Centroparietal signals track multiple rounds of evidence accumulation in an intermittent evidence task

Authors: Elisabeth Parés-Pujolràs[1], Anna Geuzebroek[1], Redmond G. O'Connell[2], Simon Kelly[1]

Affiliations:

[1]School of Electrical and Electronic Engineering, University College Dublin, Dublin, Ireland

[2]Trinity College Institute of Neuroscience and School of Psychology, Trinity College Dublin, Dublin, Ireland

Abstract

Making decisions often requires the integration of multiple pieces of information. An extensive body of research has investigated the neural mechanisms underpinning evidence accumulation in tasks where information is continuously present, but less is known about how this neural architecture supports evidence accumulation if the flow of information is paused for a variable amount of time. In two electroencephalography (EEG) experiments (N = 22, N = 21), participants made perceptual decisions about two pulses of motion evidence separated by gaps of varying duration. Our goal was to investigate 1) whether and how people integrate the two pulses across various delay durations, and 2) how a prominent EEG marker of evidence accumulation, the centroparietal positivity (CPP), reflects decisions in the context of our intermittent-pulses task. Our results are consistent with a bounded evidence accumulation process where the second pulse further informs the decision only if the bound is not reached by the first pulse. We show that whereas motor preparation signals reflected in motor beta lateralisation (MBL) maintain the cumulative evidence level reached from pulse 1 through the gap, the CPP does not sustain but rather exhibits transient deflections following each pulse, encoding momentary belief updates.

1. Introduction

Perceptual decision-making is widely thought to involve a process of evidence integration over time (Hyafil et al., 2023) and an extensive body of research has successfully linked these integration computations to multiple neural decision signals in various species (Gold & Shadlen, 2007; Hanks & Summerfield, 2017; Heekeren et al., 2008; O'Connell & Kelly, 2021; Shadlen & Kiani, 2013). While most previous work has investigated fast perceptual decisions based on a single continuously-presented stimulus, many real-life decisions involve incorporating multiple, temporally-disparate pieces of evidence. Accordingly, many lines of work have moved to examine the potentially distinct modes of operation of decision processes in tasks requiring integration of several discrete tokens (Balsdon et al., 2020, 2021; Murphy et al., 2021; Parés-Pujolràs et al., 2021, 2025; Rouault et al., 2022; Sospedra et al., 2025; Waskom & Kiani, 2018). In the present study, we examined the neurocomputational mechanisms supporting decisions in a little-studied but ecologically prevalent scenario, whereby only a short glimpse of relevant evidence is presented and it is uncertain whether and when one will receive an additional glimpse of the evidence source (Azizi & Ebrahimpour, 2023; Kiani et al., 2013; Tohidi-Moghaddam et al., 2019). For example, we often need to keep track of moving objects in traffic that may be intermittently occluded. Combining an intermittent-evidence task, human electroencephalography (EEG) and computational modelling, we investigate how prominent neural decision signals may underpin choice formation in such scenarios.

In environments where sensory statistics do not change over time, the optimal strategy is to perfectly accumulate all available information. Humans and other animals are capable of such perfect integration (Brunton et al., 2013), even in tasks where evidence can last for up to several seconds (Shadlen & Kiani, 2013; Waskom & Kiani, 2018). However, in delayed-response tasks where stimulus duration is experimentally controlled and decision reports withheld until stimulus offset, human and non-human primates often fail to follow this optimal strategy and engage in a bounded evidence accumulation process, making early decisions before evidence stops (Balsdon et al., 2020, 2021; Mazurek et al., 2003), or exhibiting primacy/recency biases (Lange et al., 2021). Such deviations from optimal integration may reduce cognitive effort at the cost of accuracy, and may be influenced by multiple factors, including the quality of information provided by a stimulus (Lange et al., 2021), or its duration (Bronfman et al., 2016). In tasks where decisions are made on the basis of noisy continuous

stimuli like random dot motion, each evidence frame contains very little information and participants need to integrate over time to make an accurate decision (Geuzebroek et al., 2023; Gherman et al., 2024; Steinemann et al., 2018). In contrast, in typical token-based paradigms each discrete stimulus contains low perceptual noise, so participants can easily discriminate the information provided by each individual token. However, given that individual tokens may be misleading, they still need to integrate across multiple tokens over longer timescales to maximise accuracy (Balsdon et al., 2020, 2021; Parés-Pujolràs et al., 2025; Wyart et al., 2012, 2015). The intermittent-evidence or "gap" task (Kiani et al., 2013) represents a middle ground between continuous-evidence and token-based paradigms: it uses noisy evidence, but groups it in two temporally discrete, sequential windows. Thus, participants are required to integrate over time to discriminate the information in each evidence pulse, but also across evidence pulses to maximise their accuracy. Previous studies using intermittent-evidence tasks have found that both evidence pulses do not tend to be utilised fully (Azizi & Ebrahimpour, 2023; Kiani et al., 2013), and although a general information leakage account could be ruled out based on the absence of any decline in accuracy with delay, the neural decision architecture through which this underutilisation came about remained unstudied.

In humans, two prominent EEG markers have emerged as key decision signals: the centroparietal positivity (CPP) and motor beta lateralisation (MBL). In conventional tasks where participants can respond at any time to a continuous stimulus such as random dot kinematograms (RDK), the CPP builds up as evidence accumulates, at a rate that scales with evidence strength, and peaks at the time of choice (Kelly & O'Connell, 2013; O'Connell et al., 2012; Steinemann et al., 2018). The CPP is always of positive polarity, regardless of the choice favoured by the evidence, and is present even in tasks where no motor output is required, suggesting it might encode an intermediary, abstract stage of evidence integration (O'Connell et al., 2012). In parallel, beta band (13-30Hz) power over motor cortex gradually decreases power with building motor preparation (Pfurtscheller & Lopes Da Silva, 1999; Stancák & Pfurtscheller, 1995), also in an evidence-dependent way (Geuzebroek et al., 2023; Steinemann et al., 2018), reaching a threshold level contralateral to the executed action, at the time of that action (Kelly et al., 2021; O'Connell et al., 2012). However, in delayed response tasks where decisions are made early but the response is deferred, MBL is maintained until that response, while the CPP shows a peak-and-fall profile, not sustaining until action execution (Twomey et al., 2016). These findings are consistent with the CPP tracking an intermediary evidence

integration process that is continuously fed to downstream motor preparation where a sustained representation of the DV is maintained (e.g.O'Connell & Kelly, 2021).

More recently, the use of delayed-response expanded judgement tasks where evidence needs to be integrated across multiple discrete tokens has shed new light onto how this neural architecture and its markers can exhibit further functional dissociations. For example, centroparietal responses evoked by discrete evidence tokens have been shown to be more closely related to the momentary evidence in each individual token than to the absolute value of an ongoing DV evolving across tokens (Wyart et al., 2012, 2015). Further, in volatile environments where the generative source of the tokens - and hence the correct answer - can stochastically flip during a trial, the CPP has been found to transiently encode normative belief updates derived from each individual token, rather than the objective stimulus information in the token or a stable representation of the ongoing DV (Parés-Pujolràs et al., 2025). In contrast, MBL behaves in a similar way to conventional continuous evidence tasks, tracking the evolving decision variable across discrete samples over multiple seconds, by stepping up or down according to the evidence in each sample and sustaining in between them (Murphy et al., 2021; Parés-Pujolràs et al., 2025). Each of these studies employed discrete, perceptually low-noise tokens that may have particularly lent themselves to individual appraisal. The "gap" task (Kiani et al., 2013) is qualitatively distinct in that each of the one or two pulses of continuous evidence is difficult to discriminate, but both indicate a consistent sensory category, in this way appearing as temporally unpredictable glimpses of a single evidence source rather than a sequence of stand-alone bits of information. Previous work has shown that in tasks where continuous evidence is temporarily stopped, the CPP exhibits transient evoked positivities after each sequential pulse, and does not perfectly sustain throughout the delay (Azizi & Ebrahimpour, 2023). However, how these pulse-evoked activities relate to an ongoing decision and which process drives the decay during the delay period remains poorly understood.

In this study, we ran two experiments involving modified versions of a motion direction discrimination task where two pulses of random dot motion were presented separated by variable temporal gaps (Kiani et al., 2013). Our primary goal was to characterise the dynamics of the CPP in the context of this intermittent evidence task, and functionally relate them to behavioural indices of integration across tokens, including any biases in their influence on choice. Here, we examined the influence of the first and second pulses on behavioural and neural signatures of decision-making and then developed a model designed to account for both.

Our results suggest that the CPP traces separate, discrete accumulation processes in response to each glimpse, computing belief updates that are fed to downstream motor areas which continuously encode a DV, as indexed by MBL. Further, we show that the weighting of evidence pulses on choice was modulated by attention, as indexed by alpha power.

2. Results

2.1. Behavioural results

In this study, participants engaged in a motion direction discrimination task where, on most trials, two pulses of partially-coherent dot motion (indicated by yellow dot colour) were separated by variable temporal 'gaps' in which motion was random, and thus uninformative (indicated by blue dot colour; *Fig. 1A*). Though coherence could differ, the direction was always consistent between the two pulses, so that they would be construed as two glimpses of the same evidence source. Our two experiments (Exp. 1: N = 21; Exp. 2: N = 22) differed only in the delays separating the two pulses and therefore also in the temporal expectations about the timing of the second pulse. In both experiments, coherence levels were titrated for each participant using a staircase procedure aiming for an average 70% accuracy following a single evidence pulse. High (Exp. 1: M = 44.73%, SD = 12.64%; Exp 2: M = 33.04%, SD = 13.41%) and low (Exp. 1: M = 26.75%, SD = 7.74%; Exp 2: M =19.82%, SD = 8.04%) coherences were then set to be, respectively, 25% higher and 25% lower than the staircase-estimated coherence (see *Methods*), and small adjustments were made between blocks if performance substantially changed (*Figure S1*).

As expected, participants' accuracy scaled with motion coherence and number of pulses. In one-pulse trials, participants were more accurate in the high coherence compared to the low coherence trials (*Eq. 1*, Exp 1: $\beta_1 = 0.61 \pm 0.11$, $p < 0.001$; Exp 2: $\beta_1 = 0.84 \pm 0.15$, $p < 0.001$, *Figure 1A*). There was a further benefit in accuracy from the second pulse when present, indicated by sensitivity to coherence in both the first (*Eq.2*, Exp 1: $\beta_1 = 0.46 \pm 0.06$, $p < 0.001$; Exp 2: $\beta_1 = 0.53 \pm 0.0245$, $p < 0.001$) and second (*Eq.2*, Exp 1: $\beta_2 = 0.25 \pm 0.05$, $p < 0.001$; Exp 2: $\beta_2 = 0.10 \pm 0.024$, $p < 0.001$) pulses. When both pulses in a two-pulse trial were the same coherence, participants performed better compared to one-pulse trials, exhibited both in high and low coherence trials in Experiment 1 (*Eq. 3*, Exp. 1: $\beta_1 = 0.22$, $p < 0.001$; $\beta_3 = -0.008$, $p = 0.89$), but only significantly so in low coherence trials in Experiment 2 (*Eq. 3*, Exp. 2: $\beta_3 = 0.053$, $p = 0.34$; $\beta_3 = 0.21$, $p < 0.001$). However, they did not perform as well as expected

from a perfect accumulator (*Figure 1B*), using performance in one-pulse trials to predict expected accuracy in two-pulse ones (see *Eq*. 5-7; see *Methods*).

The fact that the regression coefficients for the second pulse were smaller than those estimated for the first one indicates that participants' choices were guided more strongly by the first piece of evidence they received. On average, accuracies were lowest for the most difficult trials (both pulses had low motion coherence, LL), intermediate for trials with mixed coherences (LH, HL) and highest for the easiest trials (both pulses had high motion coherence, HH). However, we also observed a significant pulse-order effect: when the two presented pulses had different coherences, participants were more accurate when the 1st pulse was the high-coherence one (HL) compared to trials where the first pulse was low-coherence (LH) in both experiments (*Eq. 4*, Exp 1: $\beta_2 = -0.11 \pm 0.03$, $p < 0.001$; Exp2: $\beta_2 = -0.11 \pm 0.02$, $p < 0.001$; *Figure 1B)*. This is in line with participants' choices being more strongly influenced by the first evidence pulse.

Importantly, there was no main effect of gap duration on accuracies (*Eq. 2*, Exp 1: $\beta_3 = 0.20 \pm 0.06$, $p = 0.81$; Exp 2: $\beta_3 = 0.09 \pm 0.06$, $p = 0.11$; *Figure 1A*), suggesting an absence of leakage of information from the first pulse, in line with previous work (Azizi & Ebrahimpour, 2023; Kiani et al., 2013). However, gap duration did interact with evidence strength in different ways. In Exp 1 participants' accuracy increased with gap duration when the 1st pulse coherence was high, and decreased with gap duration when it was low (*Eq 2*, Exp 1: $\beta_4 = 0.25 \pm 0.05$, $p < 0.001$). In Exp 2, accuracies increased with gap duration when the 2nd pulse coherence was high, and decreased when it was low (*Eq 3*, Exp 2: $\beta_5 = 0.26 \pm 0.04$, $p < 0.001$). Since these interactions were inconsistent across experiments and we did not have prior hypotheses regarding such effects, we leave these to one side and focus on coherence-related effects. In what follows, we combine EEG and computational modelling to test the possibility that the greater sensitivity to the first pulse arises from the participants engaging in a bounded accumulation process where decisions are sometimes terminated early, curtailing the extent to which the second pulse benefits accuracy and driving the coherence order effects.

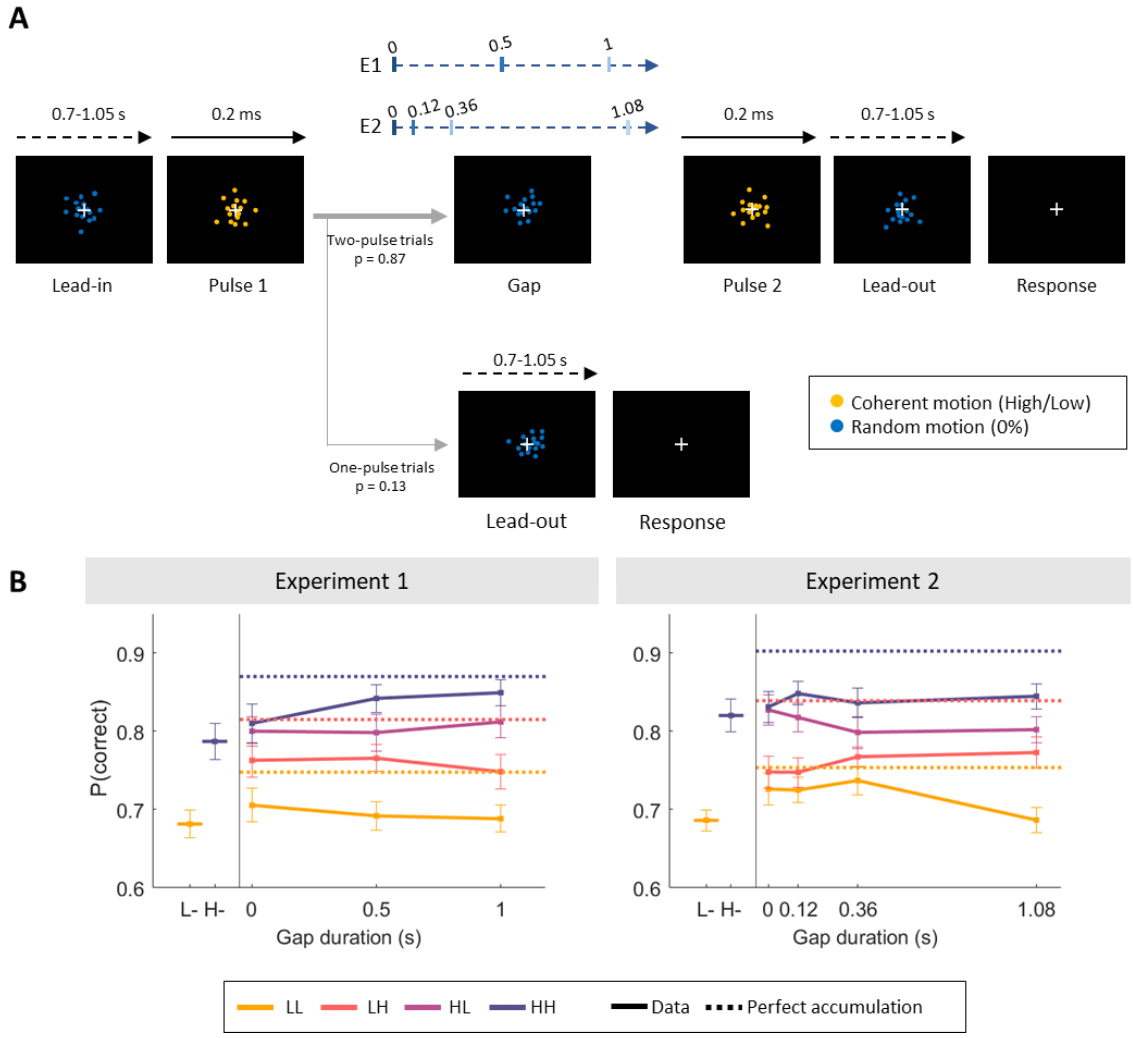

Figure 1. **Task and behavioural results.** ***A***. Experimental structure. Participants viewed a random dot stimulus that coherently moved either to the left or right during yellow evidence pulses, of which there was sometimes one, but usually two, separated by a variable gap in evidence indicated by blue dot colour. **B.** Grand-averaged (mean ± se) accuracy in two-pulse trials, sorted by pulse coherence (low, L/ high, H), order (LL, LH, HL, HH) and gap duration. Data points in the left-hand section indicate single-pulse accuracies. Overlaid horizontal dashed lines indicate the expected accuracy of a perfect accumulator, based on the grand-average performance on one-pulse trials. Note that a perfect accumulator predicts identical accuracies for the HL and LH conditions, and therefore the two lines overlap.

2.2.EEG results

2.2.1. Delay-related activity

We then turned to neural correlates of evidence accumulation during the task, specifically, the MBL and CPP. Across the two experiments and in all delay conditions, MBL diverged from zero shortly after the first pulse (*Fig. 2A*) indicating relatively immediate motor preparation, and MBL was sustained at a coherence-dependent level (see *Fig. S2; S3*) throughout the gap between the two pulses and until response. In contrast, the CPP exhibited a transient build-up following each evidence pulse, shortly falling back to baseline levels prior to the second pulse (*Fig. 2B).*

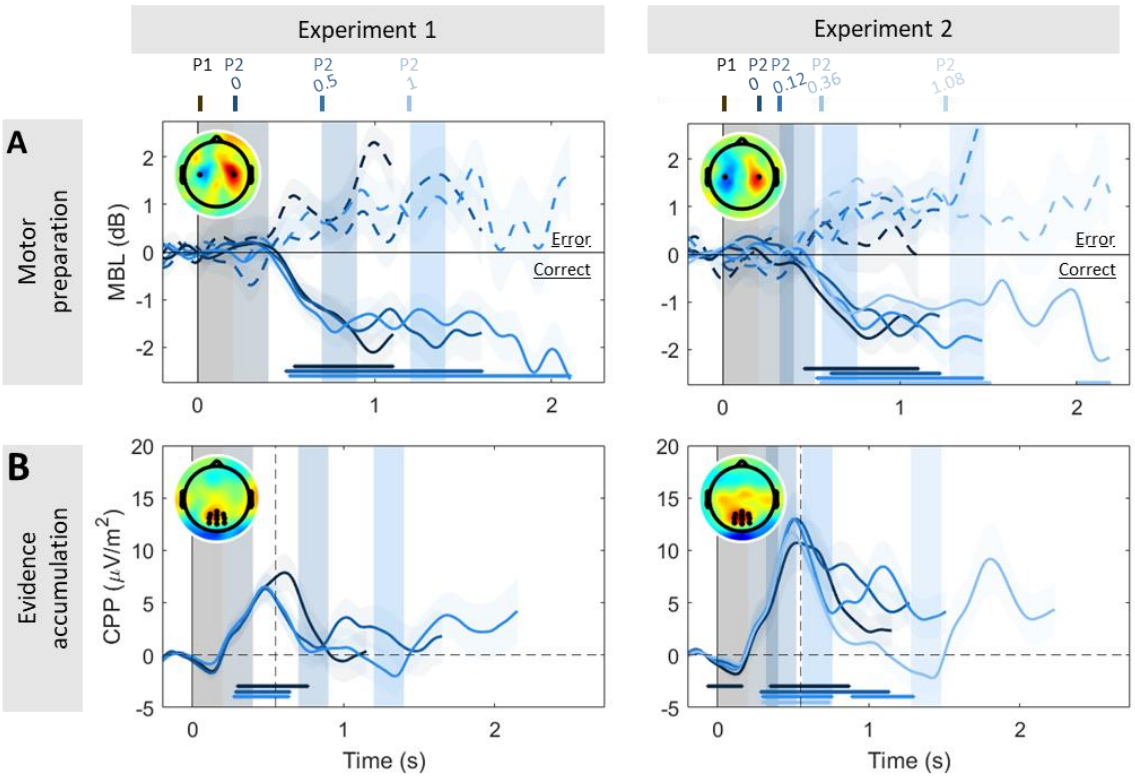

Figure 2. **Neural correlates of evidence accumulation and motor preparation across temporal gaps. A.** Grand-averaged (±SEM) motor beta lateralisation (MBL) [Contra-Ipsilateral hemispheres with respect to correct response hand], sorted by accuracy (*solid line* correct, *dashed line* error). Participants started preparing a motor response immediately after the first pulse, and MBL was sustained throughout the delay period and until the time of response across all delay conditions in both experiments. **B.** Grand-averaged (±SEM) centroparietal positivity (CPP) across all temporal gap conditions. The CPP showed transient evoked responses following each pulse. Note that the evoked amplitudes after P1 (CPP-P1) were systematically higher than those evoked after P2 (CPP-P2). In all panels, shaded areas

indicate the timing of evidence pulses for the various gap durations, which are colour coded (see top axis). Topographies illustrate MBL before response (right minus left responses, ***A***) in correct trials, and CPP 550ms post stimulus onset (*dashed line*, ***B***). In all panels, markers along the bottom indicate times where the signal significantly differed from zero (two-tailed cluster-based permutation test, p<0.005) in only correct trials for MBL for consistency of polarity (***A***), and for all trials pooled in the case of the CPP (***B***).

2.2.2. Centroparietal signals track multiple rounds of evidence accumulation during a single decision

Next, we turned to coherence-dependent effects. The CPP evoked after the first pulse (CPP-P1) scaled with the strength of evidence as expected, reaching higher amplitudes after high coherence pulses compared to low-coherence ones over a cluster of centroparietal electrodes (*Fig. 3A-D;* two-tailed cluster-based permutation test, Exp. 1: $p = 0.012$, Exp. 2: $p = 0.031$). Furthermore, single-trial variability in CPP-P1 amplitude significantly modulated the first pulse's weight on choice: in trials where evoked CPP-P1 amplitudes were particularly high, the coherence of pulse one had a stronger influence on behaviour (*Fig. S11A*).

The coherence of the second pulse modulated peak amplitude in the standard way, with stronger evoked potentials in trials with high P2 coherence, although this effect was only significant in Experiment 2 (*Fig. 3E-H;* Exp. 1: $p = 0.147$, Exp.2: $p = 0.008$; two-tailed cluster-based permutation test*)*. CPP-P2 amplitude was also influenced by the coherence of P1: on trials where the first pulse was low coherence, CPPs evoked by the second pulse reached higher amplitudes (Exp. 1: $p = 0.020$, Exp. 2: $p = 0.005$; two-tailed cluster-based permutation test; see *Fig. 3F,H*). Additionally, CPP-P2 amplitude was generally smaller than CPP-P1 amplitude (Exp. 1:$t(21) = 5.08$, $p = 0<0.001$; Exp. 2: $t(20) = 20.34$, $p < 0.001$; see *Fig. 2B*).

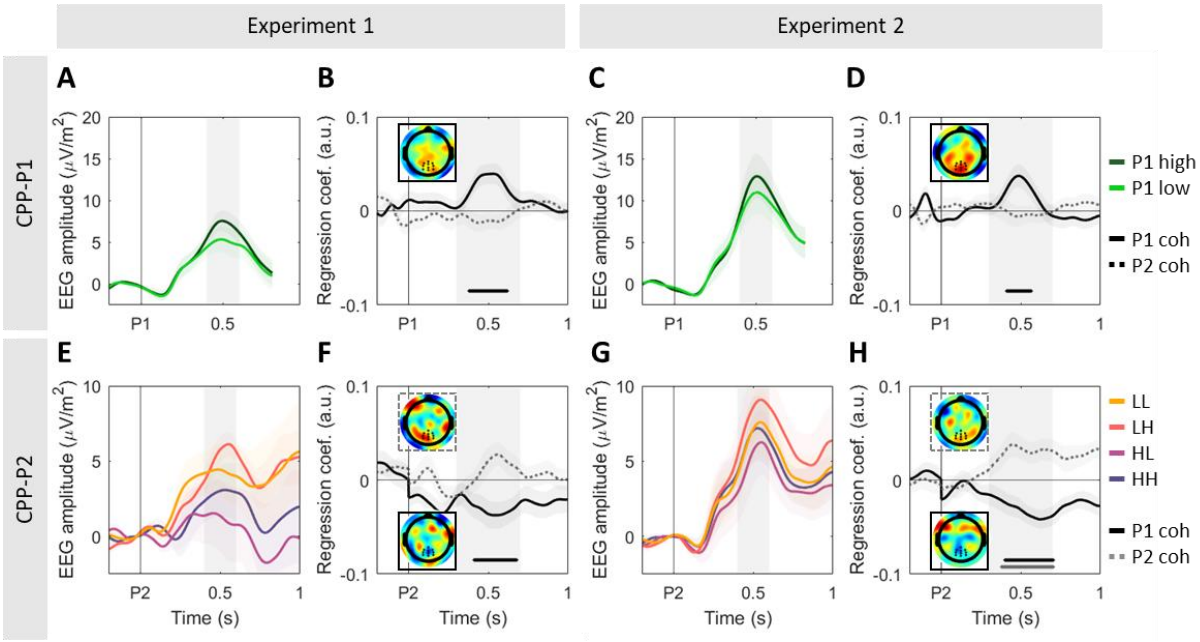

Figure 3. Sensitivity to evidence strength in the two pulses in centroparietal signals. A,C. Grand-averaged (±SEM) centroparietal activity after pulse 1 (CPP-P1), averaged across all gap conditions. Inset topographies illustrate the difference between [P1 high – P1 low] coherence pulses at [0.4-0.6s] after P1. **B,D.** Standardized regression coefficients illustrating the effect of P1 (*solid line*) and P2 (*dashed line*) coherence on evoked CPP-P1 amplitudes. **E,G.** Grand-averaged (±SEM) centroparietal activity after pulse 2 (CPP-P2), averaged over all non-zero gaps in the two experiments. For illustration purposes, we subtracted the average activity of long delay trials with the same P1 coherence from the epoched EEG traces, between [-0.2 to 1.4s] from P1 presentation, and baseline corrected traces [-50 to 0ms] before P2. This isolates the neural activity uniquely associated with the presentation of the second pulse in short gap trials (<1s), where potentials evoked by the first and second pulses overlapped, by removing the activity associated with the first pulse from the P2-evoked traces, including baseline trends (see *Fig. S4*). The qualitative patterns did not vary if this correction was not applied (*see Fig S5*). Note also the difference in the y axis scale compared to panels A,C. **F,H.** Standardized regression coefficients (mean ±SEM) illustrating the effect of P1 (*solid line*) and P2 (*dashed line*) coherence on evoked CPP-P2 amplitudes. Inset topographies illustrate the difference between [P1 high – P1 low], solid line; [P2 high – P2 low], dashed line) [0.4-0.6s] after P2. The negative effect of P1 illustrates higher amplitudes for P1-low trials.*Markers along the bottom of panels B,D,F and H indicate significant clusters where coherence effects differed from zero (two-tailed cluster-based permutation test, p<0.05). Shaded areas indicate the tested time window. See Fig. S3 for full-trial ERP traces, sorted by gap and coherence conditions.

2.2.4. A single bounded decision process can account for key behavioural and neural patterns

We then aimed to identify a computational model that could account for the key behavioural and CPP amplitude effects observed in our data. Similar to previous studies (Azizi & Ebrahimpour, 2023; Kiani et al., 2013) there was no evidence of leak in our data, given that no general decrease in accuracy was observed for longer delay durations. Instead, the primacy effect observed in our behavioural results, with higher accuracies in HL than LH trials, is consistent with a bounded process where participants sometimes terminated evidence accumulation early, failing to use all available evidence. Thus, we fitted a single bounded accumulator model with three free parameters (bound, high- and low-coherence drift rates), where evidence accumulation is terminated upon bound crossing (see *Methods* for full details). If the bound was not crossed by the end of the trial, the sign of the DV at the end of the trial determined the model's predicted choice.

This simple single-DV model could account for the order effects observed in behavioural data, where the 1st pulse had a greater bearing on choice than the 2nd one, while at the same time capturing the generally greater accuracy of two-pulse trials compared to single-pulse ones (*Fig. 4A*). In this model, this partial benefit of the second pulse arises from the fact that the DV is further influenced by it only on the trials where it has not already reached a bound. As expected, a greater proportion of trials terminate early in the P1-high trials compared to P1-low trials (*Fig. 4B*). Note that this simple diffusion model has no way to produce systematic delay-dependent effects, meaning that predicted accuracies are forced to be flat across gaps; again, since the observed delay-dependent effects were unexpected and differed in our two experiments, we did not pursue further modelling approaches to attempt to capture them.

Then, we simulated the neural activity corresponding to the CPP by taking the absolute values of the DV on single trials (in keeping with the always-positive nature of the CPP; see Afacan-Seref et al., 2018; Kelly et al., 2021) and averaging across trials. If a bound was hit, the CPP traces were set to return linearly to 0 within 100 ms, as empirically observed in action-locked traces in self-paced tasks (e.g. Steinemann et al., 2018), and it did not resume upon presentation of a second pulse. Given that the CPP traces in our grand-averaged data (*Fig. 2B*) fell fully back to baseline by the end of the longer gaps, CPP traces were also assumed to linearly return to zero as soon as the dots turned blue at the end of evidence pulses if a bound had not been hit earlier on. Such a pattern assumes that the CPP generators are active during periods where

evidence is presented (yellow dot pulses), and become inactive as soon as evidence stops (blue dot periods). Given the absence of previous empirical data on how signals may fall down without having hit a bound, we assumed the same dynamics in this case as when a bound is hit (linear return to zero within 100 ms). The simulated DV traces from our model captured the sensitivity of both CPP-P1 and CPP-P2 to motion coherence. The difference in CPP-P1 slopes for high vs. low coherence pulses is simply explained by the high vs. low drift-rate associated with these conditions (*Fig. 4E*). Note that the difference in buildup rate appears slight with respect to the differences in motion coherence used; this arises due to taking the absolute value of the DV in each trial, which, when evidence strengths are generally low, means that it builds up almost as much for error trials as correct trials (Steinemann et al., 2018). In the more complex case of CPP-P2, the evoked activity after P2 reached higher amplitudes in trials where P1 was low coherence (*Fig. 4F*). Our model can account for this due to two effects: First, because P1-Low trials are less likely to terminate early (*Fig. 4B*), and therefore fewer trials take a value of zero during P2. Second, because among those trials that do not terminate early, the level reached by the DV by the end of pulse one is on average lower for low- than high-coherence (*Fig 4C*), and therefore more evidence from P2 needs to be integrated before the bound is hit and accumulation terminated. Thus, the higher CPP-P2 amplitudes reached in trials with low P1 coherence reflected both a higher proportion of trials where participants engaged in accumulation of P2, and the fact that they had to cover a longer distance to reach the bound.

The model also predicted differences in CPP amplitude in correct/incorrect trials, which our dataset broadly recapitulated (*Fig. S8*). In particular, it predicted stronger CPP-P1 coherence-dependent effects in correct vs. incorrect trials, and it also predicted an interaction between accuracy and P2-coherence on CPP-P2 amplitude in trials with mixed coherences (HL, LH). See also *Fig. S9* for simulated DV traces in various conditions.

Finally, we assessed the extent to which some of our models' assumptions influenced the simulated CPP profile. For example, it is unclear whether the CPP would deactivate with the same temporal dynamics when a decision process terminates because a bound has been hit, as when accumulation is temporarily paused during a gap. Additional simulations illustrate how different assumptions about the decay rate (i.e. time to go back to 0) and fall-down delay (time before the signal starts going down) would affect the CPP's profile during the gap in our two experiments (*Fig. S6*). These parameters affected the CPP shape during the gap, but the patterns of coherence-dependent CPP-P1 and CPP-P2 amplitudes remained unchanged. We

also simulated an alternative scenario where the signal is not assumed to decay (*Fig. S7*), which would be predicted if the CPP's neural generators remained active throughout the delay, encoding a sustained representation of the DV, so long as the bound has not been crossed. These simulations did not recapitulate the CPP's full drop to baseline in the long gap conditions, and produced significantly smaller CPP-P2 evoked potentials relative to CPP-P1 than observed in the empirical data.

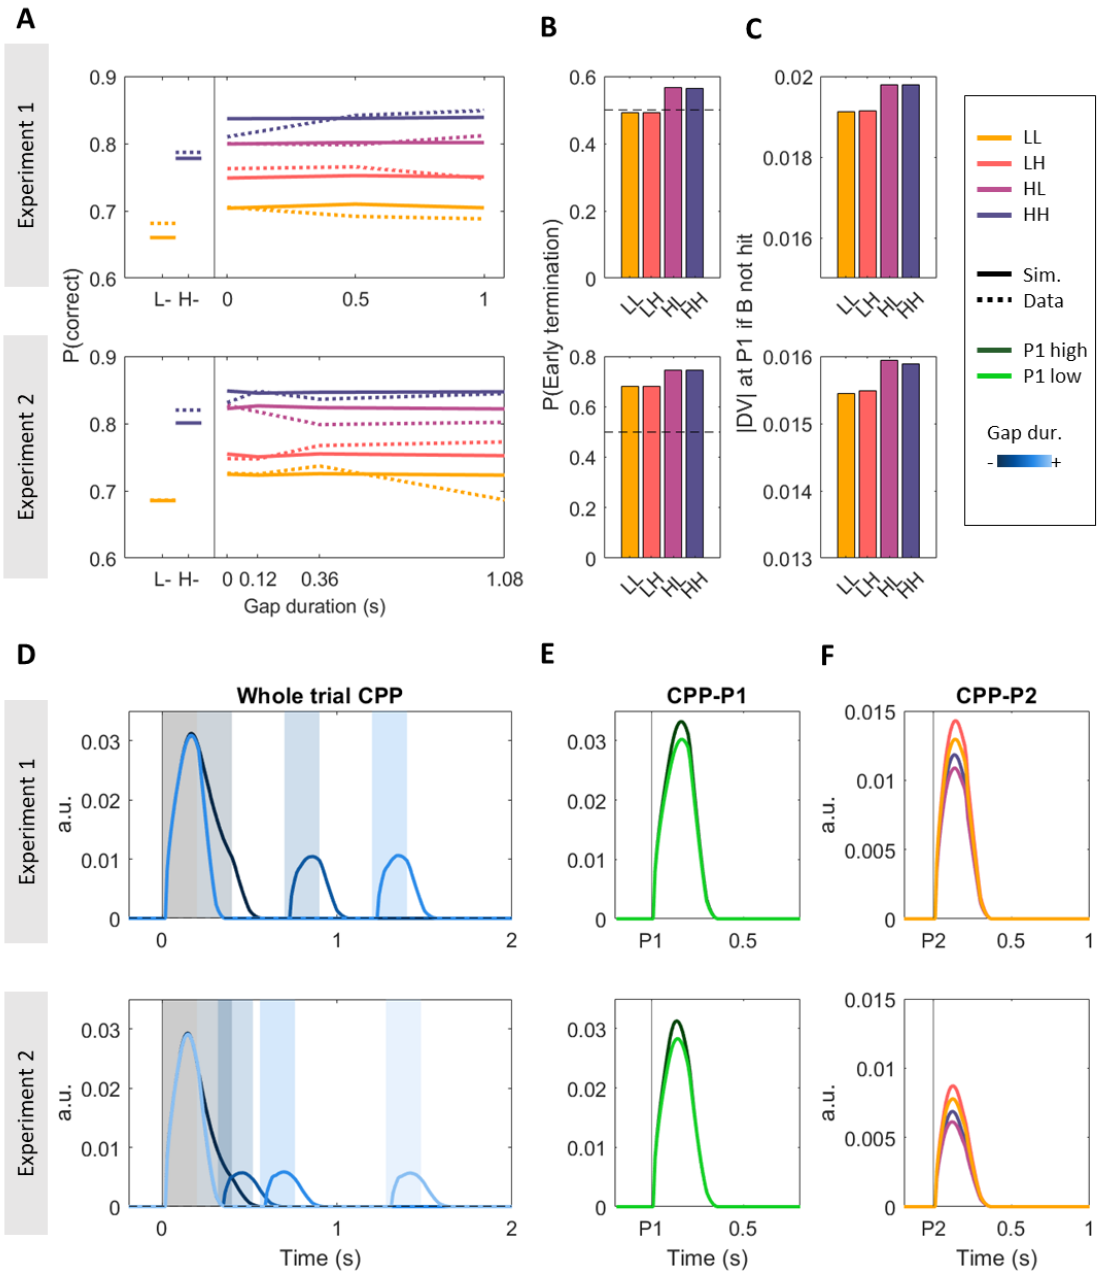

Figure 4. **Simulation results** (N = 50000) of a simple bounded accumulator model with three parameters (B (bound), d1 (high drift rate), d2 (low drift rate)), fit to all data points for each experiment separately. **A.** Grand-averaged simulated (*solid lines*) accuracy in single- (*left*) and

two-pulse (*right*) trials, sorted by pulse coherence and gap duration in experiments 1 (*top*) and 2 (*bottom*). Real data are overlaid (*dashed*) for comparison. Simulated data recapitulated the observed order effects. For simplicity, the model was not endowed with any means to account for delay-related modulations. **B.** Proportion of simulated two-pulse trials that hit a bound during the first pulse, and therefore did not use 2nd pulse evidence to make a choice. **C.** Absolute value of the decision variable (DV) at the end of pulse 1 in the subset of trials where the bound was not hit during that pulse, sorted by condition. **D.** Simulated CPP traces over the whole trial, sorted by gap duration. The simulations recapitulate the bimodal shape of the CPP across the whole trial, with the second pulse reaching lower amplitudes than the 1st one in both experiments (see *Fig. 2B*). **E.** Simulated CPP-P1, sorted by P1 coherence, and averaged across all gap durations. **F.** Simulated CPP-P2, sorted by P1 and P2 coherence, and averaged across all gap durations. These simulations recapitulate the observed coherence effects at both CPP-P1 (see *Fig. 3A,C*) and CPP-P2 (see *Fig. 3E,G*). Note that here we plot the estimated build-up at P2 assuming that activity started at 0 and not considering any potential overlap from earlier on in the trial, for consistency with the empirical traces where CPP-P1 activity is removed from the P2-evoked traces. Note also that, for the sake of simplicity, our CPP simulations do not assume any delay between stimulus presentation and signal build up. Therefore, our simulated CPPs start building up immediately after pulse onset, while empirical ones start with approximately a 0.2s delay.

2.2.3. Occipital alpha power tracks within-trial attentional modulations and temporal expectations

The simple computational model described above captured the accuracies in the various coherence conditions. However, we also observed time-dependent effects in both experiments that our model could not explain. Temporal contingencies are an integral aspect of our task, and participants' expectation about whether or not a second pulse is likely to occur should naturally change over time. Alpha power has been linked to temporal expectation encoding (Rohenkohl & Nobre, 2011), as well as being a well-known index of attentional engagement (Thut et al., 2006) and we next sought to explore whether its within trial dynamics might be related to attentional fluctuations and temporal regularities in our two experiments.

In both experiments, occipital alpha power showed transient event-related desynchronisation that was maintained for as long as evidence was on, likely indicating specific attentive processing during temporal windows containing decision-relevant information, and re-synchronised at evidence offset (*Fig. 5A*). In long gap trials, alpha power was greater before the second pulse than before the first one (i.e. power increased beyond baseline levels), suggesting that participants engaged attention at least partially less before the second pulse (two-tailed cluster permutation test, $p < 0.05$ in both experiments). One possibility is that this attentional disengagement was driven by early termination trials. However, if this was the case,

we would expect this effect to be dependent on pulse 1 coherence, as it is more likely for a DV to reach a bound after a high coherence pulse (*Fig. 4B*). Since alpha power before P2 did not systematically vary with pulse 1 coherence (*Fig.5B*), this effect cannot be entirely accounted for by early terminations. Furthermore, we observed that the magnitude of the resynchronisation at the time of the second pulse in the longest-delay trials was stronger in Experiment 2 than in Experiment 1 ($t(41) = 2.26$, $p = 0.028$; *Fig. 5C)*. Given that the distribution of gap durations differed in our two experiments, temporal expectations should exhibit distinct temporal profiles across the two datasets. In particular, in those trials where no second pulse had been presented after a 0.5s delay, the probability of it happening was 0.67 in Experiment 1, vs 0.6 in Experiment 2. Thus, the differences between our two experiments in alpha power before P2 are consistent with alpha reflecting temporal expectations. Again, given that alpha power before P2 was not dependent on P1 coherence (*Fig. 5B*), the difference between the two experiments cannot simply be explained by different probabilities of early terminations in Exp. 2 relative to Exp. 1 predicted by our model (see *Fig. 4C*).

Since pre-stimulus alpha power has been linked to drift rates (Kelly & O'Connell, 2013; McGovern et al., 2018) and evidence accumulation onset (Van Den Brink et al., 2021), we investigated whether an extended computational model with an additional free drift rate variability parameter applied to pulse 2 only could better capture our behavioural results. We found that while the fit was overall better, especially for Experiment 2, BIC did not warrant its inclusion (see *Fig S10* for full details). Finally, we explored whether single-trial fluctuations in the neural correlates of evidence accumulation and attention were linked to variability in choice behaviour (*Fig. S11*). We found that variability in phasic alpha responses to P2 modulated its weight on choice in Experiment 2 only, where estimated alpha variability was estimated to be highest and provided the greatest improvement on performance in our extended model (*Fig. S10*). Put together, these results suggest that occipital alpha power 1) is related to phasic attentional modulations during evidence pulses, 2) is consistent with temporal expectation encoding, and 3) that its single-trial fluctuations explain some variability in choice behaviour. However, further research is needed to precisely determine how these multiple processes interact.

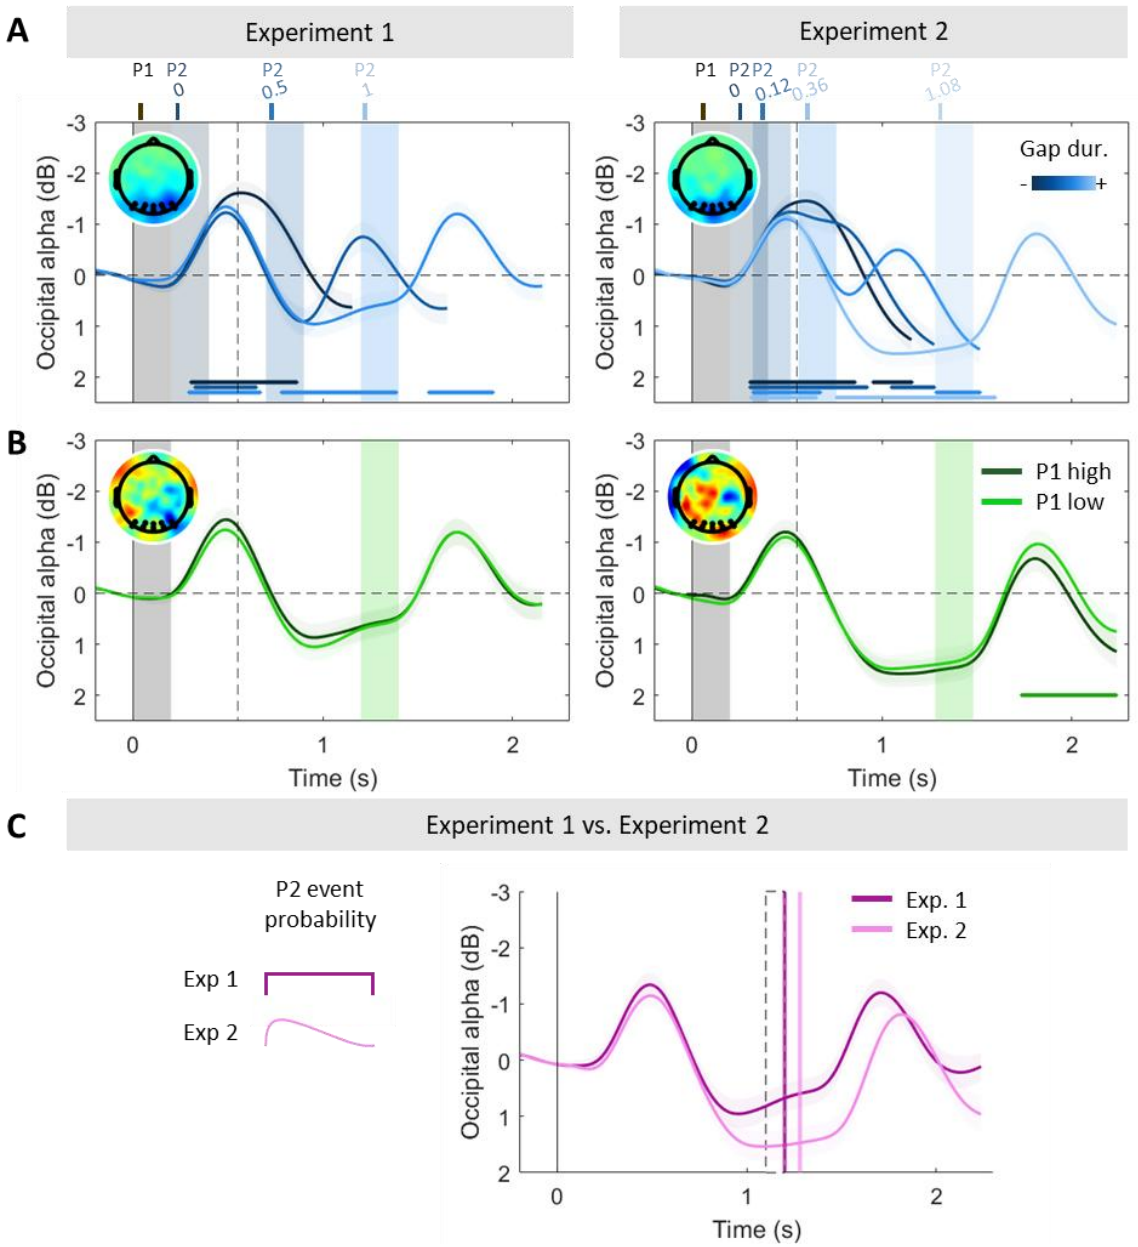

***Figure 5*. Occipital alpha power tracks attention and temporal expectations. A.** Grand-averaged (±SEM) occipital alpha power across all delay conditions, baselined before the onset of the first pulse. Occipital alpha power showed transient desynchronisation starting on pulse onset (which was clearly indicated by a change in the colour of the moving dots), and resynchronising again once the dots changed colour again to indicate non-coherent motion (see *Figure 1*). **B.** Grand-averaged (±SEM) occipital alpha power in the long delay conditions, baselined before the onset of the first pulse and sorted by P1 coherence. Occipital alpha resynchronisation did not systematically vary with pulse one coherence. Topographies illustrate the difference in alpha power between P1-high and P1-low coherence trials [-0.2 to 0s] before P2 onset. **C.** Occipital alpha power in the longest delay condition in the two experiments. In Experiment 2, where the event distribution was skewed towards short delays between the 1st and 2nd pulses, participants' occipital alpha was significantly higher

(indicating lower attention) than in Experiment 1 (* $p<0.05$; two-sample t-test for the difference in alpha power between occipital activity [-0.1 0s] before P2 and the baseline [-0.2 0s] before P1. The dotted box indicates the test period. In all panels, shaded areas indicate the timing of evidence pulses for the relevant gap durations. *Markers along the bottom of panel A indicate significant clusters where alpha power differed from zero (two-tailed cluster-based permutation test, $p<0.05$). In panel B, they indicate significant clusters where alpha power differed in P1-high vs. P1-low coherence trials.

3. Discussion

A wealth of studies has investigated the neural mechanisms underpinning evidence accumulation in tasks where information is continuously available, identifying multiple neural decision signals at various levels in the sensorimotor hierarchy (Gold & Shadlen, 2007; O'Connell & Kelly, 2021; Okazawa & Kiani, 2023). However, little is known about the mechanisms underpinning decision-making in contexts where evidence comes in temporally separate bouts. This is important not only because many real-life decisions require this kind of integration, but also because well-known neural decision signals can exhibit functional dissociations when probed in different contexts (Okazawa & Kiani, 2023; Parés-Pujolràs et al., 2025). In this study, we asked: how do neural decision signals support evidence accumulation if information availability is paused for a variable amount of time? Focussing on two human EEG decision signals, the CPP and MBL, we show that they perform distinct roles during evidence accumulation in this context. Across two experiments, we show that the CPP does not exhibit sustained activity during the gap between evidence pulses, but exhibits transient activations following each pulse that track distinct rounds of evidence accumulation. In contrast, MBL exhibited a sustained build-up throughout the duration of each trial, consistent with the encoding of a DV. Our behavioural and neural data can be explained by a simple bounded accumulation process, where participants sometimes terminate decisions early despite a cost in accuracy. Finally, we find evidence of an interaction between occipital alpha signals and evidence accumulation processes, which is suggestive of a potential attention-gating mechanism.

Our finding that the CPP exhibits transient patterns of activity tracking multiple rounds of evidence accumulation that feed into a single decision illuminates its functional role. Most previous studies using continuous evidence paradigms showed that the CPP steadily grows throughout an evidence accumulation process, in parallel with motor preparation, and is consistent with its encoding the absolute value of a DV (Kelly & O'Connell, 2013; O'Connell

et al., 2012; Steinemann et al., 2018). In contrast, in our task the CPP exhibits two distinct build-up processes that align with the two pulses presented in each trial. This suggests that the CPP builds up while accumulation is ongoing, and disengages when the accumulation process is paused or a bound is hit. This was further supported by simulations of the CPP profile under various signal decay assumptions, which showed that empirical signals were better recapitulated by a process that assumes signal decay when evidence is paused (*Fig S6*). Furthermore, we observed that the amplitude of each pulse-evoked CPP reflected its weight on choice. At the grand average level, the initial pulse systematically evoked higher amplitudes than the second one, consistent with its greater influence on choice behaviour (*Fig. 2B*). Further, single-trial regressions revealed that fluctuations in CPP-P1 amplitude for a given coherence level were linked to behavioural variability, with higher-than-average CPPs associated with a stronger weighting of the corresponding pulse on choice *(Fig S11).* A previous study using a similar paradigm showed results consistent with this interpretation. In Azizi & Ebrahimpour (2023) the first pulse had a smaller weight on choice than the second one and CPP amplitudes were biased in the same direction as the behavioural weighting. This supports the idea that the CPPs following sequential pulses capture multiple rounds of evidence accumulation that are part of a single decision, and that its amplitude reflects each pulse's partial contribution to an evolving DV. In contrast to the transient pattern of activity that characterises the CPP in our task, MBL started immediately after the first pulse, and was sustained at a coherence dependent level throughout the whole trial duration (*Fig 2 & Fig S2*). This is in line with multiple previous studies, which have identified sustained encoding of decision variables in motor preparatory signals (Murphy et al., 2021; Parés-Pujolràs et al., 2025; Wyart et al., 2012, 2015). Put together, our findings support the idea that in this task the CPP encodes discrete rounds of evidence accumulation which provide belief updates to downstream motor signals that maintain a sustained representation of the evolving DV.

In our two experiments, both the behavioural effects and the CPP profile could be reproduced with a simple bounded accumulation model. Our model assumed that participants stop accumulating evidence if a bound is hit and that accumulation is paused during the interval between pulses. By assuming that the CPP becomes inactive and falls down during the gaps, our simulations could recapitulate key CPP characteristics observed in our task including the opposite effects of the first and second pulse coherence on its amplitude. The fact that trials with low coherence first pulses were less likely to be terminated early, and had to cover a longer

distance-to-bound upon presentation of the second pulse explained why CPP-P2 reached higher amplitudes in those trials.

A key feature of our model which enabled it to account for our behavioural observations was that decisions could be terminated before all of the evidence has been sampled and accumulated. Previous studies have shown that people sometimes terminate decisions early, even when this comes at a cost to accuracy and there is no apparent benefit to performance (Balsdon et al., 2020). Given that the motion in any two sequential pulses always went in the same direction and that ours was a delayed response task, the optimal strategy to solve our task was to perfectly accumulate all available evidence, and indeed previous tasks similar to ours (Azizi & Ebrahimpour, 2023; Kiani et al., 2013) found no evidence for early terminations. Whether or not participants engage in early decision terminations may depend on various factors, including task context and difficulty (Lange et al., 2021). For example, our study had comparatively long lead-in periods before the first pulse (min. 0.7s compared to min. 0.4s of previous studies), allowing more preparation time, which may have facilitated integration of the earliest evidence in a pulse. Further, even though the probability of a single pulse trial was the same, the single-pulse staircasing procedure we used at the beginning of the task may have encouraged a strategy more focused on the information in the first pulse.

Our findings regarding the CPP's contribution to decisions based on temporally separate bouts of evidence aligns with the results of recent work in expanded judgement tasks (Parés-Pujolràs et al., 2025; Wyart et al., 2012, 2015). In expanded judgement studies, participants view a sequence of discrete easily identified tokens that support one of two choices and must integrate over multiple tokens to make an accurate decision. In these tasks, CPP responses to discrete tokens have been shown to scale with context-optimal belief updates in classic non-volatile contexts that prescribe perfect linear accumulation (Wyart et al., 2012, 2015) as well as in volatile environments where the correct answer may change within a given trial, which prescribe a different integration strategy (Parés-Pujolràs et al., 2025). That is, evoked CPP activities track how much a DV changed following each piece of evidence, but not the absolute running sum of accumulated evidence. This stands in contrast to the results from conventional decision-making tasks where evidence is presented continuously, and in the form of high-noise continuous stimuli (Geuzebroek et al., 2023; Gherman et al., 2024; Steinemann et al., 2018) rather than low-noise discrete tokens. Our task represents a middle ground between these two types of paradigms: we used noisy continuous evidence, but grouped

it in two temporally discrete, sequential windows. Thus, participants needed to integrate over time to discriminate the information in each evidence pulse, but also across evidence pulses to maximise their accuracy. This allowed us to show that, in conditions where evidence is noisy and continuous but can be paused, the CPP's profile tracks multiple rounds of evidence accumulation. When such updates occur in discrete steps as in expanded judgement tasks, or in temporally discrete bouts as in our gaps task, the CPP does not seem to keep track of the state of a DV. However, it remains unclear to what extent the CPP may behave differently in other contexts. In particular, a common feature between expanded judgement tasks and our design is the lack of temporal ambiguity. Here, we used colour cues to remove uncertainty about the presence of evidence. This external cue may have acted as a trigger to cue the initiation and termination of two rounds of evidence accumulation. However, it is plausible that in the absence of such colour cues the CPP may have exhibited a sustained pattern of activity consistent with continuous monitoring of the environment, potentially plateauing during periods with incoherent dot motion that should elicit close-to-zero, or heavily reduced, momentary belief updates. Future research is needed to further elucidate the mechanisms behind the CPP activity during low evidence periods.

Given that the effects of gap duration on behaviour were not expected and were inconsistent across the two experiments (*Fig 1B*), we did not aim to capture them in our modelling. However, we hypothesized that they might be influenced by attentional fluctuations related to the temporal structure of the task, and we ran some exploratory analyses to preliminarily investigate this possibility. A robust body of work has shown that attentional orientation in time plays an important part in information processing (Correa et al., 2005, 2006). Furthermore, both the temporal statistics and the probability of events occurring in time have a huge impact on behaviour, even in simple reaction time tasks (Grabenhorst et al., 2019, 2021). Yet, it remains unclear how exactly these temporal expectations and their associated attentional modulations influence the way in which we integrate discrete, temporally separate pieces of information over time. Occipital Alpha-band power has been linked to attention and temporal expectations (Rohenkohl & Nobre, 2011), and also to the onset (Van Den Brink et al., 2021) and build-up rate (Kelly & O'Connell, 2013) of evidence accumulation signals like the CPP. In our two experiments, the delays between the two pulses followed different distributions: in Exp 1 gaps were uniformly distributed, whereas in Exp 2 they were skewed towards shorter gaps durations. We found evidence that occipital alpha power kept track of this difference (*Fig 5*), with stronger desynchronisation during the longer gaps in experiment 1, where they were

relatively more likely. We also found that evoked alpha desynchronisation additionally modulated the weight of the second pulse on behaviour in experiment 2 (*Fig S11*), suggesting that phasic alpha activations following may influence evidence integration. However, more work is required to specifically test how tonic and phasic occipital Alpha may interact with evidence accumulation processes, or influence behaviour through other mechanisms.

In sum, our study provides novel insight into the neurocomputational mechanisms supporting decision making in a context where decisions need to be made on the basis of a single stimulus feature that is only intermittently available. Our task allowed us to identify novel functional characteristics of a well-established neural decision signal in the human EEG: the CPP. We show that in tasks where evidence is presented intermittently and each evidence pulse is difficult to discriminate the CPP does not maintain a sustained representation of cumulative evidence, but rather engages in multiple rounds of evidence accumulation, which in turn update DVs represented in downstream areas. This is consistent with recent work showing that the CPP encodes belief updates rather than the ongoing state of a DV in expanded judgment tasks involving the integration of multiple stimuli that are individually easy to identify (Parés-Pujolràs et al. 2025), and suggests that the CPP may only trace the entire evidence accumulation process underpinning decisions when the physical evidence is presented in one continuous stream (Geuzebroek et al., 2023; Kelly & O'Connell, 2013; O'Connell et al., 2012; Steinemann et al., 2018), instead representing discrete bouts of evidence accumulation when there are gaps between evidence samples. Our results thus add to growing research suggesting that the neural architecture underlying evidence accumulation for decision-making can exhibit flexible functional roles in different contexts, underpinning adaptive behaviour.

4. Methods

Participants: Twenty-three participants were recruited to take part in Experiment 1, and 23 different participants were recruited to take part in Experiment 2 of this study. In Experiment 1, one participant withdrew after a few blocks of the task because they were feeling dizzy, and thus their data were excluded from further analysis. Two participants were unable to perform the task at the minimum required level of accuracy and were excluded from analysis in Experiment 2. A total of 22 (M = 21.5 years old, SD = 1.79 years old, 10 females) and 21 (M = 22 years old, SD = 3.85 years old, 11 females) participants were thus included in the final

samples for Experiment 1 and 2 respectively. All participants provided informed written consent before participation in the study, and all procedures were approved by the University College Dublin ethics committee.

Task: Participants performed a motion direction discrimination task adapted from Kiani et al. 2013 and uses random dot kinematograms (RDK). The experiments consisted of a total of 700 (Exp 1) and 933 trials (Exp 2), divided in blocks of 70 trials each. Most trials in the experiment were two-pulse trials (85.7%) and the remaining fraction were single-pulse trials (14.3%). In single-pulse trials, the second pulse was not presented. Instead, the first pulse was followed by a 0% coherence lead-out period and then the trial terminated. The proportion of single- and two-pulse trials mirror the ones in the original Kiani et al. 2013 study.

Stimuli were presented centrally on a standard monitor, and participants sat at a 50-60cm distance from the monitor. Each trial began with a white fixation cross presented at the centre of the display on black background. Then, a patch of 0% coherence dots appeared (700 – 1050 ms) in blue colour on the centre of the screen. This was followed by a 200 ms pulse of yellow dots moving at either high or low coherence. After a gap of variable duration (Exp 1: [0, 0.5, 1] s, Exp 2: [0, 0.12, 0.36, 1.08] s) a second 200 ms pulse of either high or low coherence motion was presented. Finally, a patch of blue dots moving at 0% coherence was presented (700 to 1050 ms) in blue. The fixation cross was visible throughout the trial to encourage participants not to follow the dots with their eyes. At the end of the final blue dot sequence, all stimuli disappeared. All gap durations were equally likely, and all possible motion coherence combinations (low-low, low-high, high-low, high-high) and gap durations were randomly permuted throughout the experiment. The expectations for a long delay to occur differed in our two experiments. Given the percentage of 1-pulse trials, if the second pulse had not yet occurred 500 ms after the 1st pulse the probability of it happening was higher in Experiment 1 ($p = 0.666$ [85.7/3/100-2/3*85.7]) than in Experiment 2 ($p = 0.6$ [85.7/4/(100-3*/4*85.7)]. Participants were told they had to respond upon stimuli offset. Participants responded by pressing the “C” or “M” keys on a standard keyboard with the left and right thumb respectively. The blue and yellow colours used in the experiment were selected to be physically isoluminant, as measured on the lab monitor.

Staircase procedure: In order to estimate the required dot coherence for participants to perform at 70% accuracy across coherence levels, all participants performed a staircase procedure

consisting of 100 trials. All trials consisted of 0% coherence lead-in dots (700 to 1050 ms), a single pulse of evidence (200 ms) and a final sequence of 0% coherence lead-out dots (700 to 1050 ms). The staircase followed a 2-down, 1-up procedure (Levitt 1970), and the starting value was set to 40% for all participants. After two correct responses, the coherence was lowered. The initial step size was set to 6%, and it was halved every second time the staircase went down. This was done to achieve greater precision towards the end of the staircase. The average coherence of the dot motion in the last 40 trials of the staircase procedure were used to estimate the two coherence strengths used in the task. Specifically, high (Exp. 1: M = 44.73%, SD = 12.64%; Exp 2: M = 33.04%, SD = 13.41%) and low (Exp. 1: M = 26.75%, SD = 7.74%; Exp 2: M =19.82%, SD = 8.04%) coherences were then set to be 25% higher or lower than the staircase-estimated coherence.

To keep performance constant throughout the task, the average accuracy was checked after each block and coherence was dynamically adjusted between blocks. If accuracy was on average lower than 60% or higher than 90% across conditions, coherence was adjusted by adding or subtracting increasing or decreasing it by a tenth of its starting value (i.e. the one obtained after the staircase). This precaution was taken to avoid ceiling & floor effects. Most participants' performance was stable throughout the task, and only a few required substantial adjustments of coherence between blocks (see *Fig. S1* for individual data).

EEG recording and preprocessing: Continuous EEG was recorded from 128 scalp electrodes using a BioSemi system with a sample rate of 512 Hz. Eye movements were recorded with four electro-oculogram (EOG) electrodes, two above and below the left eye and two at the outer canthi of each for vertical (VEOG) and horizontal HEOG) eye. Additionally, we recorded electromyogram (EMG) signals from the thenar eminence using two electrodes on the left and two on the right hand. All data were analysed in a Matlab R2021a using the EEGlab toolbox (Brunner et al., 2013; Delorme & Makeig, 2004).

Continuous data were high-pass filtered using a 0.1Hz, and low-pass filtered at 35Hz using a zero-phase shift Butterworth filter. Then, stimulus-locked epochs were generated from -1.5s to 3s around the onset of the first evidence pulse. Bad channels were defined as those exceeding 50 standard deviation units, and were interpolated. Data were then re-referenced to the common average of all EEG channels. Trials where participants blinked during one of the evidence pulses were excluded from analysis. Eye movements occurring at any other time during the

trial were identified using Independent Component Analysis (ICA) and artefacts were removed from the data. Epochs were baselined to the 200 ms before the onset of the first evidence pulse, and channels showing ±150µV oscillations were interpolated on a trial-by-trial basis using the TBT toolbox (Ben-Shachar, 2018). If any given epoch had more than 20 bad channels, it was marked for rejection. Further, if any given channel was marked as "bad" on >30% of epochs, it was interpolated across the whole experiment. Finally, we applied a Current Source Density (CSD) transform (second spatial derivative) using the CSD toolbox (Kayser & Tenke, 2006) to the preprocessed data to reduce the spatial spread and thereby reduce overlap between topographic foci. Time frequency decomposition was performed using a 7-cycle morlet wavelet transform, extracted in 1Hz steps in the 1-40Hz range. Alpha and beta band power were obtained by averaging power in the 8-12Hz and 13-30Hz frequencies, respectively.

Behavioural analysis:

First, we investigated the effect of pulse coherence and gap duration on participants' accuracy, we run the following logistic regressions on one-pulse (*Eq. 1*) and two-pulse trials (*Eq. 2*):

$$Logit(Accuracy) \sim \beta_0 + \beta_1 p1coh + \left(\beta_{2,s} + \beta_{1,s} p1coh | ID\right) \quad (1)$$

$$Logit(Accuracy) \sim \beta_0 + \beta_1 p1coh + \beta_2 p2coh + \beta_3 Gap + \beta_4 p1coh * Gap + \\ + \beta_5 p2coh * Gap + \left(\beta_{6,s} | ID\right) \quad (2)$$

Next, to test whether participants were generally more accurate in two-pulse compared to one-pulse trials we run the following logistic regression (*Eq. 3*):

$$Logit(Accuracy) \sim \beta_0 + \beta_1 nPulses + \beta_2 p1Coh + \beta_3 nPulses * p1Coh \\ \left(\beta_{0,s} + \beta_{1,s} nPulses + \beta_{2,s} p1Coh \middle| ID\right) \quad (3)$$

This analysis was restricted to two-pulse trials with equal coherences in the two pulses (HH, LL trials only) in order to make one- and two-pulse trials comparable.

Then, to investigate the effect of pulse order on participants' accuracy, we run the following regression (*Eq. 4*):

$$Logit(Accuracy) \sim \beta_0 + \beta_1(p1coh + p2coh) + \beta_2(p2coh - p1coh) + \left(\beta_{0,s} + \beta_{2,s}(p1coh + p2coh) + \beta_{3,s}(p2coh - p1coh)\middle|ID\right) \quad (4)$$

where the difference coefficient $ß_2$ is expected to not be significant if the order of pulses does not influence performance.

In all regressions, coherence was coded as a factor (*high/low*) because coherence values were staircased for each participant individually and varied widely. Where possible, we included all main effects of interest as random effects to control for interindividual variability. These regressions were performed on R using the anova and lme4 toolboxes (Bates et al., 2015).

Finally, we compared participants' accuracy on two-pulse trials to the expected performance assuming perfect integration, and based on their accuracy in one-pulse trials. Following Kiani et al. 2013, we assumed that decisions are determined by the sign of the sum of evidence in the two pulses. The evidence distribution was inferred from participants' accuracy on single-pulse trials as follows (*Eq. 5 -7*):

$$e1 \sim \Phi^{-1}(P_1, 0,1)$$
$$e2 \sim \Phi^{-1}(P_2, 0,1) \quad (5)$$

Where P1 and P2 are grand-average accuracies for high- and low-coherence one-pulse trials. **Φ** is the normal cumulative distribution function:

$$\Phi(s, \mu, \sigma) \sim \int_{-\infty}^{0} N(v, \mu, \sigma) dv \quad (6)$$

and $\mathbf{\Phi}^{-1}$ is its inverse. N(v,μ,σ) is the normal probability density function with mean (μ) and standard deviation (σ). The expected accuracy for double-pulse trials was then calculated as follows:

$$P_e \sim 1 - \Phi(0, e1 + e2, \sqrt{2}) \quad (7)$$

ERP analysis:

CPP-P1 & CPP-P2 analysis: To investigate how centroparietal responses to the two sequential pulses reflected motion coherence, we extracted EEG activity from a cluster of centroparietal electrodes where activity at 500 ms after pulse 1 was maximal across conditions (see *Fig 2*). For CPP-P2 visualisation, we removed P1-related activity by subtracting the traces from the longest delay conditions from all traces. This resulted in an average EEG activity of 0 (see *Fig. S4*) preceding the second pulse, in all conditions, and rendered the key CPP-P2 findings more clearly visible by removing P1-induced EEG drifts. We report the average across all non-zero gap conditions in the main manuscript (*Fig. 3*), as well as the data sorted by gap in the supplementary material (*Fig. S3*).

To investigate the impact of dot motion coherence on both CPP-P1 and CPP-P2 amplitudes, we run a series of time-resolved regressions:

$$CPP_t \sim |p1Coh| + |p2Coh| \tag{8}$$

Where *t* indicates time from P1 for CPP-P1 analysis, and time after P2 for CPP-P2 analysis. For this analysis, we aligned all gap conditions to P1/P2 onset respectively and ran the regressions for all trials pooled, for each subject separately. For CPP-P2, data were additionally baselined using 50ms before P2 onset to minimise the impact of gap duration on average amplitudes. All EEG data were low-pass filtered (6Hz) prior to running the regression to reduce the impact of high frequency noise on these single-trial analyses.

Single-trial fluctuation analysis:

Next, to investigate how single-trial fluctuations in attention and evidence accumulation influenced performance, we run a series of time-resolved logistic regressions.

First, we investigated whether single-trial variability in centroparietal responses modulated the weight of pulses on choice. The interaction terms between pulse coherence and CPP amplitude in the following regression (*Eq. 9*) thus test whether pulses which evoked a particularly high centroparietal response were also associated with that pulse having a greater weight on choice:

$$Logit(P(Choice = right)) \sim p1Coh + p2Coh + p1Coh * CPP_t + p2Coh * CPP_t \quad (9)$$

Where *t* indicates time after P1 for CPP-P1 analysis, and time after P2 for CPP-P2 analysis.

Second, we run an analogous analysis to investigate whether fluctuations in evoked occipital alpha desynchronisation also modulated the weight of each pulse on choice (*Eq. 10*):

$$Logit(P(Choice = right)) \sim p1Coh + p2Coh + p1Coh * -Alpha_t + p2 * -Alpha_t \quad (10)$$

Where *t* indicates time after P1 for alpha-P1 analysis, and time after P2 for alpha-P2 analysis. All regressors and EEG signals were normalised across trials by timepoint. For alpha power analysis, we took the *negative* alpha power across a range of bilateral occipital electrodes (see *highlighted* electrodes in *Fig. 5*), so that positive coefficient values indicate that higher attention (i.e. lower alpha power) increase weight on choice (*Eq. 10*). For the CPP analysis, we used the signal averaged over a cluster of centroparietal electrodes for CPP (*Fig. 2B*). Note also that absolute (unsigned) values of P1 and P2 coherence regressors were taken in all regressions where CPP is the dependent variable (*Eq. 8*), because the CPP's polarity is always positive regardless of the choice alternative supported by the evidence. Regressions were run for each gap condition separately for times between [0-1s] after each pulse, and the resulting coefficients were aligned to P1 and P2 and averaged across conditions.

EEG statistics: EEG evoked potentials and regression results were statistically evaluated against zero using cluster-based permutations using the FieldTrip toolbox (Maris & Oostenveld, 2007; Oostenveld et al., 2011). All tests were two-tailed, cluster-corrected ($p<0.05$), and 10000 permutations were run. We additionally compared the magnitude of the occipital alpha resynchronisation in the longest gap conditions using an independent samples t-test in Matlab.

Model fitting & simulations: In order to account for participants' behaviour and its key neural signatures during evidence accumulation, we fit a simple diffusion model to the grand-averaged accuracy in the 4 conditions with three free parameters: bound (*b*) and two drift rates (*d1*, *d2*) for the high and low coherence conditions separately. The accumulator model followed this basic equation:

$$y_t \sim y_{dt} + x_t * d \tag{11}$$

Where yt is the accumulator output, and xt indicates evidence (i.e. motion coherence) at any given time. The evidence was set to 1 during low coherence pulses, the empirically observed high/low coherence ratio (Exp. 1: 1.67, Exp. 2: 1.77) during high coherence pulses, and 0 otherwise. Coherence was scaled by drift rate '*d'*, and Gaussian noise, N(0,σ), was added in order to represent noisy sensory evidence through the trial. Given the behavioural observation that overall accuracy did not decrease in longer gap durations, and the fact that evidence pulses were clearly indicated by a change in the dots colour, the accumulator was "turned off" and not allowed to accumulate pure noise during the 0% coherence motion delay separating the two pulses. When the accumulator crossed a bound '*b',* evidence accumulation terminated. The accumulator assumed no noise or leak during the gap, and in any trials where the bound was not crossed accumulation continued upon presentation of the next pulse.

We fit the model by minimising G^2 with the *particleswarm* function in the Global Optimisation Toolbox for Matlab. G^2 is a goodness of fit measure that in this study quantified how much predicted accuracies differed from observed ones in each condition.

Computational models were fit to the two experiments separately simulating 10000 trials for each condition, and the best fitting parameters were then used to simulate CPP traces. Akaike Information Criterion (AIC) and Bayes Information Criterion (BIC) were computed to assess the goodness of fit as follows:

$$AIC = G^2 + 2 * k \tag{12}$$

$$BIC = G^2 + k * \log(n) \tag{13}$$

Where *k* indicates the number of free parameters for a given model, and *n* represents the number of trials used to fit the model.

Neural data simulation:

It has been suggested that the CPP represents the absolute value of a signed decision variable. Following previous work, we thus simulated the CPP by taking the absolute value of the signed

accumulator traces produced by the best fitting parameters of our model (Afacan-Seref et al., 2018; O'Connell & Kelly, 2021). If the simulated decision variable hit a bound, the accumulator traces were set to fall back down to zero within 100 ms after hitting the bound, based on past observations of post-peak CPP dynamics. If the CPP had not hit a bound after the time of the first pulse, it was assumed to stay stable and resume accumulation upon the presentation of the second pulse.

Acknowledgements

This work was supported by a grant from Science Foundation Ireland (now Taighde Éireann) under grant number 19/US/3599 and an Investigator Award from the Wellcome Trust under grant number 219572/Z/19/Z (to S.P.K). The authors would also like to thank Mike Shadlen for discussions on this work, and students Liang Tong and Bowen Duan for assistance during data collection.

Author contributions

Conceptualisation: E.P-P, & S.P.K; software: E.P-P; data curation & analysis: E.P-P.; resources: S.P.K.; writing of original draft: E.P-P.; review & editing: E.P-P, R.G.O., A.C.G, & S.P.K.; supervision: S.P.K.

REFERENCES

Afacan-Seref, K., Steinemann, N. A., Blangero, A., & Kelly, S. P. (2018). Dynamic Interplay of Value and Sensory Information in High-Speed Decision Making. *Current Biology*, *28*(5), 795-802.e6. https://doi.org/10.1016/J.CUB.2018.01.071

Azizi, Z., & Ebrahimpour, R. (2023). Explaining Integration of Evidence Separated by Temporal Gaps with Frontoparietal Circuit Models. *Neuroscience*, *509*, 74–95. https://doi.org/10.1016/J.NEUROSCIENCE.2022.10.019

Balsdon, T., Mamassian, P., & Wyart, V. (2021). Separable neural signatures of confidence during perceptual decisions. *ELife*, *10*. https://doi.org/10.7554/eLife.68491

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

Bates, D., Mächler, M., Bolker, B. M., & Walker, S. C. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, *67*(1), 251–264. https://doi.org/10.18637/jss.v067.i01

Ben-Shachar, M. S. (2018). TBT: Reject and Interpolate channels on a trial-by-trial basis. *Zenodo*.

Bronfman, Z. Z., Brezis, N., & Usher, M. (2016). Non-monotonic Temporal-Weighting Indicates a Dynamically Modulated Evidence-Integration Mechanism. *PLOS Computational Biology*, *12*(2), e1004667. https://doi.org/10.1371/JOURNAL.PCBI.1004667

Brunner, C., Delorme, A., Makeig, S., & Diego, S. (2013). *EEGLAB – AN OPEN SOURCE MATLAB TOOLBOX FOR ELECTROPHYSIOLOGICAL RESEARCH*. *c*, 8–9. https://doi.org/10.1515/bmt-2013-4

Brunton, B. W., Botvinick, M. M., & Brody, C. D. (2013). Rats and humans can optimally accumulate evidence for decision-making. *Science*, *340*(6128), 95–98. https://doi.org/10.1126/SCIENCE.1233912/SUPPL_FILE/BRUNTON.SM.PDF

Correa, Á., Lupiáñez, J., Madrid, E., & Tudela, P. (2006). Temporal attention enhances early visual processing: a review and new evidence from event-related potentials. *Brain Research*, *1076*(1), 116–128. https://doi.org/10.1016/J.BRAINRES.2005.11.074

Correa, Á., Lupiáñez, J., & Tudela, P. (2005). Attentional preparation based on temporal expectancy modulates processing at the perceptual level. *Psychonomic Bulletin & Review*, *12*(2), 328–334. https://doi.org/10.3758/BF03196380

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*(1), 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009

Geuzebroek, A. C., Craddock, H., O'Connell, R. G., & Kelly, S. P. (2023). Balancing true and false detection of intermittent sensory targets by adjusting the inputs to the evidence accumulation process. *ELife*, *12*. https://doi.org/10.7554/ELIFE.83025

Gherman, S., Markowitz, N., Tostaeva, G., Espinal, E., Mehta, A. D., O'Connell, R. G., Kelly, S. P., & Bickel, S. (2024). Intracranial electroencephalography reveals effector-independent evidence accumulation dynamics in multiple human brain regions. *Nature Human Behaviour 2024 8:4*, *8*(4), 758–770. https://doi.org/10.1038/s41562-024-01824-9

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574. https://doi.org/10.1146/annurev.neuro.29.051605.113038

Grabenhorst, M., Maloney, L. T., Poeppel, D., & Michalareas, G. (2021). Two sources of uncertainty independently modulate temporal expectancy. *Proceedings of the National Academy of Sciences of the United States of America*, *118*(16), e2019342118. https://doi.org/10.1073/PNAS.2019342118/SUPPL_FILE/PNAS.2019342118.SAPP.PDF

Grabenhorst, M., Michalareas, G., Maloney, L. T., & Poeppel, D. (2019). The anticipation of events in time. *Nature Communications 2019 10:1*, *10*(1), 1–13. https://doi.org/10.1038/s41467-019-13849-0

Hanks, T. D., & Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron*, *93*(1), 15–31. https://doi.org/10.1016/J.NEURON.2016.12.003

Heekeren, H. R., Marrett, S., & Ungerleider, L. G. (2008). The neural systems that mediate human perceptual decision making. *Nature Reviews Neuroscience 2008 9:6*, *9*(6), 467–479. https://doi.org/10.1038/nrn2374

Hyafil, A., Rocha, J. de la, Pericas, C., Katz, L. N., Huk, A. C., & Pillow, J. W. (2023). Temporal integration is a robust feature of perceptual decisions. *ELife*, *12*. https://doi.org/10.7554/ELIFE.84045

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., Churchland, A. K., & Shadlen, M. N. (2013). Integration of direction cues is invariant to the temporal gap between them. *Journal of Neuroscience*, *33*(42), 16483–16489. https://doi.org/10.1523/JNEUROSCI.2094-13.2013

Lange, R. D., Chattoraj, A., Beck, J. M., Yates, J. L., & Haefner, R. M. (2021). A confirmation bias in perceptual decision-making due to hierarchical approximate inference. *PLOS Computational Biology*, *17*(11), e1009517. https://doi.org/10.1371/JOURNAL.PCBI.1009517

Maris, E., & Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*(1), 177–190. https://doi.org/10.1016/j.jneumeth.2007.03.024

Mazurek, M. E., Roitman, J. D., Ditterich, J., & Shadlen, M. N. (2003). A Role for Neural Integrators in Perceptual Decision Making. *Cerebral Cortex*, *13*(11), 1257–1269. https://doi.org/10.1093/CERCOR/BHG097

McGovern, D. P., Hayes, A., Kelly, S. P., & O'Connell, R. G. (2018). Reconciling age-related changes in behavioural and neural indices of human perceptual decision-making. *Nature*

Human Behaviour, *2*(12), 955–966. https://doi.org/10.1038/S41562-018-0465-6;SUBJMETA=2629,2649,2811,378,477,631;KWRD=COGNITIVE+NEUROSCIENCE,HUMAN+BEHAVIOUR,NEUROSCIENCE,PSYCHOLOGY,SENSORIMOTOR+PROCESSING

Murphy, P. R., Wilming, N., Hernandez-Bocanegra, D. C., Prat-Ortega, G., & Donner, T. H. (2021). Adaptive circuit dynamics across human cortex during evidence accumulation in changing environments. *Nature Neuroscience*, *24*(7), 987–997. https://doi.org/10.1038/s41593-021-00839-z

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

O'Connell, R. G., & Kelly, S. P. (2021). *Neurophysiology of Human Perceptual Decision-Making*. https://doi.org/10.1146/annurev-neuro-092019

Okazawa, G., & Kiani, R. (2023). *Annual Review of Physiology Neural Mechanisms That Make Perceptual Decisions Flexible*. https://doi.org/10.1146/annurev-physiol-031722

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J. M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*. https://doi.org/10.1155/2011/156869

Parés-Pujolràs, E., Kelly, S. P., & Murphy, P. R. (2025). Dissociable encoding of evolving beliefs and momentary belief updates in distinct neural decision signals. *Nature Communications 2025 16:1*, *16*(1), 1–14. https://doi.org/10.1038/s41467-025-58861-9

Parés-Pujolràs, E., Travers, E., Ahmetoglu, Y., & Haggard, P. (2021). Evidence accumulation under uncertainty – a neural marker of emerging choice and urgency. *NeuroImage*, *232*(May 2021). https://doi.org/10.1101/2020.06.30.179622

Pfurtscheller, G., & Lopes Da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: Basic principles. *Clinical Neurophysiology*, *110*(11), 1842–1857. https://doi.org/10.1016/S1388-2457(99)00141-8

Rohenkohl, G., & Nobre, A. C. (2011). Alpha Oscillations Related to Anticipatory Attention Follow Temporal Expectations. *Journal of Neuroscience*, *31*(40), 14076–14084. https://doi.org/10.1523/JNEUROSCI.3387-11.2011

Rouault, M., Weiss, A., Lee, J. K., Drugowitsch, J., Chambon, V., & Wyart, V. (2022). Controllability boosts neural and cognitive signatures of changes-of-mind in uncertain environments. *ELife*, *11*. https://doi.org/10.7554/ELIFE.75038

Shadlen, M. N., & Kiani, R. (2013). Decision making as a window on cognition. In *Neuron* (Vol. 80, Issue 3, pp. 791–806). https://doi.org/10.1016/j.neuron.2013.10.047

Sospedra, A., Canals, S., Marcos, E., Hernández, M., & Joan, S. (2025). A temporal gap in sensory streams amplifies the influence of subsequent input on decision making. *BioRxiv*, 2024.07.17.603868. https://doi.org/10.1101/2024.07.17.603868

Stancák, A., & Pfurtscheller, G. (1995). Desynchronization and recovery of β rhythms during brisk and slow self-paced finger movements in man. *Neuroscience Letters*, *196*(1–2), 21–24. https://doi.org/10.1016/0304-3940(95)11827-J

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*(1), 3627. https://doi.org/10.1038/s41467-018-06117-0

Thut, G., Nietzel, A., Brandt, S. A., & Pascual-Leone, A. (2006). α-Band electroencephalographic activity over occipital cortex indexes visuospatial attention bias and predicts visual target detection. *Journal of Neuroscience*, *26*(37), 9494–9502. https://doi.org/10.1523/JNEUROSCI.0875-06.2006,

Tohidi-Moghaddam, M., Zabbah, S., Olianezhad, F., & Ebrahimpour, R. (2019). Sequence-dependent sensitivity explains the accuracy of decisions when cues are separated with a gap. *Attention, Perception, and Psychophysics*, *81*(8), 2745–2754. https://doi.org/10.3758/S13414-019-01810-8/FIGURES/6

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 Brink, R. L., Murphy, P. R., Desender, K., De Ru, N., & Nieuwenhuis, S. (2021). Temporal expectation hastens decision onset but does not affect evidence quality. *Journal of Neuroscience*, *41*(1), 130–143. https://doi.org/10.1523/JNEUROSCI.1103-20.2020

Waskom, M. L., & Kiani, R. (2018). Decision Making through Integration of Sensory Evidence at Prolonged Timescales. *Current Biology*, *28*(23), 3850-3856.e9. https://doi.org/10.1016/j.cub.2018.10.021

Wyart, V., de Gardelle, V., Scholl, J., & Summerfield, C. (2012). Rhythmic Fluctuations in Evidence Accumulation during Decision Making in the Human Brain. *Neuron*, *76*(4), 847–858. https://doi.org/10.1016/j.neuron.2012.09.015

Wyart, V., Myers, N. E., & Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. *Journal of Neuroscience*, *35*(8), 3485–3498. https://doi.org/10.1523/JNEUROSCI.3276-14.2015

Supplementary material

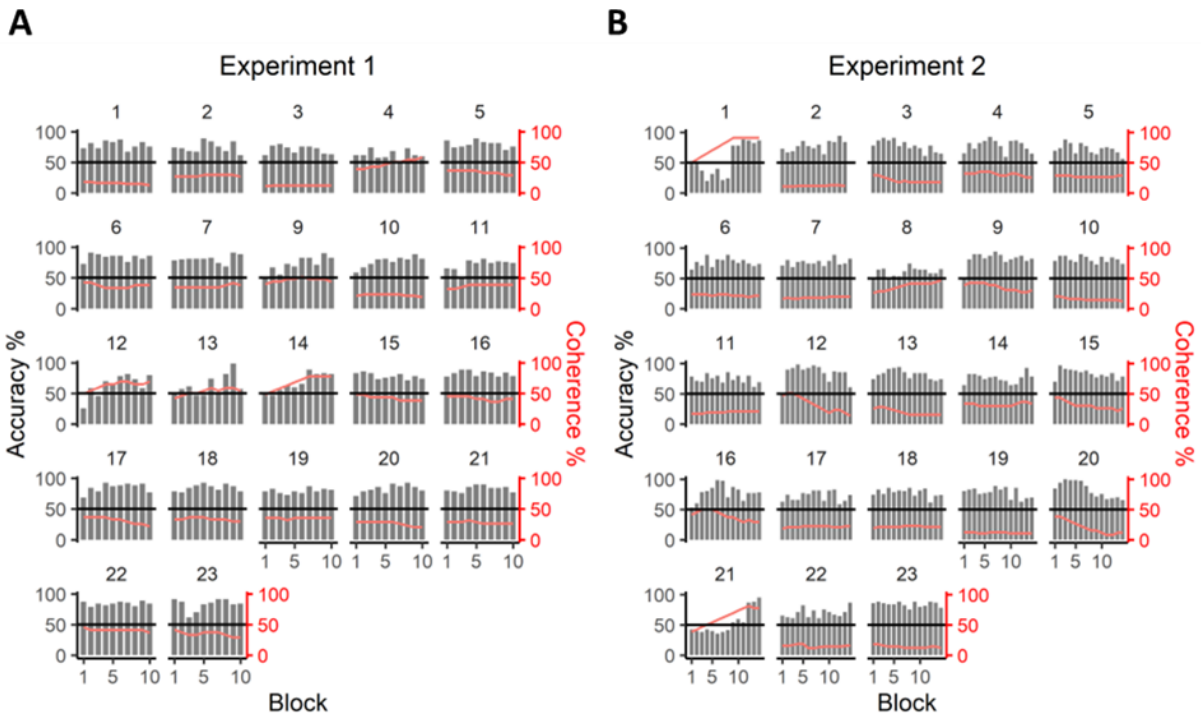

Figure S1. Individual participants' accuracy (*grey bars*) and average coherence (*red line*) for each block in Experiments 1 (**A**) and 2 (**B**). On average, most participants performed close to 70% accuracy and clearly above chance (*solid black line*), and only required minimal coherence adjustments between blocks. The data of participants 1 and 21 in Experiment 2 were excluded from further analysis due to poor task performance in a majority of blocks.

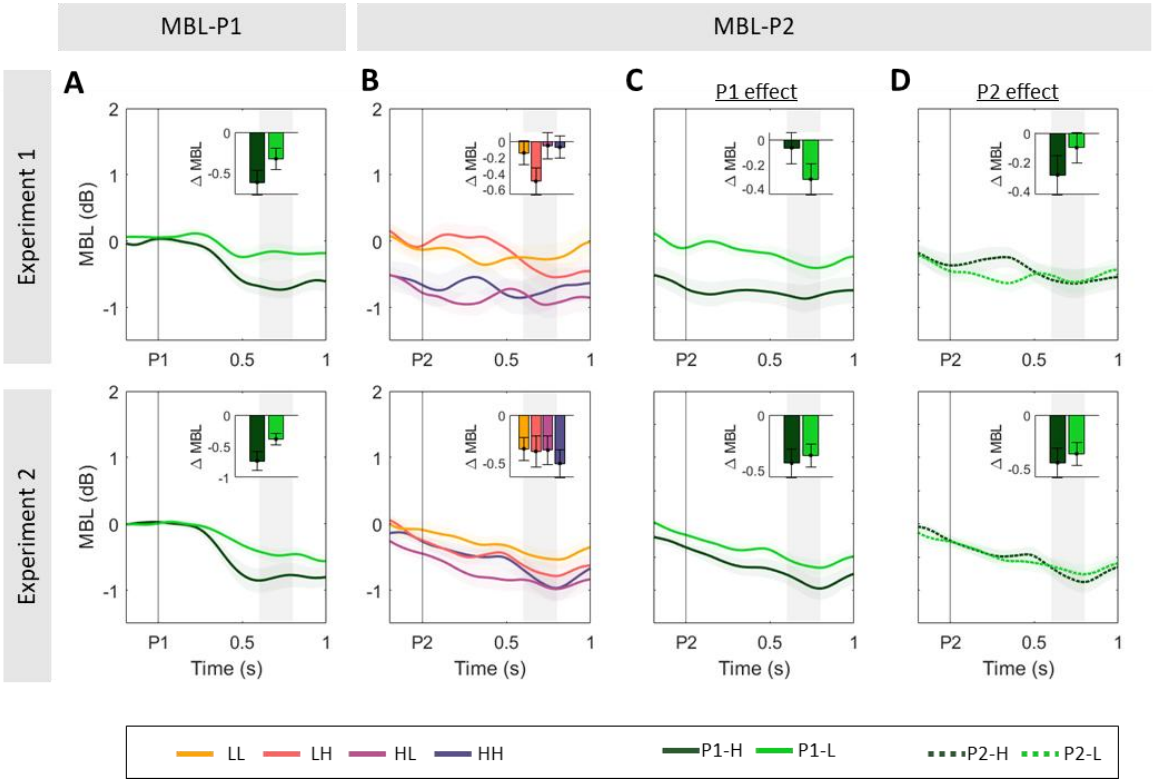

Figure S2. Sensitivity to evidence strength in the two pulses in motor beta lateralisation signals. A. Grand-averaged (±SEM) motor beta lateralisation [Contra-Ipsilateral hemispheres with respect to correct response hand] after pulse 1 (MBL-P1), averaged across all delay conditions. Motor preparation started shortly after the first pulse, and was sustained at a coherence-dependent level. **B.** Grand-averaged MBL after pulse 2 (MBL-P2), averaged over all non-zero gaps in the two experiments. **C.** Effect of P1 coherence on MBL after P2. Traces illustrate the grand-averaged (±SEM) aligned to pulse 2 (MBL-P2), averaged over all non-zero gaps and across P2 coherence. **D.** Effect of P2 coherence on MBL after P2. Traces illustrate the grand-averaged (±SEM) aligned to pulse 2 (MBL-P2), averaged over all non-zero gaps and across P1 coherence. In panels **A-D,** inset bar graphs illustrate the MBL excursion as the difference between activity at [0.6-0.8s, shaded area] after pulse onset, minus activity at baseline [-0.2-0s].

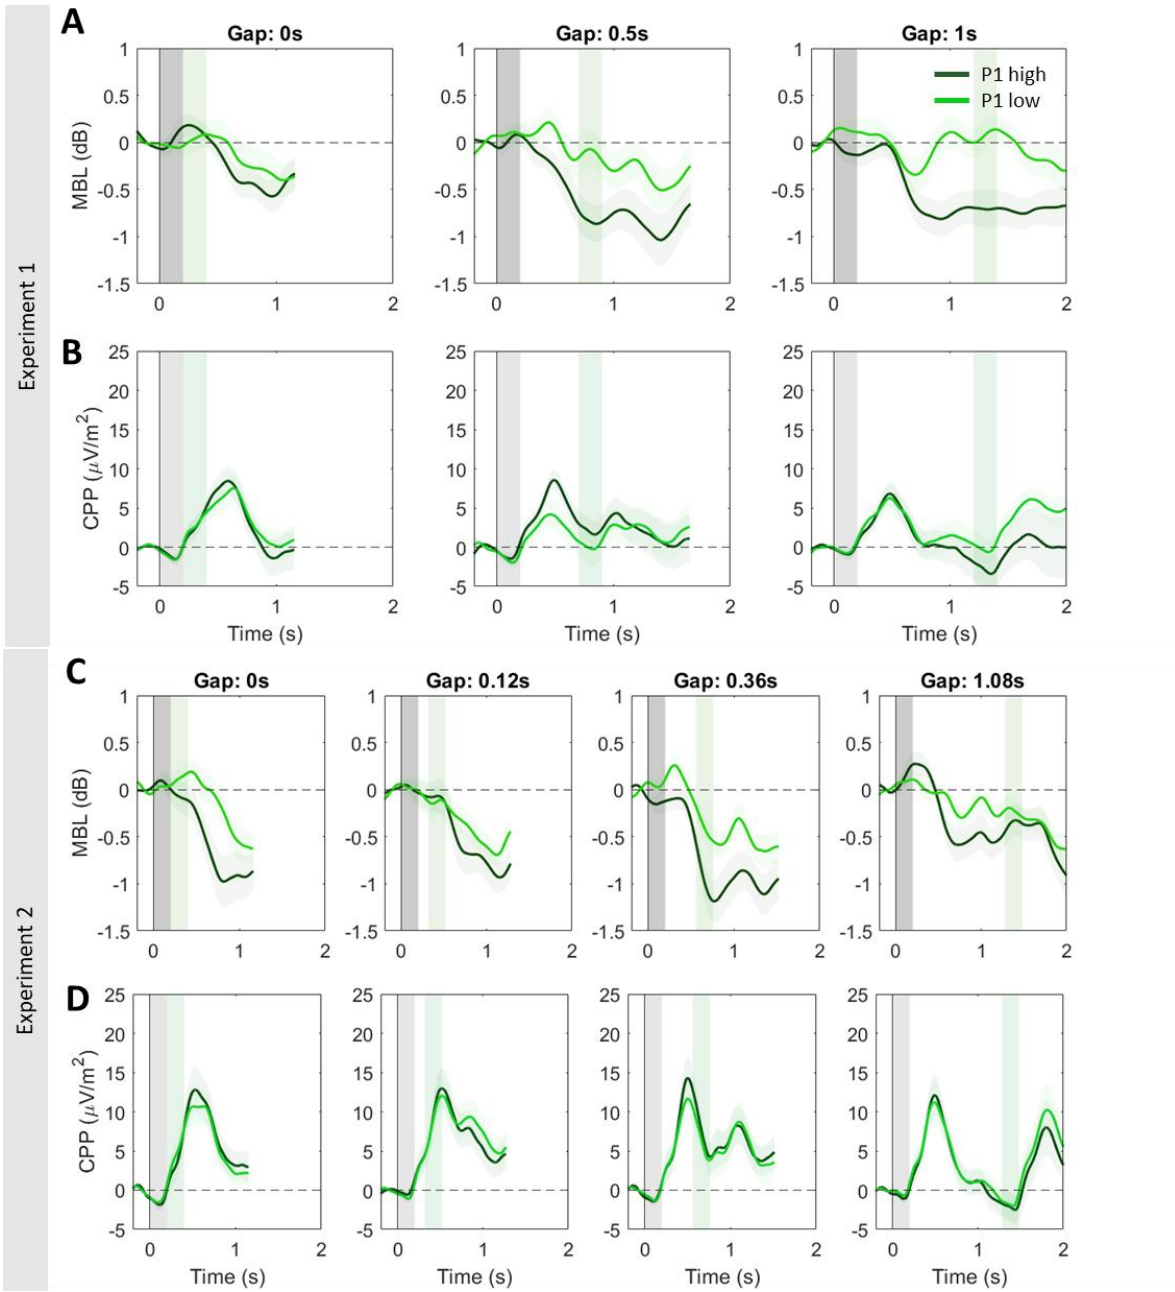

Figure S3. **A,C.** Motor beta lateralisation (MBL) for each gap condition, sorted by P1 coherence in experiments 1 (A) and 2 (C). Lateralisation is computed as a function of actual dot motion [Contra-Ipsilateral hemispheres with respect to correct response hand], sorted by gap duration and P1 coherence pooled across correct & error trials. **B,D.** Centroparietal activity for each gap condition, and sorted P1 coherence in experiments 1 (***B***) and 2 (***D***).

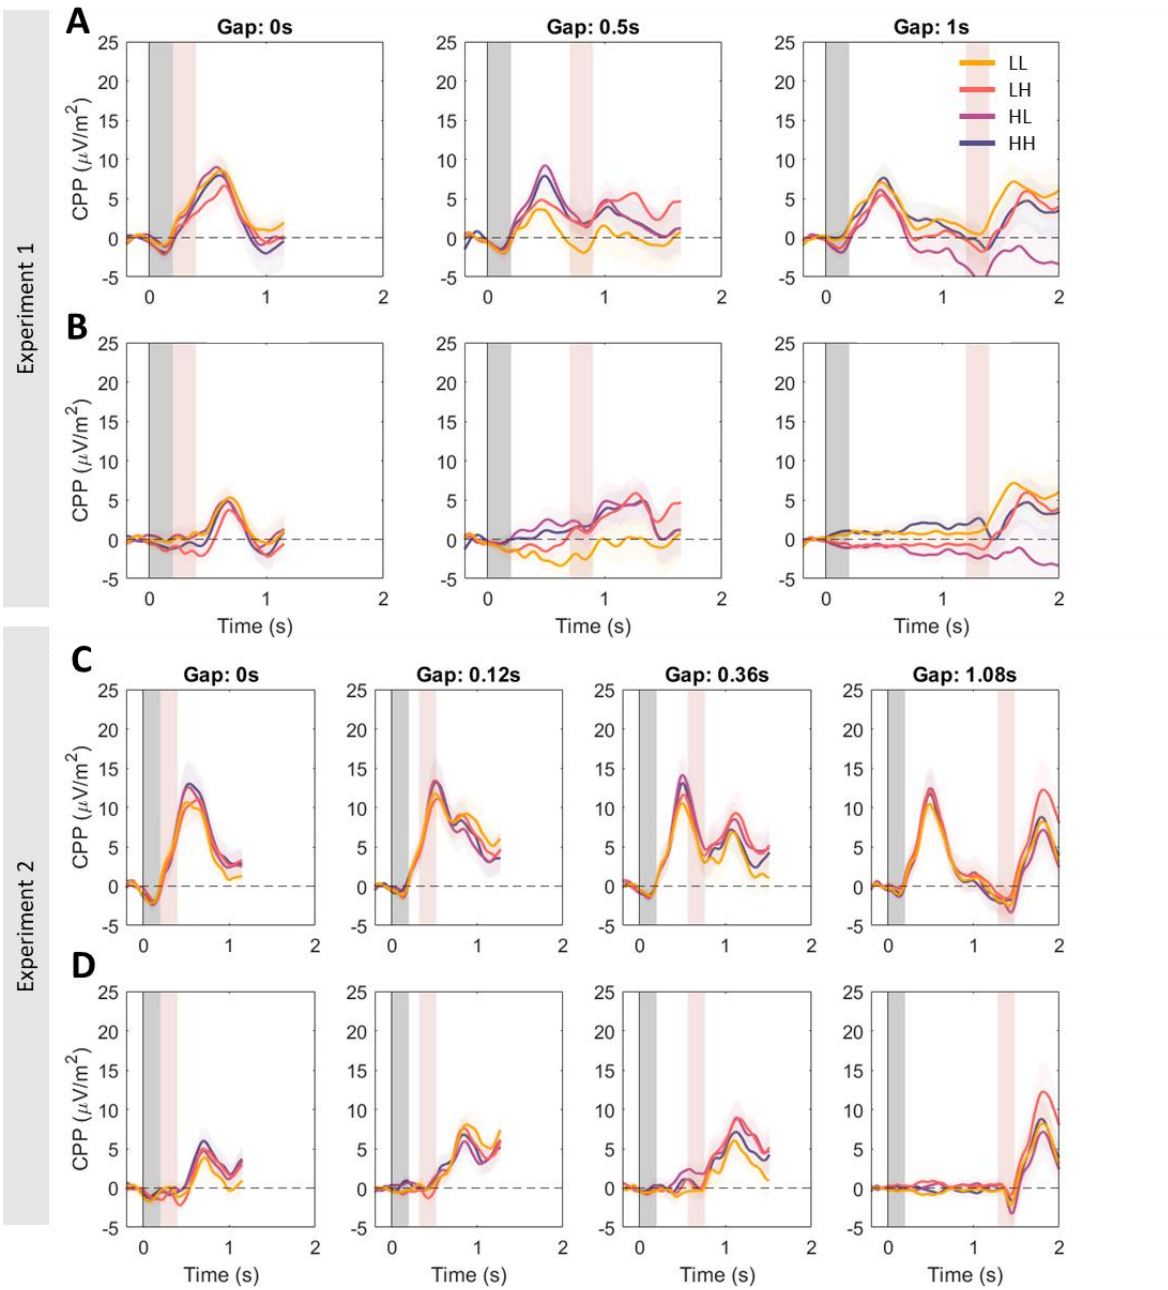

Figure S4. **A,C.** Centroparietal activity sorted by gap, P1 and P2 coherence in Experiments 1 and 2. **B,D.** Centroparietal activity after removing CPP-P1 activity. For each condition, the average of long-delay trials with the same P1 coherence was removed. That is, the average of the long-delay HH & HL conditions was subtracted from HH and HL traces. Similarly, the average activity of the long-delay LL & LH conditions was subtracted from LL and LH traces. This removes overlapping CPP-P1 potentials, and isolates centroparietal activity related to P2. Subtractions were performed between [-0.2 to 1.4s] from P1 presentation. Thus, CPP-P2 potentials in the long delay condition remain unaltered.

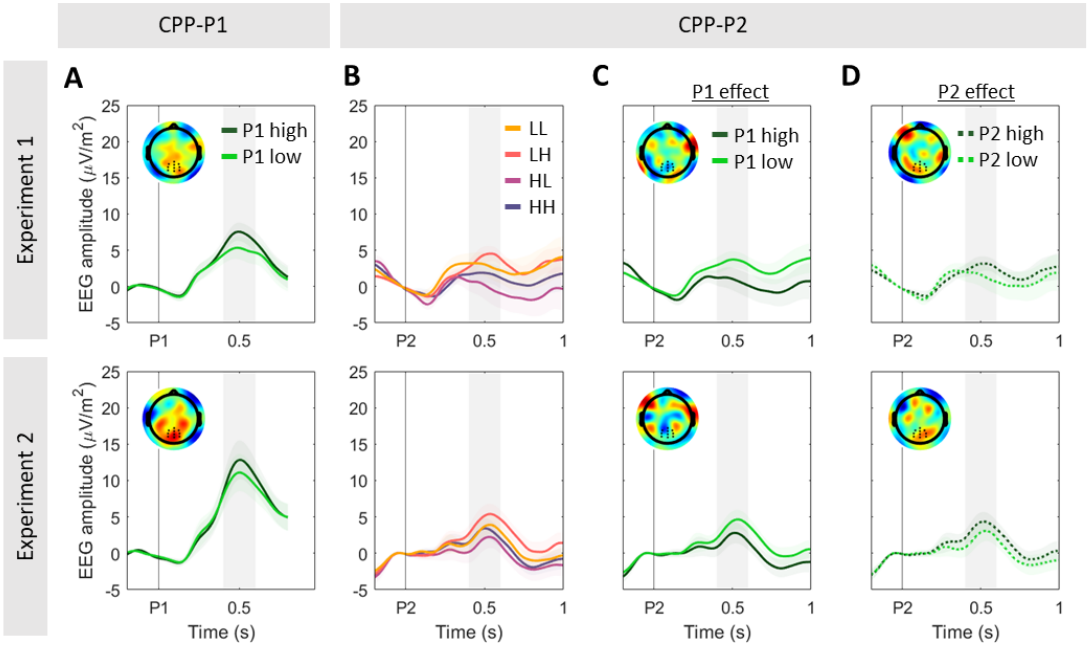

Figure S5. Replication of results in *Fig. 3*, now performing the analysis on CPP-P2 data (**B-D**) without removing CPP-P1 activity. Note that baseline activity shows a drift consistent with the overlapping potentials from the preceding pulse, which is corrected for when removing CPP-P1 activity (cf. Fig. 3). Note also that the same patterns are qualitatively present in the uncorrected data.

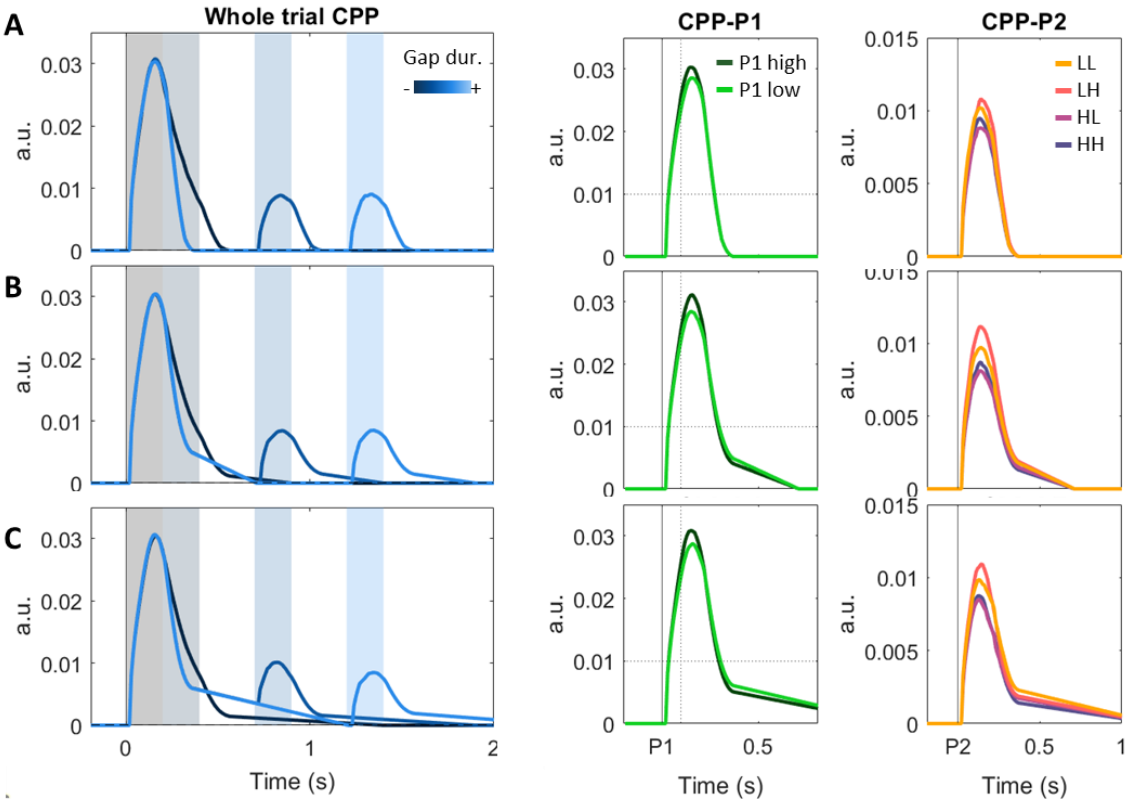

Figure S6. Simulations (N=1000) illustrating the effect of signal decay rate on CPP activity. Given that it's unclear whether the CPP fall-down to zero when evidence is paused follows the same dynamics as its fall-down when it hits a bound (where it typically reaches zero within 100ms), we simulated its CPP activity under three different assumptions about its decay rate upon evidence interruption (i.e. when blue dots appear onscreen after a pulse). We simulated the CPP assuming the signal takes 100ms (**A**), 500 ms (**B**) or 1000ms (**C**) to go back to baseline. Simulations are based on the fitted model in the main manuscript and for Exp. 1 only for illustration purposes. Longer signal decay times for trials where no bound was hit result in a protracted ramp down to baseline.

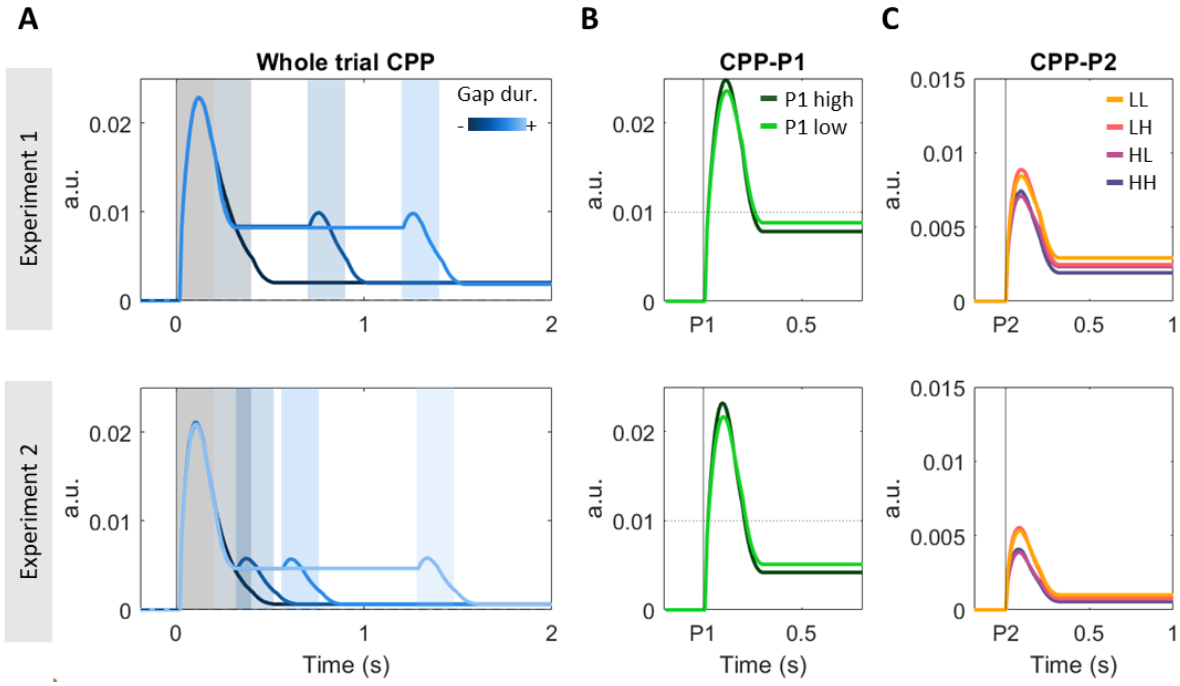

Figure S7. Simulated CPPs using the same computational model as in Fig. 4, but under the assumption that the signal only went back to baseline if a bound was hit. If the bound was not hit, the DV was set to stay stable at the final value reached by the end of pulse 1 with no decay and no added noise for the whole duration of the gap. Given that we do not model reaction times, we assumed no transmission delay between events on the screen and their impact on the accumulator. In this model, the partial fall down of the signal during the gap was driven by those trials that hit a bound following the first pulse. Conversely, the fact that it did not fall all the way to 0 on average during the gap is driven by those trials where the DV did not hit a bound during processing of pulse 1 and thus sustained through the gap. Accumulation resumed from that sustained level upon presentation of a second pulse, and the fact that on average CPP-P2 traces reach lower amplitudes than CPP-P1 evoked activities is explained by the fact that trials that were terminated during the first did not engage in accumulation again, and therefore took value 0 in our model. Note that this model provides a qualitatively worse fit to the empirical CPP traces during the delay and CPP-P2 trials. In particular, the magnitude of the CPP-P2 potentials was predicted to be significantly smaller than observed on the grand-averaged data.

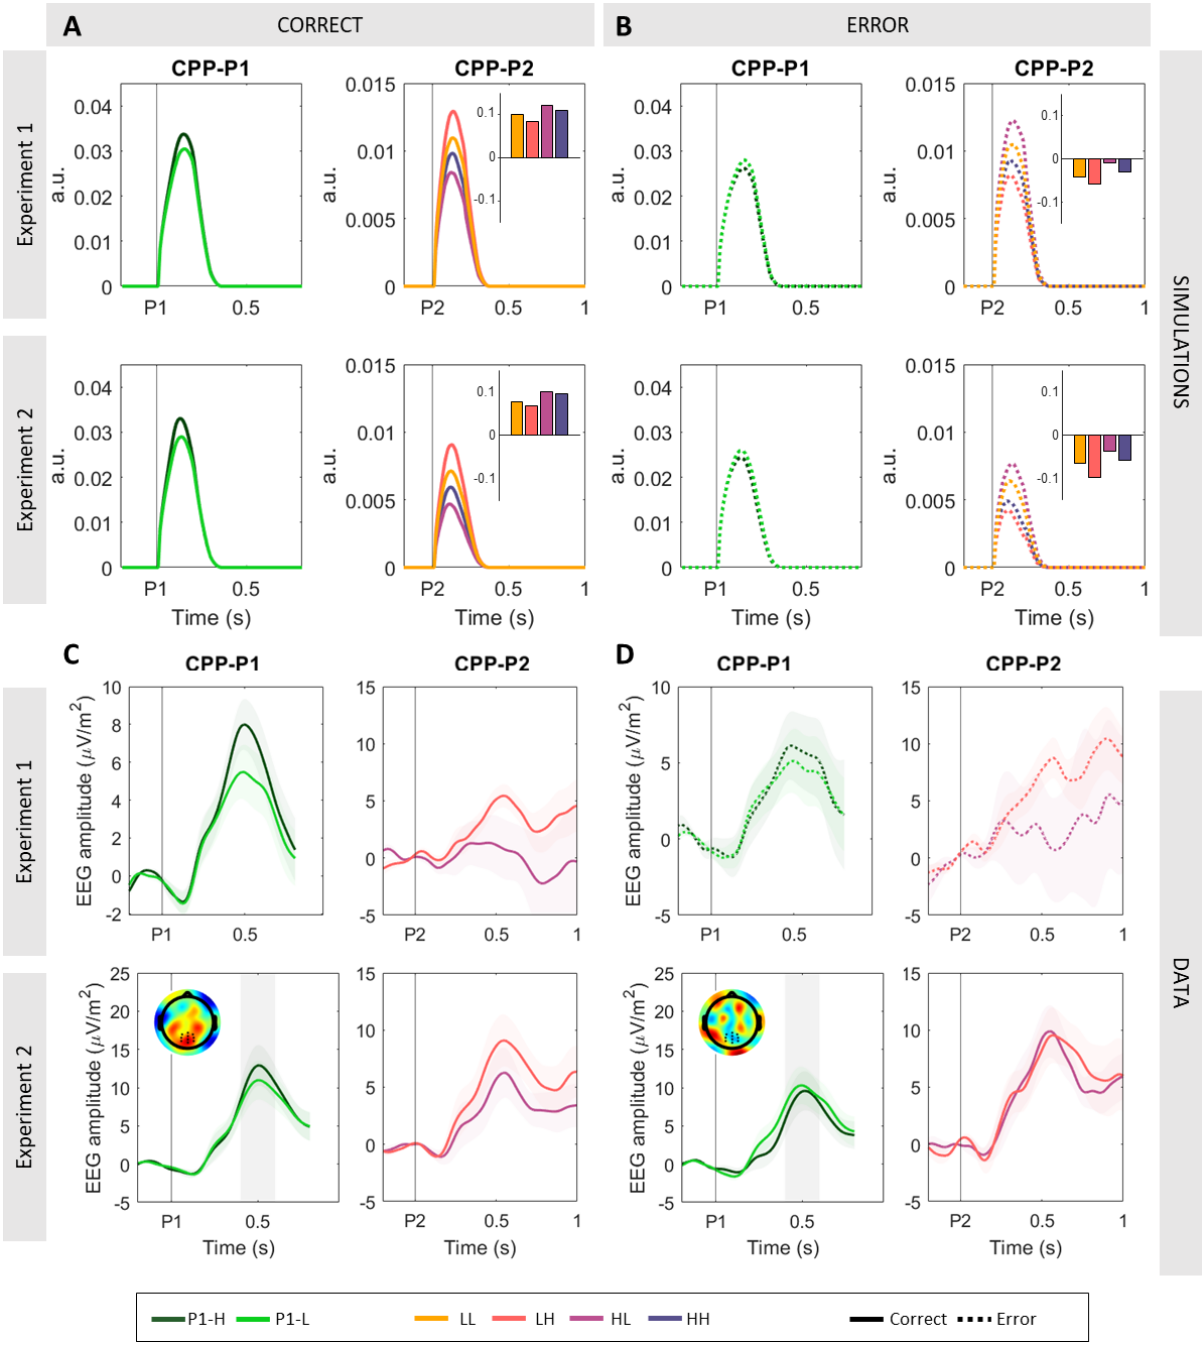

Figure S8. Model-predicted (**A-B**) and observed (**C-D**) CPP-P1 and -P2, in correct vs. incorrect responses, for Experiments 1 (*top*) and 2 (*bottom*). Bar graph insets in **A-B** indicate the state of the simulated DV at the end of pulse 1 in those trials where the bound was not hit. Positive values indicate that the DV was leaning towards the correct answer, whereas negative values indicate that it was leaning towards an error. The model predicts stronger P1 coherence effects in correct compared to error trials (A-B, *left*), with this effect also being observed in empirical data across both experiments (C-D, *left*). For the sake of clarity, we do not plot LL and HH conditions in the empirical CPP-P2 data because no clear amplitude differences are predicted by the model in correct vs. error trials, and the trial count in the HH condition was insufficient in most participants. Further, the model predicted a strong interaction in CPP-P2 amplitudes in the HH vs. HL condition. While correct trials (which were the majority of trials

in our task and the model) showed higher amplitudes at CPP-P2 for HL compared to the LH condition (A, *right*), consistent with the effect of P1 on CPP-P1 amplitudes reported in *Fig 4*, this effect reversed in incorrect trials (B, *right*). This interaction was also clearly visible in Experiment 2, although slightly less prominent in Experiment 1. Note however that Experiment 1 had fewer overall trials, and power to detect differences is likely to be insufficient after splitting conditions by accuracy.

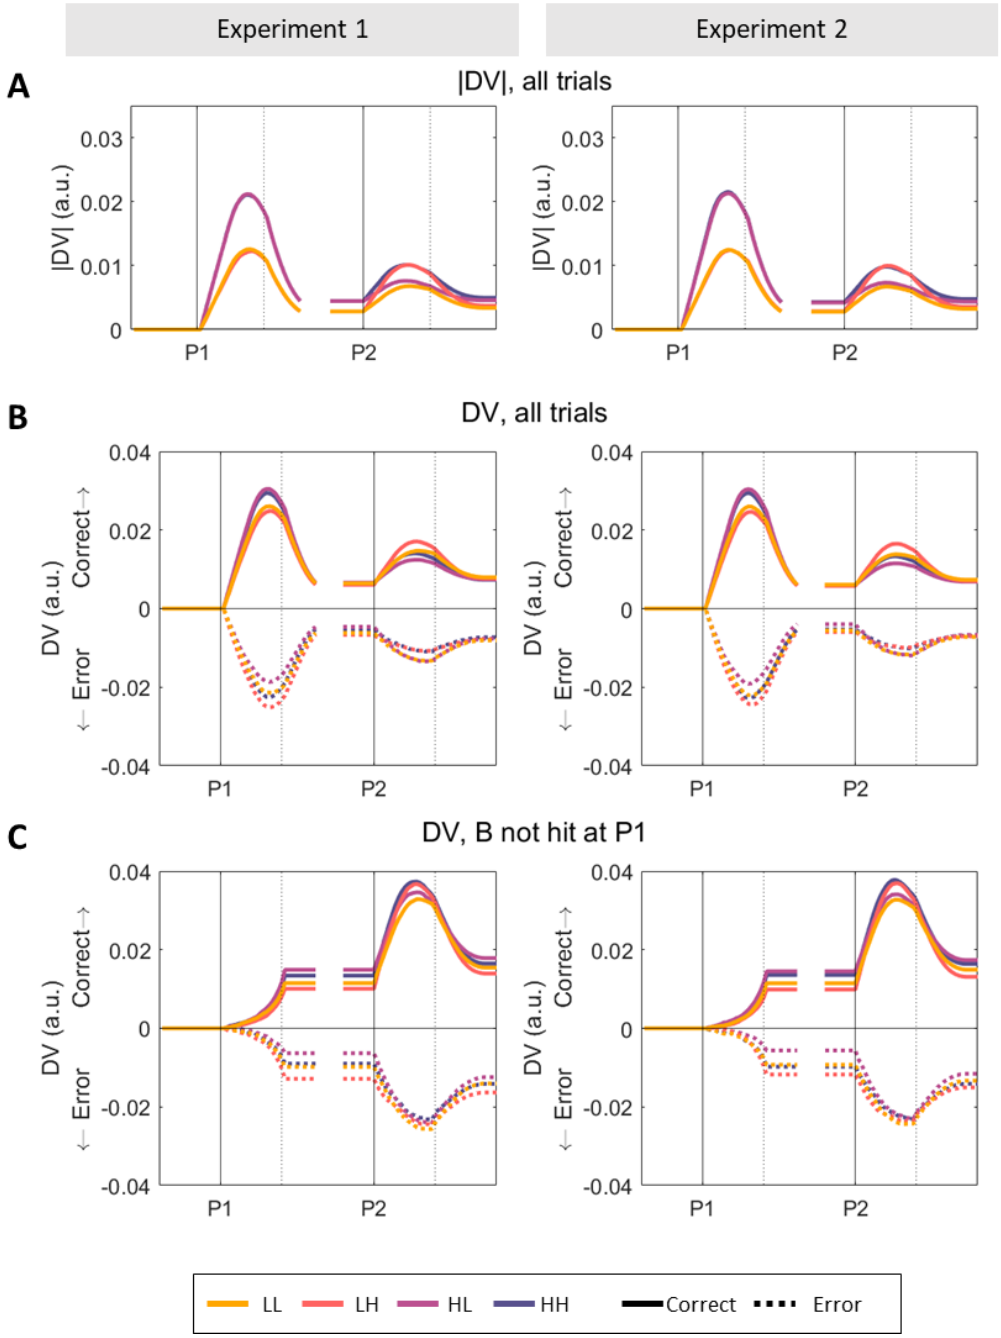

Figure S9. Simulated decision variable (DV) traces for various conditions, and in both experiments. The DV is assumed to go back to zero within 100ms of hitting a bound, but to remain stable throughout the gap otherwise. **A.** Absolute value of the simulated decision variable across all trials, pooled across accuracy, and sorted by condition. **B.** Signed decision variable across all trials, sorted by accuracy and condition. **C.** Signed decision variable in the subset of trials where no bound was hit following the first pulse. Note the difference in the state of the DV before P2 in correct (solid line) vs. error (dashed line) trials in the LH and HL conditions. The mirroring excursion of the DV after P2 explains the predicted interaction of CPP-P2 amplitudes with accuracy (see *Fig. S8*).

		Free parameters				Goodness of fit		
		b	D1	D2	Dvar_P2	AIC	BIC	G^2
Exp. 1	***Model 1***	0.041	0.103	0.187		6.38	20.79	0.38
	Model 2	0.055	0.091	0.181	0.152	8.34	27.55	0.34
Exp 2	***Model 1***	0.040	0.110	0.196		6.79	22.37	0.79
	Model 2	0.039	0.122	0.211	0.237	8.59	29.36	0.59

Figure S10. **Full model comparison**. **A. Model 1 (k = 3, b, d1, d2)**. We initially fit a model with two different drift rates for low (d1) and high (d2) coherence pulses. **B. Model 2 (k = 4: b, d_low, d_high, dvar_p2).** Finally, we tested whether a full model including two drift rates (high/low) and drift rate variability at P2 would provide a better fit to the data. This model resulted in the lowest overall error as measured by G^2, but the additional complexity was not warranted according to AIC/BIC metrics.

Supplementary note 1

Given that temporal expectations have been shown to strongly influence behaviour (Grabenhorst et al., 2019, 2021), and in an exploratory manner, we seeked to directly investigate whether single-trial fluctuations in the neural correlates of evidence accumulation and attention were linked to variability in choice behaviour. To do so, we exploited the trial-by-trial variability in the neural signals.

First, we investigated whether single-trial CPP amplitude related to choice using time-resolved logistic regressions where the weight of P1 and P2 were allowed to be modulated by single-trial CPP amplitude following each pulse (*Eq. 9; Fig. S11A*). In our task, variability in the CPP may be related to trial-by-trial variability in motion energy, as well as random noise. Since the weight on choice of P1 and P2 may trade-off with one another due to a bound being set – i.e. trials where P1 has a strong influence P2 is likely to have a weaker one - we also computed the difference between the two modulatory terms (*[p1*CPP minus p2*CPP]*). This difference index effectively quantifies whether centroparietal responses to either pulse modulate the relative weight on choice of each of the pulses. A positive index indicates a relative increase of P1 weight on choice, whereas a negative index indicates a relative increase of P2 weight on choice.

We found that samples that evoked particularly large centroparietal responses after the first pulse (CPP-P1) for a given level of evidence correlated with a stronger weight on choice for this pulse in both experiments (Exp. 1: *[p1*CPP]*: p = 0.04 *; [p1*CPP-p2*CPP]*: p = 0.038; *Exp. 2: [p1*CPP]*: p =0.008*; Fig. S11A*). Further, variability in evoked response after the second pulse (CPP-P2) also modulated its relative weight on choices, but only in the second experiment (Exp. 2: *[p1*CPP-p2*CPP]*: p = 0.042; *Fig. S11A*).

Next, we investigated whether variability in within-trial attentional fluctuations modulated weight on choice using an analogous regression approach (*Eq. 10, Fig. S11B*). We found that attentional fluctuations only modulated the weight on choice of the second pulse in Experiment 2 (Exp. 2: *[p2*CPP]: p* = 0.032; *[p1*CPP-p2*CPP]*: p = 0.038; *Fig. S11B*). This effect indicates that in trials where participants showed higher occipital alpha desynchronisation (i.e. paid particularly high attention) after the second pulse, they also showed stronger weights of that pulse on choice.

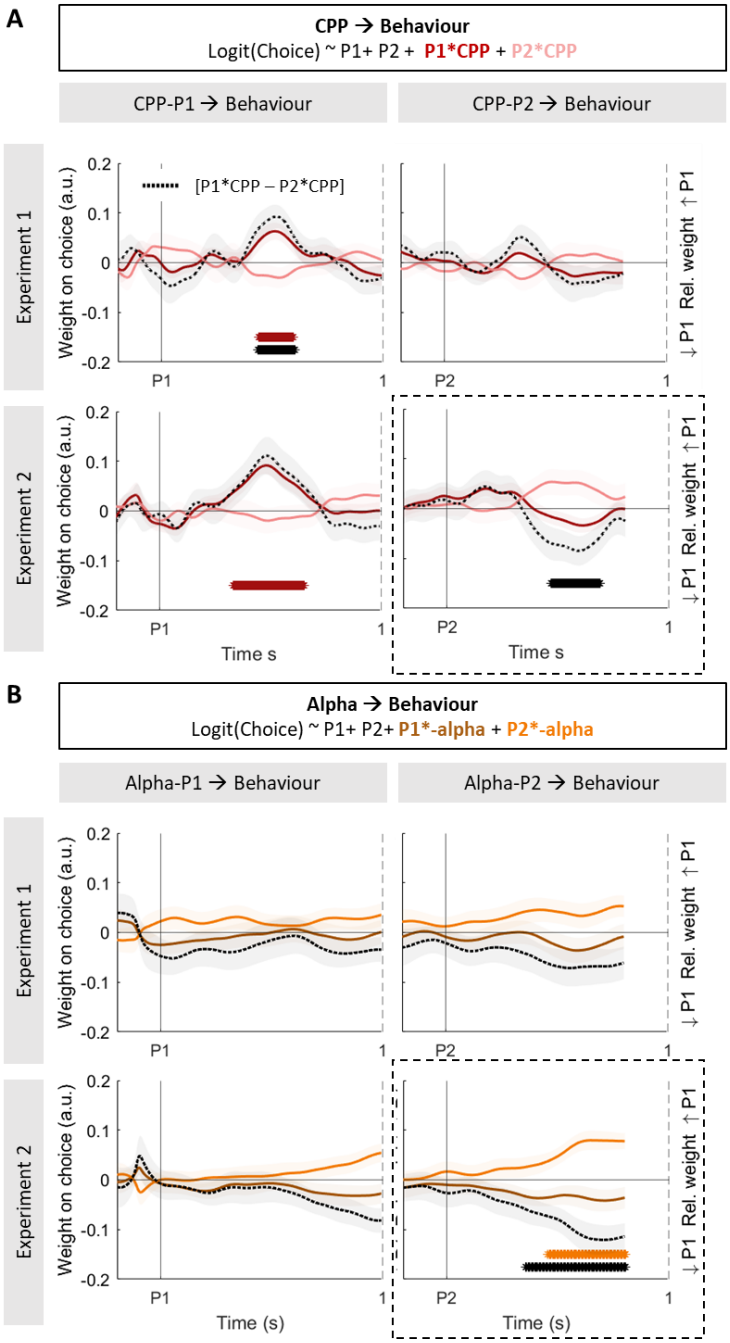

Figure S11. **Weight on choice modulations by single-trial fluctuations in neural evidence accumulation and attention markers. A.** Standardized regression coefficients (±SEM) aligned to pulse 1 (*left*) or pulse 2 (*right*), averaged across all delay conditions. The magnitude of CPP-P1 modulated the impact of the first pulse on choices in both experiments while the magnitude of CPP-P2 only modulated the relative weight of P2 in Experiment 2. Black dashed lines correspond to the difference in regression coefficients for the P1*CPP – P2*CPP modulations. **B.** Standardized regression coefficients (±SEM) aligned to pulse 1 (*left*) or pulse 2 (*right*), averaged across all delay conditions. The magnitude of occipital phasic alpha responses did not modulate the weight on choice of P1, and that after P2 modulated its weight on choice in Experiment 2 only. Black dashed lines correspond to the difference in regression

coefficients for the P1*alpha – P2*alpha modulations, which effectively indicate a change in the relative weight for P1 and P2. *Asterisks along the bottom of the panels indicate significant ($p<0.05$) effects in the test period [0.2-0.8s post event].