Centroparietal activity mirrors the decision variable when tracking biased and time-varying sensory evidence

Carmen Kohl [1,2,*], Laure Spieser [1,2], Bettina Forster [2], Sven Bestmann [3], Kielan Yarrow [2]

[1] These authors contributed equally

[2] Department of Psychology, Cognitive Neuroscience Research Unit,

City, University of London, UK

[3] Sobell Department of Motor Neuroscience and Movement Disorders, UCL Institute of

Neurology, University College London, UK

* Correspondence: Carmen Kohl

City, University of London

Northampton Square

London

EC1V 0HB

+44 (0)20 7040 8530

carmen.kohl@city.ac.uk

Decision-making is a fundamental human activity requiring explanation at the neurocognitive level. Current theoretical frameworks assume that, during sensory-based decision-making, the stimulus is sampled sequentially. The resulting evidence is accumulated over time as a decision variable until a threshold is reached and a response is initiated. Several neural signals, including the centroparietal positivity (CPP) measured from the human electroencephalogram (EEG), appear to display the accumulation-to-bound profile associated with the decision variable. Here, we evaluate the putative computational role of the CPP as a model-derived accumulation-to-bound signal, focussing on point-by-point correspondence between model predictions and data in order to go beyond simple summary measures like average slope. In two experiments, we explored the CPP under two manipulations (namely non-stationary evidence and probabilistic decision biases) that complement one another by targeting the shape and amplitude of accumulation respectively. We fit sequential sampling models to the behavioural data, and used the resulting parameters to simulate the decision variable, before directly comparing the simulated profile to the CPP waveform. In both experiments, model predictions deviated from our naïve expectations, yet showed similarities with the neurodynamic data, illustrating the importance of a formal modelling approach. The CPP appears to arise from brain processes that implement a decision variable (as formalised in sequential-sampling models) and may therefore inform our understanding of decision-making at both the representational and implementational levels of analysis, but at this point it is uncertain whether a single model can explain how the CPP varies across different kinds of task manipulation.

Key words: decision-making, centroparietal positivity, decision bias, non-stationary evidence, accumulator model

Disclosures and Acknowledgements

This work was supported by a Leverhulme Trust Research Project Grant (RPG-2014188).

K.Y., B.F. and S.B conceived the research programme. C.K., L.S. and K.Y. designed the experiment; C.K. and L.S. conducted the research and analysed the data; C.K. drafted the paper, which all authors critically revised and approved.

Declarations of interest: none.

1. General Introduction

Both mathematical modelling of cognitive processes and the analysis of neural and behavioural data have generated important insights about human cognition. Recently, the importance of combining these approaches has become increasingly apparent. This triangulation of methods (sometimes referred to as model-based cognitive neuroscience; Forstmann, Wagenmakers, Eichele, Brown, & Serences, 2011) provides several obvious advantages over traditional approaches, as neural data can inform mathematical models, while models can in turn break complex cognitive processes into separate mechanisms, which are easier to test using neural data (Turner, Rodriguez, Norcia, McClure, & Steyvers, 2016).

A variety of approaches have now been suggested to combine cognitive neuroscience and mathematical modelling (Forstmann, Ratcliff, & Wagenmakers, 2016; van Ravenzwaaij, Provost, & Brown, 2017). One field in which model-based cognitive neuroscience has been particularly fruitful is the study of perceptual decision-making (e.g. Mulder, van Maanen, & Forstmann, 2014). Perceptual decisions, in which we quickly categorise sensory stimuli, directly trigger some of our most basic but essential behaviour, and also provide a building block towards higher cognition. Such decisions can be described by sequential sampling models, a group of computational models which assume that to make a decision, we accumulate sensory evidence over time until a decision threshold is reached, at which point we typically initiate the corresponding motor response (Brown & Heathcote, 2008; Ratcliff & McKoon, 2008; Usher & McClelland, 2001).

Importantly, although these models were developed to explain behavioural data and have done so successfully in a large variety of paradigms (Huk & Shadlen, 2005; Milosavljevic et al., 2010; Ratcliff, 2002; Ratcliff, Thapar, College & Mckoon, 1992), they have been further validated by electrophysiological recordings in non-human primates, as several studies have reported accumulation-like neuronal activity while monkeys perform perceptual-decision tasks (e.g. Hanes & Schall, 1996; Shadlen & Newsome, 1996; for a review see Schall, 2002; Gold & Shadlen, 2007; Hanks & Summerfield, 2017). This connection between models and neural data has since been successfully used to directly compare electrophysiological signals with predictions made by mathematical models (e.g. Hanks, Kiani, & Shadlen, 2014; Purcell et al., 2010; Purcell, Schall, Logan, & Palmeri, 2012), and provided important insights into decision-making processes. For example, by analysing firing rates of frontal eye field neurons, Purcell and colleagues (2010) were able to evaluate different cognitive models, thereby highlighting the potential role of neural data as a model selection tool.

The study of neural substrates of the decision variable (i.e. the decision-related accumulation profile) in the human brain, on the other hand, has been advancing more slowly. One method which is commonly used to study decision-making within model-based cognitive neuroscience is functional magnetic resonance imaging (fMRI). In this field, brain activity is analysed in reference to specific model parameters, which has led to the association of different brain regions with specific sub-processes of decision making (e.g. Forstmann et al., 2010; Heekeren, Marrett, Bandettini, & Ungerleider, 2004; for a review, see Mulder et al., 2014).

In order to track the decision variable in the human brain, however, electroencephalography (EEG) or magnetoencephalography (MEG, which produces comparable data) are commonly used, due to their greater temporal resolution. A variety of different signals have been proposed to be decision-related, ranging from event-related potentials (ERPs; Philiastides, Heekeren, & Sajda, 2014; Philiastides et al., 2006; Philiastides & Sajda, 2006; Pisauro, Fouragnan, Retzler, & Philiastides, 2017; Ratcliff et al., 2009) to changes in theta-band power (van Vugt et al., 2012), and motor-related lateralised desynchronisation in beta power (Donner, Siegel, Fries, & Engel, 2009; Meindertsma, Kloosterman, Nolte, Engel, & Donner, 2017; Siegel, Engel, & Donner, 2011).

A particularly promising approach was introduced by O'Connell, Dockree, and Kelly (2012). In a series of experiments, they identified the centroparietal positivity (CPP), an ERP component which shows several key properties of the decision variable. It displays a build-up over the course of the decision, reflecting the integration of sensory evidence, and its crossing of a stereotyped level was shown to predict reaction time (RT; Kelly & O'Connell, 2013; O'Connell et al., 2012). Importantly, the CPP was shown to be independent of sensory and motor signals, as it was fully dissociable from both steady-state visual evoked responses, which provide a readout of sensory input, and contralateral beta power, which reflects motor activation. Independence from motor signals was later confirmed in a study which directly compared the CPP to motor-related beta power, and showed that while both signals build up over the course of the decision, the CPP drops back to baseline levels after a given threshold is reached, while beta activity persisted until a delayed response (Twomey, Murphy, Kelly, & O'Connell, 2016).

Interestingly, the CPP was also observed in an auditory decision-making task, highlighting its putative role as a supramodal decision signal (O'Connell et al., 2012). Following their initial series of experiments, Kelly and O'Connell (2013) provided further evidence supporting the role of the CPP as a decision variable by exploring the CPP in a perceptual decision-making task with different levels of difficulty. This manipulation is known to affect the slope at which sensory evidence is accumulated, with easier stimuli leading to a steeper evidence accumulation rate. This was confirmed in Kelly and O'Connell's study based on parameter estimates derived from the Diffusion model (Ratcliff & McKoon, 2008). The CPP build-up slope varied according to task difficulty level, qualitatively mirroring model predictions regarding accumulation rate. Hence, experimental evidence from previous studies consistently indicates that summary statistics describing the CPP (such as average slope over some arbitrary time window) correspond with the equivalent intuited or abstracted characteristics of a decision variable.

Identifying the neural substrates of human perceptual decisions is an important goal, because a compelling explanation of behaviour should marry computational plausibility with biological reality (Krakauer, Ghazanfar, Gomez-Marin, MacIver, & Poeppel, 2017; Marr, 2010). To move forward, we must go beyond a broad-brush equivalence between brain signals and model predictions, and show that the quantitative precision of sequential sampling models extends to both behaviour and brain dynamics. Although the CPP appears to be a serious candidate for bridging this divide, few studies have formally compared CPP profiles with the decision variable exactly as predicted by sequential sampling models. Building on Kelly and

O’Connell’s approach, Twomey et al. (2015) added a critical step to their analysis to allow for a direct comparison between the model decision variable and the CPP. After fitting the Diffusion model, the resulting parameters were used to simulate the mean level of evidence accumulation across time predicted by the model. The simulated accumulation profile and the CPP were in close agreement. This finding is important, as it goes beyond comparing summary measures derived from a potential neural substrate of decision-making against a set of abstract characteristics derived from intuitions about model behaviour, and instead allows for a direct comparison of the entire accumulation profile. Indeed, with more complex sequential sampling models (e.g. those incorporating inhibition or leakage; Usher & McClelland, 2001) it becomes virtually impossible to intuit how accumulation profiles may change as a function of different experimental manipulations, making detailed modelling essential (Purcell & Palmeri, 2017).

The current study fulfils this brief, going beyond previous work testing the role of the CPP as a decision variable through formal implementation of sequential sampling models. As outlined above, the CPP has only been tested in the context of a limited number of manipulations (O'Connell et al., 2012; Kelly & O’Connell, 2013), and until recently only the impact of decision difficulty has been compared to simulations based on behaviourally constrained sequential sampling models (Twomey et al., 2015). Similar analyses have since been applied to investigate the speed-accuracy trade-off (Spieser et al., 2018) and under combined conditions of extreme time pressure and value-based bias (Afacan-Seref, Steinemann, Blangero, & Kelly, 2018) but comparisons with precise model predictions remain scarce. Hence here we compared the CPP profile to exact model predictions in two separate EEG

experiments. These experiments tested both probabilistic decision biases, which, to our knowledge, have not been previously assessed using the CPP, and non-stationary evidence profiles, which we believe have not previously been examined for the CPP under conditions of speeded choice. In line with previous behavioural work (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012; Spaniol, Voss, Bowen, & Grady, 2011; Summerfield & Koechlin, 2010; Voss, Nagler, & Lerche, 2013), our estimation of model parameters revealed that decision bias affects the amount of evidence required to attain response threshold, while non-stationary evidence affects the detailed time-course of evidence accumulation. We then used the estimated parameter values to simulate the accumulation profiles as predicted by the models and compared them to the recorded CPP.

We chose two types of race accumulator models (Brown & Heathcote, 2008; Heathcote & Love, 2012) to account for our behavioural data, namely, the leaky competing accumulator model (LCA; Usher & McClelland, 2001), suggested to be one of the most neurophysiologically plausible sequential sampling models, and a simplified independent race accumulator model. Contrary to random walk models such as the Diffusion model, in which evidence is integrated in a single accumulator (Smith & Ratcliff, 2004), and which are motivated more by mathematical optimality than neurobiological plausibility (Ratcliff et al., 2016; Usher & McClelland, 2001), what we here refer to as 'race accumulator models' assume that evidence for each response alternative is integrated in separate accumulators, which race to reach a common threshold. Assuming that processes similar to these occur in the brain, with each accumulator being associated with a neural population, and given the nature of EEG, which records the sum of all underlying electrical activity from the scalp, we

propose that the CPP should be best predicted by the summed activity of both accumulators in a two-choice task. Across experiments varying the two core characteristics of accumulation-to-bound activity, namely the shape of accumulation build-up and the extent of baseline-to-bound distance, our results show that CPP dynamics can indeed closely match time-varying predictions derived under a sequential-sampling modelling framework, but that this match partly reflects the flexibility we enjoyed as a result of having several candidate models available.

2. Experiment 1: Non-stationary Evidence

Most research in the field of perceptual decision-making has focused on binary choices with stationary evidence, where information remains virtually unchanged in quality and intensity throughout the decision-making process (Gold & Shadlen, 2000; Kelly & O'Connell, 2013; Ratcliff & McKoon, 2008; Ratcliff et al., 2010). In everyday life, however, decisions typically occur in a dynamic environment, in which sensory evidence is continuously changing, and several studies have drawn attention to the fact that comprehensive models of decision-making have to be able to account for decisions with non-stationary evidence. Researchers have hence started to use decisions in response to non-stationary evidence in order to distinguish between different sequential sampling models (Bronfman, Brezis, & Usher, 2016; Nunes & Gurney, 2016; Tsetsos, Gao, McClelland, & Usher, 2012; Tsetsos, Usher, & McClelland, 2011; Zhou, Wong-Lin, & Philip, 2009), which often offer indistinguishable accounts of data from more traditional decision-making paradigms (Brown & Heathcote, 2008; Ratcliff & Smith, 2004; Teodorescu & Usher, 2013; Tsetsos et al., 2012).

Tsetsos et al. (2011, 2012), for example, conducted a series of experiments, using a paradigm in which the evidence for a given alternative changed dynamically throughout a trial to compare race accumulator (Brown & Heathcote, 2008; Usher & McClelland, 2001) and random-walk models (Ratcliff & McKoon, 2008). They found that the race accumulator model gave a better description of the data (Tsetsos et al., 2011), and was able to account for various subtleties, including a primacy effect which showed that changes in evidence had a larger impact on decisions early on in the decision-making process (Tsetsos et al., 2012). Recently, Holmes, Trueblood, and Heathcote (2016) showed that a simplified race accumulator model labelled 'piecewise LBA' could provide a good account of participants' behaviour. In that study, participants were asked to discriminate between left and right motion in a random dot motion task, in which, halfway through the decision-making process, the motion direction switched. The best-fitting race model parameters confirmed that accumulation rates were affected by the motion switch. Interestingly, while the switch led to motion in the opposite direction but equal in magnitude, estimated changes of accumulation rates were not symmetrical between the two accumulators, indicating a difference in discrimination after the switch. Incorporating a delay between the switch in evidence and the resulting change in accumulation rates was shown to improve model fit, revealing that some time is necessary to take a modification of evidence into account.

It is clear that dynamically changing evidence also has implications for any neural signal posited to reflect the decision-related accumulation of evidence. This was observed for instance in the firing rate of lateral intraparietal (LIP) neurons in non-

human primates. Huk and Shadlen (2005) demonstrated that additional positive/negative motion pulses during a random dot motion task had persistent effects on LIP activity, which increased/decreased for several hundreds of milliseconds. In humans, O'Connell et al. (2012) explored the impact of changing evidence on the CPP and motor-related beta band power. In a detection task in which stimuli gradually decreased in contrast, the CPP (and, to a lesser extent, beta power) was shown to plateau for several hundreds of milliseconds when the gradual contrast decrease was interrupted by a 450 ms increase towards the baseline. In this study, however, no comparisons were made between a simulated accumulation profile and the recorded CPP waveform.

Here, we instead utilise a choice RT task and provide detailed modelling/simulation. Participants performed a random dot motion task which required them to discriminate between motion to the left or to the right. In one third of the trials, dot motion remained unchanged throughout the trial ('continuous' condition), while in the rest of the trials, it was interrupted for a 200 ms period. In these interrupted trials, dot motion was replaced by either coherent motion in the opposite direction, before continuing in the original direction ('reverse' condition), or by random motion without any directional evidence ('random' condition; cf. Tsetsos et al., 2012). These changes in motion should affect the build-up of the accumulation profile, and be visible in any neural signal reflecting the decision variable. While we assumed that the decision variable will 'plateau' during the coherent motion interruption in the 'random' condition, predictions regarding the impact of the reversal of evidence are less clear, and are likely to depend more on the specifications of the model, such as the presence or absence of reciprocal inhibition. To determine exactly how a signal

reflecting the decision variable is affected, we simulated accumulation profiles predicted by sequential sampling models. Importantly, in order to use model specifications best resembling the underlying decision processes, we tested several models and selected the one providing the best fit to our behavioural data. We then directly compared the selected model's profiles to CPP waveforms. In so doing, we confirmed the impact of time-varying evidence on the CPP profile and showed that it corresponds closely to the modulations of evidence accumulation predicted by a leaky competing accumulator model.

2.1. Methods

2.1.1. Participants

In line with commonly reported sample sizes in the CPP literature (e.g. Kelly & O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015), a total of 21 participants (eight males) were recruited. To ensure a reasonable and distinguishable task performance at two different difficulty levels, each participant completed a staircase procedure to establish the appropriate level of difficulty (see below, 2.1.2). In line with criteria defined prior to data collection, one participant was excluded from the experiment as the calibrated level of coherence exceeded 98% for the 'easy' condition, leading to a sample of 20 participants (seven males) with a mean age of 27.55 (*SD* = 8.83). The experiment was approved by the City, University of London Psychology Department Ethics Committee.

2.1.2. Stimuli and Procedure

Participants were asked to complete a random dot motion task. The task was written in Matlab (The Mathworks, Natick, U.S.A.), making use of Psychtoolbox functions (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997). In this task, an array of white dots was presented on a black screen. A proportion of dots moved coherently either to the left or to the right, while the rest of the dots moved in random directions. Participants were instructed to indicate the perceived motion direction by pressing a button in their right/left hand for movement to the right/left. For this, digital response buttons interfaced via a 16 bit A/D card (National Instruments X-series PCIe-6323, sample rate 100,000 Hz) were held between the thumb and index finger of each hand. Participants were seated 100 cm away from a cathode ray tube screen (size: 41 x 30 cm), operating at a refresh rate of 85 Hz and with a resolution of 1240 x 786. A total of 300 dots, 0.04 x 0.04 degrees visual angle (dva) in size, were presented within a 5 dva circular aperture. During random motion, on each frame, each dot was displaced into a random direction. During coherent motion on the other hand, only a subset of dots followed this random motion, while the remaining dots (defined by the level of coherence, see below) moved uniformly either to the left or to the right, depending on the trial. Both random and coherent dot movements occurred at a speed of 3.3 dva per second. Additional to this motion, all dots were relocated to a random position on the array every five frames. This process was added so that participants could not determine the direction of the motion by following one specific dot, instead having to consider the entire motion array.

Each trial began with a central fixation cross (size: 0.33 x 0.33 dva) for 500 ms (plus a jitter of up to 1000 ms, drawn from a uniform distribution), followed by a period of

random motion (1000 ms plus a jitter of up to 1500 ms, drawn from a gamma distribution with shape parameter 1 and scaling parameter 150[1]). Since the onset of moving dots on the screen is likely to produce a visual evoked potential which would interfere with the recording of the CPP, this period of random motion was introduced to allow for the evoked potential to occur before the onset of the decision-making process. The random motion was followed by the onset of coherent motion (left/right) which continued for up to 2000 ms or until the response (see Figure 1 a).

Participants completed a minimum of 100 practice trials at high levels of coherence (i.e. > 80% of dots moving in one direction) to familiarise themselves with the task. In order to calibrate suitable levels of difficulty for 'easy' and 'hard' trials for each participant individually, a further 100 trials were completed in which the QUEST (Watson & Pelli, 1983) staircase procedure, implemented in Psychtoolbox, estimated the coherence level at which each participant responded correctly in 80% of trials. This coherence level was then used for the 'hard' condition. The 'easy' coherence level was set as 150% of the 'hard' coherence level. Participants had 1300 ms to respond, and no feedback was provided during staircase trials. Overall, the appropriate difficulty levels estimated for the remaining participants resulted in a mean of 27.70% (*SD* = 14.74) coherence for 'hard', and 40.15% (*SD* = 22.15) for 'easy' trials.

After the staircase procedure, participants were asked to complete a further 100 practice trials which included all conditions of the main experiment, including the

[1] A gamma distributed foreperiod with a shape parameter of 1 was chosen as it is associated with a uniform hazard function (Luce, 1986).

different difficulties and evidence interruptions (see below). Like in the main task (described next) participants now had 2000 ms to respond. During this training, participants were given feedback in the form of their mean accuracy and RT every 10 trials. In order to introduce a moderate speed pressure, participants were instructed to aim for a mean accuracy of at least 80% and a mean RT of less than 1000 ms throughout the task.

During the experiment, we manipulated the continuity of the evidence by introducing three motion conditions, in addition to the manipulation of difficulty (see Figure 1 a). One third of the trials, like the practice and staircase trials, were 'continuous' trials, i.e. the coherent motion began after a period of random motion and remained unchanged throughout the trial. In the 'random' condition, the coherent motion was interrupted 200 ms after motion onset and replaced by a 200 ms period of random motion (i.e., 0% coherence level), before being reinstated. Similarly, in the 'reverse' condition, the coherent motion was interrupted for the same time period, but replaced by coherent motion in the opposite direction (see Figure 1 a). Informal questioning of participants indicated that these interruptions were not perceived consciously. During the main task, the interruption condition ('continuous', 'random', or 'reverse'), motion direction (left or right) and coherence level ('easy' or 'hard') varied randomly from trial to trial in an equiprobable factorial design. Each participant completed 16 blocks of 60 trials. After each block, participants were given feedback in the form of their mean accuracy and RT. No feedback was provided for individual trials.

2.1.3. EEG Recording and Pre-processing

During the task, we recorded participants' EEG using 64 active electrodes, placed equidistantly on the scalp (EasyCap, M10 Montage) and referenced to the right mastoid. Data were recorded through a BrainAmp amplifier (BrainProducts, sampling rate: 1000 Hz).

The data were pre-processed in Matlab (The Mathworks, Natick, U.S.A.), using custom scripts and implementing functions from the EEGLAB toolbox (Delorme & Makeig, 2004). Data were re-referenced to the average reference and band-pass filtered from 0.1 (low cut-off) to 45 Hz (high cut-off), using a Hamming windowed finite impulse response filter. We then visually inspected the data to remove noisy channels and reject large artifacts, before applying independent component analysis to correct for eye blinks. Afterwards, the data were visually inspected a second time in order to manually remove any remaining noise. Lastly, we used spherical spline interpolation to reconstruct any channels that were previously removed. In line with the procedures used in previous CPP studies (Kelly & O'Connell, 2013; O'Connell et al., 2012), the data were converted to current source density (CSD) estimates to increase spatial selectivity. The CSD transformation was applied using the CSD toolbox, which uses a spherical spline algorithm, with the spline interpolation constant *m* set to its default value (*m* = 4; Kayser & Tenke, 2006).

2.1.3.1. ERP Analysis

For the ERP analysis, we extracted both stimulus-locked (-200 to 2000 ms, relative to coherent motion onset) and response-locked (-1000 to 100 ms, relative to the button press) epochs. All epochs were baseline corrected to the average over a 200

ms period preceding the coherent motion onset. As only medial electrodes were analysed, and initial observations revealed no difference depending on the direction of motion, we collapsed over ‘left’ and ‘right’ trials. Further, since high overall accuracy scores led to insufficient numbers of error trials to generate reliable ERP signals, error trials were excluded.

The appropriate electrode to generate the CPP waveform was chosen individually, by visually inspecting each participant’s averaged ERP topography to identify the centroparietal region of maximum amplitude (chosen electrodes: 1, 5, or 14, roughly equivalent to electrodes Cz, CPz, and Pz in the 10-20 system; see Figure 1 d). The activity in the selected electrode was averaged for each condition and for stimulus and response-locked signals separately.

2.1.4. Statistical Analysis

Differences between conditions for behavioural data were inferred using ANOVAs and generalized linear mixed models (GLMMs) with logistic link functions, for RTs and error rates respectively. GLMMs were chosen for the analysis of accuracy data since the non-normal distribution of such data will, at a theoretical level, always violate the assumptions of ANOVA (Jaeger, 2008). They were implemented using the Matlab fitglme command; all effects of interest (e.g. ‘Difficulty’, ‘Interruption’, and their interaction) were clustered within participants and included as random effects in the model specifications (e.g. Wilkinson notation: `Accuracy ~ 1 +`

`Interruption*Difficulty + (1+Interruption*Difficulty | Participant)`.[2]

In order to test the effects of the difficulty and interruption manipulations on the ERP, we explored both the slopes and the amplitudes of the waveforms. First, we compared the slopes between the different conditions by fitting a straight line to the CPP for each participant and each condition and measuring its slope. The resulting slopes were then compared in an 'Interruption' ('continuous', 'random', 'reverse') x 'Difficulty' ('easy', 'hard') repeated-measures ANOVA.

We compared slopes during two different time intervals in the stimulus-locked data: an early interval between 100 and 300 ms and a late interval between 300 and 500 ms relative to the onset of coherent motion. Given the interruption interval of 200 to 400 ms and the assumption of a small lag between stimulus presentation and accumulation (typically observed in the CPP, see Kelly & O'Connell, 2013; Spieser et al., 2018), we assume that the early interval reflects accumulation mainly before the interruption and the late interval reflects accumulation mainly during the interruption. However, since these intervals were primarily chosen based on visual inspection, and Kelly and O'Connell (2013) suggested a longer 200 ms delay between the evidence and its visible effect on the CPP waveform, we also repeated the analysis, defining the interruption interval as a 400-600 ms time window.

[2] This represents the 'maximal' random effects structure (Barr, Levy, Scheepers, & Tily, 2014) which makes the model as equivalent as possible to a traditional repeated-measures ANOVA, whist properly respecting the nature of the data.

Additionally, we analysed the impact of difficulty and interruption on the amplitude of the waveform. Between 0 and 1000 ms in the stimulus-locked data, and between -1000 to 0 ms in the response-locked data, we compared conditions using an 'Interruption' ('continuous', 'random', 'reverse') x 'Difficulty' ('easy', 'hard') ANOVA at each time point. The results were controlled for multiple comparisons using the false discovery rate (FDR) approach (Benjamini & Hochberg, 1995)[3].

2.1.5. Model Fit

To model the behavioural data, we used two sequential sampling models. Firstly, the independent race accumulator model which is, at least conceptually, one of the simplest sequential sampling models (Brown & Heathcote, 2008; Usher & McClelland, 2001). In this model, evidence for each response alternative is integrated in independent accumulators which race towards the decision threshold. At each time point, a given accumulator *i* accumulates the input evidence I_i supporting response alternative 'i', as well as noise *N*, drawn from a normal distribution with mean 0 and standard deviation σ, so that the quantity accumulated at each time point is described by:

$$dx_i \propto I_i + N(0, \sigma^2) \quad (1)$$

[3] In this procedure, the uncorrected *p*-values are sorted from lowest to highest (p_i refers to the *i*th lowest value out of *m* total *p*-values). The largest *i* for which $p_i < \left(\frac{i}{m}\right) \propto$ is identified and all *p*-values associated with *i*s smaller or equal to the identified *i* are considered significant.

The strength of input I_i depends on the mean accumulation rate v_i, which reflects the quality of evidence. To remain physiologically plausible, the accumulation process is restricted to positive values at each time step[4]:

$$x_i(t+1) = \max(0, x_i(t) + dx_i) \quad (2)$$

Once either of the accumulators reaches the threshold A, the corresponding response (here response 'i') is initiated. Potential variations between trials' starting states are introduced by varying accumulation starting point, which is drawn for each accumulator and each trial from a uniform distribution between 0 and S_z. The time taken to reach the threshold, in addition to a non-decision time which represents any time taken for sensory and motor processes before and after the accumulation process respectively, defines the modelled RT. The non-decision time is drawn from a uniform distribution with width S_{Ter}, centred on T_{er}.

In addition to the independent race accumulator model, we also used the more physiologically plausible LCA model (Usher & McClelland, 2001) which introduces interactions within and between accumulators. In this model, like the simpler independent race model described above, evidence for each response alternative is accumulated in separate accumulators which race towards response threshold A. Additionally, the LCA includes a leakage parameter k as well as a parameter β for

[4] Strictly, for physiological plausibility, the quantity accumulated should always be positive (as neurons cannot have negative firing rates) and also generally begin at a positive baseline (given spontaneous neural activity). Many of the models tested in this paper do begin at positive values, although this is not always the case for our LCA models (in line with conventional implementations of this model).

mutual inhibition between accumulators. Thus, in a binary decision involving the accumulators *i* and *j*, the change in activation in each accumulator is given by[5]:

$$dx_i \propto I_i - kx_i - \beta x_j + N(0, \sigma^2) \quad (3)$$

$$dx_j \propto I_j - kx_j - \beta x_i + N(0, \sigma^2)$$

Where *I* is the input into the accumulator and $N(0,\sigma^2)$ is noise drawn from a normal distribution with a mean of 0 and a standard deviation of σ. Again, the accumulation process is limited to positive numbers:

$$x_i(t+1) = \max(0, x_i(t) + dx_i) \quad (4)$$

$$x_j(t+1) = \max(0, x_j(t) + dx_j)$$

A decision is made when either of the accumulators reaches the threshold *A*, and the RT is made up of the time required to reach the threshold, and a non-decision time drawn from a uniform distribution centred on T_{er} with width S_{Ter}, which accounts for sensory and motor processes before and after the accumulation process.

To determine which model provided the best fit to our behavioural data, four independent race and four LCA models were tested. In all models, the response threshold *A* was chosen as the scaling parameter and fixed to 1. Apart from the periods of motion interruption, evidence supporting the correct response alternative was accumulated in the 'correct' accumulator at a mean accumulation rate $v_{correct}$, while evidence for the incorrect response was integrated in the 'incorrect'

[5] In our code, these equations were implemented as:

$$dx_i = \left(v_i - kx_{i,t-1} - \beta x_{j,t-1}\right)dt + N(0, \sigma^2)\sqrt{dt}$$

$$dx_j = \left(v_j - kx_{j,t-1} - \beta x_{i,t-1}\right)dt + N(0, \sigma^2)\sqrt{dt}$$

With dt = 0.01s. Hence a correction (by a factor of dt) may be required for comparison with parameters reported in some other papers based on finite difference equations.

accumulator at a mean rate $v_{incorrect}$. All models implemented a change in accumulation rates during the interruption interval (from 200 to 400 ms relative to the decision onset), but each assumed different mechanisms (Holmes et al., 2016), as described below. For consistency, 'correct' and 'incorrect' accumulator labels remained constant throughout each trial, such that, during the evidence interruption, $v_{correct}$ and $v_{incorrect}$ still referred to the correct and incorrect responses according to the initial motion direction[6]. Finally, as trial difficulty influences evidence accumulation, accumulation rates were always estimated separately for easy and hard trials.

Model 1 was an independent race model defined by eight parameters, assuming symmetrical changes in accumulation rates during motion interruption. In 'continuous' trials, evidence was accumulated at mean rates $v_{correct}$ and $v_{incorrect}$ throughout the whole trial. In 'random' trials, in which the evidence becomes random during the interruption, we assumed that only noise was accumulated during this period, i.e., $v\text{-}random_{correct} = v\text{-}random_{incorrect} = v_{incorrect}$ from 200 to 400ms after decision onset. Outside of this interval, correct and incorrect rates were set to the initial values $v_{correct}$ and $v_{incorrect}$. In the 'reverse' condition, the evidence changed direction during the interruption interval, but remained at its original strength, which may lead to a reversal of drift rates, i.e., $v\text{-}reverse_{correct} = v_{incorrect}$, $v\text{-}reverse_{incorrect} = v_{correct}$. Again, outside for the interruption interval, evidence was accumulated at mean rates $v_{correct}$ and $v_{incorrect}$. This describes a model with only four accumulation rates ($v_{correct}$ and $v_{incorrect}$, estimated separately for easy and hard decisions), as well

[6] In the 'reverse' condition, evidence during interruption supports the *incorrect* response alternative, and is integrated in the 'incorrect' accumulator.

as the parameters S_z, T_{er}, S_{Ter}, and σ^2 which were fixed between conditions (see Table 1).

Instead of symmetrical changes, model 2 assumed free variation in rates with changing evidence leading to a new set of accumulation rates for the 'random' and 'reverse' intervals. This results in a total of 12 accumulation rates: for each difficulty condition, $v\text{-}continuous_{correct}$, $v\text{-}continuous_{incorrect}$, $v\text{-}random_{correct}$, $v\text{-}random_{incorrect}$, $v\text{-}reverse_{correct}$, $v\text{-}reverse_{incorrect}$. All other parameters ($S_z$, T_{er}, S_{Ter}, σ^2) were fixed between conditions, resulting in a model of 16 free parameters (see Table 1).

Models 3 and 4 were identical to models 1 and 2 respectively, but also included a delay parameter *d* to account for a potential delay between the change in evidence and the change in the decision variable (Holmes et al., 2016). Note that the delay introduced here is different from simple sensorial delay, caught by the encoding part of non-decision time. It instead adds a time lag between the *change* in evidence and accumulation rate modulation to account for potential persistence of accumulation even when evidence has changed.

Finally, Models 5, 6, 7, and 8, were LCA models implementing the same modulations as Models 1, 2, 3, and 4 respectively (see Table 1).

For each participant, trials with RTs faster than 180 ms or slower than 2000 ms (less than 3%) were discarded. RT distributions in each condition were then summarized by five quantiles for correct trials, and by the median RT value for incorrect trials (the median was used due to the low number of incorrect trials in some cases). Best

fitting model parameters were then determined at an individual level. Modelled RTs were simulated based on the equations described above and compared to RT data using Quantile Maximum Probability Estimation (Heathcote et al., 2002). Parameter values were adjusted using a differential evolution algorithm implemented in Matlab (The Mathworks, Natick, U.S.A.; Price et al., 2005).

We compared the goodness of fit of models by calculating the mean Bayesian information criterion (BIC, Schwarz, 1978) as well as the mean Akaike information criterion (AIC; Akaike, 1977). These measures take into account the likelihood of the model, but also penalise a model for the number of parameters used in order to resolve the problem of overfitting. For our data, AIC and BIC were not in agreement regarding the best overall model. We therefore made a (somewhat arbitrary) decision to favour BIC, but to also present AIC in all tables for transparency. The model which best fitted the data according to the BIC measure was then used to generate predictions of the accumulation profile.

In addition, we also performed a recovery study to estimate the accuracy of our fitting procedure. This was done by simulating 20 RT datasets using Model 5 (LCA-symmetric with no delay, i.e., the lowest BIC model, see results). The simulated datasets were constructed as per our empirical individual data with the 3 interruption conditions and 2 difficulty levels. The number of trials also corresponded to empirical data (160 trials per condition, i.e., 960 trials in total). Results of the recovery are presented in Appendix A.

2.1.6. Model Prediction (neurodynamics)

Since EEG recordings reflect the summation of neural activity in a given area, we assumed that, if the CPP is a neural correlate of the decision variable, it represents the sum of all evidence accumulation. Although a binary choice may recruit separate neural populations to accumulate evidence, these neural populations would likely be in close proximity. An ERP component recorded at the scalp over these neural populations measures the summation of electrical activity and therefore most likely the sum of both accumulation processes. In order to compare the model prediction to the CPP, we therefore considered the sum of the correct and incorrect accumulation profiles of correct choices.

Based on the model equations described above, a total of 10,000 accumulation paths (in 10 ms time steps) were computed using individual best-fitting parameters obtained for each condition. To account for sensory processes, accumulation started after a sensory delay (fixed to 50% of non-decision time). Evidence was then accumulated until the response threshold and continued to be accumulated for a short period after the threshold was reached to account for motor processes (50% of non-decision time*;* note that we assume that accumulation continues until the offset of the stimulus, i.e. during the time to reach the threshold plus the time taken to make the motor response and thus stop the stimulus in our paradigm).

To match with EEG processing, the ‘sum of accumulations’ signal was baseline corrected by subtracting the first data point value from the whole trial. Finally, we averaged accumulation signals in each condition, locked to both the estimated onset of the decision process (stimulus-locked) and the response (response-locked). Since

the stimulus-locked signal includes varying time spans of post-decision stages, and we can only speculate about the behaviour of the accumulator after the response, we removed simulated trials from averaging after the response (i.e. after the crossing of the threshold plus 50% non-decision time). Both stimulus and response-locked individual predictions were then averaged across participants, to obtain “grand average” model predictions.

To compare the EEG signal with these model predictions, we recomputed individual stimulus-locked CPPs, by removing trials from the average once they reached the associated RT, and then recomputed the corresponding grand average. EEG signals were then low-pass filtered with a cut-off of 5 Hz for better visualisation, and downsampled to match the 10 ms time steps used in the model predictions. To quantify the similarity between the two signals, we analysed the correlations between the model predictions and the downsampled, but not low-pass filtered EEG data for each difference between conditions (stimulus-locked time-window: 0 – 1000 ms, response-locked time-window: -1000 – 0 ms).

2.2. Results

2.2.1. Behavioural Results

Behavioural data were collapsed over ‘left’ and ‘right’ trials. All trials with very short (< 180 ms) or very long (>= 2000 ms) RTs were excluded from the analysis (2.99% of trials). The remaining data are displayed in Figure 1 c.

As expected, ‘easy’ decisions were faster than ‘hard’ decisions, $F(1, 19) = 134.96$, $p < .001$, $\eta_p^2 = .88$. For the main effect of ‘Interruption’, Mauchley’s test indicated that the assumption of sphericity had been violated, $\chi^2(2) = 18.77$, $p < .001$. We therefore Greenhouse-Geisser corrected the degrees of freedom ($\varepsilon = .61$). There was a significant main effect of ‘Interruption’, $F(1.21, 23.07) = 63.45$, $p < .001$, $\eta_p^2 = .77$. Pairwise comparisons using Fisher’s Least Significant Difference (LSD) revealed that all three levels of ‘Interruption’ were significantly different from each other with 'continuous' trials leading to shorter RTs than 'random' ($p = .001$) and 'reverse' ($p < .001$) trials, and 'random' trials showing shorter RTs than 'reverse' trials ($p = .005$). There was no significant interaction, $F(2, 38) = 2.00$, $p = .15$, $\eta_p^2 = .10$.

Additionally, GLMMs showed that accuracy also differed significantly by ‘Difficulty’, $F(1, 114) = 7.19$, $p = .008$, with ‘easy’ conditions associated with higher accuracy than ‘hard’ conditions. ‘Interruption’ was also a significant predictor, $F(2, 114) = 108.88$, $p < .001$. The ‘Interruption * Difficulty’ interaction was not significant, $F(2, 114) = 2.33$, $p = .10$. In order to explore the differences between all three levels of ‘Interruption’ (‘continuous’, ‘random’, ‘reverse’), we fitted the model a second time, but setting the reference level of ‘Interruption’ to ‘random’, rather than ‘continuous’. We found that both the ‘continuous’ and the ‘random’ conditions were associated with higher accuracy scores than the ‘reverse’ condition ($p < .001$). There was no significant difference between the ‘continuous’ and the ‘random’ conditions ($p = .13$).

2.2.2. ERP Results

The resulting ERPs are displayed in Figure 1 d. The CPP displays a build-up over the course of the decision, which seems disrupted by the interruption of evidence in

relevant conditions.

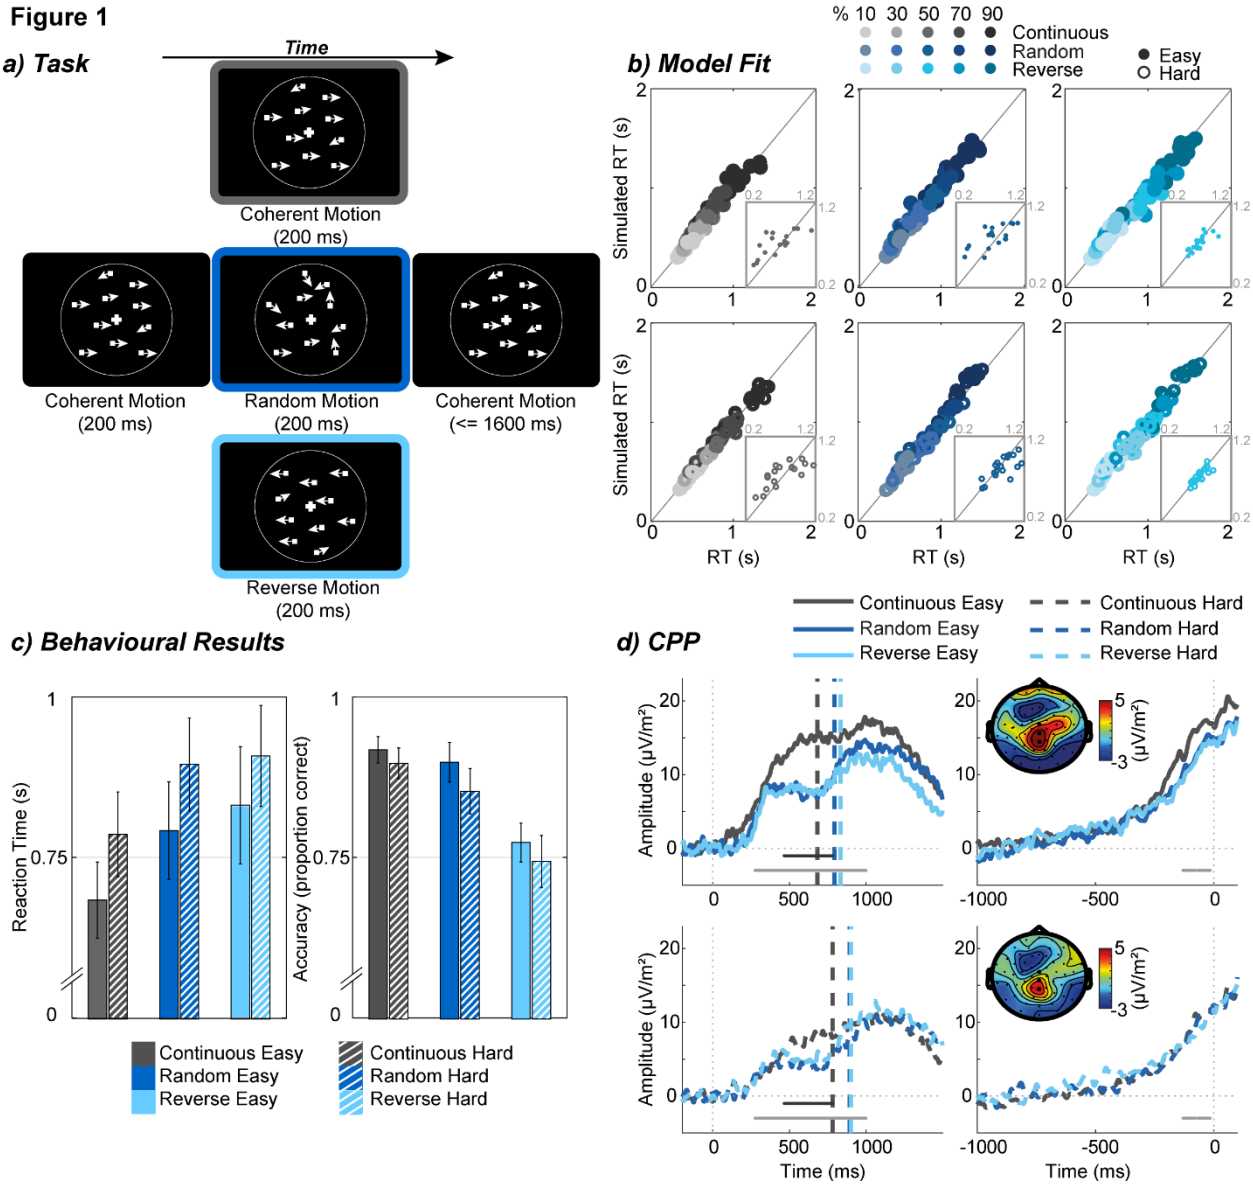

Figure 1: a) Experiment 1 random dot motion task trial procedure: in each trial, coherent motion (here: direction: right; coherence: 70%) was either continuous ('continuous' condition), or was interrupted by either random motion ('random' condition) or coherent motion in the opposite direction ('reverse' condition), before continuing in the original direction. b) Model fit: each participant's quantiles from behavioural data (x-axis) and the LCA model (Model 5) simulations (y-axis) for easy (top, filled circles) and hard (bottom, circle outlines) decisions, as well as continuous (left), random (middle) and reverse (right) conditions. Increasing quantiles (10%, 30%, 50%, 70%, 90%) are represented by increasingly dark colours. Small inserted panels show observed and simulated RT medians for error trials. c) Behavioural results: mean reaction time (left) and accuracy (right) in each condition. Error bars indicate 95% confidence intervals. d) CPP results: Stimulus-locked (left) and response-locked (right) CPP waveforms for easy (top), and hard (bottom) trials. Right panels show topography averaged over the

stimulus-locked 0 to 1000 ms interval. Electrodes used to generate the waveform are highlighted. Vertical dashed lines in the stimulus-locked CPP represent mean RTs per condition. Note that the mean RTs here are computed only from trials which were included to generate the waveform and therefore differ slightly from those displayed in c. Grey dots at the bottom of the waveforms indicate significance based on FDR-controlled comparisons of amplitude: dark grey dots indicate a significant effect of Interruption, while light grey ones indicate a significant effect of Difficulty.

First, we compared the slopes of the ERP occurring in response to evidence accruing before and during the interruption period. In the first interval (100-300 ms), analysis revealed that the slope of the CPP associated with 'easy' waveforms was higher than 'hard' waveforms, $F(1, 19) = 12.93$, $p = .002$, $\eta_p^2 = .40$. There was no main effect of 'Interruption', $F(2, 38) = 1.01$, $p = .38$, $\eta_p^2 = .05$, and no interaction effect ($p = .82$). Conversely, in the second, interruption-driven, interval (300-500 ms), the slope of the CPP was affected by the 'Interruption' condition, $F(2, 38) = 9.52$, $p < .001$, $\eta_p^2 = .33$, but not by 'Difficulty', $F(1, 19) = .19$, $p = .67$, $\eta_p^2 = .01$, with no interaction between the two factors ($p = .39$).[7] Investigating the interruption effect with Fisher's LSD post-hoc tests showed that the slope was significantly higher in the 'continuous' waveform than the 'random' and the 'reverse' waveforms, $t(19) > 3.40$, p

[7] We selected two windows for slope analysis based on the timing of our stimulus (and assumptions about the time course with which information feeds through to decision areas of the brain). This approach is consistent with previous work on the CPP, but incorporates no correction for familywise error, which might raise concerns in the absence of pre-registration for the analysis. For completeness, we attempted an analysis that varied the position of the 200 ms window used to assess slope (in steps of 1 ms) and incorporated an FDR correction for these multiple comparisons. Under this approach, the slope difference associated with difficulty (100-300 ms) remains significant, but the later slope difference associated with interruption condition (300-500 ms) fails to reach significance. However, subsequent FDR-corrected analyses of amplitude provide an alternative source of evidence regarding the impact of the interruption conditions on the CPP.

< .003. No significant difference between the ‘random’ and ‘reverse’ conditions was observed, $t(19) = .76$, $p = .46$. Since the interruption-driven time interval of 300-500 ms was chosen primarily based on visual inspection, we repeated the analysis using a time window which assumes a 200 ms delay between the evidence and its visible effect on the CPP, as suggested by Kelly and O’Connell (2013). The analysis of this time window (400-600 ms) confirmed our findings (significant main effect of ‘Interruption’, $p = .005$, no other effects $p > .24$).

CPP amplitudes (as opposed to slopes) were also compared, by performing a series of FDR-controlled ANOVAs. For brevity, only results showing a corrected *p*-value of < .05 for at least 50 ms continuously are reported. In the stimulus-locked CPP, an ‘Interruption’ effect was observed between 466 and 783 ms (corrected $p < .049$; see Figure 1 d, where asterisks denote statistical effects on amplitude, not the previously described analysis on slopes). Fisher’s LSD-corrected post hoc tests found that the ‘continuous’ waveform displayed a higher amplitude than both the ‘random’ (between 466 and 783 ms relative to the onset of coherent motion, corrected $p < .02$) and the ‘reverse’ waveforms (between 488 and 783 ms, corrected $p < .046$). There was no significant difference in amplitude between ‘random’ and ‘reverse’ conditions (corrected $p > .26$). Further, we found a significant effect of ‘Difficulty’ in the time interval between 276 and 1000 ms relative to stimulus onset, with ‘easy’ waveforms reaching higher amplitudes than ‘hard’ waveforms (corrected $p < .046$). There was no significant interaction effect (corrected $p > .34$).

In the response-locked CPP, we found only a ‘Difficulty’ effect on amplitude, with ‘easy’ trials displaying a higher amplitude than ‘hard’ trials between -229 and 0 ms

relative to response. There was no main effect of 'Interruption' (corrected $p > .07$), and no interaction effect (corrected $p > .9$).

2.2.3. Model Fit

We fitted eight sequential sampling models (four independent race and four LCA) to the RT data. In each model type, models differ by assuming either symmetrical (models 1,3 and 5,7) or free modulations (models 2,4 and 6,8) of accumulation rates during the motion interruption period, which are applied either immediately (models 1,2 and 5,6) or after a free delay (models 3,4 and 7,8). For most individual participants (90% by AIC; 85% by BIC) no model was strongly supported (AIC/BIC improvement > 10) relative to all others. We thus averaged individual BICs (Schwarz, 1978) and AICs for each model to compare goodness of fit (see Table 1). It is clear that the exact ordering of models was criterion dependent, although the overall preference for the LCA class of model was not, with a pair of 2 (model class) x 2 (presence of delay) x 2 (presence of asymmetry) repeated-measures ANOVAs on both AIC and BIC showing main effects of model class ($F(1, 19) = 21.81$, $p < .001$, $\eta_p^2 = .53$ and $F(1, 19) = 13.11$, $p = .002$, $\eta_p^2 = .41$, respectively).[8] We elected to focus on BIC. The best (lowest) BIC was obtained for model 5, an LCA model with symmetric variation for the interrupted accumulation rate and no delay (see Table 1). Following Tukey correction, this model was reliably better than models 2, 4, 6 & 8 (i.e. all models allowing free modulation of accumulation rates during the interruption

[8] For our purposes here, model comparison was a means to an end, in terms of finding a reasonable candidate for the subsequent generation of neurodynamic predictions, not an end in itself. Hence we do not present detailed results breaking down these ANOVAs, both of which included three-way interactions, but instead simply summarise all possible pairwise comparisons (see main text).

period; all $p < 0.001$). Without such correction, it additionally beat model 1 ($p = 0.018$).

Table 1: Model Comparison: BIC and AIC values for each independent race (IRA) and LCA model. The BIC and AIC values of the chosen model (Model 5) are displayed in bold.

Model	Starting point interval	Decision threshold	Accumulation rates			Delay	Leak	Inhibition	Non-decision time	Non-decision time interval	Gaussian Noise SD	Number of parameters	AIC	BIC
			continuous	random	reverse									
Model 1 (IRA)	S_Z	A	v_{corr} v_{inc}	v-random$_{corr}$= v_{inc} v-random$_{incr}$= v_{inc}	v-reverse$_{corr}$ = v_{inc} v-reverse$_{inv}$ = v_{cor}	-	-	-	T_{er}	S_{Ter}	σ^2	8	3811	3850
Model 2 (IRA)	S_Z	A	v-continuous$_{corr}$ v-continuous$_{inc}$	v-random$_{corr}$ v-random$_{inc-}$	v-reverse$_{corr}$ v-reverse$_{inc}$	-	-	-	T_{er}	S_{Ter}	σ^2	16	3811	3889
Model 3 (IRA)	S_Z	A	v_{corr} v_{inc}	v-random$_{corr}$= v_{inc} v-random$_{incr}$= v_{inc}	v-reverse$_{corr}$ = v_{inc} v-reverse$_{inc}$ = v_{cor}	d	-	-	T_{er}	S_{Ter}	σ^2	9	3791	3835
Model 4 (IRA)	S_Z	A	v-continuous$_{corr}$ v-continuous$_{inc}$	v-random$_{corr}$ v-random$_{inc-}$	v-reverse$_{corr}$ v-reverse$_{inc}$	d	-	-	T_{er}	S_{Ter}	σ^2	17	3812	3894
Model 5 (LCA)	-	A	v_{corr} v_{inc}	v-random$_{corr}$= v_{inc} v-random$_{inc}$= v_{inc}	v-reverse$_{corr}$ = v_{inc} v-reverse$_{inc}$ = v_{corr}	-	k	β	T_{er}	S_{Ter}	σ^2	9	**3789**	**3833**
Model 6 (LCA)	-	A	v-continuous$_{corr}$ v-continuous$_{inc}$	v-random$_{corr}$ v-random$_{inc-}$	v-reverse$_{corr}$ v-reverse$_{inc}$	-	k	β	T_{er}	S_{Ter}	σ^2	17	3786	3868

Model 7 (LCA)	-	A	v_{cor} v_{inc}	$v\text{-random}_{corr} = v_{inc}$ $v\text{-random}_{inc} = v_{inc}$	$v\text{-reverse}_{corr} = v_{inc}$ $v\text{-reverse}_{inc} = v_{corr}$	d	k	β	T_{er}	S_{Ter}	σ^2	10	3787	3835
Model 8 (LCA)	-	A	$v\text{-continuous}_{corr}$ $v\text{-continuous}_{inc}$	$v\text{-random}_{corr}$ $v\text{-random}_{inc}$-	$v\text{-reverse}_{corr}$ $v\text{-reverse}_{inc}$	d	k	β	T_{er}	S_{Ter}	σ^2	18	3778	3865

As expected, mean accumulation rates (*v*) for the correct accumulator were higher in easy compared to difficult conditions. In this model, interruptions and reversals in evidence were modelled parsimoniously by substituting the appropriate parameters during this interval, rather than fitting new ones. Note that the exact parameter values returned should be treated with some caution, as a recovery study (Appendix A) suggested that this LCA model has issues with identifiability, i.e., some parameters can trade off to produce equally good fits (see discussion, below). Due to these identifiability issues, we do not report the parameter estimates for this model here, but have included them in the appendix (see Table A1).

Figure 1 b shows the quality of the model fit by displaying each participant's empirical (x-axis) and modelled (y-axis) RT quantiles (10%, 30%, 50%, 70%, 90%, increasing quantiles represented by increasingly dark colours) for each interruption condition as well as easy (top) and hard (bottom) trials (for behavioural fits for all other models, see Appendix B). The overlap between empirical and modelled quantiles indicates that the model fitted the data well.

2.2.4. Model Prediction

The parameters of the chosen model were then used to estimate individual accumulation profiles for each condition. Figure 2 displays the mean resulting predictions (b) and the corresponding EEG data (a) for stimulus (left) and response-locked (right) data. The model prediction was produced by summing correct and incorrect accumulators (see methods), and these contributory signals are shown separately as insets. Visual inspection shows that the EEG and predicted profiles are qualitatively very similar. With stimulus-locking, both profiles show an initial build-up which is slower (lower slope) in ‘hard’ (dashed lines) compared to ‘easy’ (solid lines) conditions, but similar across interruption conditions. Both profiles also show that the ‘continuous’ waveforms continue the build-up, while ‘random’ and ‘reverse’ waveforms display a plateau at approximately the same time, before continuing to build up. A further similarity between the model prediction and the EEG signal is the unexpected finding of a near complete overlap of the ‘random’ and ‘reverse’ conditions during the interruption period.

Figure 2

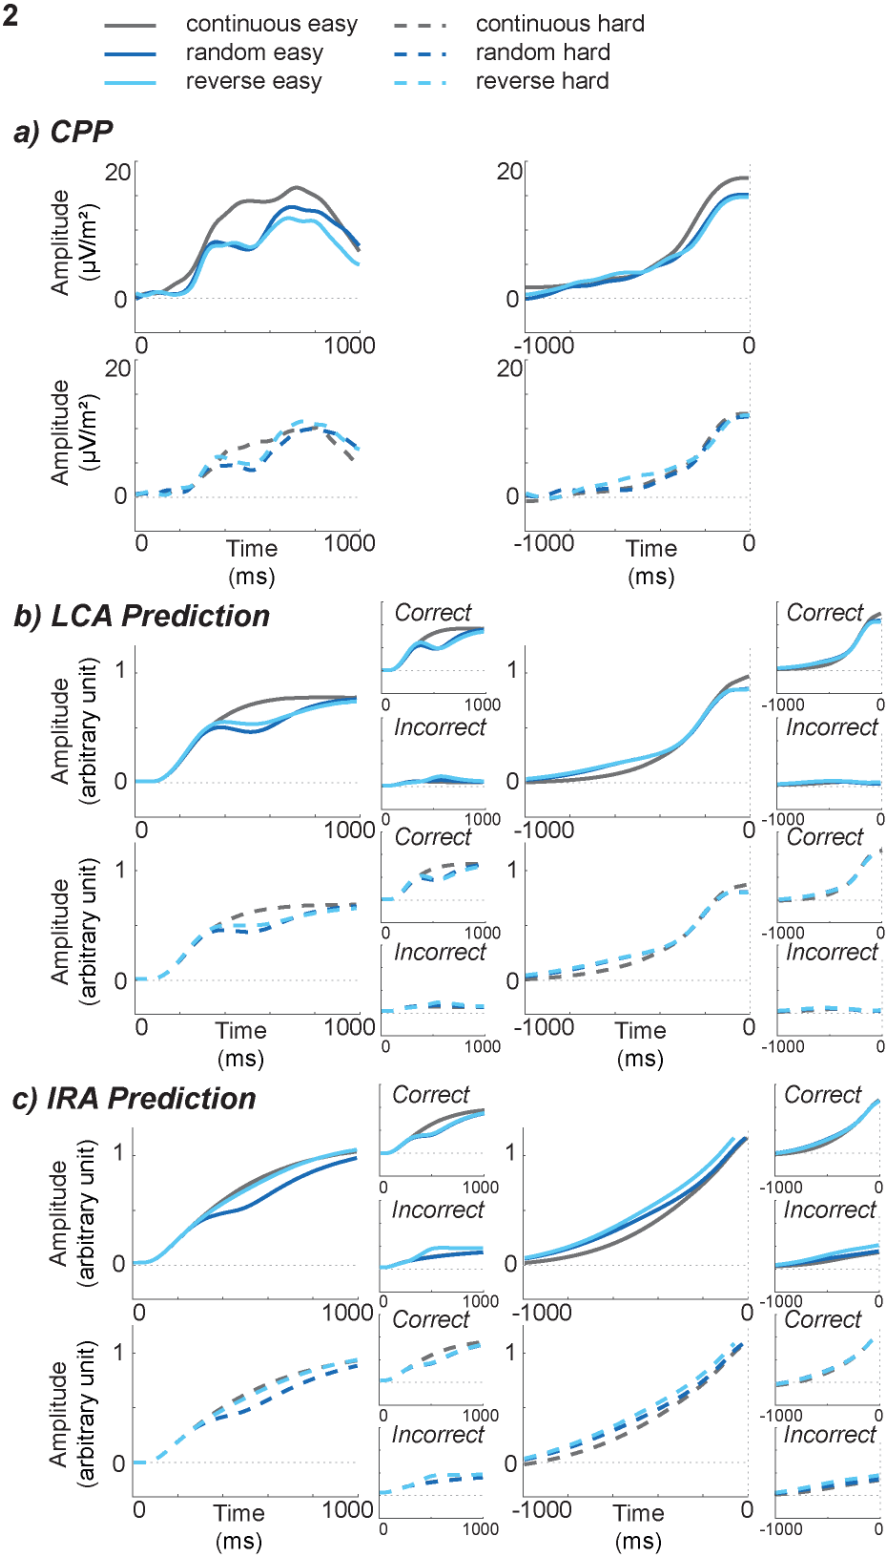

Figure 2: Decision variable (empirical and simulated): a) CPP waveform for easy (top, solid) and hard (bottom, dashed) trials, as well as stimulus (left) and response-locked (right) data. The CPP here has been filtered and downsampled to match model predictions. b) Accumulation profile (correct and incorrect accumulator summed) per Interruption condition as predicted by the best-fitting LCA model, for easy (top, solid lines) and hard (bottom, dashed lines) trials, as well as stimulus (left) and response-locked (right) data. Correct and incorrect

accumulators were summed to form the prediction, so these contributory signals are shown separately as smaller insets.c) Accumulation profile as predicted by the best-fitting independent race accumulator (IRA) model. Details as in part b.

While a degree of positive correlation over time between EEG signals and model predictions is to be expected for any ERP that grows across the RT period, the ability to predict differences between experimental conditions is more challenging and therefore more convincing. Hence, to quantify similarities between model predictions and neurodynamic data, we analysed the correlation between *differences* of conditions (differences between 'continuous – random', 'continuous – reverse', and 'random – reverse', for both easy and hard, as well as stimulus-locked and response-locked signals, resulting in a total of 12 correlations between the model predictions and the downsampled EEG; see 'Model Prediction (neurodynamics)'). We found that 9 out of 12 tests revealed significantly positive correlations ($r_{mean}(98) = .67$ $p_{mean} < .001$). All significant positive correlations remained significant after Bonferroni correction. Since 'random' and 'reverse' profiles largely follow the same trajectory, correlations between EEG and model signals reflecting the difference between these two conditions were naturally the lowest, and in fact, non-significant in some cases. The most meaningful correlations are therefore those between signals reflecting the difference between 'continuous' and 'random', and 'continuous' and 'reverse' conditions, specifically the stimulus-locked signals, as the manipulation in this experiment targeted the stimulus-locked trajectory of the accumulation. These correlations remained significant after Bonferroni correction ($r_{mean}(98) = .79$, $p_{mean} < .001$).

For reasons of concision, with eight models, our main focus when assessing the overlap between model predictions and EEG was on the model which best predicted

the behavioural data. However, we also assessed the extent to which the winning model from the other broad category (independent race model 3) could predict accumulation signals resembling the CPP. Indeed, behaviourally, this model was almost indistinguishable from LCA model 5 in terms of its ability to capture RTs. Neurodynamic predictions for independent race model 3 are shown in Figure 2 (c). As can be seen, although the global accumulation pattern is present, the independent race model does not predict the empirical observation of no difference during the interruption period between the 'random' and 'reverse' conditions. However, although for this model the raw predictions looked rather less well matched to their corresponding EEG signals, correlations based on differences between conditions followed a broadly similar pattern to that observed for LCA model 5, i.e., the best fitting independent race model also did a good job of predicting the time-varying ordering of EEG signals in different conditions (10 out of 12 tests revealed significant correlations after Bonferroni correction $r(98) = .51$, $p = .001$). This highlights that the correlations used here should not be used in isolation in order to evaluate different models.

2.3. Discussion Experiment 1

In the first experiment, we tested the impact of non-stationary evidence on the CPP, a potential neural substrate of the decision variable. Assuming that a change in evidence must necessarily induce a change in the accumulation profile, the CPP waveform should display a similar time-varying build-up in order to support its role as a decision variable signal. To test this, we observed the CPP under three different conditions: a 'continuous' condition in which the evidence was constant throughout the trial, a 'random' condition in which the evidence was stopped for a brief interval

and replaced by random noise, and a ‘reverse’ condition in which the evidence was reversed to support the opposite response alternative for a brief period. We also added a more established manipulation (task difficulty) as a positive control. We expected that the continuous condition would lead to the stereotypical, smooth build-up, while the random and reverse profiles should deviate from this build-up to varying extents. Critically, however, we went beyond intuitive predictions about the interrupted decision variable, by first using our RT data to identify and fine-tune a plausible behavioural model, and then using this model to formulate exact predictions for the CPP under the assumption that this spatially diffuse EEG component should represent a sum of accumulators within a race-model framework. As we expand below, the resulting correspondence between model predictions and CPP was striking.

Both evidence interruption and difficulty manipulations had the expected effects on participants’ performance, with faster and more accurate responses in ‘easy’ than ‘hard’ trials, and when evidence was ‘continuous’. The slowest and least accurate responses were observed in ‘reverse’ trials, while the ‘random’ condition lengthened RT compared to continuous trials, with a less clear impact on accuracy. Hence, interrupted trials led to worse performance, with evidence reversal disrupting the decision more than a simple pause. These findings are broadly in line with previous research (Holmes et al., 2016; Huk & Shadlen, 2005; O’Connell et al., 2012; Tsetsos et al., 2012).

We infer that these changes in performance were caused by modulations of decision-related evidence accumulation. It is well-established that difficulty affects

the slope of accumulation, with easier stimuli leading to steeper evidence accumulation (Brown & Heathcote, 2008; Kelly & O'Connell, 2013; Ratcliff & McKoon, 2008; Ratcliff & Rouder, 1998). The interruption of evidence, on the other hand, should lead to an interruption in accumulation. To formalise this account of the behaviour we observed, we tested several LCA and independent race accumulator models, and found that an LCA model with symmetrical changes of accumulation rates during the epoch of interruption (and for different difficulty levels) provided the best account of our RT data (although other models were viable).

We hypothesised that a pause in evidence would cause the accumulation to stop and plateau for the duration of the interruption interval. The impact of the 'reverse' condition on the accumulation profile is somewhat harder to predict intuitively, and is probably more dependent on the specifications of the model. For instance, the assumption of reciprocal inhibition between accumulators may attenuate the impact of evidence reversal. Specifically, the accumulator corresponding to the initial direction of dot motion may inhibit the accumulator receiving the reversed evidence in most trials, hence limiting accumulation growth during the reversal period. Issues like these led us to emphasise modelling in formulating predictions.

We used the estimated parameters from our best-supported LCA model to simulate the accumulation profiles (and, in particular, their sum) associated with each condition, and directly compared the resulting patterns to the CPP. We found considerable overlap between the model predictions and the neural signal, even though these profiles were not fitted to one another directly. As previously reported (Kelly & O'Connell, 2013; Twomey et al., 2015), task difficulty affected the slope of

the CPP, with lower build-up rates in ‘hard’ decisions. A very similar difference appeared in model predictions. Furthermore, we obtained novel evidence that both model predictions and the CPP showed the same gradual build-up in the ‘continuous’ condition, and interruption of this build-up (which plateaued before continuing to build up approximately 300 ms later) in the ‘random’ and ‘reverse’ conditions. Interestingly, model predictions also mimicked the CPP signal in terms of the unexpected similarity between the ‘random’ and ‘reverse’ waveforms. These patterns are particularly telling as they show an overlap between neural data and evidence accumulation which might not have been predicted based on intuitive reasoning alone. Our results build on previous research which found that the evolution of the CPP is sensitive to a brief interruption of evidence (O’Connell et al., 2012) by testing additional conditions, in a choice rather than simple RT task, and making more precise model-derived predictions. Overall, the similarities we observed seem to support the role of the CPP as a neural substrate of decision-making.

An additional finding worth noting is the delay in the disruption of the CPP build-up compared to the timing of the evidence interruption. While the interruption of motion took place between 200 and 400 ms after stimulus onset, the divergence in CPP amplitude between ‘continuous’ profiles and the two interrupted profiles was observed around 470-780 ms post stimulus. This finding supports the role of the CPP as an accumulation signal, rather than a mere sensory signal, which would arguably display a faster reaction in response to the change in evidence, suggesting that it represents a higher-level integration of evidence.

The details of our best-fitting model are somewhat suggestive regarding the way evidence accumulation follows from operations occurring in sensory regions of the brain. Holmes et al. (2016) found that a change in evidence was better explained by a new, independent accumulation rate, rather than a symmetric change of rates, even when the change in evidence itself was symmetric. We instead found that a change in evidence could be explained by a (more parsimonious) swap in accumulation rate during the interruption interval. Essentially, Holmes et al. (2016) found steeper accumulation rates after evidence reversal, while our results support a symmetrical rate change during the reversal period. In fact, some non-linearity in the sensory representation of a time-varying motion signal is to be expected (with the waterfall effect offering a well-known example of repulsive sensory after-effects, which are themselves complemented by assimilative tendencies; Addams, 1834; Yarrow, Minaei, & Arnold, 2015). However, the exact time-course of such effects are somewhat challenging to predict. The difference in findings here relative to Holmes et al. (2016) may perhaps be explained by the different task procedures, as we used brief perturbations while the evidence in their study remained reversed for the rest of the trial. It is conceivable that sensory evidence rebounds after a change, perhaps via sensory repulsion, and is thus accumulated faster, but only after some delay, which is why Holmes et al. observed it and we did not. It is interesting to note that even for our independent race models (which were more equivalent to Holmes et al.'s piecewise LBA) a symmetric change of rates proved sufficient in our experiment. Differences between our findings and those reported by Holmes et al. (2016) may further be due to methodological differences in the way the models were fitted to the data. While in the current study, we used Quantile Maximum Probability Estimation, Holmes et al. (2016) fitted reaction time distributions using hierarchical

Bayesian methods, which may be sensitive to different subtleties in the data, leading to different findings.

Another divergence between the two studies is that while Holmes et al.'s best model introduced a delay between the presentation and the incorporation of the new evidence, explaining the temporal lag between the change in evidence and its behavioural consequences, positive evidence for this delay was not observed in the current study. This difference may be explained by the type of model used. The LCA model implements reciprocal inhibition between accumulators, which presumably smooths accumulation-rate variations and produces a slow response to the change in evidence without the need for a delay parameter. In the case of independent accumulators on the other hand, as in Holmes et al.'s piecewise LBA, a delay parameter is necessary to model the slow response to changing evidence (note that our results using independent race accumulator models were consistent with Holmes et al.'s findings). We hence confirm that a change in evidence is explained by change in accumulation rates, and that some time is necessary for those changes to feed through and become visible in the decision variable. However, while a delay parameter was previously introduced to account for this temporal lag, we propose that it could naturally arise from reciprocal inhibition between accumulators, as implemented in the LCA model. Note, however, that our conclusion favouring an LCA model without any delay was dependent on our decision to elevate BIC over AIC in model comparison, and that a model with delay performed similarly well.

Finally, although LCA complexity seems advantageous in this case, it is also known to induce parameter recovery issues and has been described as a model in which

different combinations of parameters values result in similar reaction time distributions (Miletić et al., 2017). In a recovery study (Appendix A) we also observed poor recovery for several of the parameters, with the implied trade-off being consistent with that observed by Miletić et al.'s (2017). Presumably, the values of common fractions of accumulation rates, leakage and inhibition trade off, making accurate estimation of parameter values hard to achieve. Importantly, however, we additionally observed that this only had a moderate impact on CPP predictions derived from the fitted parameters, most probably because parameters also trade off in the accumulation signal. Hence, although difficulties of parameter estimation must be considered when one draws conclusions on parameter values, investigation of derived accumulation signals may be less affected.

3. Experiment 2: Decision Bias

Experiment 1 suggested that the CPP reflects the complex decision variable generated by a requirement to track time-varying sensory evidence. However, a viable neurodynamic correlate should respond appropriately to a wide range of manipulations known to affect the decision process. In Experiment 2, we went on to test the effects of decision biases on the CPP. Probabilistic decision biases are associated with strong behavioural effects, and can often be explained using sequential sampling models by varying just one parameter (Summerfield & de Lange, 2014; but see Rae, Heathcote, Donkin, Averell & Brown, 2014). In a sequential sampling process, evidence is accumulated from a given starting point towards a threshold. With the introduction of a bias towards a given alternative (e.g. a greater a priori likelihood that that alternative will be evidenced) the starting point

moves towards the respective threshold, thereby decreasing the amount of evidence required to make the choice in favour of the biased alternative (Bode et al., 2012; Gao, Tortell, & McClelland, 2011; Mulder et al., 2012; Spaniol et al., 2011; Summerfield & Koechlin, 2010; Teodorescu & Usher, 2013; Voss et al., 2013). In contrast to Experiment 1, in which the shape of the accumulation process was affected, here, we set out to investigate the impact of varying the magnitude of accumulated evidence required for a decision on the CPP waveform. To our knowledge, the impact of probabilistic decision biases on the CPP has not been tested so far.

The neurodynamics of biased decisions have nonetheless been explored before in other ways. Rorie, Gao, McClelland, and Newsome (2010) presented monkeys with a binary motion-discrimination task in which the reward for the two choices was either equal or unequal. Rewards primarily influenced LIP firing rates prior to the motion onset, with unbalanced payoffs leading to a baseline shift towards the rewarded threshold. These findings support the notion of a starting point difference in accumulation for biased decisions. No difference in the slope of the build-up in firing rate throughout the decision was observed. The same finding of a shift in baseline activity and unaltered slopes in LIP firing rates was supported when instead of unequal payoffs, predictive directional cues were used in a motion discrimination task (Rao, DeAngelis, & Snyder, 2012). Similarly, it has been shown that firing rates in neurons which show a build-up to threshold profile associated with a given choice show a reduction in baseline activity with decreasing probability of this choice (Basso & Wurtz, 1998; Dorris & Munoz, 1998), further supporting the role of starting point activity in decision biases.

Evidence regarding biased neural correlates of evidence accumulation in humans remains somewhat scarce. EEG research has focused primarily on motor signals to track decision biases. Noorbaloochi et al. (2015) recorded human EEG during a decision task with either biased or unbiased payoffs and explored the lateralised readiness potential (LRP) as a signal reflecting evidence accumulation. In line with findings from non-human primates, it was found that in biased decisions, the LRP amplitude was shifted towards the alternative associated with the higher payoff prior to stimulus onset, suggesting a starting-point difference. Additionally, de Lange et al. (2013) concluded that it is a variation in accumulation starting point which accounts for bias-related activity. Using MEG, de Lange and colleagues found that motor-related activity in the beta frequency range displayed a pre-stimulus bias in the direction associated with the biased alternative. Together, these data suggest that biases push accumulation signals prior to the accumulation onset towards the threshold, without affecting the accumulation slope. However, recently Afacan-Seref et al. (2018) have reported somewhat different results in a study recording the CPP and LRP during binary choice with strongly biased rewards and extreme time pressure. They modelled an accelerating accumulator and found effects on the slope of accumulation, with some specific predictions regarding slow, low-valued choices mirrored in the CPP. We return to this study in the discussion.

To our knowledge, the effects of probabilistic decision biases on CPP profiles have not yet been explored. We therefore set out to explore the CPP waveform under different bias conditions. We presented cues which either provided information regarding the likely direction of subsequent motion or gave no directional

information. Based on the literature summarised above, we expected that the presence of a directional cue would lead to a shift in accumulation starting point, decreasing the baseline-to-threshold distance in the accumulator corresponding to the cued response. Regarding the CPP waveform, this baseline variation would appear as a modulation in amplitude, since the CPP computation requires a baseline correction (i.e. the decreased baseline-to-threshold distance in correctly cued trials would translate to a decrease in the magnitude of the accumulation). However, if we assume that the CPP reflects the sum of both accumulators, the CPP waveform should also be affected by changes occurring in the accumulator opposed to the cue. If a decreased starting point in the non-cued accumulator co-occurs symmetrically with the increased starting point in the cued accumulator, it is possible that we would observe no difference in the waveforms. Again, fitting a sequential sampling model to the resulting behavioural data and directly comparing accumulation profiles simulations to the recorded CPP waveforms is crucial to yield insights into the role of the CPP as an accumulation signal.

3.1. Methods

Methods were, unless otherwise stated, identical to Experiment 1.

3.1.1. Participants

Twenty participants (five males), with a mean age of 30.15 (*SD* = 7.28) were recruited. All participants met the pre-defined requirement to achieve an average accuracy score of 80% in the random dot motion task at a coherence level no greater than 90% (i.e. 90% of dots moving coherently). Each participant took part in a session lasting between 2 and 2.5 hours.

3.1.2. Stimuli and Procedure

All participants first completed a minimum of 50 practice trials at a coherence level of 80%. During the practice trials, feedback was provided after each trial (‘correct’/‘incorrect’). Afterwards, each participant completed 100 trials without feedback in order to establish an appropriate level of difficulty for the experiment via a QUEST staircase procedure targeting 80% correct. The resulting average level of coherence was 32.25% (*SD* = 27.92).

For the main experiment, each participant completed 450 trials. The trial procedure is displayed in Figure 3 a. In each trial, a fixation cross was followed by a cue (500 ms) that consisted of two arrows, one pointing to the left, and one pointing to the right. In one third of the trials, both arrows were white, indicating no specific direction (‘uncued’ trials), while in two thirds of the trials, one arrow was yellow, providing a cue towards a given direction. Left-pointing and right-pointing cues were equiprobable. In each trial, the cue was followed by random dot motion, i.e. at a coherence level of 0%. After the random motion, the coherent motion started (left/right) and lasted up to 1300 ms or until the response. Note that the deadline is shorter here than in experiment 1, due to the decreased difficulty of the task. If a directionally specific cue was given, the subsequent dot motion corresponded with the cue direction 80% of the time (‘congruent’ trials), and opposed it in 20% of trials (‘incongruent’ trials). No feedback was provided after each trial, but every 60 trials, participants took self-timed breaks during which they were provided with feedback in the form of mean accuracy scores and RTs over that period.

3.1.3. Statistical Analysis

In order to analyse the impact of the different cue conditions on the ERP waveform, we again compared both the slope and the amplitude between conditions. Like in Experiment 1, we compared the build-up rate by fitting a straight line to the waveform. The chosen time intervals to which we fitted a line were 200 to 350 ms for the stimulus-locked CPP, and -200 to -150 ms for the response-locked CPP (Kelly & O'Connell, 2013). The resulting slopes were then compared using a one-way ANOVA to compare 'congruent', 'incongruent', and 'uncued' waveforms.

3.1.4. Model Fit

Again, independent race accumulator and LCA classes of models were used to model RT data. Within each class we tested a total of five different models, all accounting for bias conditions using starting point modulations, but assuming different mechanisms in order to account for different bias conditions.

Model 1 was an independent race model assuming that cues induced changes of accumulation starting point in the accumulator corresponding to the cued response. In 'cued' trials, the lower limit of the starting point distribution, Z, was increased by the bias parameter in the cued accumulator, and was set to 0 in the accumulator opposite to the cue. In 'uncued' trials, the value of Z was fixed to 0 in both accumulators. Trial-to-trial starting point variability was introduced, such that each accumulator starting point was drawn from a uniform distribution on the interval [Z $Z+S_z$]. Hence, on average, the starting point was higher in the cued accumulator than the accumulator opposite to the cue, and both accumulators in the neutral condition. Note that this results in starting point changes in the 'correct' accumulator in congruent trials and the 'incorrect' accumulator in 'incongruent' trials. Specifically,

the bias parameter should favour the ‘correct’ accumulator in ‘congruent’ trials and the ‘incorrect’ accumulator in ‘incongruent’ trials. All other parameters were fixed between conditions, resulting in a model with a total of seven parameters (see Table 2).

Models 2 and 3 were also independent race models implementing starting point variations, but now impacting both cued and opposite accumulators. Model 2 assumed symmetrical changes while model 3 assumed free variations. In model 2, again, the lower limit of the starting point distribution Z was fixed to 0 in the accumulator opposite to the cue, and was increased by the bias parameter in the cued accumulator. In this case however, in ‘uncued’ trials, the value of Z was fixed in both accumulators to half of the bias parameter value. Again, each accumulator starting point was drawn from the interval $[Z\ Z+S_z]$. Therefore, on average, starting point variations of equal magnitude but opposed sign were applied in the cued and the opposite accumulators compared to the neutral condition, leading to opposite effects on the ‘correct’ and ‘incorrect’ starting point in ‘congruent’ and ‘incongruent’ trials. Model 3 assumed similar mechanisms, with free rather than symmetrical changes in ‘cued’ compare to ‘uncued’ trials. Here again, Z was fixed to 0 in the accumulator opposite to the cue, and increased by the bias parameter in the cued accumulator. In this case however, Z was also free to vary in ‘uncued’ trials. As such, free variations of the lower limit of starting point interval occurred in ‘cued’ compared to ‘uncued’ trials. Again, note that this translated into inverse effects on ‘correct’ and ‘incorrect’ accumulators between ‘congruent’ and ‘incongruent’ trials (see Table 2) but without assuming that uncued accumulators started (on average) midway between congruent and incongruent ones. All other parameters were fixed between

conditions, resulting in a total of seven parameters in model 2 and eight parameters in model 3 (see Table 2).

Finally, models 4 and 5 tested whether cues also influenced the rate of evidence accumulation, again assuming either symmetrical or free variations. Model 4 implemented symmetrical starting point variation as in model 2, plus symmetrical changes in accumulation rates. V_{cued} was added to the cued accumulator rate, and was subtracted from the opposite accumulator rate. In model 5, assuming free changes, starting point variations were implemented as in model 3, and v_{cued} was added to the cued accumulator rate while v_{opp} was subtracted from the opposite rate. Again, note that the 'cued' accumulator was 'correct' in 'congruent' trials and 'incorrect' in 'incongruent' trials. These manipulations resulted in a total of eight and ten parameters in models 4 and 5, respectively (see Table 2).

Model 6 to 10 were LCA implementations of Model 1 to 5 respectively (see Table 2). Like in Experiment 1, best-fitting model parameters were determined at the individual level. Trials with RTs faster than 180 ms or slower than 1300 ms (less than 6%) were discarded.

3.2. Results

3.2.1. Behavioural Results

The data remaining after trimming outlying RTs (5.34%) are displayed in Figure 3 c. Statistical analyses revealed RT differences between cue conditions, $F(2, 38) = 42.72$, $p < .001$, $\eta_p^2 = .69$. Fisher's LSD corrected follow-up t-tests showed that all conditions differed from each other, with faster RTs in 'congruent' than in 'uncued',

$t(19) = 6.21$, $p < .001$, and 'incongruent' trials, $t(19) = 7.38$, $p < .001$, and in 'uncued' than 'incongruent' trials, $t(19) = 5.17$, $p < .001$.

Additionally, a GLMM revealed that the 'Cue' condition affected accuracy scores, $F(2, 57) = 18.56$, $p < .001$. To explore the differences between all three levels, we fitted the model a second time, but setting the reference level of 'Cue' to 'incongruent', rather than 'congruent'. Results showed that accuracy was higher in 'congruent' compared to 'uncued' trials, with both being higher than accuracy of 'incongruent' trials (all $p < .001$).

3.2.2. ERP Results

The CPP waveform for each condition is displayed in **Error! Reference source not found.** 3 d. In both the stimulus-locked and the response-locked CPP, the waveform associated with 'incongruent' trials displays the highest amplitude, followed by the 'uncued' and 'congruent' waveforms. Note that the interpretation of the CPP, when related to the predictions of sequential-sampling models, requires that we keep in mind the baseline correction applied to ERPs. Higher end points are consistent with greater excursions, which may be implemented in models as lower starting points, and vice versa.

Figure 3

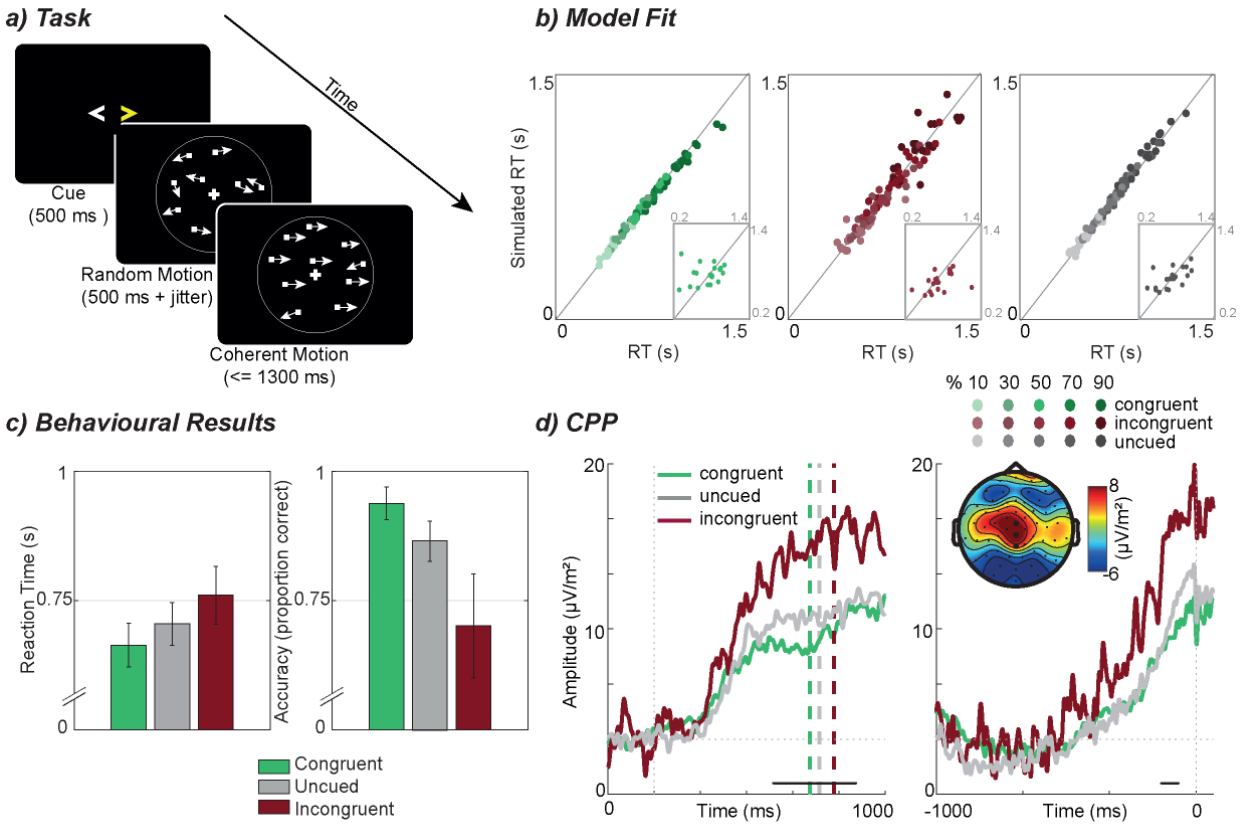

Figure 3: a) Random dot motion task trial procedure: in each trial, a cue consisting of two arrows was presented. If both arrows were white ('uncued'), no directional information was given. If one of the arrows was yellow, this cue correctly described the direction of the upcoming motion in 80% of trials ('congruent'), and was false in 20% of trials ('incongruent'). Here, the right side is cued, and the coherent motion following the random motion is to the right ('congruent'). Note that the size and number of dots have been adjusted for illustration. b) Model fit: each participant's quantiles estimated from behavioural data (x-axis) and race model simulations (y-axis) for each cue condition (from left to right: congruent, incongruent, uncued). Increasing quantiles (10%, 30%, 50%, 70%, 90%) are represented by increasingly darker colours. Small inserted panels show observed and simulated RT medians for error trials.c) Behavioural results: reaction time (left) and accuracy (right) averages for 'congruent', 'incongruent', and 'uncued' trials. Error bars indicate 95% confidence intervals. d) CPP results: Stimulus-locked (left) and response-locked (right) CPP waveforms. Electrodes used to generate the waveforms are highlighted on the topography (which shows the averaged signal over the stimulus-locked 0 to 1000 ms interval). Vertical dashed lines in the stimulus-locked CPP indicate mean RTs per condition. Note that the mean RTs are based only on trials which were included in the generation of the waveform and differ slightly from the ones displayed in c. Black dots at the bottom of the waveform indicate time points at which FDR-controlled comparisons of amplitude showed a significant 'Cue' effect.

No difference in the CPP slopes was observed across the different conditions, in either the stimulus-locked, $F(2, 38) = .39$, $p = .68$, $\eta_p^2 = .02$, or the response-locked CPP, $F(2, 38) = .40$, $p = .67$, $\eta_p^2 = .02$. We also tested the variation of amplitudes in the CPP using a series of FDR-controlled ANOVAs and found a significant effect of 'Cue' between 518 and 873 ms relative to the onset of coherent motion (corrected $p < .049$). Follow-up t-tests revealed that 'incongruent' amplitudes were higher than both the 'congruent' (for the entire duration of the main effect, corrected $p < .02$), and the 'uncued' ones (between 542 and 863 ms relative to stimulus onset, corrected $p < .05$). There was less difference between 'congruent' and 'uncued' amplitudes (corrected $p < .05$ only between 639 and 645 ms).

In the response-locked CPP, we found a significant main effect between -198 and -104 ms relative to the response (corrected $p < .047$). Post-hoc tests showed the same patterns as the stimulus-locked data, with higher amplitudes in 'incongruent' than 'congruent' trials (during the entire duration of the main effect, corrected $p < .018$) and in 'incongruent' than 'uncued' trials (between -198 and -108 ms, corrected $p < .049$). There was no difference between 'congruent' and 'uncued' trials ($p > .09$).

3.2.3. Model Fit

Ten models assuming changes in starting point across bias conditions were fitted to the data. We once again focussed on BIC to help us discriminate between them. For individual-level fits, there were no cases where a participant's data were strongly supportive of one model over all others (BIC or AIC difference > 10). The best (lowest) group-average BIC was obtained for Model 2, an independent race model with a symmetrical cuing bias affecting start points of accumulation (see Table 2). Tukey-corrected contrasts suggested that this model significantly outperformed

models 3, 5 and 10. Without correction, it additionally beat models 1, 4, 8 and 9, but not models 6 or 7[9]. This is somewhat suggestive that the additional bias and/or inhibition/leak parameters of many of the other models did not increase the quality of the fit enough to warrant the increased model complexity. However, model 6, a simple LCA model with only a positive cuing bias, performed best based on AIC. Somewhat arbitrarily, we begin by discussing accumulation profiles for model 2, but go on to consider them for model 6 as the best performer from the other model class (for behavioural fits for all models, see Appendix B).

Table 2: Model Comparison: BIC and AIC values for each independent race (IRA) and LCA model. The BIC and AIC values of the chosen model (Model 2) are displayed in bold.

Model	Starting point lower limit	Starting point interval	Response threshold	Accumulation rates		Leak	Inhibition	Non-decision time	Non-decision time interval	Gaussian noise SD	Number of parameters	AIC	BIC
Model 1 (IRA)	Neutral: $Z = 0$ Cued: $Z = bias$ Opp: $Z = 0$	$[Z\ Z+S_z]$	A	v_{corr} v_{inc}	- -	-	-	T_{er}	S_{Ter}	σ^2	7	1525	1546
Model 2 (IRA)	Neutral: $Z = bias/2$ Cued: $Z = bias$ Opp: $Z = 0$	$[Z\ Z+S_z]$	A	v_{corr} v_{inc}	- -	-	-	T_{er}	S_{Ter}	σ^2	7	**1516**	**1543**

[9] A 2 (model class) x 5 (model details) repeated-measures ANOVA gave little evidence that independent race models were generally better supported than LCA models (with no main effects) for either AIC or BIC, but did yield interactions in both cases.

Model 3 (IRA)	Neutral: $Z = Z$ Cued: $Z = bias$ Opp: $Z = 0$	[Z Z+S_z]	A	v_{corr} v_{inc}	- -	-	-	T_{er}	S_{Ter}	σ^2	8	1515	1547
Model 4 (IRA)	Neutral: $Z = bias/2$ Cued: $Z = bias$ Opp: $Z = 0$	[Z Z+S_z]	A	v_{corr} v_{inc}	v_{cued} $v_{cued} * (-1)$	-	-	T_{er}	S_{Ter}	σ^2	8	1515	1546
Model 5 (IRA)	Neutral: $Z = Z$ Cued: $Z = bias$ Opp: $Z = 0$	[Z Z+S_z]	A	v_{corr} v_{inc}	v_{cued} v_{opp}	-	-	T_{er}	S_{Ter}	σ^2	10	1516	1556
Model 6 (LCA)	Neutral: $Z = 0$ Cued: $Z = bias$ Opp: $Z = 0$	-	A	v_{corr} v_{inc}	- -	k	β	T_{er}	S_{Ter}	σ^2	8	1513	1545
Model 7 (LCA)	Neutral: $Z = bias/2$ Cued: $Z = bias$ Opp: $Z = 0$	-	A	v_{corr} v_{inc}	- -	k	β	T_{er}	S_{Ter}	σ^2	8	1515	1546
Model 8 (LCA)	Neutral: $Z = Z$ Cued: $Z = bias$ Opp: $Z = 0$	-	A	v_{corr} v_{inc}	- -	k	β	T_{er}	S_{Ter}	σ^2	9	1515	1550
Model 9 (LCA)	Neutral: $Z = bias/2$ Cued: $Z = bias$ Opp: $Z = 0$	-	A	v_{corr} v_{inc}	v_{cued} $v_{cued} * (-1)$	k	β	T_{er}	S_{Ter}	σ^2	9	1515	1550

Model 10 (LCA)	Neutral: $Z = Z$ Cued: $Z = bias$ Opp: $Z = 0$	-	A	v_{corr} v_{inc}	v_{cued} v_{opp}	k	β	T_{er}	S_{Ter}	σ^2	11	1516	1559

Table 3: Mean estimated parameter values for the chosen model (Model 2): note that the response threshold A was set to 1 as a scaling parameter, and that all lower limits of the starting point distributions were generated with just two free parameters. Note that, due to the raised starting point in the uncued condition, these parameters are not directly comparable to the ones displayed in Experiment 1 (Table A1).

Model 2: Parameters			
Lower limit starting point	'congruent'	correct	0.2598
		incorrect	0
	'incongruent'	correct	0
		incorrect	0.2598
	'uncued'	correct	0.1628
		incorrect	0.1628
Starting point variability (S_Z)			0.3389
Response threshold (A)			1
Accumulation rate (v)		correct	1.6709
		incorrect	0.2867
Non-decision time (T_{er})			0.300
Non-decision time interval (S_{Ter})			0.220
Gaussian noise SD (σ^2)			0.5698

The parameter estimates of the chosen race model are displayed in Table 3. Figure 3 b shows the quality of the model fit by displaying each participant's empirical (x-axis) and modelled (y-axis) RT quantiles in each condition. It indicates that independent race accumulator model 2, with varying starting points, can account well for our biased decision-making.

3.2.4. Model Prediction (neurodynamics)

Model parameters were used to compute the predicted accumulation profile for each condition. Figure 4 displays the resulting predictions (b) and the corresponding CPP (a) for stimulus (left) and response-locked (right) signals. Components of the prediction (correct and incorrect accumulators) are shown as insets. Visual inspection shows some qualitative similarities between the best independent race accumulator model predictions and the EEG signals. Importantly, both the model prediction and the CPP display an amplitude difference in the response-locked signal, with 'incongruent' decisions being associated with the highest values. However, this pattern is not visible in the stimulus-locked prediction, despite appearing in the corresponding EEG signal. Furthermore, the amplitude variations appear far more pronounced in the EEG signals than in the model predictions.

As in Experiment 1, we analysed the correlation between differences of conditions in both the EEG data and the model predictions (differences between 'congruent – incongruent', 'congruent – uncued', and 'incongruent – uncued', for both stimulus-locked and response-locked signals, resulting in a total of 6 correlations). We found that 3 out of 6 tests revealed significant positive correlations, all of which remained significant after Bonferroni correction ($r_{mean}(98) = .44$, $p_{mean} < .001$). Since this

experiment targeted the amplitude of the accumulation, which is visible primarily in the response-locked profiles, the correlations between response-locked signals, which were all significant ($r_{mean}(98) = .44$, $p_{mean} < .001$), are arguably most meaningful.

Finally, as in Experiment 1, we looked at predictions from the best-performing model in the alternative class (LCA model 6, Figure 4c). Here, predictions were noticeably less consistent with the EEG signal. In fact, an identical correlation analysis run on differences between conditions showed an equal tendency towards both positive and *negative* significant correlations after Bonferroni correction (three correlations revealed positive results, $r_{mean}(98) = .46$, $p_{mean} < .001$, and two showed negative results, $r_{mean}(98) = -.46$, $p_{mean} < .001$), i.e. a failure to properly order the EEG signals from the three conditions across time.

Figure 4

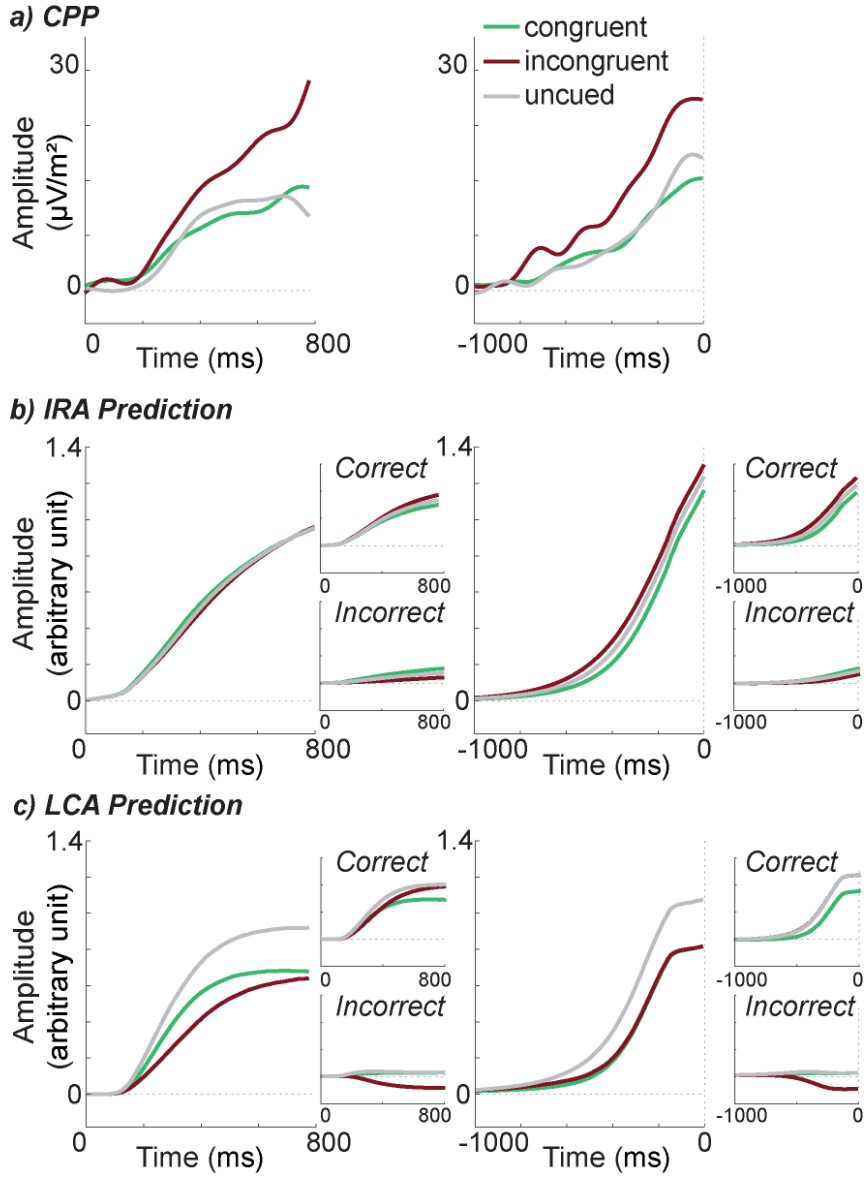

Figure 4: Decision variable (empirical and simulated): a) CPP waveform for stimulus (left) and response-locked (right) data. The CPP here has been filtered and downsampled to match model predictions. b) Accumulation profile per cue condition as predicted by the best-fitting independent race accumulator model (IRA), for stimulus (left) and response-locked (right) data. Correct and incorrect accumulators were summed to form the prediction, so these contributory signals are shown separately as smaller insets.c) Accumulation profile as predicted by the best-fitting LCA model. Details as in part b.

3.3. Discussion Experiment 2

In Experiment 2, we tested how decision biases affect the CPP waveform and, like in Experiment 1, compared its profile to model predictions. To this end, we asked participants to complete a motion discrimination task in which cues prior to each trial either gave no information about the direction of the upcoming trial ('uncued'), or indicated the upcoming direction either correctly ('congruent') or incorrectly ('incongruent'). In accordance with previous research (de Lange et al., 2013; Mulder et al., 2012; Teodorescu & Usher, 2013), we observed that participants' choices were biased towards the cued direction. Compared to 'uncued' trials, 'congruent' cues resulted in faster RTs and less errors, while 'incongruent' cues led to lower accuracy and longer RTs. Note that in order to avoid the co-occurrence of visual evoked potentials (associated with a sudden stimulus onset) with the accumulation profile, we added a period of random dot motion prior to the coherent motion but following the directional cue (Figure 3 a). This means that there was a short period of time where participants were presented with a stimulus which was potentially inconsistent with the cue, even in congruent trials, which may have weakened the effect of the cue. However, since we observed strong behavioural differences between all three conditions, we do not believe that this had a qualitative impact on our conclusions. Nevertheless, we note that this may hinder the direct comparison with versions of decision-making tasks in which the evidence immediately follows the cue.

The observed changes in behaviour were well captured by an independent race model with varying start points, and this model predicted some of the trends we observed in the CPP as decisions were being made, albeit imperfectly. However, this result may be viewed as somewhat fortuitous. Although generating predictions for

the independent race model followed a natural logic, because this model (just about) won at a behavioural level, the other class of models we considered here, with inhibition and leakage, failed to capture nuances in the CPP.

Based on previous research, we hypothesised that prior cues would affect the starting point of each accumulator (Bode et al., 2012; Gao et al., 2011; Rorie et al., 2010; Teodorescu & Usher, 2013)leading to a change in the baseline-to-threshold distance, and incorporated free parameters capable of capturing this change. For the best-fitting model, the mean starting point was higher in the corresponding cued accumulator and lower in the opposite non-cued accumulator compared to the neutral, uncued, condition. By modifying the baseline-to-threshold distance, starting point variations affect both the time required for accumulation to reach the decision threshold and the probability of attaining the threshold due to noise. In incongruent trials, for example, where the incorrect response was cued, errors occurred frequently due to the small baseline-to-threshold distance in the cued, but incorrect, accumulator, and correct RTs were slower due to the larger baseline-to-threshold distance in the opposite non-cued accumulator[10]. In line with many, but not all, previous studies, our results hence confirmed that decision biases can be accounted for by simply varying accumulation starting point (Basso & Wurtz, 1998; de Lange et al., 2013; Rao et al., 2012; but see Rae et al., 2014).

[10] In the case of the best LCA model, which incorporated a change to only the cued starting point, correct RTs would instead be slower due to the extra inhibition flowing from the boosted correct accumulator towards the non-cued accumulator.

The exact pattern these changes would evoke in the CPP waveforms however is difficult to predict intuitively. Firstly, due to the baseline correction applied to compute the CPP waveform, a starting point difference would not be observed directly, but would instead lead to a difference in amplitude, with higher starting points leading to lower ERP peaks. Secondly, and as confirmed by model parameters, prior cues induced both an increased accumulation starting point for the cued response, and a decreased starting point for the non-cued response. Since the EEG signal recorded from the scalp is the sum of all underlying neural activities, the CPP presumably reflects the sum of all accumulation in a race model. It is hence unclear how opposite effects on the activity of 'correct' and 'incorrect' accumulators affects the global activity amplitude. There are a number of possible outcomes which could, at least conceptually, be considered in line with sequential sampling models. It is therefore particularly important to directly compare a signal to predictions made through model fits, in order to comment on its similarity to an accumulation process. However, it is worth bearing in mind that the relative nature of the CPP may make it an inherently less informative signal (relative to single-cell firing rates, with meaningful zero points) for the evaluation of experimental manipulations affecting the start point of accumulation.

The pattern we observed in the CPP was somewhat similar to what might be expected for just a correct accumulator. We found a clear difference in amplitude between the conditions, but no difference in slope. The waveform associated with 'incongruent' decisions showed a greater excursion than 'congruent' or 'uncued' profiles in both the stimulus and the response-locked data. The 'uncued' CPP also seemed to build up to a slightly higher plateau than the 'congruent' waveform,

although this difference was not significant in our analysis. However, it is difficult to conceive how a non-lateralised EEG signal could represent only one accumulator – only a sum, or perhaps absolute difference of accumulators makes sense. In order to evaluate to what extent this observed CPP pattern resembled the sum of accumulation processes as predicted by sequential sampling models, we simulated accumulation profiles predicted in each condition, based on the estimated parameters of the best-fitting (independent race) model. The resulting waveforms showed that all three conditions are predicted to follow a very similar trajectory, but do differ slightly in amplitude. For response-locked signals, the order in which the amplitudes differ is identical to the one described by the CPP, with the highest amplitude seen for 'incongruent' decisions, followed by 'uncued' decisions, and 'congruent' waveforms showing the lowest amplitude.

Although both the (race-model) simulated accumulation profiles and the CPP display similar patterns, it is not immediately clear what caused them. As outlined above, while we expected this pattern for the correct accumulator, summing over the accumulators would presumably cancel differences between the conditions. To aid our interpretation, we explored the accumulation profiles in more detail. First, we found that dividing correct and error trials had an impact. In Figure 4, only correct trials are averaged to match with the CPP analysis. However, in the incongruent condition in which the mean starting point is higher in the incorrect accumulator, correct trials are primarily trials in which noise has favoured the correct accumulator, such as trials in which, by chance, the cued incorrect starting point was at the lower limit of the distribution, leading to a larger baseline-to-threshold distance.

Nonetheless, averaging the accumulation profiles over all correct and error trials still resulted in a pattern qualitatively similar to the one for correct trials alone, indicating that some additional mechanism/s must help generate the observed differences. Inspecting correct and incorrect accumulation traces separately (see figure insets) confirmed that starting-point differences resulted in opposing amplitude modulations in correct and incorrect accumulators. For correct accumulation, the highest amplitude was obtained for incongruent trials, and the lowest trace in congruent trials. The reversed pattern was observed in the incorrect accumulator. However, differences between conditions were more pronounced on correct than incorrect traces, particularly in response-locked signals. We presume that this divergence arises from the accumulation rate difference between the accumulators, which implies that correct accumulation is less affected than incorrect accumulation by noise. Accordingly, incorrect traces are flatter overall and diverge less between conditions, such that differences in the correct accumulator contribute more to the summed signal.

Regardless of the computational specifics that generate differences between our conditions, the CPP and the simulated accumulation profiles display somewhat similar patterns, suggesting similar underlying mechanisms, and supporting the role of the CPP as an accumulation signal, at least when certain classes of model are used to describe the decision process. Furthermore, these findings again emphasise the importance of a direct comparison between the CPP and model predictions, as the patterns reported here are difficult to predict based on intuitive reasoning alone. However, it is also clear that our conclusion was dependent on the models we included, and on the particular model that won at a behavioural level (although we

gave our models no capacity to adjust to the neurodynamic data, a point we return to in the general discussion).

Our findings also contrast in some respects with a very recent but highly relevant CPP study, investigating the effect of a decision bias induced through manipulating the reward value of different choices under extreme time pressure (Afacan-Seref et al., 2018). Their overall conclusion is similar to ours – both studies successfully modelled RTs via sequential sampling, and showed correspondence between predicted accumulation profiles and the CPP. However, their data supported a non-standard model incorporating sensory-level dynamics (a linearly increasing accumulation rate for a constant stimulus) and a bias affecting accumulation rates rather than starting points (leading to an initially negative relative accumulation rate for a low valued but strongly evidenced choice). We did not test such a model, which may have specific relevance in their somewhat unusual experimental context. The extreme time pressure used in their experiment is likely to influence the decision dynamics, as the urgency of the choice may accelerate the accumulation in a way that is qualitatively different from the decisions made in our experiment. In any case, we make no claims that the model we have fitted and illustrated predictions from is the only (or best) possible implementation. We do, however, argue that it is a plausible choice, and one that is consistent with both the behaviour and, to some extent, the neurodynamics that we observed.

4. General Discussion

Model-based cognitive neuroscience, which combines the analysis of neural data with mathematical modelling, has gained momentum in recent years. However, the field of human perceptual decision-making has oftentimes not made full use of this approach. Here, we aimed to explore decision-related evidence accumulation in the human brain by directly comparing predictions made by different behavioural models to the dynamics of the CPP. The CPP is a centroparietal ERP component which has previously been suggested to display decision-related accumulation of evidence independent of sensory and motor processes (Kelly & O'Connell, 2013; O'Connell et al., 2012; Twomey, Kelly, & Connell, 2016). We aimed not only to explore the effect of previously untested manipulations on the CPP, but also to evaluate the resulting waveforms using sequential sampling modelling. Neural correlates of accumulation are often evaluated by deriving summary measures, such as slope of accumulation, and comparing them with expectations made with reference to sequential sampling models. However, the dynamics of even simple models are difficult to intuit. We therefore used sequential sampling models to fit the behavioural data and compared neural data to the predicted accumulation profiles based on the estimated parameters. The CPP showed a marked degree of correspondence with certain model predictions – perhaps fortuitously, the very predictions made by the models which best explained the behavioural data in each experiment.

In Experiment 1, we investigated the impact of non-stationary evidence on the CPP waveform, under the assumption that changing evidence should affect evidence accumulation dynamics. In Experiment 2, we explored the impact of decision bias on CPP patterns. We expected that decision biases induced by predictive cues would result in shifts of accumulation starting points, hence changing the baseline-to-

threshold distance. In both experiments, we observed the anticipated behavioural changes. Furthermore, sequential sampling model fits confirmed that accumulation rates were affected during evidence interruption, while starting point shifts could account for decision biases effects. It is worth noting however that when considering only behavioural data (for which free parameters in the models could be tuned to enhance goodness of fit), Experiment 1 and Experiment 2 supported two different model architectures. While a simple independent race accumulator model provided the best fit to biased decision data, the LCA model was superior in the case of non-stationary evidence, although in neither case were the differences between models entirely compelling.

We believe that this apparent discrepancy might be explained by the nature of each task manipulation, and the universal preference for simpler models. This preference is expressed in goodness-of-fit indices such as BIC or AIC by penalising models for a higher number of model parameters. Simple independent race models may therefore be favoured compared to the more complex LCA (which has a similar basic architecture but additional parameters to capture plausible physiological processes), especially in the case of fast RTs as observed in Experiment 2, in which the influence of inhibition and leakage may be limited. Conversely with longer decisions, especially associated with dynamical modulations of each accumulator's activity as in Experiment 1, both reciprocal inhibition and leakage potentially play an important role. In this case, a model incorporating these phenomena may be preferred. In other words, inhibition and leakage may always be present to some extent, but including these parameters in the decision models improves model fit only when decisions are

slow and potentially more sensitive to interactions between accumulators[11]. Indeed, in some cases, patterns of behavioural data emerge which seem to demand the inclusion of parameters capturing crosstalk between accumulators. For example, we have recently found that when up to four manual actions are instructed by a stimulus (left/right hand pinch/power grip responses), gross differences in error rates emerge based purely on the anatomical adjacency of responses (i.e. without any correlate in the stimulus; Kohl, Spieser, Forster, Bestmann, & Yarrow, 2019).

Experiments 1 and 2 were designed to be complementary, because the two types of manipulation tested two different predictions about the decision process, each realised as a different aspect of evidence accumulation. In Experiment 1, we used non-stationary evidence to affect the accumulation process. In their initial CPP description, O'Connell et al. (2012) observed that the CPP was susceptible to a change in evidence. Our results confirmed that the CPP profile is affected by a time-varying input, a necessary feature of a signal which could reflect the accumulation of evidence, and extended this result to choice-RT settings. While continuous evidence led to a gradual build-up of the CPP waveform, interrupted evidence caused a disruption in this build-up. Surprisingly, the two different interrupted conditions, one in which evidence was stopped, and one in which evidence was reversed, gave rise to very similar waveforms, even though they were associated with different behavioural patterns. Nevertheless, the pattern of the CPP closely resembled our best-fitting model predictions. In other words, our LCA model, combined with realistic

[11] Another perspective would be that these models are all describing the same fundamental model architecture, but with certain strategies requiring additional parameters, as when a non-stable environment demands the presence of leak parameters to discount the past (Kilpatrick, Holmes, Eissa & Josić, 2019).

assumptions about the origin of the CPP signal, successfully predicted the *absence* of an effect that might have been expected based on intuition alone.

In Experiment 2, we used predictive cues to manipulate decision biases. Previous research mainly suggests that biases affect the starting point of accumulation, with the resulting effect on the EEG signal requiring further clarification (Bode et al., 2012; Gao et al., 2011; Rorie et al., 2010, but see Afacan-Seref et al., 2018). We found that the CPP differed in amplitude across bias conditions. In particular, decisions in which a directional cue was incongruent with subsequent motion were associated with higher amplitudes than both decisions in which the cue was congruent with the motion and decisions in which there was no directional cue. Once again, a sequential sampling model was able to account for all behavioural data, in this case by varying the starting points across bias conditions. Furthermore, for the best-fitting independent race model, both real and model-predicted EEG signals displayed a pattern in which profiles associated with different bias conditions differed only in amplitude, with decisions with incongruent cues showing the highest amplitude, followed by uncued decisions, and trials with congruent cues showing the lowest amplitude, at least for response-locked signals. Hence here, an independent race model successfully predicted the presence of an effect that might *not* have been predicted intuitively. The simulations revealed that these differences in amplitude were not strictly the result of baseline differences, which in fact largely cancelled out on average, but were instead caused by mechanisms such as a biased representation of variability parameters in correct trials, or interactions between accumulation rate and noise parameters.

However, a problematic feature of our results emerges when looking across experiments. In our first experiment, an LCA model best fitted the behavioural data, and provided a good match to the CPP. A simpler independent race model was slightly less successful, but nonetheless showed qualitative agreement on both counts. In Experiment 2, an independent race model best fitted the behavioural data, and provided a reasonable match to the CPP. However, the more complex LCA model failed to predict the precise ordering of conditions in the CPP signal. What are we to conclude across both experiments?

When considering this disparity, we would emphasise that our approach gave the models leeway to fit the behavioural data, but not the CPP. By exploiting free parameters to capture nuances (and even noise) in the behavioural data, models may end up producing neurally unrealistic accumulation patterns. The approach we apply here has some clear strengths – by fitting only to behaviour, a model's success in predicting the associated neurodynamics becomes all the more striking, because no flexibility is provided for achieving this match (a situation somewhat akin to cross validation, but on a second form of data). However, it is only one of several ways in which model-based cognitive neuroscience might be applied (see e.g. Turner, Forstmann, Love, Palmeri, & Van Maanen, 2017, for discussion) and it is not clear whether a subsequent comparison of models on this (unfitted) neurodynamic data is a fair one. If we accept that signals like the CPP do indeed represent evidence accumulation, an important goal for future research will be to produce a consensus method for simultaneously fitting models to both RT and EEG data (cf. Turner et al.'s "integrative" approach). This is by no means trivial, because EEG data are

autocorrelated (to an uncertain extent) which greatly complicates the estimation of likelihood when matching model predictions to data.

In fact, one might argue that our observation here, that specific sequential sampling models can predict the CPP under a particular manipulation, but that a single model may not apply under different manipulations, is the norm in a fragmented literature. Thus far, where specific models have been compared to the CPP in terms of the full time-varying profile of accumulation, researchers have tended to capture only a small subset of possible manipulations. For example, a difficulty manipulation has been modelled via a drift-diffusion model (Twomey et al., 2015); a speed-accuracy trade-off has been captured via a (reconfigured) race model (Spieser et al., 2018), albeit with an unusual take on how the brain might implement this strategic adjustment; and value-based biasing under extreme time pressure has been modelled via an accelerating accumulation model (Afacan-Seref et al., 2018). Whether one views the primary question as “does the CPP represent evidence accumulation”, or, having accepted this predicate, as “which model best captures both behaviour and neurodynamics”, it seems clear that finding a single (class of) model(s) that explains the CPP across multiple experimental manipulations should be of central concern in future research.

In line with research which is increasingly emphasising the advantages of combining behavioural data, mathematical modeling, and neural dynamics (Ditterich, 2010; Forstmann et al., 2011; Mulder et al., 2014; Purcell & Palmeri, 2017), our findings also highlight the importance of combining behavioural modeling and neuroimaging

methods and directly comparing the dynamics of the neural signals and the model predictions, as neither are easily predictable based on conceptual reasoning alone. Despite the substantial similarity between the CPP and the predicted accumulation profiles observed here, there were also differences worth noting. For example, in Experiment 2, the amplitude differences between the conditions are far more pronounced in the CPP than in the model predictions even in the response-locked signals. This is likely to represent some degree of error in either our choice of models or assumptions regarding exactly how accumulators combine to form the CPP (something about which there is currently no consensus). However, it is important to note that the CPP is unlikely to ever replicate model predictions exactly for a number of reasons. Firstly, any model can, at best, be an approximation of true biological processes. A second reason for differences between the CPP and the model predictions lies in the nature of EEG recordings. EEG is measured from the scalp and can only record the sum of all electrical activity underneath each electrode, which has presumably been subject to complex filtering by intervening biological substrates. Furthermore, since the brain is constantly performing computations unrelated to accumulation, the signal-to-noise ratio is low. Most of these computations are unlikely to be time-locked to the decisions and are therefore averaged out, and the impact of conducted activation from more distal sources is reduced using the current source density transform which increases the spatial selectivity of the data. Nevertheless, noise and systematic distortions likely remain. For reasons like these, the degree of similarity between the CPP and predicted accumulation profiles derived from a class of models originally intended to predict only behaviour remains remarkable.

4.1. Conclusions

In summary, we provide further support for the role of the CPP as a neural substrate of the decision variable, but also highlight how researcher flexibility regarding which models to consider and apply might give a false degree of assurance on this front. We have shown that the CPP is sensitive to two manipulations which influence decision-making behaviour, namely non-stationary evidence and decision biases. Importantly, we fitted sequential sampling models to the behavioural data and simulated the resulting accumulation profiles. We found that the CPP waveform resembled the modelled accumulation in important ways when models were selected in a principled, but perhaps somewhat fortuitous, manner. In our opinion, the CPP probably reflects the accumulation of evidence and remains a highly plausible correlate of the decision variable. Indeed, it may now be time to move beyond mere validation of the CPP, to a point where we can instead use it as an additional metric to help differentiate competing models of speeded choice.

References

Addams, R. (1834). An account of a peculiar optical phænomenon seen after having looked at a moving body. The London and Edinburgh Philosophical Magazine and Journal of Science, 5(29), 373–374

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

Akaike, H. (1977). On entropy maximization principle. In P. R. Krishnaiah (Ed.), *Applications of Statistics* (pp. 27–41). https://doi.org/10.1007/s10955-006-9121-z

Basso, M. A., & Wurtz, R. H. (1998). Modulation of neuronal activity in superior colliculus by changes in target probability. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *18*(18), 7519–7534.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: a practicaland powerful approach to multiple testing. *Journal of the Royal Statistical Society*, *57*(1), 289–300.

Bode, S., Sewell, D. K., Lilburn, S., Forte, J. D., Smith, P. L., & Stahl, J. (2012). Predicting Perceptual Decision Biases from Early Brain Activity. *Journal of Neuroscience*, *32*(36), 12488–12498. https://doi.org/10.1523/JNEUROSCI.1708-12.2012

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*, 433–436. https://doi.org/10.1163/156856897X00357

Bronfman, Z. Z., Brezis, N., & Usher, M. (2016). Non-monotonic Temporal-Weighting Indicates a Dynamically Modulated Evidence-Integration Mechanism. *PLoS*

Computational Biology, *12*(2), 1–21. https://doi.org/10.1371/journal.pcbi.1004667

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153–178. https://doi.org/10.1016/j.cogpsych.2007.12.002

de Lange, F. P., Rahnev, D. A., Donner, T. H., & Lau, H. (2013). Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *Journal of Neuroscience*, *33*(4), 1400–1410. https://doi.org/10.1523/JNEUROSCI.1094-12.2013

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

Ditterich, J. (2010). A comparison between mechanisms of multi-alternative perceptual decision making: Ability to explain human behavior, predictions for neurophysiology, and relationship with decision theory. *Frontiers in Neuroscience*, *4*(NOV), 1–24. https://doi.org/10.3389/fnins.2010.00184

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

Dorris, M. C., & Munoz, D. P. (1998). Saccadic probability influences motor preparation signals and time to saccadic initiation. *The Journal of Neuroscience*, *18*(17), 7015–7026.

Forstmann, B. U., Anwander, A., Schafer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., … Turner, R. (2010). Cortico-striatal connections predict

control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences*, *107*(36), 15916–15920. https://doi.org/10.1073/pnas.1004932107

Forstmann, B U, Ratcliff, R., & Wagenmakers, E.-J. (2016). Sequential Sampling Models in Cognitive Neuroscience: Advantages, Applications, and Extensions. *Annual Review of Psychology*, *67*, 641–666. https://doi.org/10.1146/annurev-psych-122414-033645

Forstmann, Birte U, Wagenmakers, E.-J., Eichele, T., Brown, S., & Serences, J. T. (2011). Reciprocal Relations Between Cognitive Neuroscience and Cognitive Models : Opposites Attract ? *Trends in Cognitive Sciences*, *15*(6), 272–279. https://doi.org/10.1016/j.tics.2011.04.002.Reciprocal

Gao, J., Tortell, R., & McClelland, J. L. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS ONE*, *6*(3). https://doi.org/10.1371/journal.pone.0016749

Gold, J. I., & Shadlen, M. N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, *404*(6776), 390–394. https://doi.org/10.1038/35006062

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

Hanes, D. P., & Schall, J. D. (1996). Neural Control of Voluntary Movement Initiation. *Science*, *274*(5286), 427–430.

Hanks, T. D., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. *ELife*, *2014*(3), 1–17. https://doi.org/10.7554/eLife.02260

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

Heathcote, A, Brown, S., & Mewhort, D. J. K. (2002). Quantile maximum likelihood estimation of response time distributions. *Psychonomic Bulletin and Review*, *9*(2), 1–31. https://doi.org/10.3758/BF03196299

Heathcote, Andrew, & Love, J. (2012). Linear deterministic accumulator models of simple choice. *Frontiers in Psychology*, *3*(AUG), 1–19. https://doi.org/10.3389/fpsyg.2012.00292

Heekeren, H. R., Marrett, S., Bandettini, P. A., & Ungerleider, L. G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, *431*(7010), 859–862. https://doi.org/10.1038/nature02966

Holmes, W. R., Trueblood, J. S., & Heathcote, A. (2016). A new framework for modeling decisions about changing information: The Piecewise Linear Ballistic Accumulator model. *Cognitive Psychology*, *85*, 1–29. https://doi.org/10.1016/j.cogpsych.2015.11.002

Huk, A. C., & Shadlen, M. N. (2005). Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. *J. Neurosci.*, *25*(45), 10420–10436. https://doi.org/10.1523/JNEUROSCI.4684-04.2005

Kayser, J., & Tenke, C. E. (2006). Principal components analysis of Laplacian waveforms as a generic method for identifying ERP generator patterns: II. Adequacy of low-density estimates. *Clinical Neurophysiology*, *117*(2), 369–380. https://doi.org/10.1016/j.clinph.2005.08.033

Kelly, S. P., & O'Connell, R. G. (2013). Internal and external influences on the rate of

sensory evidence accumulation in the human brain. *J Neurosci*, *33*(50), 19434–19441. https://doi.org/10.1523/JNEUROSCI.3355-13.2013

Kilpatrick, Z. P., Holmes, W. R., Eissa, T. L., & Josić, K. (2019). Optimal models of decision-making in dynamic environments. Current opinion in neurobiology, 58, 54-60.

Kleiner, M., Brainard, D. H., Pelli, D. G., Broussard, C., Wolf, T., & Niehorster, D. (2007). What's new in Psychtoolbox-3? *Perception*, *36*, S14. https://doi.org/10.1068/v070821

Kohl, C., Spieser, L., Forster, B., Bestmann, S., & Yarrow, K. (2018). The Neurodynamic Decision Variable in Human Multialternative Perceptual Choice. *Journal of Cognitive Neuroscience*, 1–16. https://doi.org/10.1162/jocn

Krakauer, J. W., Ghazanfar, A. A., Gomez-Marin, A., MacIver, M. A., & Poeppel, D. (2017). Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron*, *93*(3), 480–490. https://doi.org/10.1016/j.neuron.2016.12.041

Luce, R. D. (1986). *Response Times: Their Role in Inferring Elementary Mental Organization*. New Yorj: Oxford University Press.

Marr, D. (2010). *Vision*. Cambridge: MIT Press.

Meindertsma, T., Kloosterman, N. A., Nolte, G., Engel, A. K., & Donner, T. H. (2017). Multiple Transient Signals in Human Visual Cortex Associated with an Elementary Decision. *The Journal of Neuroscience*, *37*(23), 5744–5757. https://doi.org/10.1523/JNEUROSCI.3835-16.2017

Milosavljevic, M., Malmaud, J., & Huth, A. (2010). The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. *Judgement and Decision Making*, *5*(6), 437–449. https://doi.org/10.2139/ssrn.1901533

Mulder, M. J., van Maanen, L., & Forstmann, B. U. (2014). Perceptual decision neurosciences - a model-based review. *Neuroscience*, *277*, 872–884. https://doi.org/10.1016/j.neuroscience.2014.07.031

Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience*, *32*(7), 2335–2343. https://doi.org/10.1523/JNEUROSCI.4156-11.2012

Noorbaloochi, S., Sharon, D., & McClelland, J. L. (2015). Payoff Information Biases a Fast Guess Process in Perceptual Decision Making under Deadline Pressure: Evidence from Behavior, Evoked Potentials, and Quantitative Model Comparison. *J. Neurosci.*, *35*(31), 10989–11011. https://doi.org/10.1523/JNEUROSCI.0017-15.2015

Nunes, L. F., & Gurney, K. (2016). Multi-alternative decision-making with non-stationary inputs. *Royal Society Open Science*, *3*(8). https://doi.org/10.1098/rsos.160376

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

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spatial Vision*, Vol. 10, pp. 437–442. https://doi.org/10.1163/156856897X00366

Philiastides, M. G., Heekeren, H. R., & Sajda, P. (2014). Human Scalp Potentials Reflect a Mixture of Decision-Related Signals during Perceptual Choices. *Journal of Neuroscience*, *34*(50), 16877–16889. https://doi.org/10.1523/JNEUROSCI.3012-14.2014

Philiastides, M. G., Ratcliff, R., & Sajda, P. (2006). Neural Representation of Task Difficulty and Decision Making during Perceptual Categorization: A Timing Diagram. *Journal of Neuroscience*, *26*(35), 8965–8975. https://doi.org/10.1523/JNEUROSCI.1655-06.2006

Philiastides, Marios G., & Sajda, P. (2006). Temporal characterization of the neural correlates of perceptual decision making in the human brain. *Cerebral Cortex*, *16*(4), 509–518. https://doi.org/10.1093/cercor/bhi130

Pisauro, M. A., Fouragnan, E., Retzler, C., & Philiastides, M. G. (2017). Neural correlates of evidence accumulation during value-based decisions revealed via simultaneous EEG-fMRI. *Nature Communications*, *8*(May), 15808. https://doi.org/10.1038/ncomms15808

Price, K. V., Storn, R. M., & Jouni, L. A. (2005). *Differential Evolution: A Practical Approach to Global Optimization*. https://doi.org/10.1038/155531c0

Purcell, B. A., Heitz, R. P., Cohen, J. Y., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2010). Neurally constrained modeling of perceptual decision making. *Psychological Review*, *117*(4), 1113–1143. https://doi.org/10.1037/a0020311.Neurally

Purcell, B. A., & Palmeri, T. J. (2017). Relating accumulator model parameters and neural dynamics. *Journal of Mathematical Psychology*, *76*, 156–171. https://doi.org/10.1016/j.jmp.2016.07.001

Purcell, B. A., Schall, J. D., Logan, G. D., & Palmeri, T. J. (2012). From Salience to Saccades: Multiple-Alternative Gated Stochastic Accumulator Model of Visual Search. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(10), 3433–3446. https://doi.org/10.1523/JNEUROSCI.4622-11.2012

Rao, V., DeAngelis, G. C., & Snyder, L. H. (2012). Neural correlates of prior expectations of motion in the lateral intraparietal and middle temporal areas. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *32*(29), 10063–10074. https://doi.org/10.1523/JNEUROSCI.5948-11.2012

Ratcliff, R. (2002). A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychonomic Bulletin & Review*, *9*(2), 278–291. https://doi.org/10.3758/BF03196283

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. https://doi.org/10.1162/neco.2008.12-06-420

Ratcliff, R., Philiastides, M. G., & Sajda, P. (2009). Quality of evidence for perceptual decision making is indexed by trial-to-trial variability of the EEG. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(16), 6539–6544. https://doi.org/10.1073/pnas.0812589106

Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, *9*(5), 347–356.

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333–367. https://doi.org/10.1016/j.pestbp.2011.02.012.Investigations

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

Ratcliff, R., Thapar, A., College, B. M., & Mckoon, G. (1992). Effects of aging and IQ

on item and associative memory. *Journal of Experimental Psychology*, *140*(3), 464–487. https://doi.org/10.1037/a0023810

Ratcliff, R., Thapar, A., & McKoon, G. (2010). Individual differences, aging, and IQ in two-choice tasks. *Cognitive Psychology*, *60*(3), 127–157. https://doi.org/10.1016/j.cogpsych.2009.09.001

Rorie, A. E., Gao, J., McClelland, J. L., & Newsome, W. T. (2010). Integration of sensory and reward information during perceptual decision-making in Lateral Intraparietal Cortex (LIP) of the macaque monkey. *PLoS ONE*, *5*(2). https://doi.org/10.1371/journal.pone.0009308

Schall, J. D. (2002). The neural selection and control of saccades by the frontal eye field. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *357*(1424), 1073–1082. https://doi.org/10.1098/rstb.2002.1098

Schwarz, G. E. (1978). Estimating the dimension of a model. *Annals of Statistics*, *6*(2), 461–464. https://doi.org/:10.1214/aos/1176344136.MR468014

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: seeing and deciding. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(January), 628–633. https://doi.org/10.1073/pnas.93.2.628

Siegel, M., Engel, A. K., & Donner, T. H. (2011). Cortical Network Dynamics of Perceptual Decision-Making in the Human Brain. *Frontiers in Human Neuroscience*, *5*(February), 1–12. https://doi.org/10.3389/fnhum.2011.00021

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. https://doi.org/10.1016/j.tins.2004.01.006

Spaniol, J., Voss, A., Bowen, H. J., & Grady, C. L. (2011). Motivational incentives

modulate age differences in visual perception. *Psychology and Aging*, *26*(4), 932–939. https://doi.org/10.1037/a0023297

Spieser, L., Kohl, C., Forster, B., Bestmann, S., Yarrow, K., & Kohl, C. (2018). *Neurodynamic Evidence Supports a Forced- Excursion Model of Decision-Making under Speed / Accuracy Instructions Abbreviated title : Neural signals support forced-excursion SAT model Correspondence :* https://doi.org/10.1523/ENEURO.0159-18.2018

Summerfield, C., & de Lange, F. P. (2014). Expectation in perceptual decision making: neural and computational mechanisms. *Nature Reviews Neuroscience*, *15*(11), 745–756. https://doi.org/10.1038/nrn3838

Summerfield, C., & Koechlin, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience*, *4*(November), 208. https://doi.org/10.3389/fnhum.2010.00208

Teodorescu, A. R., & Usher, M. (2013). Disentangling decision models: from independence to competition. *Psychological Review*, *120*(1), 1–38. https://doi.org/10.1037/a0030776

Tsetsos, K., Gao, J., McClelland, J. L., & Usher, M. (2012). Using time-varying evidence to test models of decision dynamics: Bounded diffusion vs. The leaky competing accumulator model. *Frontiers in Neuroscience*, *6*(JUN), 1–17. https://doi.org/10.3389/fnins.2012.00079

Tsetsos, K., Usher, M., & McClelland, J. L. (2011). Testing multi-alternative decision models with non-stationary evidence. *Frontiers in Neuroscience*, *5*(MAY), 1–18. https://doi.org/10.3389/fnins.2011.00063

Turner, B. M., Rodriguez, C. A., Norcia, T. M., McClure, S. M., & Steyvers, M. (2016). Why more is better: Simultaneous modeling of EEG, fMRI, and

behavioral data. *NeuroImage*, *128*, 96–115. https://doi.org/10.1016/j.neuroimage.2015.12.030

Twomey, D. M., Kelly, S. P., & Connell, R. G. O. (2016). Abstract and Effector-Selective Decision Signals Exhibit Qualitatively Distinct Dynamics before Delayed Perceptual Reports. *The Journal of Neuroscience*, *36*(28), 7346–7352. https://doi.org/10.1523/JNEUROSCI.4162-15.2016

Twomey, D. M., Murphy, P. R., Kelly, S. P., & O'Connell, R. G. (2015). The classic P300 encodes a build-to-threshold decision variable. *European Journal of Neuroscience*, *42*(1), 1636–1643. https://doi.org/10.1111/ejn.12936

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: The leaky, competing accumulator model. *Psychological Review*, *108*(3), 550–592. https://doi.org/10.1037/0033-295X.108.3.550

van Ravenzwaaij, D., Provost, A., & Brown, S. D. (2017). A confirmatory approach for integrating neural and behavioral data into a single model. *Journal of Mathematical Psychology*, *76*, 131–141. Retrieved from http://ees.elsevier.com/jmp/reviewer_current.asp?currentpage=1%0Apapers3://publication/uuid/180B3732-94AD-47D9-920C-CB0E2822FF88

van Vugt, M. K., Simen, P., Nystrom, L. E., Holmes, P., & Cohen, J. D. (2012). EEG oscillations reveal neural correlates of evidence accumulation. *Frontiers in Neuroscience*, (JULY), 1–13. https://doi.org/10.3389/fnins.2012.00106

Voss, A., Nagler, M., & Lerche, V. (2013). Diffusion models in experimental psychology: A practical introduction. *Experimental Psychology*, *60*(6), 385–402. https://doi.org/10.1027/1618-3169/a000218

Watson, A. P., & Pelli, D. G. (1983). QUEST: A Bayesian adaptive psychometric method ANDREW. *Perception & Psychophysics*, *33*(2), 113–120.

Yarrow, K., Minaei, S., & Arnold, D. H. (2015). A model-based comparison of three theories of audiovisual temporal recalibration. *Cognitive Psychology*, *83*, 54–76. https://doi.org/10.1016/j.cogpsych.2015.10.002

Zhou, X., Wong-Lin, K., & Philip, H. (2009). Time-varying perturbations can distinguish among integrate-to-threshold models for perceptual decision making in reaction time tasks. *Neural Computation*, *21*(8), 2336–2362. https://doi.org/10.1162/neco.2009.07-08-817

Appendix A

Parameter identifiability issues have been reported in the LCA model (Miletić et al., 2017). Hence, we conducted a recovery study to assess the accuracy of parameter estimation in Experiment 1. The mean parameter estimates of the chosen LCA model (Model 5, LCA-symmetric with no delay) are displayed in Table A1[12]. Based on this model, we simulated 20 RT datasets with all 3 interruption conditions and 2 difficulty levels. We simulated 160 trials in each condition, leading to 960 trials in total (i.e., corresponding to the size of one participant's RT data). Parameters values for each of the 20 simulated datasets were drawn from a uniform distribution around mean empirical values.

Table A1: Mean estimated parameter values for the chosen model (Model 5), note that the response threshold A was set to 1 as a scaling parameter.

Model 5: Parameters			
Decision threshold (A)			1
Accumulation rate (v)	easy	correct	6.0154
		incorrect	1.4110
	hard	correct	5.0199
		incorrect	1.5039
Leakage (k)			5.2706
Inhibition (β)			65.7646
Non-decision time (T_{er})			0.3574
Non-decision time interval (S_{Ter})			0.2763

[12] Parameter values are only comparable across studies if the same scaling parameter is used. Here we fixed the decision threshold but let noise vary, yielding a larger than typical Gaussian noise SD (and thus amplified values for many parameters).

Figure A1

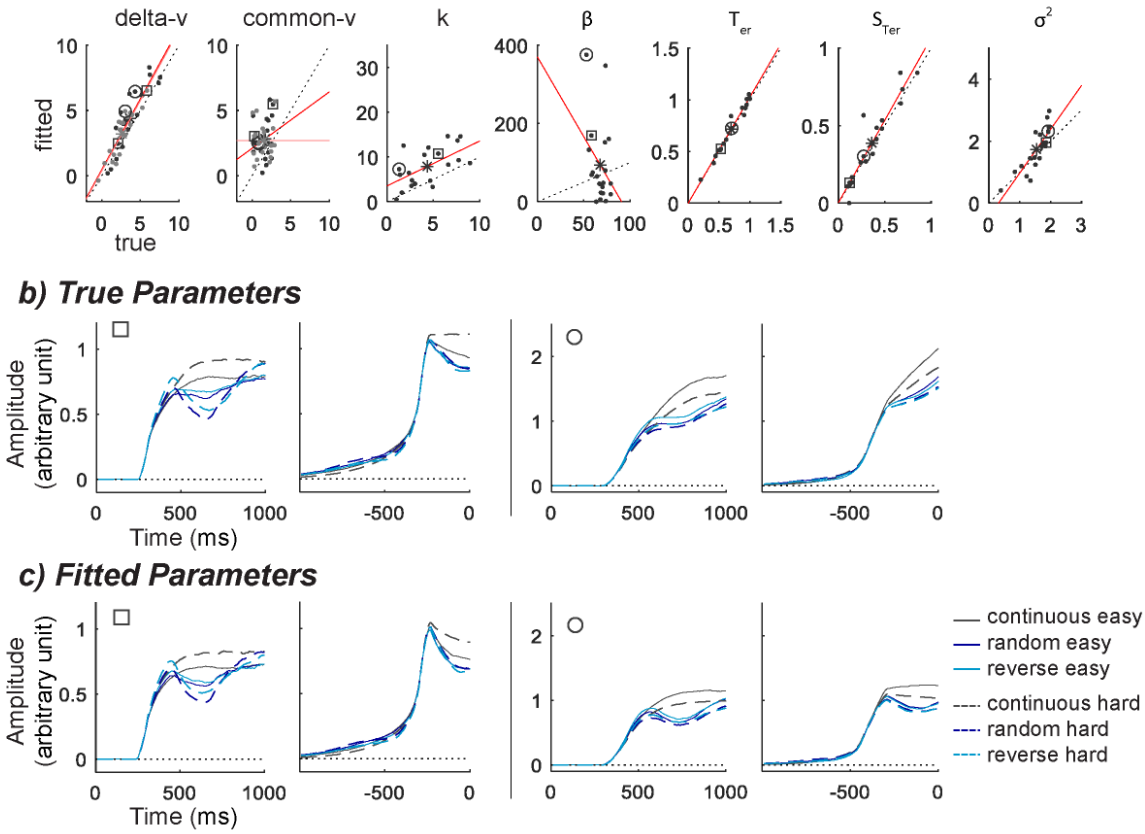

Figure A1: a) Parameter recovery: fitted parameter values as a function of true values, for 20 simulations of individual RTs. Dots show the 20 individual fit values and asterisks show mean fitted value as a function of mean true value. Dotted lines show ideal recovery of fitted from true parameters. Red lines show linear regressions between true and fitted values. Rate parameters are decomposed in delta-v and common-v (see details in text), and both easy (dark) and hard (light) conditions are shown. Circles and squares identify parameter sets used to compute predictions in b and c. b) and c) CPP predictions for 2 sets of parameters, computed based on true values (b) and fitted values (c). Both parameters sets are identified in a) by circles (predictions on left panel) and squares (predictions on right panel). Stimulus-locked (left) and response-locked (right) predictions are shown.

Figure A1 shows the obtained fitted values as a function of true values for each parameter. Note that accumulation rates are decomposed into *delta-v* and *common-*

v, corresponding respectively to the difference and the common components of correct and incorrect rates (i.e., *delta-v* equals v_{corr} minus v_{inc}, and *common-v* equals v_{inc}). Ideally, values recovered from the fit would equal the true parameters, falling on the black dotted line. Red lines show best-fitting linear regressions between true and fitted parameter values. To assess the accuracy of parameter estimation at a group level, we also represented the average of fitted values as a function of the mean true value. Consistent with a previous report (Miletić et al., 2017), we observed good recovery for *delta-v*, T_{er} and σ^2, as well as S_{Ter}, and poor recovery for *common-v*, *k* and β parameters. At the group level, however, the mean fitted parameter values were still a good estimation of mean true values (asterisks in Figure A1a).

Finally, and critically, in order to assess the impact of parameter estimation accuracy on derived CPP predictions, we computed predictions based on true and fitted parameters values. Predictions are shown for two sets of parameters in Figure A1. They have been selected as being both representative of our general findings (across all 20 simulations) and illustrative of cases where recovered parameters appear to have traded off, and thus differ from true parameters. As can be seen, the global pattern is retrieved in fitted parameter predictions, even in those cases where *common-v* and *beta* parameters were not estimated accurately.

Appendix B

In both experiments, many of the models performed somewhat similarly. For completeness, the behavioural fits of all models are displayed in Figures B1 (Experiment 1, see Figure 1 b), and B2 (Experiment 2, see Figure 3 b).

Figure B1

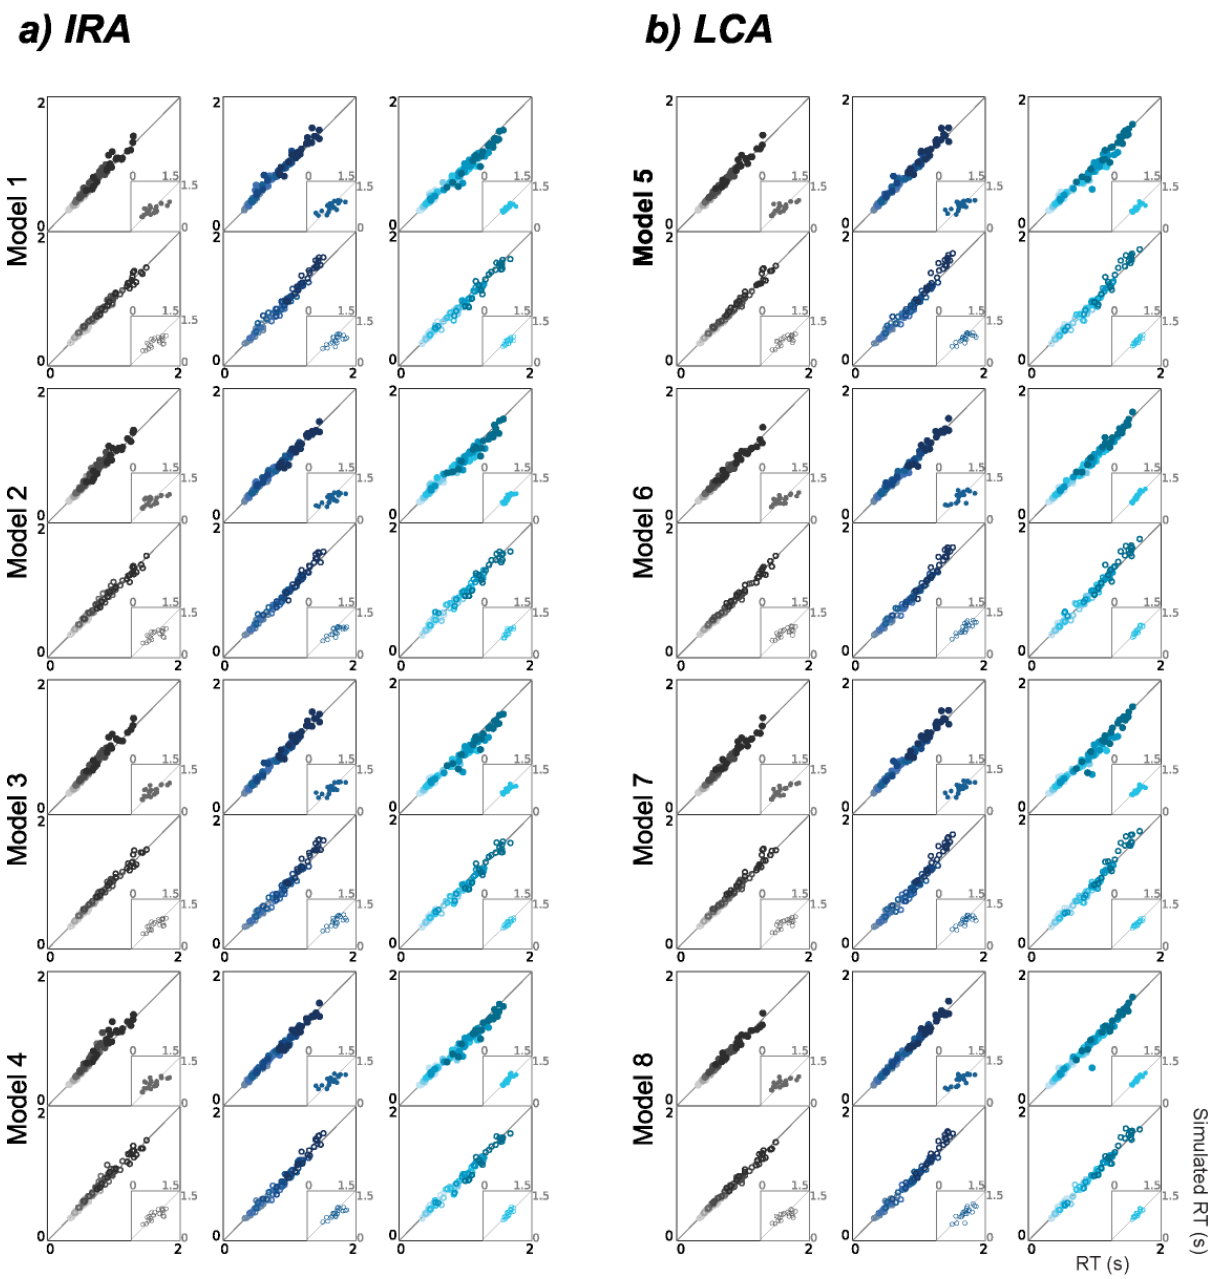

Figure B1: Experiment 1, Behavioural fits for all models: RT quantiles from behavioural data (x-axis) and simulations (y-axis) in seconds for each independent race (IRA, 1 to 4) and LCA (5 to 8) model for easy (filled circles, top rows) and hard (empty circles, bottom rows) decisions. Small inserted panels show observed and simulated RT medians for error trials.

Figure B2

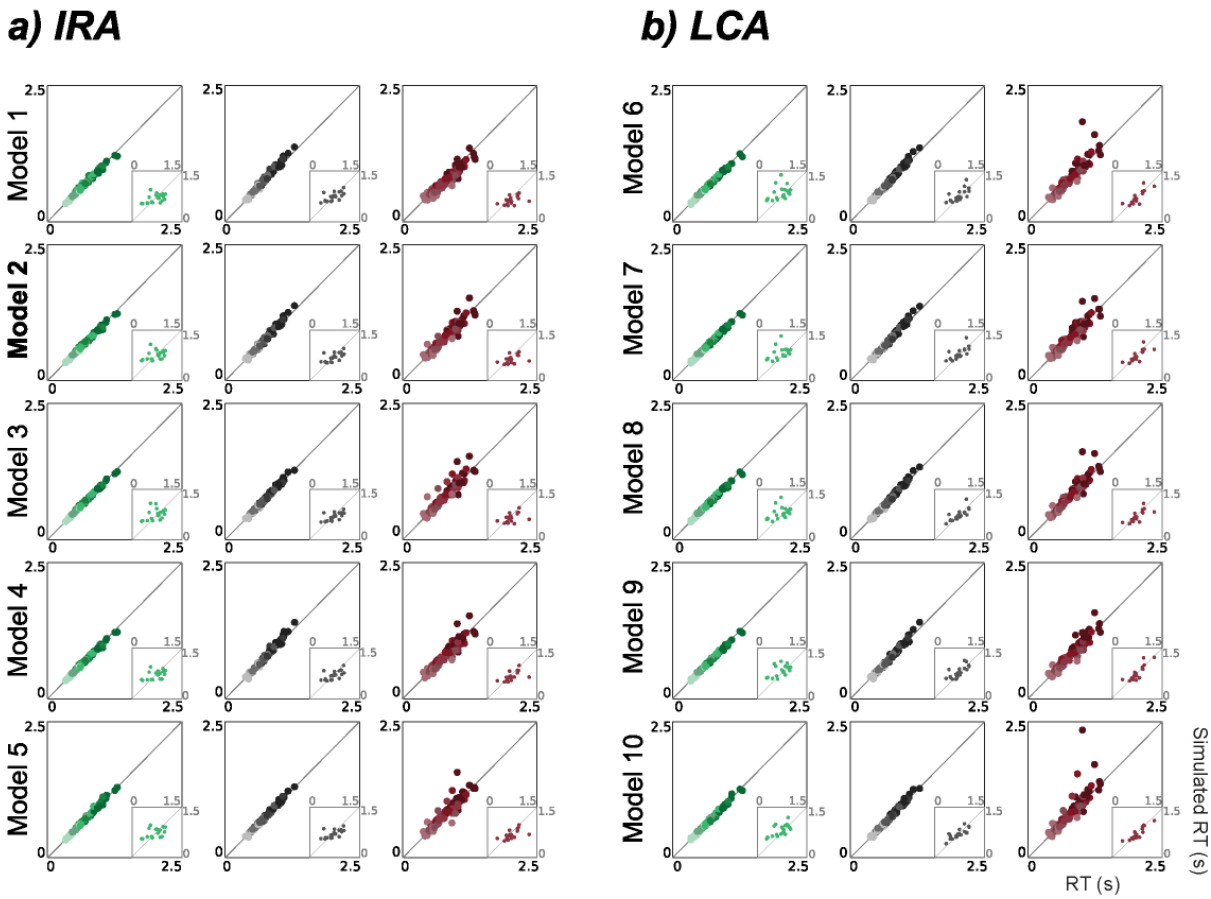

Figure B2: Experiment 2, Behavioural fits for all models: RT quantiles from behavioural data (x-axis) and simulations (y-axis) in seconds for each race (IRA, 1 to 5) and LCA (6 to 10) model. Small inserted panels show observed and simulated RT medians for error trials.