



Confidence is predicted by pre- and post-choice decision signal dynamics

John P. Grogan^{a,b}, Wouter Rys^a, Simon P. Kelly^c, Redmond G. O'Connell^{a,b}

^aSchool of Psychology, Trinity College Dublin, Dublin, Ireland.

^bTrinity College Institute of Neuroscience, Trinity College Dublin, Dublin, Ireland.

^cSchool of Electrical and Electronic Engineering and UCD Centre for Biomedical Engineering, University College Dublin, Dublin, Ireland.

Corresponding Author: John P. Grogan (john.grogan@tcd.ie)

ABSTRACT

It is well established that one's confidence in a choice can be influenced by new evidence encountered after commitment has been reached, but the processes through which post-choice evidence is sampled remain unclear. To investigate this, we traced the pre- and post-choice dynamics of electrophysiological signatures of evidence accumulation (Centro-parietal Positivity, CPP) and motor preparation (μ /beta band) to determine their sensitivity to participants' confidence in their perceptual discriminations. Pre-choice CPP amplitudes scaled with confidence both when confidence was reported simultaneously with choice, and when reported 1 second after the initial direction decision with no intervening evidence. When additional evidence was presented during the post-choice delay period, the CPP exhibited sustained activation after the initial choice, with a more prolonged build-up on trials with lower certainty in the alternative that was finally endorsed, irrespective of whether this entailed a change-of-mind from the initial choice or not. Further investigation established that this pattern was accompanied by later lateralisation of motor preparation signals toward the ultimately chosen response and slower confidence reports when participants indicated low certainty in this response. These observations are consistent with certainty-dependent stopping theories according to which post-choice evidence accumulation ceases when a criterion level of certainty in a choice alternative has been reached, but continues otherwise. Our findings have implications for current models of choice confidence, and predictions they may make about EEG signatures.

Keywords: decision-making, confidence, ERP, CPP, meta-cognition

1. INTRODUCTION

Mathematical modelling and neurophysiological investigations of perceptual decision making suggest that choice confidence evolves over the course of deliberation, and is informed by several factors, including the strength of evidence favouring the chosen alternative (Bang & Fleming, 2018), the difference in evidence between the available options (Li & Ma, 2020), and the time taken to decide (Kiani et al., 2014). Confidence judgements can also be updated after a decision has been made, taking account of evidence accumulated while the decision-reporting action is still being executed (Resulaj et al., 2009) and/or new evidence encountered after response completion (Fleming

et al., 2018; Moran et al., 2015; van den Berg et al., 2016; Yu et al., 2015). What form such post-commitment accumulation processes take remains unclear, and several alternative possibilities have been proposed.

According to some models, confidence levels are updated by a continuation of the same evidence accumulation process that informed the initial choice (Moran et al., 2015; Pleskac & Busemeyer, 2010; Yu et al., 2015) but other accounts invoke a distinct metacognitive process that evaluates the accuracy of the preceding choice (Desender, Ridderinkhof, et al., 2021; Fleming & Daw, 2017). Another key point of distinction between existing models of confidence in perceptual decisions is whether

Received: 13 June 2023 Accepted: 16 June 2023 Available Online: 19 July 2023



The MIT Press

© 2023 Massachusetts Institute of Technology.
Published under a Creative Commons Attribution 4.0
International (CC BY 4.0) license.

Imaging Neuroscience, Volume 1, 2023
https://doi.org/10.1162/imag_a_00005

they assume that post-choice evidence accumulation operates until confidence is probed (Yu et al., 2015), or is terminated when a criterion level of confidence (Pleskac & Busemeyer, 2010) or elapsed time (Desender, Donner, et al., 2021) has been reached ("optional stopping"). Adjudicating among these accounts has been difficult as post-choice evidence accumulation has rarely been observed or probed in the brain with sufficient temporal precision. Electrophysiological research in monkeys using opt-out paradigms has established that choice confidence can be read-out jointly from the pre-choice firing rate of decision-variable encoding neurons and deliberation time (Kiani & Shadlen, 2009), but this work has not yet examined the neural underpinnings of post-choice confidence reports.

One promising signal for probing post-choice confidence representations in humans is the Centro-parietal Positivity (CPP), which tracks sensory evidence accumulation during decision formation (O'Connell & Kelly, 2021). Whereas effector-selective decision signals previously reported in humans and other species reach a stereotyped amplitude immediately prior to response execution, the CPP's amplitude varies with several factors known to influence choice confidence, for example, having greater amplitudes on correct trials and reduced amplitudes for trials with longer RTs (Kelly et al., 2021; Steinemann et al., 2018). Thus far, only a few studies have directly investigated the CPP's relationship with confidence, and mainly using stimulus-locked CPPs which are potentially confounded by faster RTs for higher confidence responses. Most of these studies reported that CPP reaches a higher amplitude after evidence onset on trials rated highly confident (Davidson et al., 2021; Gherman & Philiastides, 2015, 2018; Herding et al., 2019; Tagliabue et al., 2019) apart from one study which found no such link (Rausch et al., 2020). Additionally, one study that used response-locked CPPs also found no confidence effect (Feuerriegel et al., 2022).

Elsewhere, it has been reported that a post-choice centro-parietal signal with identical topography to the CPP, but traditionally labelled as the Error Positivity (Pe), also scales with choice confidence but in the opposite direction to that reported for the pre-choice CPP. The Pe is seen after erroneous choices that are explicitly reported to be incorrect (Falkenstein, 1990; Nieuwenhuis et al., 2001; Steinhauser & Yeung, 2010), and its amplitude increases the more confident participants are that they have made an error (Boldt & Yeung, 2015; Feuerriegel et al., 2022). The Pe has also been shown to exhibit a build-to-threshold relationship with error signalling reports (Murphy et al.,

2015), and to predict subsequent post-error slowing and post-choice information-seeking (Desender, Boldt, et al., 2019; Desender, Murphy, et al., 2019).

While the Pe literature seems to suggest an "error accumulation" process that selectively gathers evidence indicating that a preceding choice should be reversed (Boldt & Yeung, 2015; Desender, Donner, et al., 2021), much of this research involved studies with no post-choice evidence available (Boldt & Yeung, 2015; Desender, Donner, et al., 2021) or where participants only responded when errors were detected (Murphy et al., 2015). Additionally, recent work has highlighted that the use of baseline correction to an interval just before the initial response may have caused pre-choice amplitude differences to be transferred to post-choice amplitude measurements (Feuerriegel et al., 2022). Finally, some studies modelling delayed confidence responses have assumed that people accumulate during the entire post-choice evidence presentation (Pleskac & Busemeyer, 2010; Yu et al., 2015), while other models propose a time- or confidence-dependent stopping rule (Desender, Donner, et al., 2021; Moran et al., 2015). By allowing accumulation to be traced while post-choice evidence remains available, EEG offers a means of testing for early stopping during the delay period between choice and confidence reports. To our knowledge, there has not been a direct comparison of the post-commitment dynamics of neural evidence accumulation signals with versus without post-choice evidence presentation.

We present two experiments investigating the relationship between CPP and confidence. Experiment 1 used simultaneous confidence and choice reports to investigate whether previous confidence effects on the stimulus-locked CPP would be found in the response-locked signal. It also served as a benchmark for Experiment 2, in which participants withheld their confidence reports for a 1-second post-commitment delay, and randomly varied across trials whether the physical evidence remained on screen or was extinguished. Pouget et al. (2016) distinguish between the related concepts of certainty and confidence, where certainty derives from the perceived probability distribution of stimulus variables irrespective of choice, while confidence is the probability that a choice—whether overtly completed or still covertly evolving—is correct. In tasks with a single choice between two stimulus alternatives, as in Experiment 1, these cannot be distinguished (Bang & Fleming, 2018); however, when confidence is reported after the initial choice and allows for changes of mind, as in Experiment 2, we can cast that report in two alternative ways to aid our analyses: as "confidence in initial choice," whose lowest value is

attained when certain that an error was made, or as “final certainty,” which ranges from “maybe” to “certain” regarding current perception irrespective of the initial choice. Our results across these two experiments indicate that when there is no further evidence provided, the CPP amplitude at the time of initial choice scales with the confidence in that choice. However, when physical evidence remains available during the interval between the choice and confidence reports, post-choice accumulation of this evidence terminates early when participants achieve high final certainty, irrespective of whether or not this entails a reversal of their initial choice.

2. METHODS

2.1. Ethics

The study was approved by Trinity College Dublin ethics committee and carried out in accordance with the Declaration of Helsinki, and EU GDPR. Written informed consent was given before the start of the first session.

2.2. Participants

Participants were between 18 and 32 years old, with normal or corrected-to-normal vision, no history of neurological or psychiatric disorders, epilepsy, or unexplained fainting. Experiment 1 recruited 27 participants (13 females, 14 males), with two excluded following artefact rejection ($N = 25$, see below). Experiment 2 recruited a different group of 30 participants (18 females, 12 males), five of whom were excluded due to insufficient data retained after artefact rejection ($N = 25$). Participants were paid for their time (€35 in Experiment 1, €45 in Experiment 2), plus a bonus depending on their performance (Experiment 1: up to €6.50, mean = €4.40, $SD = 1.10$; Experiment 2: up to €16, mean = €8.50, $SD = 1.50$).

2.3. Experimental design

2.3.1. Experiment 1—simultaneous choice and confidence reports

Experiment 1 was a random dot-motion direction discrimination task in which participants simultaneously reported the direction of coherent motion and their choice confidence (Fig. 1). The task was programmed in MATLAB (R2013b) and Psychtoolbox-3 (Kleiner et al., 2007). Testing took place in a dark, sound-attenuated room, with participants seated 57 cm from a CRT monitor (51 cm, 65.2 cd/m^2 luminance, 75 Hz refresh rate, 1024 x 768 resolution), with their head on a chin rest.

Trial onsets were self-paced, beginning when participants pressed the space button. A white central fixation cross was presented for 400 ms, followed by the random dot kinematogram composed of 75 white dots (0.16° diameter) in an aperture (8° diameter) on a grey background (fixation cross remained on-screen until feedback). The dots moved with zero coherence for an initial lead-in period of 1500 ms to prevent visual-evoked potentials elicited by stimulus onset from overlapping with choice-related signals. A proportion of the dots began moving coherently after this 1500 ms lead-in, with the proportion individually titrated to achieve a criterion discrimination accuracy level (see below). Dot positions were updated on every frame, with a proportion (matching the coherence) randomly selected to move either left or right on each trial (equal probability) by 0.2° relative to their location three frames earlier, to give an overall motion speed of 15°s⁻¹. The remaining dots were moved to a random new location every frame, and the coherent dots were re-selected every frame to prevent people tracking individual dots. The coherent motion was presented for 350 ms, 500 ms, or 750 ms (equal probability), and was followed by the appearance of a 6-point response scale. Participants rested the first three fingers of each hand on the response keys of a keyboard and were instructed to click the “s,” “d,” or “f” key with their left hand to indicate “certain left,” “probably left,” or “maybe left,” respectively, and clicking “j,” “k,” “l” with their right hand to report “maybe right,” “probably right,” “certain right,” respectively. Participants were instructed to withhold reporting their choices until the appearance of the confidence scale, and only responses within 0–1500 ms of scale onset were recorded. Visual feedback was then provided for 500 ms in the form of “correct,” “error,” “too fast” ($RT < 0$ ms), or “too slow” ($RT > 1500$ ms). Participants were shown their mean accuracy and response time at the end of each block and were informed of their bonus winnings at the end of the experiment (see scoring rule in Experiment 2 section).

Participants completed two consecutive days of testing, with training on the behavioural task and task difficulty titration taking place on day 1. Training started with 50 trials at high coherence, until participants performed close to 100% accuracy on the direction decision and were comfortable rating their confidence simultaneously. The coherence was then titrated in blocks of 30 trials to achieve approximately 75% discrimination accuracy using a staircase procedure that increased the coherence 1% after every error and decreased it by 1% after three consecutive correct responses (average titrated coherence 11.19%, $SD = 4.66$). The EEG testing session took place on day 2




Fig. 1. Task designs for experiments 1 and 2. Experiment 1 was a dot motion direction discrimination task, with simultaneous direction and confidence reporting. Each trial started with an initial 1500 ms lead-in of incoherent motion to allow visual-evoked potentials associated with stimulus onset to resolve, offering a clear view on choice-related signals. Participants withheld reporting their choice and confidence until the appearance of a 6-point confidence scale following stimulus offset (350/500/750 ms stimulus duration). In Experiment 2, participants were presented with overlaid gratings and reported whether the grating tilted to the left or right was higher in contrast. Again, trials started with an initial zero-evidence lead-in period during which the gratings appeared at equal contrast. Initial left/right choices had to be reported within a 1500 ms deadline. Participants were then probed to report their choice confidence using the same scale as in Experiment 1 after a post-choice delay of 1000 ms. During the delay period, the evidence either remained on screen (as shown in the figure) or was extinguished with equal probability. Note that stimulus sizes in all panels are not to scale.

and consisted of 8 blocks of 72 trials with a short rest break in between blocks. Only the data from day 2 are included in the analyses reported below.

2.3.2. Experiment 2—delayed confidence reports with or without post-choice evidence

Experiment 2 (Fig. 1) was a contrast discrimination task in which participants made an initial speeded two-alternative choice and were subsequently cued to report their final choice confidence after a 1-second delay. In randomly interleaved trials, evidence was either extinguished immediately after the initial choice report or continued throughout the delay period. The same physical set-up was used as for Experiment 1, except the monitor was a 40.5 cm CRT monitor.

After pressing the space bar to begin the trial, participants saw a central fixation for 250 ms, followed by two overlaid gratings tilted at 45° from vertical in each direction (spatial frequency = 2 cycles per degree) in an annulus

shape around a central fixation (inner radius = 1° , outer radius = 6° ; the fixation-cross remained on-screen until the confidence-scale appeared). The two gratings were initially presented at 50% contrast. Evidence onset occurred after 1000 ms, with one grating increasing by a criterion amount (see below) and the other grating decreasing by the same amount. To allow reverse correlation analyses (analyses not reported here), we introduced small frame-to-frame variations in the contrast-difference between the two stimuli with values drawn from a Gaussian noise distribution ($SD = 1.5\%$ which corresponds to around 10% of the mean evidence strength, maximum variation was capped at 3 SD). Throughout the stimulus presentation, the two gratings were flickered at 15 Hz in anti-phase, which allowed us to recover a 15 Hz steady-state visual evoked potential (SSVEP) driven by the contrast difference between alternating gratings, whose phase indicated the direction of the difference. In addition, activity at 30 Hz indexed the summed visual response to the two gratings, offering a general index of engagement. Participants indi-

cated which grating had the higher contrast with a button press (“f” or “j” for left or right) as soon as they liked, with a deadline of 1500 ms. Following the initial response, the evidence either Continued at the same mean strength or was Extinguished and replaced with a fixation cross. These trial types were randomly interleaved and occurred with equal probability. 1000 ms after the initial response, the confidence rating scale was shown on the screen, and participants responded as in Experiment 1—importantly, the scale still ran from “certain left” to “certain right.” Participants were instructed to report their final choice confidence in light of all the evidence they had viewed, as opposed to retrospectively reporting on their confidence in their initial choice. Feedback was given as in Experiment 1.

Participants completed 16 blocks of 80 trials while undergoing EEG recordings, with half of the trials completed on the first day of testing, and half on the second (consecutive) day. These two days of data were combined for the analyses. The first day also included training and staircasing, using a two-down one-up staircase to reach 70% accuracy (step-size starting at 6% and decreasing by 1% point each step until they reached 1%; mean contrast = 13.19%, SD = 4.61). A bonus was paid out based on their accuracy and confidence in the main task, based on a quadratic scoring rule (points = $100 * (1 - (\text{initial accuracy} - \text{confidence-in-initial-choice})^2)$), scaled to €0–€16. Initial accuracy in this equation is coded as 1 for correct and 0 for incorrect, and confidence-in-initial-choice is expressed relative to initial choice on each trial, taking one of 6 equally spaced values ranging from 0, corresponding to “certain” in the initially unchosen option, to 1, corresponding to “certain” in the initially chosen option. Given there is no guarantee of additional evidence being presented after the initial response, this scoring rule incentivises initial accuracy as well as the accuracy of confidence responses (Staël von Holstein, 1970). It also orthogonalises confidence from expected reward; more points are given for high-confidence than low-confidence initially correct responses, while, at the same time, more points are given for low-confidence than high-confidence initially error responses. We checked that participants understood this rule by asking them how many points they would get if they incorrectly responded “left” and then pressed “certain left” (the correct answer to which is zero).

2.4. EEG acquisition and pre-processing

EEG was recorded with a BioSemi ActiveTwo system (BioSemi, Netherlands), with 128 scalp electrodes at 512 Hz. Vertical EOG was recorded from electrodes above

and below the left eye. Data were processed and analysed using custom MATLAB scripts that drew on routines from EEGLab (Delorme & Makeig, 2004). EEG data were linearly detrended across each session, low-pass filtered at 40 Hz (FIR filter), and epoched in intervals of -1000:3000 ms relative to evidence onset. Trials were baseline-corrected to a period after the initial zero-evidence stimulus appeared, but before the evidence appeared (the 400 ms before coherent motion onset in Experiment 1, and the 200 ms before the contrast-difference in Experiment 2). The same pre-evidence baseline was used for pre- and post-choice CPP analyses in Experiment 2, except for the pre-response baseline investigation which used the 100 ms before initial RT (see Supplementary Materials). Channels with extreme variance or high artefact counts were interpolated (mean = 4.48 channels, SD = 2.88, range = 0-12).

Epochs with any scalp electrode voltages over 100 μV , or with bipolar VEOG voltages over 200 μV (250 μV in Experiment 2) between the baseline and the response onset (confidence cue onset in Experiment 2) were flagged as artefactual and removed. Participants with more than half their trials removed were excluded from the analysis (2 in Experiment 1, 5 in Experiment 2). Experiment 1 had a mean of 490.56 out of 576 trials included (SD = 63.30, range = 337-567), and Experiment 2 had a mean of 929.88 out of 1280 trials included (SD = 151.65, range = 653-1242). Response-aligned epochs were also extracted from the evidence-aligned epochs, using the interval -1000 ms:500 ms relative to response in Experiment 1, and -1000 ms:1000 ms in Experiment 2. The voltages were transformed into Current Source Density (CSD) using CSDToolbox (Kayser & Tenke, 2006) with the default parameters ($\lambda = 1 \times 10^{-5}$, m-constant = 4; see Supplementary Materials for non-CSD key data).

2.5. Analysis

For both experiments, trials with initial choice reaction times less than 100 ms were excluded (Experiment 1: mean = 10.12, SD = 12.69, range = 0-52; Experiment 2: mean = 2.04, SD = 3.31, range = 0-12). Because in Experiment 2 participants were able to change their minds between the initial choice and the final confidence report, we scored their final reports in two different ways. Firstly, to test the hypothesis that the post-choice CPP reflects a distinct accumulation process selectively gathering evidence calling for a revision of the initial choice, we compared waveforms according to “confidence-in-initial-choice” as defined above. This provides a directional measure of the

participants' finally reported confidence in their earlier choice, allowing for changes-of-mind (CoM). For example, if the subject initially reported "right," then a final report of "certain right" corresponds to the highest level of confidence-in-initial-choice ("certain no-CoM"), while "certain left" corresponds to the lowest value ("certain CoM") and the "maybe" reports lie in the middle. In later analyses, we also tested whether the CPP scaled with the level of confidence they had in their final choice, irrespective of whether this involved a CoM or not, a factor we labelled "final certainty" to reflect its independence of the initial choice.

Data were analysed with trial-wise linear mixed models (LMM), after z-scoring all variables. The regression coefficients from these models are standardised by the z-scoring, and the degrees of freedom represent the total number of trials minus the degrees taken up by the factors included in the model. Generalised LMM were used for logistic regression when binary variables such as accuracy or change-of-mind were the dependent variable. Reaction times (RTs) were measured from stimulus offset in Experiment 1 (to account for different stimulus durations), and from evidence onset in Experiment 2. Experiment 2 final RTs were measured from the confidence-scale onset. Experiment 1 RTs and Experiment 2 final RTs were log-transformed for statistical analysis (Experiment 2 initial RTs were approximately normal). For Experiment 1, Stimulus Duration was included in the LMM. The main factors of interest were initial accuracy and final certainty (maybe, probably or certain), and Experiment 2 also had Post-choice Evidence Condition (Continued or Extinguished) and Change-of-mind. We ran a model comparison to select the best random-effects structure to use, using a backward-selection likelihood ratio test method ($\alpha_{LRT} = 0.2$) starting from the maximal random-effects structure (Matuschek et al., 2017). We ran this once, using the pre-response CPP amplitude effect in Experiment 2, with fixed-effects of Post-choice Evidence Condition and final certainty, and the best fitting model was one with only random intercepts (see Supplementary Materials). Rather than running this selection in each of the 36 main LMMs and the >200 LMMs used when analysing the time-bins across each waveform figure, we used this intercept-only random-effect for all LMMs. This keeps the interpretation of the models consistent across analyses.

"Maybe" responses were less common than the others, especially when evidence was extinguished in Experiment 2, so a minimum number of 10 trials for each Final-Certainty*Evidence combination was applied. This led to the exclusion of "maybe" extinguished-evidence trials from 5 participants (35 trials in total; their other trials were kept in).

Similarly, 7 participants in Experiment 1 had fewer than 10 trials for some Final-Certainty*Stimulus Duration combinations, so these trials were excluded (46 in total). Other trials from these participants were kept in the analysis, as the LMM allows for missing data—excluding these participants entirely did not change the pattern of results.

The added stochastic contrast variation in Experiment 2 intended for reverse correlation analyses was unable to replicate previous effects of initial choice accuracy, indicating that the variance was too low, and so also had no detectable effects of confidence ratings, so is not included here.

2.6. EEG signals

To identify appropriate electrodes for measuring the CPP, we examined the grand-average ERP topographies (i.e., averaged across all included trials, thus across all confidence ratings) and covering the time range from -150:-50 ms before the initial choice report. A cluster of 5 Centro-parietal channels centred on the focus of the CPP topography were selected and, for each individual, we identified the channel within this cluster that exhibited the largest pre-choice amplitude in order to extract CPP measurements. The pre-choice mean amplitude was taken within -150:-50 ms before initial response in Experiment 1, and -140:-50 ms for Experiment 2 (in order to capture an integer number of cycles of the SSVEP signal). We used these same electrodes for the post-choice CPP, and examining this indicated that this signal dropped partially toward baseline soon after the initial choice report before undergoing a second build-up in advance of the presentation of the confidence cue, which differed between Continued and Extinguished Post-choice Evidence Conditions towards the end of the delay (see Supplementary Fig. 1). Based on this observation, we measured the amplitude of the post-choice CPP in the window 700:1000 ms after the initial choice. A 10 Hz low-pass filter is applied for figure plotting only in order to remove the 15 Hz and 30 Hz SSVEP components. We additionally examined the effect of a pre-response baseline (-100:0 ms before initial response) on the post-choice CPP both in our time-window, and an earlier window (200:350 ms) as reported in a previous paper (Feuerriegel et al., 2022; see Supplementary Materials). In addition to the *a priori* mean amplitude windows analysed, we analysed mean amplitude within each 100 ms bin with LMM, across the 1000 ms before initial-response and the 1000 ms between initial-response and final response in Experiment 2. Effector-selective motor preparation was analysed via mu/beta band

(8-30 Hz) activity. Fast-Fourier transforms were performed on the CSD data, with 256-sample windows (~500 ms) moving in 8-sample steps (20 ms). Clusters of three electrodes in each hemisphere were chosen, centred on the strongest ipsilateral minus contralateral pre-response (-150:-50 ms) amplitude in the grand-average topography. Within these clusters, the channel with the largest difference in amplitude was chosen for each person. Mean amplitudes were calculated for the interval -250:-100 ms before initial choice in Experiment 1, and -300:0 ms in Experiment 2, along with a post-choice measurement in Experiment 2 from 700:1000 ms (i.e., -300:0 ms before confidence cue appears).

In Experiment 2, the two gratings each flickered at 15 Hz in anti-phase generating distinct phase-tagged contrast-dependent SSVEP responses to each grating. When both gratings are at equal contrast, the anti-phase ensures that the signals for each grating will cancel out on the scalp. When one grating has higher contrast, the 15 Hz signal will be more strongly phase-aligned to that grating, providing an index of the encoding of the sensory evidence for this experiment (i.e., the differential contrast of the two gratings). The 15 Hz phase-tagged signals were calculated by convolving the CSD data with a 15 Hz sinewave (171 samples length, 334 ms, 5 cycles), and the phase was projected onto the mean phase 400 ms after the initial response, when the 15 Hz signal was steady. The real component of this projection reflects the strength of the alignment with this phase, with negative values corresponding to the opposite phase, which we term the *differential SSVEP*. Mean topographies were used to pick the best channel from a cluster centred on Oz, per person. We took the mean value of this differential SSVEP from 700:1000 ms after initial response. We could only analyse the continued evidence condition, as there was no post-choice visual stimulus when evidence was extinguished.

The 30 Hz SSVEP, reflecting overall stimulus engagement to both stimuli, was taken from the same Fourier transformed data described above, from the Oz electrode, and normalised to the adjacent frequencies. The mean amplitude was taken -400:-200 ms before the initial response, and 700:1000 ms after it.

3. RESULTS

3.1. Pre-choice CPP predicts choice accuracy and simultaneously reported confidence

In Experiment 1, participants made two-alternative forced choice dot motion direction discrimination decisions and

simultaneously reported their confidence in that choice (Fig. 1). As expected, Accuracy ($\beta = 0.15$, $t(11765) = 6.75$, $p < .0001$) and Confidence ($\beta = 0.06$, $t(11765) = 7.73$, $p < .0001$) increased with Stimulus Duration, as shown in Figure 2a-d.

RT relative to stimulus offset decreased with longer Stimulus Durations ($\beta = -0.23$, $t(11759) = -33.32$, $p < .0001$), and with greater Confidence ($\beta = -0.11$, $t(11759) = -15.60$, $p < .0001$), and was faster for correct responses ($\beta = -0.33$, $t(11759) = -43.47$, $p < .0001$). There were also interactions of Accuracy and Stimulus Duration ($\beta = -0.03$, $t(11759) = -4.51$, $p < .0001$), and Accuracy and Confidence ($\beta = -0.03$, $t(11759) = -4.80$, $p < .0001$), as the effects of Stimulus Duration and Confidence on RT were stronger on correct trials.

As previously observed in studies of perceptual choice, the CPP exhibited a steady build-up during deliberation (Fig. 2e), reaching a peak just before response execution. The CPP reached a significantly higher amplitude prior to Correct responses compared to erroneous ones ($\beta = 0.04$, $t(11759) = 3.93$, $p < .0001$), and on trials with higher reported Confidence ($\beta = 0.05$, $t(11759) = 4.88$, $p < .0001$) with no significant interactions ($p > .5$). Post-hoc contrasts indicated that CPP amplitude was significantly larger for “certain” compared to “maybe” ($F(1,11761) = 8.87$, $p = .0029$) and “probably” ($F(1,11761) = 5.32$, $p = .0211$) confidence ratings, but did not differ reliably between “maybe” and “probably” confidence ratings ($F(1,11761) = 0.39$, $p = .53$).

3.2. Pre-choice CPP only predicts delayed confidence reports when evidence is extinguished following commitment

In Experiment 2, participants made two-alternative forced choice contrast discriminations and reported their confidence after a delay of 1000 ms (Fig. 1), during which the physical evidence either Continued on-screen or was Extinguished. Discrimination accuracy increased from initial-choice to final choice ($\beta = 0.15$, $t(42494) = 11.90$, $p < .0001$, Fig. 3a) with a significant Response (initial vs final) by Post-choice Evidence (Continued vs Extinguished) interaction ($\beta = 0.05$, $t(42494) = 4.12$, $p < .0001$), as accuracy increased more when evidence presentation Continued throughout the delay period.

We first analysed the participants' confidence-in-initial-response (i.e., whether their delayed confidence report showed a change-of-mind or not). Reported confidence-in-initial-choice was higher following a Correct initial choice (Fig. 3c; $\beta = 0.45$, $t(21245) = 80.02$, $p < .0001$), and lower when Evidence presentation Continued throughout the

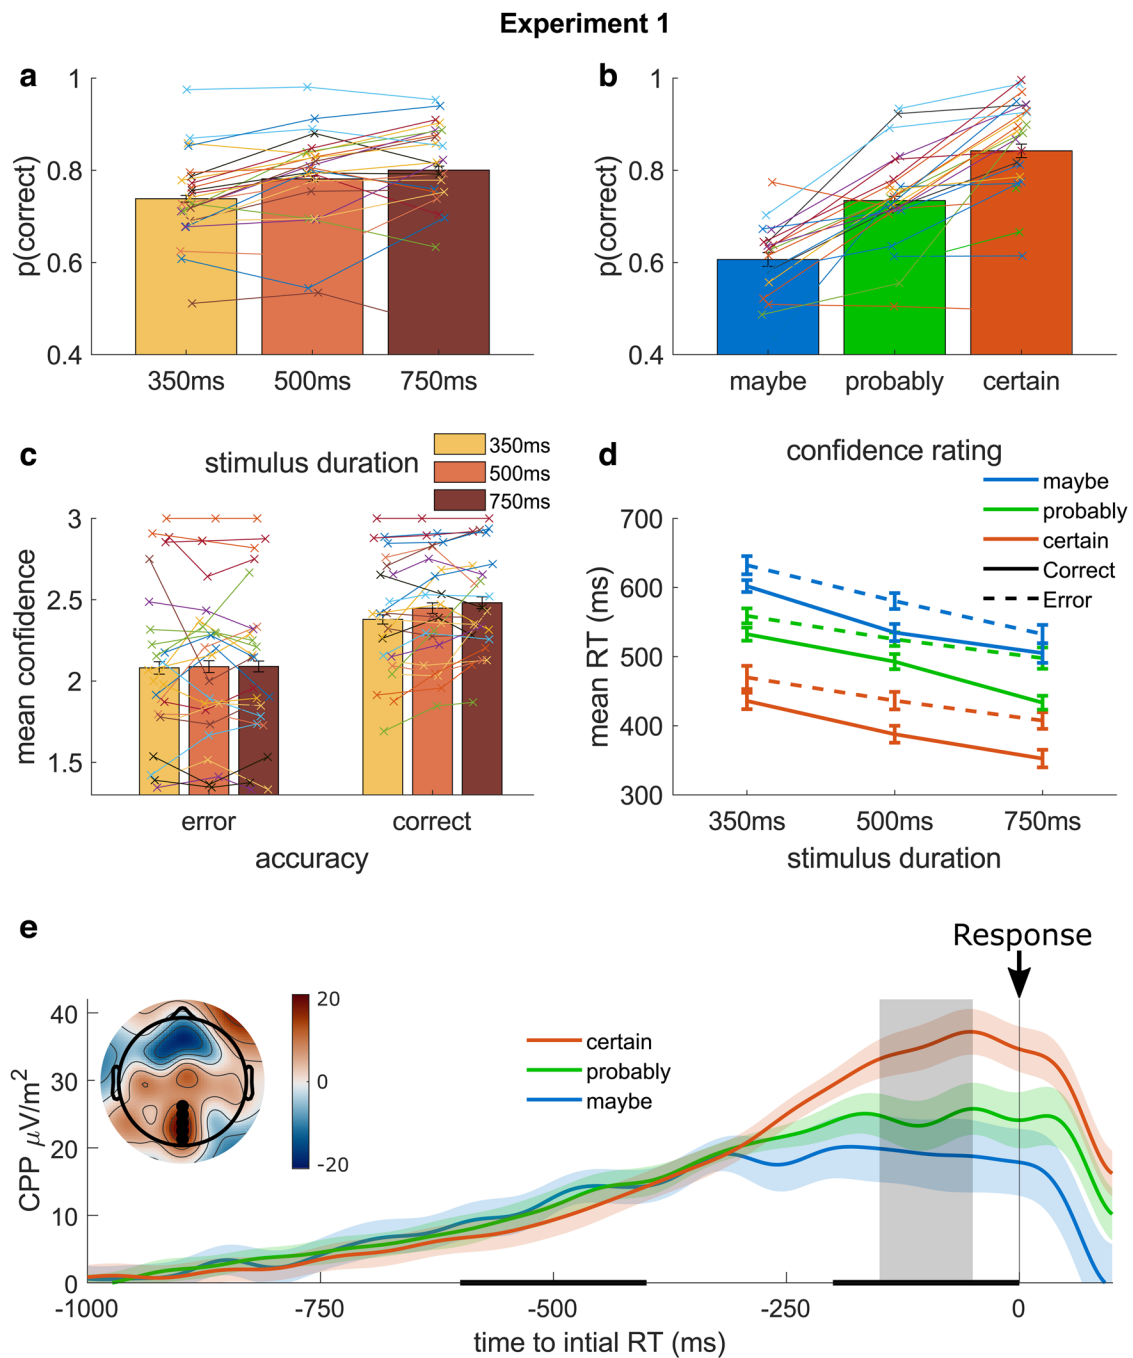


Fig. 2. Pre-choice CPP predicts confidence reported simultaneously with direction in Experiment 1. (a) Mean accuracy increases with stimulus duration. (b) Mean accuracy is higher for trials rated higher confidence. (c) Mean confidence rating (1 = maybe, 2 = probably, 3 = certain) increases with stimulus duration only for correct responses. (d) Mean post-offset RT is quicker for longer stimulus durations, and for higher confidence ratings, but is slower for incorrect trials (dashed lines). (e) Mean pre-choice CPP traces in Experiment 1, leading up to response time, split by Confidence (within each stimulus duration and averaged across durations). Experiment 1 pre-choice CPP amplitude within the grey time-window is higher for “certain” responses (black bar = effect of Confidence, $p < .05$ within each 100 ms time-bin). The inset topography shows the mean activity within the time-window (red = positive, blue = negative) for the grand-mean over all confidence levels, with the black dots showing the electrodes used for CPP selection.

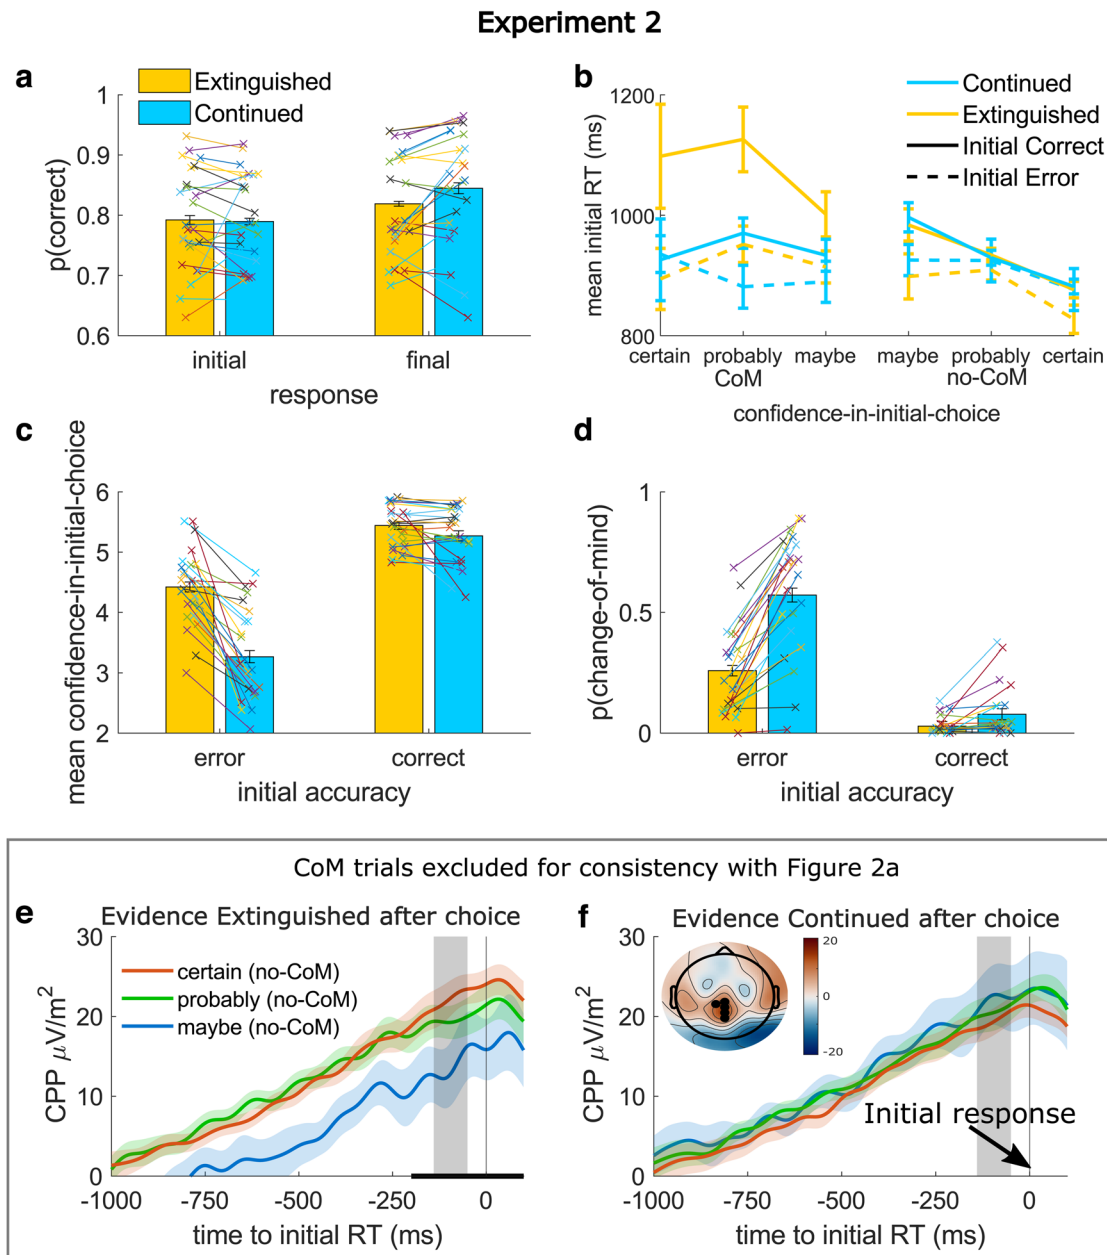


Fig. 3. Pre-choice CPP in Experiment 2 predicts delayed confidence-in-initial-choice only if evidence is extinguished after initial response. (a) Choice accuracy increases between initial and final responses, with a much greater increase if evidence continues. (b) Mean RT for initial responses is quicker for trials rated higher confidence-in-initial-choice (i.e., dependent on whether this involved a change-of-mind or not), slower for correct responses, with an interaction of the two. (c) Mean confidence-in-initial-choice (6 = certain no-CoM, 1 = certain CoM) is higher after correct responses, but slightly lower if evidence continues. (d) Changes-of-mind are far more common following errors, and continued evidence increases changes-of-mind, with greater effect on error trials. (e-f) To allow better comparison with Figure 2e, trials with changes-of-mind are excluded from the waveforms in panels (e) & (f) only (see Supplementary Fig. 2 for those trials). (e) Pre-choice CPP amplitude is greater on trials later rated as higher confidence-in-initial-choice, but only when the evidence is extinguished after the initial response (black bar shows significant effects of confidence on mean amplitude within 100 ms time-bins, $p < .05$). (f) If post-choice evidence instead continues, this disrupts the link between pre-choice CPP and the delayed confidence-in-initial-choice rating. The inset topography shows the grand-mean activity within the time-window across all trials (both evidence conditions, all no-CoM confidence ratings), with the black dots showing the electrodes used for CPP selection (same scale as inset in Fig. 2).

delay period ($\beta = -0.13$, $t(21245) = -23.53$, $p < .0001$), with an interaction of the two ($\beta = 0.15$, $t(21245) = 26.29$, $p < .0001$) as there was much greater confidence resolution (i.e., the difference between confidence to correct and error trials) when Evidence Continued ($\beta = .57$, $t(12104) = 71.18$, $p < .0001$) than was Extinguished ($\beta = 0.29$, $t(9141) = 39.27$, $p < .0001$).

Investigating this trend further, we found that Continued evidence led to a significant increase in change-of-mind rates (Fig. 3d; $\beta = .54$, $t(16930) = 12.40$, $p < .0001$) following correct initial response ($\beta = .54$, $t(16930) = 12.40$, $p < .0001$) as well as erroneous ones ($\beta = .80$, $t(4315) = 20.88$, $p < .0001$), although the effect was significantly greater for initial errors (Accuracy by Post-choice Evidence interaction ($\beta = -0.10$, $t(21245) = -4.34$, $p < .0001$). Thus, while post-choice evidence mainly causes corrective changes-of-mind leading to greater final accuracy overall, it also causes a reduction in confidence ratings and increase in choice reversals following correct initial responses (Fig. 3d). This observation likely stems from the fact that, based on confidence levels reported for correct trials in Experiment 1 and in the extinguished evidence condition of Experiment 2, participants were likely close to the maximum level of confidence that could be reported at the time of initial commitment, reducing the scope for any further measurable increases.

Turning to the CPP, we first sought to test the relationship between the delayed confidence reports and its amplitude measured immediately prior to the initial choice. To allow direct comparison with the pre-choice CPP effects observed in Experiment 1 (where participants rated confidence alongside their choice and therefore could not report changes-of-mind), we examined pre-choice CPP amplitude as a function of confidence-in-initial-choice, but excluded any trials with changes-of-mind—thus, it reflects the confidence relative to the initial choice, but rated at the final response (similar effects were observed when changes-of-mind were included, see Supplementary Fig. 2 for statistics). Consistent with the results reported for Experiment 1, pre-choice CPP amplitude increased with confidence-in-initial-choice ($\beta = 0.01$, $t(18512) = 2.06$, $p = .0394$), although the effect of choice Accuracy did not reach significance ($\beta = 0.01$, $t(18512) = 1.02$, $p = .31$). There was also a Confidence-in-initial-choice*Post-choice Evidence interaction ($\beta = -0.02$, $t(18512) = -2.24$, $p = .0251$), and separate LMMs in each Post-choice Evidence Condition indicated that pre-choice CPP amplitude increased with confidence-in-initial-choice when evidence was Extinguished during the

delay period (Fig. 3e; $\beta = 0.04$, $t(8485) = 2.74$, $p = .0062$), but not when Evidence Continued ($p = .9699$; Fig. 3f). This accords with the observation that post-choice evidence had a substantial influence on the confidence reports—if decision processes are updated during the post-choice interval, then this would naturally reduce the degree to which pre-choice CPP amplitudes would predict the confidence level reached by the end of the trial, even if no change-of-mind occurred.

3.3. Post-choice CPP scales inversely with final choice certainty irrespective of change-of-mind

Next, we looked at how the CPP evolves during the post-choice delay period in Experiment 2 (Fig. 4), including in change-of-mind trials. When evidence was Extinguished, the grand-average CPP returned gradually back to baseline over the next 1000 ms (Supplementary Fig. 1) but, when evidence presentation continued, the CPP decayed at a markedly slower rate overall, and exhibited a positive build-up during the delay period on certain trial types.

To match analyses conducted in previous studies of confidence-related modulations of post-choice ERP signals (e.g., Boldt & Yeung, 2015; Feuerriegel et al., 2022), we first analysed post-choice CPP amplitudes (measured at the end of the delay period) as a function of confidence relative to the initially chosen option (“confidence-in-initial-choice”). Here, we observed an interaction of Post-choice Evidence and confidence-in-initial-choice ($\beta = -0.02$, $t(21073) = -3.084$, $p = .0020$), and separate LMM in each evidence condition showed this was driven by higher post-choice CPP amplitudes for lower confidence-ratings when evidence presentation Continued during the delay period ($\beta = -.03$, $t(12014) = -3.59$, $p = .0003$), an effect that was absent when evidence was Extinguished ($\beta = 0.00$, $t(9059) = 0.11$, $p = .91$). The absence of any relationship with confidence on Extinguished trials appears to be in direct disagreement with results reported by Boldt and Yeung (2015). However, that study used the interval immediately prior to the initial choice for its baseline correction which could cause pre-choice amplitude differences to be transferred to post-choice measurements. When we applied the same pre-choice baseline correction, we found the same pattern of results as Boldt and Yeung, with the post-choice CPP now decreasing with confidence-in-initial-choice in the Extinguished condition both in our measurement window and in the one used by Boldt and Yeung (Supplementary Fig. 3). This last observation also indicates that,

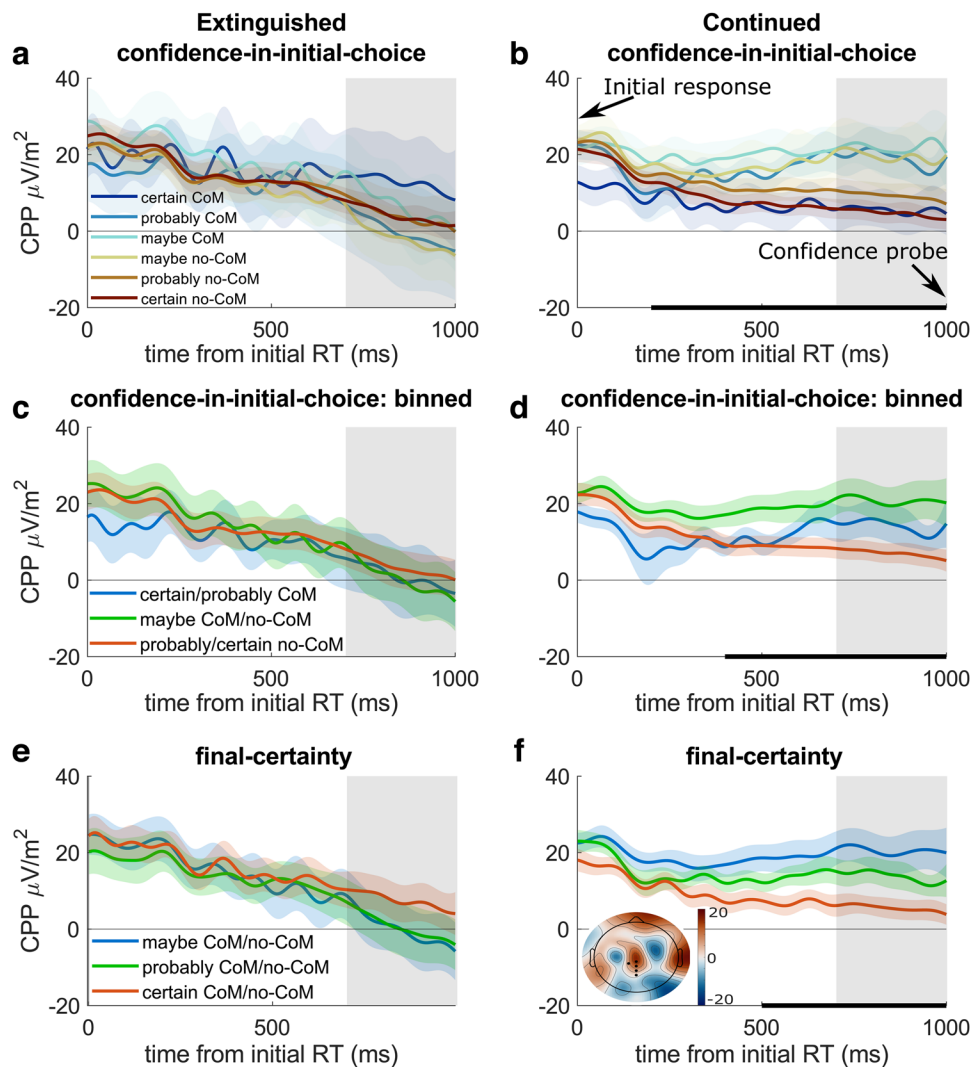


Fig. 4. Post-choice CPP is sensitive to post-choice evidence and increases when participants have low final certainty. Post-choice CPP refers to the CPP time-course between the initial choice response and the final choice with accompanying confidence 1000 ms later in Experiment 2. (a) CPP waveforms as a function of the 6-point “confidence-in-initial-choice” rating, taking into account whether the trial included a change-of-mind (CoM) or not, indicating whether the participant believed the initial response was wrong or correct. The mean CPP (700:1000 ms) decreases after the initial choice if the evidence is Extinguished at that time, with no significant difference depending on confidence-in-initial-choice. (b) When Evidence Continued after the initial response, the post-choice CPP plateaued, and rose again on trials where subjects’ final ratings translated to their initial choice being “maybe” or “probably CoM.” The effect of confidence-in-initial-choice became significant from about 400 ms after the initial decision (solid black bars show $p < .05$ in 100 ms time-windows). (c) The same data as above, but with confidence-in-initial-choice binned in pairs to increase the trial-numbers within one waveform; again, Extinguished trials have no differences, but (d) Continued trials show a significant effect, which is non-linear, as the medium bin (maybe CoM & no-CoM) is highest. (e) Extinguished evidence trials did not differ by final certainty (i.e., maybe/probably/certain rating for the final option chosen, regardless of CoM or no-CoM) responses. (f) Continued-evidence trials had greater post-choice CPP amplitudes on trials rated “maybe” than “probably” or “certain.” The topography inset in panel (f) shows the mean amplitude within the grey window for the low final-certainty trials (red = positive, same scale as Figures 2 & 3, black dots are CPP electrodes), and the solid black bars along the bottom show which 100 ms time-windows have significant effects in the CPP waveforms ($p < .05$, none were significant for the Extinguished conditions).

while the present study used confidence scales that were mapped to the choice alternatives (left vs right) and Boldt and Yeung (2015) used a scale mapped to the accuracy of the initial choice (correct vs incorrect), these differences did not alter the post-choice dynamics of the CPP qualitatively, at least in the Extinguished condition.

Although there was an effect of confidence-in-initial-choice on CPP amplitude when evidence presentation continued, Figure 4b shows that CPP amplitudes were very similar for “certain CoM” and “certain no-CoM” trials where the difference in confidence-in-initial-choice is greatest ($t(22) = 0.119$, $p = .907$, $BF_{01} = 4.543$, indicating moderate evidence for the null). CPP amplitudes were also highly similar for “maybe CoM” compared to “maybe no-CoM” trials ($t(24) = 0.140$, $p = .890$, $BF_{01} = 4.701$, indicating moderate evidence for the null) with differences only apparent between “probably CoM” and “probably no-CoM” ($t(23) = 2.247$, $p = .035$, $BF_{01} = 0.573$ indicating anecdotal evidence for the alternative hypothesis).

To investigate this pattern further, we grouped the confidence ratings such that trials were labelled according to the subjects’ confidence in the option they finally chose, regardless of whether this involved a CoM from their initial choice or not (termed “final certainty” to reflect its independence from the initial-choice). Here, we observed a Final-Certainty*Post-choice Evidence interaction ($\beta = -0.29$, $t(21073) = -4.22$, $p < .0001$), driven by the fact that when Evidence Continued (Fig. 4b) higher final certainty was associated with a *smaller* post-choice CPP amplitude ($\beta = -0.02$, $t(12014) = -2.27$, $p = .0231$), while no such relationship was observed when evidence was extinguished (Fig. 4a; $\beta = 0.01$, $t(9059) = 0.75$, $p = .45$). BIC slightly favoured final certainty as a predictor of CPP amplitude over confidence-in-initial-choice ($\Delta BIC = -3$).

Taken together, these results suggest that the post-choice CPP scales with the participant’s final certainty, irrespective of whether or not the ultimately chosen alternative differs from the initially chosen one. While the CPP scaled positively with the participants’ initial choice confidence, post-choice CPP amplitude was *inversely* related to final certainty in Experiment 2, with “maybe” trials now having the highest amplitude. One potential explanation for this pattern, inspired by a previously reported mathematical model (Pleskac & Busemeyer, 2010), is that the duration of post-choice evidence accumulation is certainty-dependent, such that participants are more likely to terminate the accumulation process when highly certain in a particular alternative, after which the CPP decays back to baseline as seen when evidence is Extinguished

(Fig. 4a). Assuming the certainty-dependent stopping rule is stable across trials, it would result in lower amplitudes in our measurement window (which was toward the end of the delay period) for trials with high final certainty because of the earlier peak and decay of the CPP. An example of this kind of effect was reported in Twomey et al. (2016) where the CPP was found to peak and decay earlier on easier trials when participants were required to withhold reporting dot motion direction decisions until stimulus offset (see also Rogge et al., 2022; Tagliabue et al., 2019). This pattern resulted in larger average CPP amplitudes toward the end of the delay period for trials with weaker physical evidence. In other words, when CPP amplitude measurements are taken within a fixed time window within a delay period, they will scale with the proportion of trials on which evidence accumulation was still ongoing during that window (see Discussion for illustration). In the following sections, we report a number of analyses designed to test the hypothesis that participants were reaching commitment earlier on trials with high final certainty.

3.4. Earlier post-choice CPP peak latency and confidence reports on trials with high final certainty

We ran a post-hoc exploratory analysis on the peak latency of the post-choice CPP. The surface plots in Figure 5a highlight substantial cross-trial variability in post-choice CPP peak latency in the continued evidence condition. Post-choice CPP peak latencies were significantly later when Post-decision Evidence Continued (Fig. 5b, $\beta = 0.28$, $t(21245) = 4.15$, $p < .0001$), but earlier when final certainty was higher ($\beta = -0.21$, $t(21245) = -2.71$, $p = .0067$), with a significant interaction of the two ($\beta = -0.24$, $t(21245) = -3.50$, $p = .0005$). This was due to final certainty only having a significant association with peak latency when Evidence Continued ($p = .0036$), but not when it was Extinguished ($p = 0.40$).

If commitment is being reached earlier on trials with high final certainty, then responses for the final report should be made more quickly. Analysis of final-response RTs confirmed this, with significantly faster final RTs when final certainty was higher (Fig. 5c; $\beta = -0.22$, $t(21245) = -32.44$, $p < .0001$). Final RTs were also significantly faster when evidence was Extinguished ($\beta = 0.30$, $t(21245) = 50.07$, $p < .0001$) and there was a significant final certainty by Evidence interaction ($\beta = -0.01$, $t(21245) = -2.17$, $p = .0298$). Examining this interaction, there were significant effects of final certainty both when Evidence Continued and was Extinguished ($p < .0001$), although the effect

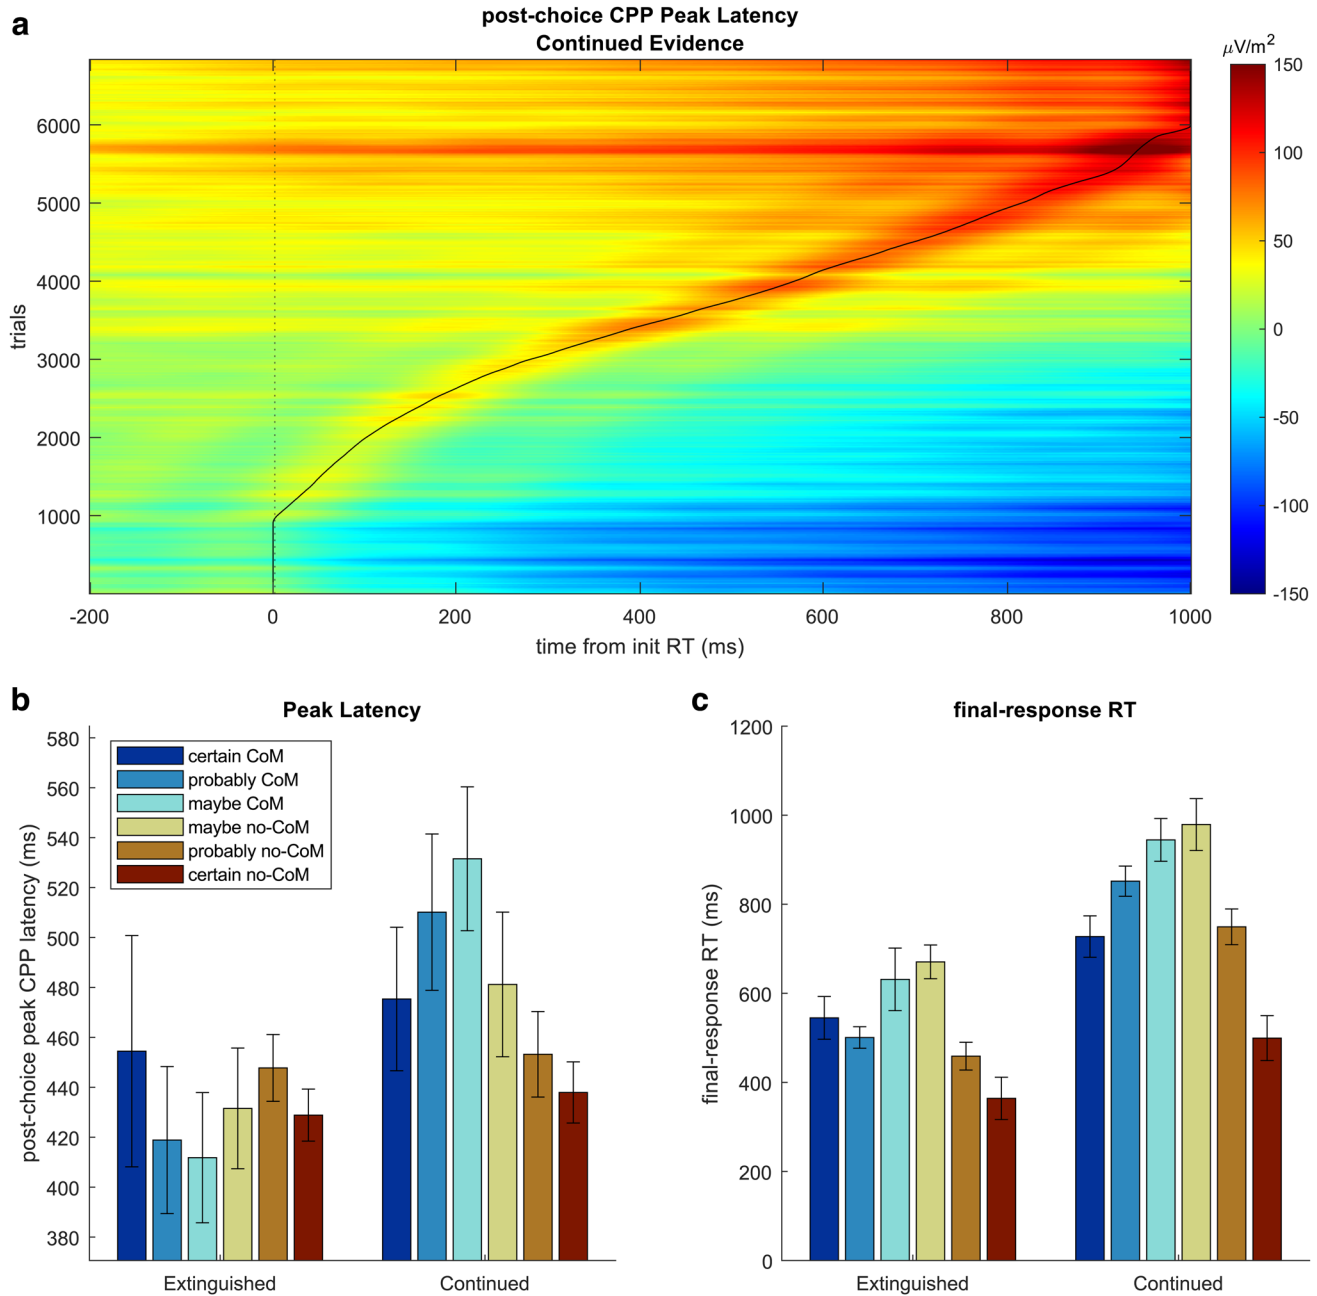


Fig. 5. Post-choice CPP peak latency and final-response RT are quicker for high final-certainty responses. (a) Single-trial post-choice CPP surface plot, sorted by peak latency (curved black line) for the Continued Evidence trials (smoothed over 100 trials). The post-choice CPP decreased after the peak, rather than remaining elevated. (b) The mean post-choice CPP peak latency was earlier on trials rated higher final certainty, but only when Evidence Continued—giving the inverted U-shaped pattern when considering confidence-in-initial-choice. (c) Final-response RTs showed a similar pattern when Evidence Continued, with faster final RTs for higher final-certainty reports. Please note that final-response RTs were cued-responses after the 1000 ms delay, so did not occur around the same time as the peak latencies.

was stronger when Evidence Continued ($\beta = -0.23$ vs -0.20). As Figure 5c shows, when Evidence Continued, the “certain CoM” and “certain no-CoM” had the fastest final-response RTs, and these were almost as fast as trials where the evidence was Extinguished (only 182 ms

and 135 ms slower on “Certain CoM” and “Certain No-CoM” trials respectively, vs >290 ms all the other trial types). Taken together, these results fit with the idea that deliberation stopped earlier when highly certain in that choice.

3.5. Faster motor lateralisation on trials with high final certainty

We reasoned that, if a certainty-dependent stopping rule were being implemented, then preparation of the final response should emerge earlier in the delay period on trials with higher final certainty. We first examined levels of motor preparation in the same measurement window used to analyse the post-choice CPP.

There was a significant effect of final certainty on mean μ/β lateralisation (700:1000 ms window; $\beta = -0.07$, $t(21018) = -10.16$, $p < .0001$), which significantly interacted with Post-choice Evidence ($\beta = -0.02$, $t(21018) = -3.43$, $p = .0006$). This interaction was due to a stronger effect of final certainty when Evidence Continued (Fig. 6f; $\beta = -0.09$, $p < .0001$) than when Extinguished (Fig. 6e; $\beta = -0.04$, $p < .0001$).

Motor lateralisation also varied significantly as a function of confidence-in-initial-choice (mean within 700:1000 ms; $\beta = -0.04$, $t(21018) = -4.90$, $p < .0001$), although without a significant interaction of Post-choice Evidence ($p > .4$). Figure 6b highlights a similar but inverted pattern to that observed for the CPP, with lateralisation increasing with final certainty in both CoM and no-CoM trials. Comparing the LMM including confidence-in-initial-choice to that including final certainty, we found that the latter again gave a better fit to the neural data ($\Delta\text{BIC} = -90$). Thus, greater delay period motor preparation was observed on trials with higher reported final certainty. Although this is the opposite pattern to that seen for the CPP, it is equally consistent with the certainty-dependent stopping-rule account because whereas the CPP would be expected to drop to baseline once the accumulation has halted, lateralised motor preparation would be expected to remain at its extreme near-threshold level since the response has yet to be executed, as seen in some delayed-response tasks (Rogge et al., 2022; Twomey et al., 2016).

To further test the hypothesis that participants were committing to and preparing their responses earlier on high final-certainty trials, we measured mean β lateralisation slopes in 200 ms windows after the initial choice and ran separate analyses on CoM and no-CoM trials (due to the large inversion of lateralisation that occurs on CoM trials). In both CoM and no-CoM trials, trials with higher final-certainty ratings had significantly stronger slopes than lower final-certainty trials, and this effect was apparent from 0-800 ms after initial choice in CoM trials and 200-600 ms in no-CoM trials ($p < .05$, see blue & red bars in Fig. 6b, respectively). This suggests an earlier build-up of motor lateralisation and accords with the final-RT results

reported above, suggesting that motor preparation occurred earlier in the delay period on high final-certainty trials, which would fit with a certainty-dependent stopping.

3.6. Sensory evidence signals increase with confidence and track changes-of-mind

We also examined the post-choice dynamics of the differential SSVEP which indexes the encoding of the sensory evidence, i.e., the relative contrast of the two grating stimuli. The differential SSVEP response following the initial choice was greater on trials with higher confidence-in-initial-choice ($\beta = 0.08$, $t(12102) = 7.62$, $p < .0001$), which interacted with initial accuracy ($\beta = 0.04$, $t(12102) = 5.85$, $p < .0001$), as the differential SSVEP was higher following correct responses subsequently rated as “certain no-CoM” ($\beta = 0.11$, $t(9607) = 8.77$, $p < .0001$), but lower following errors subsequently rated as “certain no-CoM” (Fig. 7a-b; $\beta = -0.03$, $t(2495) = -2.15$, $p = 0.0313$). Unlike the CPP and μ/β signals above, confidence had a monotonic effect; as the evidence always favoured the correct response, stronger evidence encoding was associated with higher confidence in correct initial choices, and lower confidence in incorrect initial choices, i.e., greater confidence in the correct option. Final certainty was significantly associated with higher differential SSVEPs ($\beta = 0.06$, $t(12102) = 7.08$, $p < .0001$), but this had no significant interaction with initial accuracy ($p = .42$). BIC favoured confidence-in-initial-choice over final certainty as a predictor ($\Delta\text{BIC} = -18$), in contrast to the post-choice CPP and motor preparation signals reported above.

Whereas the differential 15 Hz SSVEP traced the sensory-level representation of the relative contrast of the two gratings (i.e., the evidence), the 30 Hz SSVEP offered a metric of the overall visual response to both stimuli combined. There was no significant relationship between the post-choice 30 Hz SSVEP and final certainty ($p > .05$), although there was a negative main effect of confidence-in-initial-choice ($\beta = -0.03$, $t(11979) = -2.09$, $p = .0369$), but this did not interact with initial accuracy ($p > .05$) like the differential SSVEP did (Supplementary Fig. 4). BIC favoured the confidence-in-initial-choice predictor only very slightly ($\Delta\text{BIC} = -1$). In order to interpret this lack of initial-accuracy*confidence-in-initial-choice effect in the 30 Hz signal, we directly compared it with the differential 15 Hz SSVEP. A significant three-way interaction was observed ($p = .0013$), indicating that the two SSVEPs differed significantly in their sensitivity to the combination of confidence and accuracy.

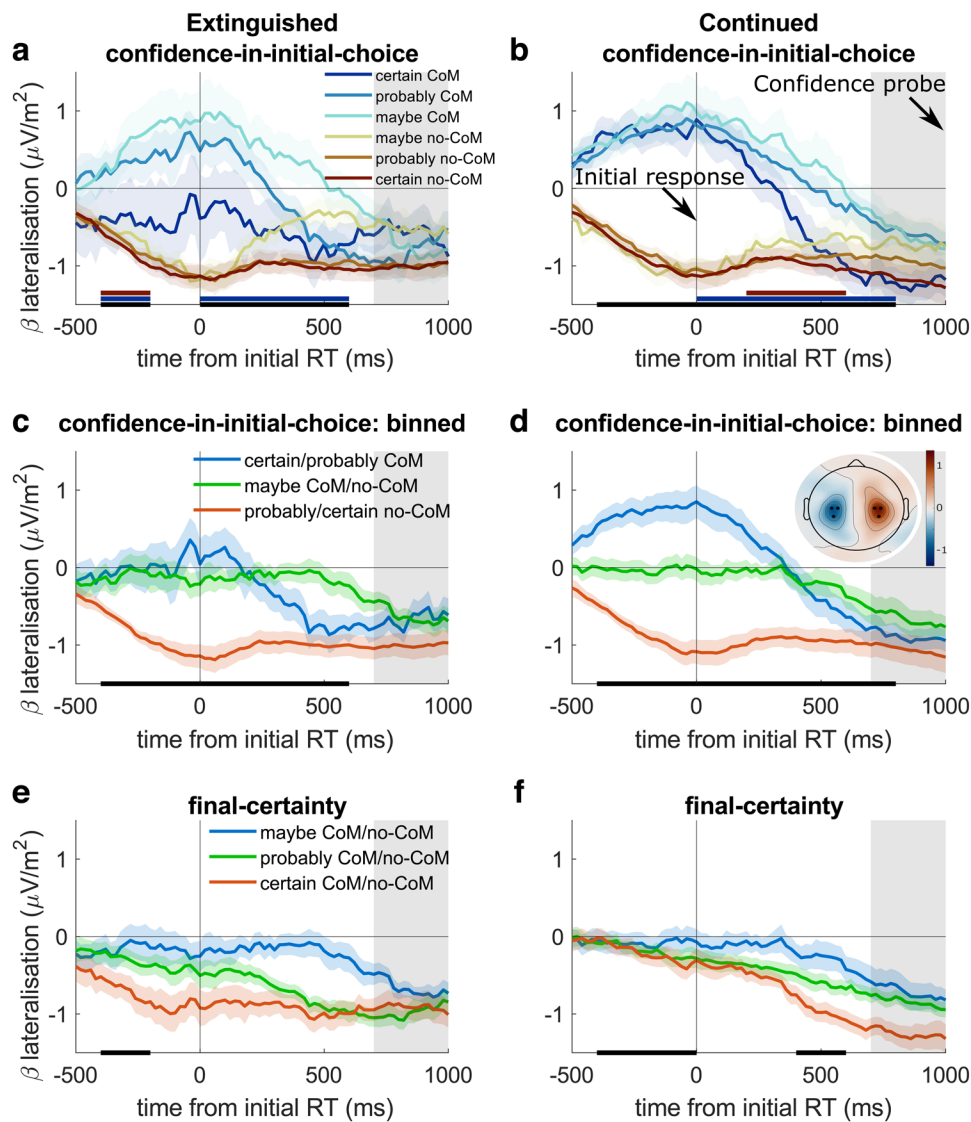


Fig. 6. Motor lateralisation signals invert with changes-of-mind, and scale with final certainty. Mu/beta power lateralisation index relative to the hand used for the final response (therefore it is inverted at the time of the initial response on trials where participants changed their minds). Topography in panel (d) shows the grand-mean motor lateralisation index within the grey window, averaged across all conditions and trials (red = positive, black dots = selected channels). The solid black bar shows the 200 ms time-bins with significant effects ($p < .05$) of that factor on the **slopes** of the lateralisation, while blue/red solid bars show the same for effects of final certainty on **slopes** within change-of-mind/no-change in the top row. Mean mu/beta lateralisation (within the grey window, 700:1000 ms after initial RT) is stronger when people have greater confidence-in-initial-choice, whether evidence is extinguished (a) or continued (b), and the slopes are steeper for high-confidence responses from around soon after initial RT when evidence Continues (blue/red solid bars). When confidence-in-initial-choice is binned into adjacent pairs (c & d), it is clearer that the “maybe” trials have less motor preparation, especially when evidence Continues (d). Using this coding, the green line involves trials that do and do-not change responses, so the average is close to zero at the initial response time; however when evidence is Extinguished (c) this is also seen in the blue line, which is due to the “certain CoM” trials having slightly negative motor preparation in this time-period, perhaps reflecting motor execution errors. (e-f) Mu/beta lateralisation is stronger for trials with greater final certainty, and weakest for “maybe” trials, with the difference in slopes appearing around 400 ms after initial response.




Fig. 7. Sensory post-choice evidence signals show greater stimulus engagement with higher confidence and predict changes-of-mind. Differential SSVEP strength (arbitrary units) measures the difference in sensory signal strength for the target and non-target (i.e., the strength of sensory evidence encoding) to the post-choice continued evidence. The phase of the stimuli reset at initial RT, so the signal builds from zero then. The topography in panel (b) shows the grand-mean differential SSVEP within the grey window, averaged across all Continued Evidence trials (red = positive, black dots show electrodes used for selection). Black solid bars show which 100 ms time-bins have a significant effect of the factor for that initial-accuracy ($p < .05$). The mean differential SSVEP (within the grey time-window, 700:1000 ms) shows an interaction with confidence-in-initial-choice and initial accuracy as: (a) following an error response, the differential SSVEP is higher on trials where participants change their minds (blue lines), and lowest for trials in which they stick with their initial choice with high confidence (dark red); (b) following a correct response, the opposite pattern is seen, with a linear increase in differential SSVEP for trials with higher confidence-in-initial-choice. (c-d) When confidence-in-initial-choice is binned (to increase trial numbers for the rarer responses), the same pattern is seen. (e-f) When split by final certainty, there was no significant interaction between initial error and final certainty; there was a significant positive relationship following initial correct responses, and a non-significant positive relationship following initial errors.

4. DISCUSSION

While the role of neural evidence accumulation processes in forming perceptual decisions is now well-established, the extent to which these same processes continue to operate after an initial choice to inform subsequent confidence judgements has remained an open question. The

results of Experiment 1 demonstrate that when confidence is reported simultaneously with a perceptual choice, the pre-choice amplitude of a motor-independent EEG signature of evidence accumulation (CPP) increases monotonically with confidence (Fig. 2e), and this effect was replicated in a different sample and task in Experiment 2 in which confidence reports were delayed by

1000 ms with no intervening evidence (Fig. 3e). Experiment 2 also allowed us to examine the post-choice dynamics of the CPP, and whether they were influenced by the continued availability of physical evidence.

Multiple elements of our data suggest that commitment to the final confidence report tended to be reached very soon after the initial choice in the Extinguished condition: compared to trials with continued evidence, participants had faster final RTs (Fig. 5c), and no positive build-up or confidence-dependent modulation in the trial averaged post-choice CPP (Fig. 4). Nevertheless, participants' final reports were significantly more accurate than their initial choices on these trials (Fig. 3a) suggesting that some further accumulation did take place, potentially drawing on evidence still in the perceptual pipeline at the time of commitment (Resulaj et al., 2009).

When the physical evidence remained on screen, the CPP did continue to exhibit confidence-dependence, but now its amplitude at the end of the delay period was better explained by final certainty than confidence-in-initial-choice (Fig. 4b & f), as there were little differences between CoM and no-CoM trials. It is perhaps not surprising that the upcoming (final) decision better explains the CPP than the earlier initial decision, but the suggestion that this is invariant to whether the final choice involves a CoM or not is novel, as previous studies find large effects of CoM, although these studies did not present evidence during the interval between choice and confidence report (Boldt & Yeung, 2015; Feuerriegel et al., 2022). Additionally, the CPP-confidence relationship has inverted, with "certain" ratings corresponding to greater CPP amplitude before the initial choice but lower amplitudes before the final choice.

Previous observations led us to suspect that this trend arose from certainty-dependent variations in the duration of the post-choice evidence accumulation. Specifically, in previous studies in which participants withheld perceptual reports until the provision of a delayed response cue, the CPP has been shown to peak and decay long before response cue onset on easy trials (Rogge et al., 2022; Twomey et al., 2016). Under a standard decision model, this would occur because the boundary is reached, after which the decision variable stops accumulating and presumably decays back to baseline. As stronger evidence drives earlier decision termination, trial-averaged CPP amplitude measured late in the post-stimulus window will scale inversely with evidence strength (Rogge et al., 2022; Tagliabue et al., 2019; Twomey et al., 2016). We reasoned, therefore, that in the present data, post-

choice accumulation may also be terminated as soon as a criterion level of certainty in one of the choice alternatives has been reached. A post-hoc exploratory analysis found that the post-choice CPP did peak significantly earlier on trials with higher final certainty, and the peak was followed by a return to baseline consistent with the process being terminated (Fig. 5a-b).

Aside from the CPP data, several other results are consistent with earlier termination of the decision process on trials with higher final certainty. First, final-response RTs were significantly faster on trials with high final certainty (Fig. 5c), consistent with actions being selected long in advance of the response cue. Second, effector-selective motor lateralisation signals exhibited earlier and greater lateralisation in favour of the ultimately chosen effector on trials with higher final certainty (Fig. 6b & f). The question remains why participants opt to stop accumulating while evidence remains available. Plausible explanations include a reduction in effort or energetic costs (Sharot et al., 2023) and/or reducing time-on-task by allowing final responses to be prepared before the cue appears. These observations have important implications for mathematical models of choice confidence which have tended to disagree with the stopping rules applied to post-choice accumulation. In particular, our results appear to be at odds with models that assume a purely time-based rule and others assuming that accumulation continues until confidence reports are probed (Pleskac & Bussemeyer, 2010; Yu et al., 2015). The use of free-response confidence ratings during EEG recordings in future work will allow the relationship between post-choice accumulation dynamics, confidence, and elapsed time for the more finely probed.

While we have referred to this pattern as "certainty-dependent" stopping, our data do not allow us to pinpoint the precise nature of the stopping rule that is being applied and several alternative proposals exist in the literature. According to the "optional stopping model" proposed by Pleskac and Bussemeyer (2010), the decision variable is translated into confidence (relative to the initial-choice) based on its proximity to "confidence thresholds" that are tied to the distinct confidence levels that can be reported, and evidence accumulation is immediately terminated when the decision variable reaches the extreme confidence thresholds (corresponding to "certain CoM" and "certain no-CoM" in our task). This model can account for a range of behavioural effects (Moran et al., 2015), and a version of this model was recently shown to produce superior behavioural fits to a time-based stopping rule (Desender, Donner, et al., 2021).

Figure 8 illustrates that the observed CPP trends can be reproduced by such a decision process with fixed certainty-thresholds, and there are two key features of the CPP that account for this. First, the CPP is positive-going irrespective of the alternative that is favoured by the cumulative evidence (O'Connell & Kelly, 2021), thus the signal is expected to rise on average even when the evidence favours the alternative that was not initially endorsed. Second, the CPP decays back to baseline following a choice commitment which ends the trial and accumulation (Fig. 2e). Assuming the same occurs with a latent certainty-threshold crossing that terminates accumulation, this would predict lower average amplitudes later in the delay period on trials which reached the threshold earlier (e.g., Twomey et al., 2016). Thus, the more trials which terminate early in the delay period, the more the signal decay will dominate the average signal, while the signal will build positively on average in trial types with later and/or fewer threshold crossings. Importantly, this account does not require that evidence accumulation signals evolve in the same way regardless of changes-of-mind (e.g., “certain CoM” vs “certain no-CoM”); these trials can differ in the time taken to reach their respective confidence thresholds and still have decreasing average accumulation signals as long as the decaying signal from trials which have stopped

accumulating outweigh the increasing signals from trials still accumulating.

Detailed modelling, not feasible given the relatively low trial numbers and fixed evidence strength in the current dataset, will be required to clarify exactly how such a stopping rule is implemented in our task. For example, in the model of Pleskac and Busemeyer (2010), while accumulation immediately terminates upon reaching the extreme confidence bounds, a lower probability of stopping is also assigned when passing intermediate thresholds that correspond to the intermediate confidence levels that can be reported. Alternatively, rather than terminating accumulation probabilistically upon reaching an intermediate confidence threshold, the confidence level required to terminate post-choice accumulation may decrease as a function of time (Moran et al., 2015), similar to a collapsing bound or dynamic urgency effect (Hanks et al., 2014; Yau et al., 2021). Finally, it is possible that certainty-dependent stopping could occur at the time of the initial choice, so that trials where participants were already highly confident in their initial response would have no additional accumulation at all, a possibility that would require measuring confidence at the initial choice time to investigate.

A certainty-dependent stopping rule can also account for key observations in the literature on the Error Positiv-

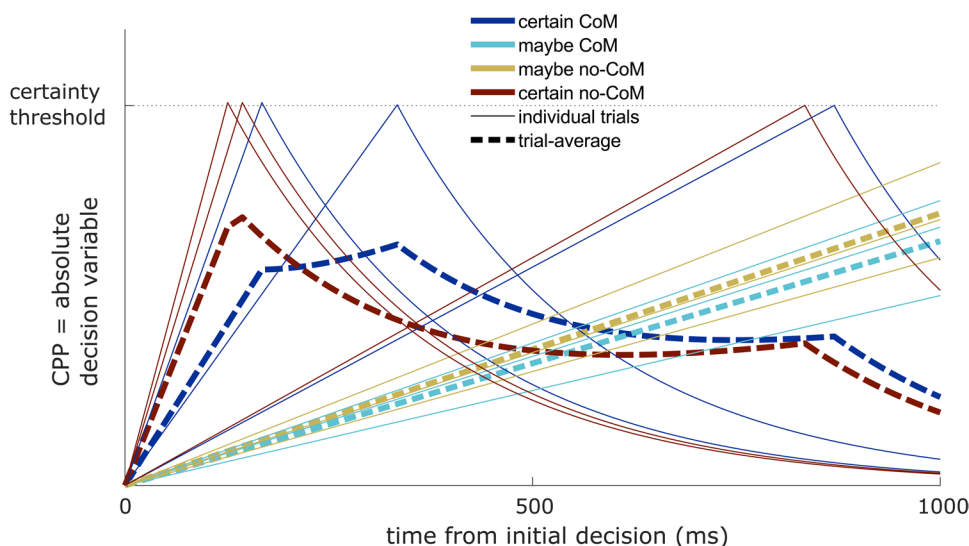


Fig. 8. Illustration of how halting accumulation upon early certainty-boundary crossing can give a decreasing signal. Trials (thin solid lines) which cross the certainty-threshold (and are thus classed as “certain” for that option; dark red+blue) exhibit a decay afterwards, giving a decreasing trial-average signal (thick dashed lines). However, trials which do not reach this certainty-threshold by the end of the delay period (and thus are classed as “maybe” in this simplified illustration; cyan+yellow) do not decay, thus giving a shallow positive-going trial-average signal. The CPP is represented here by the absolute decision variable, meaning that CoM and no-CoM trials both have a similar signal, although differences in the slopes (and therefore crossing-times) can still exist.

ity (Pe), a signal that shares many functional similarities with the CPP, including its polarity and topography (Murphy et al., 2015). The Pe is typically examined in the context of tasks in which errors primarily occur due to a strong prepotency being established for one of the choice alternatives, either through differences in choice outcome probability, as in Go/No-Go tasks (e.g., Endrass et al., 2012; Niessen et al., 2017; Shalgi et al., 2009), or due to response-biases such as the antisaccade task (Falkenstein, 1990; Nieuwenhuis et al., 2001) or flanker task (Overhoff et al., 2021; Selimbeyoglu et al., 2012). It has been well established using these kinds of tasks that the Pe is elicited by erroneous choices that are explicitly detected by participants and is greatly diminished or absent following correct choices and undetected errors (Endrass et al., 2012; Nieuwenhuis et al., 2001; O'Connell et al., 2007; Steinhauser & Yeung, 2010). These observations have fuelled the theory that the Pe reflects the operation of a process that is designed specifically to detect action errors (Desender, Ridderinkhof, et al., 2021). That is, rather than being referenced to the original choice alternatives, post-choice evidence is accumulated in a new reference frame representing the probability that the preceding choice was incorrect. However, our results highlight another plausible functional account in which the Pe can be understood as a continuation of the same process indexed by the CPP, i.e., an evidence accumulation process mapped to the choice alternatives that is subject to a certainty-dependent stopping rule. Taking the example of the Go/No-Go task, Go trials have much greater probability than No-Go trials and therefore the decision bounds would be expected to be much higher for the latter than the former. Consequently, for a participant to change their mind with high certainty following an error of commission on a No-Go trial, post-choice evidence would need to accumulate to the higher No-Go bound, leading to a larger signal following errors compared to correct Go responses.

However, our data do not exclude the possibility that the post-choice CPP reflects a metacognitive operation that is distinct from the evidence accumulation process that it traces prior to the initial choice. While we have presented the stopping-rule as a continuation of the initial decision process, it can be equally applied to a metacognitive accumulation process that evaluates the initial choice accuracy (e.g., Desender, Donner, et al., 2021; Fleming & Daw, 2017), yet stops earlier when highly certain that the response was either correct or erroneous. These two possibilities give very similar predictions for the neural and behavioural data, especially in the current paradigm

and dataset, so we are not able to distinguish between them. Future experiments measuring confidence at initial and final choices, and manipulating post-choice evidence, will aim to directly compare these two theories.

Two previous studies have examined the post-choice CPP on tasks in which a delay was imposed between initial choice and confidence reporting, but with no intervening evidence (Boldt & Yeung, 2015; Feuerriegel et al., 2022). In the case of Boldt and Yeung (2015), the post-choice parietal ERP was found to scale negatively and monotonically with confidence-in-initial-choice, contrary to the present results. However, in that study, signals were baseline-corrected to an interval immediately prior to the initial choice, which potentially confounds pre- and post-choice variations in cumulative evidence. Interestingly, their confidence effects remain when using a pre-stimulus baseline, although they appear to manifest only in the trials rated as errors (corresponding to CoM trials in the present study). Comparison with our results is difficult, as waveforms with pre-stimulus baselines are not presented to allow examination of effects before or after their amplitude measurement window (250-350 ms). When Feuerriegel et al. (2022) applied a pre-stimulus baseline correction, these post-commitment confidence effects on amplitude were substantially diminished and observed only following erroneous choices. We found a similar pattern here, with the application of a pre-response baseline period producing significantly greater post-choice CPP amplitudes for low confidence-in-initial-choice ratings, which were not observed with a pre-stimulus baseline (see Supplementary Fig. 3).

We also found that the encoding of the sensory evidence (differential contrast) in early visual areas, as indexed by the differential SSVEP, was highly sensitive to confidence (Fig. 7). Unlike the CPP and motor lateralisation signals, the differential SSVEP scaled most strongly with the participant's confidence relative to their initial choice, with stronger responses on initially correct trials that were rated as likely to be correct and weaker responses on trials rated as likely to be incorrect (and vice-versa following an initial-error). Previous studies have found sensory ERPs are stronger on higher confidence trials (Squires et al., 1973; Zakrzewski et al., 2019), but to our knowledge this is the first finding of the SSVEP showing a similar scaling. The differential SSVEP signals remained at a relatively constant level throughout the post-choice delay period even on trials with high final certainty, suggesting that early termination of the decision process did not result in a disengagement of

visual processing resources from the stimulus. It was not possible in the present study to establish whether these signal modulations by confidence manifest before or after commitment to the final confidence report. Again, the use of speeded confidence reports would be useful in this regard.

While the focus of this study was on probing post-choice evidence accumulation, Experiment 1 found that pre-choice CPP amplitudes increased with simultaneously rated confidence, consistent with previous analyses of stimulus-locked CPP amplitudes (Davidson et al., 2022; Gherman & Philiastides, 2015, 2018; Herding et al., 2019; Tagliabue et al., 2019). This shows that neural signatures of confidence are available at the time of the initial choice whether confidence is rated simultaneously or 1 second later without post-choice evidence continuing. Possible mechanisms to explain why a neural signature of cumulative evidence measured prior to the choice report would scale with confidence include that the initial choice is subject to a collapsing boundary (e.g., Kelly et al., 2021; Steinemann et al., 2018) and/or starting-point variability. Future investigations can explore whether including these in post-choice confidence models can account for the patterns we report here, such as confidence-effects being present before initial-responses occur, or even before response-cues appear, as in Experiment 1.

In our task, delayed confidence reports did not correlate with pre-choice CPP amplitudes if post-choice evidence was presented. The weakening of this relationship can be explained by the fact that the neural and behavioural data indicate that a greater amount of post-choice evidence accumulation occurred in this condition and, thus, the amplitude of the CPP at the time of the initial choice would be less predictive of the participant's final confidence level. A previous study in which stimuli were extinguished during a delay period between initial choice and confidence reports found no link between pre-choice CPP and post-choice confidence (Feuerriegel et al., 2022). One possible explanation for this discrepancy with our own results is that Feuerriegel et al. (2022) presented participants with brief static stimuli and shorter response deadlines which may have promoted a greater degree of post-choice evidence accumulation from iconic memory. As was the case in our continued evidence condition, this further accumulation may have degraded the relationship between pre-choice CPP amplitude and final confidence. The stimulus variables that determine whether evidence accumulation will continue in the absence of a physical stimulus is an interesting question for future research.

We found that “certain” responses were substantially more common than “maybe,” and CoM responses were relatively rare, especially when the initial response was correct. We used a quadratic scoring rule in combination with trial-by-trial accuracy feedback, which incentivises maximising initial-accuracy and the accuracy of the confidence responses, and should push people to accurately rate their confidences (Staël von Holstein, 1970). While monetary incentives have been shown to improve meta-cognitive performance (Lebreton et al., 2018), they have also been reported to induce overconfidence (Lebreton et al., 2018), and quadratic scoring rules can cause confidence responses to cluster at the edges of confidence scales (certain correct and certain error; Hollard et al., 2016). A different way of incentivising accurate confidence reports may have given a more even spread of responses, and thus lower variances for those rarer responses, facilitating more detailed investigation of the more extended accumulation dynamics on those trials.

Our analyses used the CSD-transformed data, as this minimises volume conductance between separate components (Kayser & Tenke, 2006), and reduces the influence of frontocentral preparation signals on CPP amplitudes (Kelly & O'Connell, 2013). A previous study found the CSD transformation was necessary in order to distinguish confidence effects on frontocentral components and the CPP (Feuerriegel et al., 2022). Here too, we found that the CSD transform influenced the extent to which CPP-confidence effects were observed. A look at the non-CSD data for Experiment 1 (Supplementary Fig. 6a) shows that the scaling of the pre-choice CPP with confidence was relatively unchanged. However, in Experiment 2, the relationship between pre-choice CPP and confidence no longer differed between Extinguished and Continued condition, with both showing confidence-effects. In addition, the removal of the CSD transform eliminated both the effects of post-choice evidence on post-choice CPP amplitude and its relationship with confidence-in-initial-choice. This suggests the CSD transform is needed to uncover post-choice evidence and confidence effects on the CPP.

In conclusion, our data suggest that evidence accumulation persists following initial choice commitment when physical evidence remains available and is terminated in a certainty-dependent manner. Whether post-choice evidence accumulation is mapped to the choice alternatives or is reframed to evaluate the accuracy of the initial choice remains an open question. In general, our study highlights some important methodological considerations when measuring post-commitment CPP activity; trial-averaged amplitudes can vary as a function of the duration of the

accumulation process as well as the bounds associated with each of the choice alternatives. Using paradigms in which participants make speeded, rather than delayed, post-choice confidence reports may facilitate the acquisition of CPP measurements that are more reflective of the state of the decision process at the time the participant commits to their final confidence report.

DATA AND CODE AVAILABILITY

Task scripts and anonymous data are available at <https://osf.io/4dqkz/> and MATLAB code for analysing this data is available at <https://doi.org/10.5281/zenodo.7550910>.

AUTHOR CONTRIBUTIONS

John P. Grogan: Data curation, Formal analysis, Investigation, Visualization, Writing—original draft, and Writing—review & editing. Wouter Rys: Conceptualization, Data curation, Investigation, Methodology, Project administration, and Writing—review & editing. Simon P. Kelly: Conceptualization, Writing—review & editing. Redmond G. O’Connell: Conceptualization, Funding acquisition, Supervision, and Writing—review & editing.

DECLARATION OF COMPLETING INTEREST

The authors declare no competing financial interests.

ACKNOWLEDGEMENTS

J.R.G. and R.G.O. were supported by Horizon 2020 European Research Council Consolidator Grant IndDecision 865474. S.P.K. was supported by Science Foundation Ireland (15/CDA/3591) and The Wellcome Trust (219572/Z/19/Z). The funders had no involvement in the study.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available with the online version here: https://doi.org/10.1162/imag_a_00005.

REFERENCES

- Bang, D., & Fleming, S. M. (2018). Distinct encoding of decision confidence in human medial prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 115(23), 6082–6087. <https://doi.org/10.1073/pnas.1800795115>
- Boldt, A., & Yeung, N. (2015). Shared neural markers of decision confidence and error detection. *Journal of Neuroscience*, 35(8), 3478–3484. <https://doi.org/10.1523/JNEUROSCI.0797-14.2015>
- Davidson, M. J., Macdonald, J. S. P., & Yeung, N. (2021). Alpha power and stimulus-evoked activity dissociate metacognitive reports of attention, visibility and confidence in a visual detection task. *BioRxiv*. <https://doi.org/10.1101/2021.11.23.469669>
- Davidson, M. J., Macdonald, J. S. P., & Yeung, N. (2022). Alpha oscillations and stimulus-evoked activity dissociate metacognitive reports of attention, visibility, and confidence in a rapid visual detection task. *Journal of Vision*, 22(10), 20. <https://doi.org/10.1167/jov.22.10.20>
- Delorme, A., & Makeig, S. (2004). EEGLAB: An open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134, 9–21. <https://doi.org/10.1016/j.jneumeth.2003.10.009>
- Desender, K., Boldt, A., Verguts, T., & Donner, T. H. (2019). Confidence predicts speed-accuracy tradeoff for subsequent decisions. *eLife*, 8, 1–25. <https://doi.org/10.7554/eLife.43499>
- Desender, K., Donner, T. H., & Verguts, T. (2021). Dynamic expressions of confidence within an evidence accumulation framework. *Cognition*, 207, 104522. <https://doi.org/10.1016/j.cognition.2020.104522>
- Desender, K., Murphy, P., Boldt, A., Verguts, T., & Yeung, N. (2019). A postdecisional neural marker of confidence predicts information-seeking in decision-making. *Journal of Neuroscience*, 39(17), 3309–3319. <https://doi.org/10.1523/JNEUROSCI.2620-18.2019>
- Desender, K., Ridderinkhof, K. R., & Murphy, P. R. (2021). Understanding neural signals of post-decisional performance monitoring: An integrative review. *eLife*, 10(e67556), 1–21. <https://doi.org/10.7554/eLife.67556>
- Endrass, T., Klawohn, J., Preuss, J., & Kathmann, N. (2012). Temporospatial dissociation of Pe subcomponents for perceived and unperceived errors. *Frontiers in Human Neuroscience*, 6(2012), 1–10. <https://doi.org/10.3389/fnhum.2012.00178>
- Falkenstein, M., Hohnsbein, J., Hoorman, J., & Blanke, L. (1990). “Effects of errors in choice reaction tasks on the ERP under focused and divided attention,” in *Psychophysiological Brain Research*, eds C. H. M. Brunia, A. W. K. Gaillard, and A. Kok. (Tilberg, Netherlands: Tilberg University Press).
- Feuerriegel, D., Murphy, M., Konski, A., Mepani, V., Sun, J., Hester, R., & Bode, S. (2022). Electrophysiological correlates of confidence differ across correct and erroneous perceptual decisions. *NeuroImage*, 259, 119447. <https://doi.org/10.1016/j.neuroimage.2022.119447>
- Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision performance: A general Bayesian framework for metacognitive computation. *Psychological Review*, 124(1), 91–114. <https://doi.org/10.1037/rev0000045>
- Fleming, S. M., van der Putten, E. J., & Daw, N. D. (2018). Neural mediators of changes of mind about perceptual decisions. *Nature Neuroscience*, 21(4), 617–624. <https://doi.org/10.1038/s41593-018-0104-6>
- Gherman, S., & Philiastides, M. G. (2015). Neural representations of confidence emerge from the process of decision formation during perceptual choices.

- NeuroImage*, 106, 134–143. <https://doi.org/10.1016/j.neuroimage.2014.11.036>
- Gherman, S., & Philiastides, M. G. (2018). Human VMPFC encodes early signatures of confidence in perceptual decisions. *eLife*, 7, 1–28. <https://doi.org/10.7554/eLife.38293>
- Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *eLife*, 3, 1–17. <https://doi.org/10.7554/eLife.02260>
- Herding, J., Ludwig, S., von Lutz, A., Spitzer, B., & Blankenburg, F. (2019). Centro-parietal EEG potentials index subjective evidence and confidence during perceptual decision making. *NeuroImage*, 201, 116011. <https://doi.org/10.1016/j.neuroimage.2019.116011>
- Hollard, G., Massoni, S., & Vergnaud, J. C. (2016). In search of good probability assessors: An experimental comparison of elicitation rules for confidence judgments. *Theory and Decision*, 80(3), 363–387. <https://doi.org/10.1007/s11238-015-9509-9>
- Kayser, J., & Tenke, C. E. (2006). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: I. Evaluation with auditory oddball tasks. *Clinical Neurophysiology*, 117(2), 348–368. <https://doi.org/10.1016/j.clinph.2005.08.034>
- Kelly, S. P., Corbett, E. A., & O’Connell, R. G. (2021). Neurocomputational mechanisms of prior-informed perceptual decision-making in humans. *Nature Human Behaviour*, 5(4), 467–481. <https://doi.org/10.1038/s41562-020-00967-9>
- Kelly, S. P., & O’Connell, R. G. (2013). Internal and external influences on the rate of sensory evidence accumulation in the human brain. *Journal of Neuroscience*, 33(50), 19434–19441. <https://doi.org/10.1523/JNEUROSCI.3355-13.2013>
- Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice certainty is informed by both evidence and decision time. *Neuron*, 84(6), 1329–1342. <https://doi.org/10.1016/j.neuron.2014.12.015>
- Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science*, 324(5928), 759–764. <https://doi.org/10.1126/science.1169405>
- Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., Broussard, C., & Cornelissen, F. (2007). What’s new in Psychtoolbox-3? A free cross-platform toolkit for psychophysics with Matlab & GNU/Octave. *Perception*, 36(14), 1. <http://www.psychtoolbox.org>
- Lebreton, M., Langdon, S., Sliker, M. J., Nooitgedacht, J. S., Goudriaan, A. E., Denys, D., Luigjes, J., & van Holst, R. J. (2018). Two sides of the same coin: Monetary incentives concurrently improve and bias confidence judgements. *Science Advances*, 4(eaaq0668), 1–13. <https://doi.org/10.1101/099382>
- Li, H., & Ma, W. J. (2020). Confidence reports in decision-making with multiple alternatives violate the Bayesian confidence hypothesis. *Nature Communications*, 11(2004), 1–11. <https://doi.org/10.1038/s41467-020-15581-6>
- Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing type I error and power in linear mixed models. *Journal of Memory and Language*, 94, 305–315. <https://doi.org/10.1016/j.jml.2017.01.001>
- Moran, R., Teodorescu, A. R., & Usher, M. (2015). Post choice information integration as a causal determinant of confidence: Novel data and a computational account. *Cognitive Psychology*, 78, 99–147. <https://doi.org/10.1016/j.cogpsych.2015.01.002>
- Murphy, P. R., Robertson, I. H., Harty, S., & O’Connell, R. G. (2015). Neural evidence accumulation persists after choice to inform metacognitive judgments. *eLife*, 4, e11946. <https://doi.org/10.7554/eLife.11946>
- Niessen, E., Fink, G. R., Hoffmann, H. E. M., Weiss, P. H., & Stahl, J. (2017). Error detection across the adult lifespan: Electrophysiological evidence for age-related deficits. *NeuroImage*, 152, 517–529. <https://doi.org/10.1016/j.neuroimage.2017.03.015>
- Nieuwenhuis, S., Ridderinkhof, K. R., Blom, J., Band, G. P. H., & Kok, A. (2001). Error-related brain potentials are differentially related to awareness of response errors: Evidence from an antisaccade task. *Psychophysiology*, 38(5), 752–760. <https://doi.org/10.1111/1469-8986.3850752>
- O’Connell, R. G., Dockree, P. M., Bellgrove, M. A., Kelly, S. P., Hester, R., Garavan, H., Robertson, I. H., & Foxe, J. J. (2007). The role of cingulate cortex in the detection of errors with and without awareness: A high-density electrical mapping study. *European Journal of Neuroscience*, 25(8), 2571–2579. <https://doi.org/10.1111/j.1460-9568.2007.05477.x>
- O’Connell, R. G., & Kelly, S. P. (2021). Neurophysiology of human perceptual decision-making. *Annual Review of Neuroscience*, 44(1), 495–516. <https://doi.org/10.1146/annurev-neuro-092019-100200>
- Overhoff, H., Ko, Y. H., Feuerriegel, D., Fink, G. R., Stahl, J., Weiss, P. H., Bode, S., & Niessen, E. (2021). Neural correlates of metacognition across the adult lifespan. *Neurobiology of Aging*, 108, 34–46. <https://doi.org/10.1016/j.neurobiolaging.2021.08.001>
- Pleskac, T. J., & Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117(3), 864–901. <https://doi.org/10.1037/a0019737>
- Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: Distinct probabilistic quantities for different goals. *Nature Neuroscience*, 19(3), 366–374. <https://doi.org/10.1038/nn.4240>
- Rausch, M., Zehetleitner, M., Steinhäuser, M., & Maier, M. E. (2020). Cognitive modelling reveals distinct electrophysiological markers of decision confidence and error monitoring. *NeuroImage*, 218, 116963. <https://doi.org/10.1016/j.neuroimage.2020.116963>
- Resulaj, A., Kiani, R., Wolpert, D. M., & Shadlen, M. N. (2009). Changes of mind in decision-making. *Nature*, 461(7261), 263–266. <https://doi.org/10.1038/nature08275>
- Rogge, J., Jocham, G., & Ullsperger, M. (2022). Motor cortical signals reflecting decision making and action preparation. *NeuroImage*, 263, 119667. <https://doi.org/10.1016/j.neuroimage.2022.119667>
- Selimbeyoglu, A., Keskin-Ergen, Y., & Demiralp, T. (2012). What if you are not sure? Electroencephalographic correlates of subjective confidence level about a decision. *Clinical Neurophysiology*, 123(6), 1158–1167. <https://doi.org/10.1016/j.clinph.2011.10.037>
- Shalgi, S., Barkan, I., & Deouell, L. Y. (2009). On the positive side of error processing: Error-awareness positivity revisited. *European Journal of Neuroscience*, 29(7), 1522–1532. <https://doi.org/10.1111/j.1460-9568.2009.06690.x>
- Sharot, T., Rollwage, M., Sunstein, C. R., & Fleming, S. M. (2023). Why and when beliefs change. *Perspectives on*

- Psychological Science*, 18(1), 142–151. <https://doi.org/10.1177/17456916221082967>
- Squires, K. C., Hillyard, S. A., & Lindsay, P. H. (1973). Vertex potentials evoked during auditory signal detection: Relation to decision criteria. *Perception & Psychophysics*, 14(2), 265–272. <https://doi.org/10.3758/BF03212388>
- Staël von Holstein, C. S. (1970). Measurement of subjective probability. *Acta Psychologica*, 34, 146–159. [https://doi.org/10.1016/0001-6918\(70\)90013-2](https://doi.org/10.1016/0001-6918(70)90013-2)
- Steinemann, N. A., O’Connell, R. G., & Kelly, S. P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nature Communications*, 9(3627), 1–13. <https://doi.org/10.1038/s41467-018-06117-0>
- Steinhauser, M., & Yeung, N. (2010). Decision processes in human performance monitoring. *Journal of Neuroscience*, 30(46), 15643–15653. <https://doi.org/10.1523/JNEUROSCI.1899-10.2010>
- Tagliabue, C. F., Veniero, D., Benwell, C. S. Y., Cecere, R., Savazzi, S., & Thut, G. (2019). The EEG signature of sensory evidence accumulation during decision formation closely tracks subjective perceptual experience. *Scientific Reports*, 9(1), 4949. <https://doi.org/10.1038/s41598-019-41024-4>
- Twomey, D. M., Kelly, S. P., & O’Connell, R. G. (2016). Abstract and effector-selective decision signals exhibit qualitatively distinct dynamics before delayed perceptual reports. *Journal of Neuroscience*, 36(28), 7346–7352. <https://doi.org/10.1523/JNEUROSCI.4162-15.2016>
- van den Berg, R., Zylberberg, A., Kiani, R., Shadlen, M. N., & Wolpert, D. M. (2016). Confidence is the bridge between multi-stage decisions. *Current Biology*, 26(23), 3157–3168. <https://doi.org/10.1016/j.cub.2016.10.021>
- Yau, Y., Hinault, T., Taylor, M., Cisek, P., Fellows, L. K., & Dagher, A. (2021). Evidence and urgency related EEG signals during dynamic decision-making in humans. *Journal of Neuroscience*, 41(26), 5711–5722. <https://doi.org/10.1523/JNEUROSCI.2551-20.2021>
- Yu, S., Pleskac, T. J., & Zeigenfuse, M. D. (2015). Dynamics of postdecisional processing of confidence. *Journal of Experimental Psychology General*, 144(2), 489–510. <https://doi.org/10.1037/a0039293>
- Zakrzewski, A. C., Wisniewski, M. G., Iyer, N., & Simpson, B. D. (2019). Confidence tracks sensory- and decision-related ERP dynamics during auditory detection. *Brain and Cognition*, 129, 49–58. <https://doi.org/10.1016/j.bandc.2018.10.007>