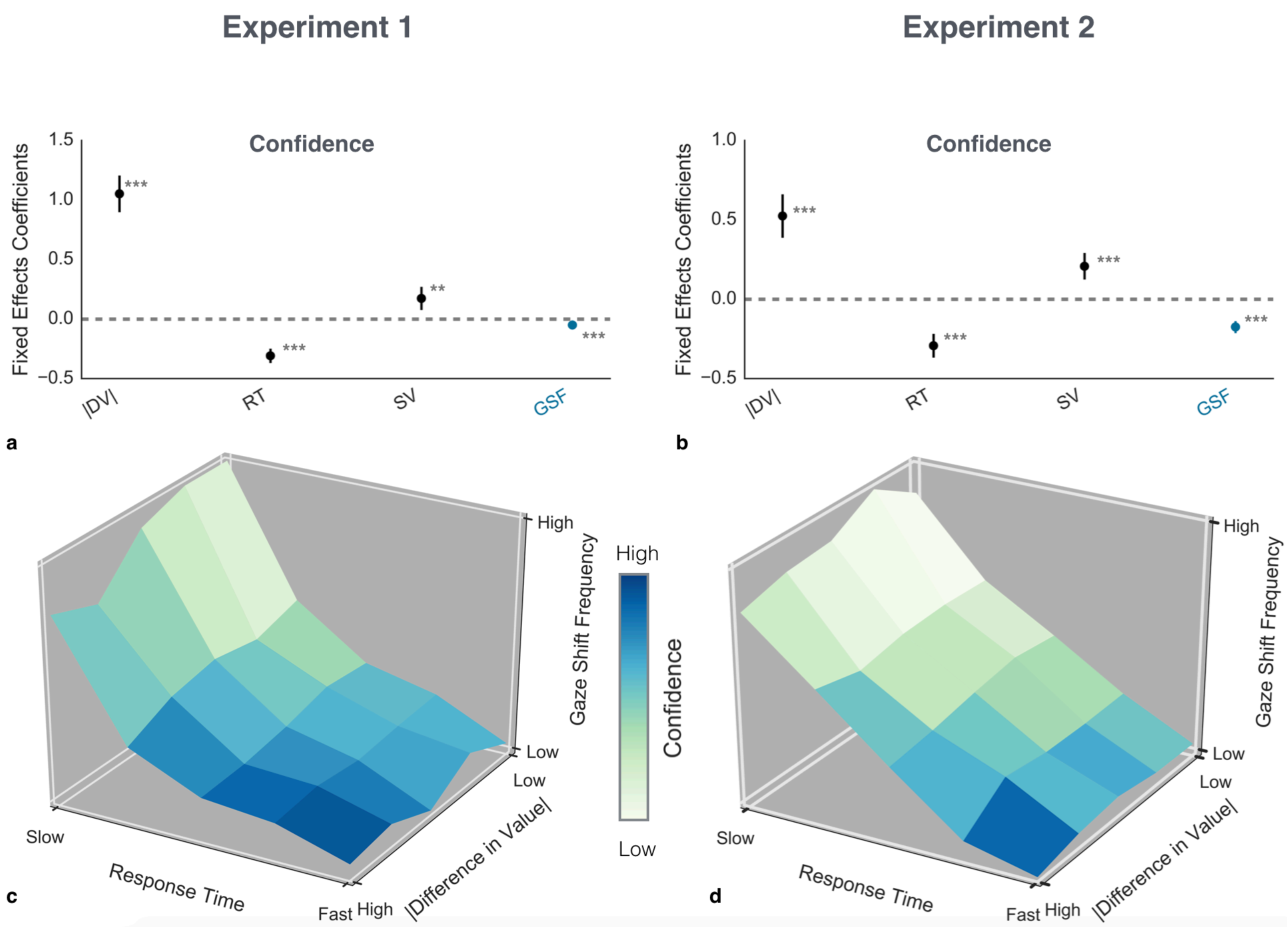

Figure 2. Factors that contribute to confidence

*__(a-b)__ Fixed-effect coefficients in hierarchical regression models predicting confidence for experiment 1 and 2, respectively. Error bars show 95% CIs. *** = p < .001; ** = p < .01; * = p < .05 (two-sided). (IDVI= absolute difference in value; RT = reaction time; SV= summed value; GSF = Gaze Shift Frequency). __(c–d)__ 4-D heat maps showing mean z-scored confidence as a function of subject specific quantiles of response time, absolute difference in value and gaze shift frequency. The sample size for experiment 1 was 28 participants (each completing 240 trials); the sample size for experiment 2 was 24 participants (each completing 144 trials)*

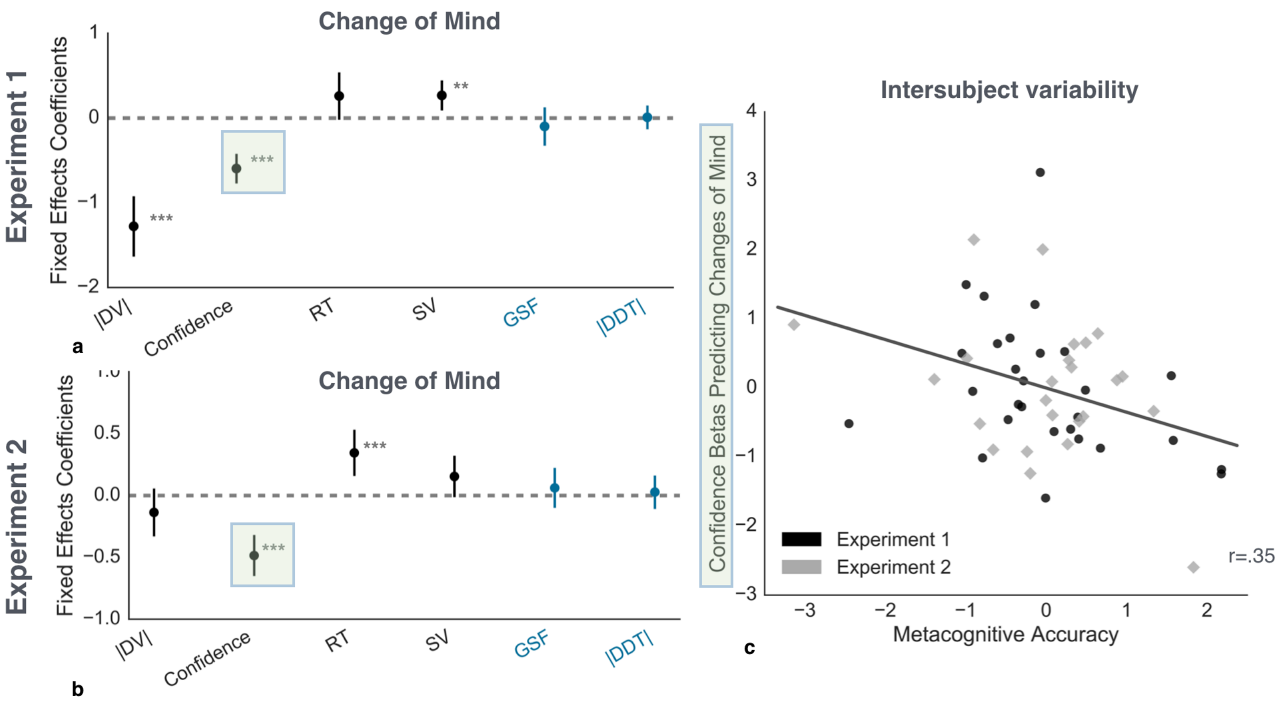

Figure 3. Confidence predicts change of mind.

(a-b) *Fixed effects coefficients from hierarchical logistic regression models predicting future changes of mind. Error bars show 95% CIs. *** = p < .001; ** = p < .01; * = p < .05 (two-sided). (|DV|= absolute difference in value; RT = reaction time; SV= summed value; GSF = gaze shift frequency; |DDT|= absolute difference in dwell time)* ***(c)*** *Correlation between metacognitive accuracy and the coefficients for confidence ratings predicting future changes of mind (highlighted in pale green). Participants with greater metacognitive accuracy are more likely to change their mind following a low-confidence judgment; note that the correlation is negative because the relationship between confidence and changes of mind is itself negative (lower confidence increases the probability of subsequent changes of mind). Participants from experiment 1 are represented by black dots; participants from experiment 2 are represented by grey diamonds. Both axes (x and y) are z-scored for each experiment separately. The sample size for experiment 1 was 28 participants (each completing 240 trials); the sample size for experiment 2 was 24 participants (each completing 144 trials)*

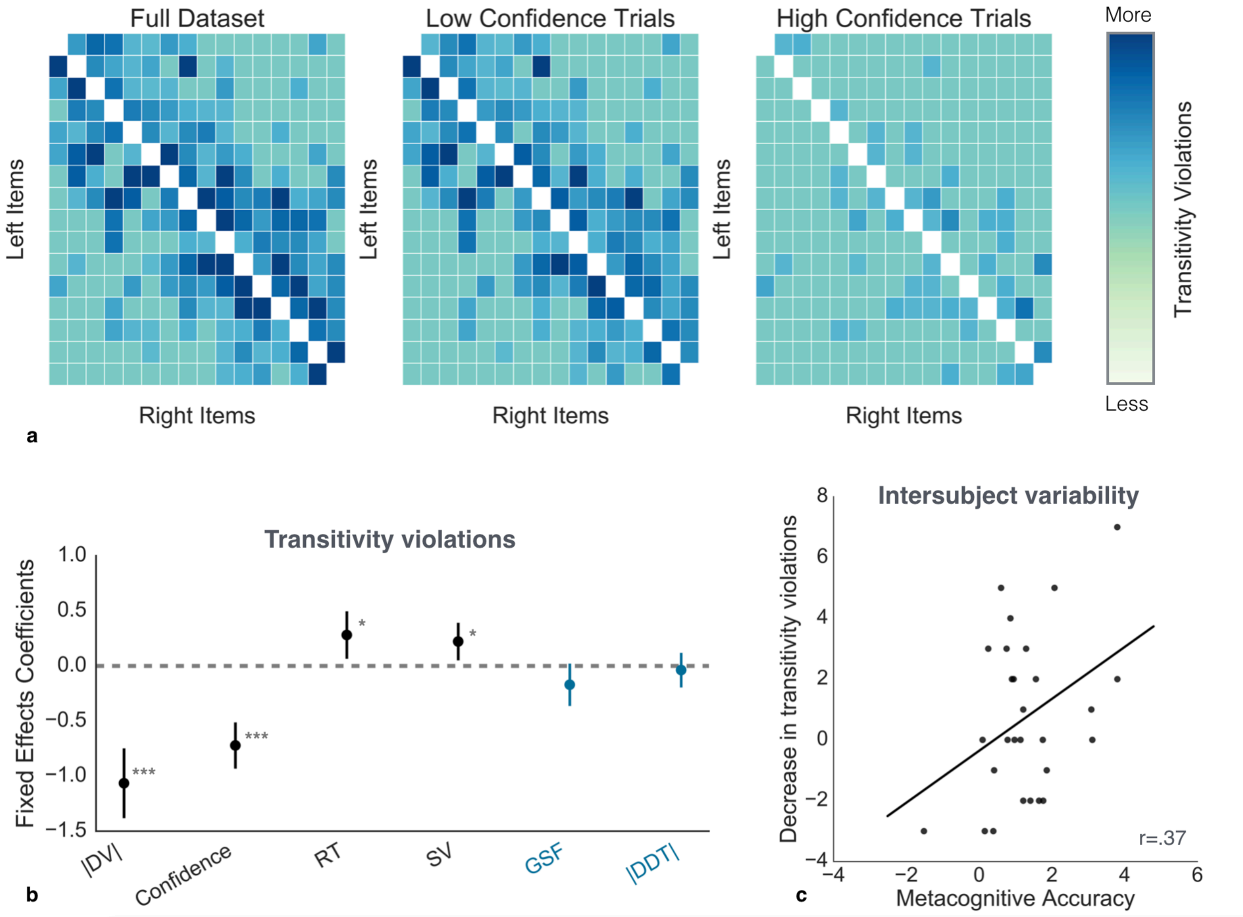

Figure 4. Link between confidence and transitivity

(a) *Heat maps showing the number transitivity violations for the full sample and for high and low confidence trials (median split). The middle diagonal line is empty because no item was ever paired with itself. Note most transitivity violations took place on low-confidence trials.* ***(b)*** *Fixed effects coefficients from a hierarchical logistic regression model predicting transitivity violations Error bars show 95% CIs. *** = p < .001; ** = p < .01; * = p < .05 (two-sided). (lDVl= absolute difference in value; RT = reaction time; SV= summed value; GSF = gaze shift frequency; lDDTl= absolute difference in dwell time).* ***(c)*** *Decreases in transitivity violations between the first and second presentation for each participant, as a function of metacognitive accuracy. The graph shows that participants who are more metacognitively accurate tend to become more transitive over time. The sample size for experiment 1 was 28 participants (each completing 240 trials).*