Trinity College Dublin
Coláiste na Tríonóide, Baile Átha Cliath
The University of Dublin

An Exploration of the Applications of Neurally-Informed Models of Perceptual Decision Making

Cian Judd,
School of Psychology,
The University of Dublin, Trinity College.

A dissertation submitted for the degree of Doctor of Philosophy to the School of Psychology at the University of Dublin, Trinity College, Ireland.

2023

Declaration

I declare that this thesis has not been submitted as an exercise for a degree at this or any other university and it is entirely my own work.

I agree to deposit this thesis in the University's open access institutional repository or allow the Library to do so on my behalf, subject to Irish Copyright Legislation and Trinity College Library conditions of use and acknowledgement.

I consent to the examiner retaining a copy of the thesis beyond the examining period, should they so wish (EU GDPR May 2018).

_____

Signed,
Cian Judd,
19th December, 2022

Summary

Every second of every day we are bombarded by a plethora of sensory information, which we must rapidly and accurately integrate to make timely decisions. Perceptual-decision making is the cornerstone of cognition, with research in this field attempting to account for how the mind translates these sensory signals into action. Computational models of the perceptual-decision making such as the drift-diffusion model allow us to utilise behavioural information to tap into the latent processes which underpin these decisions (Ratcliff & McKoon, 2008). However, as mathematical models are abstract, they may not accurately represent how decisions are actually formed. This leads to issues such as model mimicry, where models can provide equally strong fits to the data by employing a combination of parameters which make fundamentally different predictions for how the decision unfolds in the brain (O'Connell et al., 2018). Furthermore, these models must assume a scaling parameter which remains consistent across individuals, groups and conditions, even in the face of evidence which shows substantial inter-group and inter-condition variability in these scaling parameters (Dosher & Lu, 2017; Manning et al., 2015; Tibber et al., 2015). Electrophysiological data can be used to address these issues by providing a set of signals that delineate how perceptual decision-making unfolds in the brain (Devine et al., 2019; O'Connell, Dockree, et al., 2012; O'Connell & Kelly, 2021; Twomey et al., 2015). Testing the degree to which a model can account for these signals offers a new means of arbitrating between competing variants (Kelly et al., 2021). Furthermore, features of neural data can be used to constrain parameters (McGovern et al., 2018; Corbett et al., in press), acting as data-driven scaling parameters and allowing for the exploration of increasingly complex models. This nascent field of neurally-informed modelling aims to synthesise these two distinct sources of information, allowing us to overcome some limitations of computational modelling efforts alone. Through this we can create data-driven, biologically plausible models of the perceptual decision-making process, providing new insights into our core perceptual processes. Exploring and investigating these potential applications is the aim of this body of work.

Chapter One provides a review of the relevant literature. It outlines the key benefits of computational modelling and electrophysiological research for the study of decision-making while identifying key issues that can be overcome through neurally-informed modelling.

Chapter Two investigates the benefits of neurally-informed modelling to improve between-group comparison. Examining age differences in a well-established Speed-Accuracy Tradeoff paradigm, it utilised neural data to inform choices regarding which parameters to include in a drift-diffusion model (DDM) and to directly estimate certain parameters. This furnished new insight into the well-studied paradigm, indicating that older adults may not show a more cautious decision style (Forstmann et al., 2011; Starns & Ratcliff, 2010), but rather a more consistent one, with smaller boundary adjustments in response to speed pressure manipulations, and a more stable boundary elevation throughout the time course of each decision. This study highlights the benefits of neurally-informed modelling in providing novel insights into well-established phenomena.

Chapter Three used neurally-informed modelling to explore the role of within-trial noise reductions as a contributing driver of perceptual learning. Psychophysical work has long identified reductions in internal perceptual noise as a cause of improved behaviour over time on visual-learning tasks (Dosher & Lu, 2017). To date, analysis of this within a DDM framework has been unattainable, as within-trial noise is traditionally fixed as a scaling parameter that remains constant across individuals and time. This study administered a fast behavioural estimate of internal noise, previously validated in the psychophysics literature (Tibber et al., 2014). Constraining boundaries to remain fixed across sessions and allowing noise to vary, we showed consistency between this behavioural measure and the freely-estimated model parameters. Using both this behavioural metric and pre-stimulus mu/beta levels to inform and constrain the model, our neurally-informed and constrained models pointed to reductions in within-trial noise with learning. We also showed the potential learning of task demands as a contributory factor to improved performance over time. This has not been studied in DDM work to date and was enabled through neurally-informed modelling.

Chapter Four aimed to explore methods to improve the reliability of neural signals for neurally-informed modelling across groups. Significant between-group differences in signals are often seen, however, we cannot always be certain that this reflects true behaviourally-relevant effects or are the result of irrelevant features such as overlapping signals, or differences in signal conduction which are borne out differently across groups. A prime example of this is the P3b, analogous to the CPP, which exhibits an anterior frontalisation and less positive amplitudes as we age (Kuruvilla-Mathew et al., 2022; O'Connell, Balsters, et al., 2012; Polich, 1997b; van Dinteren et al., 2014a, 2014b). However, the nature of the tasks used to study these effects may contribute to overlapping sensory-evoked potentials or other ERP components (O'Connell et al. 2012; West et al., 2010). This study examined frontalisation across a variety of tasks, using ICA to extract an evidence-dependent CPP component which showed no frontalisation, but a lower amplitude in older adults. Through this, the study presents a novel means for extracting neural indices of the decision-making process which may serve to improve neurally-informed modelling efforts

Chapter Five concludes with a general overview of this body of work. It outlines the key contributions of these experiments. It then discusses some potential limitations and avenues for further research.

Acknowledgements

In many ways, I am not sure where to begin, and even then, I am not sure where I would end. To say I could never have achieved this submission without the unwavering support of an entire community would be an understatement, but you have to start somewhere.

Redmond, your support and guidance throughout this process has been invaluable. Since I began my journey into research nearly eight years ago as a research assistant, you have shown me endless encouragement and patience, showing me not only what it means to be a great researcher, but how this can be achieved with kindness and compassion.

Elaine, working closely with you has inspired me. You have given me endless advice, answered all of my modelling questions and allowed me to develop skills I never thought attainable.

To the Irish Research Council for supporting this research, my thesis committee, the TCIN community, Amelia, Tina, and all of my participants who gave up their time to make this happen. Thank you all.

O'Connell Lab members past and present. Having spent nearly eight years in this narrow office, I've gotten to know you each so well. We have shared too many memories to count, and too many smiles that will stay in my heart forever. Kev, Wouter and Jess. You helped me through my early years and became close friends. I hope to continue our friendship when I join you all in the big bad world. Hannah, Jade, Dasha, John, Pat and Katerina. You have put up with my mumbled singing in the office and provided some welcome relief to the long days in TCIN. Harvey, you have become a true friend, a friendship I know I will treasure forever. You have helped me through it all, and for that, I am eternally grateful.

My friends. Anna, for never-faltering support, listening to my five-minute voice notes and indulging in late-night visits to Italian restaurants. You are my rock and always

will be. Laura, I can see your smile and hear your laughter now. Thank you for McFlurry trips in the freezing-cold Punto. Ashleigh, although you are on the other side of the world now, I wouldn't be who I am today without you. Your outlook on life and positive attitude has always been an inspiration. Eimear, a sentence or two could never describe how much you mean to me, so I will just come around for tea and try to tell you in person (although I know we would both start laughing before I got a word out). You each mean the world to me. To have one good friend is a blessing. To have more has been the greatest gift of all.

My family. Al, you have lit candles to St. Joseph the Worker for me since primary school, and they seem to have gotten me through. You have always been my number one supporter, and I love you. Judy, you instilled in me a passion for education that I know has taken me down this path. Thank you for guiding me in my earliest days, and I know you are with me still. Dad, a friend and constant support, always willing to receive a late-night call-in, and always with the kettle on. Sineád, with you I share my most resounding laughter. Your smile brightens my day, and I am thankful to have a sister like you.

Mam. I suppose some might see it as trite to end my acknowledgements with a dedication to you, but I am not sure any other way would feel fitting. You have shown me how to approach life with humour, honour and above all, love. You believed in me long before I believed in myself. I can never begin to repay you for all that you have done for me, but I will try, every day, for as long as I live.

Table of Contents

Chapter One:

An Introduction to Neurally-Informed Modelling of Perceptual Decision-Making

1.1: Introduction

To study how the mind makes simple perceptual decisions is to investigate human experience at its most fundamental level. Each day, and indeed each instant, the brain is tasked with integrating a diverse array of sensory inputs to reach decisions about which actions to make in response to a rich and constantly changing environment. Perceptual decision-making is the bedrock upon which all subsequent cognitive processes are built. Without the brain's accumulation and evaluation of sensory inputs, the complex cognitive operations characteristic of our everyday life would be impossible. This includes the types of cognitive operations more readily recognised as 'decisions' in the standard sense - i.e., value-based choices including what to eat for dinner, and social choices including choosing a partner (Shadlen & Kiani, 2013).

While the area of perceptual decision-making has been studied extensively through a range of disciplines and methodologies, the complex mechanisms underpinning these operations have yet to be fully understood (Shadlen & Kiani, 2013). To date, attempts to study the decision-making process have focused on either behavioural or neurophysiological data, whilst using computational models of reaction time and accuracy to extract latent features of the processes that translate sensation into action (e.g., Brown & Heathcote, 2008; Ratcliff et al., 2016; Ratcliff & McKoon, 2008; Starns & Ratcliff, 2010). Analysis of neurophysiological data has in turn provided a range of neural indices which trace unique stages of the decision-making process and reflect parameters of computational models (Fischer et al., 2018; Kelly et al., 2021; Nunez et al., 2019; O'Connell, Dockree, et al., 2012; Twomey et al., 2015). Although each of these approaches have provided valuable insight and been successful in their own right, thus far the field still lacks a unified and consistent account of perceptual decision-making phenomena which can account for neural as well as behavioural observations. The mission to provide this unified account has also been hampered by the sets of unique drawbacks associated with each respective approach that will be discussed in the coming sections - drawbacks which are further compounded when we attempt to investigate group differences.

For example, oftentimes, competing model variants can provide similarly well-fitting simulations of the data, however, they achieve this by invoking fundamentally different accounts for how behaviour is unfolding in the brain (Malhotra et al., 2018; Ratcliff et al., 2016). Neural data provides an array of direct readouts of each stage of the decision process as it happens in the brain, allowing us to test the assumptions of the model against neural data, and arbitrate between competing accounts (Kelly et al., 2021; McGovern et al., 2018). We can also utilise this neural data to inform our model construction in cases where it provides evidence for the inclusion of novel parameters to represent constructs not traditionally allowed for in standard models (e.g., urgency and visual encoding time, Ghaderi-Kangavari et al., 2022; Hanks et al., 2014). Furthermore, we can constrain parameter values based on the neural data, allowing us to investigate increasingly complex models and thus more acutely investigate the nuanced processes which underpin decision-making (Kelly et al., 2021). In combining neurophysiological and computational modelling accounts of decision-making, neurally-informed modelling can reveal novel insights into perceptual phenomena and create more biologically grounded and neurophysiologically plausible models of the decision-making process.

The purpose of the present work is to examine the benefits of this neurally-informed modelling. This chapter begins with a review of current sequential sampling models of decision-making, highlighting their benefits and applications, and exploring the electrophysiological work which supports them. Methodological drawbacks of current computational models will be discussed, in particular, aspects which limit the capacity of sequential sampling models of decision-making to enable efficacious between-group comparisons. The emerging technique of neurally-informed modelling and its potential to assist in overcoming some of these limitations will then be examined and discussed. This technique is not without its own issues, and potential complications and drawbacks in the use of electrophysiology will also be addressed. From this review of the current literature, a rationale for the empirical studies contained within this thesis will be provided, with the aim of highlighting this work's potential to address some core fundamental issues in the extant literature. Collectively, this body of work aims to

advance neurally-informed modelling methodology and our understanding of perceptual decision-making in general.

1.2: Computational Models of Perceptual Decision-Making

1.2.1: Sequential Sampling Models of Decision-Making

Traditional experimental protocols for studying decision-making extract two primary sources of data: reaction time (RT) and choice accuracy. In psychophysical models such as Signal Detection Theory, complex task manipulations are applied, with psychometric functions fitted to accuracy data under differing conditions, giving us estimates of an observer's sensitivity to stimuli or any response biases they may have (Lynn & Barrett, 2014). However, in reality, decisions almost always operate with respect to time. When driving we continuously sample the environment around us to allow us to make correct and safe decisions. However, at a certain point, an observer must decide whether enough evidence has been accumulated that they are satisfied to commit to a particular course of action (Forstmann et al., 2016). If a possible obstruction appears on the road, we cannot continue to sample evidence indefinitely for as long as it takes us to gain full certainty of its presence or absence, because, by the time this perceptual decision has been made, it may be too late to avoid it. Therefore, in situations like driving, and a host of other time-contingent processes, we must make rapid decisions in which speed is prioritised at the expense of accuracy. However, if the decision in question requires more precision, for example when a surgeon is performing a difficult and risky operation, they will likely apply a slower and more steady approach, resulting in decisions which, while slower, are ultimately more accurate. These examples illustrate two important features of perceptual decision-making which are addressed in sequential sampling models: the continuous sampling of evidence over time, and the necessity for some adjustable stopping criterion, or boundary to determine when commitment will be reached and the decision process terminated (Ratcliff & Smith, 2004).

Sequential sampling models allow for the investigation of the latent processes which underpin perceptual decision-making (Forstmann et al., 2016). Each

process takes the form of a parameter in the model, and these parameters are algorithmically optimised to produce parity between the real and simulated data. These parameters can then be compared under different conditions, or across different groups, to investigate the effects of various phenomena on unique components of the decision process (Forstmann et al., 2016; O'Connell et al., 2018). For example, when people have to make faster decisions, as in the driving example above, they may lower the level of evidence required before committing to a decision (Heitz, 2014). Sequential sampling models allow for such parameters and effects to be extracted from simple choice and reaction time behaviour, meaning they provide a detailed account of the processes which determine both the timing and accuracy of our perceptual choices.

A variety of sequential sampling models exist. While all models within this family contain the core agreement of evidence accumulation to a bound, they differ in the number and nature of the parameters that they invoke. For example, Linear Ballistic Accumulator Model (LBA) is among the most parsimonious variants. Within a two-choice task, as modelled by the LBA, evidence is accumulated with no noise into two independent accumulators, with the rate of each build-up depending on the strength of the physical evidence (Brown & Heathcote, 2008; Forstmann et al., 2008). These accumulators continue until one crosses a bound, at which point a response is made. One significant drawback of the LBA is that it assumes the perceptual process at play is one which would be produced by an "ideal observer" throughout the course of a given trial, with only between-trial variability in the rate of evidence accumulation and starting point of the accumulators allowed (Brown & Heathcote, 2008). Therefore, this model might provide suitable fits to behaviour, with analytic solutions that dramatically reduce the required number of trials and computation time (Brown & Heathcote, 2008), but it does not truly reflect how decision-making takes place in the brain. That is to say, this simplified characterisation of the decision-making process overlooks the fundamental biological limits of human perceptual processing, overestimating the efficiency of the human brain, particularly given that the biological systems underpinning perception are stochastic in nature. In early perceptual stages, transduction of sensory stimuli through the retina and perceptual cortices accumulates variance at several points (see Faisal et al., 2008 for review). This

ongoing, within-trial noise in the perceptual system is not allowed for in the simplified LBA as evidence is forced to accumulate at a fixed rate across a trial. In the LBA we see an example of where a computational model may provide a fitting mathematical explanation of behaviour, but this does not necessarily mean it is providing a direct reproduction or reflection of the biological underpinnings of cognition.

1.2.2: The Drift Diffusion Model

By far the dominant sequential sampling variant is the drift-diffusion model (DDM), within which the biological reality of within-trial noise is addressed and incorporated to some extent. In the DDM, the decision variable is assumed to build over time at a rate proportional to the mean evidence strength, or drift rate, which reflects a combination of observer ability and task difficulty (Ratcliff et al., 2016; Ratcliff & McKoon, 2008). Gaussian noise is added at each sampling point until a boundary is crossed (Ratcliff & Smith, 2004). As in the LBA, this boundary represents the amount of evidence required before a decision is made. An increased boundary results in increased accuracy, as more evidence is accumulated, but with slower reaction times, while the lowering of bounds leads to faster reaction times at the cost of accuracy (Ratcliff & McKoon, 2008). The DDM also includes a non-decision time parameter, which reflects a combination of the time required for sensory encoding and motor response execution (Ghaderi-Kangavari et al., 2022). In the full version of the DDM, additional variability parameters are also applied to the starting point, non-decision time and drift rate (Ratcliff et al., 2018). The additive effects of these noise sources represent true noise in the biological system and are reflected elegantly in double-pass procedures (Awwad Shiekh Hasan et al., 2012; Vilidaite & Baker, 2017). Here, identical stimuli are presented several times at various stages throughout an experiment. Without variability in the neurocognitive system, uniformly consistent responses should be observed, yet exact repetitions of stimuli elicit variable behavioural responses in both accuracy and response time. Given that *external* sources of variability in stimulus complexity or perceptual features remain constant across repetitions, behavioural variance can only be attributed to variability along the observer's perceptual chain. Starting point

variability accounts for Gaussian between-trial variability in the decision's initial point. This is mathematically equivalent to trial-by-trial variability in the decision bound. Normally distributed drift rate variability allows the mean strength of evidence to vary from trial to trial, while non-decision time variability adds uniformly distributed variability to the added motor and sensory encoding time (Ratcliff et al., 2018). This “full drift-diffusion model” therefore, allows for a range of potential sources of variability, both in the ongoing evidence accumulation process and at the inter-trial level.

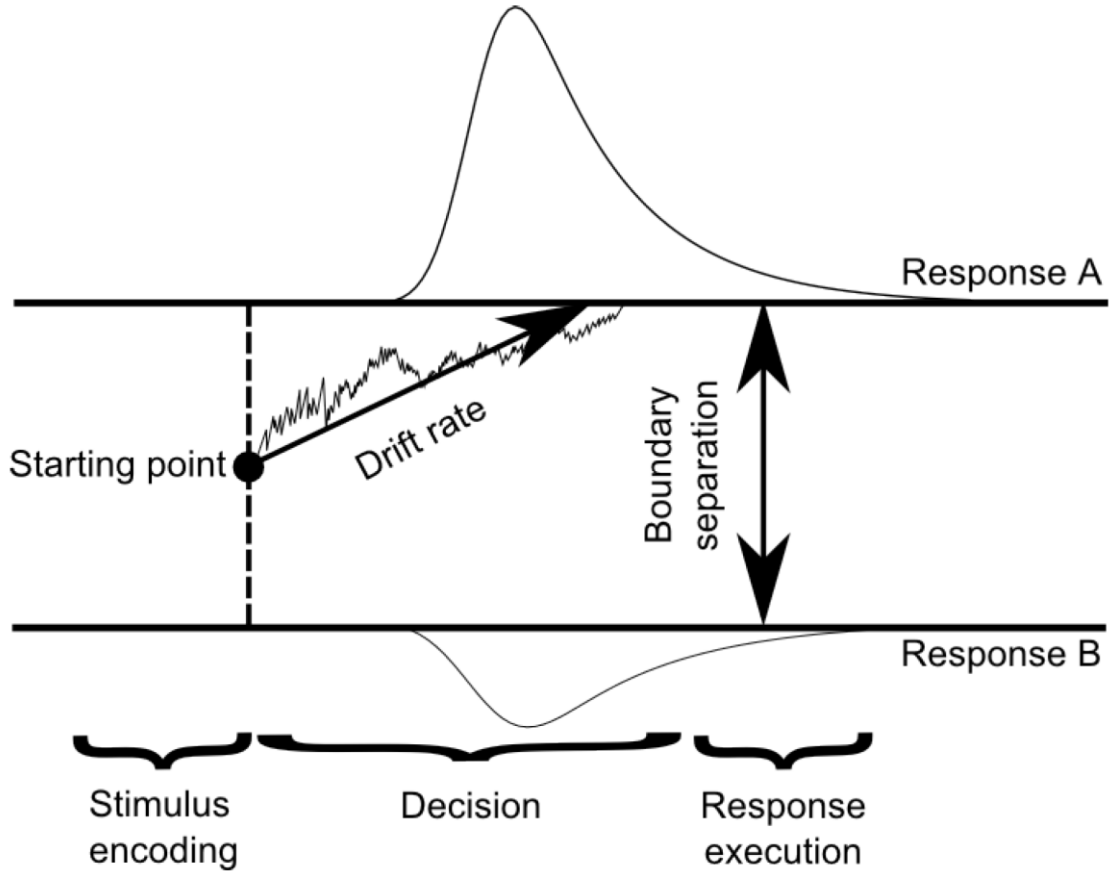

Figure 1.1: Illustration of a simple drift-diffusion model. The decision variable begins accumulating from its starting point at a mean rate set by the drift rate. Gaussian noise is added at every-time step until one of two boundaries is crossed, at which point accumulation stops and a decision is made. Non-decision time allows additional time for stimulus encoding and response execution. Between-trial variability is applied to drift rate, non-decision time and starting point. Reprinted from Schubert, A.-L., Frischkorn, G., Hagemann, D., & Voss, A. (2016). Trait Characteristics of Diffusion Model Parameters. Journal of Intelligence, 4(3), 7. https://doi.org/10.3390/jintelligence4030007

To date, the DDM has been applied extensively to the study of perceptual decision-making, giving a testable model that enables the study of how cognition

may change under different conditions, and there are several prepackaged software toolboxes across a range of platforms which enable its use and application (Voss & Voss, 2007; Wiecki et al., 2013). It provides excellent fits to behavioural data across a wide range of perceptual and cognitive tasks (see Ratcliff & McKoon, 2008; Ratcliff et al., 2016 for review) and, in its extraction of latent parameters beyond simple behavioural measurement, it provides insight into the effects of a particular task or the decision-making differences between particular conditions *beyond* those attainable from the study of behaviour alone. This makes it a useful, widely available and widely utilised tool for the study of the effects of various demands on decision-making. For example, recent work has applied the DDM in explanations of value-based decision-making, finding that high-value trials involve a raising of decision bounds (Afacan-Seref et al., 2018). The DDM has also been used to link response caution on a current trial to an individual's confidence in their performance on a previous trial (Desender et al., 2019). Furthermore, perceptual learning has been studied using the DDM, showing behavioural improvements over time to be driven by a combination of refinements to the perceptual representation via increased drift rates, and adjustments to the decision boundary (Ivanov et al., 2022; Zhang & Rowe, 2014, see Chapter Three for a more detailed discussion of this literature).

The DDM has also been applied to the study of between-group differences. For example, in ageing research, the robust finding of age-related reaction time slowing has often been attributed to a general information processing slowdown (Finkel et al., 2007; Salthouse, 1996), whereas the application of DDMs showed this behaviour was more likely due to age-related adaptations to a decision-bound and response caution (Starns & Ratcliff, 2010, see Dully et al., 2018 for review and Chapter Two for a detailed discussion). This provides a fundamentally different viewpoint, where older adults do not exhibit a generalised cognitive processing impairment, but rather employ strategic, although potentially maladaptive, adjustments to their decision-making criteria. Similarly, DDMs have been applied to the study of perceptual performance in people with autism (Pirrone et al., 2017). Here, as in older adults, poorer performance on a task was driven not by any impairment to the perceptual processing of the stimuli, but by increased non-decision times and response caution. Finally, increased reaction-

time variability in children with attention deficit/hyperactivity disorder (ADHD), is driven by reduced drift rates and faster but more variable non-decision times (Karalunas et al., 2014; Karalunas & Huang-Pollock, 2013). In summary, it is clear that the DDM is a valuable tool for providing explanatory accounts of the effects of a range of disorders on perceptual decision-making processes, with the examination of latent parameters allowing us to go beyond a simple assertion of impaired behaviour.

While the DDM is a powerful model, it is at its core entirely mathematically derived. That means its parameters are based solely on a predefined set of assumptions as to how decision-making unfolds. As a result, it may fail to capture processes that are not explicitly allowed for in the standard DDM. For example, significant debate exists in the literature as to whether decision boundaries remain fixed at the same level for the duration of a trial, or change dynamically across time (Hawkins et al., 2015). While this may seem like an arbitrary distinction, it has substantial potential to impact the model's power to explain behavioural features and the parameters to which we ascribe this power, impacting the accuracy with which other parameters are estimated. Behaviourally, participants often exhibit a higher proportion of slow errors. Traditionally the DDM attributes a higher proportion of slow errors to increased drift rate variability, where randomly lower drift rates result in slower responses that are more likely incorrect (Ratcliff et al., 2016). In contrast, the same behavioural features can be alternatively explained by a collapsing bound. Within the latter account, a boundary which collapses over time results in a higher degree of slow errors, as the response criterion continues to reduce for later reaction times, and decisions are made on less and less evidence over time. In effect, this represents an increasing desire to respond, or “urgency”. This urgency may also arise intrinsically from the participant as a means of maximising their “reward rate” (Bogacz et al., 2006). Here a participant's desire to increase the ratio of correct responses over time generates urgency to respond more quickly in order to move on to the next trial and next reward. These collapsing bounds have been shown to be optimal (Malhotra et al., 2018), especially in tasks where participants are faced with a hard and predictable deadline for response (Frazier & Yu, 2007). Here we are presented with a significant issue in modelling

research. The DDM, with fixed bounds, may be able to account for features of behaviour such as slow errors through between-trial variability parameters, yet an alternative account can explain the same behaviour just as reasonably, viably and potentially more optimally. Traditionally, model comparison relies solely on how well the simulated model data fit the real behaviour, and the number of parameters needed to achieve these fits (Vrieze, 2012; Wagenmakers & Farrell, 2004), meaning that if models with and without collapsing bounds can provide equally good fits to data, it becomes difficult to arbitrate between them. This is a significant issue for sequential sampling models, as they may provide accurate fits to the data, but achieve these fits in a way that does not accurately reflect the decision process underneath, or how it evolves in the brain. These issues can be addressed through neurally-informed modelling, as will be discussed. However, it is first beneficial to assess evidence which has supported the DDM as a viable account for how the brain makes decisions.

1.2.3: Electrophysiological Signatures of the Decision-Making Process

In addition to its success in modelling behavioural data, the DDM has been supported as a reasonable approximation for how decision-making unfolds in the brain by a range of neural studies (Shadlen & Kiani, 2013, see O'Connell et al., 2018 for review). Electroencephalography (EEG) provides perhaps the best method for studying rapidly evolving decisions (O'Connell et al., 2018), due to its high temporal resolution, which allows for the extraction of event-related potentials (ERPs) at a latency unattainable through other neuroimaging techniques such as functional magnetic resonance imaging (fMRI). However, the spatial resolution of EEG is limited due to volume conduction through the brain and skull, which means that signals measured at the scalp are a summation of many neural sources (van den Broek et al., 1998), the possible combinations of which are infinite. EEG activity also reflects mostly cortical activity and, while possible through decomposition techniques, it is not as readily equipped to accurately capture activity in deep-lying, subcortical brain areas (Krishnaswamy et al., 2017). As such, while the technique of EEG offers little in the way of brain

activity localisation, its rapid timescale provides a valuable lens through which to study the emergent cognitive and behavioural processes that underpin perceptual decision-making.

To date, electrophysiological research has offered a significant degree of empirical support for the DDM. Elegantly designed studies have linked several key neural markers with distinct stages of the decision-making process, from sensory encoding to motor preparation and response. In animal research, spike rates of neurons in the monkey lateral intraparietal (LIP) cortex have been shown to exhibit accumulation to bound dynamics which scale in proportion to evidence strength and reaction time (Huk & Shadlen, 2005; Huk & Meister, 2012; Rorie et al., 2010; Shadlen & Newsome, 1996), mirroring the core predictions of the DDM. These accumulation-to-bound dynamics have also been observed in areas such as the frontal eye field (FEF; Hanes & Schall, 1996) and superior colliculus in the midbrain (Ratcliff et al., 2003, 2007). Collectively these lend support to the DDM, showing neuronal firing which accumulates to a threshold before response in several regions throughout the brain.

Efforts to find similar signals in the human brain have been complicated due to the necessity to use non-invasive techniques. Direct-brain recordings are often not attainable in a general population, as they rely on intracranial placements of electrodes which is not achievable unless the skull is opened for a clinical purpose (e.g., surgery for intractable epilepsy, see Lachaux et al., 2003 for review). Instead, we must rely on scalp-recorded EEG which, due to the effects of volume conduction outlined above, makes it more difficult to study specific brain areas in isolation. However, recently identified EEG signals present easily identifiable and reliable neural markers of the decision process in humans. For example, the centroparietal positivity (CPP) reflects evidence accumulation for any stimulus type in any sensory modality even when no response is required (Kelly & O'Connell, 2013; O'Connell, Dockree, et al., 2012). Consistent with predictions of the DDM, and analogous single unit signals observed in monkeys, the CPP builds at an RT-predictive, evidence-dependent rate and reaches its peak at the time of the decision report, reflecting an effector-independent readout of accumulated evidence over time (Kelly & O'Connell, 2013). It accumulates

positively, irrespective of which alternative the evidence favours, indicating it is not choice-selective, but rather reflects the absolute value of cumulative evidence, and appears over the same brain areas where monkey research has identified decision signals (O'Connell, Dockree, et al., 2012). In a study by Twomey and colleagues (2015) this signal was shown to be equivalent to the classic P300 or P3b. This component is perhaps the most intensively researched ERP signals of the past 60 years across a range of tasks. It shows a peak after approximately 300ms and was previously linked to a range of potential cognitive functions, including context updating (G. McCarthy, & Donchin, 1981), working memory (Polich, 2007), and response facilitation (Nieuwenhuis et al., 2005; see Verleger, 2020 for a review of theories of P3b function). The P3b is classically elicited on oddball tasks, where target stimuli appear irregularly among non-target stimuli. However, in their 2015 study, Twomey and colleagues used an auditory oddball with varying difficulty levels to show that the P3b reached a consistent amplitude before response and a build-up that correlated with detection difficulty and reproduced the decision variable of a diffusion model. This was an important advancement, indicating the evidence accumulation function of a well-studied ERP component, and showing it to be functionally equivalent to the CPP. Through this it further established the CPP as a marker of evidence accumulation, providing a useful readout of cumulative evidence which appears across a range of cognitive tasks.

Time-frequency analysis can also provide unique neural indices of the decision-making process. Mu/beta activity (8-30Hz) measured bilaterally over the motor cortex has been shown to be a valuable marker of effector-selective response preparations. Mu/beta activity over the motor cortex contralateral to the chosen hand desynchronises towards a stereotyped amplitude irrespective of reaction time or task difficulty (Donner et al., 2009; Fischer et al., 2018; Kelly et al., 2021; Murphy et al., 2016; O'Connell, Dockree, et al., 2012; Rogge et al., 2022). While this threshold-crossing effect at response appears consistent across conditions, pre-stimulus levels of mu/beta activity in both hemispheres have been shown to vary systematically as a function of different contextual variables. For example, under speed pressure, both hemispheres show lower pre-evidence mu/beta activity, decreasing the excursion to the motor threshold and reflecting indicating

increased motor preparation and pre-evidence adaptation to the starting point of the decision process (Kelly et al., 2021; Steinemann et al., 2018). Steady-state visual evoked potentials (SSVEPs) can also be used to track perception at its fundamental sensory encoding level (O'Connell, Dockree, et al., 2012). Using either frequency or phase tagging to bind stimulus features to specific measurable traits, such as setting the flicker rate for a stimulus at a specific frequency, we can extract the representation of that frequency in the visual cortex by examining the power of that band over occipital electrodes. This provides a readout of the sensory encoding of the stimulus in the brain. Under speed pressure, the SSVEP amplitude for the target grating is boosted, likely due to top-down attentional control, potentially reflecting increased drift rates at the earliest sensory level (Steinemann et al., 2018).

EEG signals may also be useful in decomposing the distinct contributors to non-decision time in the DDM. For example, motor potentials, measured as EEG activity over the motor cortex contralateral to the hand chosen for response, show an inflexion point approximately 100ms before response, from which the signal begins to ascend positively, indicating a marker of commitment to the motor action, and estimating the motor execution component of non-decision time (Kelly et al., 2021). Nunez and colleagues (2019) have also shown a relationship between the N200, a marker of the onset of visual-sensory processing, and the visual-encoding portion of non-decision time. Here, neural data can be used to identify distinct components of non-decision time which are usually compiled into one value in the DDM. Collectively these studies present a rich roster of neural signals which are uniquely related to distinct components of the perceptual decision-making process.

1.3: Neurally-Informed Modelling

To date, a substantial body of fMRI, magnetoencephalography (MEG) and EEG work has been done that attempts to link parameters of sequential sample models to neural data. For example, Jia and colleagues (2018) linked drift rates on an LBA model of learning to blood-oxygen-level dependent (BOLD) activity in the supplementary eye field (SEF), FEF, and intraparietal sulcus (IPS). Similarly, using a speed-accuracy tradeoff paradigm, Forstmann and colleagues (2008) correlated increased activity in the striatum and pre-supplementary motor area to decreased decision bounds under speed emphasis. These studies are beneficial in providing a means of directly associating brain activity with model parameters. Through this, they provide neural evidence for various processes in sequential sampling models.

While beneficial, these examples are ultimately unidirectional in that they do not utilise the observed neural data to inform their models. Neurally-informed modelling is a relatively novel approach, which aims to leverage neural data to refine, constrain and validate existing decision models and, in so doing, tackle some of the limitations of behavioural modelling (O'Connell et al., 2018; O'Connell & Kelly, 2021). Neural data can lend support to the inclusion of novel parameters that more accurately reflect how decisions are formed in the brain, helping us to tackle some of the issues of model mimicry outlined above. Furthermore, they provide means for constraining specific model parameters in a reasoned way, allowing for increased model parsimony and enabling the investigation of increasingly complex effects (Kelly et al., 2021). Finally, they provide a data-driven means for estimating scaling parameters across groups and conditions, which can help to overcome some of the limitations of the DDM in its use of fixed scaling parameters. The following section will introduce some novel neurally-driven means for addressing these issues. It will explore some studies which have successfully applied neurally-informed modelling to grant new insights into the perceptual decision-making process, before outlining some potential issues which may impede neurally-informed modelling itself.

1.3.2: Applications: Inform Model Parameters

As outlined, sequential sampling models have provided parameters that are linked to psychological constructs and neural indices (Forstmann et al., 2008a; Jia et al., 2018; Nunez et al., 2019). However, thus far, the models have been almost exclusively evaluated based on their ability to fit behavioural data. While a model may be well-fitting, it may not accurately represent the key neural processes that determine this behaviour in the brain. As a result, a range of different models may successfully reproduce observed behaviour, but each model can be driven by fundamentally distinct parameters (O'Connell et al., 2018). These distinctions are important as each parameter represents unique psychological constructs, and therefore each implies different explanations for how a decision unfolds in the brain. An example of this competition between different model accounts is outlined above in the case of the static vs collapsing bounds (Hawkins et al., 2015; Ratcliff et al., 2016, see section 1.2.2). One potential benefit of neurally-informed modelling is that close examination of neural markers of the decision-making process can grant valuable information that allows us to successfully arbitrate between competing models and guide which parameters should and should not be included in a model. This results in models which are more reflective of how the brain makes decisions and can shed new light on how novel or existing parameters contribute to decision-making.

For example, while behavioural modelling alone has struggled to resolve the debate between static or collapsing bounds (Malhotra et al., 2018; Ratcliff et al., 2016), neurophysiological research has consistently pointed to a role for dynamic urgency. The CPP builds to a lower amplitude at response for slower reaction times indicating a lower level of evidence accumulated before the response is made (Steinemann et al., 2018), in line with predictions of slower errors caused by a collapsing bound. Recent work from Devine (2019) has indicated the frontocentral negativity (FCN) or contingent negative variation (CNV) to be a marker of dynamic urgency. The CNV has previously been indicated as a marker of perceptual timing or anticipation (Kononowicz & Penney, 2016). However, in their work Devine (2019) showed the CNV to be a time-dependent, evidence-independent marker which shows static pre-stimulus adjustments to amplitude

under speed emphasis, reflecting a marker of increasing urgency to respond over time. This supports previous modelling efforts showing dynamic urgency by Boehm and colleagues (2014) and neural evidence from the macaque's LIP (Hanks et al., 2014) which also showed an evidence-independent urgency marker which was steeper under speed pressure. A similar urgency component is observed in mu/beta signals, where increased speed-pressure bilaterally increases the pre-evidence level of desynchronisation, beginning its desynchronisation to threshold in advance of any evidence appearing on screen (Kelly et al., 2021). Here we see ample evidence for neural signatures of a key behavioural feature which is often not allowed for in a standard DDM.

Collectively, the examination of this neural data provides evidence for the inclusion of a dynamic urgency parameter in decision models. The benefits of incorporating novel parameters based on neural data were demonstrated in the study by Kelly and colleagues (2021) outlined above. Here, the inclusion of a freely estimated urgency parameter, which set a rate for the linear collapse of the boundary across time significantly improved model fits, both in terms of numeric model comparison score, and the model’s ability to successfully reproduce observed neural data.

1.3.2: Applications: Constrain Model Parameters

Neurally-informed modelling can go a step beyond correlating brain activity with neural information or leveraging this information to add new parameters to a model. Utilising established neural markers of the decision-making process, we can constrain model parameters based on observed expectations for how these parameters should behave based on the neural data. For example, Turner and colleagues (2013) jointly fit neural and behavioural data to a Hierarchical LBA model. This allowed the model parameters to be constrained in their estimation by the neural data, however, these constraints may be imperfect as they assume a straightforward relationship between slow-moving BOLD activity and LBA parameters. In contrast, the neural EEG signatures outlined in section 1.2.3 provide empirically characterised and validated indices whose features are uniquely related to different stages of the decision process and behaviour. This

affords a more nuanced means of directly estimating and constraining specific parameters, which, as will be described, create testable models which furnish significant advancements to our understanding of decision-making.

To reduce the potential for overfitting models to specific data sets, model comparison scores such as Akaike's Information Criterion (AIC) or Bayesian Information Criterion (BIC) penalise model complexity (Vrieze, 2012; Wagenmakers & Farrell, 2004). In this way, every additional parameter introduced increases the comparison score. While this has the benefit of reducing the potential for bloated, overfit models with enough parameters to fit any data set, it may impede the investigation of paradigms which require the addition of novel parameters such as urgency. Similarly, it becomes difficult to assess the within or between-group effects of different model parameters in tandem, as each parameter that is allowed to vary by condition incurs another cost. Constraining models based on neural data provides a data-driven estimation of the parameter value, allowing an additional degree of freedom to the model, and enabling the investigation of more complex model variants without a cost to the model comparison score. These neurally-constrained models can then shed new and unique insight into a range of phenomena which is not achievable through standard DDMs alone.

For example, Hanks and colleagues (2014) highlighted start point biases in the macaque LIP under speed-emphasis on a motion discrimination task. These start point increases had the effect of lowering the decision bound. An increasing, evidence-independent but time-dependent build-up of the LIP activation was also observed. They used this as evidence for an increased urgency component under speed pressure and calculated an urgency function based on these firing rates for each monkey. Using these values to constrain their model, they indicated that both pre-evidence adjustments to the starting point and a linearly increasing urgency component were required to provide the best fits to behaviour. In this study, we not only see an example of neural evidence being used to support the inclusion of an urgency parameter, but the utilisation of specific, neurally derived constraints to test these inclusions.

A recent study by Corbett and colleagues (in press) aimed to constrain urgency in a similar way in their model of value-based decision-making in humans. Here, when participants were instructed that a certain choice was more valuable than others, a significant pre-evidence adjustment to mu/beta motor preparation was observed. For low-value trials, however, the onset of motor preparation was slower but built at a significantly steeper rate once it began. This suggests a lower pre-evidence anticipation, but an increased urgency to respond once the evidence became available. They constrained two elements of their model: the pre-evidence adjustments to the starting point by pre-evidence beta amplitude, and the rate of urgency by the mean slope of these signals for each value condition. This successfully explained both the increased accuracy in the high-value condition and slow and fast errors in the low-value condition. This neurally-constrained model was then used to simulate motor signals which successfully reproduced distinct pre and post-evidence biases that the standard DDM could not achieve.

Constraints of the decision process based on pre-evidence adjustments observed in the neural data were also explored by Kelly and colleagues (2021) in their study introduced above. Here, they constrained the start-points of the decision process based on the pre-evidence mu/beta levels, which were shown to exhibit significantly more motor preparation under speed-emphasis. While initial standard DDMs only implicated adaptations of the decision boundary under speed emphasis, this neurally-constrained model allowed for the investigation of a more complex model. That is to say, the boundary adjustments were accounted for by the neural constraint, freeing the model to investigate the effects of speed-emphasis on other parameters, without incurring additional costs to model comparison scores that penalise for complexity. Their model was further supported by simulating evidence accumulation and motor preparation signals based on their neurally-constrained model. These simulations successfully reproduced the observed neural signals, showing simulated motor-preparation signals that closely resembled beta and an evidence accumulation signal, which mirrored the CPP. The close alignment between neurally derived and model-simulated indices of the decision process lends further evidence in support of the neurally-constrained model. Overall, these studies highlight the potential

applicability of neurally-informed modelling, granting information which can guide the construction of a model. They also illustrate the benefit of simulating model-estimated decision signals as comparing these signals against neural data can provide additional validation and evidence for or against a given model.

Much in the same way as the inclusion of urgency in models helps the model better represent the brain, neural constraints have led to the exploration of other novel parameters which are not allowed for in the standard DDM. For example, the DDM accounts for fast errors as increased start-point variability (Ratcliff et al., 2016) where the decision variable begins randomly closer to one of the decision bounds. However, in tasks with a period of random, no-evidence stimulus lead-in, where participants may not fully know whether evidence has begun or not, participants are often seen to begin their evidence accumulation signal (CPP) before any evidence becomes available on screen (Devine et al., 2019). This early build-up forms an important part of the decision process and is associated with faster responses and errors (Devine et al., 2019), suggesting a period where the decision has begun to form without any evidence on screen. This accumulation can only be noise as no evidence is available, yet the DDM cannot allow for any accumulation of noise in advance of evidence onset. Instead, evidence always accumulates alongside noise from the model's starting point.

Recent neurally-informed modelling has shown that early accumulation of noise is a distinct phenomenon which may need to be incorporated into models on some tasks, and which behaves differently under different conditions. In their neurally-informed model, Kelly and colleagues (2021), utilised motor potentials measured over the motor cortex contralateral to the chosen response to estimate the time at which a motor response was committed to. Constraining the motor-response execution portion of non-decision time allowed the researchers to produce two novel and distinct freely estimated model parameters for sensory encoding time and accumulation-onset time. Here, only Gaussian noise was accumulated from accumulation-onset time until the sensory encoding time had elapsed, while sensory encoding time accounted for the delays in stimulus encoding. This pre-evidence accumulation was shown to vary and begin before

the evidence could have borne an influence on the decision process. This allowed the model to produce simulated decision variables which captured the pre-evidence accumulation of the CPP. Here neurally-informed modelling provides evidence for features of decision-making which are not incorporated into the DDM, showing two uniquely adaptable components of non-decision time, and granting new and valuable insights into perceptual decision-making.

Work from Ghaderi-Kangavari and colleagues (2022) used neurally-informed DDMs to investigate the role of visual encoding time (VET) as a subcomponent of non-decision time. Studying the effects of spatial cueing on performance on a visual discrimination task, they constructed a DDM which constrained non-decision time by a linear connection to the single-trial estimate of visual-evoked N2 latency. Through this they separated a distinct VET from a motor execution time (MET), improving the model's fit to the real data. A standard DDM would suggest that improved performance was due to drift rate biases introduced by the cue. However, the neurally-informed model indicated that spatial cueing reduced both VET and MET, with no adjustments to the drift rate. While this study did not show unique influences of spatial queuing on the two subcomponents of non-decision time, this decomposition would be unachievable using standard DDMs, which penalise the addition of parameters to separate non-decision time into its constituent parts. Here we see another example where the neural data can indicate psychological constructs which are not traditionally accounted for in the DDM. The introduction of neural constraints allows us to explore novel parameters which enable the DDM to more closely resemble how the decision unfolds in the brain. This should be a priority. Through this, the models move from purely mathematical algorithms to more realistic reflections of how the complex chain of perceptual decision-making unfolds.

Beyond introducing novel parameters in within-subject designs, neural constraints also grant us new insights into between-group differences phenomena which are not achievable through standard modelling alone. For example, in recent work by McGovern and colleagues (2018), neurally-informed modelling was used to give novel insight into decision-making in older adults. Older adults showed higher hit rates on a contrast-change detection task. Initial

behavioural modelling attributed improved performance by older adults to a combination of higher drift rates and higher decision bounds. However, analyses of CPP and mu/beta amplitudes immediately prior to the participants' decision reports suggested no between-group differences in their decision bound or rate of evidence accumulation. With this in mind, a new model was implemented in which both the boundary and drift rate were constrained to be equal across both groups. This improved model comparison scores, with the hard constraint resulting in only a slight increase in error score. Importantly, the new neurally-constrained model indicated a different parameter as driving the behavioural difference. Here, there was increased drift rate variability in younger adults, rather than lower drift rates and bounds, which resulted in poorer performance than older adults. This model was supported by neural data, where examination of posterior alpha-band activity, a well-established marker of attentional engagement (Dockree et al., 2007), and response-aligned CPP slopes, as a marker of evidence accumulation, showed increased variability in younger adults. This suggests that older adults achieved better performance through increased attentional engagement, leading to more consistent responses. Again, it is shown that neurally-informed and constrained modelling can lead to substantially different assertions as to the drivers of behavioural differences across groups, but ones which are supported by the neural data.

1.3.3: Applications: Data-Driven Scaling Parameters

Within applications of the DDM, for the model to converge, one parameter is fixed at an arbitrary value around which all other parameters are scaled (Ratcliff & McKoon, 2008). Other model parameters are then calculated in reference to this value, meaning for example, that a doubling of the scaling parameter would cause a doubling of all other parameters. Without a scaling parameter, the model search space would be massive. Fixing a scaling parameter reduces this set of possible parameter combinations drastically, allowing the model to converge more fully without falling into local minima. By convention, in the DDM the standard deviation of the within-trial Gaussian noise parameter added at each time step to the drift rate is used as the scaling parameter and is typically set to 0.1. This means that, at each time-point, a sample of evidence is taken from a

Gaussian distribution with a mean of the model estimated drift rate and a fixed standard deviation of 0.1. This within-trial noise parameter, therefore, sets the scale of values for drift rate, bound, non-decision time, and all variability parameters.

A scaling parameter is therefore useful when examining an individual subject, as it assists convergence by allowing all parameters to be searched on the same scale. However, when we attempt to model data in which we cannot assume that the process reflected in the scaling parameter is stable across *groups* of participants we encounter significant issues. For example, recent work has shown that perceptual internal noise (Dosher & Lu, 2017, see Chapter Three for detailed discussion) is seen to vary significantly across a range of clinical samples, such as individuals with autism and schizophrenia (Chen et al., 2014; Park et al., 2017). Similar group differences in ongoing noise levels are often seen in neural data. For example, while older adults show equal behavioural measures of internal noise relative to younger adults (Pardhan et al., 1996), they show decreased neural variability over long time scales and increased variability at a moment-to-moment level (Kumral et al., 2020; Sleimen-Malkoun et al., 2015), with this increased noise in older adults linked to poorer behavioural performance and working-memory decline (Voytek et al., 2015). Group differences in neural variability across different timescales are also seen in individuals with ADHD (Gonen-Yaacovi et al., 2016; Woltering et al., 2012), autism (Hecker et al., 2022; Milne, 2011), and schizophrenia (Kim et al., 2018; Yang et al., 2014). While some neural noise may be behaviourally-irrelevant (Ribeiro & Castelo-Branco, 2022), these behavioural and neural markers of ongoing noise in the neural system may indicate important and clinically relevant group differences which underpin behavioural impairments across a range of populations (Dinstein et al., 2015). The DDM would allow us to directly probe the relationship between intra-trial noise and behaviour, however, in its current form, it often does not allow for this, as within-trial noise is fixed as a scaling parameter across groups. Given the growing body of evidence for potential group differences in ongoing variability, it is clear that a fixed within-trial noise level across participants of different clinical groups may not be suitable and may result in the misattribution of group differences to fundamentally incorrect parameters.

Within-trial noise is also fixed as a scaling parameter at an individual level, meaning intra-trial variability in the decision process is not allowed to vary as a function of any experimental manipulation or effect. This is problematic, as psychophysical estimates of internal noise are shown to reduce with perceptual learning (Dosher & Lu, 1998, 1999; 2017). As perceptual learning can contribute in part to all perceptual decision-making tasks, this is a significant issue in the standard DDM. One method to overcome this is to fix an alternative parameter as the scaling parameter (Donkin et al., 2009). However, in doing this we preclude the possibility of the newly fixed parameter varying across conditions or groups. For example, an alternative solution to the study of perceptual learning would be to fix the boundary at a constant level across sessions. However, perceptual learning has also been shown to induce boundary adaptations (Liu & Watanabe, 2012; Zhang & Rowe, 2014). Here, the DDM cannot investigate changes to the scaling parameter without fixing another potentially important parameter and thus prohibiting it to vary. In essence, scaling parameters are an imperfect solution to model convergence, as they enforce at least one parameter to remain fixed, irrespective of the evidence for group differences, or changes across time or tasks.

A method capable of estimating scaling parameters in a more principled way would allow for increased model precision. Neurally-constrained models use neural data to set a numeric value for a parameter (Ghaderi-Kangavari et al., 2022; Hanks et al., 2014; Kelly et al., 2021; McGovern et al., 2018, Corbett et al., in press). In effect, this value acts as a scaling parameter around which the remaining parameters can be estimated. While not attempted in these papers, such a technique allows for within-trial noise to be freely estimated, or to be estimated differently across conditions. This presents a data-driven neurally-informed solution to the scaling problem in the DDM. An alternative method would be to utilise behavioural estimates of internal noise from psychophysical work as a means of hard constraining the within-trial noise either at the individual or group level (Dosher & Lu, 2017; Tibber et al., 2014). This allows for adaptations in within-trial noise across individuals, conditions or groups while allowing for parameters such as bound to be freely estimated. Both of these methods present

compelling means for potentially addressing the scaling problem in DDM literature and would engender models with the ability to reflect the complexity of the brain’s perceptual decision-making machinery more accurately. Again, we see a limitation of DDMs is presented which can potentially be addressed through the use of neurally-informed modelling.

1.4: Methodological Challenges for Neurally-Informed Modelling

While neurally-informed modelling presents an exciting opportunity to refine and improve our models of perceptual decision-making, there are certain key considerations which must be considered to ensure that neurally-informed models are sufficiently useful. These issues must be addressed in order to develop a concrete and reliable methodology which accurately informs and constrains decision models.

1.4.1: Inter-Individual and Inter-Group Difference in Signal Presentation

Of particular importance to neurally-informed modelling is our ability to confidently assert that the neural signals we are measuring accurately depict the psychological process of interest. As outlined above, myriad works have used intricate paradigms to associate neural indices of the decision-making process with DDM parameters (see section 1.2.3). However, there exist significant behaviourally-irrelevant inter-individual and inter-group differences in signals which may impede the reliability of these signals. For example, there exists substantial inter-individual variability in skull thickness and conductivity, which may directly influence the amplitudes and presentations of various signals measured at the scalp (Antonakakis et al., 2020). When these signals are averaged at a group level, individuals with larger amplitudes due to task-irrelevant differences in head shape or skull thickness may dominate a group's average topography. These inter-individual differences can be exacerbated when measuring across multiple sessions. Here, issues such as cap placement can introduce additional intra-subject variability unless special attention is paid to ensure precise and consistent EEG set-up, and these concerns must be addressed for us to be able to fully rely on a neural signal as a true indicator of a particular underlying psychological construct.

The effects of behaviourally-irrelevant differences in signals are even more apparent when we compare across groups where the effect of non-decision related factors which influence the appearance of the signals, but not the processes they represent may become compounded. For example, skull thickness has been shown to change with age, meaning that the same signal measured in the same location across two age groups may vary in amplitude simply due to reduced skull conductivity in older adults (Antonakakis et al., 2020). Furthermore, recent evidence from Ribeiro and colleagues (2022) has suggested that some components of neural variability may not be behaviourally relevant but may impact an ERP's amplitude at a trial-to-trial level. Here, older adults were shown to exhibit reduced ongoing variability in the CNV when compared to younger adults. However, this variability was not related to behavioural performance. When this was controlled for, older adults' CNV amplitudes were more closely aligned with behaviour, suggesting that this non-decision-related background activity bore an impact on key signals of decision-making. Given an increasing reliance on neural indices as markers of the decision process, efforts must be made to identify and correct for potential behaviourally-irrelevant differences in EEG signals. Failing to do so may mean that researchers could incorrectly infer an adaptation to some component of the decision process that is in fact driven by another non-decision-related group difference.

For example, in comparisons between the younger and older brain, there is an observed increase in mu/beta-band power at rest and increased desynchronisation during movement in older age (Rossiter et al., 2014; Stacey et al., 2021). This may be caused by increased GABA inhibition across the motor cortex, yet whether this represents any functional impairment remains unclear (Muthukumaraswamy et al., 2013). Pharmacological studies which artificially increase levels of GABA in younger adults show similar increases in beta amplitude but are not associated with any behavioural impairments to movement (Muthukumaraswamy et al., 2013). This suggests that elevated beta amplitudes seen in older adults may reflect background sources which hold little influence on task performance. This could in turn lead to incorrect conclusions regarding the effects of ageing on a decision-making process, such as pre-evidence starting point adjustments (Kelly et al., 2021). Oftentimes, however, individual

differences can be corrected through normalisation. For example, mu/beta amplitude is shown to desynchronise toward the same threshold at response, independent of the difficulty of a task or reaction time (Kelly et al., 2021, see section 1.2.3). This means that differences in observed spectral power can be controlled through normalisation of the signal in a manner that maintains inter-group and inter-condition differences in amplitude and build-up (see Chapters Two & Three).

This signal normalisation may create an elegant solution for the comparison of accumulation-to-threshold signals such as mu/beta activity across groups. However, this normalisation cannot be applied to signals such as the CPP or CNV. While some studies have shown consistent CPP thresholds at response across RT and difficulty level (Kelly & O'Connell, 2013; Twomey et al., 2015), others have shown reduced amplitude at response for more difficult trials and slower RTs (Steinemann et al., 2018). Here the CPP amplitude at response may be influenced by parameters such as dynamic urgency, meaning we cannot assume a stereotyped boundary of the CPP across all trials in the same way we do for mu/beta. Without this normalisation, we may be unsure of the extent to which differences in CPP presentation across conditions or groups represent changes to underlying psychological constructs, or rather a combination of task-irrelevant effects on scalp potentials. This is especially important for the study of ageing, where evidence has suggested lower P3b/CPP amplitudes in older adults (see van Dinteren et al., 2014b and Chapter Four for detailed discussion). Without a way to normalise these signals, it is difficult to ascertain if these differences represent any true functional deficit in older adults. This presents an issue for neurally-informed modelling. However, comparing the CPP signal to its model-simulated cumulative evidence counterpart may provide a way forward. Through this, we can see whether the two align and show concordance. Failure for one to reflect the other may signify a model which does not represent how the decision unfolds in the brain. However, if the model successfully reproduces behaviour and simulates other signals such as the mu/beta motor-preparation signal well, it may instead indicate that the observed neural signals are affected by behaviourally-irrelevant issues such as volume conductance. Through this, steps can be taken that attempt to normalise the observed neural signal,

potentially utilising model-derived parameters to rescale the ERPs in a reasoned, data-driven way.

1.4.2: Signal Overlap

A further issue with the study of ERP signals such as the CPP and the CNV is the fact that scalp-recorded EEG is an overlapping mix of a range of neural sources (van den Broek et al., 1998). This means that an ERP measured through electrodes recording from the scalp may not always be a reliable depiction of one signal of interest in isolation. This can lead to significant issues, as signals of interest such as the CPP may become overshadowed by other nearby signals such as the CNV (Kelly & O'Connell, 2013). Furthermore, paradigms can often elicit sensory-evoked potentials which may in turn overlap with emerging decision-related signals. For example, the auditory oddball task traditionally used to study the P3b/CPP (e.g., Polich, 1997; van Dinteren et al., 2018) uses sudden-onset stimuli which elicit an auditory-evoked potential. This presents an issue as it is unclear whether the observed effects of age on P3b/CPP amplitude may be driven by ageing-effects on the sensory-evoked potential (Anderer et al., 1996, 1998; Gajewski et al., 2018; Bertoli et al., 2005; Golob et al., 2001; Kuruvilla-Mathew et al., 2022; Polich, 1997a). To reduce the influence of these sensory-evoked potentials on decision signals of interest, the O'Connell lab has designed and implemented paradigms in which sensory evidence emerges gradually over time, removing sudden sensory signals and allowing for a clearer picture of the decision-relevant signals such as the CPP in isolation (McGovern et al., 2018; O'Connell, Dockree, et al., 2012; Steinemann et al., 2018, see Chapter Four)

Scalp spatial filters, such as surface Laplacian or current source density (CSD) transforms, offer another potential means of reducing the effects of signal overlap (Kayser & Tenke, 2015). These filters transform scalp-recorded voltage into estimates of current flowing from the brain, aiming to undo a portion of this volume conduction (Tenke & Kayser, 2005). This has the effect of reducing signal overlap by providing more spatially resolute topographies. The application of CSD has been shown to be beneficial to the study of the CPP, removing the influence of the neighbouring, but negative going CNV (Kelly & O'Connell, 2013).

However, in many EEG studies, notably those investigating the effects of ageing on P3b/CPP amplitude, these filters are not applied, and thus the influence of these overlapping signals may remain (e.g., Polich, 1997; van Dinteren et al., 2018, O'Connell, Balsters, et al., 2012). Resultantly, we may inappropriately infer some effect of ageing on the CPP, and therefore on the decision process, when the differences may be driven by age-related changes in another signal entirely.

Independent Component Analysis (ICA) offers another solution for addressing signal overlap. The technique allows for the selection of spatially-independent but temporally correlated components which reflect ongoing neural activity and is often used to remove components that reflect artifactual activity in the data during preprocessing stages (Jung et al., 1997). However, it can also be utilised as a means for extracting neural markers of cognitive and perceptual processes, presenting a valuable tool for the extraction of more refined indices of perceptual decision-making. This technique has been applied to varying degrees of success in the extraction of P3/P3b components (e.g., Debener et al., 2005, 2005; Porcaro et al., 2019; van Dinteren et al., 2018, see Chapter Four). However, to date studies have often relied on potentially biassed manual selection features based on peak stimulus-locked amplitudes and latencies using oddball tasks. As a technique, ICA offers a unique opportunity to extract evidence-dependent neural indices of decision formation, leveraging the predictable difference in CPP build-up for different evidence levels and reaction times to automatically select CPP-like components. This can provide a spatially independent neural index of evidence accumulation that reduces the influence of overlapping signals, giving a more precise representation of how decisions are formed in the brain. This would be of significant benefit to neurally-informed modelling, providing a more refined neural index of decision-making. This technique will be explored more thoroughly in Chapter Four.

1.5: Summary and Current Research

This review illustrates the exciting possibilities of using neurally-informed modelling to enable more biologically grounded and reliable accounts of the perceptual decision-making process. Current modelling and electrophysiological techniques in isolation face a set of fundamental challenges which limit their potential. Computational modelling is often performed without specific reference to observed biological realities, and as a result, substantial model mimicry can be observed, where various accounts can explain behaviour equally well (O'Connell et al., 2018). This complicates our ability to rely on modelling accounts, and efforts must be made to routinely use neural markers of the decision process to arbitrate between biologically reasonable model variants. Furthermore, as models must assume a fixed scaling parameter (Ratcliff & McKoon, 2016), significant issues exist when comparing model parameters across groups. If computational models are to have significant clinical value, novel methods for addressing these problems must be found.

Neurally-informed modelling may allow us to tackle these issues. Using neural data allows us to guide parameter selection, giving evidence for the inclusion of parameters such as urgency (Devine, 2019, Kelly et al., 2021). It also allows us to arbitrate between competing models, giving us evidence in favour of, or against, specific models based on how well they reflect what occurs in the brain (Corbett et al., in press, McGovern et al., 2018). Furthermore, using neural data to provide hard constraints for model parameters allows for the exploration of novel and increasingly complex model variants which may more accurately reflect the true underlying cognitive processes (Ghaderi-Kangavari et al., 2022; Hanks et al., 2014; Kelly et al., 2021; McGovern et al., 2018, Corbett et al., in press), without increasing the number of free parameters and thus penalising model comparison scores. These constraints may also allow us to replace standard scaling parameters with neurally derived constraints, allowing for the investigation of the relationship between within-trial noise and observed behaviour (see Chapters Two & Three, Dosher & Lu, 2017). Utilising these techniques has already proved fruitful in providing novel insights and understanding into cognitive phenomena. However, as has been demonstrated

so far, the use of electrophysiological indices to constrain models is complicated by significant task-irrelevant, inter-individual, and group differences which can significantly impact how these signals present themselves (Antonakakis et al., 2020; M. Ribeiro & Castelo-Branco, 2022; Rossiter et al., 2014; Stacey et al., 2021). While signals such as mu/beta which have consistently been shown to accumulate to a stereotyped bound can be normalised (e.g., Kelly et al., 2021), a more refined and considered approach is needed to ensure signals such as the CPP are comparable across groups. Alongside this, a clearer understanding of the effects of grouping (e.g., by age) on the localisation and presentation of signals is needed, exploring spatially independent means of extracting neural signals to reduce the influence of overlapping signals.

It is the purpose of this present body of work to attempt to address some of these issues. Chapter Two explores the potential for neurally-informed modelling to overcome issues in between-group comparison. Within this chapter, age differences in a classic speed-accuracy tradeoff are assessed using novel neurally-informed modelling methods to give new insight into the well-researched phenomenon. Chapter Three provides methods for investigating the role of within-trial noise in visual perceptual learning through the use of both behavioural and neural constraints. In this, it shows a lowering of within-trial noise through perceptual learning, with a neurally-constrained model indicating additional possible learning of task timings. Finally, Chapter Four highlights and explores core issues in the use of neural signals to constrain and refine models. Here, the well-studied effects of ageing on P3b/CPP localisation are explored, with specific attention paid to the often-overlooked influence of overlapping signals. From this, a means of selecting spatially independent neural markers of decision formation through ICA are derived. Chapter Five concludes with an overarching view of the findings of this body of work. It is hoped that through this body of experiments, an exciting new approach to the synthesis of behaviour and electrophysiology can be exhibited. With this new knowledge, the work concludes with an assessment of some remaining questions and issues for neurally-informed modelling, alongside avenues for further research.

Chapter Two:

Neurally-Informed Modelling of Ageing Effects on the Speed-Accuracy Trade-Off

2.1: Introduction

Oftentimes when making decisions, there exists a tradeoff between two key considerations; the speed with which you can make the decision, and the level of accuracy you wish to achieve. The mind must flexibly adapt and reweight the relative importance of each under different circumstances. For example, if we are driving at speed, it is more important to react quickly and apply the brakes than to be certain of the nature or presence of an obstruction because a crash at speed is risky and the benefits of slowing down to avoid an obstruction at all, even if there was none, is more important than being entirely certain. Conversely, when accuracy is of paramount importance, for example when choosing which car to buy in the first place, we may accumulate more evidence over a longer period of time, leading on average to more accurate but slower responses. Our ability to actively assess and adjust our response styles under different conditions is a core facet of decision-making behaviour. These adjustments to emphasise speed or accuracy, known as the speed-accuracy tradeoff (SAT), have been intensively researched (see Heitz, 2014 for a review), including several studies that have examined their susceptibility to natural ageing (Forstmann et al., 2011; Starns & Ratcliff, 2010, for a review see Dully, et al., 2018). A consistent finding in the literature on cognitive ageing is an apparent diminished capacity of older adults to flexibly adapt their decision strategies and response times when placed under increased speed pressure (Forstmann et al., 2011; Starns & Ratcliff, 2010). This is a key insight for our understanding of the brain across the lifespan, as it represents a potential reduction in the brain's capacity to adapt optimally to task-relevant demands with age (Starns & Ratcliff, 2012). However, it has alternatively been argued that this finding may instead represent a strategic difference, whereby older adults choose to pursue a more deliberate and cautious response style, either out of preference or in order to compensate for impaired perceptual or motor processes (Dully et al., 2018). Arbitration between these alternative accounts has important ramifications for our understanding of the brain as we age, and significant progress has already been made in teasing apart these explanatory accounts, using a combination of mathematical modelling and neurophysiological analyses.

2.1.2: Mathematical Models of the Speed-Accuracy Tradeoff

Sequential sampling models, as introduced in Chapter One, allow us to utilise reaction time (RT) and accuracy data to investigate mathematical approximations of latent cognitive processes and to assess how those processes may change under different conditions (Forstmann et al., 2016; Ratcliff & McKoon, 2008; Ratcliff & Smith, 2004). Sequential sampling models posit that for a two-choice decision, evidence accumulates noisily across time until the crossing of some decision bound, with added non-decision time to allow for response execution and sensory encoding (Brown & Heathcote, 2008; Ratcliff & McKoon, 2008, see Chapter One section 1.2). In sequential sampling models under speed emphasis, decision bounds are lowered to facilitate faster responses at the cost of accuracy. Behaviourally, this lowering of bounds under speed pressure has been modelled extensively across a range of tasks and species (see Heitz, 2014 for review). A mathematically equivalent alternative to this lowering would be an increase in the starting point of the decision variable, which while performing functionally the same as a boundary reduction, implies a strategic pre-stimulus adaptation (Kelly et al., 2021). For example, Steinmann and colleagues (2018) have shown that beta-band desynchronisation, which builds to a threshold at response (Donner et al., 2009; Fischer et al., 2018; Murphy et al., 2016, see Chapter One 1.2.3), is lower following a speed emphasis cue well in advance of stimulus onset, indicative of an increased pre-evidence motor preparedness under speed pressure. In modelling terms, this shift in starting beta levels is reflective of starting point adjustments, where the decision process begins closer to its bound. Similar pre-evidence adaptations under speed pressure have also been observed in the monkey frontal eye field (FEF) and lateral intraparietal area (LIP; Hanks et al., 2014; Heitz & Schall, 2012). Functional magnetic resonance imaging (fMRI) studies in humans have linked increased starting level adjustments under speed pressure with activity in the pre-supplementary motor area (pre-SMA; Forstmann et al., 2008), with the capacity to adapt these thresholds governed by cortical connectivity between the pre-SMA and the striatum (Forstmann et al., 2010). This neural evidence is valuable, as it shows that this adaptation to the decision process begins well in advance of any evidence appearing on screen.

A range of behavioural experiments have investigated the SAT specifically in older adults. As mentioned, older adults are routinely shown to be less capable of adapting these decision bounds under speed emphasis (Forstmann et al., 2011; Starns & Ratcliff, 2010). In their work, Starns and Ratcliff (2010) calculated the Reward Rate Optimal Boundaries (RROB), which indicate the optimal boundary setting for an ideal observer under speed or accuracy emphasis and compared these to model estimated boundaries for younger and older adults across the two regimes. Here, older adults show suboptimal bound settings, consistently setting higher bounds than necessary, and failing to adapt them sufficiently under speed emphasis to more closely align with the optimal bound. Neural data suggests that this impairment may be underpinned by the degradation of white matter tracts connecting the pre-supplementary motor area and striatum in older adults, which correlated significantly with more cautious boundary setting (Forstmann et al., 2011).

While some previous studies of the SAT have used simpler linear ballistic accumulator models (e.g., Forstmann et al., 2008, 2010, 2011), the drift-diffusion model (DDM) is uniquely valuable, as it proposes a more comprehensive set of parameters which are relatable to psychological constructs and distinct stages of the decision-making process that are supported by neural data (see Chapter One, section 1.2). However, extant studies of the SAT in ageing have failed to link DDMs to observed neural data (e.g., Starns & Ratcliff, 2010, 2012). As a result, the model may not directly reflect how the decision unfolds in the brain and may overlook important components of the decision process, such as urgency. Neural data grants us a means of linking model parameters to psychological constructs, allowing us to deliberate between competing models, and providing another line of evidence for novel parameters that may better capture how decisions are formed (O'Connell et al., 2018, see Chapter One: 1.3). For example, neural data suggests that the DDM may not account for important components of the decision process which are important in speed-accuracy paradigms, such as urgency (Hanks et al., 2014; Steinemann et al., 2018; Devine, 2019). To date, the role of urgency adaptations in older adults under speed emphasis has not been investigated, and so the extent to which the DDM provides a comprehensive account of the effects of ageing on the SAT in older

adults is uncertain. As a result, current models may ascribe the effects of ageing to the wrong parameters, leading to fundamentally different conclusions on the effects of ageing on decision-making and ultimately, different therapeutic approaches.

Analysis of electroencephalography (EEG) signals which reflect each distinct stage of perceptual decision-making can enable us to deliberate more astutely between computing models, potentially furnishing evidence for the inclusion of often overlooked parameters such as urgency. Furthermore, utilising neural constraints allows us to investigate increasingly complex models without incurring a cost to model comparison scores (Kelly et al., 2021, Corbett et al., in press). This can allow us to explore the potentially more nuanced effects of ageing on the SAT, whilst also allowing for the investigation of between-group differences in the traditionally fixed within-trial noise scaling parameters, which itself may change as a function of age (Pardhan et al., 1996; Kumral et al., 2020; Sleimen-Malkoun et al., 2015). Using a neurally-informed modelling approach may reveal novel insights into how older adults adapt their decision-making under speed-emphasis, granting a more biologically grounded, neurally-driven account for the effects of ageing on the SAT.

2.1.3: Issues in Modelling the SAT

To date, the majority of sequential sampling models rely on the assumption of a static decision bound, meaning one which remains set at a constant rate from the beginning to the end of a trial (Ratcliff & McKoon 2016; see Chapter One: 1.2.2). Importantly a growing body of evidence suggests that rather than a fixed bound, the amount of evidence required to reach commitment may vary throughout the course of a decision (Hawkins et al., 2015; Voss et al., 2019). In this way, the bounds "collapse" and the amount of evidence required to make a decision reduces over time, resulting in slow errors that reflect an urgency to respond before a trial ends. However, some authors suggest that these changes can be explained without the implementation of collapsing bounds through between-trial variations in drift rate (Ratcliff et al., 2016; Ratcliff & McKoon, 2008). Furthermore, some have suggested that even when optimal, collapsing bounds may not be an ecologically valid construct, and may only be the result of highly specific task dynamics, such as reward rates or the large amounts of practice completed in advance of data collection (Hawkins et al., 2015). In this way, they suggest that collapsing boundaries may not represent a meaningful dynamic of the decision process in the real world, but rather a fallout from the highly constructed nature of psychological tasks. Irrespective of the ecological validity of collapsing bounds, the fact remains that in perceptual decision-making tasks, trials are often constructed with these strict cut-off points and deadlines for responding. Behavioural evidence has shown collapsing bounds to be optimal, especially in difficult tasks, or those with predictable deadlines such as the paradigms traditionally used in the SAT (Frazier & Yu, 2007; Malhotra et al., 2018). Here it is demonstrated that an optimal observer would be better to respond using the evidence they have accumulated rather than suffering the cost of failing to respond entirely (Malhotra et al., 2018). These urgency dynamics become increasingly important when participants are instructed to prioritise speed. With speed emphasis, the cost of missing is often greater, and the rate of reward is often yoked to the reaction time, making a more urgent decision style optimal.

Evidence in support of collapsing bounds and urgency dynamics has been shown in the neural data. For example, in a study by Hanks and colleagues (2014), the level of firing in the LIP neurons of macaques was shown to exhibit start point adjustments under speed pressure and contain an additive, evidence-independent component which was more prominent under speed emphasis. Using this component to constrain an urgency parameter in their model provided stronger fits to the behavioural data, highlighting the role of dynamic urgency in the task. In humans, Kelly and colleagues (2021), increased motor preparedness under speed emphasis was shown to reflect pre-evidence starting point adjustments, suggestive of an evidence-independent urgency component that brings the decision-making process increasingly closer to its bound well in advance of any evidence appearing on screen. As in Hanks and colleagues (2014) study, including this urgency in the model greatly improved model fits. Notably, it also improved the concordance between model-simulated and real neural signals of the decision process, highlighting again the utility of using neural data to refine computational models of decision-making. Recent work from Devine (2019) illustrated the contingent negative variation (CNV) as another potential neural marker of urgency. Here, it was shown to build at an evidence-independent, but time-dependent rate that increases with speed pressure. An important distinction between the CNV and mu/beta motor preparation as a marker of urgency is that motor preparation scales with evidence strength, while the CNV does not. That means the CNV may provide a more accurate readout of urgency, unaffected by any effects of evidence strength. Finally, in Steinemann and colleagues (2018), the amplitude of the CPP at response was shown to reduce for later reaction times, indicating a lower amount of evidence accumulated by the time of response, and reflecting a reduction of the decision boundary across time. Collectively, these neural features indicate and are mathematically equivalent to collapsing bounds, highlighting an urgency to respond which is increased under speed pressure in a range of signals across both monkey and human work.

While these urgency dynamics have been exhibited both in behavioural modelling and neural data, current DDMs often fail to incorporate collapsing bounds into their simulations. Specifically, the influence of these dynamics in

older adults remains uninvestigated, especially in the case of SAT tasks, where urgency dynamics may play a particularly important role. This is important, as failing to incorporate urgency as a parameter raises questions as to the validity of findings in the modelling studies of cognitive ageing. For example, while current literature points to a failure to adapt their boundary setting (Forstmann et al., 2011; Starns & Ratcliff, 2010), behavioural differences in ageing under speed pressure may instead reflect a failure to adapt their urgency dynamics to ensure a response is made in time. Even if not ascribing the adaptation solely to changes in urgency, failing to incorporate urgency as a parameter when it is a true psychological phenomenon may impact other parameters. Here, a model without urgency is forced to adapt other parameters to account for the behavioural effects of a collapsing bound. Therefore, urgency dynamics must at least be explored in the context of SAT with ageing, especially as traditional SAT tasks rely on hard deadlines and time-based rewards under speed emphasis, which may elicit increased urgency (Devine, 2019; Kelly et al., 2021; Steinemann et al., 2018). Examining neural signals to establish the potential presence of urgency can provide additional evidence for its necessity, and its inclusion in a model may reveal new insights into the SAT in older adults.

Beyond the inclusion of novel parameters, neurally-informed modelling offers an improvement to the use of computational modelling alone and may help to advance our understanding of the effects of ageing on decision-making. Through this, the comparison of model parameter estimates and neural signatures across age groups may provide valuable information into the complex cognitive processes which give rise to the effects of ageing on a range of phenomena. For example, recent work from McGovern and colleagues (2018) showed older adults outperformed younger adults on a contrast-change detection task. A standard DDM attributed this to higher drift rates and decision bounds in the older adult group. However, this did not align with the observed neural data. In contrast, the CPP as a marker of cumulative evidence was shown to accumulate at the same rate to the same threshold at response across older and younger adults. This contradicted the model, which estimated higher drift rates and bounds in older adults. The neural data was then used as a constraint to enforce consistent bounds and drift rates in a neurally-informed model. This model then showed that

the poorer performance in younger adults was driven by increased trial-by-trial drift rate variability. This novel finding was in turn supported by the neural data, which showed increased variability in younger adults' CPP slopes. Furthermore, this variability was attributed to increased variability in attentional engagement in the task, which was supported by increased variability in between-trial levels of posterior alpha-band activity, an indicator of attentional engagement (Dockree et al., 2007). Here, a standard DDM attributed the performance difference in older adults to a more cautious response criterion and better sensory representation, while the neural data created and supported a model which indicated that the improved performance was likely driven by increased attentional engagement in the older cohort. These insights were unachievable through behavioural modelling alone.

Collectively, this presents compelling evidence for the application of neurally-informed modelling to the study of the SAT in ageing, allowing us to investigate the role of novel parameters such as urgency, to constrain our models using observed neural data in order to investigate more nuanced models, and to deliberate on competing model accounts based on their capacity to reflect neural indices of decision making. This has the potential to lead to significant advancement in our understanding of the effects of ageing on key facets of decision-making such as the SAT.

2.1.4: Issues in Between-Group Comparison of Neural Data and Model Scaling Parameters

Neurally-informed modelling is, therefore, an extremely beneficial means for investigating the effects of ageing on decision-making. However, to enable accurate and reliable insights, we must be certain that the signals we use for constraint and comparison represent true group differences and are not unduly affected by behaviourally-irrelevant features that may affect their presentation uniquely in each cohort. For example, neural signals can present with varying strengths across groups for reasons which may be completely unrelated to the decision-making process. This may be of particular importance in the study of ageing, where changes in physiological features such as skull thickness with age have been shown to result in reduced skull conductance (Antonakakis et al., 2020). This may have a direct influence on how signals present on the scalp, without reflecting any changes to the underlying psychological construct that the signal is thought to represent. For example, baseline beta power has also been shown to be consistently increased in older adults (Stacey et al., 2021). In a recent study by Rossiter and colleagues (2014), older adults showed greater beta power at rest alongside a greater degree of beta desynchronisation during movement. The authors suggest that this is due to stronger levels of GABAergic inhibition in the older adult motor cortex. However, as outlined in Chapter One, pharmacologically elevating GABA levels in younger adults increases baseline beta levels but has no effect on their actual movement (Muthukumaraswamy et al., 2013). This may suggest that elevated beta-amplitudes in older adults may not reflect any functional difference but are due to behaviourally-irrelevant ageing effects. Failing to account for these may lead to incorrect conclusions as to age differences in, for example, pre-evidence motor preparation, a feature of particular importance for studies of the SAT (e.g., Steinemann et al., 2018). Fortunately, as mu/beta is shown to exhibit stereotyped accumulation to threshold dynamics (Donner et al., 2009; Fischer et al., 2018; Kelly et al., 2021; Murphy et al., 2016; O'Connell, Dockree, et al., 2012; Rogge et al., 2022), this signal can be normalised across conditions and groups, reducing the influence of behavioural-irrelevant group differences (Kelly et al., 2021).

The DDM may also face scaling issues which are not traditionally addressed and may present significant challenges to between-group comparisons of model parameters. As outlined in Chapter One, behavioural models rely on a scaling parameter, a fixed value around which other values are estimated (section 1.3.3). This is traditionally pinned as the standard deviation of the Gaussian within-trial noise added at each time point, which is fixed by convention at 0.1 across conditions, groups and individuals (Ratcliff & McKoon, 2008; Ratcliff & Smith, 2004; Voss et al., 2004, see section 1.3.3). Any comparison of parameter estimates across groups therefore inherently relies on the assumption that within-trial noise levels do not change. However, the validity of this assumption is called into question by a range of recent evidence, especially when studying the effects of ageing on decision-making. For example, while behavioural estimates of internal noise have pointed to no differences in internal noise levels across older and younger adults (Pardhan et al., 1996), EEG signals in the older adult brain have been shown to exhibit higher complexity in short time scales both at rest and throughout a task (Sleimen-Malkoun et al., 2015). However, when considered across longer time scales, older adults exhibit reduced variability in both EEG and fMRI BOLD signals (Kumral et al., 2020; Sleimen-Malkoun et al., 2015). This may be indicative of increased within-trial noise, contributing to increased ongoing neural complexity, with reduced inter-trial variability in parameters across time. If this increased ongoing variability in older adults is true, any behavioural modelling which assumes within-trial noise values to be consistent across age groups is false and may attribute behavioural effects to differences in other parameters, when in fact they are driven by group differences in within-trial noise. However, due to their use of within-trial noise as a consistent scaling parameter across groups, the DDM is not currently equipped to investigate these differences.

An alternative may be to fix the decision bound as the scaling parameter (Ratcliff & Tuerlinckx, 2002), yet this would force the boundary to remain constant across speed-accuracy regimen and may also obscure the routinely evidenced cross-condition and age effects on decision bound setting (Forstmann et al., 2011; Starns & Ratcliff, 2010). This is a key limitation of standard modelling approaches, as they rely on the assumption that one parameter will not change

between conditions, and we can rarely do this with certainty. One solution to this issue is to use neural data to constrain a parameter. As it is a fixed value, this neurally-constrained parameter then acts as a scaling parameter, thereby freeing up within-trial noise to vary, while maintaining space for potential adjustments to the decision bound. Recent work from Kelly and colleagues (2021) utilised pre-evidence mu/beta motor signal levels to great effect to constrain diffusion models. As model comparison scores penalise for model complexity, these constraints allowed the model to capture the decision-bound adjustments without the need for them to be freely estimated by the model, allowing the model to take on additional free parameters and detect a wider range of parameter effects than previously possible. Similar results were achieved in the McGovern and colleagues (2018) ageing study outlined above, however, neither of these studies utilised the neural constraints in place of traditional scaling parameters. As such, they did not investigate potential changes in within-trial noise as a driver of behavioural performance.

The use of these constraints may be uniquely beneficial to the study of between-group effects such as the SAT in older adults. Here, the groups could be linked using scaling parameters directly derived from their neural data, allowing for more comparable models and parameter estimates. It also has the added benefit of freeing up within-trial noise to be estimated as a free parameter, allowing the impact of ageing on this parameter to be reasonably assessed. However, given the expressed differences in beta-power outlined above (Rossiter et al., 2014), a means of rescaling mu/beta across groups is required, which retains the relative differences between conditions. If this could be achieved, the use of a neurally-derived scaling parameter would represent a truly unique and beneficial tool for between-group modelling, which has direct implications not only for the study of ageing but could easily be extended to the study of a range of clinical disorders such as autism, schizophrenia and ADHD.

2.1.5: Present Study

The present study aimed to thoroughly investigate the speed-accuracy tradeoff in older adults using behavioural modelling informed by electrophysiologically

derived indices of the decision process in a data set previously collected by Dully (2020). Neurally-informed modelling presents a unique opportunity to investigate these effects, enabling parsimonious explanations which synthesise behavioural and biological evidence sources. Using neural signals to establish plausible and biologically likely effects allows for the examination of more complex models without increasing the number of free parameters and thus penalising model comparison scores. Neural constraints also have the benefit of providing neurally derived and applicable scaling factors which allow for direct between-group comparison. Through this it was hoped to provide a more biologically grounded model of the speed-accuracy tradeoff in ageing, giving new insight into the effects of ageing on this well-studied phenomenon.

2.2 Methods

2.2.1: Participants

A total of 30 younger adults, (18-34, mean age = 23 ± 4.4 years), and 30 older adults (62-77, mean age = 70 ± 3.5 years) participated in this study. Eleven participants were excluded from the analysis for failing to achieve above-chance accuracy (>60%) in either of the Speed-Accuracy conditions (6 younger adults, and 5 older adults). A post-hoc power analysis was performed using G*Power (Faul et al., 2007). This indicated that given a sample size of 49 across two groups and two conditions, with a significance criterion of 0.05 and a power of 80%, the study would have sufficient power to detect effects of small to medium effect sizes (f = 0.2) . Participants had normal, or corrected-to-normal, vision and were screened for personal and familial history of epilepsy, psychiatric and neurological disorders, and personal history of traumatic brain injury and photosensitivity. The groups were matched for years of formal education (Young: Mean = 17.2 +/- 3.1 years; Older: Mean = 17 +/- 3.8 years; p = 0.38) and all participants were required to have a score over 24 points on the Mini-Mental State Exam. The National Adult Reading Test (NART; Nelson, 1982) was used as a measure of premorbid IQ. Age groups did not differ on this measure. All participants were required to give written informed consent prior to any data collection and were remunerated €30 in lieu of their time and travel expenses. The study was approved by TCD's School of Psychology Ethics Committee, in line with the Declaration of Helsinki.

Upon granting informed consent, participants completed the demographic, NART and MMSE components. Participants then completed the task as outlined below in a sound-attenuated, darkened EEG booth. The task was presented on a 40.5cm CRT monitor, 50cm from the participants, who sat in headrests to decrease movement and keep viewing distance consistent. The task was presented at a refresh rate of 100Hz on a screen with a resolution of 1024 x 768.

2.2.2: Contrast Change Discrimination Task:

Participants were asked to fixate on a central point which appeared for 500ms at the beginning of each trial, followed by a cue indicating the required strategy (“go fast!” / “be accurate”). This cue was displayed for 500ms, followed by a further 500ms display of the fixation point. Overlaid left/right gratings flickering at either 20/25 Hz were displayed with equal contrast (50%) for 400ms. After this period, one grating increased in contrast to 60% or 66% (Low/High contrast condition), as the other reduced by the same amount. Participants held a mouse in both hands and were asked to indicate which of the two gratings had increased in contrast, a left click with the left hand for leftward tilted gratings, and a right click with the right hand for rightward tilted gratings (See Figure 2.1). Each block contained 30 low-contrast trials (60%) and 30 high-contrast trials (66%). The flicker frequency of the target grating was balanced. Participants completed 3 blocks of 60 trials under Accuracy emphasis, and 3 blocks of 60 trials under Speed emphasis. The order of these blocks was alternated (e.g., Speed, Accuracy, Speed etc.) and the condition of the first block in the sequence was counterbalanced by participants. A point system was used to encourage adaptation to the required response style. Under speed emphasis participants were rewarded a diminishing number of points over time, reducing from 100 points at a rate of 75 points/second. Errors also incurred a point deduction, rising from 20 points at a rate of 4 points/second. Failure to respond before the deadline incurred a deduction of 116 points. In the accuracy condition, 60 points were awarded for correct and 60 were deducted for incorrect responses regardless of response time. There was no penalty for missed responses (0 points). Feedback was given after each trial, including the number of points awarded. Before commencing the testing session, participants completed an initial practice block of 20 trials at which targets were presented at 100% contrast. They then proceeded to a further practice block in which stimuli were presented at the required 60/66% contrasts until the participant had fully understood the task.

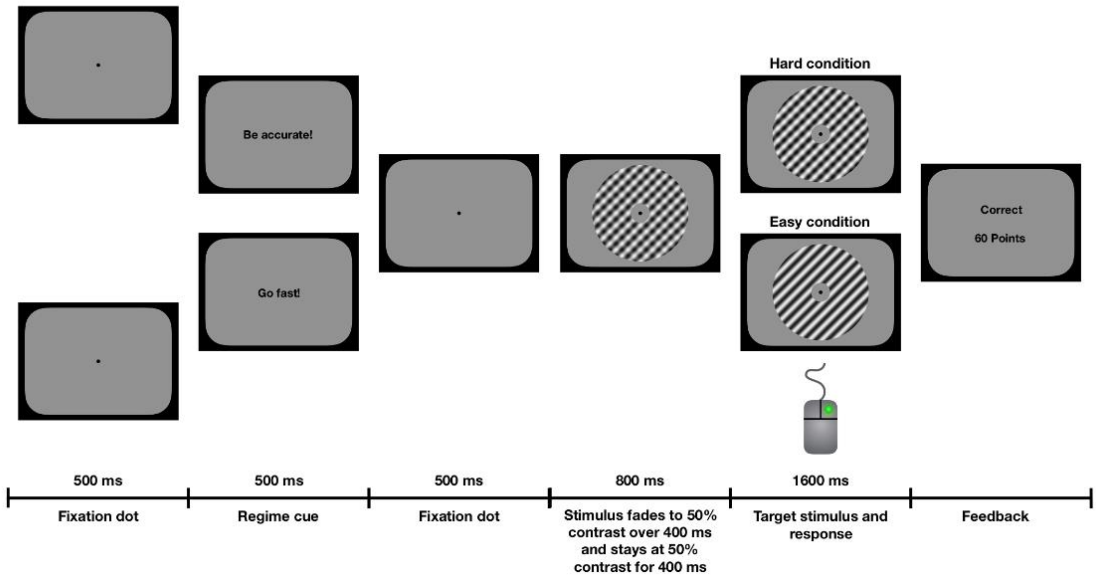

Fig. 2.1: Task schematic: from Dully (2020). Participants are shown 500ms of a fixation dot, followed by a regime cue. Another 500ms of the fixation dot is followed by the stimulus fading into 50% contrast across 400ms. The stimulus remains at 50% contrast before the target stimulus emerges, at which point participants make a button press to indicate the direction of the high-contrast grating. Participants are given feedback at the end of every trial. Reprinted from Dully, J. (2020). *The Impact of Natural Aging on Perceptual Decision Making* [Thesis, Trinity College Dublin. School of Psychology. Discipline of Psychology]. http://www.tara.tcd.ie/handle/2262/94134

2.2.3: EEQ Acquisition and Pre-processing

The task was performed while 128-channel EEG was recorded using an ActiveTwo system (BioSemi), with a 512Hz sampling rate. Data were analysed using a combination of customised MATLAB scripts implemented using some features of the EEGlab toolbox (Delorme & Makeig, 2004). A Low-Pass filter at 35Hz was applied to the data offline. To enable the reliable extraction of slow-emerging waveforms such as the CPP, no High-Pass filter was used. Data were detrended to remove linear slow drift and channels with excessive noise (>= 3 *SD* away from normalised joint probability) were interpolated. A secondary visual inspection was performed to identify and interpolate any additional noisy channels. These channels were marked and interpolated. After interpolation, an average reference was applied across all channels and time points. The data were epoched in windows from -500:1800ms centred on target onset. Baseline correction was applied, subtracting the mean activity across each electrode in the time period of -200:-100ms from the target onset. This baseline was chosen

to allow for the pre-stimulus build-up of decision signals. Trials with channel activity exceeding 100uv between target onset and 100ms after response were labelled as artefacts and removed from EEG analyses and plotting. Response-locked epochs were extracted in a time range of -600 to 200ms centred around response. As a result of having two overlaid gratings flickering at either 25 or 20 Hz on every trial, all ERP plots (CPP, CNV) below have a 5Hz Boxcar filter applied for plotting. This gives a clearer insight into the development of the signal over time, removing fast-moving, task-irrelevant oscillations. All statistical analyses of the CPP and CNV have been performed with these filters applied.

2.2.4: EEG Signal Electrode Selection

Extant literature indicates an anterior shift of the P3b/CPP topography in older adults, potentially due to changing brain function, with the compensatory recruitment of more frontal regions for cognition in older age (Li, Gratton, Fabiani & Knight, 2013). In order to ensure appropriate electrode sites were chosen to extract the CPP signal, unique electrodes were chosen for individual subjects. Electrodes were selected as the site with the maximum amplitude -50ms to 50ms from the response (see Figure 2.2.A) along a central “line” of electrodes covering traditional central-parietal CPP sites, along with more frontal electrodes. This central line was chosen to minimise the influence of lateralised motor activity, with the most frontal electrode sites excluded from the selection procedure as they may be uniquely affected or contaminated by extraocular muscle activity.

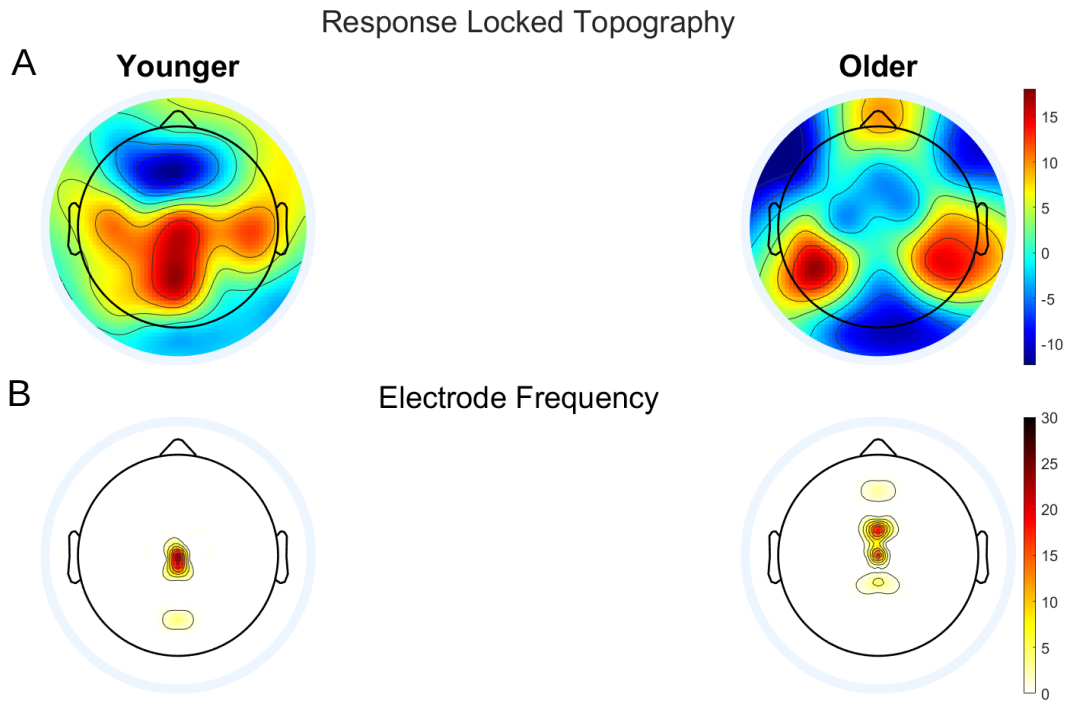

Fig. 2.2. A) Scalp topography of 100ms window centred on response showing more dispersed and frontal activity in older adults. B) Frequency map of chosen electrodes for CPP analysis showing more frontal presentation in older adults.

A similar individual electrode selection procedure was used for selecting CNV electrodes as was used for the CPP. CNV electrodes were selected as the two electrodes with the most negative going slope from -400 to -100 from the response (see Figure 2.3.A), within a cluster of frontal electrodes in line with previous literature (Devine et al., 2019). Selected electrode frequencies are shown in Figure 2.3.B.

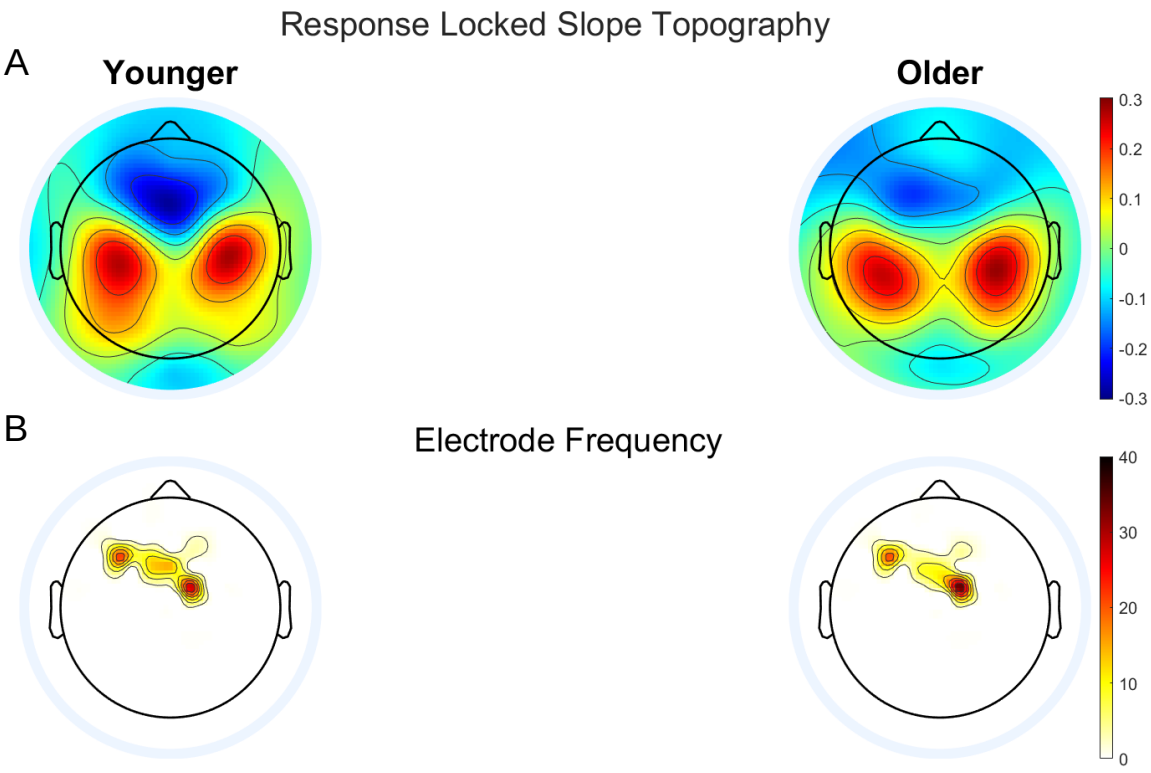

Fig. 2.3.A) Scalp topography of slopes in the -400:-100ms window centred on response. B) Frequency map of chosen electrodes for CNV analysis.

Frequency tagging of target vs. non-target stimuli at either 25 or 20Hz, allowed for the extraction of a steady-state visual evoked potential, which tracks the encoding of sensory evidence in the occipital cortex across time. A fast-Fourier transform of window length 400ms, iterating in 20ms-step windows was used to extract the power for each frequency. To further discretise the 20 and 25hz frequency bands, the power of each frequency was normalised by the mean power of its immediately neighbouring frequency bands. The two frequency bands were then collapsed into target and non-target trials, averaging the power across the two signals on trials where they were either the target or non-target. Two electrodes were selected for each group, choosing the site with maximal target vs non-target power difference over the occipital cortex based on the stimulus-locked topography (300:500ms post-evidence onset) shown in Figure 2.4.

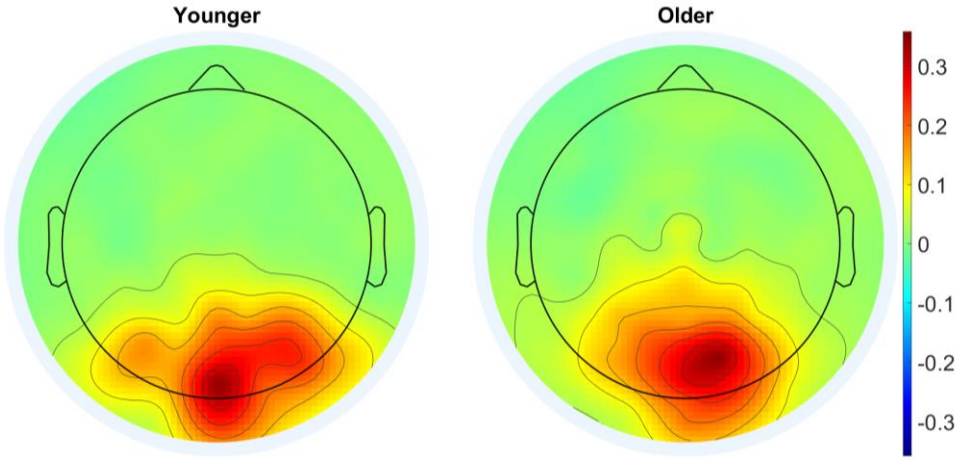

Fig. 2.4. A) Scalp topography of differential SSVEP in 300:500ms from evidence onset showing similar occipital presentation for both groups.

Activity in the mu/beta frequency range of 8-30Hz was used to index emerging motor response selection (see section 1.2.3). A fast-Fourier transform of window size 400ms and step size of 25ms was used. Mu/beta electrodes were selected for each group based on the grand average topographies shown in Figure 2.5.

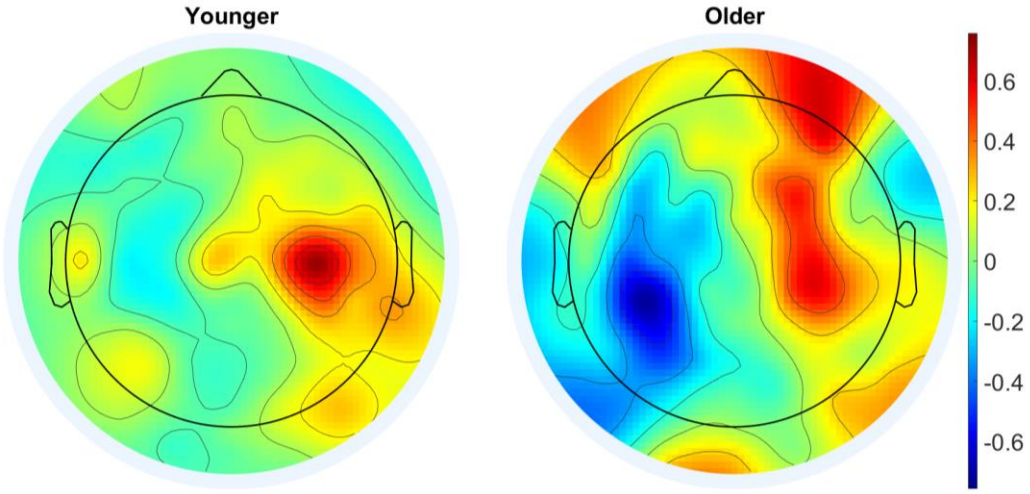

Fig. 2.5. Scalp topography of difference in mu/beta amplitude for left minus right trials in the -50:50ms centred on response.

2.2.5: Analysis Procedure

Statistical analyses were performed using JASP, and plotting was completed using MATLAB custom codes and JASP. Error bars shown represent the standard error of the mean. Where no interaction terms are reported the interaction was non-significant. Where Mauchly's test of sphericity indicates the assumption of sphericity is violated, Greenhouse-Geisser corrected degrees of freedom and statistics are reported.

2.2.6: Modelling Procedure

Individual data were pooled within each age group for initial model refinement procedures. The goodness of fit was quantified as the chi-square error between real and simulated group reaction-time quintiles (0.1, 0.3, 0.5, 0.7, 0.9). Parameters were optimised using a Bayesian Adaptive Direct Search algorithm (BADS; Acerbi & Ma, 2017), which combines local Bayesian search optimisation with a polling stage which uses a slower mesh grid search. Akaike's Information Criterion was calculated and compared to enable model selection models, penalising more complex models in favour of more parsimonious ones. To find initial parameter starting vectors, parameter values were drawn randomly within pre-established probable ranges, and 1000 trials were simulated. Behavioural features of simulated data were compared to the true data and parameter sets were accepted for use as starting points if real and simulated mean RTs differed by no more than 20% and real and simulated standard deviations by 0.25. Once 30 starting parameter vectors were found, individual trials were simulated by a factor of 100 for each quantile, giving 7,200 trials simulated per iteration, and these probable starting parameters were optimised using BADS.

2.3 Results

2.3.1: Behaviour

Analysis of behavioural data indicated a significant main effect of Speed emphasis on accuracy ($F_{(1,47)} = 28.424$, $p < 0.001$), but no significant main effect of age ($F_{(1,47)} = 1.547$, $p = 0.220$) or interaction term ($F_{(1,47)} = 1.250$, $p = 0.269$), indicating a comparable lowering of mean accuracy under increased speed pressure in both groups ($Mean_{Speed} = 81.56\%$, $SD = 9.8$, $Mean_{Acc.} = 85.71\%$, $SD = 9.28$, Fig 2.6.a).

A significant shortening of reaction time was observed under increased speed emphasis ($F_{(1,47)} = 46.574$, $p < 0.001$). There was no significant main effect of Age ($F_{(1,47)} = 0.155$, $p = 0.696$) on reaction time but a significant Condition-Age interaction was observed ($F_{(1,47)} = 7.88$, $p = 0.007$). Post-hoc analyses indicated an effect of Condition on RT on both age groups ($t_{Older} = 6.882$, $p = 0.001$, $t_{Younger} = 2.811$, $p = 0.029$) but the RT adjustment was more than twice as large in the younger group ($Mean\ Difference_{Younger} = -0.13$, $SD = 0.11$, $Mean\ Difference_{Older} = -0.05$, $SD = 0.07$, Fig 2.6.b).

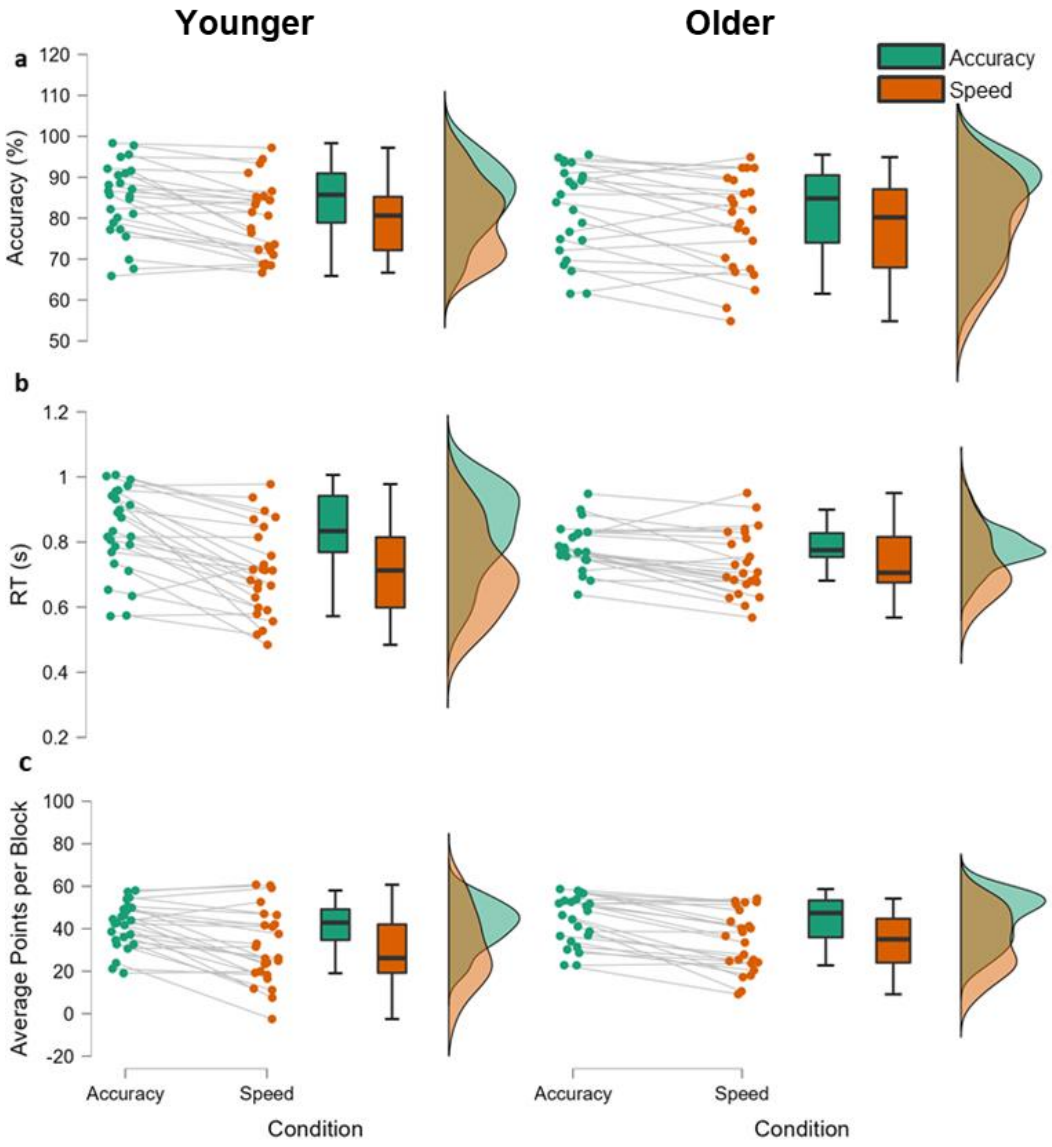

Fig. 2.6 Plots of behaviour showing raincloud plots, boxplots and probability distributions. The left column shows younger adults, the right column shows older adults. a) Accuracy % b) Reaction time in seconds c) Average points scored per block. Older adults did not decrease their RT as much under speed emphasis as younger adults. There were no group differences in accuracy. There was also no main effect of age on the number of points scored.

Analysis of the effects of Speed emphasis on points scored indicated a significant effect of Condition ($F_{(1,47)}$ = 48.544, p <0.001), but no significant effect of Age ($F_{(1,47)}$ = 0.583, p = 0.449) or interaction ($F_{(1,47)}$ = 0.003, p = 0.956), indicating a lower overall score in the speed Condition that was consistent across the older and younger groups ($Mean_{Acc.}$ = 42.846, SD = 11.135, $Mean_{Speed}$ = 32.476, SD = 15.691, Fig, 2.6.c).

2.3.2: Standard Model Fits

Preliminary analyses used a full DDM, with between trial variability parameters for non-decision time. Initially, a “no-change” model with no parameter varying by condition was run. Subsequent variants of this model were run, allowing non-decision time or bound to vary by Speed/Accuracy condition using the best fitting parameters from the no-change model as starting parameters. To allow the drift rate to vary by condition, a freely estimated drift rate was calculated for the Accuracy condition, with a multiplicative ‘drift-boost’ parameter applied to this drift rate to calculate the drift rate of the Speed Condition. Initial fits of the full DDM were poor for all model variants (AIC Range: 35.9 - 66.47, see Table 2.1) The model in which non-decision time was free to vary provided the best fit for both groups (Young: G^2 = 30.54, AIC = 46.54, Older: G^2 = 19.9, AIC = 35.9). The different models provided very similar fits in the older group with the non-decision time-varying model only slightly outperforming the others. This model suggests that older adults show shorter non-decision times in the accuracy condition, higher bounds overall and a smaller adaptation to non-decision time under speed emphasis (see Table 2.2). From an examination of the real-vs-simulated reaction time distributions and conditional accuracy functions, it is evident that the full-diffusion model fails to capture fast errors in both groups, despite the inclusion of starting point variability (see Figure 2.7). The far larger proportion of slow errors under speed pressure in the younger adult group has not been captured even with the inclusion of drift rate variability.

Table 2.1: Table of standard DDM fits, with each model variant allowing each of the parameters in the leftmost column to vary by Speed-Accuracy condition. The green highlight indicates the winning model with the lowest AIC.

Standard DDMs		Younger		Older	
Parameter Varying	N parameters	G^2	AIC	G^2	AIC
No Change	7	50.41	64.41	23.05	37.05
Boundary	8	37.57	53.57	20.42	36.42
Non-Decision Time	8	30.54	46.54	19.90	35.90
Drift Rate	8	50.47	66.47	23.05	39.05

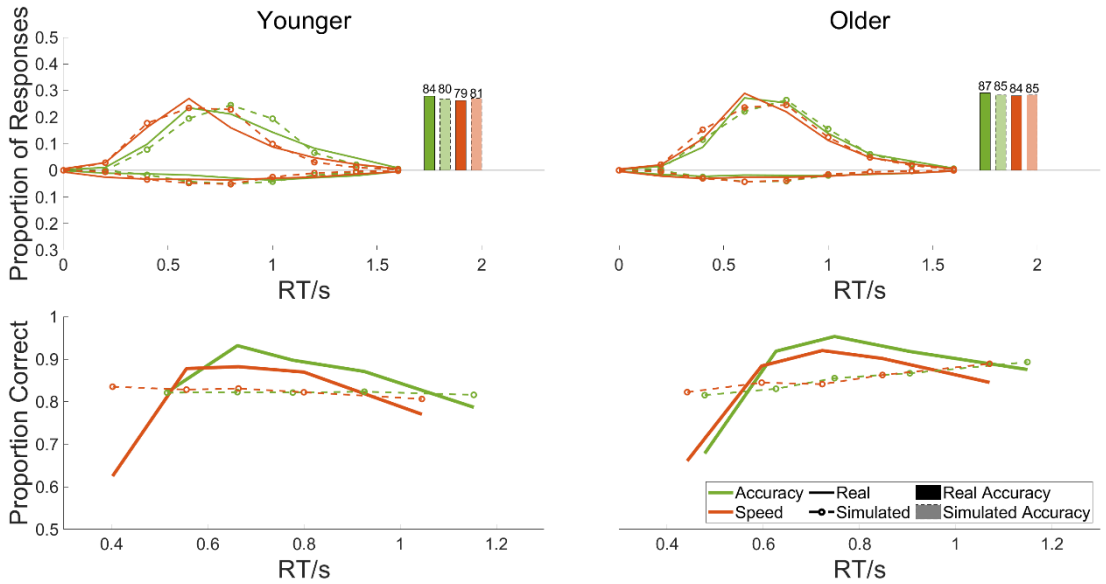

Fig 2.7: Model with non-decision time varying by condition. Top) Real vs model simulated RT distributions, with correct responses going up and errors going down. Bars show real accuracy per condition and semi-transparent bars show model-simulated accuracy. Bottom) Real vs simulated conditional accuracy functions. Responses are binned into quintiles by RT and the proportion of correct responses for each quintile bin. Full diffusion model incapable of capturing the u-shaped distribution of conditional-accuracy function over time.

Table 2.2: Model estimated parameters for the best-performing full diffusion model variant. For condition-specific parameters, the requisite condition is indicated in subscript.

	Bound	$Drift_{Low}$	$Drift_{High}$	Ter_{Acc}	Ter_{Spd}	Drift Var	Ter Var.	S.P. Var	G^2	AIC
Younger	0.056	0.148	0.242	0.61	0.476	0.127	0.308	0.028	30.54	76.71
Older	0.06	0.153	0.223	0.539	0.488	0.006	0.297	0.043	19.9	65.87

Collectively, initial applications of the full DDM suggest a winning model which is failing to capture a significant proportion of responses. It also produces a model which suggests changes to non-decision time under speed emphasis, rather than a well-documented boundary adjustment. However, the second leading model was one in which the boundary changed by condition. This model only scored 0.52 AIC points more than the winning non-decision time model in the older group and would be in line with the majority of extant literature. The results of this model are shown in Table 2.3. Here older adults showed larger bounds overall, with less adjustment and higher drift rates. This aligns with the findings of previous studies (Forstmann et al., 2011; Starns & Ratcliff, 2010)

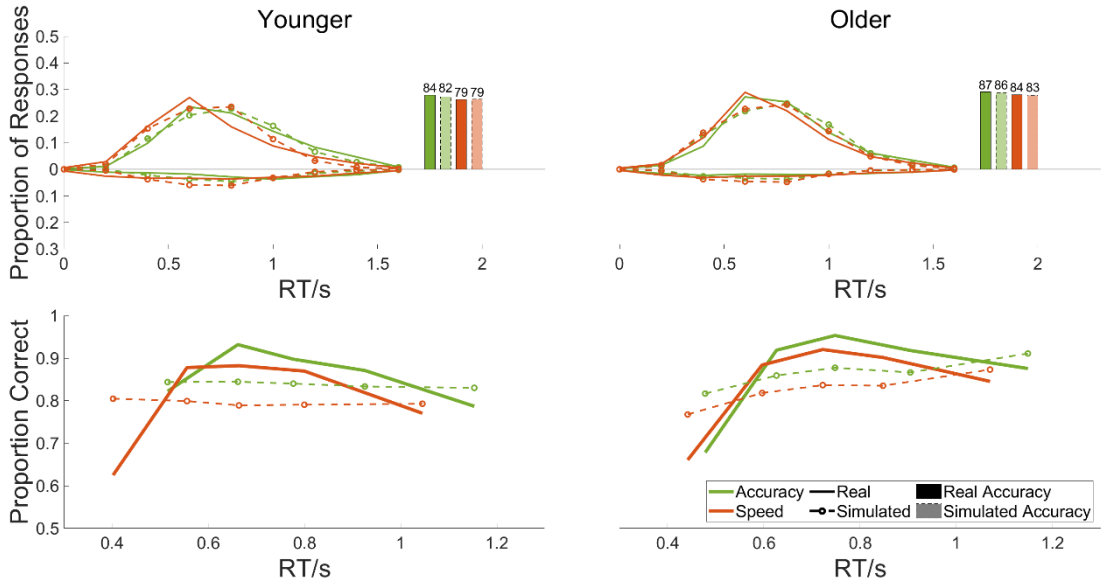

Fig 2.8: Model with boundary varying by condition and no parameters varying by condition. Top) Real vs model simulated RT distributions, with correct responses going up and errors going down. Solid bars show real accuracy per condition, and semi-transparent bars show model-simulated accuracy. Bottom) Real vs simulated conditional accuracy functions. Responses are binned into quintiles by RT and the proportion of correct responses calculated for each quintile bin.

Table 2.3: Model estimated parameters for the second-best performing DDM variant with boundary varying by condition. For condition-specific parameters, the requisite condition is indicated in subscript.

	$Bound_{Acc}$	$Bound_{Spd}$	$Drift_{Low}$	$Drift_{High}$	Ter	Drift Var	Ter Var.	S.P. Var	G^2	AIC
Younger	0.061	0.047	0.116	0.19	0.513	0.06	0.322	0.005	37.57	53.57
Older	0.063	0.058	0.18	0.269	0.547	0.047	0.326	0.051	20.42	36.42

2.3.3: Neurally-Informed Model Construction

In order to improve these modelling efforts and to create a model which reliably captured behaviour in a biologically plausible way, key neural signatures of the decision process were analysed.

2.3.3.1: Sensory Encoding

Figure 2.9 plots this difference in amplitude between target and non-target frequencies across time, averaged across trials, for each experimental condition.

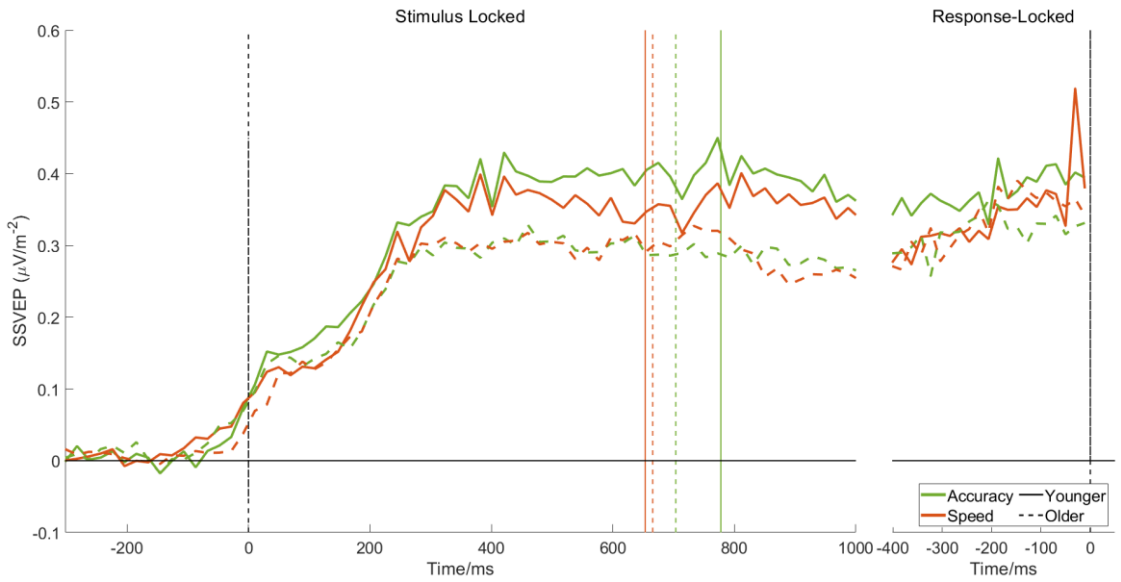

Fig. 2.9: Target minus non-target SSVEP activity in stimulus and response aligned waveforms showing largely similar patterns of emergence across groups and conditions

SSVEP traces appeared larger for younger adults than older adults. A mixed-methods factorial analysis of variance (ANOVA) indicated no significant

differences in pre-response amplitude (Age: $F_{(1,47)}$ = 0.006, p=0.937; Speed-Accuracy: $F_{(1,47)}$ = 2.341, p = 0.133; Interaction $F_{(1,47)}$ = 0.5124, p = 0.478). Analysis of slope in this window also indicated no significant main effect of Condition ($F_{(1,47)}$ = 0.038, p = 0.847), with no significant effect of Age ($F_{(1,47)}$ = 1.998, p = 0.164) or interaction ($F_{(1,47)}$ = 0.22, p = 0.641). There were no significant main effects of Condition or Age-Condition interaction on stimulus-locked amplitude or slope measure 300:600ms after evidence onset (all p > 0.257) This suggests consistency across groups and conditions in the integration of sensory information at this earliest perceptual processing stage.

2.3.3.2: Evidence Accumulation

A mixed-methods ANOVA was used to examine the effects of Age and speed-accuracy Condition on CPP features. Analysis of CPP amplitude in the 100ms window centred on response indicated that older adults had significantly less positive amplitudes but there was no effect of speed-accuracy Condition, or significant interaction (Age: $F_{(1,47)}$ = 12.604, p < 0.001; Condition: $F_{(1,47)}$ = 0.481 p = 0.491; Age-Condition Interaction $F_{(1,47)}$ = 0.226 p = 0.636). Analysis of the pre-response slope in the window of -400 to -100ms from response indicated a significant effect of Age with shallower slopes in the older group but there was no effect of speed-accuracy Condition, or significant interaction (Age: $F_{(1,47)}$ = 13.7, p < 0.001; Speed-Accuracy: $F_{(1,47)}$ = 0.004, p = 0.95; Age-Condition Interaction $F_{(1,47)}$ = 0.443 p=0.509).

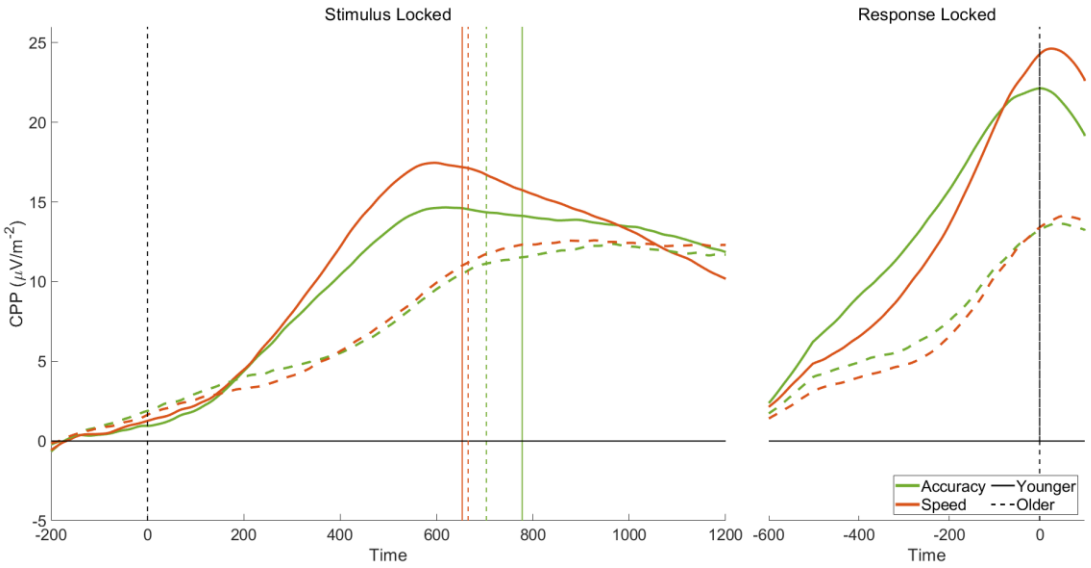

Fig. 2.10: Stimulus and response aligned CPP traces over time, with vertical coloured lines indicating mean reaction time for a given condition. Older adults show significantly lower response-locked amplitudes

Previous examinations of the CPP on similar tasks have indicated a substantial build-up of the CPP, several hundred milliseconds before any contrast change occurs, suggesting pre-evidence decision formation (Devine et al., 2019, see section 1.2.2). Here, the contrast grating is on screen with no difference in contrast for 400ms in advance of the contrast change, yet the CPP appears to begin its accumulation by evidence onset in advance of evidence or any temporally feasible input from sensory areas. This is indicative of some pre-contrast formation of the decision before any evidence becomes available. This feature is not accounted for in standard DDMs. To assess if these pre-evidence slopes were significantly different from 0, indicative of pre-evidence decision formation, a series of independent samples t-tests against 0 were used to investigate the early build-up of the CPP prior to coherent motion onset (-200ms to 0ms). A Bonferroni corrected p-value for 4 comparisons of 0.0125 was used. The slope of this signal in this time window was shown to be significantly greater than 0 for older adults under both accuracy ($t_{(23)}$ = 2.79, p = 0.01) and speed emphasis ($t_{(23)}$ = 3.589, p = 0.001). There were no significant differences for younger adults (all p > 0.23). Analysis of pre-evidence build-up of the CPP was performed to assess if groups differed in the rate of this accumulation process and whether this differed under speed emphasis. This indicated no significant main effect of Age, Condition or interaction on pre-evidence amplitude (Age: p =

0.365; Condition: p = 0.481; Interaction p = 0.92) or on the pre-evidence slope (Age: p = 0.5; Condition: p = 0.139; Interaction: p = 0.856).

Collectively this suggests pre-evidence accumulation of noise in the older adult group. The rate of this noise accumulation did not differ significantly across conditions. Considered in conjunction with the high proportion of fast errors seen in the behavioural data, this encouraged the inclusion of a parameter which would allow for some pre-evidence accumulation of noise. To incorporate this observation into our neurally-informed model, a drift-onset parameter replaced traditional non-decision time. The accumulation process was allowed to commence at evidence onset, accumulating only Gaussian noise until the drift-onset time was reached.

2.3.3.3: Starting Point Adjustments and Urgency

Pre-evidence mu/beta has been marked as a useful proxy for pre-evidence starting point adjustments (Kelly et al., 2021). Initial visual inspection of mu/beta signals indicated substantial group differences in amplitude (see Figure 2.11).

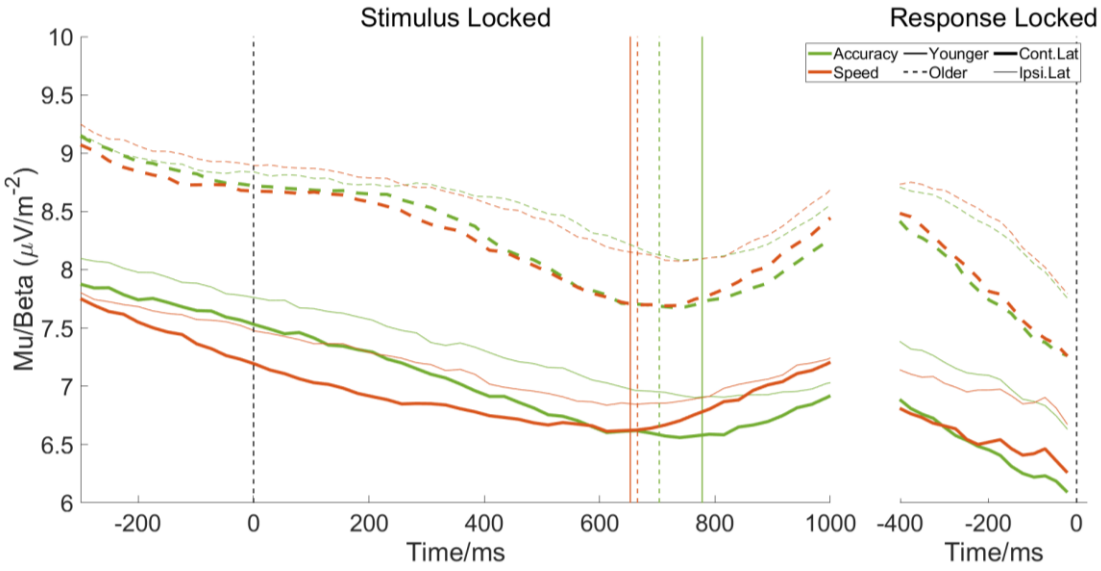

Fig. 2.11: Non-normalised contralateral and ipsilateral mu/beta-traces over time. A substantial difference in average amplitude across the whole trial between older and younger adults is observed

However, these differences were equally evident at the time of the contrast change and at the time of response execution suggesting that they may not have any bearing on the decision-making process. To test this, trials were split by median RT and analysed using a 2 x 2 factorial ANOVA (Condition X RT Bin) for each group, allowing us to test whether mu/beta showed a consistent desynchronisation to a threshold. There was no significant effect of Condition ($F_{(1,49)}$ = 1.125, p = 0.294) or Reaction Time ($F_{(1,49)}$ = 0.742, p = 0.394) on contralateral mu/beta at the response in either group consistent with motor signals reaching an action-triggering threshold in both groups (see Figure 2.12).

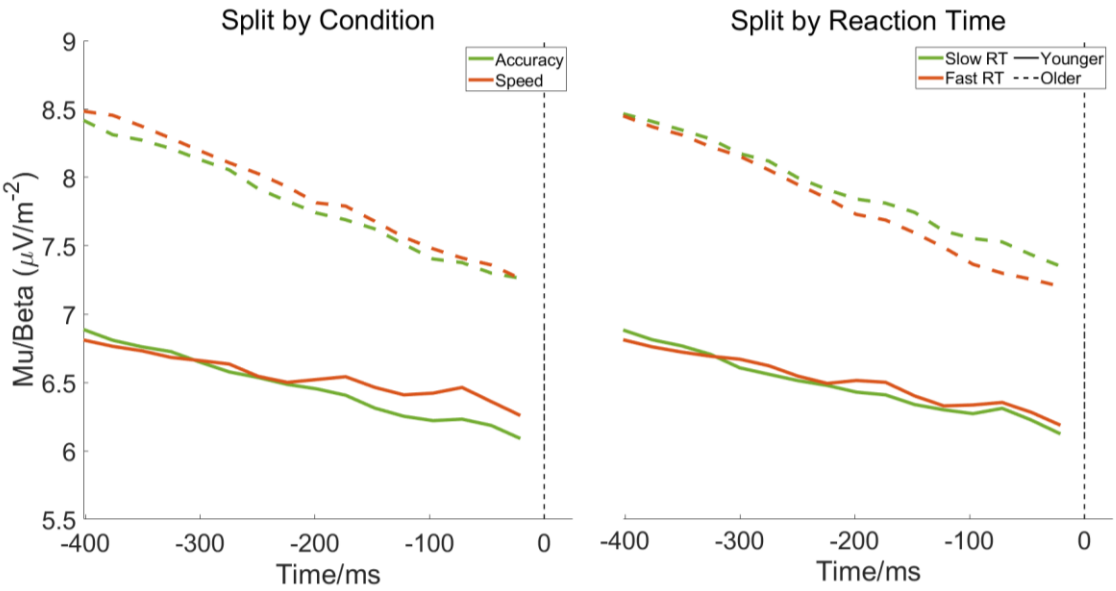

Fig. 2.12: Contralateral mu/beta by coherence level (left) and RT bin (right). Contralateral mu/beta accumulates to a threshold irrespective of RT or trial difficulty.

This suggested that the group differences in mu/beta amplitude were not consequential for the decision process and could reflect general age effects on the background EEG in this frequency range, in line with recent work from Stacey and colleagues (2021) who showed a considerable age-related increase in beta-spectral power in older age. In order to eliminate irrelevant differences in spectral amplitude and to facilitate group comparisons, we normalised the mu/beta signals for each subject relative to its contralateral amplitude prior to response execution (-200ms) within each task condition (see Figure 2.13).

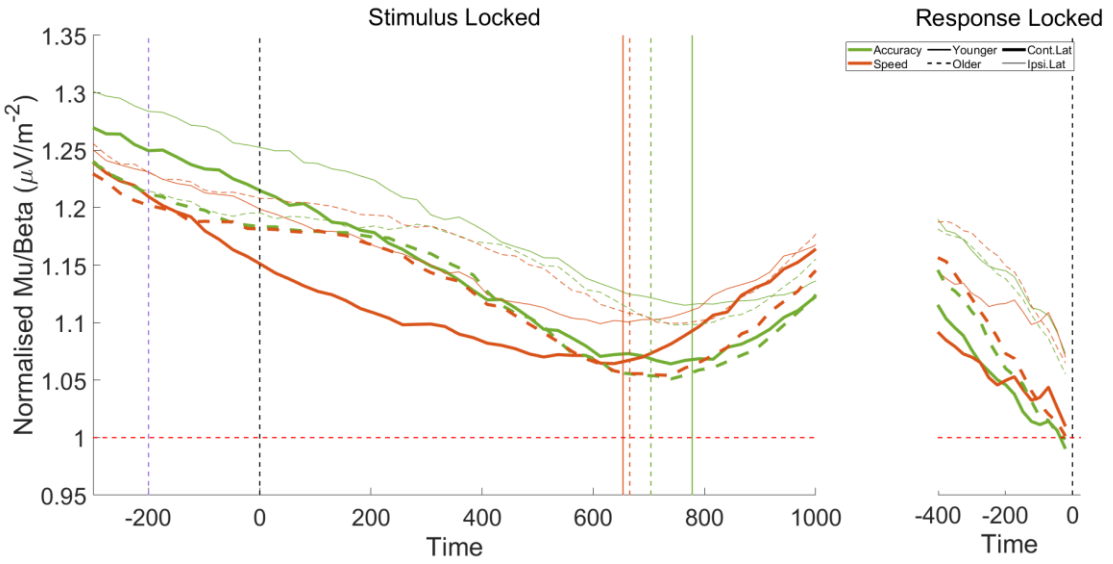

Fig. 2.13: Stimulus and response-locked normalised contralateral and ipsilateral mu/beta. Vertical dashed lines indicate the mean RT for that group and condition. A vertical dashed purple line indicates the time point used for pre-evidence mu/beta calculation.

In order to test for any dynamic pre-evidence adjustments to the decision bound a series of independent samples t-tests were used to investigate whether mu/beta slope averaged across contralateral and ipsilateral cortexes showed any significant build-up prior to coherent motion onset (-300 to -100ms). As the test was performed for each age group and condition, a Bonferroni corrected p-value for 4 comparisons of 0.0125 was used. Significant negative build-up was shown for younger adults in both the Accuracy ($t_{(24)}$ = -1.415., p_1 =0.007), and Speed conditions ($t_{(24)}$ = -2.58., p_1 =0.006). Similar results were seen in the older cohort, with significantly negative pre-evidence slopes in the Accuracy ($t_{(24)}$ = -1.332., p_1 =0.007) and Speed conditions ($t_{(24)}$ = -2.54., p_1 =0.006). A mixed-factorial ANOVA indicated a significant main effect of Condition ($F_{(1,47)}$ = 4.416, p = 0.04), suggesting significantly more negative pre-evidence mu/beta slopes under speed emphasis. There was no significant effect of Age (p = 0.99) or Age-Condition interaction (p = 0.91). These analyses suggest a motor signal that is building towards its threshold before any evidence appears on the screen, and which builds at a steeper rate under speed emphasis, indicating potential urgency dynamics.

Pre-evidence mu/beta is shown in Figure 2.16 and is seen to show adjustments to its amplitude in advance of the contrast difference onset. To assess for age or

condition differences in this pre-evidence motor preparation, a mixed factorial analysis of variance of pre-evidence mu/beta amplitude in the -300:-100ms before stimulus onset was performed on mu/beta signals averaged across contralateral and ipsilateral hemispheres to the ultimately chosen direction. No significant effect of Condition (p = 0.159), Age (p = 0.695) or Condition by Age interaction (p = 0.114) was observed.

The CNV was then analysed as a further marker of urgency. Analysis of amplitude in response aligned traces indicated no significant main effects of Condition, Age or interaction term (Time Window= -50:50ms, Age: $F_{(1,47)} = 0.907$, p = 0.346, Condition: $F_{(1,47)}$ = 1.829, p = 0.183, Interaction: $F_{(1,47)}$ = 2.363, p = 0.131). Analysis of pre-response slope in the window of -400:-100ms indicated no significant main effect of Age ($F_{(1,47)}$ = 0.114, p = 0.737) Condition ($F_{(1,47)}$ = 2.696, p = 0.591) or Age-Condition interaction ($F_{(1,47)}$ = .2.696, p = 0.107, see Figure 2.14).

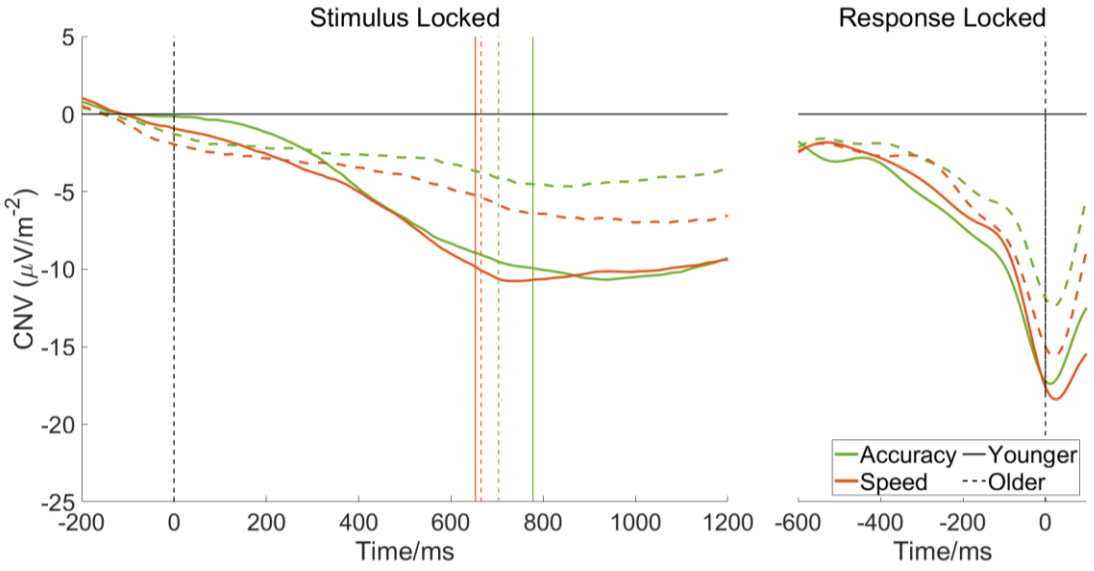

Fig. 2.14: Stimulus and response aligned CNV traces over time with vertical coloured lines indicating mean response time per condition

CNV features were analysed in the pre-target window of -100:0ms. There was no significant main effect of Age on the pre-target CNV slope ($F_{(1,47)}$ = 2.166, p = 0.148). The main effect of Condition was significant ($F_{(1,47)}$ = 5.489, p = 0.0234), with a non-significant interaction term ($F_{(1,47)}$ = 1.8952, p = 0.255), indicating a significant increase in pre-stimulus CNV slope under Speed emphasis (*Mean* = -

0.0184, *SD* = 0.048,) relative to Accuracy (*Mean* = -0.008, *SD* = 0.044, see Figure 2.14).

In sum, mu/beta signals show a clear build-up of motor preparation towards a threshold in advance of evidence-onset which is stronger under speed emphasis, consistent with an evidence-independent, time-dependent bound collapse. Furthermore, the CNV shows a negative going slope which is more pronounced under speed emphasis. Both of these features suggest the presence of urgency in the decision process. This could explain key features of the behavioural data, which contained a high proportion of slow errors. These slow errors can be accounted for by an increasing urgency to respond in advance of the end of a trial. Collectively this encouraged the inclusion of a dynamic urgency component in the model.

2.3.4: Neurally-Informed Modelling

The EEG signatures outlined above were used initially to inform the model construction. Analysing these features gave insight into potential model parameters which should be incorporated. This provided a foundational model framework which would then be constrained by reliable neural estimates of the parameter values.

Examination of the shape of the conditional accuracy function suggested the presence of potential urgency dynamics, with low accuracy in later quantiles. Alternatively, late errors may be attributable to variable drift rates, whereby trials with lower drift rates would tend to yield a higher proportion of slow errors. However, the inclusion of this parameter in the full-diffusion model above did not successfully capture this data feature (see Figure 2.7). In light of these behavioural features, and with the evidence of urgency from both the CNV and pre-evidence mu/beta signals, urgency was incorporated as a freely estimated model parameter. It took the form of a linearly collapsing bound. Furthermore, to account for the early accumulation of the CPP and the high number of fast errors which were not captured by starting point variability in the full DDM, noise was allowed to accumulate from time point 0, until a drift-onset time which was freely estimated by the model from which point the drift rate was also incorporated. To increase model parsimony, efforts were made to reduce the parameter space by removing the variability parameters. These model variants with a reduced parameter space were found to be capable of capturing behaviour while improving model comparison scores relative to the full DDMs.

These adaptations improved model fits substantially across model variants with each of the key parameters varying by condition for both groups (see Table 2.4). For example, the simplest model with no parameter varying by condition and no variability parameters reduced from an AIC of 64.41 in the Standard DDM for the younger group to 41.58 in the neurally-informed model, even with the removal of between-trial variability parameters. Notably, with the introduction of a collapsing bound, the most efficient model in both raw error score and AIC for both groups became one in which the boundary varied by condition (G^2 Younger = 14.707,

AIC = 26.71, G^2 Old = 3.823, AIC = 15.82). Attempts to fit a more complex Weibull-function boundary collapse function did not lead to substantial increases in raw error scores and, in increasing the number of required parameters, significantly reduced model comparison indices. A model which allowed urgency to vary was a close second for both the younger and older groups, lending further support to the importance of including urgency in the model.

Table 2.4: Table of neurally-informed model fits, with each model variant allowing each of the parameters in the leftmost column to vary by Speed-Accuracy condition. The green highlight indicates the winning model with the lowest AIC in both groups. These models added an urgency parameter and allowed for noisy-accumulation.

Neurally-Informed Models		Younger		Older	
Parameter Varying	N Params.	G^2	AIC	G^2	AIC
None	5	31.58	41.58	6.65	16.65
Boundary	6	14.71	26.71	3.8	15.82
Drift Onset	6	29.86	41.86	6.23	18.23
Drift-Boost	6	31.24	43.24	6.48	18.48
Urgency	6	20.64	32.64	4.28	16.28

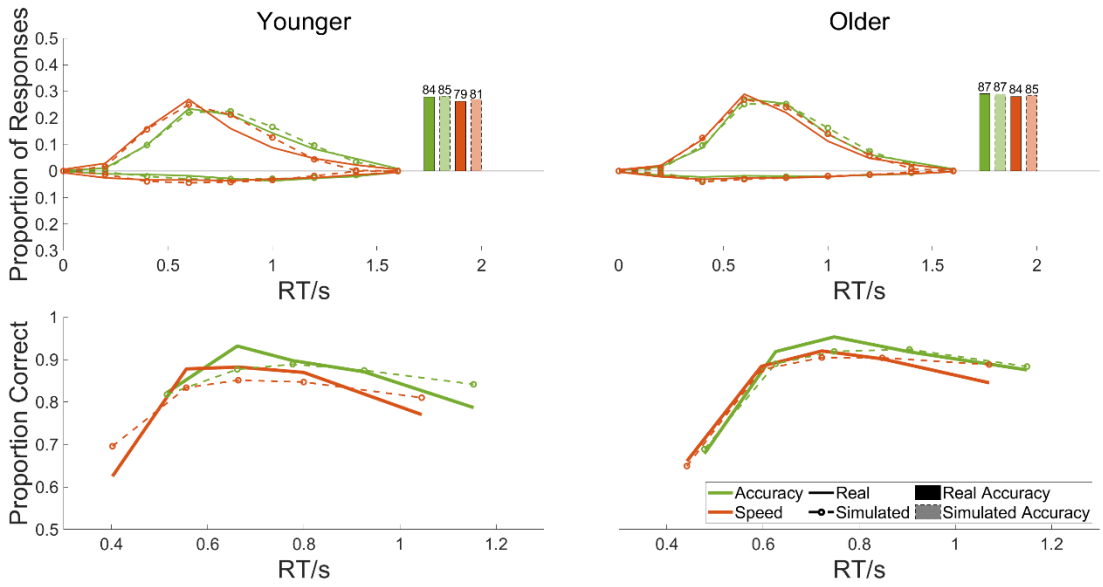

Fig. 2.15: Real and simulated behaviour for the best fitting neurally-informed model, where the boundary was allowed to vary by session. The top row shows real vs simulated reaction time distributions and accuracies. The bottom row shows real and simulated accuracy quantiles for five RT quintile bins.

Table 2.5 Model estimated parameters for the neurally-informed model with boundary varying by condition. For condition-specific parameters, the requisite condition is indicated in subscript. Older adults show lower boundaries overall and less adjustment across conditions, with higher drift-rates. Drift On = Drift-onset time. Urg = Urgency component.

	$Bound_{Acc.}$	$Bound_{Spd}$	$Drift_{Low}$	$Drift_{High}$	Urg	Drift-On	G^2	AIC
Younger	0.186	0.164	0.097	0.156	0.111	0.275	13.11	47.74
Older	0.172	0.163	0.136	0.197	0.1	0.348	3.78	38.26

2.3.5: Neural Constraints

As outlined, neurally-informed modelling allows for the exploration of a wider range of possible group effects without sacrificing model parsimony (Kelly et al., 2021). To constrain the current data by the neural features, following the procedure employed by Kelly and colleagues (2021) pre-evidence mu/beta was used as an index for starting-point adjustments to bound, which are mathematically equivalent to adaptations of the decision bound. While the differences in pre-evidence mu/beta failed to reach statistical significance (section 2.3.3.3), the observed trends followed predictions for increased motor preparation under speed pressure for both groups, with less adaptation of motor preparation in older adults (see Figure 2.16). The group and condition with the largest excursion of normalised mu/beta (Figure 2.13) from the pre-evidence window (-200ms, dashed purple line) to the mu/beta threshold at response was used as the baseline value and assumed to show no starting point adjustments. In this dataset it was the Younger-Accuracy condition, meaning this condition was assumed to have 0 starting point adjustments. The remaining conditions for each group were scaled in reference to this baseline, giving the values in Table 2.6. To encourage parameter values more numerically similar to the non-constrained model, bounds were fixed at 0.1 rather than 1, and starting points were scaled to reflect this (e.g., Younger-Speed Condition: 0.161/10 = 0.0161). Starting point adjustments were then subtracted from a bound of 1 to give a bound setting for each condition for each age group.

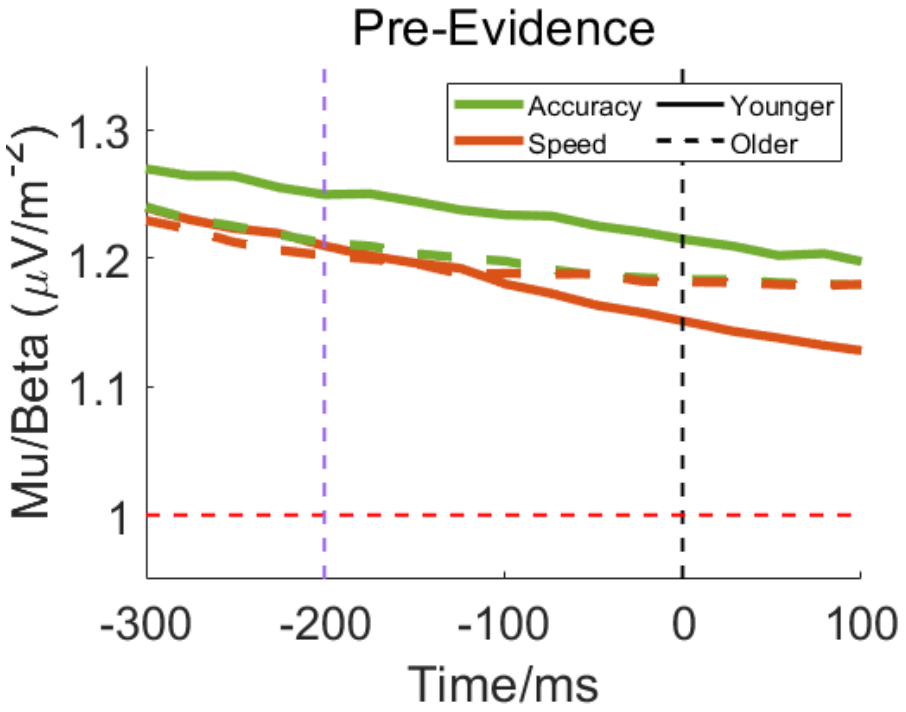

Fig. 2.16. Pre-evidence mu/beta levels averaged across electrode sites contralateral and ipsilateral to the chosen response. The dashed-purple line indicates the time at which pre-evidence mu/beta levels were calculated.

Table 2.6 Mu/beta derived starting points for each group and condition. These values are subtracted from a set bound of 0.1 for each condition, meaning an equal probability of accumulation to a correct or error response.

	Accuracy Starting Point	Speed Starting Point	% Adjustment: Accuracy to Speed
Younger	0	0.0161	-16%
Older	0.0151	0.0190	-5%

These constraints acted as a scaling parameter allowing for within-trial noise to be estimated for each group without increasing the number of free parameters. Again, multiple variants allowing each of the key parameters to vary by condition were simulated (See Table 2.3). Through this, it was shown that a model allowing urgency to vary by condition benefitted the younger group most significantly (G^2 = 7.21, AIC = 19.21), while the older adults showed good fits using only these neural constrained static bound adjustments (G^2 = 5.17, AIC = 15.17, Table 2.7). These fits were improvements on the neurally-informed models (Table 2.4),

highlighting the benefit of neural constraints for allowing more complex explanatory models.

Table 2.7. Error scores for each neurally-constrained model variant. All models have starting points set by the beta-derived values in Table 2.6. The additional parameters indicated were free to vary by speed-accuracy condition. The green highlight indicates the winning model for each group. A model with boundary constrained and no other parameter varying provided the best fits for older adults, while one with urgency varying fit best for the younger cohort.

Neurally-Constrained Models		Younger		Older	
Free Parameters	N parameters	G^2	AIC	G^2	AIC
None	5	12.26	22.26	5.17	15.17
Urgency	6	7.21	19.21	3.46	15.46
Drift-onset	6	10.16	22.16	3.67	15.67
Drift Boost	6	9.75	21.75	3.52	15.52

Although group-level fitting does not allow for statistical comparison of the model-estimated parameters, Table 2.8 outlines the model-estimated parameters for the neurally-constrained model with urgency varying by condition. Here stronger drift rates are observed in the older adult group, with roughly equivalent levels of within-trial noise across groups. Levels of boundary collapse across conditions seem largely consistent in the older adult group, but younger adults show a more pronounced boundary collapse component under speed pressure. The winning model for the older group was one in which there was no urgency adaptation across conditions (Table 2.7), providing further evidence that older adults may not adapt their urgency dynamics significantly under speed emphasis. A later drift-onset time is also observed in the older adult group.

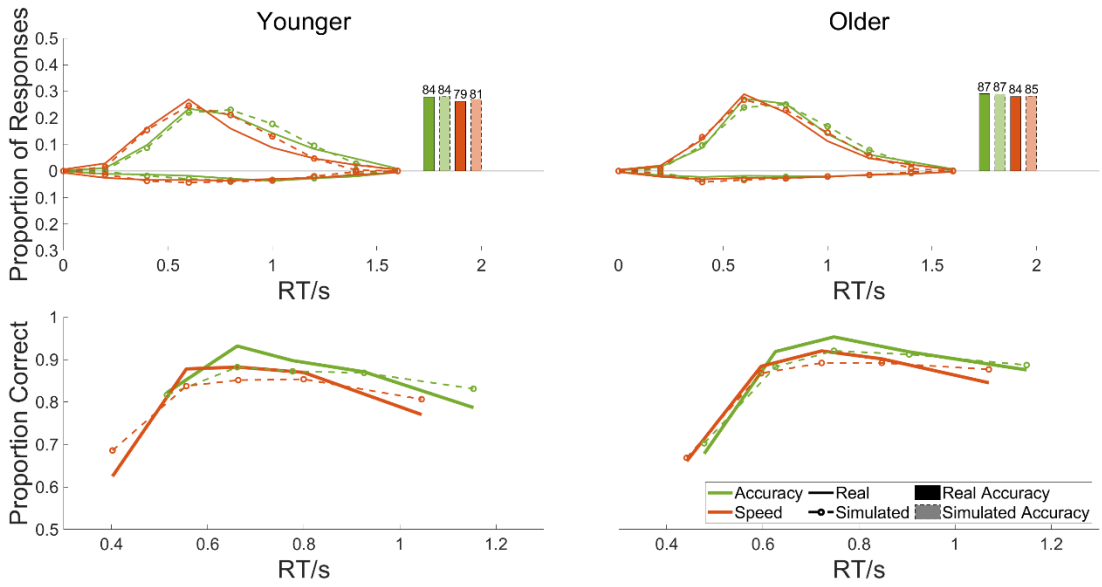

Fig. 2.17. Real and simulated behaviour for the best fitting neurally-constrained model, where the boundary was fixed, and urgency allowed to vary by session. The top row shows real vs simulated reaction time distributions and accuracies. The bottom row shows real and simulated accuracy quantiles for five RT quintile bins.

Table 2.8. The model estimated parameters for the neurally-constrained model with urgency varying by condition. For condition-specific parameters, the requisite condition is indicated in subscript Drift-On = Drift-Onset time, Urg. = Urgency component.

	$Bound_{Acc.}$	$Bound_{Spd}$	$Drift_{Low}$	$Drift_{High}$	Urg_{Acc}	Urg_{Spd}	Drift. On	Noise	G^2	AIC
Younger	0.1	0.084	0.051	0.081	0.061	0.055	0.276	0.052	13.11	47.74
Older	0.085	0.081	0.066	0.096	0.049	0.050	0.342	0.049	3.78	38.26

2.3.6: Model Validation Through Neural Signals

In order to establish whether our winning neurally-constrained model was in concordance with the observed neural data, key features of the model were simulated and compared to our neural indices following the procedure employed by Kelly and colleagues (2021). These comparisons were the evidence accumulation variable vs the CPP, the model simulated motor response variables vs contra/ipsilateral mu/beta, and the linearly collapsing urgency component vs the CNV.

2.3.6.1 Evidence Accumulation

Model estimated decision variables from the winning neurally-constrained model (Mu/Beta Constrained: Urgency Varying) were simulated and plotted to track model concordance with our neural marker of evidence accumulation, the CPP.

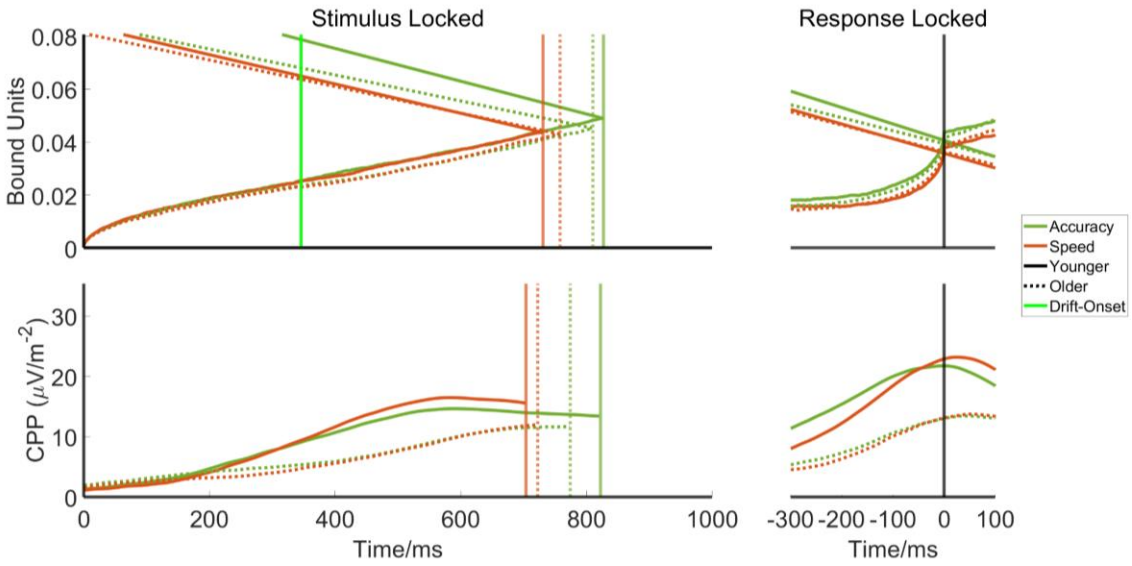

Fig. 2.18: Real vs simulated decision variables. The top row shows stimulus and response-locked simulated cumulative evidence, with diagonal lines showing collapsing bounds for the neurally-constrained model with urgency varying. The bright-green vertical line indicates the time at which the drift rate began. The bottom row shows real CPPs for each group and condition. Vertical lines indicate the mean RT for each group and condition. Real CPP shows a lower amplitude for older adults which is not apparent in the simulated decision variable.

Figure 2.18 shows the model simulated stimulus and response-locked cumulative decision variable for the model. Descending coloured dashed lines illustrate the models' collapsing bounds, with the dashed green line indicating

drift-onset time. The bright-green solid line represents drift-onset time. Simulated decision variables show a consistent amplitude across conditions in both groups at response. However, the real CPP is shown to be significantly lower in older adults. This is not in line with model predictions

2.3.6.2: Motor Preparation

Motor signals were simulated as the negatively signed, cumulative evidence for either the chosen or unchosen response with the negative going urgency component added. These simulations show a strong concordance with the neural data. Simulated motor signals are constrained to match mu/beta-starting levels meaning they inherently align early at the start of the trial, however, the simulated signals exhibit the same pattern of desynchronisation across time and are highly similar by response execution. The large contra-vs-ipsilateral difference early in the trial in the younger group in the accuracy condition is not captured by the simulated motor preparation signal.

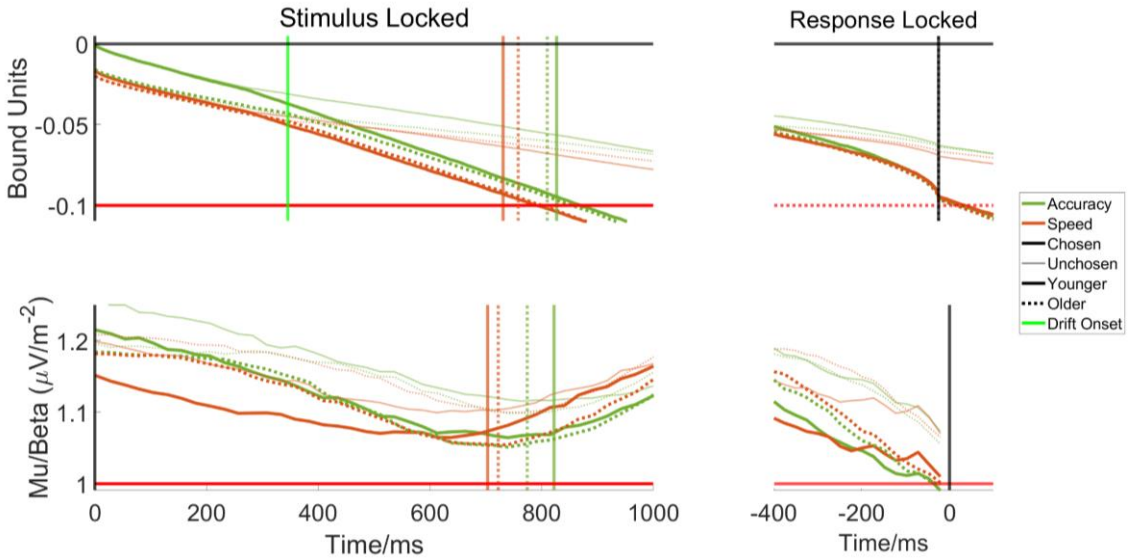

Fig. 2.19: Real vs simulated motor signals. The top row shows stimulus and response-locked simulated motor signals, with diagonal lines showing collapsing bounds for the neurally-constrained model with urgency varying. The bright-green vertical line indicates the time at which the drift rate began. The bottom row shows real mu/beta amplitudes for each group and condition. Vertical lines indicate the mean RT for each group and condition. Real and simulated signals show a strong concordance.

2.3.6.3: Urgency

Urgency signals were simulated as the time-dependent but evidence-independent urgency component for each condition, with pre-evidence adjustments to the decision-bound subtracted at the beginning of the trial.

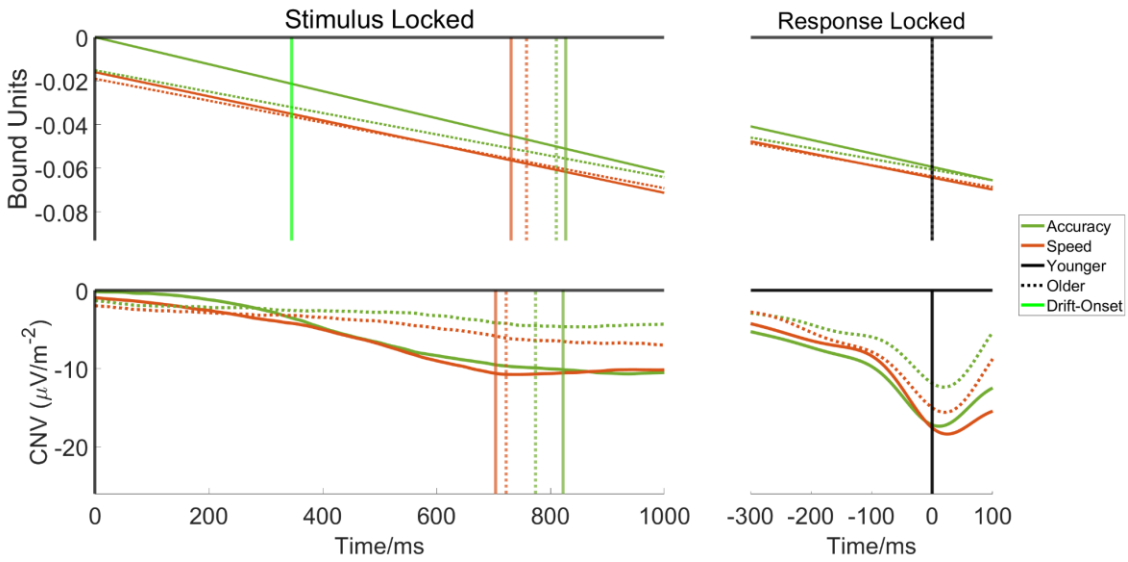

Fig. 2.20: Real vs simulated urgency signals for the neurally-constrained mode with urgency varying. The top row shows stimulus and response-locked simulated urgency signals for each group and condition. The bright-green vertical line indicates the time at which the drift rate began. The bottom row shows real CNV amplitudes for each group and condition. Vertical lines indicate the mean RT for each group and condition. CNV amplitude does not capture speed-accuracy differences at response in the younger group, but strong concordance before drift-onset time is seen.

Similar patterns across simulated urgency components and the CNV are observed. Simulated urgency components capture differences across conditions and groups in the early stages of the trial (see Figure 2.20). Older adults show a less negative CNV by response than younger adults in the response-locked signal. Simulated urgency components do not show this. Similarly, after an initial differentiation between speed and accuracy, with a less negative CNV in the accuracy condition, younger adults show a similarly negative CNV amplitude across both conditions at response. In the simulated urgency component, a more negative urgency component is observed in the speed condition and this difference is maintained through to response.

2.3.6: Model Fits to Individual Subject Data

Group level fits provided strong evidence for differences in the strategies used by older and younger adults to adapt to changing speed versus accuracy demands, with older adults showing lower boundaries and urgency, and less adaptation of these parameters across regimes. To enable statistical comparison of parameter values across groups, non-neurally-constrained versions of the winning Boundary and Urgency varying by condition, and only Boundary varying by condition models were fit at the individual level. Due to the high degree of inter-subject variability in mu/beta-presentation, neural constraints were not appropriate, and models were fit following the same procedure outlined in section 2.3.4 above. Model fits showed an average AIC of 64 (*SD* = 15.38) for the model with Boundary alone varying, and a mean AIC of 63 (*SD* = 14.77) for the model with Boundary and Urgency varying. There was no significant main effect of Age group ($F_{(1,47)}$ = 3 p = 0.09) or Model Type ($F_{(1,47)}$ = 1.73, p = 0.195) on AIC scores across the two models. A significant Age-Model interaction ($F_{(1,47)}$ = 7.85, p = 0.007) was observed, with post-hoc analyses indicating that younger adults showed a significant lowering of AIC with the introduction of a varying urgency parameter in younger adults (t = -2.94, p = 0.03), while the difference did not reach statistical significance in the older adult group (t = -1.04, p = 0.491). The difference in magnitude of these AIC scores is illustrated in Figure 2.21. Notably, it is evident that a subset of approximately nine individuals in the younger adult group benefited extremely from the inclusion of an urgency component that varies by condition. This suggests that urgency adaptation may be an elective strategy employed by certain younger individuals.

Fig. 2.21. Violin plots of AIC differences between the simple neurally-constrained model and one allowing urgency to vary by condition, showing substantial improvements in allowing urgency to vary by condition for a subset of younger adults.

2.3.6.2: Individual Subject Modelling: Parameter Differences

Group differences in parameter estimates on the Bound and Urgency Varying model were then compared (see Figure 2.22). A repeated measures ANOVA indicated no significant main effect of Age on bound ($F_{(1,47)}$ = 0.537, p = 0.467), but a significant main effect of Condition ($F_{(1,47)}$ = 50.05, p < 0.001), suggesting lower boundaries under speed-emphasis (see Figure 2.22). There was also a significant Age-Condition interaction term ($F_{(1,47)}$ = 14.81, p < 0.001). Post-hoc analyses showed a significant reduction of bound under speed emphasis only in the younger group (t = -7.8, p = < 0.001, Mean Reduction = 0.045, SD = 0.036) but not in the older group (t = -2.258, p = 0.163, Mean Reduction = 0.013, SD = 0.028, see Figure 2.21).

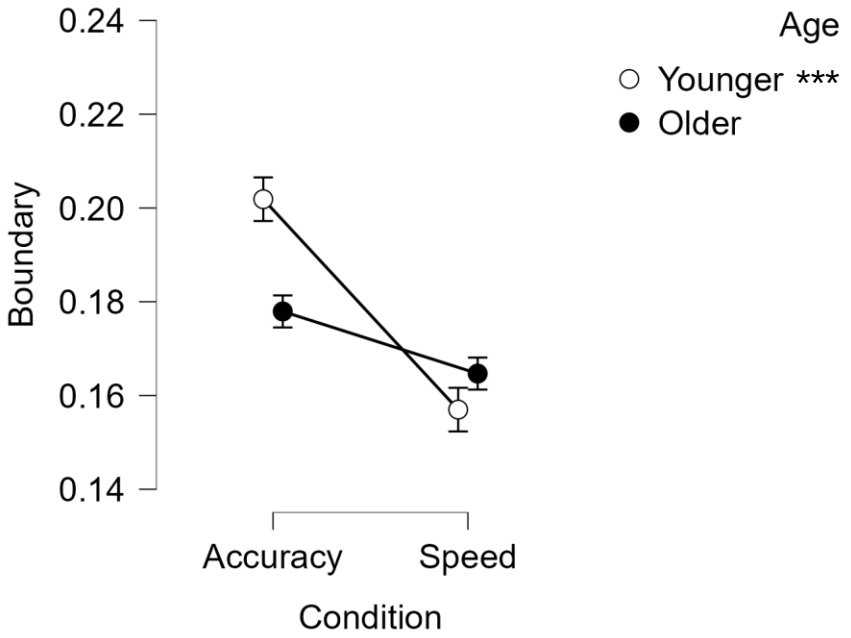

Figure 2.22 Boundary adaptation across speed-accuracy conditions by age group, showing significant adaptation of bound in the younger group but not in older adults.

A similar analysis was performed to investigate group differences in urgency adaptation. Here no significant main effect of Age was observed ($F_{(1,47)} = 1.625$, $p = 0.209$), with a significant effect of Condition ($F_{(1,47)} = 14.735$, $p < 0.001$) and Age-Condition interaction ($F_{(1,47)} = 10.265$, $p = 0.002$). Post-hoc analyses indicated that the younger adults showed a significant reduction in urgency under speed pressure ($t = -5.03$, $p < 0.001$), while this difference was non-significant for older adults ($t = -0.444$, $p = 0.1$, see Figure 2.23).

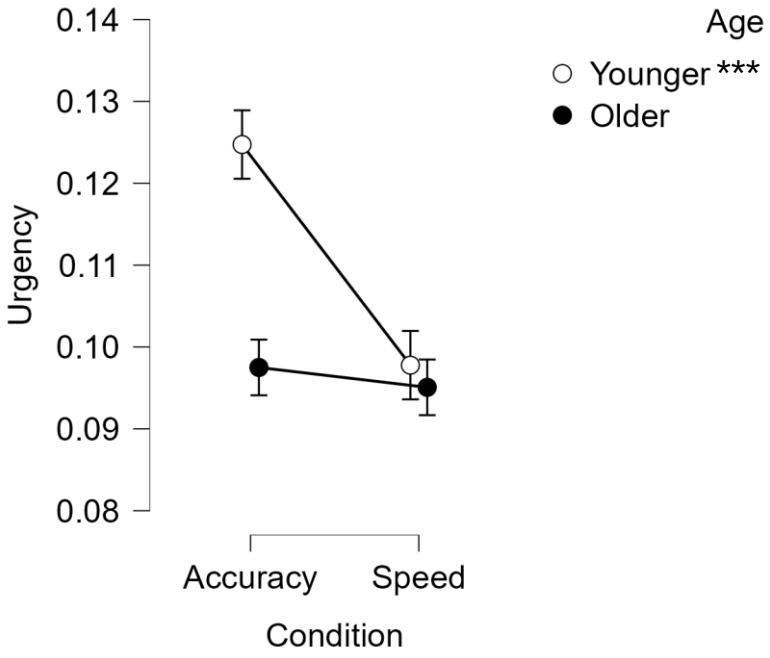

Fig. 2.23: Urgency adaptation across speed-accuracy conditions by age group, showing significant adaptation of urgency in the younger group, but not in older adults.

Analysis of drift-rate differences in low and high contrast trials indicated no significant main effect of Age ($F_{(1,47)} = 1.131$, $p = 0.29$) or Age-Contrast interaction ($F_{(1,47)} = 0.066$, $p = 0.798$), but an expected significant effect of Contrast level ($F_{(1,47)} = 102.272$, $p < 0.001$), with higher drift-rates for high-contrast trials. This suggests that the groups do not differ significantly in drift rate across high and low-contrast trials.

An independent samples t-test was used to compare drift-onset times across groups. A significant difference in drift-onset time was shown ($t_{(1,47)} = -2.042$ $p = 0.047$), indicating a later drift-onset time in the older group (Mean = 0.332, $SD = 0.086$) relative to the younger group (Mean = 0.283, $SD = 0.8$, Figure 2.24).

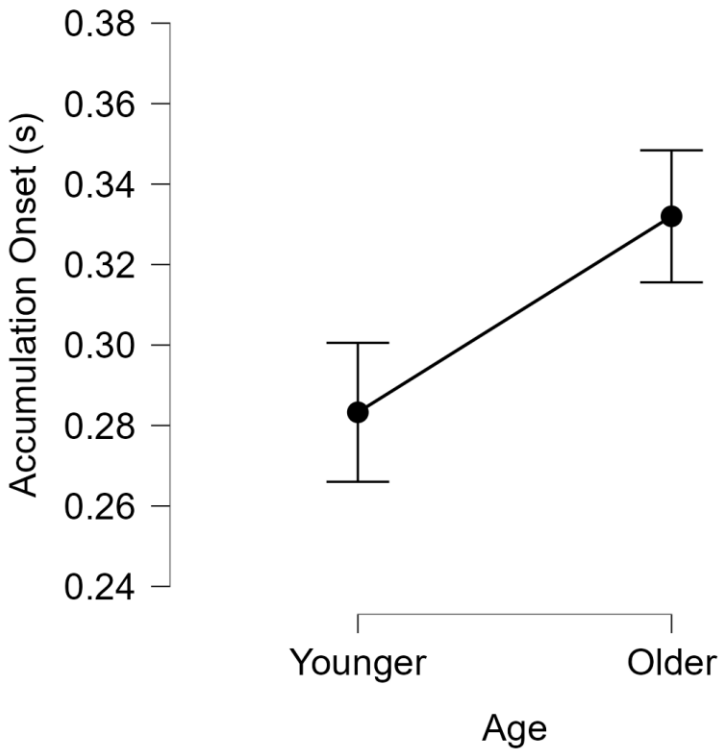

Figure 2.23, Model simulated drift-onset times in seconds by age group showing slower drift-onset times in the older group than younger group.

2.3.7: CPP Scaling by Model Parameters

As outlined in section 2.1.4, older adults may be subject to a range of task-irrelevant differences in the presentation of electrophysiological signals. In this study, the CPP, whose amplitude at response is thought to reflect the cumulative evidence for the chosen alternative, was shown to reach a much lower boundary than the model simulated decision variable would predict (see Figure 2.18). While this may suggest the model is not successfully representing the decision process, it is also possible that CPP amplitude measurements are affected by behaviourally-irrelevant group differences which result in a less positive CPP for older adults. Features of simulated models may be potentially useful in addressing this issue. By constraining features of the CPP to match what is observed in the computational model, we may obtain a more representative illustration of the CPP's emergence over time. To explore this, the CPP was rescaled relative to the model-derived bounds at response (marked with horizontal lines in Figure 2.21). The mean CPP at response was taken in the window -100:-60ms from the response, to allow for non-decision-related motor

response times (time window highlighted in purple in Figure 2.25). The entire signal ERP was then scaled by multiplying each time point by the model-derived bound at response divided by the mean CPP at response giving the signals shown in Figure 2.21. Through this both groups show signals on similar scales, illustrating more clearly an earlier build-up of the decision formation signal seen in older adults that is reflective of their later drift-onset times.

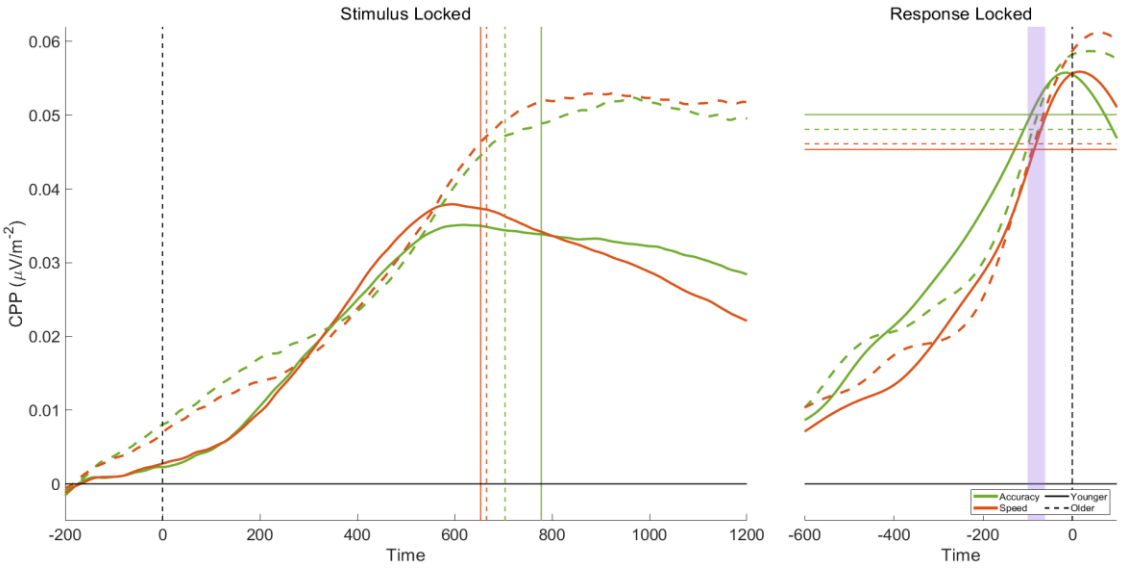

Fig. 2.25) Stimulus and response-locked CPPs rescaled by neurally-constrained model bounds at response, The purple bar in the response-locked signal highlights the -100:-60ms window where mean CPP at response is calculated, while horizontal bars show the value of the model-derived bound at response.

To test for the statistical significance of these observed differences, a series of ANOVAs on CPP features were performed. While the un-rescaled CPP showed a significantly shallower slope in older adults, model estimated drift-rates showed no group differences. To assess if the rescaled CPP could better represent this, pre-response slopes were analysed. There was no significant effect of Age (p = 0.3), Condition (p = 0.924) or Age-Condition interaction (p = 0.554), reflecting the similar drift rates across groups seen in modelling efforts. To assess if the rescaled signal better captured the later drift-onset time, and thus the higher proportion of accumulated noise in the older adult group, pre-evidence CPP slopes were analysed. There was no significant effect of Age (p = 0.08), Condition (p = 0.423) or Age-Condition interaction (p = 0.774). There was also no significant effect of Age (p = 0.141), Condition (p = 0.613) or Age-Condition interaction (p = 0.91) on pre-response amplitudes. While non-significant, visual

analysis of the rescaled CPP suggests more pre-evidence accumulation in the older adults, which may better reflect the later drift-onsets seen in the older cohort, and similar build-ups reflecting the similar model-estimated drift rates across groups.

2.4: Discussion

The speed-accuracy tradeoff has been extensively studied in both younger and older adults. Using a DDM, an ideal observer is assumed to make this tradeoff by lowering their decision bound or raising their starting point in response to speed pressure (Heitz, 2014). Older adults have been shown to be less capable of making these adjustments at the behavioural level, reflected by a higher decision bound and reduced lowering of decision boundaries (Forstmann et al., 2011; Starns & Ratcliff, 2012). However, the data presented here suggest that these alternate decision strategies do not come at any functional cost to older adults, given their parity in point scores. Furthermore, reliance on the DDM often overlooks core components of the decision process which are observable in the neural data such as urgency and early evidence accumulation. This study aimed to leverage neurally derived proxies of key parameters of the diffusion model to inform and constrain decision models. This led to a range of considerable and important new insights into the effects of ageing on the speed-accuracy tradeoff. Notably, older adults were shown to exhibit lower boundaries, but less adaptation of urgency dynamics. This suggests that rather than simply increasing their bounds overall, older adults are adopting a more cautious and *consistent* decision style with less collapse of these bounds over time in order to maximise their points scored.

Initially, the standard full DDM indicated poorer regulation of non-decision time across conditions in older adults as the best-fitting model. This was not in line with accounts of the speed-accuracy tradeoff in an ageing population, which suggest adjustments to boundaries as driving the behavioural differences across conditions (Forstmann et al., 2011; Starns & Ratcliff, 2010). However, a model with boundary varying by condition performed similarly well and was consistent with previous reports which highlight higher decision bounds and poorer regulation of these bounds in older adults. This presented two similarly well-performing models, which model comparison scores alone could not conclusively adjudicate between. Analysis of the neural data indicated the presence of a range of factors which encouraged the revision of the DDM. Behaviourally, both groups were shown to exhibit a negative parabolic function of accuracy across reaction

time. This indicated the presence of a large proportion of fast and slow errors, which standard DDMs have difficulty capturing despite the presence of drift rate and starting point variability parameters. Fast errors may be driven by the accumulation of noise prior to evidence onset, with slow errors driven by an increased urgency to respond as the time limit for response approaches. Importantly these features were reflected in the neural data. The CPP evidence accumulation signal was seen to have accumulated before any temporally feasible feedforward of information from sensory areas in older adults, indicating its accumulation of noise in advance of the target contrast. This accumulation of noise could therefore help to explain the large proportion of fast errors in early reaction time bins. Urgency dynamics were also observed in the neural data, where mu/beta signals were seen to begin their negative descent towards the bound in advance of the evidence. Furthermore, the CNV, a marker of dynamic urgency, was shown to exhibit a more pronounced negative slope under speed pressure before the target appeared.

Using this information to inform our models led to the inclusion of a linearly collapsing bound to reflect urgency while allowing for noisy accumulation until some time point. The inclusion of these parameters greatly improved model fits relative to the standard DDMs, even with inter-trial variability parameters removed. Importantly, the best-fitting model was one in which the boundary was allowed to vary by condition, in line with previous literature (Forstmann et al., 2011; Starns & Ratcliff, 2010). Here, the neurally-informed model provided a clearer adjudication as to the winning model than the standard decision model, with a model with boundary varying by condition outperforming other variants. In a departure from previous work suggesting increased boundaries in older adults (Forstmann et al., 2011; Starns & Ratcliff, 2010), older adults exhibited a lower boundary in the accuracy condition (0.172) than younger adults (0.186), with a less pronounced reduction in these boundaries under speed-emphasis (Older = -0.009, Younger = -0.022). While the reduced boundary adjustment in older adults replicates previous studies, the lower boundary in the accuracy condition is in direct contrast to previous work which has indicated higher decision bounds across conditions in older adults. Though the group-level modelling precluded statistical analysis of these differences, they were reflected in the pre-evidence

levels of mu/beta, where older adults showed more pre-evidence motor preparation, equivalent to lower bounds, than the younger in the accuracy condition and a less pronounced adaptation of motor preparation under speed emphasis. Here, the neurally-informed model is not only capable of better capturing the behavioural data but also produces results which are more in line with observable neural metrics.

Using these start-points to constrain the model allowed for the investigation of increasingly complex models without penalty to the model comparison scores. Neurally-constrained models provided better fits than the neurally-informed models while providing a more nuanced account of the adaptation. Here it was indicated that individuals regulate both their urgency dynamics across conditions, with a more pronounced urgency adaptation in younger adults. This is a highly important extension beyond traditional models, which indicate older adults adopting more cautious decision styles driven by boundary adjustments alone. Instead, the model indicates a smaller boundary adjustment and no urgency adaptation in older adults, which enables more consistent responses over time. These conclusions may have been unattainable using standard modelling efforts, as a more complex model would have been penalised heavily for an increased number of parameters. Here, neural constraints allow us to circumvent this penalisation to show a second adaptation difference between older and younger adults, in the degree to which they change their levels of urgency across regimes. Using standard diffusion models without reference to the neural information would have led to completely different conclusions. To date, the majority of the literature has done just that, implicating higher boundaries in older adults (Forstmann et al., 2011; Starns & Ratcliff, 2010, 2012). This work presents the key benefit of utilising neural information, allowing us to deliberate between similarly well-performing models and create models that more accurately describe the effects of ageing on the speed-accuracy tradeoff.

Importantly, these constraints allowed us to investigate group differences in within-trial noise, which standard models do not readily allow for. There were found to be no substantial group differences in levels of within-trial noise. This, combined with evidence for equal drift rates and no group differences in the

amplitudes or build-up rate of the sensory-encoding SSVEP signals suggests an absence of a deficit in the core information-processing stages of the decision process. This represents an important advancement as it offers combined neural and behavioural evidence for unimpaired evidence encoding in older adults, both in terms of mean evidence representation (drift rate) and the incumbent noise incorporated at each time point (within-trial noise). A task which estimates a behavioural measure of internal noise (Dosher & Lu, 1998; Tibber et al., 2014), would be an additional useful method for assessing group differences in within-trial noise levels. The absence of a group difference in within-trial noise levels may have been due to the nature of the stimulus, which emerged gradually over time. In contrast, a stimulus with more ongoing stimulus variability and noise, such as a random dot motion task, may elicit different results. Furthermore, as our models were reduced, removing variability parameters for parsimony, they could not investigate potential differences in inter-trial variability parameters. Older adults have been shown to exhibit lower ongoing neural variability, correlated with impairments in cognitive flexibility (Kumral et al., 2020; Sleimen-Malkoun et al., 2015), which may be reflected in maladaptive lower levels of intertrial variability.

Finally, the group-level parameter differences were supported statistically by fitting the data at the individual level, suggesting the group-level constraints were capable of capturing true between-group differences. Older adults were shown to exhibit substantially slower drift-onset times, which may be reflective of a combined slowing of evidence accumulation-onset time and increased motor times. Our capacity to tease apart the relative proportion of evidence onset time and motor response delays was one limitation of this study. However, analysis of the CPP indicated significant pre-evidence build-up of the decision formation signal in the older adult group but not the younger. This lends support to the idea that at least some of this increased accumulation of noise may have been driven by an increased pre-evidence accumulation of noise in the older adult group. Future work may aim to separate these components as individual parameters constrained by neural signals. For example, Kelly and colleagues (2021) used the lateralised readiness potential as an index of the time at which commitment to a motor action was made. This neurally-informed model constrained the motor

execution portion of non-decision time, allowing the model to freely estimate an accumulation-onset time. This would allow the model-estimated drift-onset time to represent sensory encoding more directly.

Individual model fitting also indicated that older adults were less likely to adapt their urgency dynamics in response to speed pressure than younger adults. However, a close examination of the individual model fits indicated that a large minority of younger individuals may have been driving the observed model improvements by allowing urgency to vary. This suggests that adapting urgency dynamics is a strategy employed by some, but not all, individuals. Future work may benefit from finding reliable methods for applying neural constraints to individual model data. These individualised neurally-informed models may be better capable of representing individual variability in strategy and parameter differences. With these individual differences accounted for, any true between-group effects may become more apparent. Collectively, however, the individual-level modelling appears to reinforce the idea of a strategic adjustment of decision bounds, with less pronounced collapsing of boundaries over time in older adults. A combined less pronounced adaptation of the decision bound and the collapse of this bound over time may reflect a more cautious decision strategy utilised to offset potential differences in motor execution or evidence accumulation-onset times. However, this evidence supports differences in strategic adaptation rather than any processing deficits and is further bolstered by the equivalence in points scored between older and younger adults.

Our attempts at feature scaling reveal key considerations for future work comparing older and younger adults. Notably, the normalisation of mu/beta by its threshold at response allowed us to bring our groups onto equal scales while maintaining our capacity to investigate inter-condition effects. This is pertinent given the considerable evidence of increased overall beta amplitude in older adults (Stacey et al., 2021). Extant research often fails to account for these observations, meaning erroneous conclusions regarding group differences in beta may be drawn. When comparing the simulated decision variable to the real CPP, we saw that the CPP reached a lower amplitude at response for older adults. However, simulated decision variables reached similar amplitudes across

groups. This presented a disconnect between model predictions and observed neural signals. Given the potential for behaviourally-irrelevant ageing effects on volume conduction (Antonakakis et al., 2020), attempts were made to rescale the CPP signal based on model-derived parameters. Here, model-based estimates of bound at response were used to rescale stimulus and response-locked CPPs. This rescaling brought the CPP in line with model expectations, showing more early build-up of the CPP in the older group, in line with their slower drift-onset times. However, rescaling appeared to result in higher amplitudes of the CPP both stimulus and response-locked in older adults than the model would predict. Future work may benefit from combining behavioural and electrophysiological data in the modelling procedure, to find a model-derived parameter at the individual level which could be used to rescale signals like the CPP and overcome the influence of individual, and therefore, group differences in signal scaling.

Our ability to utilise signals such as the CPP and CNV as neural constraints may be limited by a series of factors. Firstly, the location of electrophysiological signals has been shown to vary significantly with age. Notably, the P3, or CPP has been shown to exhibit a pronounced frontalisation as we age (Fjell & Walhovd, 2004; O'Connell, Balsters, et al., 2012). While the direct cause of this is not well understood, it has been theorised to reflect the recruitment of largely intact frontal regions to compensate for cortical thinning (Davis et al., 2008; van Dinteren et al., 2018). Here, individual electrode selection based on peak amplitude showed a more frontalised CPP, however, this may have resulted in the measurement of potentially overlapping signals (Kelly & O'Connell., 2013; West et al., 2010). Furthermore, the P3 component has also shown reduced amplitude with ageing (Tsolaki et al., 2015; van Dinteren et al., 2014b), although the extent to which this reflects a poorer evidence integration system or task-irrelevant changes to brain structure is less well understood, as some studies have shown these reduced amplitudes without group differences in behaviour (O'Connell, Balsters, et al., 2012). A spatially independent method for isolating the CPP's pure decision formation component, which reduces the impact of localisation and overlapping signals, may provide a more appropriate neural

metric of the model-simulated decision variable. This will be explored in Chapter Four.

While our modelling efforts provided strong fits to the data at both the individual and group level, future research may benefit from the use of hierarchical models capable of estimating both group and individual level parameters in parallel and in relation to each other (Turner et al., 2013). However, given the novel nature of the mu/beta constraint technique, and the inherent individual variability in mu/beta signals, such a model was not achievable for this work. One potential avenue for further exploration would be to normalise mu/beta at the individual level and constrain by the individual degree of adaptation to motor levels across conditions. This could provide a potentially valuable method for constraining boundary adaptation while eliminating the potential influence of individual variability in signal power.

In conclusion, the present study provides not only a new insight into speed-accuracy tradeoff dynamics in older adults but an application of a new approach to behavioural modelling, which ultimately grounds mathematical models in observable neural data to provide a synthesised and more nuanced understanding of cognition.

Chapter Three:

Modelling the Role of Internal Noise Suppression in Perceptual Learning

3.1: Introduction

Perceptual learning is a core cognitive capacity, representing the ability to learn and adapt to the demands of a novel perceptual task. Across an array of perceptual domains, including audition, olfaction, and vision, perceptual learning has proven to be a popular phenomenon for researchers hoping to understand the brain's capacity to adapt and its underlying neural plasticity (see Gilbert et al., 2001; Maniglia & Seitz, 2018 for review). The mechanisms underpinning perceptual learning are not well understood and it is not clear what sensory and/or cognitive adaptations may underlie it. One dominant psychophysical theory of perceptual learning has indicated that a key underlying process may be the progressive suppression of sensory noise (Dosher & Lu, 1998, 1999; 2017). In essence, the perceptual processing of a stimulus can be simplified to the extraction of a signal relative to some degree of noise (Dosher & Lu, 2017; Levi et al., 2005). Importantly, these sources of noise exist not only in the external sensory environment but also internally in the brain at each stage of information processing, which can operate over different timescales from milliseconds to minutes and hours (Faisal et al., 2008; Ratcliff et al., 2018). In sequential sampling models, such as the drift-diffusion model (DDM) however, learning is often attributed not to reductions in noise or variability parameters, but to specific mechanisms in the decision process like accelerated drift rates or decision-bound adjustments (Ivanov et al., 2022; Petrov et al., 2011; Zhang & Rowe, 2014). This may, in part, be driven by the fact that in order to fit behavioural data these models must fix a scaling parameter, and by convention that parameter is typically within-trial noise (Ratcliff & McKoon, 2008). Consequently, in studies of perceptual learning, within-trial noise is set to a fixed value across all sessions, preventing the models from identifying any changes to the very parameter that has been heavily indicated in the psychophysics literature. The aim of this study, therefore, was to leverage a combination of neurophysiological evidence and neurally-informed computational modelling to allow for adaptations of within-trial noise with learning, further elucidating the neural substrates of perceptual learning on a visual perception task.

3.1.2: Psychophysical Models of Perceptual Learning

Psychophysical models of the learning process provide a framework from which the changes in the perceptual system which result in learning can be studied. In their perceptual template model, Dosher and Lu (2012; 2017) suggest that perceptual learning occurs through the suppression of internal sources of additive noise and the improved filtering of external noise in the physical stimulus. To study these, they apply equivalent noise paradigms which estimate the internal noise levels of individual subjects by finely manipulating levels of external noise in a stimulus. When levels of external stimulus noise are high, performance on a task is governed by improved filtering of this external noise equivalent to improved representations of the sensory stimulus. However, when stimulus noise is low, external noise filtering is of less importance and performance variability is primarily attributable to the levels of intrinsic, or internal, noise in the observer. Using this principle, participants complete visual tasks at varying levels of noise, and an equivalent noise function is fit to estimate the observer's threshold level of internal noise (Lu & Dosher, 2008). This approach has been used to show significant reductions in internal noise with perceptual training on motion direction discrimination tasks or texture orientation tasks, where the perceptual representation may already be optimised (Dosher & Lu, 2006). Conversely, for more complex stimuli, learning may be driven predominantly by improvements to the exclusion of stimulus noise, causing better perceptual representations in tasks such as Gabor orientation discrimination or facial discrimination (Gold et al., 1999; Lu & Dosher, 2004). As such, perceptual learning may be driven by distinct mechanisms depending on stimulus complexity, however, the majority of studies have shown a combination of reduced internal noise and improved sensory representations of the stimulus as drivers of learning (Dosher & Lu, 1998, 1999; 2017).

3.1.3: Sequential Sampling Models of Perceptual Learning

Dosher and Lu's perceptual template model provides a potential schematic for the causes of perceptual learning. However, drift-diffusion models (DDM) may provide more comprehensive alternatives for investigating perceptual learning,

allowing for exploration of the effects of learning on a range of measurable and testable parameters, which in turn reflect core perceptual and cognitive constructs that can be linked to neural data. To date, DDMs of perceptual learning have reported that learning is primarily driven by a combination of boundary adjustments and increased drift rates. For example, recent work by Ivanov and colleagues (2022) examined learning of orientation discrimination using the DDM. Here, participants were trained to indicate the orientation of a Gabor grating using a joystick and then were required to repeat the task using eye movements. Fitting of the DDM showed learning was associated with both increased drift rates and decision bounds. However, these effects were reduced when the effector was changed, suggesting that visual perceptual learning may be yoked to effector-dependent integrators and therefore, that the enhanced perceptual sensitivity that characterises perceptual learning cannot arise exclusively from changes in the early stages of sensory processing. Zhang & Rowe (2014) investigated perceptual learning in a speed-accuracy tradeoff task across multiple testing sessions. Early sessions did show increases in drift and bound separation under accuracy emphasis, but only boundary adjustments were observed in later sessions, with no further improvements to drift rate. Here, the authors argue that initial behavioural improvements may be due to refinements of the signal representation via top-down attentional shifting, but the prolonged changes associated with perceptual learning arose from strategic changes in setting the decision criteria.

However, a growing body of evidence suggests that some of these adaptations to the decision bound may be reflective of growing familiarity with task timings. For example, modelling by Liu and Watanabe (2012) attributed the reductions in RT and increased accuracy exhibited by participants after three days of training on a motion discrimination task to a combination of an increased drift rate and a *reduced* decision boundary. They suggest that the lowered decision bound is characteristic of a learned speed-accuracy tradeoff, where participants' improving sensory representations of the stimulus, as evidenced by their stronger drift rates, allow them to make equally reliable decisions while accumulating fewer sensory samples. This grants them an improved "reward rate" (Bogacz et al., 2006), where more correct responses due to better sensory representations

can be made in a shorter amount of time. Petrov and colleagues (2011) have also reported that participants showed improvement not only in their sensory processing but also in their knowledge of task timings. Applying the DDM, the authors found that participants showed increases in drift rate and reductions in non-decision time variability across four testing sessions of a motion discrimination task. As non-decision time influences RT alone, this may reflect learning of the task timings and improved timing of the initiation of the evidence accumulation process.

The above studies disagree markedly on the nature of the boundary adjustments accompanying perceptual learning, with some highlighting progressive boundary increases and others, decreasing boundaries with training. These divergent findings may have arisen from differences in methodology. The studies that showed increased boundary separation with learning had random evidence-onset times and intertrial intervals (Ivanov et al., 2022; Zhang & Rowe, 2014) making learning these task timings difficult even with prolonged exposure. However, both the Liu and Watanabe (2012) study and the Petrov and colleagues (2011) study used tasks preceded by a fixed 500ms interval. This means that over time participants could become aware of and learn when to commence accumulating evidence and when the deadline for response may be due.

Each of these studies implemented models in which the decision bounds were assumed to remain at a constant level across each trial. However, as outlined in Chapter One, and evidenced in Chapter Two, there is growing support for the inclusion of dynamic urgency in perceptual decision-making tasks (e.g., Malhotra et al., 2018; Frazier & Yu, 2007). The possibility of urgency-related changes as a component of perceptual learning has received little attention thus far but represents a promising avenue for further investigation. For example, studies showing boundary collapse adjustments with learning (Liu & Watanabe, 2012; Ivanov et al., 2022; Zhang & Rowe, 2014), may in fact be reflective of changes to dynamic urgency as participants become attuned with task timings across sessions. Thus far, the literature has failed to determine whether the observed boundary adjustments were exclusively made to the starting points of the decision process and/or are implemented dynamically during decision formation

at a rate which itself changes as we learn, meaning a potential contribution to perceptual learning has been overlooked.

3.1.4: Neurophysiological Research on Perceptual Learning

Sequential sampling models are typically fit to behaviour alone, but more recent work has attempted to directly associate adaptation of parameters to changes in the brain. For example, recent work by Jia and colleagues (2018) studied learning on a random dot-motion task using a linear ballistic accumulator model (LBA; see Chapter One: 1.2.1) while fMRI data were recorded. Learning was associated with both increased drift rates and increased 'decision caution', which within this model is calculated as the difference between the starting point and decision bound. These behavioural improvements were associated with increased signal for the trained direction in the frontal eye field (FEF) and ventral premotor cortex (PMv), with enhanced connectivity between sensory areas V3A and the PMv. This result corresponds with the tuning of connections between early visual areas and decision processes proposed in Dosher and Lu's (2012) model. However, as discussed in, LBA models are often too simplistic when compared to the DDM and fail to provide a comprehensive account of constructs such as perceptual learning. Importantly, in order to reduce computing times during model fits, they do not include a within-trial noise parameter at all.

Several electrophysiological studies suggest that learning may be instantiated in the progressive reweighting of connections between early sensory areas such as V1 and later decision-processing areas (Dosher et al., 2013). For example, Law and Gold (2008) trained monkeys on a visual motion task, while recording neurons in the lateral intraparietal (LIP) and medial temporal (MT) visual area. LIP activity offers a neurophysiological index of the evolving decision process as sensory evidence is converted into a motor plan. The authors demonstrated that these LIP neurons exhibited faster build-up rates as the monkeys became increasingly well-practised on the task. Conversely, MT neurons which encode motion direction did not show any increase in their sensitivity. This suggests improvement, not to the encoding of the sensory representation itself, but to the decision readout which serves to reduce the influence of internal noise. Similar

evidence has also been shown in humans. Diaz and colleagues (2017) studied ERP components of perceptual learning on a visual discrimination task. Using single-trial multivariate discriminate analysis they decomposed the electroencephalography (EEG) signal into an Early (~=170ms) component, which they associated with the basic encoding of sensory evidence, and a Late (~=300ms) component, thought to reflect higher-level decision processing. While the Early component appeared to be unaffected by perceptual learning, the Late component showed increases in amplitude as a function of learning. This may be reflective of increased decision bounds (Ivanov et al., 2022; Zhang & Rowe, 2014), however, in this study ERP changes were not directly linked to any model parameters. Collectively, while not directly linking neural data to any model, these studies (Diaz et al., 2017; Law & Gold, 2008) suggest that the behavioural changes exhibited in perceptual learning may be driven by improvements in the latter stages of the decision process, not changes in early sensory encoding.

There are relatively few perceptual learning studies which directly link the DDM to EEG data. Recent work from the O'Connell lab investigated the effect of perceptual learning on several key electrophysiological markers of perceptual decision-making (Devine, 2019). Using an established neurophysiological index of sensory encoding, the steady-state visual evoked potential (SSVEP), the author found significant boosting in the representation of the contrast evidence with increasing exposure to the task, providing some evidence that perceptual learning may manifest in modulations at the earliest stages of perceptual processing. Furthermore, there was a significantly steeper pre-response build-up rate of the CPP and higher amplitudes at response with learning. In close agreement with Ivanov and colleagues (2022), these results indicate that perceptual learning may arise due to a combination of increased decision bound (higher CPP at response) and drift rate (steeper SSVEP and CPP slopes). Similarly, Song and colleagues (2005) reported that perceptual learning with both complex and simple stimuli was associated with modulation of the N1, reflecting improvements to early-stage sensory processing. However only learning with complex stimuli was associated with increased amplitude of P3. This may reflect an increase in decision bound, where participants accumulate more evidence before committing to a decision for a complex stimulus. However, no behavioural

modelling was performed in either of these studies (Devine, 2019; Song et al., 2005), meaning these links between neural data and model parameters could not be directly assessed. Without modelling, it becomes difficult to directly concretely attribute changes to these signals to underlying psychological constructs. A combined EEG and modelling effort may assist in determining the true effects of learning on the decision process. As demonstrated in Chapter Two, EEG data may provide a valuable tool for improving models of perceptual decision-making. By informing model construction we may find evidence for the necessary incorporation of parameters such as dynamic urgency. Furthermore, by constraining these neurally-informed models, we may investigate more complex models in order to glean new insights into the dynamic adjustments that emerge throughout perceptual learning.

3.1.5: Internal Noise and the DDM

As outlined, efforts to model perceptual learning using the DDM have found improvements to the drift rate and some form of adaptation of the decision bound (Ivanov et al., 2022; Petrov et al., 2011; Zhang & Rowe, 2014). However, as exhibited in Dosher and Lu's work (2017), perceptual learning can cause not only improved representations in the stimulus, which can be accounted for in a DDM by improved drift rates, but also significant reductions in an observer's internal noise. In the DDM, an analogue of this internal noise would be the Gaussian 'within-trial noise' applied at each time point to the decision variable. In standard DDMs, inter-trial variability is instantiated in a range of parameters (drift rate, start-point and non-decision time), but the Gaussian within-trial noise is the only source of intra-trial variability (Ratcliff & Tuerlinckx, 2002). Standard DDMs cannot account for changes to this within-trial noise with learning, as it is typically set as the scaling parameter (see Chapter One: 1.2.3), forcing it to remain at a fixed common value for all participants and conditions, or in the case of perceptual learning, testing sessions. If one were to allow within-trial noise to vary by session, this would require another parameter to be fixed across conditions. The most commonly used alternative scaling parameter is the decision bound, yet this too is problematic as decision bounds have consistently been shown to vary as a function of perceptual learning (Ivanov et al., 2022; Liu

& Watanabe, 2012; Zhang & Rowe, 2014). Therefore, to allow noise to vary as a function of learning, would necessitate omitting a known component of the adjustments that characterise perceptual learning. Conversely, forcing within-trial noise to remain consistent across testing sessions may render the model incapable of capturing critical alterations to the formation of the decision variable across time. Either case may lead to erroneous conclusions about the effects of learning on the freely estimated parameters. For example, as all parameter values are expressed in proportion to the scaling parameter, what are reported to be boundary changes with training could potentially arise in reality from a reduction in within-trial noise. One method for investigating training effects on within-trial noise would be to directly estimate the value of another parameter from neurophysiological measurements, thus ensuring that the previously used scaling parameter can vary across testing sessions. With this constraint, within-trial noise values can be freely fit to the behavioural data. As outlined in Chapter Two and previous work (Kelly et al., 2021; Steinemann et al., 2018, Corbett et al., in press), it is now well-established that pre-evidence mu/beta amplitude measurements over the premotor cortex offer valid estimates of the decision variable's starting points. Constraining the boundary to match mu/beta adjustments allows within-trial noise to be freely estimated, as was shown in Chapter Two. For studies of perceptual learning, this would afford the opportunity for within-trial noise to be estimated across sessions, allowing for investigation of how this parameter changes as a function of learning while maintaining the model's capacity to represent any changes to the boundary over time.

Psychophysics may offer another principled method to overcome this dilemma by providing a means of estimating and individually constraining a participant's internal noise (Tibber et al., 2014). However, classic methods for estimating internal noise require double-pass paradigms (e.g., Awwad Shiekh Hasan et al., 2012; Gold et al., 1999), where the same stimulus is repeated several times, or lengthy tasks with varying levels of external stimulus noise (Dosher & Lu, 1999, 2006; 2017). This makes estimating an individual's internal noise time-consuming, meaning it is cumbersome to achieve in tandem with already lengthy EEG data collection sessions. However, a novel fast version of this equivalent noise (EQN) procedure, estimates maximum levels of tolerable noise and noise-

free perceptual thresholds, allowing for the estimation of internal noise across 150 trials, making it easily administrable in conjunction with other behavioural paradigms (Tibber et al., 2014). Indeed, group differences in internal noise levels estimated using the EQN task have been reported across a range of clinical groups when observing a range of stimuli (Chen et al., 2014; Manning et al., 2015; Tibber et al., 2014), although this specific version of the task has not yet been applied in a study of perceptual learning. Using this tool, and in line with Dosher and Lu's (2012) account of perceptual learning, we would expect to see a lowering of internal noise estimates as participants learn. The EQN's measure of internal noise may then present a convenient proxy for constraining within-trial noise levels in DDMs across sessions. This would allow a model which accounts for changes to within-trial noise, allowing it to more successfully pick up on the effects of learning in other parameters

3.1.6: Present Study

The primary aim of this study was to develop a model that can more comprehensively account for all of the decision process adjustments made during perceptual learning. Initially, we aimed to utilise neural data to inform model construction and create a more biologically grounded model of perceptual learning. Furthermore, due to the necessity of a scaling parameter, to date, sequential sampling models have been unable to test for changes to within-trial noise alongside changes to other parameters. To achieve this, decision bounds were constrained to match motor preparation signal measurements, allowing for the model to account both for changes to the bound and within-trial noise with learning. In doing this, we aimed to assess if model-estimated measures of within-trial noise were consistent with the behavioural EQN internal noise estimate. We then utilised the EQN internal noise estimate as a further constraint, allowing for a wider range of possible parameter effects to be explored. In particular, this was devised with the intention of testing for possible dynamic bound adjustments as participants become more familiar with the temporal sequence of the trials. Using this combination of behavioural and electrophysiological constraints, we sought to provide a biologically grounded behavioural account of perceptual learning that was sufficiently sensitive to

capture more subtle changes amidst the more dominant adjustments previously reported in perceptual learning studies.

3.2: Methods

3.2.1: Participants

Sixteen participants took part, four males and twelve females, with a mean age of 20.72 years (SD = 2.49). Two participants failed to complete all sessions and as a result, were not included in any of the following analyses. A post-hoc power analysis was performed using G*Power (Faul et al., 2007). This indicated that given a sample size of 14 across three sessions, with a significance criterion of 0.05 and a power of 80%, the study would have sufficient power to detect effects of medium to large effect sizes (f = 0.36) Participants all had normal or corrected to normal vision with no personal or family history of neurological illness or injury and were free from medication for psychiatric conditions. Participants were recruited via the School of Psychology's online research portal. Participants were reimbursed with research credits. Participants provided informed written consent before testing.

3.2.2: Stimulus Features

Stimuli were generated through MATLAB using PsychToolbox and presented on a 51cm gamma-corrected CRT monitor, 1024x768 resolution, with a refresh rate of 60Hz. Both tasks used moving dot stimuli which appeared around a fixation dot (4 pixels) on a grey background in a circular aperture of 15°. Each stimulus consisted of 100 white dots each with a diameter of 0.44 degrees. Dots updated their position every 3 frames. Dots moved at a rate of 3°/s and had a lifetime of 6 frames.

3.2.3: Equivalent Noise Task

An efficient version of the Equivalent Noise Task was used (EQN: Tibber et al., 2014; Manning et al., 2015). Rather than using several noise levels, this fast version of the EQN uses a zero-noise condition to constrain an Equivalent Noise function at their noise-free perceptual threshold, and a high-noise condition to estimate the participant's maximum tolerable noise level. An equivalent noise

function is then fit to estimate two parameters: internal noise and sampling efficiency. In the high-noise condition, dots moved at a mean direction of +/- 45° to the right or left of a vertical line through the fixation point. The standard deviation of the dots' direction was titrated using a QUEST system to find the maximum level of tolerable noise at which participants achieved 82% accuracy (Figure 3.1.B: Right). In the Zero-Noise condition, the standard deviation of dot-motion was 0, meaning there was no stimulus noise while the mean direction was incrementally decreased to find the value at which participants achieved 82% accuracy (Figure 3.1.B: Left), estimating a perceptual threshold with no external stimulus noise. These trial conditions were interleaved, with participants indicating their chosen direction using a mouse click, left or right. Seventy-five trials of each condition were presented. Upon completion, these values for maximum tolerable noise and minimal sensitivity to directional discrimination were used to fit the equivalent noise function shown in Figure 3.1.C. Prior to their first run of the task, participants completed a practice version of 30 trials with increasing difficulty.

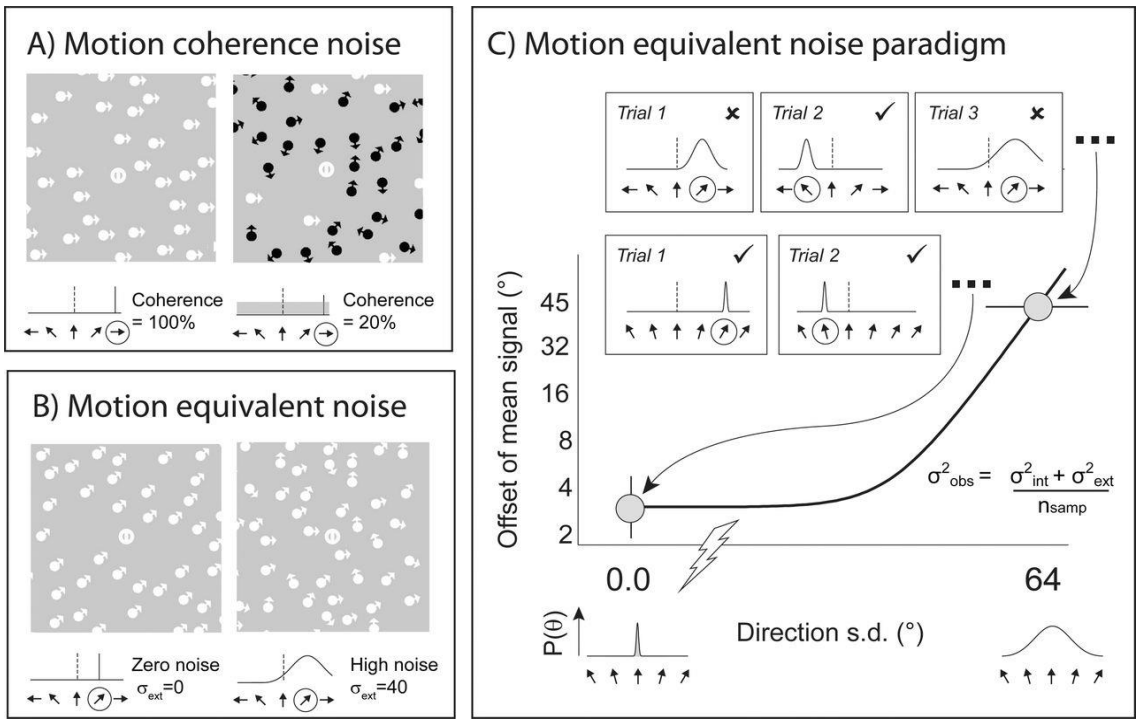

Figure 3.1: Schematic of Equivalent Noise function from Tibber and colleagues (2014). A) Shows coherent motion task, where portions of dots move to either right or left, while the remaining dots move randomly. B) *Left:* Zero-noise condition, where all dots move in the same direction and participants are required to indicate their motion either to the left or right of a vertical dissecting line. The difficulty is increased by reducing the

degree of offset from this vertical. *Right:* High-noise condition where all dots move with a mean direction of ±45° from a vertically dissecting line. The difficulty is increased by increasing the standard deviation of the dot motion, increasing the stimulus noise. C) Shows the psychometric Equivalent Noise function, constrained using individually titrated high and zero-noise conditions. Reprinted from: Tibber, M. S., Kelly, M. G., Jansari, A., Dakin, S. C., & Shepherd, A. J. (2014). An Inability to Exclude Visual Noise in Migraine. *Investigative Ophthalmology & Visual Science*, *55*(4), 2539–2546. https://doi.org/10.1167/iovs.14-13877.

3.2.3: Coherence Task

The random-dot motion task was adapted from the coherence task introduced in Tibber and colleagues' work (2014). Participants identified the direction of a subset of coherently moving dots while the remaining dots moved randomly. A lower proportion of these coherently moving dots, therefore, constitutes a lower coherence level and a more difficult trial (lower signal-to-noise ratio). Coherent dots moved horizontally either right or left and participants indicated their chosen direction using a mouse click. The coherent motion was preceded by 1000ms of random motion, in which all dots moved randomly. This was done to prevent the overlap of visually evoked potentials from stimulus onset with task-relevant decision signals. The coherent motion was presented for a total of 1500ms, during which a participant could respond at any time. There was an inter-trial interval of 750ms before the next trial began, during which only the fixation point was presented. To encourage learning, a points system was used. Participants were granted 20 points for a correct response, 0 points for an incorrect response and -10 points for deciding before the coherent motion began (too fast) or for failing to decide the coherent motion ended (miss). After every 60 trials participants were given a self-timed break, during which they could see their total points score vs the total points achievable. Trial-to-trial feedback was only given for misses and too-fast responses; for misses, the central fixation point flashed red, and for responses before coherent motion onset, the point flashed blue. Participants were informed there would be a point-based monetary reward to incentivise improved performance, however, all participants were given the full reward amount (€8 per session).

3.2.4: Procedure

All procedures complied with the Declaration of Helsinki and were approved by Trinity College Dublin's School of Psychology Research Ethics Committee. Participants were required to attend 3 sessions across one week, with efforts to keep the time of each session consistent. Participants were seated in a darkened, sound-attenuated room, using a chin rest to prevent excess movement placed approximately 51cm from the screen. In each session, participants began by

completing one block of the equivalent noise task to estimate internal noise. In session one this was followed by a QUEST-titrated staircase of the coherence task to obtain a coherence threshold which elicited an estimated 70% accuracy. Once the coherence value was chosen, participants completed 5 blocks (180 trials each) of the coherence task at the given coherence level, giving 900 trials total for each session. The session ended with the completion of a second internal noise estimation. This procedure was repeated, without the staircase procedures in the following two sessions. EEG data were recorded only for the coherence task, as the speed of stimulus presentation in the equivalent noise precluded the ability to extract event-related potentials free from visually-evoked signals.

3.2.5: EEG Acquisition and Pre-processing

Electrophysiological data were recorded using a continuous 128-channel Biosemi ActiveTwo system at a sample rate of 512Hz. Movements of the external eye musculature were recorded using two vertical electrooculogram (VEOG) electrodes placed above and below the left eye. The data were analysed using custom code in MATLAB utilising features of the EEGLAB toolbox (Delorme & Makeig, 2004). Data were detrended to remove the first-order linear trend. To remove slow drift while allowing for the emergence of slow moving-EEG signals, a conservatively low High-Pass filter was applied with a cut-off of 0.01Hz. A Low-Pass filter with a cut-off frequency of 38Hz was also applied. To remove a 20Hz flicker frequency evoked by the dot updating, a 20Hz notch-band filter was applied for the plotting of all EEG signals. Statistical analyses were performed on unfiltered data. Noisy channels were identified using EEGLAB's bad channel-detect function and interpolated. An average reference was then applied, and data were epoched into stimulus-locked (-500ms pre-coherent motion onset to 2000 post coherent motion onset) and response-locked (-800ms pre-response and 300ms post-response). Baseline correction was applied to both epochs, subtracting the average activity across all sessions in the -500ms to -400ms window prior to target onset. This allowed for potential pre-coherent motion build-up of EEG signals during the 500ms of random motion. Artefact trials were identified as trials with over 100 microvolts of activity in any electrode in the

stimulus-locked epoch from -200ms pre-coherent motion onset until 100 samples after response and for differences in voltage above 175 microvolts across the two VEOG channels in the same window. These trials were removed from analyses. To enhance signal localisation a Laplacian current-source density (CSD) transform was applied. To perform mu/beta power analysis a short-time Fourier transform (STFT) was performed, using a boxcar window size of 400 samples, stepping forward in increments of 25 samples.

3.2.6: EEG Signal Electrode Selection

Following an approach adopted by Kelly et al (2021), we measured the onset of the motor potential contralateral to the responding hand in order to estimate the point at which the decision bounds were crossed. This was calculated to allow us to identify an appropriate window for measuring the CPP's amplitude at the time of choice commitment. Motor potential electrodes were chosen for each hemisphere as the electrode site with the maximal positive difference in amplitude between trials in which a left versus a right response was made and vice versa, based on a topography centred -150ms to 0ms on response (see Figure 3.11).

The CPP was chosen as a marker of evidence accumulation. Based on the timing of the motor execution signal of 119ms before response, a window of -135 to -75ms around the response was analysed to select a search window for individual CPP electrode selection. The search space for selecting CPP electrodes was constrained to a predefined cluster of electrodes based on a topography of the same window shown in Figure 3.4. Individual CPP electrodes were then selected for each participant as the electrode with the maximal amplitude in the same window.

Mu/beta amplitudes over the motor cortex were used as an index of response preparation, pre-evidence starting point adjustments, and urgency (Kelly et al., 2021). Electrodes were chosen for each individual based on the identification of the electrodes contralateral to the responding hand that had the largest excursion from their pre-coherence starting levels (-200 to 0ms, see Figure 3.5) to the time

of response execution. The CNV was examined as a further, evidence-independent marker of dynamic urgency (Devine et. al, 2019). Electrodes were selected individually for subjects as the electrode with the most negative going slope in the period of -400ms to -100ms before response (see Figure 3.8). The search range was selected based on the topography of the slopes of the ERP over electrodes in the same time window.

3.2.7: Analysis Procedure

Statistical analyses were performed using JASP, and plotting was completed using MATLAB custom codes and JASP. Error bars shown represent the standard error of the mean. Where no interaction terms are reported the interaction was non-significant. Where Mauchly's test of sphericity indicates the assumption of sphericity is violated, Greenhouse-Geisser corrected degrees of freedom and statistics are reported.

3.2.8: Modelling Procedure

G^2 error was calculated to assess the similarity of real and simulated reaction-time distributions, binned into 5 equally spaced quintiles. As in Chapter Two, error scores were optimised using a Bayesian Adaptive Direct Search algorithm (BADS, Acerbi & Ma, 2017) which combines local Bayesian optimisation with a slower grid exploration. Model fitting procedures for the standard full DDM, neurally-informed and neurally-constrained variants are described below. To ensure the reliability of the G^2 estimate, each model was run using the winning parameters but 100 random seeds. Akaike's Information Criteria were calculated to enable model comparison, penalising more complicated models with additional parameters (Vrieze, 2012). Akaike Weights were used to compare the conditional probability of the best-performing models (Wagenmakers & Farrell, 2004). These compare the log-likelihood of a model to the mean log-likelihood of all models. The probability of the winning model compared to another model can then be estimated (see Table 3.6).

3.3: Results

3.3.1: Behaviour

To examine changes in behaviour on the EQN task across and within sessions a 3x2 factorial repeated measures was used (Session x EQN Measurement). A significant reduction in EQN estimate across sessions was observed (Session: $F_{(2,28)}$ = 4.643, p = 0.019). Bonferroni corrected post-hoc tests indicated that this reduction was significant only between sessions one and two (t = 3.03, p = 0.016), with a non-significant increase between sessions two and three (t = 1.797, p=0.252, Figure 3.2.A). There was no significant change in EQN measurement within a session (Estimate: $F_{(1,13)}$ = 0.991, p = 0.338).

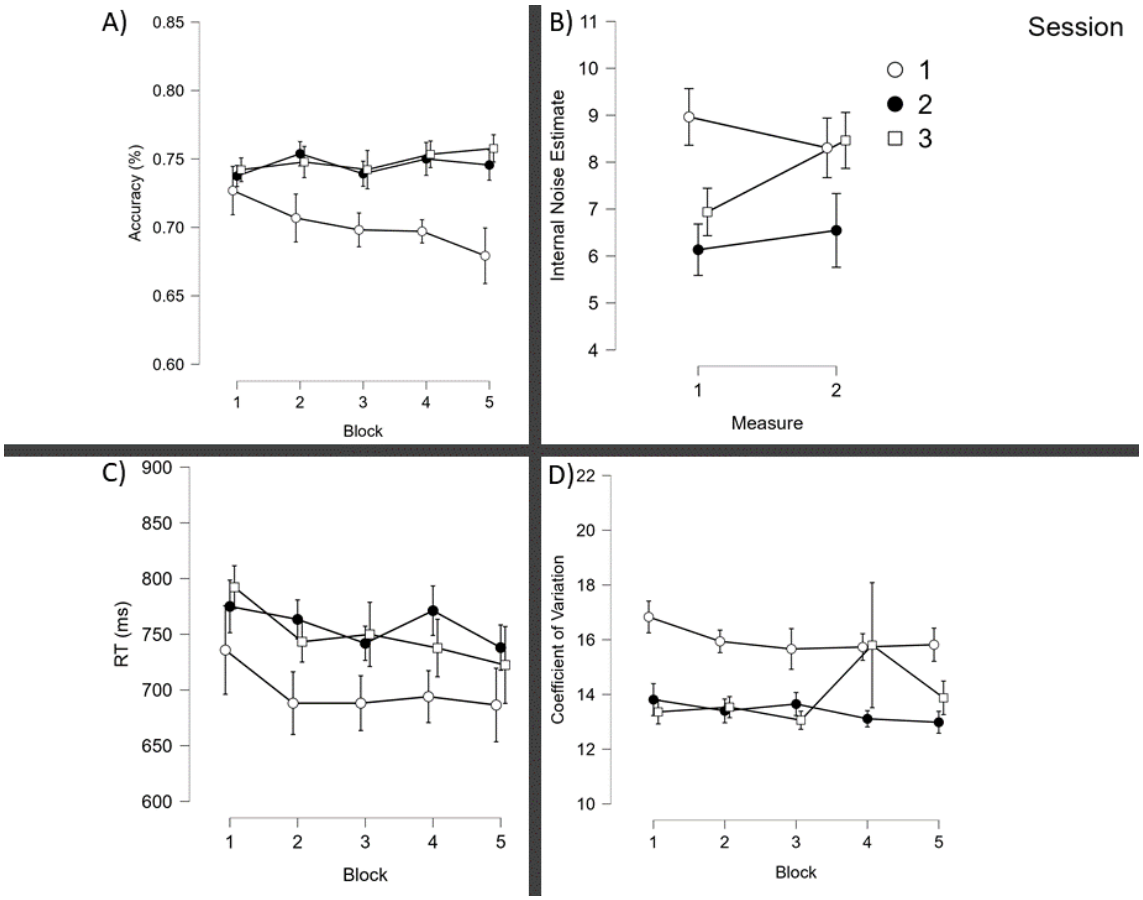

Fig. 3.2.A) Internal noise estimates by session, showing a significant reduction which plateaus with a slight rebound in session 3. B) Significant increase in accuracy between sessions one and two, but no increase into session three. C) Significant reduction in reaction time variability as measured by the coefficient of variation (mean reaction time/standard deviation). No reduction in session three. D) significant reduction in median RT across blocks within one session, but no effect across sessions.

Accuracy on the random dot motion task increased significantly across sessions (Session: $F_{(1.336, 16.034)} = 14.147$, $p < 0.001$, Figure 3.2.B). Post-hoc tests indicated that similar to the EQN, this increase occurred from session one to two (Mean Difference = 4.4%, $t = 4.424$, $p < 0.001$) but plateaued between sessions two and three (Mean Difference = 0.3%, $t = 0.346$, $p = 1$). No significant improvement within a session was observed (Block: $F_{(4,48)} = 0.375$, $p = 0.825$). Median reaction time was shown to significantly reduce across blocks within a session (Block: $F_{(4,48)} = 2.867$, $p = 0.033$, Figure 3.2.C), but not across sessions (Session: $F_{(2,24)} = 1.958$, $p = 0.163$). The reaction time coefficient of variation (ratio of the standard deviation to the mean), was shown to reduce significantly across sessions (Session: $F_{(2,24)} = 8.963$, $p = 0.001$, Figure 3.2.D), but not within sessions (Block: $F_{(4.047, 20.29)} = 0.621$, $p = 0.522$). Again, post-hoc analyses indicated that this effect was significant only between sessions one and two ($t = 4.008$, $p = 0.002$), not from two to three ($t = 0.824$, $p = 0.418$). Overall behavioural metrics indicate substantial behavioural improvements between sessions 1 and 2 that were maintained, but not increased in session 3.

3.3.2: Standard Drift-Diffusion Models

To investigate the effect of learning on latent psychological variables, standard full DDMs were fit to the data, allowing each of the bound, drift and non-decision time to vary by session. To investigate potential changes in within-trial noise by session, in one model variant, the boundary was fixed at 0.1 for all sessions as the scaling parameter, with noise estimated independently for each session.

These model runs gave the results shown in Table 3.1. Here despite considerable trial-to-trial flexibility granted by variability parameters, model fits are poor (all $G^2 > 135$). A model allowing non-decision time to vary by session outperformed others in both raw error ($G^2 = 135.58$), and model comparison scores (AIC = 151.58). An increase in non-decision time from 487ms in session one to 578ms in session three was observed, with a reduction to 563ms by session three (Table 3.2).

Table 3.1) Model variants for initial standard full diffusion model fits, showing the number of parameters, raw error score and model comparison score. The column on the left indicates which parameter if any was allowed to vary by session. All remaining parameters were fixed across sessions.

Parameter Varying	**N params**	**G^2**	**AIC**
None	6	219.64	231.64
Bound	8	192.82	208.82
Drift	8	216.59	232.59
Non-Decision Time	8	135.58	151.58
Within-Trial Noise	8	188.60	204.60

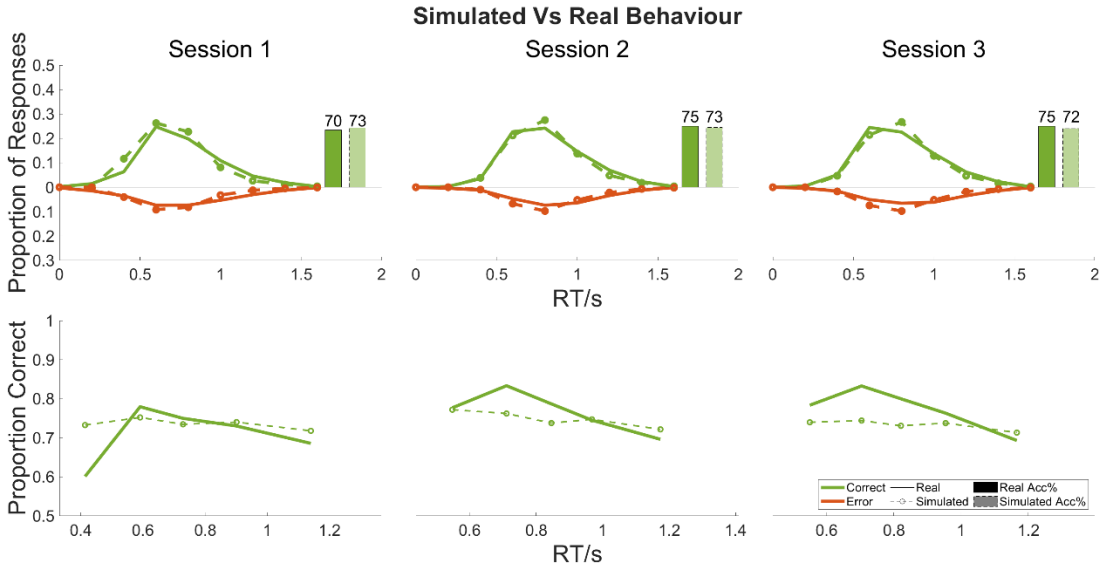

Figure 3.3) Real and simulated behaviour for the best fitting standard DDM, where non-decision time was allowed to vary by session. The top row shows real vs simulated reaction time distributions and accuracies. The bottom row shows real and simulated accuracy quantiles for five RT quintile bins.

Table 3.2) Shows model estimated parameter values, error scores and model comparison scores for winning standard DDM with non-decision time-varying. Non-decision time shows an increase across sessions. Bnd = Bound. N.D = Non-decision Time, S.P = Start-Point. G^2_{100} = G^2 error averaged across simulations with 100 random seeds.

Bnd	Drift	N.D. Time 1	N.D. Time 2	N.D. Time 3	Drift Var.	N.D. Time Var.	S.P. Var.	G^2	BIC	AIC	G^2_{100}
0.052	0.13	0.5	0.581	0.576	0.13	0.24	0.005	135.6	199	152	135.7

The models failed to accurately capture the shape of reaction time distributions, in particular substantially underestimating the rate of fast and slow errors, visible in the quintile plots shown in Figure 3.3. Here the real data presents an inverted u-shape that the simulated data cannot reproduce. This is a similar outcome to that observed for the standard DDM fits outlined in Chapter Two. To improve these models, we turned to electrophysiological markers of the decision-making process. Through this it was hoped to identify key psychological features which were not captured in standard DDM, allowing for a better fitting and more comprehensive model.

3.3.3: Neurally-Informed Model Construction

3.3.3.1: Early Accumulation

Based on evidence for the pre-evidence build-up of noise in advance of stimulus onset (Devine et. al, 2019, see section 1.2.2 and Chapter Two), we tested for this in order to determine if a period of premature accumulation should be allowed for in the model. To this end, the centroparietal positivity was used as a marker of effector-independent evidence accumulation.

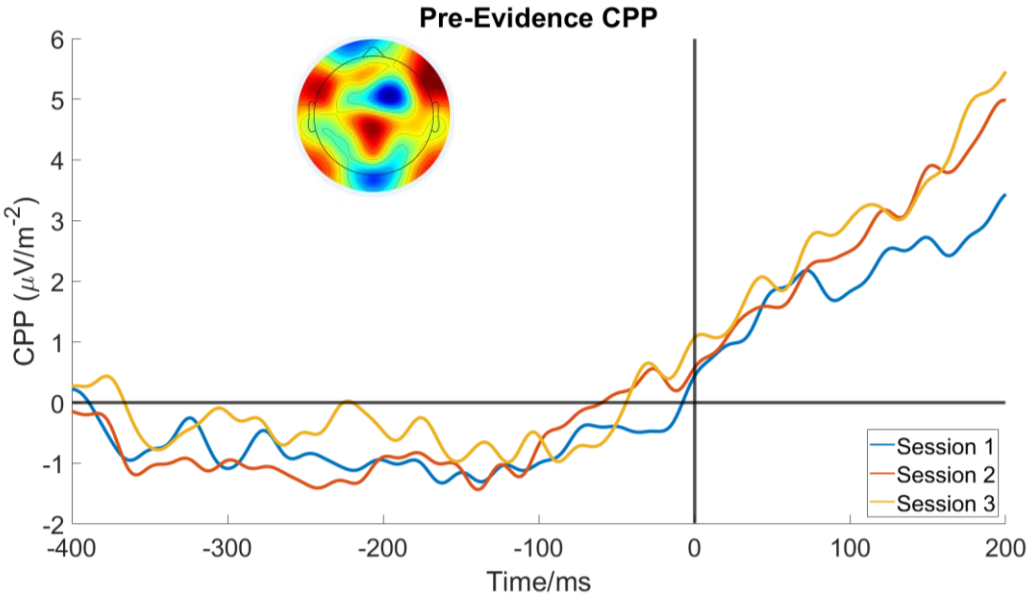

Fig. 3.4) Pre-evidence CPPs by session showing significant pre-evidence build-up in all sessions. CPP electrodes were chosen for individuals from the response-locked topography (-135ms to -75ms) shown.

A series of independent samples t-tests were used to investigate the early build-up of the CPP in advance of and in the early stages of coherent motion onset (-200ms to 100ms). A Bonferroni corrected p-value for 3 comparisons of 0.017 was used. The slope of this signal in this time window was shown to be significantly greater than 0 across all 3 sessions (df = 13, t_1 = 5.435, p_1 < 0.001, t_2 = 4.297, p_2 < 0.001, t_3 = 3.408, p_3 = 0.005). A repeated measures ANOVA indicated that slopes in this time range did not differ significantly across sessions ($F_{(2,26)}$ = 0.356, p = 0.704). Consistent with previous observations using paradigms involving a zero evidence lead-in between stimulus and evidence onset (e.g., Devine, 2019, Chapter Two), these results indicate that participants commenced accumulating sensory information before the physical evidence became available or sensory encoding time could have elapsed. This neural feature

would account for the high proportion of fast errors observed in the behavioural data. The full DDM could not accommodate this despite its inclusion of starting point variability, as it does not allow accumulation to begin until non-decision time has elapsed.

3.3.3.2: Starting Point Adjustments and Urgency

We aimed to evaluate evidence for the presence of pre-evidence starting point adjustments. In a DDM these starting point adjustments are mathematically equivalent to a shift in the bound and can be used as measures of pre-evidence adaptations to the decision boundary. Pre-evidence mu/beta build-up has also been shown to reflect urgency dynamics, with the pre-evidence adjustments thought to reflect a decision bound already on its way to collapse (Kelly et al., 2021). Mu/beta (8-30Hz) activity over the premotor cortex was analysed as an indicator of motor preparation and urgency.

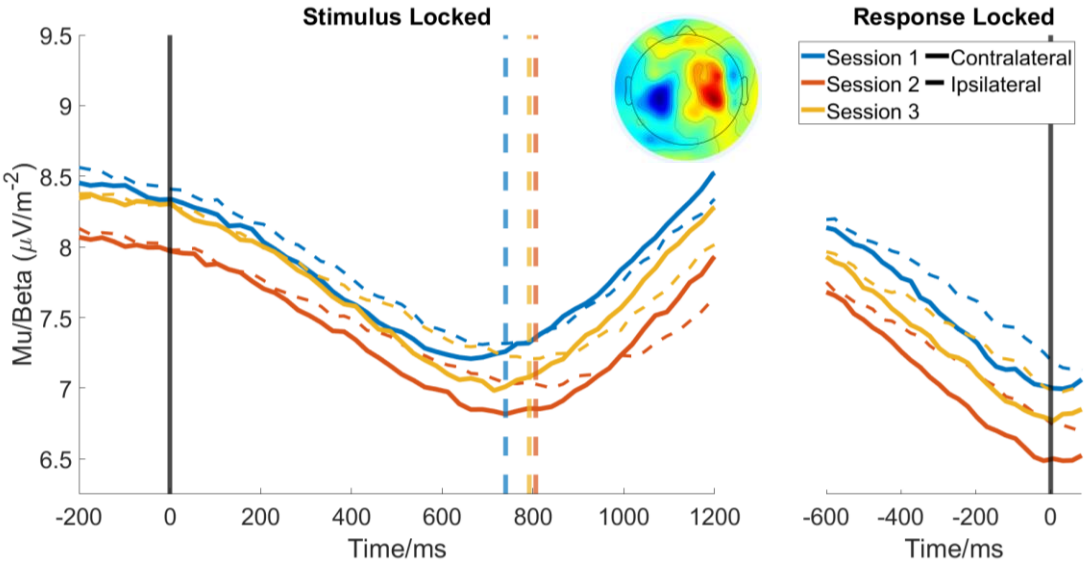

Fig. 3.5) Stimulus and response locked non-normalised mu/beta by session. Mu/beta electrodes were chosen for individuals from the response-locked topography subtracting trials on which a left response vs a right response was made (-100:0ms, shown above). Observed power differences potentially due to task-irrelevant features, explored in Figure 3.6.

Initial analyses indicated substantial average amplitude differences across the whole trial between sessions across stimulus and response-locked epochs (see Figure 3.5). Mu/beta however has been shown to desynchronise to a consistent threshold across responses, irrespective of evidence strength or reaction time

(Kelly et al., 2021). This difference in amplitude across sessions aroused suspicion as to whether these observed differences reflected true adjustments to mu/beta average amplitude or were the effects of behaviourally-irrelevant measurement error across sessions such as cap placement. To test this, trials were split by median RT and analysed using a 3 x 2 factorial ANOVA (Session X RT Bin), allowing us to test whether, within a session, mu/beta showed a consistent desynchronisation to a threshold. Here, there was no significant within-session difference in contralateral mu/beta at response between fast and slow trials ($F_{(2,13)}$ = 0.085, p = 0.775, see Figure 3.6). This shows that within a session, mu/beta is desynchronising to a threshold independent of reaction time and that the observed amplitude differences were likely not associated with differences in the decision process. To eliminate these differences, the mu/beta traces for each session were normalised by subtracting the mean mu/beta-activity over contralateral electrodes at response for each session, giving the signals shown in Figure 3.7.

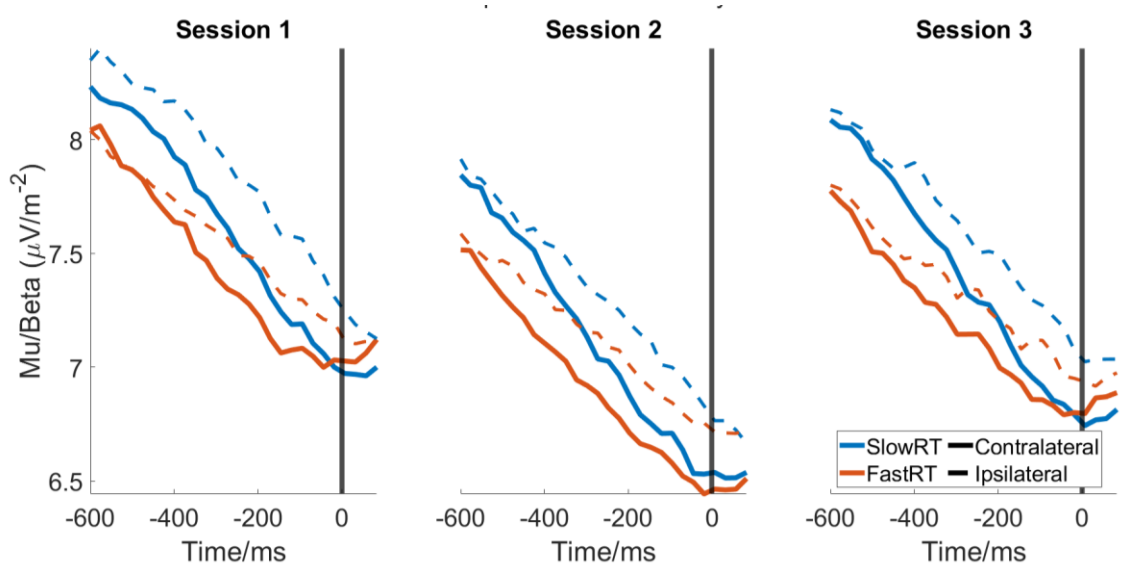

Fig. 3.6) Contralateral and ipsilateral mu/beta traces split by median reaction time for each session. No significant difference between contralateral mu/beta at response for fast or slow responses in any session, suggesting an accumulation to a motor-bound within a session. This suggests mu/beta acts as a bounded signal, legitimising an inter-session baseline.

In order to test for pre-evidence adjustments to the starting point, a repeated measures ANOVA was used to analyse pre-evidence mu/beta power in the window of -200 to 0ms preceding coherent motion onset averaged across contralateral and ipsilateral hemispheres. Here there was shown to be no significant difference across sessions in pre-evidence amplitude ($F_{(2,26)}$ = 0.813, p = 0.454), or slope ($F_{(2,26)}$ = 0.665, p = 0.523). Visual examination of the pre-evidence window shown in Figure 3.7, however, implied a potential decrease in motor preparation by session.

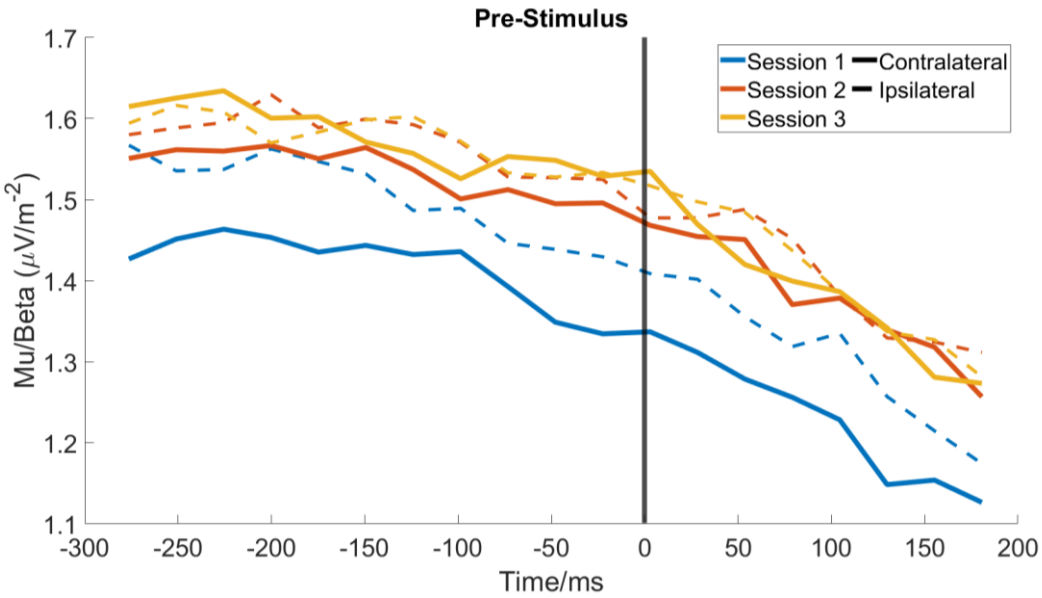

Fig. 3.7) Contralateral and ipsilateral mu/beta traces in the pre-evidence window. A significant negative slope is observed in the pre-evidence window for session one. While differences in pre-evidence amplitude are non-significant, the trend suggests a decrease in pre-evidence motor preparation across sessions.

In order to test for any dynamic pre-evidence adjustments to the decision bound a series of independent samples t-tests were used to investigate whether mu/beta slope averaged across contralateral and ipsilateral cortexes showed any significant build-up prior to coherent motion onset (-200ms to 0ms). As the test was performed for each of the three sessions, a Bonferroni corrected p-value for 3 comparisons of 0.017 was used. Significant negative build-up was shown only in Session 1 ($t_{(13)}$ = -3.681., p_1 =0.003), but not for the remaining sessions (all p >= 0.09). These analyses suggest a motor preparation in session one that is building towards its threshold before any evidence appears on the screen.

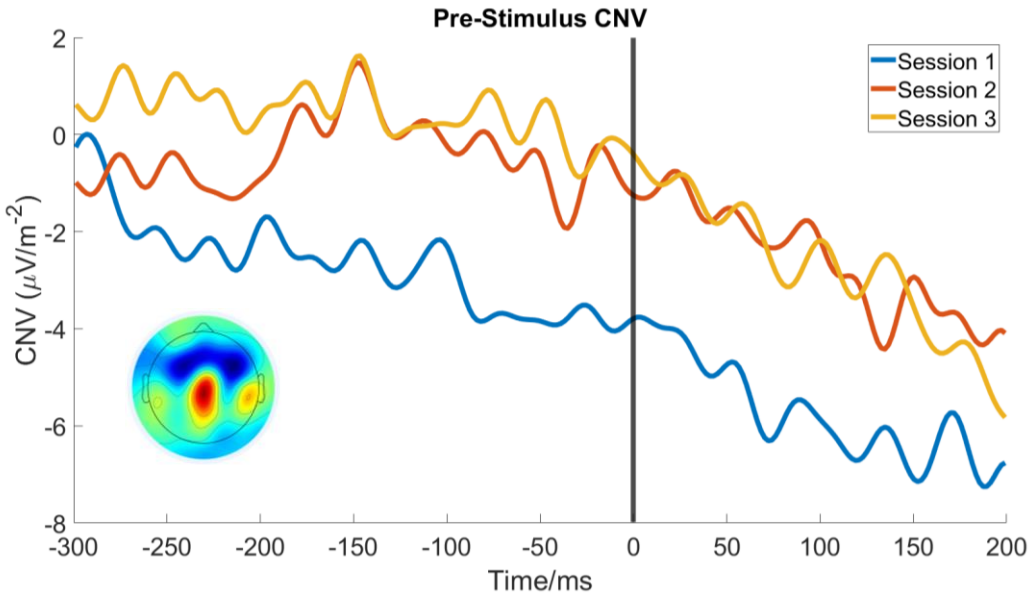

Fig. 3.8) Pre-evidence CNV. Session one shows a significantly negative slope, indicating a build-up of the CNV before evidence onset. Response-locked topography shows the slope in the -400:-100ms before response. This was used to select electrodes with the most negative going slope in this window.

The Contingent Negative Variation was investigated as a further source of evidence for dynamic urgency. Analysis of the pre-evidence window (-300ms to 0ms before coherent motion) indicated no significant difference in amplitude ($F_{(2,26)}$ = 1.719, p = 0.199) or slope ($F_{(2,26)}$ = 1.462, p = 0.250). As was the case for mu/beta, one-sample t-tests against 0 indicated slopes were significantly different from 0 only in session one ($t_{(13)}$= 2.873., p = 0.013, all other p > 0.166) There was striking correspondence between the CNV waveforms and those observed for mu/beta with both exhibiting significant build-up prior to evidence onset in Session 1. Combined with the evidence from pre-evidence mu/beta, this provides evidence in support of pre-coherence urgency and motor preparation by the time of evidence onset, lending support for the inclusion of dynamic urgency in the model.

3.3.4: Neurally-Informed Modelling

Using the evidence from the neural data, we constructed neurally-informed models that more accurately reflected the decision processes. Given electrophysiological results indicating substantial motor preparation in advance of evidence onset, alongside significant pre-evidence CNV build-up, an urgency

component was added. As in Chapter 2, this was implemented as a linear decrease of the bound over time starting at evidence onset. Similarly, given the potential pre-evidence build-up of the CPP in all sessions, standard non-decision time was replaced by a drift-onset parameter where only Gaussian noise was accumulated from evidence onset until a model estimated “drift-onset time”, from which the drift rate was also incorporated at each time point. This allowed for noisy accumulation, equivalent to a pre-evidence build-up of the CPP, allowing the model to more accurately mimic the neural signals and capture fast errors. Changes in drift rate, boundary, urgency and within-trial noise were tested by running separate models in which each of these parameters vary by session, alongside a model in which no parameter was allowed to vary. In all of these, within-trial noise was used as the scaling parameter. However, to test for changes in within-trial noise with learning, one model run used the boundary as the scaling parameter, set to 0.1 across sessions as outlined in section 3.3.2 above. These neurally-informed models gave the fits shown in Table 3.3. The inclusion of early accumulation and urgency greatly improved model performance for all variants, with even the worst performing neurally-informed model (No-Change: AIC = 121.5) outperforming the best standard DDM (Non-Decision Time: AIC = 152). Of these neurally-informed models, a model allowing noise to vary by condition provides the lowest error and model comparison score (G^2 = 23.07, AIC = 41.07).

Table 3.3) Showing error and model comparison scores for neurally-informed models. Non-decision time is replaced by a drift-onset time, where noise accumulates from 0 until drift-onset time. Urgency is included as a linearly collapsing bound. The green highlight indicates the winning model where boundary is constrained across sessions and within-trial noise is allowed to vary.

Parameter Varying	N parameters	G2	AIC	Scaling Parameter
None	7	107.57	121.57	Within-Trial Noise
Bound	9	64.11	82.11	Within-Trial Noise
Drift	9	103.8	121.8	Within-Trial Noise
Urgency	9	91.95	109.95	Within-Trial Noise
Drift-Onset	9	61.34	79.34	Within-Trial Noise
Within-Trial Noise	9	23.07	41.07	Boundary

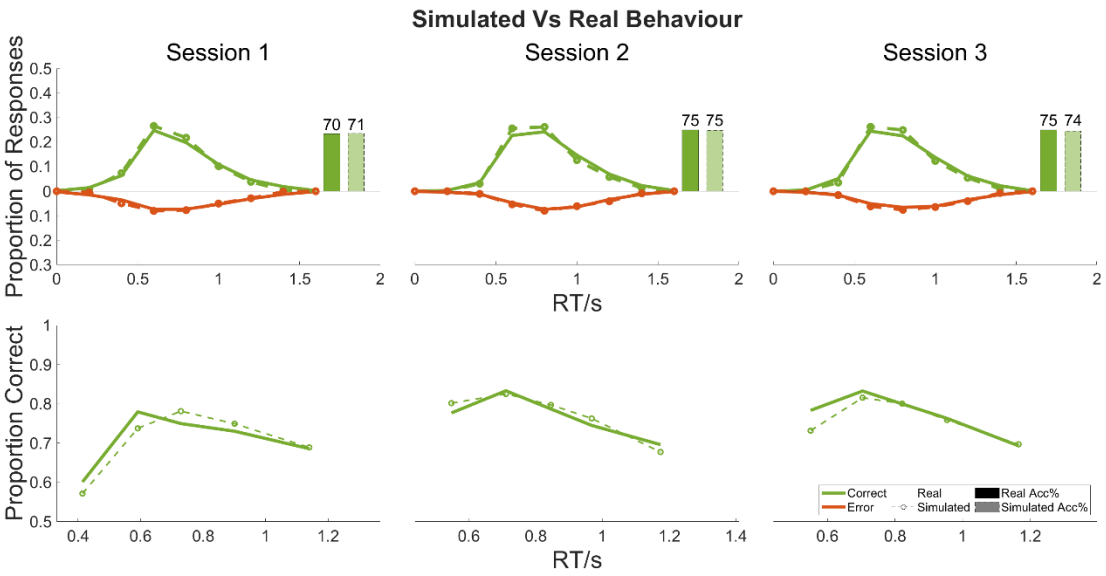

Figure 3.9) Real and simulated behaviour for best fitting neurally-informed model, where within-trial noise was allowed to vary by session. The top row shows real vs simulated reaction time distributions and accuracies. The bottom row shows real and simulated accuracy quantiles for the five RT quintile bins.

Table 3.4) Shows model estimated parameter values and comparison scores. G^2100 indicates the average G^2 when the same parameters are used to simulate trials with 100 fresh random number seeds. Var = Variability, S.P = Start Point, W.T. Noise = Within-Trial Noise.

Drift	Bound	Drift-Onset	Urg	Drift Var	Drift-Onset Var	S.P. Var	W.T. Noise1	W.T. Noise 2	W.T. Noise 3	G^2	BIC	AIC	G^2100
0.1	0.1	0.462	0.1	0.134	0.133	0.0011	0.054	0.039	0.041	23	94	41	24.3

The plotted accuracy by reaction time quantiles indicated that the best-fitting model (variable noise) was successful in capturing behaviour in each session (see Figure 3.9). These adaptations allow the model to capture the fast and slow errors and the change in these proportions across sessions. Notably, the model estimated values of noise track largely with the behaviourally estimated marker of internal noise, which indicates a reduction in within-trial drift-rate noise from session one to two, and a plateau of this reduction by session three (see Figure 3.10). In Figure 3.10, the model-estimated noise values have been normalised relative to their value in session one to bring them onto a similar scale for plotting. This forces parity between behavioural and model-derived noise estimates in session one, but session two shows near-perfect agreement with the model-simulated noise estimate. The model also shows agreement with the EQN in identifying an increase in noise from Session 2 to Session 3 although the EQN rebound is more pronounced.

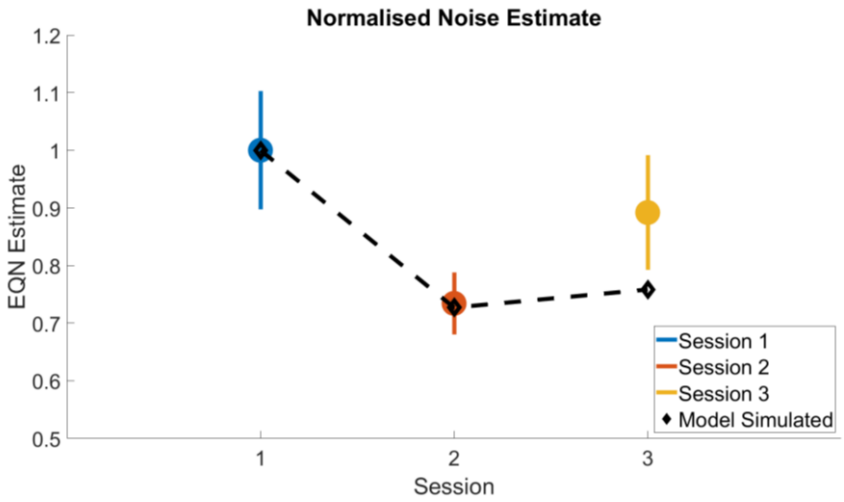

Fig. 3.10) Comparison of model simulated within-trial noise values (coloured markers) for the neurally-informed model with a fixed boundary to behavioural estimates of internal noise (dashed black line) for each session. Error bars show the standard error of the mean. Significant concordance is observed for sessions one and two.

3.3.5: Neurally-Constrained Modelling

Efforts were then made to constrain the models using observable electrophysiological indices, in line with the methods outlined in Chapter Two. This allows for the exploration of more complex models with minimal impact on model comparison scores by not increasing the number of free parameters.

3.3.5.1: Boundary Adjustment Constraint

Although in these data, the differences in pre-evidence mu/beta were non-significant, pre-evidence, mu/beta has been shown to be a good estimate of pre-evidence starting point adjustments to decision bound (Kelly et al., 2021). As such, pre-coherence mu/beta was used to constrain boundary estimates by session. Pre-evidence starting point adjustments were calculated as the mean pre-evidence mu/beta across ipsi and contralateral channels in the -200 to 0ms before evidence onset (see Figure 3.7). Values were normalised by the lowest starting point (Session 1) and divided by ten to bring them in line with the boundary constrained in the Standard-Noise model above 0.1.

3.3.5.2: Motor Time Constraint

Following the method used by Kelly and colleagues (2021), a motor commitment time point was estimated. This has the benefit of enabling the drift-onset parameter to accurately reflect noisy accumulation. Without constraining the time required for motor response, the relative contribution of true “noisy build-up” relative to the motor time in a drift-onset parameter cannot be determined. This motor execution time was taken as the pre-response inflection point of ERP activity measured over the motor cortex contralateral to the chosen direction (see topography in Figure 3.11), calculated as the last point prior to response at which there is a non-negative slope. Although motor potential onset appears closer to response execution in later sessions, these differences were not statistically significant ($F_{(1.418,\ 18.438)}$ = 0.8, p = 0.424, see Figure 3.11). The average motor time across sessions was estimated to be 119ms. This motor time of 119ms was added to the simulated RTs to account for motor execution time. Non-decision

time variability was then applied to this motor-time parameter, adding uniformly distributed variability.

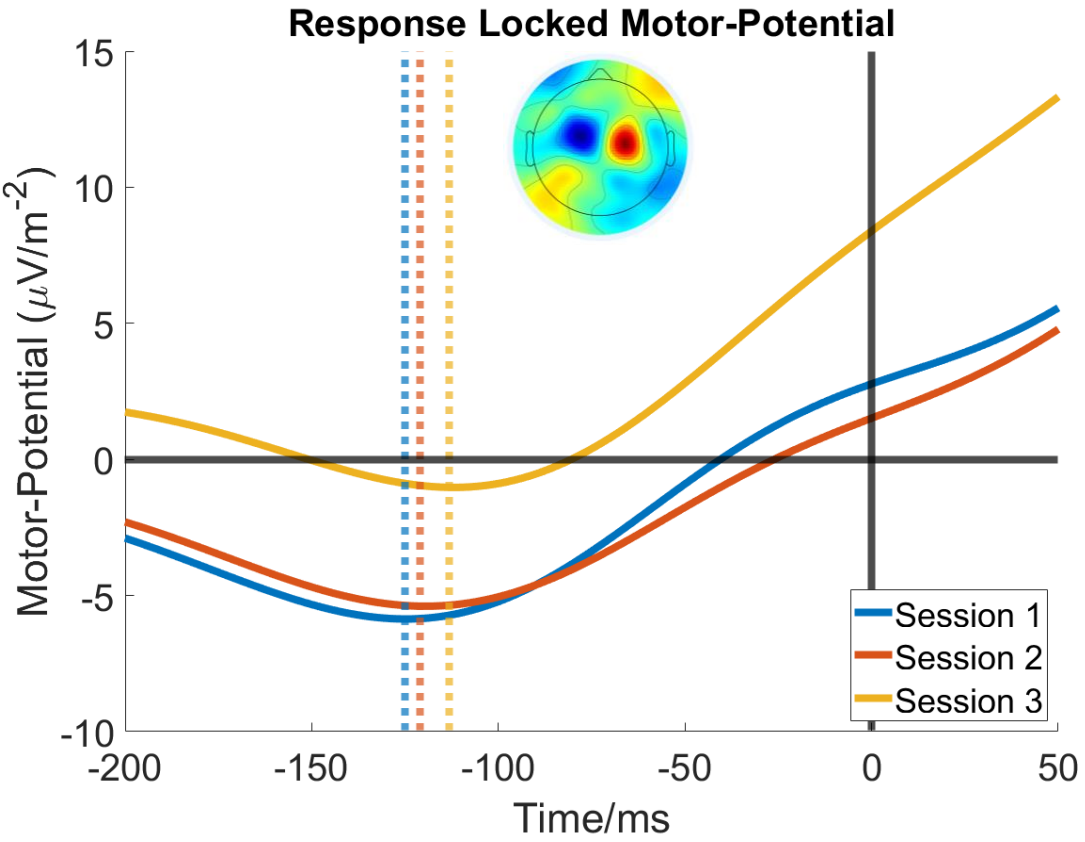

Fig. 3.11) Contralateral pre-movement ERP over the motor cortex in the pre-response window. Dashed lines indicate the inflection points for these signals after which only a positive slope is observed. These were not significantly different across sessions (p = 0.424). Topography shows the difference between trials in which a left versus a right response was made. It maps the response-locked activity in the -150:0ms window centred on the response used for electrode selection.

3.3.5.3: Equivalent Noise Constraint

As there was no observed difference in internal noise estimate within a session ($F_{(1,13)}$ = 0.991, p = 0.338), the average of the two internal noise estimates for each session were used as estimates of within-trial noise for that session. As with the boundary constraint, noise values were averaged across subjects and normalised relative to the session with the lowest noise estimate (Session 1). These were then multiplied by 0.1 to bring them in line with the traditional scaling parameter value for within-trial noise.

3.3.5.2: Neurally-Constrained Model Fitting Procedure

Constrained versions of the neurally-informed models were then run. These were constrained by either mu/beta or internal noise estimate alone, or by both simultaneously. Where both constraints were used, a freely estimated parameter was added to scale the noise value in session one. This was done to ensure that any arbitrary differences between the magnitude of the noise and mu/beta constraints did not unduly affect the search space. In effect, this means that the within-trial noise value for Session 1 was freely-estimated and the values for Sessions 2 and 3 were fixed to a proportion of the Session 1 value, using the percentage changes derived from the EQN estimates.

For each of the constraint methods, models were run allowing each of either boundary, drift, urgency or drift-onset time to vary by session. In models where only one of beta or noise was constrained, a model variant allowing the remaining unconstrained value to vary across sessions was performed (e.g., if the boundary was constrained and the noise was not, a model allowing noise to vary by session was run).

To allow for the possibility of distinct pre-evidence accumulation times and drift-onset times, variants that allowed for both an accumulation-onset time and a drift-onset time were fitted. Here, either accumulation-onset or drift-onset time, or both were freely estimated (earliest possible onset of accumulation-onset time set to -500ms relative to evidence onset). Again, for each constrained model family, variants which allowed accumulation-onset or drift-onset to vary by session were also fit. In general, these models did not provide better fits to the data than one with only drift-onset time estimated.

Table 3.5) Shows a complete breakdown of included parameters in each of the model types. Green ticks mean the parameter was included, and red shaded x's indicate the parameter was not. Variants of all models which included either/both a drift-onset and noisy-accumulation-onset time were also run. In this bottom row, a yellow "-" indicates the parameter was either included or not included depending on the model variant and family. W.T. = Within-Trial Noise. Mot.-Time = motor execution time constrained by the motor potential. Acc.-On = accumulation-onset time, or time-point at which the drift rate began.

	Standard				Variability			Neurally-Informed			Neurally-Constrained			
	Bound	Drift	Ter	W.T. Noise	Drift	Ter	Z	Urg.	Drift-On	Acc.-On	Mot. Time	Bound	W.T. Noise	Noise-Scaling
Standard	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗	✗	✗	✗
Neurally-Informed	✓	✓	✓	✓	✓	✓(mot.-time)	✓	✓	✓	✗	✓	✗	✗	✗
Beta-Constrained	✗	✓	✗	✓	✓	✓(mot.-time)	✓	✓	✓	✗	✓	✓	✗	✗
EQN-Constrained	✓	✓	✗	✗	✓	✓(mot.-time)	✓	✓	✓	✗	✓	✗	✓	✗
EQN/Beta-Constrained	✗	✓	✗	✗	✓	✓(mot.-time)	✓	✓	✓	✗	✓	✓	✓	✓
Accumulation-Onset	✓	✓	~	✓	✓	✓	✗	~	✓	✓	~	~	~	~

3.3.5.3: Constrained Model Fits

The top two performing models for each constraint family are presented in Table 3.6. Overall, the neurally-informed but unconstrained model with freely estimated noise was shown to provide the best fit to the data, outperforming standard DDMs and the constrained models. Models constrained by both mu/beta and internal noise estimate with urgency varying provided the next best fits.

Table 3.6) Showing two best model fits for each of the variations of neural constraints. A non-constrained model with noise varying remains the best model fit. Green shows the winning model: Neurally-informed with Within-Trial Noise varying, orange shows the second-best model: EQN & Mu/Beta-Constrained with urgency varying, and red shading shows the third best model: EQN & Mu/Beta-Constrained with accumulation-onset time varying. Conditional probabilities comparing the probability of the winning model to each individual model are calculated based on the Akaike Weights (Wagenmakers & Farrell, 2004).

Free Parameter	**N Params.**	**G^2**	**AIC**	**BIC**	**$G^2_{100Seed}$**	**AIC: Difference from Winning**	**Probability**
Standard							
Non-Decision Time	8	135.58	151.58	198.5	134.67	110.51	9.93e+23
Within-Trial Noise	8	188.6	204.6	251.52	192.74	163.53	3.24E+35
Neurally-Informed							
Within-Trial Noise	9	23.07	41.07	93.85	24.33	0	1
Accumulation-Onset	9	30.49	48.49	101.27	31.53	7.42	40.85
EQN-Constrained							
Drift-Onset	9	35	53	105.78	35.69	11.93	389.55
Bound	9	37.27	55.27	108.05	39.33	14.2	1211.97
Beta-Constrained							
Within-Trial Noise	9	39.55	57.55	110.33	43.09	16.48	3789.54
Drift-Onset	9	44.88	62.88	115.66	44.38	21.81	54447.92
EQN and Beta-Constrained:							
Urgency	9	25.32	43.32	96.1	24.96	2.25	3.08
Accumulation-Onset and Urgency	11	26.19	48.19	112.7	25.86	7.12	35.16

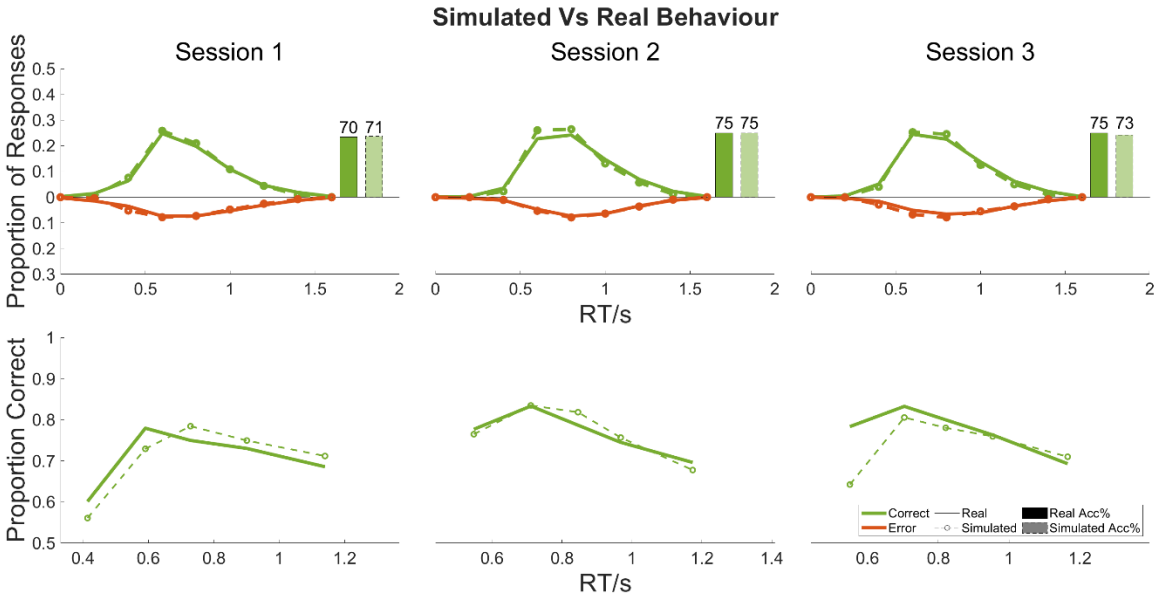

Figure 3.12) Real and simulated behaviour for the best fitting neurally-constrained model, where bound and noise are constrained with urgency allowed to vary by session. The top row shows real vs simulated reaction time distributions and accuracies. The bottom row shows real and simulated accuracy quantiles for five RT quintile bins.

Table 3.7) Shows model estimated parameter values, error scores and model comparison scores for winning EQN & Mu/Beta constrained model. Boundary and within-trial noise are constrained by behavioural/neural constraints and were not freely estimated. A noise scale parameter was freely estimated to ensure the model was not limited by arbitrary differences between the constrained boundary and EQN measure. Participants show higher bounds and increasing urgency across sessions. W.T Noise = Within-Trial Noise. Urg. = Urgency component. S.Pnt. = Start Point. Var = Variability

Bound 1	Bound 2	Bound 3	W.T. Noise 1	W.T. Noise 2	W.T. Noise 3	Drift Rate	Drift On	Urg. 1	Urg. 2	Urg. 3
0.1	0.106	0.107	0.069	0.051	0.061	0.11	0.3409	0.065	0.078	0.076
Drift Var.	Mot Time Var	S.Pnt Var	Noise Scale	G^2	AIC	G^2_{100}				
0.144	0.078	0.007	0.069	25	43	24.96				

While not outcompeting the neurally-informed model with noise varying, models constrained by both mu/beta and the EQN provide a more complex picture of the data. These models indicate an increase in noise and decision bound across sessions, along with an increase in urgency. However, the accuracy-quantile figure shows a poor fit in session three, especially in the proportion of fast errors made.

The best-performing model was a non-constrained but neurally-informed model with boundary fixed and noise reducing across sessions. The second two best-performing models were the Mu/Beta and EQN-constrained models. However, the winning model only outperformed the second-best neurally-constrained model by 2 AIC points. To assess how much more probable the winning model was compared to the neurally-constrained model, Akaike Weights were calculated. These were calculated for all models shown in Table 3.6. These give an estimate of how much more probable the winning model is compared to each individual model (Wagenmakers & Farrell, 2004). Shown in the rightmost column is an estimate of how much more probable the winning, non-constrained model is compared to the model in each column. With this, it can be roughly estimated that the winning neurally-informed model with fixed bounds and noise-varying is roughly 3 times more probable than the winning neurally-constrained Model. However, these models were significantly more complex and had either urgency varying or both noisy-accumulation-onset time and urgency varying by session. This allowed us to investigate more complex effects.

Overall, it is suggested that the main driver of the behavioural improvements were adaptations to within-trial noise. This is seen in the neurally-informed model where the boundary is fixed across sessions. This is similarly observed when boundary adjustments are constrained by pre-evidence mu/beta-constrained alone, where within-trial noise levels are also shown to reduce by session (see Figure 3.13). Collectively, this provides substantial support for reductions in within-trial noise as a contributory factor to perceptual learning. Secondary improvements may be driven by the learning of task-timings, which increase urgency and reduce the premature onset of accumulation prior to evidence appearing.

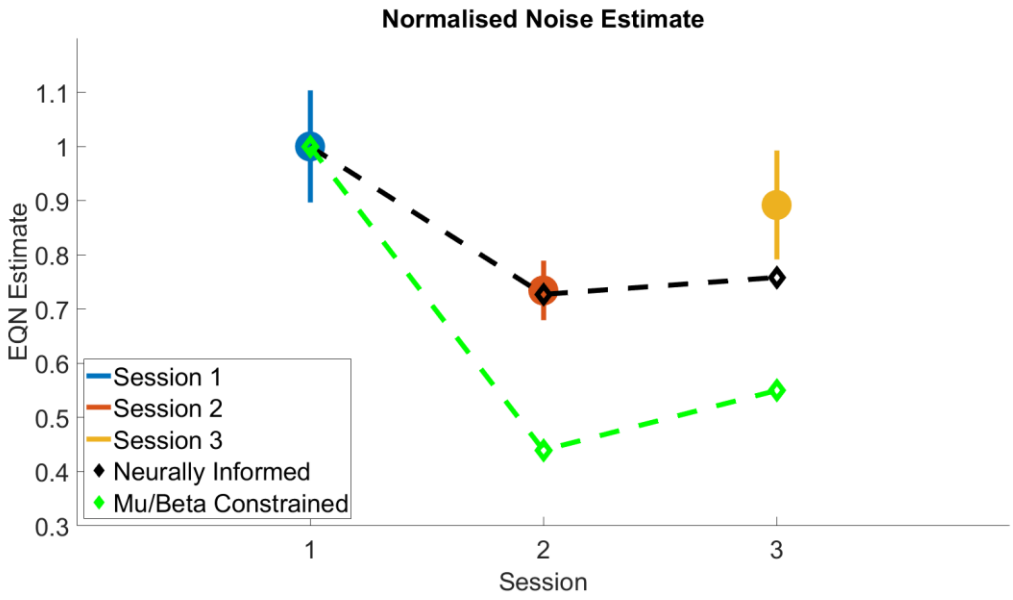

Fig. 3.13) Comparison of model simulated within-trial noise values for the neurally-informed model with a fixed boundary (dashed black line) and mu/beta constrained model with within-trial noise varying (dashed green line) to behavioural estimates of internal noise (coloured markers) for each session. Error bars show the standard error of the mean. Significant concordance is observed for sessions one and two.

3.3.6: Model Validation through Independent Neural Data

In order to establish whether our winning neurally-constrained model was in concordance with the observed neural data, key features of the model were simulated and compared to our neural indices. These comparisons were the evidence accumulation variable vs the CPP, the model simulated motor response variables vs contra/ipsilateral mu/beta, and the linearly collapsing urgency component vs the CNV.

3.3.6.1: Evidence Accumulation

Model estimated decision variables from the winning neurally-constrained model (Mu/Beta & EQN Constrained: Urgency Varying) were simulated and plotted to track model concordance with our neural marker of evidence accumulation, the CPP.

Figure 3.14 shows the model simulated stimulus and response-locked cumulative decision variable for the Mu/Beta and EQN-constrained model with

urgency varying. Descending coloured dashed lines illustrate the models' collapsing bounds, with the dashed green line indicating drift-onset time. The dashed green line represents drift-onset time.

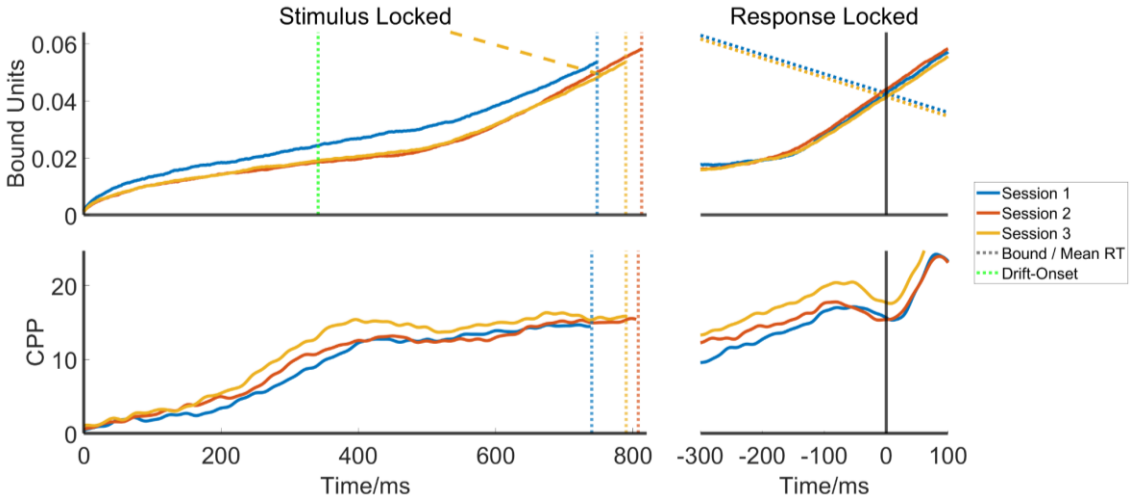

Fig. 3.14) Stimulus and response locked CPP signals by session, compared to model simulated urgency components for the neurally-constrained model with urgency varying. Dashed coloured lines represent the mean RT for each session. Green dashed lines represent drift-onset time. There is little concordance between observed and simulated decision variables.

Comparing these features to the model simulated data, we can see that while providing strong fits to the behavioural data, the models do not seem to successfully reproduce key features of the CPP, including CPP amplitude at response and the increasing build-up by session visible in the stimulus-locked traces.

We then compared our neural data to the winning unconstrained neurally-informed model with Within-Trial Noise varying (see Figure 3.15). Here there is a fixed boundary across sessions. Again, there is little concordance between the simulated decision variable and the CPP.

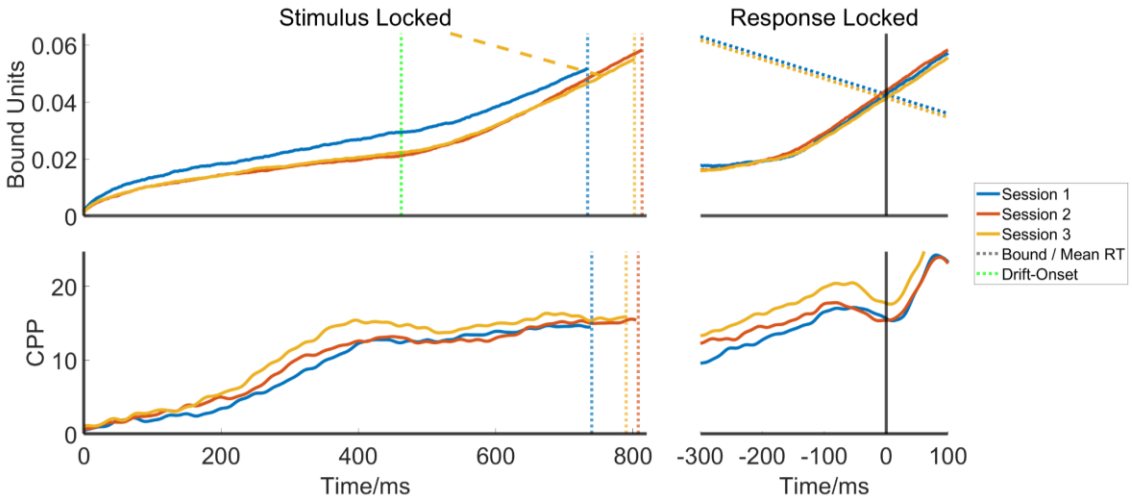

Fig. 3.15) Stimulus and response locked CPP signals by session, compared to model simulated urgency components for the neurally-informed model with within-trial noise varying. Dashed coloured lines represent the mean RT for each session. Green dashed lines represent drift-onset time. The boundary is fixed across sessions. There is little concordance between observed and simulated decision variables.

3.3.6.2: Motor Signal

Motor signals were simulated as the negatively signed, cumulative evidence for either the chosen or unchosen response with the negative going urgency component added. Notably, the neurally-constrained simulations show a strong resemblance to the observed neural data. While constrained to match mu/beta-starting levels, the simulated signals exhibit the same pattern of desynchronisation across time (see Figure 3.16).

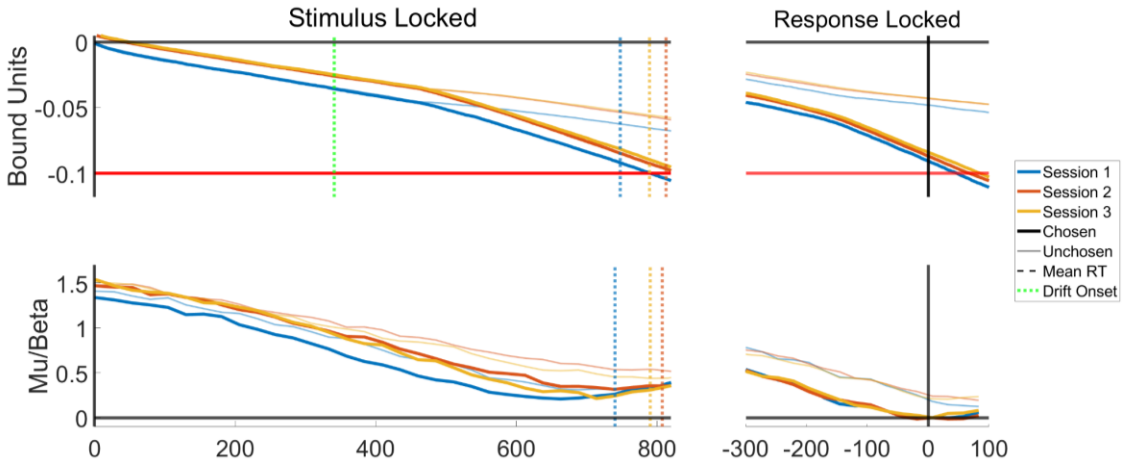

Fig. 3.16 Stimulus and response locked mu/beta signals by session, compared to model simulated urgency components for the neurally-constrained model with urgency varying. Dashed coloured lines represent the mean RT for each session. Green dashed lines represent drift-onset time. Significant concordance is seen, especially around evidence onset. The Green dashed line represents drift-onset

We then compared our neural data to the winning unconstrained neurally-informed model with within-trial noise varying (see Figure 3.17). Here there is a fixed boundary across sessions. Even without pre-evidence starting point adjustment, the signals show significant concordance. The increased build-up of the motor signal for the chosen response in earlier sessions is driven by the higher levels of noise.

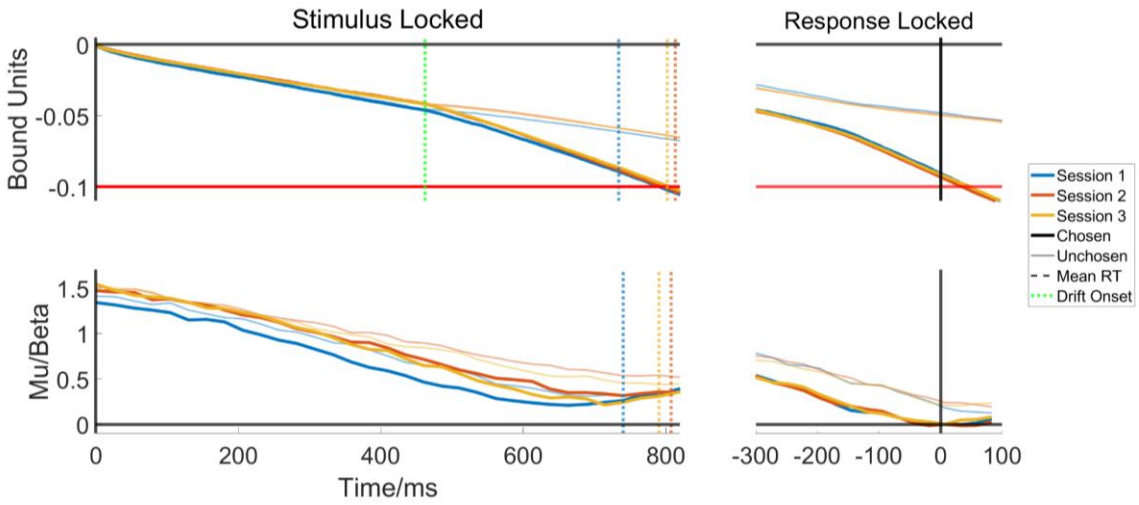

Fig. 3.17) Stimulus and response locked mu/beta signals by session, compared to model simulated urgency components for the neurally-informed model with within-trial noise varying. Dashed coloured lines represent the mean RT for each session. Significant concordance is seen, especially around evidence onset. The green dashed line represents drift-onset time.

3.3.6.3: Urgency

The raw urgency components, without added evidence, were extracted and compared to the CNV as a marker of urgency. The lower amplitude of the CNV at evidence onset response aligns with the starting-point adjustments before evidence onset. Similarly, the simulated urgency slopes align with the CNV slopes across sessions (see Figure 3.18). Concordance between the neural signals and model simulated variables at response is less apparent.

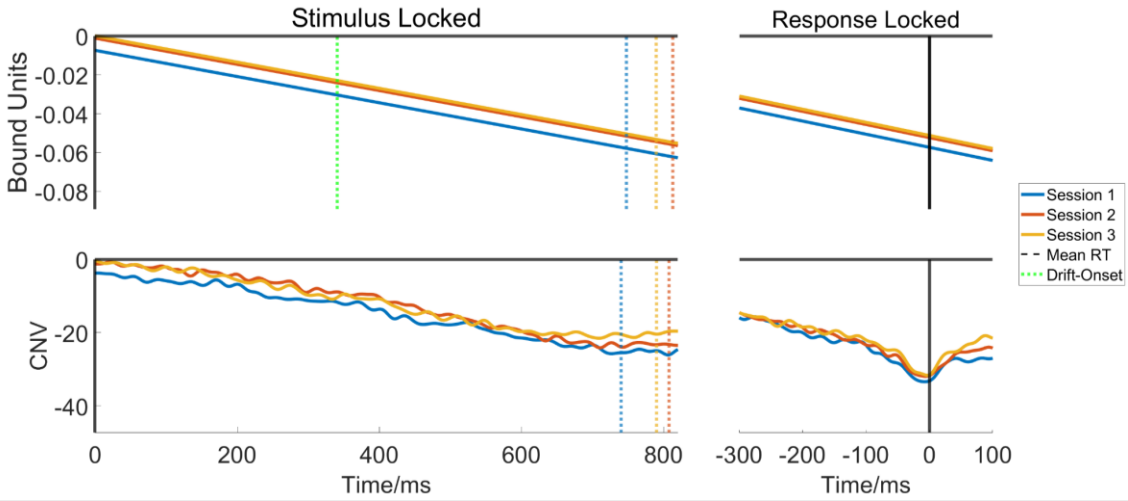

Fig. 3.18) Stimulus and response locked CNVs by session, compared to model simulated urgency components for the neurally-constrained model with urgency varying. Dashed coloured lines represent the mean RT for each session. Green dashed lines represent drift-onset time. Significant concordance is seen, especially around evidence onset.

We performed the same comparison for the winning neurally-informed model. Here as there is a fixed urgency component across sessions, and the simulated urgency component is unable to capture the early differences in CNV amplitude around evidence onset (see Figure 3.19).

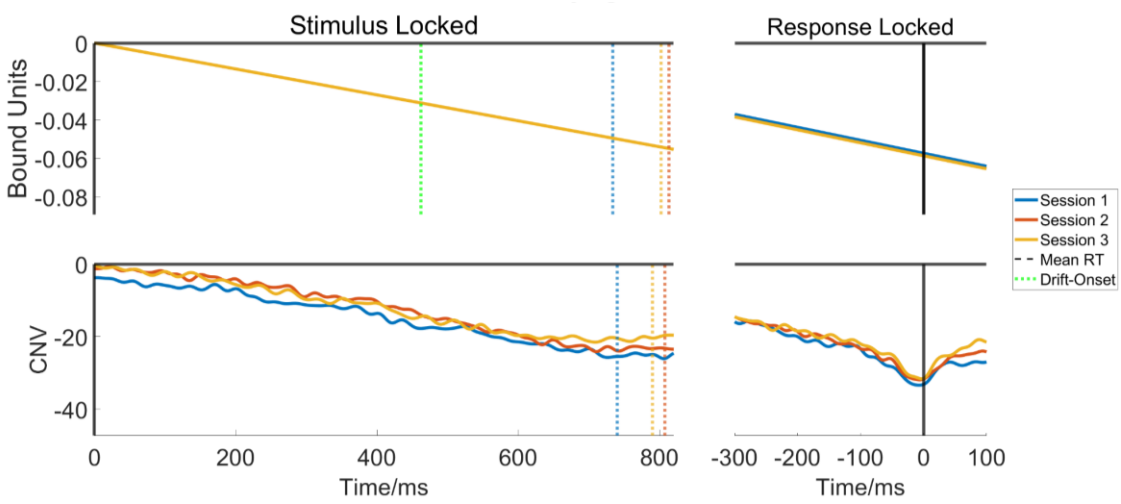

Fig. 3.19) Stimulus and response locked CNVs by session, compared to model simulated urgency components for the neurally-informed model with within-trial noise varying. Dashed coloured lines represent the mean RT for each session. Green dashed lines represent drift-onset time. Urgency is fixed across sessions.

3.4: Discussion

The psychological factors which contribute to perceptual learning remain poorly understood. These improvements may arise from a combination of increasingly precise expectations about the task demands (e.g., stimulus timing, Liu & Watanabe, 2012; Petrov et al., 2011) and/or a sharpening of psychophysical sensitivity born from the more efficient encoding of the sensory evidence (Dosher et al., 2013). This improved sensory representation can be achieved through the suppression of internal noise, or improved exclusion of external noise, resulting in improvements to sensory encoding (see Dosher & Lu, 2017 for a review). These adaptations can be represented in sequential sampling models, like the DDM, as strategic adjustments to the decision bounds and increases in drift rate, respectively (Ivanov et al., 2022; Petrov et al., 2011; Zhang & Rowe, 2014). However, the standard DDM has overlooked potential reductions in within-trial noise as a driver of visual perceptual learning. This study aimed to assess the contributory factors leading to visual perceptual learning, using neural data to inform and constrain existing models. It also aimed to explore psychophysical estimates of internal noise as potential estimates of model within-trial variability.

Over the course of repeated exposure to the task, participants exhibited some of the behavioural hallmarks of perceptual learning: significant reductions in reaction-time variability and increases in accuracy. In addition, we observed a reduction in the psychophysical estimate of internal noise using the EQN. However, the behavioural effects of time-on-task tapered off after Session 2. This may have been due to the nature of the task, which was a basic motion discrimination task. Performance on these “low-level” perceptual judgement tasks is shown to exhibit relatively small improvements in comparison to learning on more complex tasks such as facial discrimination which require input from a broader range of perceptual processes (Fine et. al, 2002). However, previous studies of perceptual learning on contrast-discrimination tasks have shown learning which continues linearly across five sessions (Devine, 2019). Similar work on a frequency discrimination task indicated that improvements were only shown in participants who completed 900 trials per session, while no improvements were shown in participants completing 360 trials per day,

suggesting a high number of trials per session may be needed to elicit consistent learning on basic perceptual judgments (Wright & Sabin, 2007). Participants in this study completed 900 trials per session, and it is possible that with additional training days further improvements may have been observed.

Initial standard DDM fits indicated changes in non-decision time as being the primary driver of the observed perceptual learning effects. However, this did not align with previous research, and the model provided poor fits to the behavioural data. Analysis of key neural signatures of the decision process indicated that the model was missing important parameters which meant the DDM was not reflecting the decision process accurately. For example, urgency dynamics were readily apparent in mu/beta activity and the CNV for Session 1, which indicated that response preparation began in advance of evidence onset and could account for the slow errors in the data. Furthermore, there were signs of pre-evidence accumulation of noise in all sessions in the early ascent of the CPP. This would account for the high rate of fast errors seen in the behavioural data, which despite the inclusion of starting point variability, the DDM could not accommodate. With this in mind, the DDM was adapted to include urgency as a linearly collapsing bound, while allowing for the accumulation of noise up to the drift-onset parameter. Model variants which freely estimated both an accumulation-onset time and a drift-onset time were tested, however, they did not provide better fits to the data when corrected for the number of parameters. These neurally-informed models included one model where the boundary was fixed across sessions as a scaling parameter and within-trial noise was allowed to vary, allowing us to assess the effects of learning on within-trial noise in a DDM. The neurally-informed models provided a significant improvement on standard models in their fits to behaviour and revealed a novel parameter effect. Here, the winning model implemented a consistent bound across sessions, attributing the perceptual learning effects to reductions in within-trial noise. Critically, these models simulated estimates of within-trial noise were validated by the closely aligned behavioural EQN estimates, grounding the parameter estimate in an independent empirical index of the efficiency of sensory encoding. Collectively, this highlights the importance of referring to the neural data in order to aid in constructing a more representative model of the decision process.

Examination of mu/beta signals over the motor cortex indicated potential pre-evidence adjustments to response preparation levels. These indicated an increase in motor excursion across sessions, where beta has further to travel from its initial point to its threshold at response. As previously outlined, these pre-evidence adjustments to mu/beta levels have been indicated as analogous to starting point adjustments in the DDM, which are in turn mathematically equivalent to boundary changes (Kelly et al., 2021). While the effects observed here were non-significant, the observed boundary increases were consistent with previous DDM studies of perceptual learning which have shown an increasing boundary with learning (Ivanov et al., 2022; Zhang & Rowe, 2014). The failure of these adjustments to reach statistical significance may have been driven by the relatively small sample size of the current study and the resulting low statistical power. However, given previous evidence in support of decision boundary increases with learning (Ivanov et al., 2022; Zhang & Rowe, 2014), pre-evidence mu/beta levels were used to constrain an increasing decision bound across sessions. From this, a series of neurally-constrained models using pre-evidence beta as a proxy for boundary adjustments and the EQN measurement of internal noise in place of within-trial noise were run. The best performing of these neurally-constrained models was one which allowed urgency to vary, with urgency shown to increase across sessions. These complementary mechanisms may work in concert: as elevated decision bounds lead to more accurate responses, reduced within-trial noise limits RT variability and a growing urgency component across sessions may account for increasing familiarity with task timings (Liu & Watanabe, 2012).

Our winning neurally-constrained model provided slightly worse fits to the data than the neurally-informed model with decreasing noise. However, the differences in AIC were small. Analysis of the Akaike Weights (Wagenmakers & Farrell, 2004) for the winning models suggested that the winning neurally-informed but un-constrained model was approximately 3 times more likely than the winning neurally-constrained model with urgency varying. In order to assess the extent to which each of these two models reflected the neural data more accurately, we simulated the core components of each model and compared them to the observed neural signals. The simulated motor preparation and

urgency signals of the neurally-constrained models showed a high degree of correspondence with their empirical neural signatures. That is, although the mu/beta and urgency build-up rates did not vary significantly between sessions, they exhibit a highly similar numerical trend to that of the simulated data. Of particular note was the concordance between the CNV and simulated urgency component, which tracked the ramping of urgency across the trial, alongside its pre-evidence adjustments. The neurally-informed model with fixed boundaries and no urgency varying by condition was incapable of recreating these similarities. Neither the neurally-informed nor neurally-constrained models simulated decision variable reflected the CPP, however, the inclusion of noisy build-up until a drift-onset captured the early kick-off of the CPP shown in the neural data. Clearly, however, the discrepancies raise questions. At a minimum, the winning model may not be giving a full account of what is happening in the brain. There may be a drift rate increase alongside noise decrease that the model is not able to disentangle as they have similar effects on behaviour. There are also question marks about how different noise levels should translate to CPP amplitude. One possibility is that when there is a lot of sensory noise, the sensory evidence is down-weighted at the evidence accumulation stage, causing a shallower build-up. This may be thought of as an extension of the work of Law and Gold (2008) who showed increased spike rates in the choice-selective neurons of the monkey LIP with training, but not the motion-sensitive MT. However, the MT neurons were shown to be increasingly predictive of choice across sessions, indicative of a reweighted readout of sensory information into the decision process with time. Here we suggest that decreasing within-trial noise at the visual encoding level may result in an increased reweighting of the evidence at the decision-formation stage. While this study did not measure early visual neural representations directly, previous work in the O'Connell lab supported this (Devine et al., 2019), where an increased differential SSVEP signal with perceptual learning indicated an improved sensory representation of the stimulus. This was coupled with increases in CPP amplitude and build-up across sessions, suggestive of adaptation to both the decision formation stage and the earliest neural representation of a stimulus. This may be indicative of reduced ongoing perceptual noise in these early sensory stages, which results in a stronger reweighting of evidence at the decision-formation stage. This may

have resulted in a model which, while capturing behaviour, does not successfully reflect how these changes to within-trial noise manifest in neural indices of evidence accumulation.

The neurally-constrained model did not perform as well as the neurally-informed model with a fixed boundary and noise varying freely across sessions. The fact that the differences across sessions in pre-evidence mu/beta were non-significant may have been driven by low statistical power. Instead, it may suggest that adaptations to the decision bound were sufficiently small that the improvement in the data fits brought about by allowing for a varying bound was outweighed by the penalisation applied for increasing the model's complexity. The EQN estimate was also shown to estimate a lower internal noise value than a model-simulated noise for session 3 (Figure 3.13). This may have been driven by measurement error in the constraints which could have contributed to the slightly poorer fits in the neurally-constrained models which incorporated this parameter. For example, the EQN internal noise estimate is a faster version of traditional estimates of within-trial noise, requiring only 150 trials (Tibber et al., 2014). A higher number of trials may have yielded better estimates of internal noise that may have more closely resembled the model. Increasing the trial count, for example to 250 trials, would not substantially lengthen the time of the overall testing session and may be beneficial in obtaining more reliable estimates of internal noise. Furthermore, as behavioural internal noise estimates may potentially encapsulate sources of both between and within-trial variability, the measure may not directly reflect within-trial noise in isolation. Future work may benefit from collecting internal noise estimates across a large sample, using the DDM to assess the extent to which internal noise loads onto between-trial or within-trial noise. Alternatively, neurally-derived means of estimating ongoing noise in the decision process may provide a more suitable constraint, constraining variability within-trial noise in a given trial in isolation.

One future avenue for research may be to analyse within-trial variability features of the CPP, such as signal entropy (Sleimen-Malkoun et al., 2015), to attempt to find a neural correlate for within-trial noise. This could provide further neural evidence for reductions in intra-trial decision noise with learning and would align

with Dosher and Lu's (2017) suggestion that learning occurs through the increasingly efficient weighting of visual inputs into a decision formation signal. At a methodological level, identifying reliable means of estimating within-trial noise via the CPP would present an easily obtainable neural proxy for measuring individual differences, distinct from any between-trial variability. This would provide a novel and convenient mode for creating neurally-informed models, eliminating the need for behaviourally estimated internal noise measurements which add to testing time such as the EQN (Tibber et al., 2014). Furthermore, neurophysiologically-derived estimates of within-trial noise would provide a principled parameter constraint that may prove crucial in allowing neurally-informed models to capture key subtleties of the decision process, which have so far gone undetected. This advance could have profound consequences for basic perceptual research and our understanding of the neural underpinnings of many clinical disorders. If the CPP is to be used as a neural constraint for the diffusion model, some potential limitations must be addressed. As outlined in Chapters One and Two and evidenced in the necessity to rescale mu/beta-signals across sessions in this study, there may be substantial individual, non-task-related variability in signal presentation (Antonakakis et al., 2020). This may also have contributed to the failure of our model-simulated decision variables to align with our electrophysiological signals. Potentially novel methods for addressing this via Independent Component Analysis will be assessed in the coming chapter.

Collectively, the neural and behavioural constraints do offer a promising possibility for improving models, but one which could not be fully validated in this study. However, the neurally-constrained models did succeed in producing simulations which were far more in line with the observed neural data than the winning non-constrained model. A larger sample size may have more clearly indicated these effects, pointing more conclusively to adaptations both to the levels of within-trial noise and urgency dynamics. However, this study is the first to devise methods of accommodating changes to within-trial noise in a sequential sampling model. These reductions are observed across both neurally-constrained and unconstrained behavioural modelling and are supported by psychophysical estimates of internal noise via the EQN. Overall, this presents

compelling evidence for the reduction of ongoing within-trial noise in visual perceptual learning.

While the neurally-constrained model did not outperform the winning neurally-informed model, this work identified potential increases in urgency in perceptual learning. Including an urgency parameter significantly improved model fits, allowing it to capture the slow errors viewed in the behavioural data. This was supported by evidence from the CNV and mu/beta. The winning Mu/Beta and EQN-constrained model may suggest possible improved learning of task timings, which occurs alongside the improved representation of the sensory stimulus through the suppression of noise. Urgency may be suggestive of a high degree of specific, non-generalisable learning, which may not be transferable to other domains or tasks. This may be in part contributing to the high degree of domain and task specificity observed in the perceptual learning literature (Fahle, 2005). The reductions to within-trial noise, however, may reflect an improved sensory representation, such as a selective tuning of visual neurons (Dosher & Lu, 2017) that may be more generalisable to other tasks and domains (McGovern et al., 2016). It was beyond the scope of this study to test for transferability of learning, but this question could be addressed in future research using similar approaches. On tasks with variable timings, the effects of urgency on learning should be diminished. Such a paradigm could also test across various stimuli to assess the extent to which lowering of within-trial noise represents a highly specific tuning of visual resources on a specific task or a more global lowering of internal noise, potentially through the boosting of top-down attentional resources.

In conclusion, this work presents an in-depth model-based analysis of the factors underpinning perceptual learning. Through a novel synthesis of behavioural modelling, psychophysical testing and electrophysiology, neurally-informed and constrained models point to reductions in within-trial noise. It is the first study to use neurally-informed modelling to highlight the reduction in internal noise across learning and also demonstrates the value of the EQN as a potential behavioural proxy for an individual's average level of within-trial noise. Neurally-constrained models also indicated potential boundary increases and increased urgency as potential contributors to the observed behavioural improvements. Future work

may explore the application of these constraint methods to individual-level modelling to further our understanding of the complex cognitive processes underpinning perceptual learning and decision-making.

Chapter Four:

Evaluating the Influence of Age-Related Frontalisation of the P3b on Comparison of Evidence Accumulation Signals Across Age-Groups

4.1: Introduction

The high temporal resolution of electrophysiological data makes it possible to track the emergence of a range of neural signals which reflect distinct stages of the decision-making process (O'Connell, Dockree, et al., 2012), granting us an increased understanding of how perceptual decision-making unfolds in the brain. As explored in previous chapters, a growing body of work has attempted to directly link aspects of these signals to computational model parameters with considerable success (Chapters One, Two & Three; Kelly et al., 2021, Turner et al., 2013, McGovern et al., 2018, Corbett et al., in press). Neural signals provide a valuable tool for improving standard sequential sampling models of decision-making, enabling us to create biologically-informed models which can produce more detailed explanatory accounts. Neurally derived markers of the decision process can also be used to directly constrain certain model parameters, allowing us to explore increasingly complex models without the risk of overfitting behavioural data and removing the need for arbitrary scaling parameters, a major limitation for between-group modelling (Chapter Two, Chapter Three, Kelly et al., 2021, McGovern et al., 2018, Corbet et al., 2018).

While these advancements are promising, our ability to rely on these signals as constraints for modelling relies on the extent to which they reflect specific decision processes accurately. These assumptions are strongly tested when we attempt to compare signals taken from two fundamentally distinct groups. Here, differences in signal topography and amplitude may well be attributable to differences in the underlying decision-making processes, but they may just as well reflect task-irrelevant features of the neural data which are borne out differently across different groups. These may include group differences in signal location (e.g., O'Connell, Balsters, et al., 2012; van Dinteren et al., 2014b), or changes in volume conduction caused by physiological differences in skull thickness (e.g., Antonakakis et al., 2020). For example, in Chapter Two and previous studies (Stacey et al. 2021), elevated beta-power in older adults is observed. As mu/beta is seen to reliably desynchronise to a threshold at response, irrespective of response time and difficulty (Kelly et al., 2021), these differences do not seem to be related to the decision process and the signal can

be normalised (see Chapters Two). However, not all signals cross a threshold at response; therefore, this normalisation approach cannot be widely applied. For example, the centro-parietal positivity (CPP) presents one of the most potentially valuable signals to a neurally-informed model. It has been shown to represent evidence accumulation (e.g., O'Connell, Dockree, et al., 2012; Twomey et al., 2015), and to exhibit strong correspondence with cumulative evidence timecourses simulated directly from sequential sampling models (e.g., Kelly et al., 2021; Corbett et al., in press). The CPP presents an important method for validating decision models alongside behavioural data. For instance, in Kelly and colleagues (2021) two model variants produced highly similar fits, but only one of them was able to successfully recapitulate the CPP. This was used as evidence in support of their neurally-informed model. Furthermore, the observation of consistent CPP bounds at response across older and younger adults was used as a means of constraining the decision bound in a model by McGovern and colleagues (2018).

As such, the CPP presents a vital tool for neurally-informed modelling. However, while the CPP builds to a peak at response, its peak amplitude varies under different conditions, with lower amplitudes for faster reaction times and more difficult trials (Steinemann et al., 2018). This is likely due to urgency, where an increasing pressure to respond causes a lowering of the decision bound over time and a lower amount of evidence accumulated by response. Without a distinct and predictable accumulation to threshold across decision-making conditions, the CPP cannot be easily rescaled as was done for mu/beta, meaning behaviourally-irrelevant influences on its amplitude and slope cannot readily be controlled for. This makes it difficult to compare these signals across groups, as we cannot be certain that observed differences in CPP amplitude are due to differences in the underlying decision-making process being studied or behaviourally-irrelevant effects on signal amplitude. This may have contributed to the discordance between model-simulated decision variables and the CPP in Chapter Two. Here, model-simulated decision variables in older adults predicted far higher cumulative evidence levels at response than was indicated by the CPP (see section 2.3.6.1: Figure 2.18), making it difficult to assess whether the model simulation or CPP better represents differences in decision-making across

groups. Without a reliable means of rescaling these signals we cannot be certain that the lower CPP amplitudes reflect a true ageing effect on decision formation, or rather some physiological difference in the older adult sample (e.g., reduced skull conductivity, Antonakakis et al., 2020), which affects the presentation of the CPP at the scalp, but not the core underlying decision-making process.

Significant between-group differences in neural data across older and younger adults are observed in a wide variety of tasks and brain signals. One of the most well-studied examples of these between-group differences in the ageing brain is the ageing-frontalisation effect, where both EEG and functional magnetic resonance imaging (fMRI) research has indicated an increased level of activity in frontal areas in the ageing brain (Davis et al., 2008; O'Connell, Balsters, et al., 2012; Van de Vondervoort & Hamlin, 2016). Here the P3, which has been shown to be analogous to the CPP (Twomey et al., 2015), shows significant frontalisation with ageing, potentially representing compensatory recruitment of frontal areas (Kuruvilla-Mathew et al., 2022). However, this frontalisation may be driven by potentially task-irrelevant reasons, such as the influence of overlapping signals or aforementioned differences in skull thickness (Antonakakis et al., 2020; Kelly & O'Connell, 2013; O'Connell, Balsters, et al., 2012)

This frontalisation effect bears important influence on attempts to incorporate the CPP into neurally-informed models of ageing. With increased frontalisation, we do not know where it is best to measure signals such as the CPP on the scalp in older adults, as it remains unclear as to whether this increased frontal presentation reflects a true shift of the evidence-dependent CPP or an ancillary process which results in increased frontal amplitudes but is unrelated to the decision-making process. As such, we cannot be sure of the extent to which these group differences in signal amplitude or build-up represent any true changes to the underlying signal, reflecting changes to the process itself, or just measurement error due to the movement of its location. Importantly, studies of ERP-frontalisation in EEG typically use tasks that elicit sensory-evoked potentials (Fabiani et al., 1998; van Dinteren et al., 2014b) that may have different amplitudes in older adults (Kolev et al., 2002; O'Connell, Balsters, et al., 2012; Anderer et al., 1996, 1998; Gajewski et al., 2018; Bertoli et al., 2005; Golob et

al., 2001; Kuruvilla-Mathew et al., 2022; Polich, 1997a). Furthermore, other neighbouring signals such as the contingent negative variation (CNV), a signal over frontocentral electrodes which peaks *negatively* at response (Devine et al., 2019, see Chapter One: 1.1.3), may also overlap with and potentially alter the overall topography and amplitudes of the CPP without the signal itself shifting position (Kelly & O'Connell, 2013). Collectively, these factors could lead to faulty interpretation of the decision process differences between older and younger adults.

This study aimed to examine these issues by analysing pre-existing data from three tasks; a classic auditory oddball typically used to elicit the P3b, a task designed not to elicit any sensory-evoked potentials to remove the potential influence of these overlapping signals, and a random dot-motion task with varying difficulty (McGovern et al., 2018). In addition, this chapter sought to establish the degree to which the CPP can be better isolated from surrounding activity through a novel application of independent-component analysis (ICA), leveraging the known functional characteristics of the CPP in the selection of spatially-independent components. Through the combination of tasks that exclude sensory-evoked potentials from the ERP and the application of ICA to isolate CPP activation profiles, the study sought to determine whether the frontalisation effects reported in the literature are due to a topographical shift in the CPP itself, or age-related changes to other signals operating on a similar timecourse. Ultimately, this serves to provide a more reliable means of evaluating the CPP across groups, improving its utility in neurally-informed modelling.

4.1.2: Brain Activity Frontalisation in Older Adults

A consistent finding in studies of neural data in older adults is a posterior to anterior shift with ageing (PASA; Davis et al., 2008). This has been shown extensively in the fMRI literature, where increased prefrontal cortical (PFC) activity is consistently observed as we age (Davis et al., 2008; Grady et al., 1994). The onset of Alzheimer's Disease has been associated with an increase in this frontalisation effect, highlighting it as a potentially clinically relevant feature of the ageing brain (P. McCarthy et al., 2014). Several competing hypotheses have

been proposed to explain this, with the most dominant of these being that frontalisation represents a compensatory mechanism. The PFC is associated with executive functions such as top-down attentional control and working memory (van Dinteren et al., 2018). As such, increased activation in these areas may be reflective of the recruitment of these resources to compensate for the effects of ageing on early-stage brain areas and their cognitive processes. This is posited by the Compensation-Related Utilisation of Neural Circuits Hypothesis (CRUNCH, Reuter-Lorenz & Cappell, 2008), which suggests that heightened activity in frontal regions may reflect a top-down compensation for increased noise in the early perceptual areas or reduced hippocampal memory capacity. Similarly, the Hemispheric Asymmetry Reduction in Older Adults theory (HAROLD, Cabeza, 2002) argues that the ageing brain becomes less hemispherically lateralised to counteract cognitive decline. This results in bilateral recruitment of the PFC and thus increased overall activation in those areas (Grady, 2012). These theories have been supported by fMRI data which show increased PFC activity irrespective of task difficulty or subjective confidence, with this frontalisation associated with a decrease in activation in occipital areas and better task performance (Davis et al., 2008).

While this data presents a compelling case for a frontalisation of brain activity driven by increased recruitment of the PFC, work by Morcom and Henson (2018) has called the compensation hypothesis into question. Here a model-based decoder was used to assess the predictive power of models that incorporated activity in the PFC, the posterior visual cortex (PVC) or a combination of the two across two working memory tasks. While frontalisation was observed as increased activity over the PFC in older adults, incorporating PFC voxel activity into their predictive model did not improve its predictive power. The authors suggest that the increased activity was therefore not performing a compensatory role on the task. Instead, they argue that increased activation may reflect impaired functional efficiency in the area, where more blood oxygen-level-dependent (BOLD) activity is needed to achieve the same level of performance.

Electroencephalography (EEG) studies of the PASA phenomenon have also been performed. The majority of these studies have focused on the P3, a signal

analogous to the CPP. The P3 is shown to exhibit an increased latency of peak amplitude with age, with peak amplitude reaching a maximum in adulthood before a steady decline in later life (O'Connell, Balsters, et al., 2012; Polich, 1997b; van Dinteren et al., 2014b). It has also been established that the P3 undergoes an anterior shift in later life, with increased positive amplitudes over frontal electrodes for older adults, alongside decreased amplitude over parietal electrodes (e.g., Kuruvilla-Mathew et al., 2022; O'Connell, Balsters, et al., 2012), with the magnitude of this frontalisation linked to poorer executive functioning (Fabiani et al., 1998; West et al., 2010).

The P3 is often subdivided into two subcomponents with distinct functional characteristics; the P3a, with a maximal presentation over frontocentral electrodes; and the P3b with a maximal amplitude over parietal areas (Polich, 2007). These components have been shown to represent two distinct processes, with the frontal P3a component specifically elicited by unexpected or novel stimuli, and the parietal P3b being elicited by any goal-relevant stimulus (Richardson et al., 2011). A study by O'Connell and colleagues (2012) linked EEG and fMRI data on a visual oddball task to investigate frontalisation of both the P3a and P3b with age. Here, they used target, novel and distractor stimuli to elicit both the P3a and P3b components. No behavioural differences between the groups were observed, however, both distractor P3a and target P3b amplitudes showed smaller amplitudes over parietal electrodes and higher amplitudes over frontal electrodes. fMRI data indicated that the frontalisation of the P3a in older adults was driven by increased activity in the left inferior frontal cortex and right cingulate cortex associated with response inhibition and cognitive control (M. T. Diaz et al., 2011; Swick et al., 2008). Conversely, the P3b frontalisation was driven by increases in the right dorsolateral prefrontal cortex, associated with working memory and response selection (Huettel & McCarthy, 2004). For both components, older adults showed increased activation of the right hippocampus, which is associated with memory and difference detection (Yamaguchi et al., 2004). This study illustrates frontalisation of both components, indicating a complex mix of brain regions involved in the frontalisation of each, while directly linking EEG frontalisation to fMRI data. Notably, frontalisation of the P3b component was associated with increased recruitment of frontal attentional

control regions, suggesting that older adults may increase their top-down attentional regulation to achieve parity in performance.

Frontalisation of the P3b as a compensatory mechanism was also supported in a recent study by Kuruvilla-Matthew and colleagues (2022). Here older adults showed increasing frontalisation of the P3b on more auditory number recognition tasks, with increased frontalisation associated with poorer behavioural performance. The authors suggest that frontalisation of the P3b is indicative of increased compensatory recruitment of frontal areas, which is more pronounced when task demands are greater. Similarly, van Dinteren and colleagues (2014a) traced age-related changes to the amplitude of the P3b following an infrequent target on an auditory oddball task, measured at either parietal electrodes or frontal electrodes. As there were no novel stimuli, they could only investigate frontalisation of the P3b in response to target stimuli. When ERPs were averaged by age group across 1572 participants, frontal-P3b amplitude was seen to increase with age. The authors suggest that increased activity in prefrontal brain areas reflects the PASA, as frontal regions are recruited to compensate for impairments in early stimulus processing.

While the studies outlined above suggest compensatory frontalisation of the P3b with age, other work has shown none. For example, West and colleagues (2010), assessed age-related frontalisation on both the P3a and P3b components using a combined oddball and novelty detection task. Using a Partial-Least Squares (PLS) data reduction analysis, they revealed two latent ERP variables which were compared for each group across conditions. The first reflected the N2, P2 and parietal P3b and showed significant differences between oddball and standard stimuli for both groups. The second, a frontally distributed component which contained the novelty P3a distinguished the novel stimuli only in younger adults but was observed for both novel and oddball stimuli in older adults. This increased novelty-P3a in response to both novel and oddball stimuli was more strongly observed in older adults with poorer performance on tests of executive function. Furthermore, when elicited, this P3a amplitude was associated with an attenuation of P3b amplitude. The authors suggest that this reflects older adults' inability to adapt to the oddball stimuli, which may be processed similarly to novel

stimuli, especially in poorer-performing older adults. Comparatively, the target-detection P3b variable showed reduced amplitudes in response to target stimuli over parietal sites in older adults, but no significant frontalisation, suggesting that ageing may impact each P3 component uniquely. It also highlights a potential issue with P3-frontalisation studies using oddball tasks alone. If older adults process oddball stimuli similarly to novel stimuli, it is difficult to assess if there is any true frontalisation of the P3b, or whether more positive amplitudes over frontal electrodes in older adults is simply just an expression of a P3a component that younger adults do not show.

Recent work has established that the P3b is in fact analogous to the CPP, with both components exhibiting the same polarity, topography, response-alignment and evidence accumulation dynamics (Twomey et al., 2015, see Introduction 1.1.3). Therefore, assessing whether the P3b undergoes a progressive frontalisation with age has direct implications for efforts to use CPP measurements to infer ageing effects on decision-making processes, as it is unclear whether activity over frontal electrodes may better represent the evidence-dependent CPP in older adults. Collectively, however, the evidence remains unclear as to whether this P3b/CPP frontalisation is compensatory, or if the observed anterior shift is better thought of as a distinct P3a component in older adults that does not arise as strongly in the younger (West et al., 2010). This is important, as for appropriate between-group comparison of the P3b/CPP we must be sure we are measuring the signal at the appropriate site for each group. This is of special importance to neurally-informed modelling, as it is the decision-formation function, not novelty detection, that is of interest to the field. If the brain areas involved in evidence accumulation shift, yet we choose the same electrode sites across groups, we may not be accurately measuring a marker of evidence accumulation. This may lead to erroneous conclusions on differences in P3b/CPP features across groups. However, the new understanding of the P3b as analogous to the CPP allows us to overcome the issue of two functionally distinct P3 components that may be affected differently by age. By identifying the link between the P3b and CPP as a marker of evidence accumulation using tasks which focus solely on this feature (e.g., random-dot motion, Chapter Three: 3.1.4), we can remove the potential influence of a novelty detection component.

While this is a speculative assertion, two-forced-choice alternative tasks with varying difficulty levels may enable us to more concretely assess the extent to which the P3b/CPP as a marker of decision formation shows any frontalisation when other potential P3 components are controlled for.

The studies above indicate that when measured across different electrode sites, P3b activity may represent fundamentally different cognitive processes which are affected in distinct ways by ageing (van Dinteren et al., 2014a; West et al., 2010). However, comparing amplitude measurements from the same electrode site across individuals and groups is common practice in the EEG literature (McGovern et al., 2018), and it is evident that significant care must be taken to ensure that the appropriate location is chosen for each age group to ensure that the cognitive process of interest is being measured. A method capable of extracting a spatially independent neural signal based on the functional properties of the cognitive process of interest could resolve this issue. Fortunately, the link between the P3b and CPP enables this, as we can utilise the functional differences in the P3b/CPP under different evidence strengths to extract spatially independent representations of evidence accumulation. This would remove any potential ambiguity as to the function of the component being examined, allowing us to focus on the P3b/CPP. This process will be further explained in the coming sections.

4.1.3: Spatial Overlap in EEG Signals

The studies above indicate a mixed picture of the frontalisation effect in older adults' EEG data. While this may be due to true shifts in the signal or frontal recruitment, it is important to acknowledge that they may not be studying the P3 in isolation. EEG is not recorded directly from the brain, instead using electrode arrays placed on the scalp. This means that the recorded data is subject to the effects of volume conduction (van den Broek et al., 1998). As electrical activity travels through the brain and skull, it conducts through biological tissue with different conductivities, becoming a mix of activity from a number of different brain areas (van den Broek et al., 1998). Scalp-recorded ERPs thus represent an overlapping picture of several distinct neural sources (Vidal et al., 2015),

meaning what we measure at the scalp is a waveform-composite of a range of different processes. This is of special importance to comparisons of EEG signals across groups. For example, older adults show changes in skull thickness, which are unrelated to cognitive processing (Antonakakis et al., 2020). As a result, the location or amplitudes of observed signals in the older adult brain may be due to changes in volume conductance as a result of physiological changes which are not directly relatable to decision-making itself. This makes it more challenging to utilise ERP data in its initial form as constraints for neurally-informed modelling, because observed differences across groups or conditions may be driven by task or group effects on an unknown or irrelevant signal.

One such potential issue with studies of the P3b/CPP is potential overlap from visual or auditory evoked potentials. These appear when a stimulus is presented suddenly on screen, eliciting a sensory-evoked signal which may overlap with signals of interest (O'Connell, Dockree, et al., 2012). For example, the visual N2, elicited by the sudden appearance of a visual stimulus, presents in occipital-parietal areas and shows a more negative amplitude with age (Kolev et al., 2002; O'Connell, Balsters, et al., 2012). If overlapping with the P3b/CPP, this may result in a more negative P3b/CPP amplitude in older adults when measured over parietal areas. In auditory tasks, such as the auditory-oddball tasks typically used in studies of P3b frontalisation (e.g., O'Connell, Balsters, et al., 2012; van Dinteren et al., 2014a), an auditory N1 is elicited by the sudden onset of a stimulus and is maximal over frontocentral electrodes approximately 100ms after the stimulus (ElShafei et al., 2022). Research in the field is divided as to whether it shows a larger (Anderer et al., 1996, 1998; Gajewski et al., 2018) or smaller amplitude (Bertoli et al., 2005; Golob et al., 2001; Kuruvilla-Mathew et al., 2022; Polich, 1997a) in older adults, however, in a general sample, it is the latency of this signal which correlates with RT, not its amplitude (Ribeiro et al., 2016), suggesting that group-differences in amplitude that, while not being relevant to behaviour, may potentially overlap with the behaviourally relevant P3b/CPP. In either case, the sudden onset of a stimulus can lead to sensory-evoked potentials that may impact measurements of the P3b/CPP. For example, a less negative overlapping sensory-evoked potential over frontal electrodes may artificially inflate the amplitudes of a frontally-measured P3b/CPP in older adults compared

to younger adults. Without tasks designed to remove the influence of sensory-evoked potentials, the effects of frontalisation on the P3b/CPP in the absence of an auditory/visual-evoked potential are not clear. The O'Connell lab has developed tasks in which the target appears gradually over time, removing these sudden-onset sensory signals and providing a clearer picture of the CPP's evolution during decision formation (O'Connell, Dockree, et al., 2012). This is important in the comparison of older and younger-CPP profiles, where group differences in the amplitudes of these sensory-evoked potentials may impact signals of interest such as the P3b/CPP.

To further combat signal overlap, surface Laplacian filters, such as the current source density (CSD) transform, are often applied to electrophysiological data. These convert recorded voltages, to estimates of current flow (Kayser & Tenke, 2015). Positive measurements indicate the movement of current from the brain toward the scalp, while negative values indicate currents flowing from the scalp towards the brain. A key benefit of these transformations is that they are successful in teasing apart spatially overlapping ERP components, providing a more distinct topography and an increased spatial resolution (Kayser & Tenke, 2015). For example, the contingent negative variation (CNV, see Chapters Two & Three, Devine 2019), is often measured at electrodes close in location to the P3b/CPP. This CNV presents many of the same features as the CPP but differs in some fundamental ways. It builds during decision formation and is effector-independent, as is the CPP, but is evidence-independent and shows steeper build-up rates under increased speed pressure (Devine 2019). Importantly, however, it is a *negative* going signal, which peaks at a negative amplitude at response, in contrast to the positive going CPP. In a study by Kelly and O'Connell (2013), an overlapping CNV was shown to "pull down", or make unduly negative, the amplitude of the CPP. In this study, the application of CSD was shown to significantly sharpen topographies, isolating CPP and removing the effects of the CNV on CPP amplitude at response, highlighting a key benefit of these spatial filters. However, the studies of P3 frontalisation to date are typically performed on standard ERP data (e.g., O'Connell, Balsters, et al., 2012; van Dinteren et al., 2014). This presents a fundamental issue, as we cannot be confident that the observed frontalisation with age is due to a true movement of the signal, or

perhaps a larger CNV in younger adults as seen in Chapter 2, which pulls down frontal amplitudes to a greater extent, thus causing older adults to exhibit a relative positivity over frontal sites.

Combined with the influence of sensory-evoked potentials and a potentially overlapping P3a, our understanding of any true between-group differences in P3b/CPP amplitude may be clouded. These issues are less pertinent for within-subject designs, where differences in skull thickness or topographical shifts cannot play a role. However, they are crucial to address for between-group comparisons. Without attempting to disentangle the CPP from overlapping signals or sensory-evoked potentials, erroneous conclusions as to the effects of, for example, ageing on the underlying process the signal represents could be drawn. Resultantly, the signals may not be reliable as neural constraints for computational modelling across groups. In order to establish the true effects of frontalisation on the CPP in older adults, attempts must be made to minimise this signal overlap through novel task design and the application of spatial filters such as CSD. This present work aimed to leverage CSD and a task in which evidence emerged gradually over time, to test for frontalisation where no potential overlapping sensory-evoked potential was present, and the impact of overlapping signals such as the P3a or CNV is minimised through spatial filtering.

4.1.4: Tackling Spatial Mixing via Independent Components Analysis

While the application of CSD is beneficial, it is only a filter and does not fully isolate overlapping signals. A further method capable of extracting and detecting spatially independent components would lend a significant advantage to our study of the brain. Independent component analysis (ICA) is a method which can achieve this (Onton et al., 2006). ICA assumes that activity across the scalp is a composite mix of activity from a range of independent components. Over time, activity coming from different sources should be temporally independent (Onton et al., 2006). ICA decomposes EEG signals into these temporally and spatially independent components. It is traditionally used in the data preprocessing stage, where artifactual components such as blinks, muscle activation or background electrical activity are identified and removed from the data (Marriott Haresign et

al., 2021). To this end, ICA has been successful. However, the removal of artifactual components such as blinks may not always be the best practice within the study of perceptual decision-making. When endeavouring to link the brain to behaviour, there is an assumption that participants have viewed the task in its entirety. The removal of artifactual blinks through ICA removes the blink component but leaves the remaining EEG data intact. Here, if participants blink and are not viewing the stimulus the trial is still counted. For neurally-informed modelling, where the model assumes continuous accumulation at every timestep, it is more fitting to exclude trials in which a blink or other artefact appears. This is common practice in decision-making research and is the approach adopted in this body of work.

Beyond its use in data pre-processing, ICA enables the isolation of discrete ERP components. ICA can be applied to ERPs in order to extract and identify components that reflect signals such as the P3 (Makeig et al., 1999). These activations are statistically independent from other components, minimising the influence of any potential overlapping signals in raw ERP data. Back-projecting these components by multiplying the data by the inverse of their mixing matrix provides an estimate of how the activation would appear on the scalp, independently of other components. This approach allows us to overcome the spatial mixing of standard ERPs. This is beneficial to the study of frontalisation in the ageing CPP, as, if the P3b really does undergo a frontalisation, then this should be observable in a more frontal activation of the associated independent components. Extracting a P3b/CPP component via ICA would allow us to more concretely assess if this frontalisation represents a true topographical shift, and thus whether it is reasonable to measure the CPP at the same site across groups. Beyond this, ICA is of unique benefit to neurally-informed modelling. If reliable components reflecting the P3b/CPP build-up dynamics in line with model predictions are identifiable, this procedure can be used to isolate components which provide more precise estimates of evidence accumulation in perceptual decision-making. These components would then serve as principled data-driven neural representations of model parameters, providing both useful information and more precise neural constraints.

To date, some studies have attempted to examine the P3b through ICA. For example, Makeig and colleagues (1999) applied ICA to 31-channel EEG on a visual spatial attention task. Using PCA to reduce the 31 extracted ICs, they identified a component which appeared visually similar to P3b in terms of its peak stimulus-locked latency (circa 300ms following stimulus onset) and parietal focus. However, this study was performed on a low-density EEG array whose dimensionality was further reduced by PCA, limiting the possible number of components. This is a concern as such reduced dimensionality risks the possibility that distinct cognitive processes may be lumped together. Work from Debener and colleagues (2005) used ICA to identify distinct P3a and P3b components on an auditory oddball task in a general adult population. Here they indicated that an overlapping P3b component may artificially increase observed P3a ERP amplitudes when participants were instructed to attend to rare stimuli over oddball stimuli, indicating the potential influence of overlapping P3 components on each other (West et. al, 2010).

A recent study by van Dinteren and colleagues (2018) applied this ICA approach to the study of P3 frontalisation in ageing. Using a hierarchical group-level ICA they manually selected 4 ICs which bore a visual resemblance to the P3b in their topography and stimulus-locked activation time course in older and younger adults across an auditory oddball task, a reaction time test and a test of continuous attention. They then used source localisation to show evidence of a frontalisation effect with age in only one of these components. The group differences were attributed to differences in the precentral and parahippocampal gyri, with increased amplitudes of this component associated with faster RTs, but only on the continuous attention test. No other relationships between this component's features and behaviour on any other task were found. The authors suggest this to be evidence for a frontalisation effect in ageing, reflecting compensatory recruitment of frontal areas. However, the fact that two of the identified P3b components did not exhibit frontalisation suggests that not all P3b-like activity may be subject to an anterior shift with age. Furthermore, as the component was only linked to behaviour on a test of continuous performance, it is unclear whether such frontalisation would be observed when examining an evidence-dependent decision formation P3b/CPP. This study was also

performed with a low-density electrode array (26-channel), further limiting the potential reliability of the decomposition.

One significant limiting factor of traditional means for selecting ERP-like ICA components is that they often rely on manual selection by the experimenter. Here, ICA activation profiles are compared to ERPs for visual similarity in windows of interest, presenting a potential source for bias in selection. A recent study by (Porcaro et al., 2019) used functional source separation (FSS) on data from an episodic memory task to identify a parietal P3b amplitude as a marker of cognitive decline in ageing. FSS is similar in principle to ICA but rather than purely extracting components based on their temporal independence, it attributes neural activity into components in order of its similarity to some predefined functional marker. In this study's case, it extracted components with maximal amplitude in the 320:480ms after stimulus onset time. The extracted component's amplitude was shown to be reduced in older adults, and a strong predictor of performance on cognitive tasks performed outside the EEG booth by older adults. The authors suggest that the application of this FFS approach allowed them to identify a spatially independent P3b component amplitude which acts as a marker of cognitive decline. More recent work has also used machine learning to cluster back-projected ICAs based on their topographies, mean amplitudes and latencies (Jervis et al., 2020). This was applied to data on an auditory oddball task performed by adults with and without confirmed Alzheimer's Disease (AD). The derived components were then able to successfully classify participants with AD, where standard ERP signals could not, indicating the potential clinical utility of such an approach.

Again, however, these EEG analyses relied only on topography and stimulus-locked peak amplitude and latency. Importantly, the P3b/CPP has been shown to be response-aligned, so its amplitude is most appropriately measured immediately before response execution (O'Connell, Dockree, et al., 2012). Traditionally the P3b has been measured at its stimulus-aligned peak (e.g., van Dinteren et al., 2014a/b, Porcaro et al., 2019, Jervis et al., 2020), but these measures may be confounded by between-group or condition differences in RT variability. Furthermore, the discovery that the P3b/CPP in fact traces an

evidence accumulation process points to a far more specific set of functional characteristics that can be used to aid its isolation than analysis of amplitude alone (O'Connell, et al., 2012; Twomey et al., 2015). The response-aligned CPP emerges with stronger build-up rates for stronger evidence and faster RTs (Twomey et al., 2015), giving two predictable indicators for how the signal should emerge under different conditions. Using this information, we can select and cluster components which reflect these dynamics, allowing us to rely not on visual similarity or peak amplitudes, but on core functional characteristics of the evidence accumulation signal. This presents an exciting opportunity for identifying a spatially independent CPP component, which minimises potential spatial overlap and may more concretely assess the effects of ageing on this signal.

4.1.5: Present Study

The present study aimed to analyse existing data sets (McGovern et al., 2018) to assess the frontalisation of the CPP. Three tasks which elicit distinct electrophysiology were analysed. The first, a simple auditory odd-ball task allowed us to assess the P3 frontalisation effects in line with previous literature. We then assessed if this frontalisation was robust to the use of a spatial filter (CSD) to reduce the impact of overlapping signals. We then aimed to assess the effect of sensory-evoked potentials on frontalisation. The second task used a contrast change detection paradigm where evidence emerges gradually, which provides a P3b/CPP unimpeded by potentially overlapping sensory-evoked potentials. This allowed us to investigate frontalisation in the absence of any sudden-onset sensory-evoked potentials. Finally, to select a spatially-independent CPP through ICA, a random dot-motion task with two coherence levels was used. This was chosen as it provides unique predictions as to the emergence of the CPP in the two different coherence conditions. High-coherence trials should exhibit steeper build-ups, as should trials with faster RTs. This grants two reliable metrics from which to automatically select independent components which reflect the evidence accumulation process. These components would be ranked and averaged across groups, investigating differences in component

activation topographies between older and younger adults to assess if frontalisation persisted after this spatial unmixing.

Collectively, these analyses had the goal of identifying if classic studies of frontalisation were robust to spatial unmixing techniques, with the ultimate aim of providing more precise neural metrics of evidence accumulation.

4.2: Methods

4.2.1: Participants and Procedure

The study reanalysed data presented in McGovern and colleagues' paper (2018). Seventy-eight participants took part of which 40 were classified as older (Mean Age = 74.5, *SD* = 5.2) and 38 as younger (Mean Age = 21.6, *SD* = 5.2). Older participants were recruited through the Trinity College Institute of Neuroscience volunteer panel and reimbursed for their travel expenses and time at a rate of €10/hour. Younger participants were recruited through the School of Psychology and word of mouth. Students were granted one research credit for each hour completed. Participants were required to report no history of addiction, brain injury or neurological illness, and no family history of the same. Participants were matched on key demographic features of gender, handedness, and years of education. All participants were required to score higher than 26 on the Mini-Mental State Examination, performed upon their arrival. One older participant was excluded from the analysis on the auditory oddball task due to a high proportion of EEG artefacts (>= 66% of trials; 38 younger adults, 39 older adults). Twelve participants were excluded from the analysis on the random-dot motion task due to a high proportion of EEG artefacts (>= 66% of trials) and technical error giving a sample of 35 younger adults and 31 older adults. A post-hoc power analysis was performed using G*Power (Faul et al., 2007). This indicated that given a sample size of 66 across two groups and three electrode locations, with a significance criterion of 0.05 and a power of 80%, the study would have sufficient power to detect effects of small to medium effect sizes (f = 0.16). Participants completed an auditory oddball, contrast change detection, and random-dot motion in a pseudo-randomised order across participants, while 128-channel EEG data was recorded. Tasks were completed in a sound-attenuated room, with stimuli presented on a 51cm CRT monitor positioned 55cm away from the participant. A refresh rate of 85Hz was used and the resolution of the screen was 1024x768.

4.2.2: Auditory Oddball

Participants were required to indicate their detection of a target tone of 1500 Hz among a series of standard tones (1000Hz). All tones were 60ms in duration with a pseudo-randomised inter-tone interval of between 1300ms and 2200ms. A randomised number of between 3 and 5 standard tones played in sequence prior to a target tone. Participants indicated the presence of a target tone by making a right mouse click. Participants were instructed not to respond to non-target tones, and to maintain fixation throughout. Each block consisted of 54 target tones, with participants completing one block per session. A shortened practice version of the task was performed prior to recording.

4.2.3: Contrast Change Detection

Participants monitored a circular checkerboard pattern flickering at 21.25Hz continuously for a decrease in contrast from 70% to 40%. This decrease occurred gradually across 1600ms, returning to its original 70% contrast over the following 800ms. Participants indicated their detection by pressing the right mouse button with their right index finger. The annulus had an inner radius of 3 degrees, with an outer radius of 8 degrees. The inter-target interval varied pseudorandomly on each trial as either 3, 5 or 8 seconds. Participants were instructed to respond as quickly as possible but only once they were sure, completing 4 blocks composed of 24 targets each.

4.2.4: Random Dot Motion

Participants monitored a circle of 150 randomly moving dots (each dot was 6 x 6 pixels) presented in an aperture of 8 degrees around a central fixation point. Dots moved randomly at a speed of 6 degrees per second for either 3 or 7 seconds. A proportion of the dots then moved either to the left or right at either 30% coherence (low) or 60% coherence (high). The coherent motion was presented for a maximum of 10 seconds, or until 500ms after the participant had responded. There was a variable intertrial interval of between 3 and 7 seconds. In total, 30 targets were presented in each block, with participants completing 6 blocks.

Participants completed a practice block to familiarise themselves with task demands in which they received feedback for accuracy, misses and false alarms.

4.2.5: EEG Acquisition and Pre-processing

As in previous chapters, EEG data were collected using the Biosemi ActiveTwo system at a sample rate of 512Hz. Vertical electrooculogram electrodes placed above and below the left eye recorded eye movements. All data were analysed using a combination of custom MATLAB code and features of the EEGLAB toolbox (Delorme & Makeig, 2004). A low-pass filter of 40Hz was applied to all data, and linear detrending was also applied to remove slow drifts. Noisy channels were identified manually and interpolated. An average reference was applied to the continuous data. Each dataset was epoched into unique time lengths based on stimulus and inter-stimulus durations. Data for the odd-ball task were divided into stimulus-aligned epochs from -500:1300ms centred on the auditory tones. Response-locked epochs were also extracted from -650ms:300ms around response. Data were baseline-corrected by removing the average activity in the -250ms:0ms before stimulus onset. Data for the random dot motion task were stimulus-locked to -750:2500ms to allow for slower reaction times and potential pre-stimulus build-up of the CPP during the period of random dot motion. Data were baseline-corrected using an average of -600ms:-500ms prior to coherent dot motion. Response-locked signals were measured in the -900ms:300ms around response. For all tasks, artefacts were identified from the beginning of each epoch to 150ms following response. Any trial in which any channel had an activity exceeding 100 microvolts was removed. Trials in which the absolute difference in voltage between the upper and lower VEOG channels exceeded a 200-microvolt threshold were labelled as blinks and removed. To investigate the potential frontalisation of the CPP, key candidate electrodes were selected and plotted based on previous literature (O'Connell, Balsters, et al., 2012) and the observed topographies. The electrodes selected were the Frontal electrode (Fz) the Central Electrode (Cz) and the Parietal electrode (Pz), alongside their immediate left and right neighbours. The activity was averaged across the three electrodes at each site and the peak stimulus-locked amplitude and latency, and response locked-amplitudes and slopes were analysed across

each location. Response-locked amplitudes were measured in the -50:50ms centred on the response while stimulus-locked peak amplitude and latency were measured from 200:600ms from stimulus onset for the oddball and random-dot motion task. To account for the slower RTs in the contrast-change detection task, a later search window from 500:1200ms was used. Response locked slope was measured in the -300:-100ms before a response for the oddball task and the contrast change detection task, while it was measured in the -400:-100ms around response for the random dot motion task to account for slower reaction times. A mixed-factorial ANOVA was then performed, examining the effects of Group and Location on each of these signal features. For the auditory oddball task, the data was analysed with CSD and non-CSD transformed data to investigate the effects of this spatial filter on frontalisation. A CSD transform of spline size 4 was then applied to the data for the remaining two tasks.

4.2.6: ICA Decomposition

To perform ICA on the data from the dots task, a minimally cleaned data set was used. VEOG channels were not used to remove blinks, and a more liberal threshold of 250uvs was used to reject artefactual trials. Sixty-four component ICA was then performed on this minimally cleaned data set with artefact, line-noise, eye, or movement components identified automatically using the IClabel toolbox (Pion-Tonachini et al., 2019) and removed from subsequent analysis. To identify components with activity reflective of the CPP, each component was back-projected to the raw stimulus and response-locked data and only components with a mean positive going signal over the central and parietal electrodes at response (-500:-100ms relative to response) and post-stimulus (200:600ms post coherent motion onset) were selected for further analysis. A set of CPP features were then used to rank ICA activation profiles. The slope for each coherence level and each reaction-time bin (fast and slow, based on a median split in each coherence level and averaged across coherence levels) were calculated in the -500:-100ms before response. The components were then ranked based on the difference in their activation build-up rate (slope) for high versus low coherence trials with components with more positive differences ranking more highly. The components were also ranked based on the difference

in their activation build-up rate for fast vs slow reaction times, again with positive differences ranking more highly. These two rankings were then averaged and the three activations with the highest mean ranks were selected for each individual. Each winning activation was back-projected to the raw data and a CPP electrode for each activation was chosen based on the same ICA selection criteria above, using maximal build-up rate difference for high vs low coherence and fast vs slow RT. This search and ranking was performed from a midline set of electrodes based on the back-projected topographies. Activity in the highest ranked electrode for each of the top 3 back-projected ICA's was averaged to reconstruct each participant's CPP waveform.

Topographical differences between age groups were tested using cluster-based permutation statistics. Electrode clusters were identified using Fieldtrip's (Oostenveld et al., 2010) *prepare_neighbour* function to cluster electrodes based on their 3-D positions, selecting electrodes within 50mm of each other. This gave an average cluster size of 5.2 electrodes. These were then compared using a cluster-based permutation independent samples t-test using the Mass Univariate ERP Toolbox in MATLAB (Groppe et al., 2011). Activity for each electrode averaged across the -50:50ms around response for each subject in each group was submitted to a two-tailed, two-samples test with a family-wise alpha level of 0.05. Across 100,000 permutations each subject was randomly assigned to one of the two age groups and a t-test was performed at each electrode site. T-values associated with p-values of less than 0.05 were summed within a cluster, giving each cluster a "mass". The distribution of t-values in the cluster with the most extreme "mass" was used as an approximate distribution of the null hypothesis. T-score distributions from each cluster were then compared to the approximate null hypothesis distribution, with any cluster giving a p-value of less than 0.05 considered significant.

4.3 Results

We first tested for group differences in CPP topographies across the three tasks using cluster-based t-tests and Age-Location ANOVAs, following the same methods used in the literature examining age-related P3 frontalisation (O'Connell et al., 2012).

4.3.1: Auditory Oddball

Behavioural analyses indicated no significant difference in hit rate across groups (Mann–Whitney U = 704.5, p = 0.671), but a significant group difference in RT (Mann–Whitney U = 310, p < 0.001) with slower RTs in older adults (Median = 423.3), than younger adults (Median = 335.69).

4.3.1.1: Topography Comparison

Figure 4.1 shows response-locked topographies of electrode activity in the -50:50ms centred on response, with no CSD transform applied (upper row), and CSD of spline-size 4 applied. In both data sets, group differences in the auditory odd-ball response-locked topographies were observed over frontal electrodes, where older adults showed more positive amplitudes, and over centroparietal electrodes where older adults showed more negative amplitudes compared to the younger (see Figure 4.1, electrode-clusters with a significant group difference marked in black). Cluster-based t-tests were used to compare group differences in topographies with no CSD and CSD applied. The application of CSD reduces the number and spread of between-group differences, revealing three distinct effects: a more negative amplitude over parietal electrodes in older adults, a central negativity and a frontal positivity. This indicates that the application of CSD may be successful in providing a more spatially distinct centroparietal positivity, while the significantly more positive frontal activity in older adults is maintained.

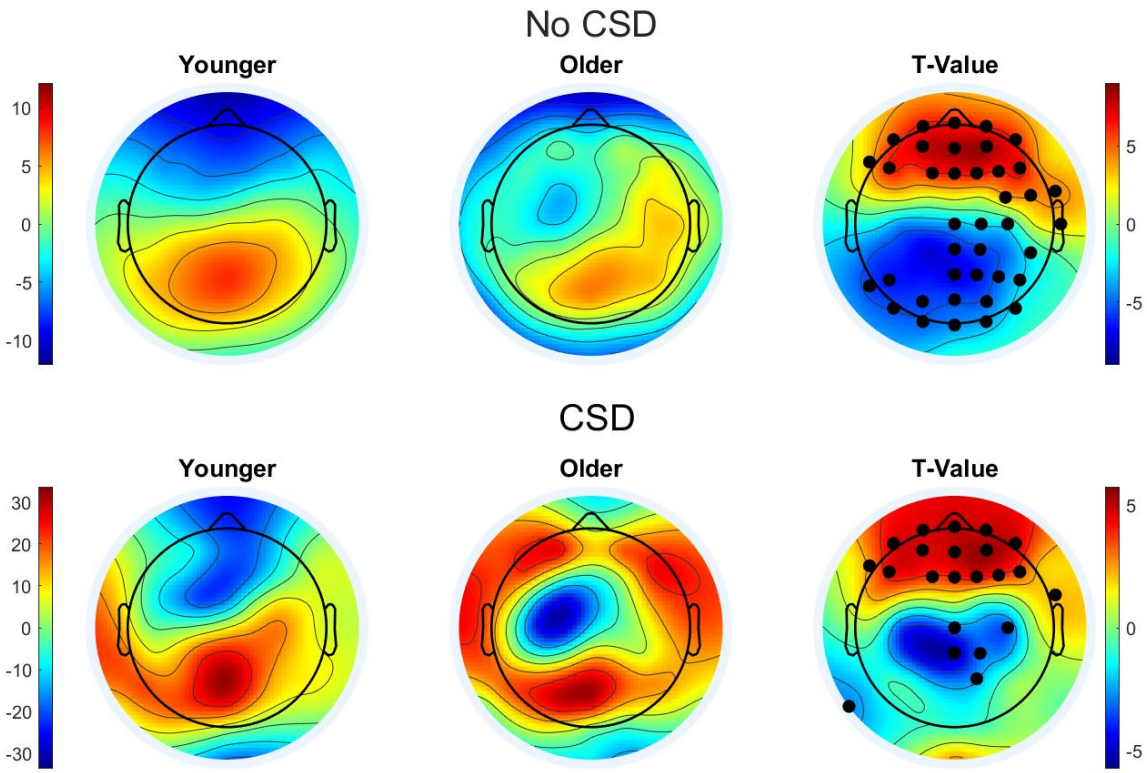

Fig. 4.1: Response-locked topographies from -50:50ms with no CSD and CSD of spline size 4 applied, showing each group and the difference between younger and older on the right.

4.3.1.2: ERP by Location

To replicate previous findings of frontalisation effects in an auditory oddball task, a Location-by-Group, 3x2 factorial ANOVA was performed to investigate the effect of electrode location on non-CSD transformed stimulus-locked peak amplitudes and latencies. The results were closely in keeping with the previous literature with a significant effect of Location on stimulus-locked peak amplitude ($F_{(2, 75)}$ = 62.63, p <0.001), a significant main effect of Age ($F_{(2, 75)}$ = 14.53, p <0.001) and Age-Location Interaction ($F_{(2, 75)}$ = 39.99, p <0.001). Post-hoc analyses confirmed a significantly more positive amplitude in older adults over frontal electrodes (t = 4.778, p <0.001), but a less positive amplitude in older adults over central electrodes (t = -3.808, p<0.001) and parietal electrodes (t = -7.566, p = 0.036). Analysis of the latency of these peaks showed a significant effect of Location ($F_{(2, 75)}$ = 9.14, p <0.001), Age ($F_{(2, 75)}$ = 10.21, p <0.001) and Age-Location interaction ($F_{(2, 75)}$ = 3.22, p = 0.043). Frontal latencies were no

different across age groups at frontal electrodes (p = 1) and central electrodes (p = 0.2) but slower for older adults over parietal electrodes (t = -3.48, p = 0.007).

Similar results were seen for response-locked amplitudes. There was a significant effect of Location on response-locked amplitude ($F_{(2, 75)}$ = 123.05, p <0.001), a significant main effect of Age ($F_{(2, 75)}$ = 5.81, p = 0.018) and Age-Location Interaction ($F_{(2, 75)}$ = 51.77, p <0.001). Post-hoc analyses confirmed a significantly higher amplitude in older adults over frontal electrodes (t = 7.72, p <0.001), but a lower amplitude in older adults over central electrodes (t = -3.83, p<0.001) and parietal electrodes (t = -7.21, p < 0.001). There was no significant effect of Age on the response-locked slope (p = 0.618), but a significant effect of Location ($F_{(2,75)}$ = 22.63, p < 0.001) and Location-Age interaction ($F_{(2,75)}$ = 43.08, p < 0.001). Post-hoc analyses indicated a significantly less positive slope in younger adults over frontal electrodes (t = -7.49, p < 0.001), but no significant differences over central (p = 0.15) or parietal electrodes (p = 0.193, see Figure 4.2).

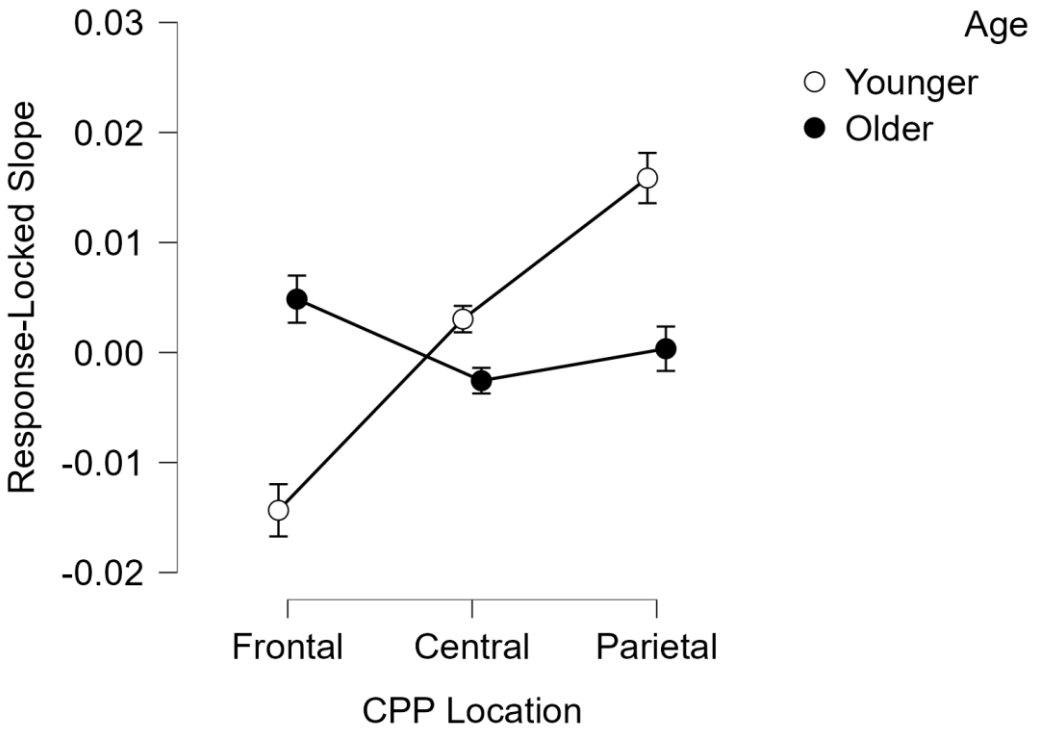

Fig. 4.2: Response-locked slope by electrode location for non-CSD transformed data. Separate lines indicate different age groups.

Collectively these results support previous studies on non-CSD transformed data, which indicate higher amplitudes over frontal electrodes for older adults, and lower amplitudes over parietal areas. Only parietal areas were associated with slower latencies in older subjects, which may be associated with their slower RTs.

To examine whether these topographical differences were altered by unmixing overlapping signals through spatial filtering, the same analysis was performed on CSD-transformed ERP data. Again, a broadly similar pattern of results was observed with a significant effect of Location on response-locked amplitude ($F_{(2, 75)}$ = 69.041, p < 0.001) and a significant Location by Age interaction ($F_{(2, 75)}$ = 33.951, p < 0.001), although this time there was no significant main effect of Age (p = 0.398). Post-hoc analyses were consistent with the non-CSD transformed data and confirmed a significantly higher amplitude in older adults over frontal electrodes (t = -6.594, p < 0.001), but a lower amplitude in older adults over central electrodes (t = 5.369, p < 0.001) and parietal electrodes (t = 2.463, p = 0.036, see Figure 4.3).

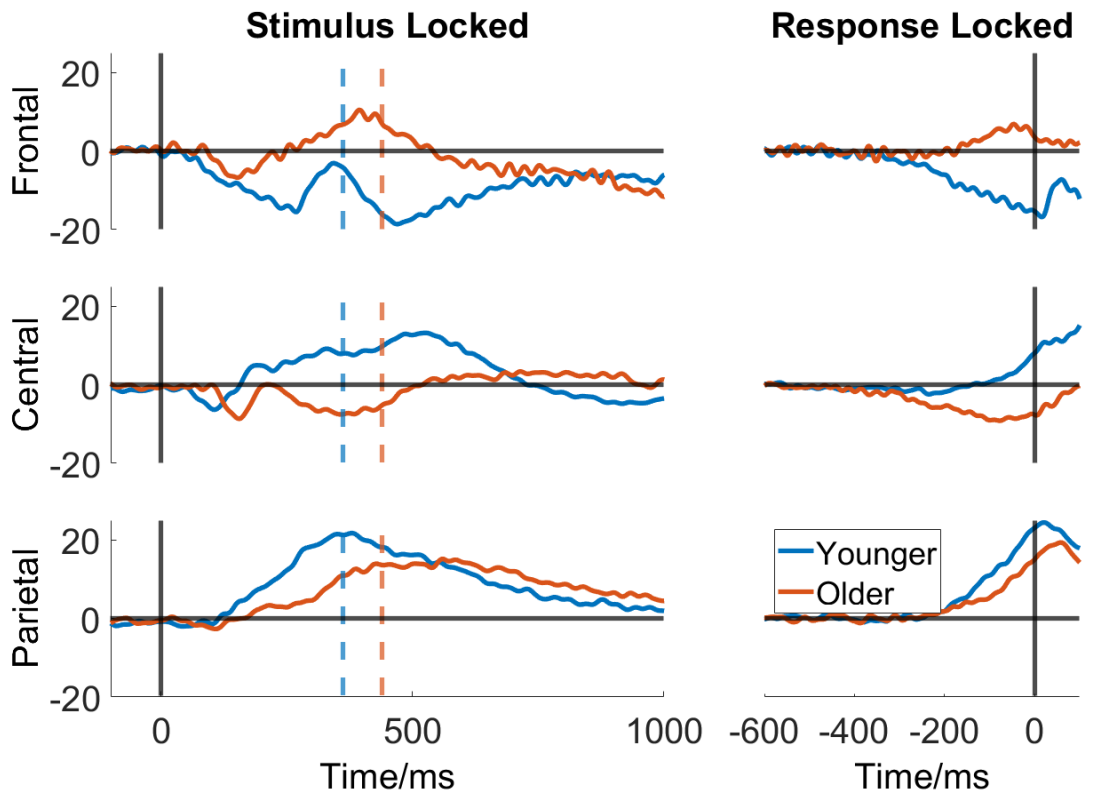

Fig. 4.3: Stimulus and response locked CPP amplitudes for the auditory oddball task, using parietal (Pz), central (Cz) and frontal (Fz) electrodes and their immediate neighbours. The negativity observed in younger adults over frontal electrodes appears more central in older adults.

There was no significant effect of Age on the response-locked slope (p = 0.54), but a significant effect of Location ($F_{(2,75)}$ = 24.73, p < 0.001) and Location by Age interaction ($F_{(2,75)}$ = 24.5, p < 0.001). Post-hoc analyses indicated a significantly less positive slope in younger adults over frontal electrodes (t = -6.4, p < 0.001), but a more positive slope over central electrodes (t = 3.429, p < 0.001), with no significant differences over parietal electrodes (p = 0.193, see Figure 4.4).

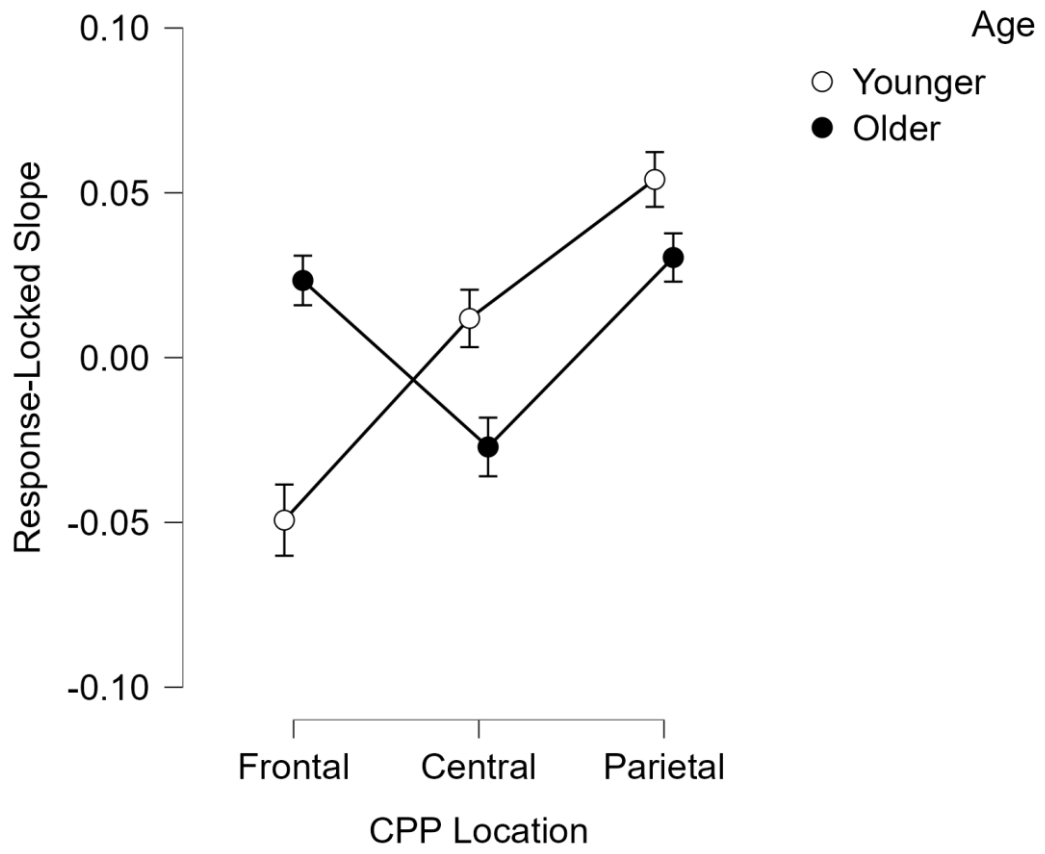

Fig. 4.4: Response-locked slope by electrode location for CSD transformed data. Separate lines indicate different age groups. CSD transform shows a frontal and parietal positivity in older adults, with a central negativity.

As the P3b/CPP has been indicated to be a response-aligned signal, and highly similar frontalisation effects are observed for response and stimulus-locked measurements, only response-locked statistics are reported in the main body of the paper for the remaining studies. Stimulus-locked results are included in the Appendices. These analyses indicated that the application of a Laplacian spatial filter may reduce the main effect of Age on CPP amplitudes, by unmixing overlapping signals (see Table 4.1), however, a frontal positivity with age is still observed. For the remainder of the studies, results from the CSD-transformed data are reported. The effects from non-CSD transformed data on response-aligned features are also presented in the Appendices.

Table 4.1: Comparison of F-values and effect sizes (η^2) of ANOVAs of peak stimulus-locked amplitude and latency, and response-locked amplitude analysing the effect of electrode location and age. The effect size difference shows a reduction in the effect size of the effect of location on all signals with the application of a CSD transform. * = p <0.05, ** = p < 0.01, *** = p < 0.001.

	No CSD		CSD		
	F	η^2	*F*	η^2	η^2 difference
Peak Amplitude					
Location	62.63***	0.282	24.61***	0.14	-0.142
Age	14.53***	0.033	3.18	0.012	-0.021
Interaction	40***	0.18	25.77***	0.147	-0.033
Latency:					
Location	9.14***	0.068	3.47 *	0.029	-0.039
Age	10.2**	0.042	8.43 **	0.031	-0.011
Interaction	3.22*	0.024	3.61*	0.031	0.007
R-locked Amplitude					
Location	123.05***	0.453	69.04***	0.343	-0.001
Age	5.81*	0.006	0.723	0.001	-0.005
Interaction	51.77***	0.191	33.95***	0.169	-0.022

4.3.2: Contrast Change Detection

Behavioural analysis indicated a significant main effect of Age on hit-rate (Mann-Whitney U = 477.5, p = 0.005), with a significantly higher hit rate in older adults (Median = 87.5%) than younger (Median = 77%). Analysis of the effect of Age on RT showed no significant group difference (p = 0.925). On this task older adults outperformed younger adults, unlike the auditory oddball in which there was no difference in hit rate.

4.3.2.1: Topography Comparison

The same series of analyses were performed on the continuous contrast change detection task with CSD transformation applied to minimise signal overlap. This task involves the detection of a gradual and subtle stimulus feature change that consequently does not elicit any sensory-evoked potential which would otherwise overlap with and obscure decision-related signals. Here again, distinct frontocentral negativities and parietal positivities are observed in both groups (see Figure 4.5.B). To detect any significant difference in the topographical presentation of the CPP in the oddball versus the contrast change task, a cluster-based permutation t-test was performed to compare response-locked topographies of the auditory oddball vs the contrast change detection task for each group. There were no significant differences at any cluster for the older adults (all p > 0.053). No differences were observed in the younger group over any of the electrodes where the CPP is most strongly expressed. There were significant differences at just 5 electrode sites over the right central and anterior frontal scalp in younger adults, where less negative activity was observed in the contrast-change detection task (see Figure 4.5.B: electrode-clusters with a significant group difference marked in red). Such differences may reflect the influence of overlapping auditory-evoked potentials on the oddball task, however, it was not within the scope of this study to thoroughly investigate this. Importantly, neither group showed any inter-task differences over centro-parietal electrodes, as is evidenced by the lack of a significant difference in the area for either group between the two tasks (cluster-based t-tests shown in Figure 4.5.B). This indicates that the frontalisation on the auditory oddball task cannot be purely driven by overlapping sensory-evoked potentials.

Cluster-based t-tests were also used to compare between-group differences in ERP topographies on the contrast-change detection task. Similarly, to the auditory odd-ball task, this analysis showed more positive amplitudes over frontal electrodes for older adults, and more negative over central electrodes compared to the younger (see Figure 4.5.B, electrode-clusters with a significant group difference marked in black)

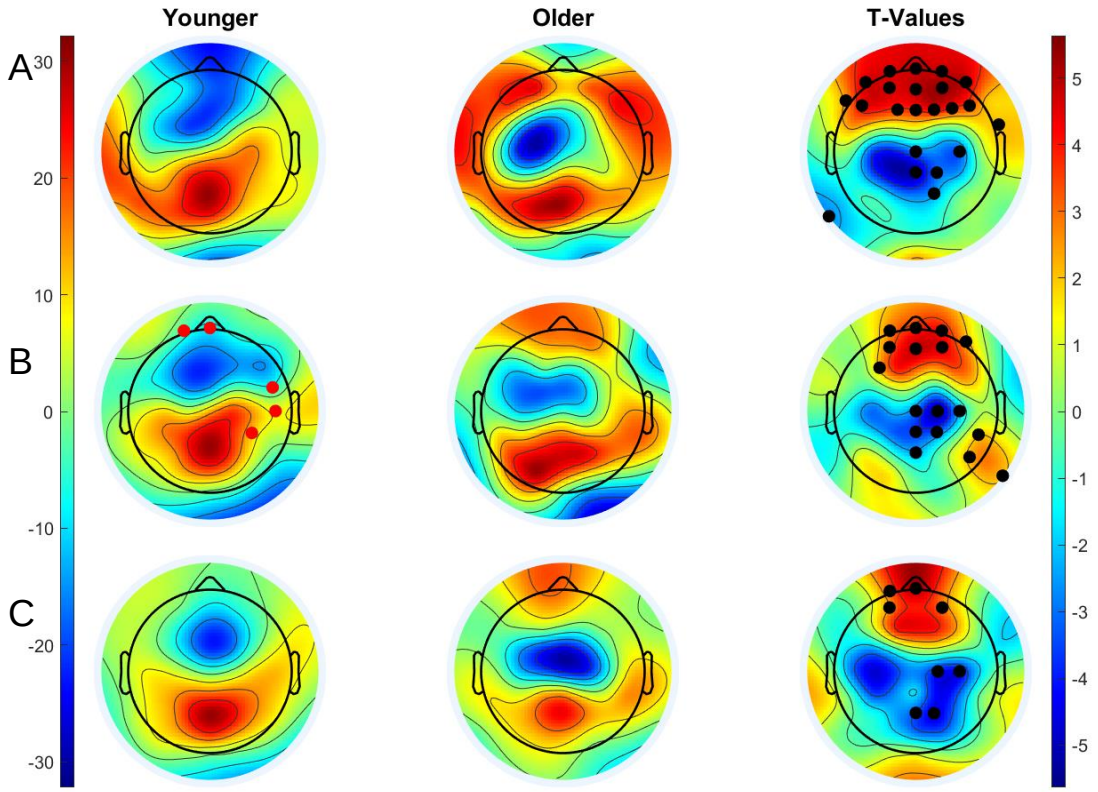

Fig 4.5 Younger and older response locked topographies (-150ms to 0ms), with the rightmost column showing t-values differences for cluster-based t-test between groups for each study (Old - Young). Electrodes with a significant difference ($p < 0.05$) are marked with a black marker. The west colourbar is for younger and older ERP topographies. The east colourbar is for t-values. A) Auditory-oddball task. A significant frontalisation is seen in older adults with higher amplitudes over frontal electrodes. Significantly lower amplitudes are seen over central electrodes. B) Contrast-change detection task. Red markers indicate significant within-group differences between the auditory oddball which evokes a sensory-evoked potential and the contrast-change task which does not. There are no significant differences for older adults and 3 over the right hemisphere for younger adults. Group differences show significant frontalisation, with higher amplitudes over frontal electrodes, and lower amplitudes over central and parietal. C) Random-dot motion task. No significant differences are observed at frontal electrodes, with lower amplitudes over central and parietal areas for older adults.

4.3.2.2: ERP by Location

ERPs plotted along each of the selected electrode bands are shown in Figure 4.6. Analysis of these data indicated a significant effect of location on the response-locked slope ($F_{(1.76,76)}$ = 64.724, p < 0.001) and a significant Age by Location interaction ($F_{(1.76,133.722)}$ = 19.667, p < 0.001). Post-hoc analyses indicated that older adults exhibited significantly more positive amplitude values over frontal electrodes (t = -4.956, p < 0.001), but less positive amplitude values over central electrodes (t = 3.030, p < 0.001). However, there was no significant difference over parietal electrodes (t = 0.657, p = 1, see Figure 4.6). The lack of a significant difference in P3b/CPP amplitudes over parietal electrodes may reflect the fact that older adults had a higher hit rate than the younger cohort. This indicates that older adults may not have a lower parietal P3b/CPP amplitude under all circumstances.

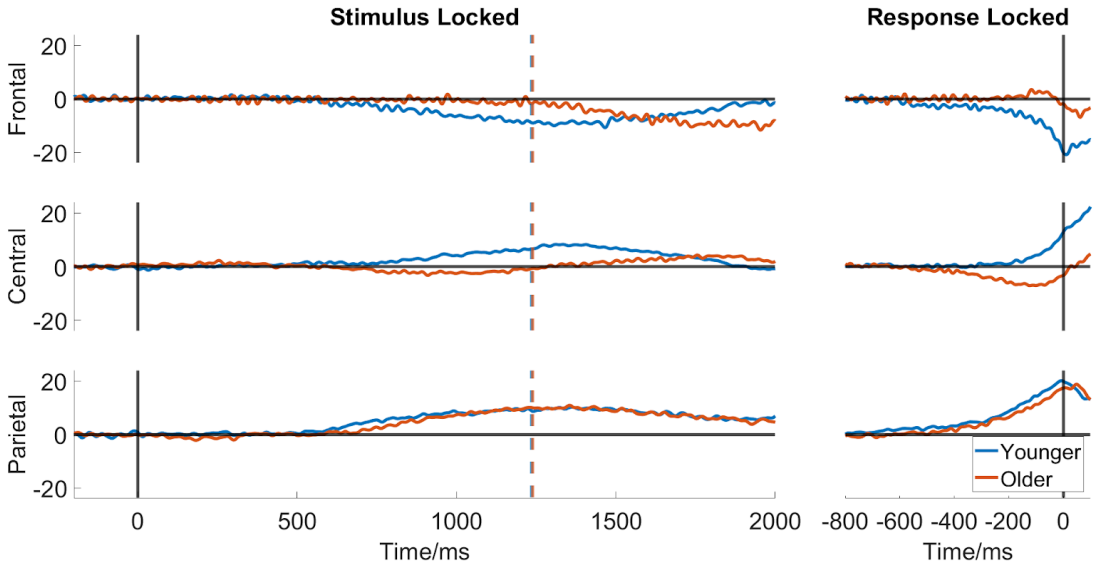

Fig. 4.6: Stimulus and response locked CPP amplitudes for contrast-change detection task, using parietal (Pz), central (Cz) and frontal (Fz) electrodes and their immediate neighbours. Mean reaction times for each group are shown using vertical lines. The negativity observed in younger adults over frontal electrodes appears more central in older adults.

Analysis of response-locked slopes indicated a significant main effect of Location ($F_{(2,76)}$ = 36, p < 0.001), no significant effect of Age (p = 0.373), but a significant Age by Location interaction ($F_{(2,76)}$ = 16.13, p < 0.001). Post-hoc analyses indicated a significantly less positive slope in younger adults over frontal

electrodes (t = -4.638, p < 0.001), but a more positive slope over central (t = 4.1, p < 0.001), with no significant differences over parietal electrodes (p = 0.318, see Figure 4.7). Non-CSD transformed data showed significantly more positive slopes for younger adults over parietal electrodes (t = 4.9, p < 0.001, see Appendices), suggesting that CSD-transformation may encourage the isolation of a distinct positive going parietal P3b, reducing the influence of overlapping signals.

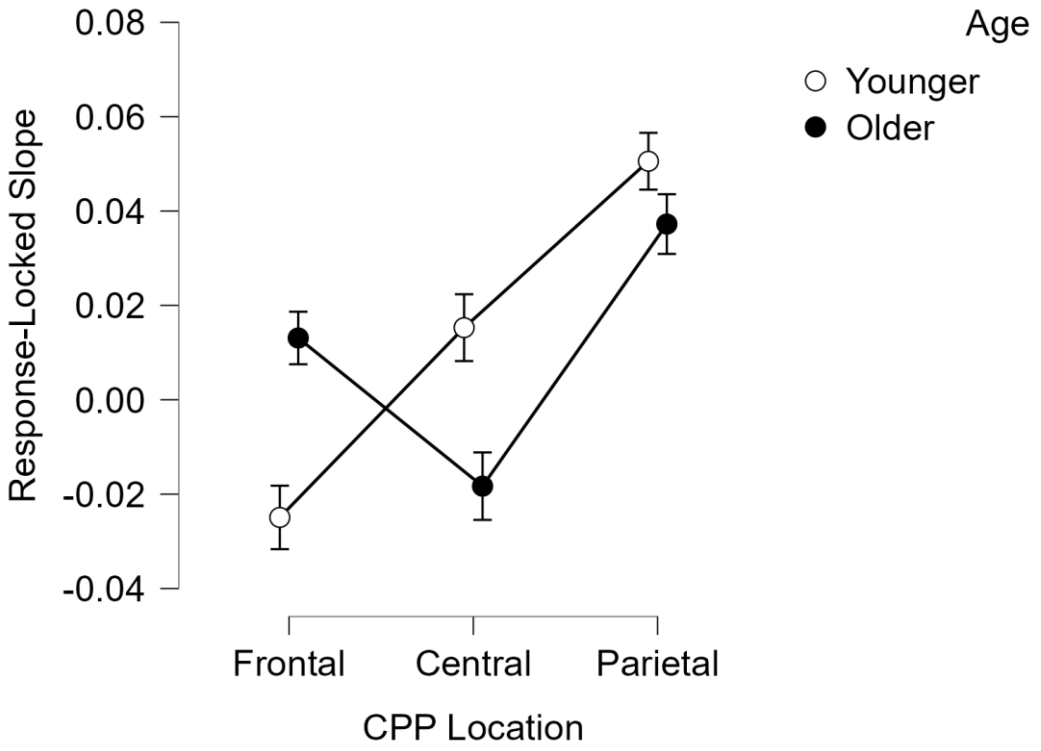

Fig. 4.7) Response-locked slopes by Age and Location for the contrast-change detection task, showing negativity at frontal electrodes in younger adults which moves more centrally for older adults.

4.3.3: Random-Dot Motion

Levene's test for homogeneity of variance was violated, meaning non-parametric equivalents were employed. These indicated a significant effect of Coherence on accuracy (Wilcoxon Z = 3.61, p < 0.001), indicating lower accuracy for low-coherence trials (Median = 93%) than high-coherence trials (Median = 97%). There was a significant main effect of Age (Mann-Whitney U = 887.1, p < 0.001), indicating lower accuracy for older adults (Median = 91%) than younger adults (Median = 97%)

There was a significant main effect of Coherence on RT (Wilcoxon Z = 7.06, p < 0.001) indicating slower RTs for low-coherence trials (Median= 879.5ms) than high-coherence trials (Median= 653.3ms). There was also a significant main effect of Age on RT (Mann-Whitney U = 171, p < 0.001), indicating slower RTs for older adults (Median = 887.8) than younger adults (Median = 648.1).

4.3.3.1: Topography Comparison

The results from the previous two tasks suggest that P3 frontalisation remains evident in older adults even with the application of a spatial filter and when using stimuli that eliminate overlapping sensory-evoked potentials. While the target detection tasks above are in line with previous examples of studies investigating the P3 frontalisation effect in older adults (see van Dinteren et al., 2014 for review), these tasks do not include difficulty manipulation and therefore do not allow measurement of a core characteristic of the P3b/CPP: its evidence-dependent build-up. As such, it becomes difficult to fully isolate the CPP from potentially overlapping background activity. Analysis of a random-dot motion task containing two coherence levels, and thus, two evidence strength levels, allows for a more intricate exploration of CPP localisation and inter-group differences.

CSD-transformed topographies showed a focalised parietal positivity and a frontal negativity in both (see Figure 4.5). This negativity appeared more centralised in older adults and more frontal in younger adults. Cluster-based t-tests were also used to compare group differences in ERP topographies. A frontal

positivity was observed in older adults, however, compared to the previous two tasks, there were fewer significant group differences at frontal electrodes. Significant group differences were observed over centro-parietal electrodes, where older adults show significantly smaller amplitudes (see Figure 4.5, electrode-clusters with a significant group difference marked in black).

4.3.3.2: ERP by Location

ERPs plotted along each of the selected electrode bands are shown in Figure 4.8. Analysis of this data indicated a significant main effect of Location ($F_{(1.76,64)}$ = 67.336, p < 0.001), Age ($F_{(1,64)}$ = 12.184, p < 0.001) and Coherence ($F_{(1,64)}$ = 33.602, p < 0.001) on response-locked amplitude. There was a significant Age by Location interaction ($F_{(1.76,64)}$ = 21.867, p < 0.001). Post-hoc analyses of the Age by Location interaction indicated similar trends as those observed in the previous tasks, with a significantly less positive in younger adults over frontal electrodes (t = - 4.423, p <0.001), but a more positive amplitude over central (t = 5.534, p < 0.001), and parietal electrodes (t = 3.364, p = 0.004, see Figure 4.8). This reduced amplitude over parietal electrodes may be consistent with the finding for lower accuracies in older adults on this task. Comparatively, on the contrast-change detection task, there were no group differences over parietal sites and higher hit rates for older adults.

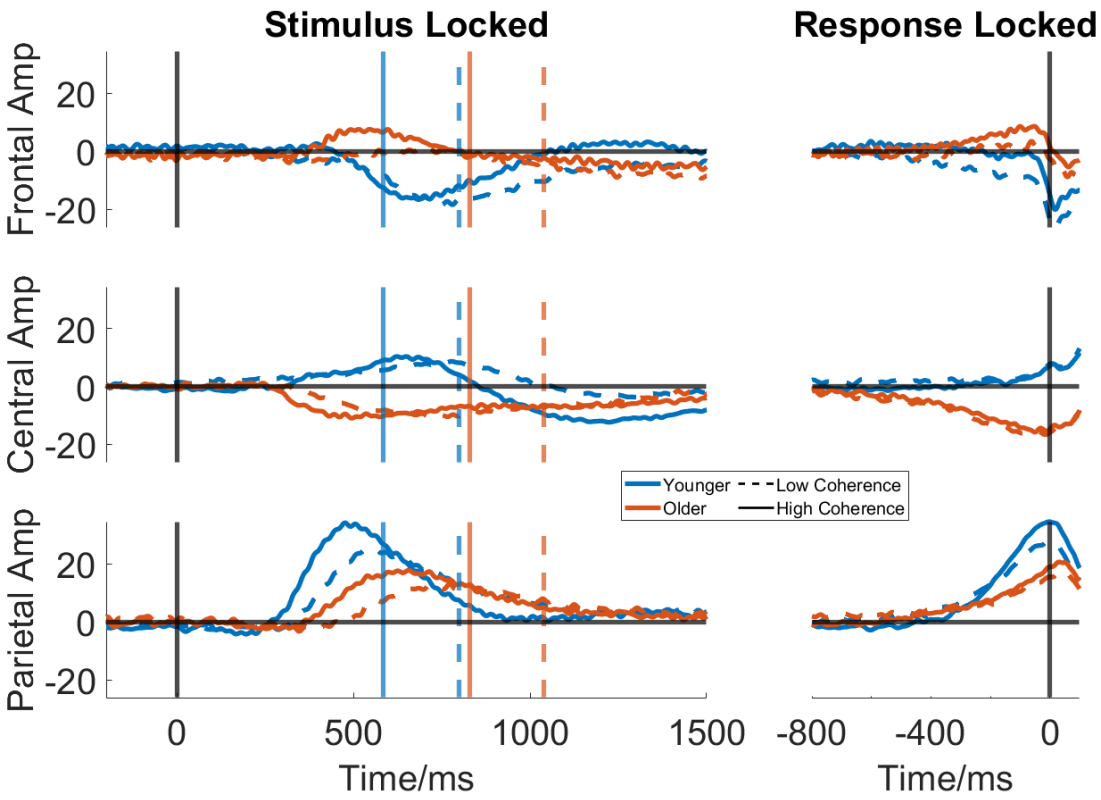

Fig. 4.8: Response-locked slope by electrode location for non-CSD transformed data. Separate lines indicate different age groups.

There was also significant Location by Coherence interaction ($F_{(1.587,64)}$ = 4.232, p = 0.025). The Location by Coherence interaction is plotted in Figure 4.9. Post-hoc analyses indicated that parietal areas show a higher amplitude for high coherence trials in both groups (t = 3.601, p = 0.002). The same is seen for frontal electrodes (t = 4.437, p < 0.001). There is no difference in amplitude across central electrodes (t = 0.251, p = 1). There was no significant Coherence by Age interaction (p = 0.174) or three-way interaction (p = 0.451). Similar trends were observed for non-CSD transformed data, reported with stimulus-locked analyses in the Appendix.

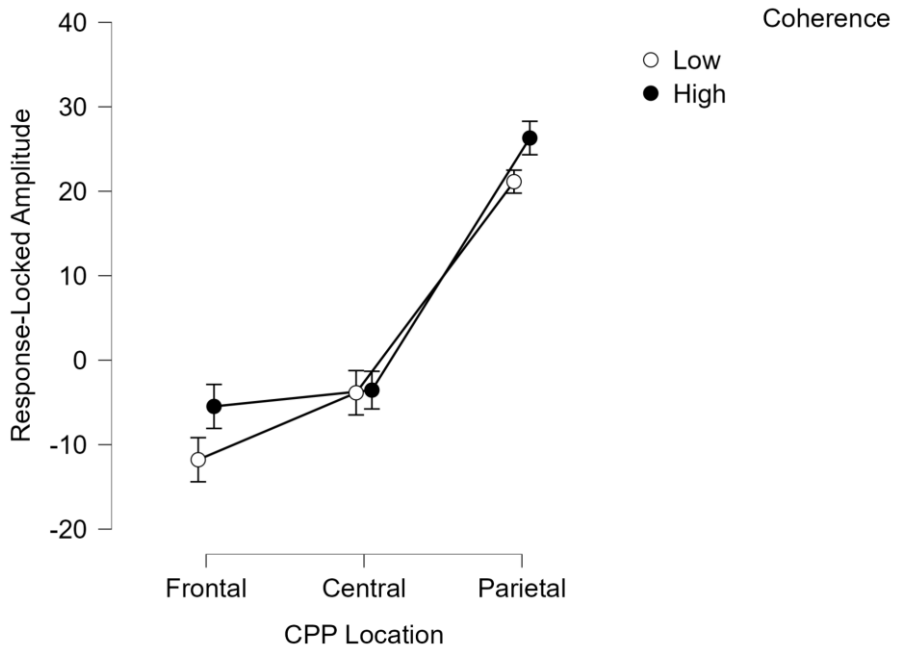

Fig. 4.9. Response-locked amplitude by electrode site and coherence. Significant coherence effects are observed at frontal and parietal sites, with more positive amplitudes for high coherence trials, but not over central electrodes.

Analysis of the effect of Age, Location and Coherence on the pre-response slope was also performed. There was a significant effect of Location ($F_{(1.76,64)} = 41.926$, $p < 0.001$), Age ($F_{(1,64)} = 15.101$, $p < 0.001$) and Coherence ($F_{(1,64)} = 31.097$, $p < 0.001$) on response-locked slope. There was a significant Age by Location interaction ($F_{(1.76,64)} = 19.499$, $p < 0.001$), but no significant Location by Coherence interaction ($p = 0.167$), Coherence by Age interaction (p = 0.548) or three-way interaction (p = 0.385). Figure 4.10 plots the significant Age by Location interaction on the response-locked slope. Here younger adults show a pattern which goes from more negative slopes at frontal areas, to increasingly more positive slopes through central and parietal areas. Post-hoc analyses indicated significant age differences at all sites (Frontal: $t = -3.848$, $p < 0.001$, Central: $t = 4.346$, $p < 0.001$, Parietal: $t = 4.709$, $p < 0.001$), with the significant interaction driven by a cross-over effect. Over central electrodes, older adults have negative slopes while the younger have positive slopes. Over parietal electrodes, both groups have positive slopes (see Figure 4.10).

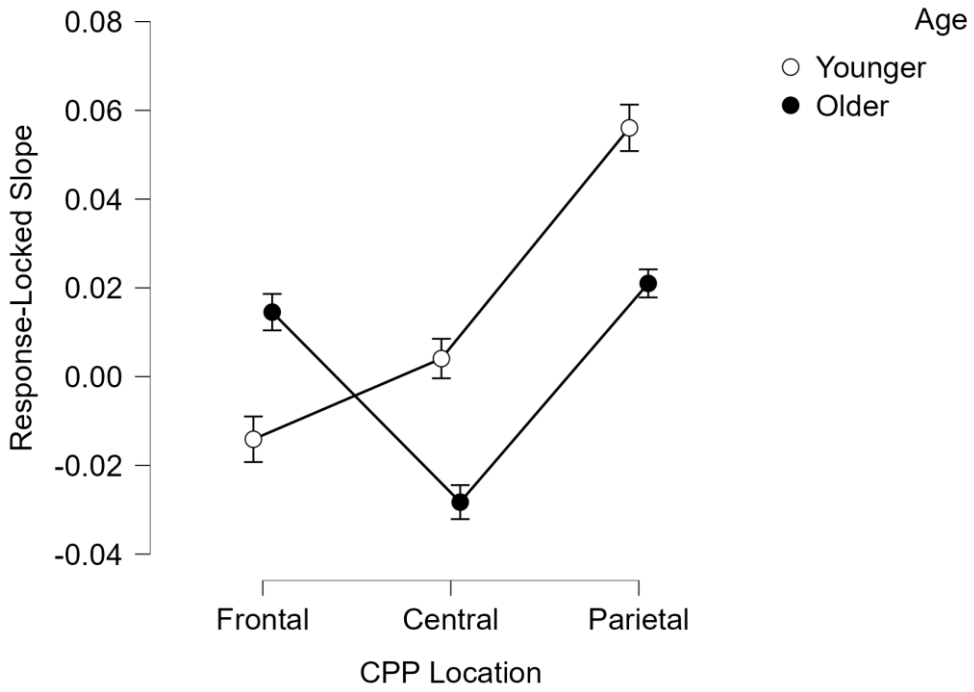

Fig. 4.10. Response-locked slope by electrode site. Significant differences are observed at all sites, but not over central electrodes.

The CPP is shown to accumulate at slower rates to lower amplitudes for slower RTs. To investigate this, trials were binned into fast and slow RTs using a median split for each coherence level. There was a significant main effect of RT bin on response-locked amplitude ($F_{(1,64)}$ = 44.81, p < 0.001) suggesting more negative amplitudes for the slower RT bin (Mean = 1.49 *SD* = 7.99), than the Fast (Mean = 5.94, *SD* = 7.85, see Figure 4.11), with no significant interactions between RT bins and other effects (all p > 0.078). There was no significant effect of RT bin on response locked-slope (p = 0.537), with no significant interactions between RT bins and other effects (all p > 0.247).

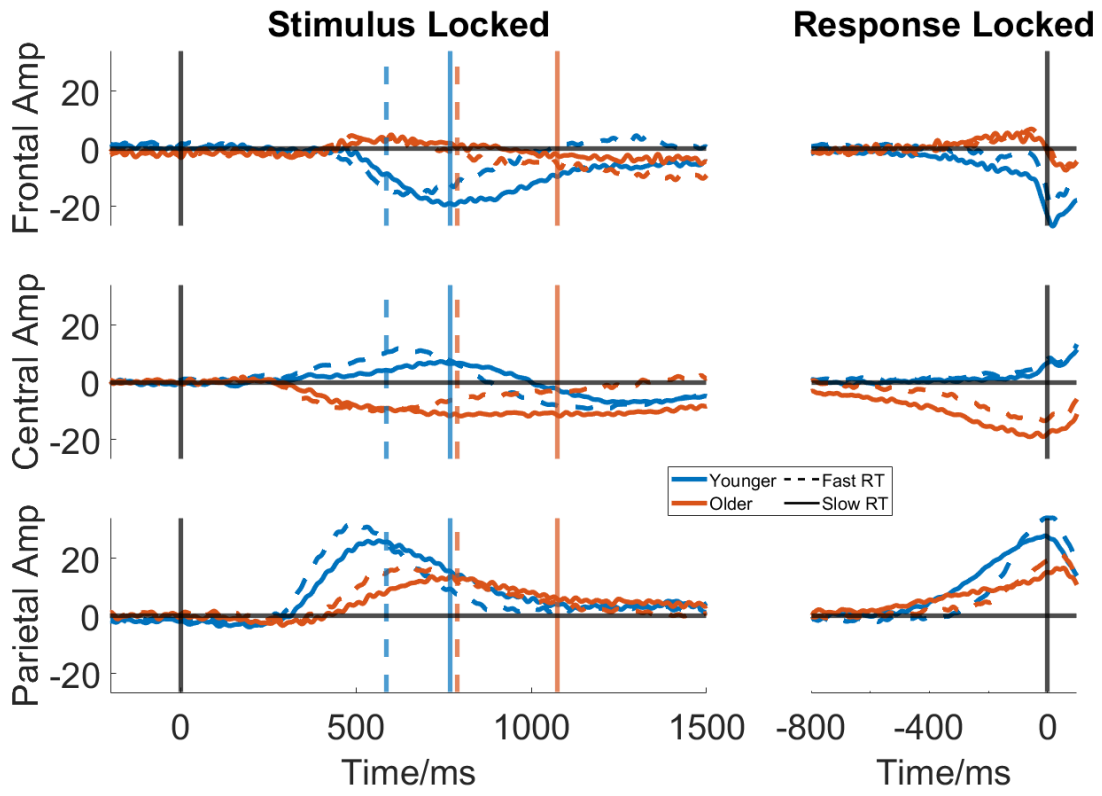

Fig 4.11. Stimulus and response-locked P3b/CPP amplitudes by location, binned by reaction time. Vertical Lines show the 25th (Fast) and 75th (Slow) percentile RT for each group. Significantly more negative amplitudes are observed for slower RTs.

4.3.3.3: CNV by Age and Reaction Time

Collectively, these studies indicate more positive amplitudes over frontal electrodes for older adults. They also suggest a negative going signal that is more present in frontal electrodes for younger adults but appears more centrally for older adults (see Figure 4.12). To investigate this potential shift an Age-Location ANOVA was performed across these two locations, using Frontal (Fz) and Central (Cz). As discussed, the frontal CNV is a likely candidate signal behind these negativities. This signal is shown to be evidence-independent but RT-dependent. As such, trials were binned into fast and slow responses. These ERPs are plotted in Figure 4.12, showing a negative going signal for younger adults over the frontal electrode, and a similar negative going signal for older adults over the central electrode.

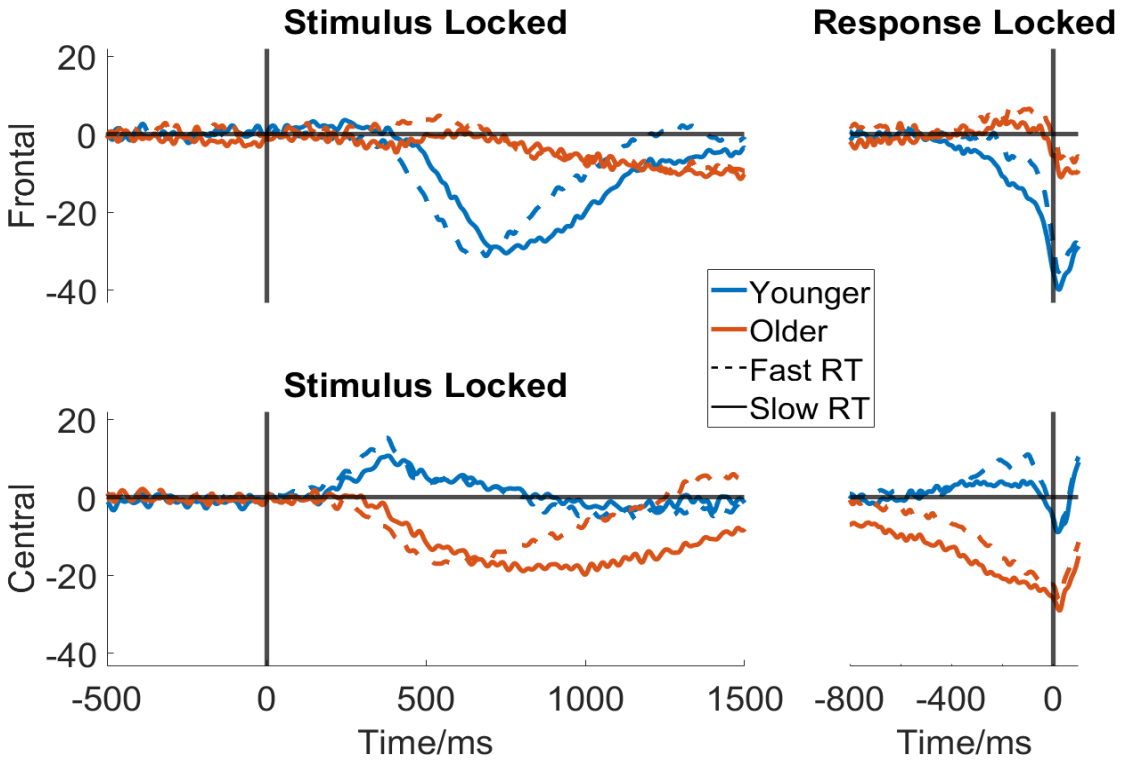

Fig. 4.12: Response-locked slope by electrode location for non-CSD transformed data. Separate lines indicate different age groups.

There was no significant main effect on response-locked CNV amplitudes for Age (p = 0.463), but a significant main effect of RT Bin ($F_{(1,64)}$ = 6.69, p = 0.012) with lower amplitudes for slower RTs (Mean = -17.1, SD = 16.9), than fast (Mean = -13.24, SD = 19.17). There was a significant Age by Location interaction ($F_{(1,64)}$ = 19.78, p < 0.001), but no Location by RT bin interaction (p = 0.415), or three-way interaction (p = 0.67). Post-hoc analyses of the Age by Location interaction indicated significantly more negative amplitudes for younger adults over the frontal site (t = -3.92, p < 0.001), with significantly more negative amplitudes for older adults over the central electrode (t = -2.99, p = 0.014).

Analysis of response-aligned slopes showed a significant main effect of Age ($F_{(1,64)}$ = 7.32, p = 0.009), a significant main effect of RT Bin ($F_{(1,64)}$ = 9.84, p = 0.003) and Location ($F_{(1,64)}$ = 17.77, p < 0.001). There was no significant three-way interaction (p = 0.38) or Location by RT Bin interaction (p = 0.247). There was a significant Age by Location interaction ($F_{(1,64)}$ = 8.97, p = 0.004). Post-hoc analyses indicated significantly more negative slopes in younger adults over frontal electrodes (t = -4.04, p < 0.001), but no significant age difference in slope

across central electrodes (p = 0.75). There was no significant difference between younger adults' slopes between frontal and central electrodes (p = 0.754), however, older adults showed significantly more negative slopes when over central electrodes than frontal electrodes (t = -4.95, p < 0.001).

There was also a significant Age by RT bin interaction ($F_{(1,64)}$ = 4.72, p = 0.034). Post-hoc analyses indicated that younger adults showed a significantly more negative slope for slow RTs than fast RTs (t = -3.87, p = 0.001), while older adults showed no significant difference in slope across RT bin (p = 0.81).

Collectively, this data shows a negative going signal over central locations in older adults, while this negativity appears more frontally in the younger group. The response-locked slopes and amplitudes of these signals did not differ in their RT dependence based on location. As such, it is inconclusive as to whether it represents an age-related shift of the CNV. Future work designed to induce other specific effects on the CNV such as through a speed-accuracy emphasis manipulation would be needed.

4.3.4: ICA Decomposition to Isolate the CPP

Sixty-four component ICA was performed on minimally cleaned data and selected as outlined in the Methods section 4.2.6 above. Components were ranked based on the degree to which their build-up rates scaled with evidence strength (high > low coherence) and RT (fast > slow). IC activation profiles and topographies are shown in Figure 4.13. Both groups show highly similar CPP-component locations with a posterior parietal peak. Older adults show a slightly more posterior presentation than younger adults. There appears to be no real frontalisation in the older adult group, and the anterior positivity shown in the older adults' ERP topographies has been removed. These differ from the ERP topographies shown in Figure 4.5. The central negativity seen in older adults is also largely removed but is still observable.

4.3.4.1: IC Activations

Activation topographies were compared between the two groups, using a cluster-based permutation statistic to correct for multiple comparisons (see Methods). No reliable group differences between any group at any cluster point were found (all $p > 0.382$) suggesting no between-group differences in IC activation topographies (see Figure 4.13.A).

Analysis of the activation profiles revealed that there was a significant effect of Coherence ($F_{(1,64)} = 20.859$, $p < 0.001$) and Age group ($F_{(1,64)} = 5.859$, $p = 0.018$) on response-locked activation amplitude, with no significant Age by Coherence interaction ($p = 0.267$), indicating higher IC amplitudes in high coherence in both groups, with the younger groups showing higher amplitudes overall (see Figure 4.13.B). On the pre-response slope, there was an expected significant effect of Coherence ($F_{(1,64)} = 165.938$, $p < 0.001$), with higher slopes for high coherence trials and a significant effect of Age ($F_{(1,64)} = 19.789$, $p < 0.001$), showing higher slopes on average for younger adults. There was a significant Age by Coherence interaction ($F_{(1,64)} = 6.123$, $p = 0.016$), driven by the stronger Coherence effect in younger (Mean Difference = 0.006, $t = 1.479$, $p < 0.001$) than older adults (Mean Difference = 0.004, $t = 1.003$, $p < 0.001$).

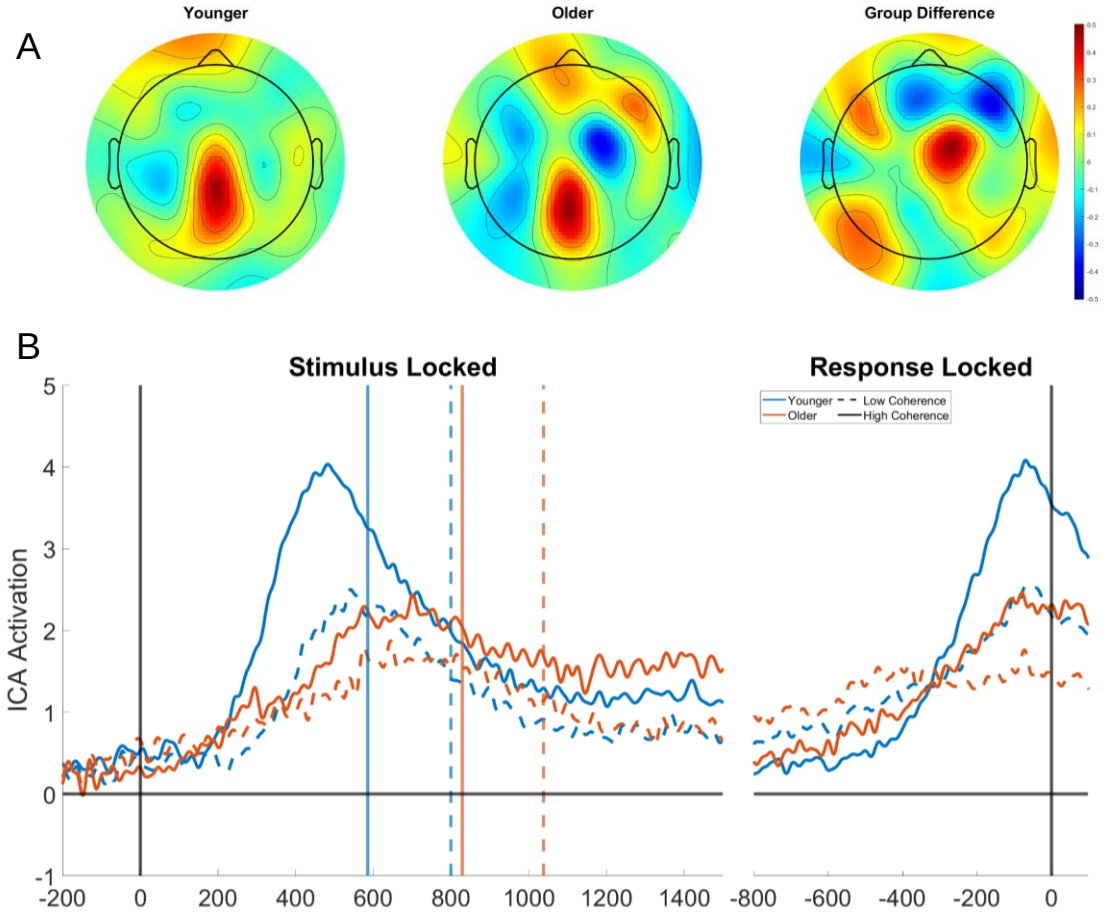

Fig 4.13.A) Response locked ICA topographies -50:50ms of top-3 ranked ICA's based on CPP features. ICA activations with negative amplitude over CPP sites at response or 200:600ms stimulus-locked are sign-inverted for plotting, with group differences between older and younger. B) Stimulus and response locked ICA activation profiles for top-3 ranked ICA's. Vertical lines show mean reaction time by group and condition.

4.3.4.1: Back-Projected ERP

Back-projected topographies were compared across groups, using a cluster-based permutation statistic to correct for multiple comparisons (see Methods). No reliable group differences between any group at any cluster point were found (all $p > 0.624$, see Figure 4.14.A). In order to assess the strength of evidence in favour of the null hypothesis, a Bayesian independent sample t-test was performed on activity averaged across clusters in the same time window. Along the central line of clusters (POz, Pz, CPz, Cz, FCz Fz), there was found to be moderate to anecdotal evidence in support of the null hypothesis (Range: 0.253 at POz to 0.818 at AFz), suggesting no group differences in back-projected response-locked topographies.

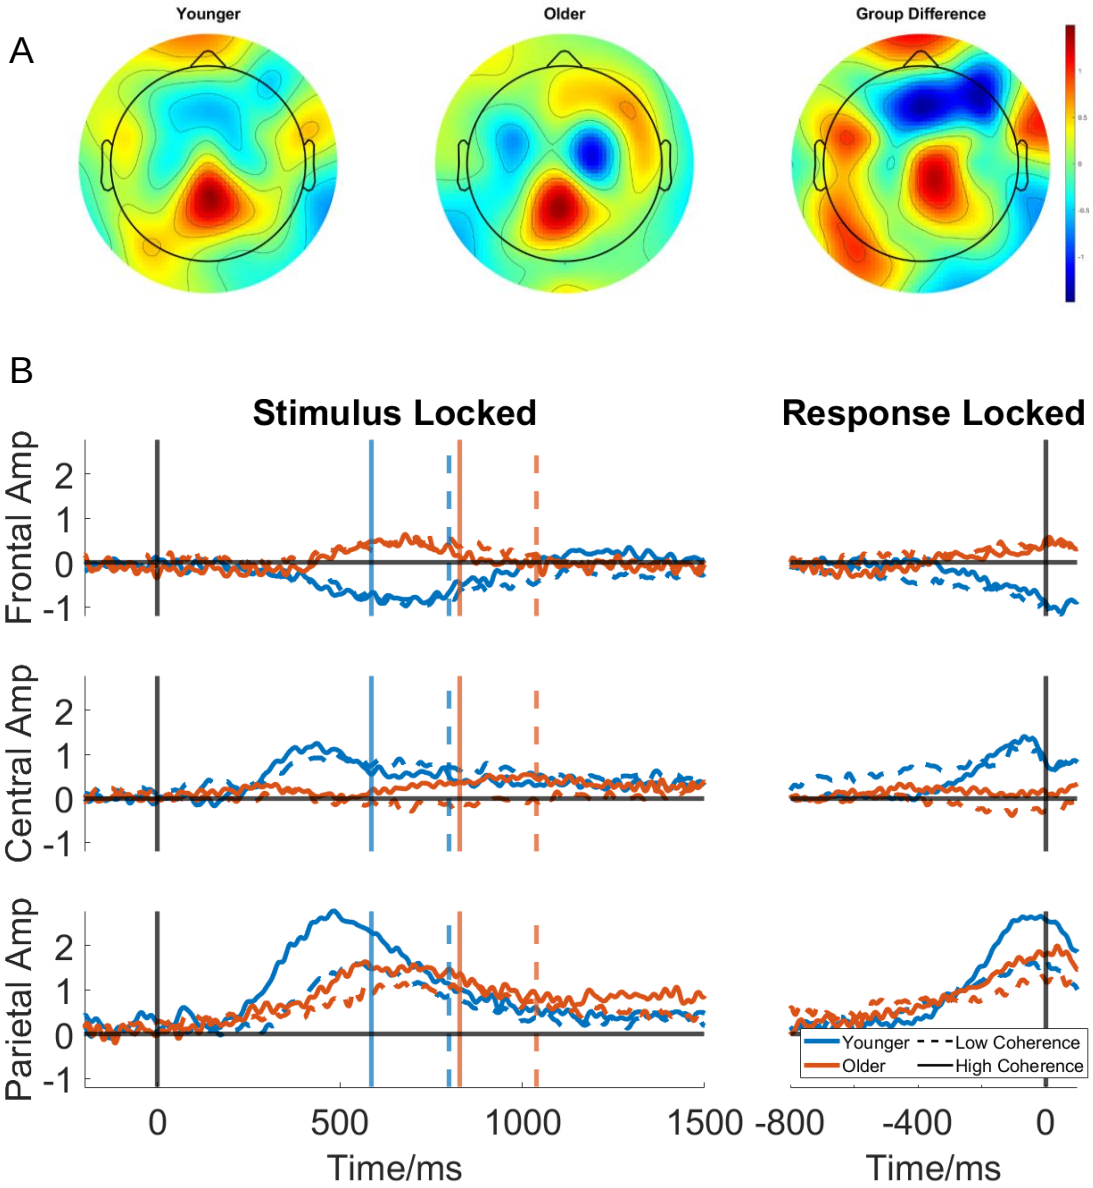

Fig 4.14 A): Topographies of average response locked back-projected ERPs across each of the top 3 components per subject -50:50ms, with group difference between older and younger. B) Stimulus and response locked average ERPs of top-3 back-projected components at their CPP location. Vertical lines show average response times for each group and condition

Response-locked amplitudes and slopes were compared across the same Location sites used in the CPP analysis for the back-projected data. Analysis of response-locked amplitude showed a non -significant effect of Age (p = 0.842), but a significant effect of Location ($F_{(1,64)}$ = 10.384, p < 0.001) and Coherence ($F_{(1,64)}$ = 13.393, p < 0.001). There was no significant Age-Location-Coherence interaction (p = 0.479), or Coherence-Age interaction (p = 0.747) but a significant Location-Coherence interaction ($F_{(1,64)}$ = 3.382, p = 0.037), with post-hoc

analyses indicating no significant coherence effect at frontal (p = 0.722), or central sites (p = 0.722), but significantly lower amplitudes for low-coherence trials over parietal electrodes (t = -4.11, p < 0.001, see Figure 4.15). There was also a Location-Age interaction ($F_{(1,64)}$ = 3.496, p = 0.033), with post-hoc analyses indicating no significant age differences at any location (all p > 0.225). The interaction term was likely driven by a cross-over effect where negative amplitudes observed over frontal electrodes for younger adults which is not present for older adults (see Figure 4.15)

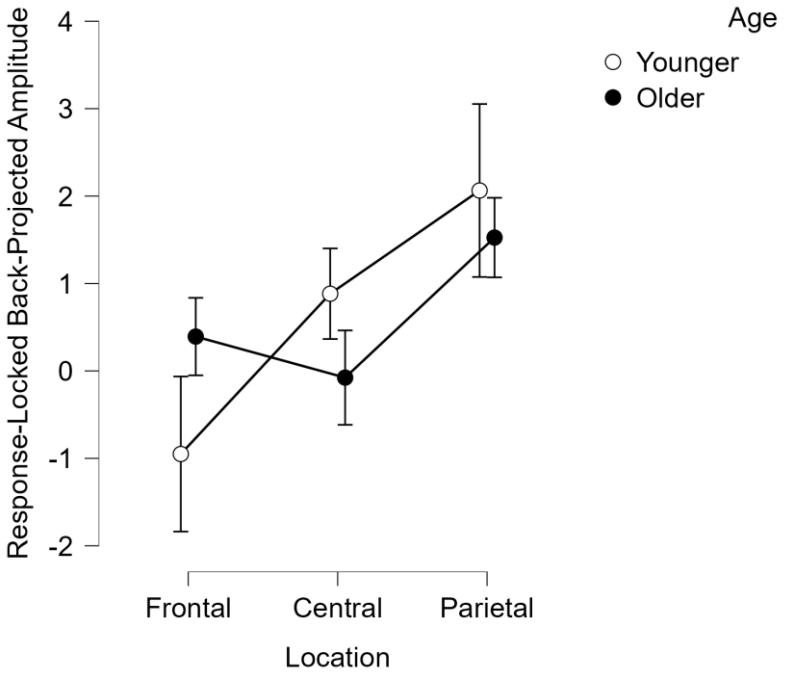

Fig 4.15: Response-locked back-projected CPP amplitude, showing a negativity over frontal electrodes in younger adults that is not present for older adults.

To assess if the positive-going, coherence-dependent back-projected CPP was localised solely to the centroparietal sites, or if coherence effects were also shown at frontal sites in either age group, analysis of the response-locked slope was performed. Analysis of the response-locked slope showed a predicted significant main effect of Coherence ($F_{(2,64)}$ = 32.941, p < 0.001) steeper slopes for high-coherence trials, Age ($F_{(2,64)}$ = 5.940, p = 0.018) with steeper slopes for younger adults, and Location ($F_{(2,64)}$ = 6.275, p = 0.003). There was a significant Age-Location interaction ($F_{(2,64)}$ = 4.466, p = 0.013), however post-hoc analyses indicated no significant group differences in response-locked slope at any electrode site (all p > 0.18). This interaction was likely driven again by the

negative going signal seen in frontal electrodes for younger adults, which is not observed in older adults (see Figure 4.14.B). There was also a Location-Coherence interaction ($F_{(2,64)}$ = 3.685, p = 0.028). Post-hoc analyses indicated significantly higher slopes for high compared to low coherence at central sites (t = 4.485, p = 0.022) and parietal sites (t = 5.821, p < 0.001, see Figure 4.14.B). However, there was no coherence effect at frontal sites (p = 1). There was no Age-Location-Coherence effect, (p = 0.062), suggesting consistent coherence effects across location in both age groups. Collectively, this analysis indicates that the evidence-dependent CPP is localised to centroparietal areas for both groups.

4.4: Discussion

Previous studies have shown significant frontalisation of brain activity between older and younger adults, potentially representing compensatory mechanisms as the brain ages (Davis et al., 2008; van Dinteren et al., 2014b). This presents a challenge for the comparison of signals such as the P3b/CPP across groups, as their location and profile may be different. This may be due to a true frontalisation but may also be driven by a number of signals that are known to evolve on an overlapping timescale such as the P3a, CNV and sensory-evoked potentials (Kelly & O'Connell, 2013, West et al., 2010). As a result, changes in P3b/CPP amplitude at response could be reported as evidence for reductions in decision bound, when in fact the differences are due to potentially functionally irrelevant differences in signal presentation (Antonakakis et al., 2020; Morcom & Henson, 2018). As neural signals are becoming increasingly relied upon as estimates of underlying psychological processes (Kelly et al., 2021; McGovern et al., 2018), studies may then erroneously attribute reductions in P3b/CPP amplitude to changes in some cognitive construct. It is evident that the origins of CPP frontalisation must be addressed in order to enable accurate between-group comparison of signals and extract more accurate neural indices of the decision-making process. This study aimed to achieve this while establishing a novel, ICA-driven method for extracting reliable markers of evidence accumulation.

The application of surface Laplacian filters such as the CSD has been shown to significantly improve the spatial resolution of EEG data, disentangling ERP components from each other (Kelly & O'Connell, 2013). In this study, the use of these filters removed the effect of age on stimulus-locked peak amplitude and response-locked amplitude but retained the significant interaction and age-by-location interaction effects. Cluster-based t-tests comparing older and younger topographies, and age-by-location ANOVAs on both the non-CSD and CSD transformed data indicated more positive amplitudes over frontal electrodes in the older group and less positive amplitudes over parietal areas, supporting the idea that changes to the P3b amplitude are location-specific (van Dinteren et al., 2014a). However, the application of this spatial filter reduced the spread and the number of these differences, with CSD-transformed data showing more focalised

amplitude differences in both parietal and frontal areas. Furthermore, the CSD transform revealed two distinct signals; a centroparietal positivity, and a frontocentral negativity, which may appear more centrally in older adults. This represents a significant advancement as the majority of previous work was performed on data sets in which no CSD transform was applied (e.g., O'Connell, Balsters, et al., 2012; van Dinteren et al., 2014b). These results indicate that spatial filters are beneficial to the study of between-group differences in signals, providing more spatially resolute topographies.

The frontalisation effect was then examined on a task which removes the impact of sensory-evoked potentials on the resulting ERPs. The contrast-change detection task also showed a significantly lower amplitude in younger adults over frontal electrodes, but a more positive amplitude over central electrodes. Comparison of the older adult topography across these two tasks indicated no significant differences at any electrode cluster, indicating that the frontalisation effect cannot be explained by overlapping sensory-evoked potentials alone. Similar frontalisation effects were observed in the random-dot motion task. However, one consistency across tasks was the retention of a centroparietal positivity in both groups in similar parietal locations regardless of task demands. Notably, response-locked parietal amplitudes did not differ across groups in the contrast-change detection task where they outperformed younger adults yet showed lower amplitudes in the remaining two tasks where older adults performed more poorly. This provides tentative evidence that, in spite of an observed frontal positivity in older adults, CPPs may be best measured over parietal sites as this area seems most linked to task performance.

Of particular note, if the P3b/CPP frontalisation in older age is driven by compensatory frontal recruitment (Kuruvilla-Mathew et al., 2022; van Dinteren et al., 2014a), it should be most apparent in the random-dot motion task with varying difficulty levels. Cluster-based topography comparisons indicated that this was not the case, as frontalisation was more strongly observed for the contrast-change and auditory oddball detection task. This provides tentative support for West and colleagues (2010) findings, which suggest no frontalisation of the P3b, but an increased P3a in older adults for a novel or oddball stimuli. As both the

auditory-oddball and contrast-change detection tasks required target detection, the observed differences between group topographies may be driven by an increased P3a for novelty detection with age, while the P3b as a marker of evidence accumulation appears largely consistent in location across age groups.

In order to further spatially isolate these signals, ICA was performed on data from a continuous monitoring version of the random-dot motion task that included two coherence levels. A proposed evidence accumulation signal such as the CPP builds at a rate that scales with evidence strength and reaction time (e.g., Twomey et al., 2015, Steinemann et al., 2018), functional characteristics that were leveraged here to cluster relevant components. The selected components, and their back-projected ERPs, exhibited topographies with a highly similar centroparietal focus for both older and younger adults. There was no observable age-related frontalisation of these components, with the elderly group, exhibiting a potentially more posterior maximum. This provides further evidence that centro-parietal sites are a reliable location to measure the CPP for both older and younger adults.

Overall, the current results suggest that older and younger groups have highly similar P3b/CPP topographies. When CSD is applied to an auditory-oddball task, both groups show distinct parietal positivity and frontal negativity in both groups, with a second, more anterior positivity in older adults. The use of a gradual-evidence emergence task to remove sensory-evoked potentials showed similar results. Finally, extracting an evidence-dependent CPP-IC to further minimise signal overlap confirmed that there was a highly similar parietal focus of the P3b/CPP across both age groups. When these CPP-ICs were back-projected to the raw EEG, a Location by Age interaction was no longer observed, suggesting that the CPP retains its parietal focus into senescence. Furthermore, a significant effect of coherence was only observed over parietal electrodes, indicating that parietal sites are the most appropriate areas from which to measure the P3b/CPP across age groups.

One outstanding question is whether the CNV moves more centrally in older adults. Comparing group topographies, we see a potential shift of negativity to

more central sites in older adults, although these group differences were non-significant. However, examining the ERPs plotted by location across all tasks we can see a potential negativity in central electrodes for older adults that appears more frontally in the younger groups. Evidence for a potentially more central negativity in the older adult group was also seen in the IC decomposition. Close examination of the activation and back-projected topographies indicates the potential presence of a central negativity which appears more pronounced in the older adult cohort (4.13.A). In the back-projected ERP, a potential negative going signal over frontal electrodes is observed in the younger adult data, with a positivity over central electrodes. However, older adults show no positivity over central electrodes, instead presenting a negative going signal, especially in low-coherence trials (see Figure 4.14.B). This may have been driven by the IC selection criterion. As the CPP and frontocentral CNV align closely in their emergence over time and reaction-time dependence, the clustering method utilised may have resulted in the incorporation of some CNV activity. As such, while successful in extracting a CPP component, this component may contain some residual activity from similarly evolving but distinct neural processes.

The aim of this chapter was to exclude the CNV component from the CPP measurements rather than to isolate the CNV itself. As such, it contained no manipulation, such as a speed-accuracy tradeoff, which would elicit specific changes to the CNV. A task which would successfully manipulate levels of dynamic urgency, such as the speed-accuracy paradigm outlined in Chapter Two, would allow for CNV-ICs to be extracted based on their increased build-up under speed pressure, and pre-evidence adaptations to the signal. This would give two functional constraints to the ICA selection of a CNV-component. Activation and back-projection topographies for distinct CNV and CPP-components could then be examined and compared across age groups in order to establish the extent to which the CNV signal shifts as a function of ageing, and whether it may contribute to the observed central negativity in older adults. These results would also help to more concretely establish the extent to which CPP frontalisation effects may be driven by changes to alternative signals, however, this was beyond the scope of this preliminary study.

As discussed, functional source separation has also been applied as a means of isolating functionally relevant components from EEG data (Porcaro et al., 2019). To date, no work has been done to incorporate the evidence-dependent nature of the CPP into the functional source cost function, instead relying on the signal's stimulus-locked amplitude alone. As signal amplitude can be influenced by a range of sources (Kelly & O'Connell, 2013), future work may do well to build upon the methods introduced here, utilising the evidence-dependent nature of the CPP build-up as additional constraints on an FSS analysis's cost function. This would provide a clearer separation than ICA, extracting components which are more reflective of evidence-accumulation. This would also preclude the need to rank components as was done in this study and would allow for a more data-driven isolation of components. This technique could also be applied to the extraction of a CNV-component using the characteristics described above. Through this, novel functionally extracted components could provide even more fitting indices of cognitive function and may be more suitable for us in refining and constraining neurally-informed models of decision-making.

As discussed at several points throughout this thesis, older adults are shown to exhibit significantly higher decision bounds than younger adults (Forstmann et al., 2011; see Dully et al., 2018 for review). However, P3b/CPP amplitudes are shown to be consistently lower in older adults (Chapter Two; O'Connell, Balsters, et al., 2012; van Dinteren et al., 2014b for meta-analysis; West et al., 2010). In Chapter Two, the modelling presented showed lower bounds in older adults, however, the degree of this boundary reduction was not of the same magnitude as the between-group difference in CPP (see section 2.3.6.1: Figure 2.18). Even through ICA, the extracted evidence-accumulation components depicted here show significantly smaller amplitudes at response in older adults. The reduced amplitude at response of the CPP is therefore difficult to reconcile with modelling efforts such as those shown in Chapter Two. It suggests the possibility that the models may be incorrect, failing to provide a good estimate of how evidence accumulation occurs in the brain. However, it may instead lend evidence to the idea that reduced CPP amplitudes may not be due to a functional difference but may be contributed to at least partly by task-irrelevant differences such as differences in voltage conduction. For example, changes in skull thickness with

age result in a significant reduction of signal conductivity through the skull (Antonakakis et al., 2020). Here, a completely task-irrelevant feature of physiology which is affected by ageing contributes to signal presentation. When comparing groups these effects are compounded, meaning we may conclude differences in a cognitive process when in fact the differences are due to some task-irrelevant feature. This highlights the benefit of conducting behavioural modelling in tandem with EEG recordings. Taking EEG indices at face value as markers of some underlying decision process or parameter in a model may result in erroneous conclusions. To combat this, and to produce a neural index capable of reflecting model parameters such as boundary at response and drift rate, a method for scaling the CPP which retains within-group differences while bringing the signals into a more comparable scale is required. In Chapter Two, novel methods for addressing these scaling issues were introduced, yet significantly more work is needed to provide a data-driven solution to this signal scaling problem.

Collectively, this study indicated no significant age-related frontalisation of the evidence-dependent CPP components. Application of CSD to ERP data was shown to significantly reduce the effects of age on signal amplitude, but examination of IC activations showed consistent reductions in component amplitude in older adults. This suggests that these decision-related sources produce signals with truly lower amplitudes and that some, but not all, reductions in amplitude can be explained by signal mixing. Of particular note, the CNV may present more centrally for older adults, potentially overlapping with and reducing the amplitude of the neighbouring CPP. Future work may benefit from exploring manipulations which allow for functional predictions for both the CPP and CNV, allowing for the extraction of these components via ICA. This would allow this assertion to be tested more concretely. Overall, the use of ICA to derive neural components based on core functional properties of decision signals presents a potentially novel and compelling tool for the extraction of neural indices of the decision process. Future work may also apply these techniques to other signals of interest such as the motor potentials used as neural constraints in Chapter Three. This may serve to provide more refined indices of decision-making processes for use in neurally-informed modelling. In conclusion, this study

presents novel methods for improving between-group comparisons of the P3b/CPP between older and younger adults, while highlighting that the evidence-dependent P3b/CPP remains best measured at centroparietal sites into senescence.

Chapter Five:

General Discussion

5.1: Overview

The field of perceptual decision-making has made critical contributions to our understanding of core cognitive and perceptual processes through a combination of computational modelling and the analysis of neural data. While each of these methods has proven fruitful in its own right, allowing us to probe the latent processes which give rise to behaviour, an emerging body of work has pointed to the overwhelming need to synthesise these approaches into a neurally-informed modelling framework (Kelly et al., 2019, 2021; McGovern et al., 2018; O'Connell et al., 2018, Chapter One, section 1.3). A range of recent work has identified a core set of neural indices capable of reflecting distinct stages of the decision-making process (e.g., O'Connell et al., 2012; Twomey et al., 2015; see Chapter One, section 1.2.3). Constraining the parameters of decision models, such as the drift-diffusion model (DDM), with these neural signatures of the decision process has proven useful in detecting subtle and complex adjustments to mechanisms underlying ageing (McGovern et al., 2018) and the implementation of prior knowledge (Kelly et al., 2021). While in its early stages, this nascent field has already contributed significantly to our understanding of the brain, allowing us to create biologically-grounded models that both reflect and support neural data, and granting an added benefit in allowing us to explore more nuanced models of the complex perceptual decision-making process (e.g., Kelly et al., 2021; Corbett et al., in press, Ghaderi-Kangavari et al., 2022). However, as outlined throughout this thesis, there remain some fundamental shortcomings of sequential sampling models which had yet to be addressed thoroughly by neurally-informed modelling, such as their inability to accommodate potential changes to their scaling parameters across groups (Chen et al., 2014; Park et al., 2017), or under conditions such as learning (Dosher & Lu, 1998, 1999; 2017). Furthermore, some underlying issues in the isolation of reliable neural indices may undermine their capacity as neural markers of the decision-making process, especially when comparing these signals across distinct groups (Stacey et al, 2021; Rossiter et al., 2014, Antonakakis et al., 2020).

The purpose of this body of work was to address these issues while making efforts to characterise and explore the key benefits of a neurally-informed

approach to computational modelling of perceptual decision-making. Across three studies this thesis illustrated the potential applications of neurally-informed modelling to perceptual decision-making data and developed new techniques which can be applied to future work. This final chapter will provide an overview of the research presented, summarising and outlining each study's novel contributions to our understanding of perceptual decision-making. Finally, potential avenues for future research will be described. In conclusion, a review of this work indicates that neurally-informed modelling offers an exciting new means for studying the complex process which underpins some of our most basic perceptual processes.

5.2: Novel Insights into Perceptual Decision-Making

5.2.1: Speed-Accuracy Tradeoff in Older Adults

Chapter Two applied neurally-informed modelling to the well-studied field of speed-accuracy adaptations in older adults. Existing research into these effects suggests that older adults may exhibit poorer performance due to an inability to adapt their decision bounds under speed emphasis (e.g., Forstmann et al., 2011; Starns & Ratcliff, 2010). This study was distinct from previous work in some key ways which allowed us to challenge these assertions. Firstly, this study used a points-based system, explicitly rewarding participants for adaptation to the regime. Here, it was shown that older adults achieved an equal number of points overall as younger adults. This indicates that any differences in behaviour did not come at a direct cost to their overall performance. This was further supported by modelling which indicated no difference between older and younger adults in their rate of evidence accumulation or drift rates. This is suggestive that not all differences between older and younger adults can be attributable to overall impairments with age (see Dully et al., 2018 for review), at least not on this simple two-choice contrast detection task.

To probe for differences in boundary adjustments between older and younger adults we applied a DDM. However, the standard DDM failed to capture key features of behaviour, such as fast and slow errors, for either group in either condition, despite allowing for flexibility through a range of inter-trial variability parameters. Exploration of the neural data indicated that pre-evidence boundary adjustments were present for both groups, alongside dynamic urgency represented in pre-evidence mu/beta and the CNV. Furthermore, pre-evidence decision formation was also observed in the build-up of the CPP. Pre-evidence mu/beta adjustments were used to constrain the neurally-informed model which allowed for both urgency and a period of pure-noise accumulation before the onset of the drift rate. Fitting this model at both the group and individual level suggested that older adults exhibited lower boundary levels than younger adults in the accuracy condition, with a far smaller degree of adaptation to their decision bound under speed emphasis. The model also revealed adaptation to dynamic

urgency across regimens in younger but not older adults. Individual model fitting suggested that this was driven by a subset of younger people who showed significant urgency adjustments under speed pressure.

Collectively these results point to a novel interpretation of the adaptations employed by older adults under speed emphasis. Specifically, older adults did not show an overall increased level of response caution across conditions as reported in previous studies (Forstmann et al., 2011; Starns & Ratcliff, 2010), but a less pronounced adjustment of these bounds and their urgency dynamics. While the extent to which these are voluntary strategic adjustments or represent an inability to adapt to task demands remains unclear, the fact that older adults were able to achieve an equal number of points on the task suggests that this may indeed be a voluntary strategy. Here, more consistent boundaries allow for more consistent responses across time in older adults. We show a novel insight into the speed-accuracy tradeoff in ageing revealed by neurally-constrained modelling that could not have been achieved through the traditional DDM alone.

5.2.2: The Role of Noise in Models of Perceptual Learning

Chapter Three examined the effects of perceptual learning on within-trial noise in a DDM. Existing psychophysical theories of perceptual learning suggest that it is at least partly driven by the suppression of internal noise in the perceptual system (Dosher & Lu, 1998, 1999; 2017). Sequential sampling models instead suggest that learning is driven by increases in the decision bound (Zhang & Rowe, 2014) and improved evidence accumulation through larger drift rates with training (Ivanov et al., 2022; Zhang & Rowe, 2014). Other modelling work has also suggested that some degree of learning may be driven by improved familiarity with task timings, instantiated by decreased non-decision time variability and reduced bounds with training (Liu & Watanabe, 2012; Petrov et al., 2011). However, as traditional DDMs fix within-trial noise as a scaling parameter, they have as yet been unable to test for changes to this parameter, as choosing a different parameter such as boundary, precludes their ability to investigate potential changes to within-trial noise in conjunction with the established changes to bound.

This study used neurally-informed modelling alongside a behavioural measure of internal noise (Tibber et al., 2014) to assess the degree to which within-trial noise changes as a function of learning. Here, learning was shown to level off after session two, with significant improvements to accuracy, alongside reductions to reaction time variability and a psychophysical estimate of internal noise (Manning et al., 2015; Tibber et al., 2014). Again, behavioural and neural data pointed to the necessary inclusion of urgency and some period of noisy accumulation. The inclusion of these parameters greatly improved fits, and this neurally-informed model indicated that a model with a fixed boundary but reducing within-trial noise across sessions best explained the behavioural data. Notably, model-simulated within-trial noise levels closely resembled behavioural measures of within-trial noise, suggesting a strong degree of overlap between the two values. Pre-evidence mu/beta was again used to neurally-constrain the model, suggesting non-significant but increasing boundaries across sessions, while constraining motor execution time via motor potentials. Behavioural measures of internal noise were also used as estimates of within-trial noise across sessions. While not outperforming the neurally-informed model with within-trial noise varying, the neurally-constrained model was a vital component of the study, as it allowed us to test for potential reductions in within-trial noise, whilst also allowing for the routinely evidenced adaptations to the boundary. Here it showed potential adaptations to dynamic urgency across sessions, revealing a novel insight into the possible learning of task demands and timings as a driver of visual-perceptual learning. Importantly, these neurally-constrained models successfully reconstructed neural markers of motor preparation and urgency. Overall, the modelling data points to the important role of within-trial noise reductions in internal noise, a feature which is not readily addressed in the standard DDM but can be investigated via neurally-informed and constrained models.

5.2.3: Novel Methods for Extracting Model Indices

One core assumption for neurally-informed modelling is that the signals we use to guide and constrain our models directly reflect the process of interest. As has been shown in previous work, however, isolating these signals is not always straightforward. They can be influenced by distinct neural processes with

coalesced scalp topographies (e.g., West et al., 2010; Kelly & O'Connell, 2013), and may present differently for different groups (e.g., van Dinteren et al., 2018; O'Connell, Balsters et al., 2012; Polich, 1997b). This is especially pertinent when we compare the CPPs of older and younger adults. Both fMRI and EEG data suggest a frontalisation of these signals with age (Davis et al., 2008; van Dinteren et al., 2014b), potentially reflecting age-related recruitment of frontal brain areas to compensate for impairments in early-stage visual processing (Reuter-Lorenz & Cappell, 2008). On the other hand, it had not yet been established definitively whether this frontalisation really reflects a change to the P3b/CPP itself or, alternatively, alterations to other signals that evolve on a similar time scale and overlap on the scalp (e.g., Debener et al., 2005; Kelly & O'Connell, 2013; West et al., 2010). The final experimental chapter of this thesis showed that the application of a CSD transform to neural data reduced the difference between older and younger adults' amplitude while retaining a frontalisation effect across a number of tasks. ICA was applied to extract a coherence and RT-dependent CPP component. Thus, rather than selecting components visually (e.g., Makeig et al., 1999; van Dinteren et al., 2018), or based on stimulus-locked peak amplitudes and latencies (e.g., Porcaro et al., 2019) as has been typical in the literature, a novel approach was applied in which ICs were selected based on their capacity to reflect key choice-relevant functional characteristics of the CPP. Analysis of this component showed no significant frontalisation in older adults, with a centroparietal focus in both groups, but a significantly lower response-locked amplitude in the older cohort. This indicates that not all age differences in CPP amplitude are due to overlapping signals, suggesting that the amplitude, but not the topography of the CPP changes with age. Through its extraction of an evidence and RT-dependent CPP component, this study also presented a novel technique for selecting a spatially independent CPP signal which may be applied in future work examining the CPP or other neural indices of decision-making such as the CNV.

5.3: Key Findings & Future Research

5.3.1: Neurally-Informed Modelling Assists Model Construction

An important feature of the neurally-informed models in both Chapters Two and Three was the inclusion of dynamic urgency i.e., a linearly collapsing bound. To date, there is significant debate as to whether the decision bound remains at a fixed level throughout the duration of a trial or collapses to reflect a ramping urgency to respond (Hawkins et al., 2015; Ratcliff et al., 2016). A hallmark of urgency in behavioural data is a large portion of slow errors. The standard DDM accounts for these errors with increased start-point variability. However, others have argued that urgency represents a very real facet of human decision-making which must be accounted for in the model (Kelly et al., 2021; Malhotra et al., 2018; Murphy et al., 2016; Steinemann et al., 2018). Here, when there are strict deadlines, participants must make a response based on lower levels of evidence to avoid missing the response deadline. In the absence of any deadline, participants may make responses based on less evidence in order to improve their "reward rate", increasing the proportion of correct responses made in a given time window (Bogacz et al., 2006). Even in tasks where committing to the response does not move the trial onto the next one any more quickly, urgency may be beneficial as once a response is chosen, participants can disengage from the stimulus, reducing the burden of energy and attention committed over time. Our neurally-informed models were able to test for the presence of dynamic urgency both in Chapters Two and Three. Behavioural features of both data sets indicated the presence of urgency, with substantially lower accuracy rates in later RT bins, reflective of a lower level of evidence accumulated by the time of response. However, even with the inclusion of starting point variability, the standard DDM was failing to capture a significant proportion of these slow errors. Turning to the neural data, the role of urgency was supported in both cases, with the CNV and mu/beta as markers of urgency beginning their negative descent in advance of any evidence presented. When this parameter was included in the model, it provided substantial improvements to model fits in both studies, with the new, neurally-informed models being able to more accurately capture the shape of the accuracy distributions over time. Here we see the benefit of neurally-

informed modelling in allowing us to incorporate a new parameter, not alone because it enables us to achieve better model fits, but because it is evidenced in the neural data.

These studies support the inclusion of dynamic urgency or at least strongly encourage the need to explore its inclusion on certain tasks. This is shown perhaps most clearly in the models of the speed-accuracy tradeoff in Chapter Two. Here the neurally-informed model was able to show not only the parameter's necessity, but that we may adapt our urgency dynamics differently across different age groups. Older adults were shown to not exhibit any strong adaptation of urgency across regimes, while younger adults did. Examining individual model fits showed that even within the younger adult group, urgency adaptation was only employed by a minority of individuals, suggesting it may be a specific strategy employed by a subset of younger individuals. Without neurally-informed models which supported the inclusion of urgency in the first instance, and then neural constraints which granted sufficient flexibility to explore a more complex model, these adaptations to urgency dynamics may not have been captured.

Collectively, it is clear that urgency plays a key role in at least certain kinds of perceptual decisions. The extent to which this role is limited to tasks such as the ones employed here, where there are strict deadlines, negative consequences for failing to reach these deadlines, and consistent trial lead-in times, is less well understood. That said, work from Devine (2019) has indicated that urgency dynamics may still play a role in tasks where evidence onset time cannot be predicted. The extent to which urgency affects real-world decisions made in ecologically valid contexts is also uncertain. In psychological experiments, participants complete trial after trial, becoming highly practised and learning the timings of the task, which could promote a role for urgency as a strategic adjustment (Hawkins et al., 2015). The extent to which urgency affects decisions made over a longer period of time, and without the added pressure of the knowledge of forthcoming consecutive decisions, remains unclear. Furthermore, in tasks of more cognitive complexity such as facial recognition or working memory, it may be in fact optimal to continue accumulating until entirely certain,

reducing the optimality of a lowering decision criterion across time. It is therefore difficult to directly ascertain whether urgency is a generalised feature of decision-making or represents a strategic adaptation to the highly constrained and ordered nature of traditional psychological experiments (Hawkins et al., 2015). However, while the inclusion of a parameter in a model may appear initially as an arbitrary attempt to include a mathematical construct which provides better fits, the work presented here suggests that failure to include a relevant parameter will not only prevent us from attaining a well-fitting model but potentially obscure or misattribute an observed behavioural difference to a completely different underlying mechanism of the decision process. Utilising neural signals as evidence for the inclusion of parameters is therefore of key importance, providing an additional line of support for urgency's inclusion or exclusion in a model for a given task.

5.3.2: Neural Constraints Can Address Scaling Issues

Sequential sampling models of perceptual decision-making are limited by their need to fix a scaling parameter. As discussed in Chapter One, this is a parameter that remains the same, at a fixed value, across experimental conditions and groups. This is beneficial as it allows the model to converge more easily, constraining the number of possible parameter combinations. By convention, within-trial noise, or the amount of Gaussian noise accumulated by the decision variable at every timestep is used as this scaling parameter, fixed across groups and conditions (Ratcliff et al., 2016). However, our ability to rely on a parameter which remains consistent across groups has been called into question by recent work. Psychophysical work has suggested that perceptual learning is driven at least partly by reductions in internal noise levels (Dosher & Lu, 1998, 1999; 2017), which may be equivalent to within-trial noise in a DDM. Similarly, different levels of noise are seen in different groups. For example, differences in neural and behavioural measures of noise are seen for older adults and people with ADHD, autism and schizophrenia (Alba et al., 2016; Chen et al., 2014; David et al., 2016; Park et al., 2017; Tibber et al., 2015). However, the effects of potential differences in within-trial noise on behaviour have thus far gone unexamined in a DDM framework. While the DDM does have the capacity to test for changes to

within-trial noise, fixing an alternative parameter such as boundary precludes its capacity to allow the new scaling parameter to change across sessions, conditions or groups.

Neurally-informed modelling presents an elegant solution to this. Parameters can be fixed based on neural evidence which shows some adaptation to a specific parameter (Kelly et al., 2021; Corbett et al., in press). In both Chapters Two and Three, pre-evidence mu/beta was used to constrain model bounds. This allowed each study to directly investigate differences in within-trial noise. Chapter Two probed differences in noise between older and younger adults, showing no appreciable difference in noise levels at the group level, suggesting an equally stable evidence accumulation signal in older age. Chapter Three went further, creating models which directly supported the role of within-trial noise in perceptual learning. Here with a fixed boundary across sessions, within-trial noise was shown to reduce. These reductions were directly related to a behavioural measure of internal noise. When using pre-evidence mu/beta to constrain the model, the best model fit was again one where noise varied by session. While not providing the best fit to the data, the neural constraints act as a data-driven scaling parameter, allowing for both adjustments to decision boundary and within-trial noise. Collectively, this study showed not only the value in allowing for noise changes across perceptual learning but also highlighted a potentially novel behavioural estimate for constraining these within-trial noise estimates; the EQN estimate of internal noise.

The neural constraints in each of these studies allowed for an investigation into changes in scaling parameters across or within groups. While promising, there is always room for improvement. For example, grand-average signals were used as neural constraints to improve the accuracy of the measures in both studies. Future experiments would benefit substantially from identifying individual-level constraints, allowing for analysis both at the individual and group level, allowing for more complex models and the accommodation of individual differences in traditionally used scaling parameters. For example, as no scaling parameters were available for the individual model fits applied in Chapter Two, we could not

assert whether the observed non-difference in noise levels between the older and younger adults within-trial noise levels was significantly different.

An individual-level scaling parameter may also be a useful way of scaling within-trial noise. A large body of work has pointed to a significant degree of individual and group differences in neural variability, potentially reflective of ongoing differences in within-trial noise (Dinstein et al., 2015; Gonen-Yaacovi et al., 2016; Hecker et al., 2022; Kumral et al., 2020; Sleimen-Malkoun et al., 2015; Yang et al., 2014). While the degree of association between neural measures of ongoing variability and model-estimated values of within-trial noise has not yet been ascertained, a measure of reliably capturing this parameter at an individual subject level would be beneficial. This would allow not only for constraints of within-trial noise across distinct groups but also a means of estimating trial-to-trial changes in noise over time within a session. One prime candidate for this would be the CPP. Given its established relationship with evidence accumulation (O'Connell et al., 2012), it is possible that a feature of ongoing variability in this signal may be directly relatable to within-trial noise. Future work may benefit from investigating the CPP as such a marker. Allowing for variable levels of stimulus noise, or through a double-pass procedure that presents the same stimulus multiple times (Awwad Shiekh Hasan et al., 2012), the effects of external stimulus noise could be controlled. From this, ongoing variability features of the CPP could be linked to variability in performance, giving a potential neural constraint for within-trial noise. Such an analysis was beyond the scope of this study but would prove beneficial to computational models of decision-making.

5.3.3: Neural Indices of Decision-Making are Refinable

One core assumption for neurally-informed modelling is that the signals we use to guide and constrain our models directly reflect the process of interest. As has been shown in previous work, however, these signals are not always reliable. They can be influenced by a range of potentially overlapping signals (e.g., Debener et al., 2005; Kelly & O'Connell, 2013; West et al., 2010), and may present differently across groups due to behaviourally-irrelevant reasons (e.g., Stacey et al., 2021, Rossiter et al., 2014, Antonakakis et al., 2020, Ribeiro et al.,

2022). However, this body of work indicated a range of techniques which can improve their reliability as neural metrics and constraints.

Mu/beta was normalised in Chapters Two and Three. In both studies, mu/beta was shown to desynchronise to the same threshold irrespective of RT. This suggested that overall differences in observed amplitude did not contribute to the decision process. Resultantly, behaviourally-irrelevant individual and group-level differences in mu/beta could be accounted for through normalisation, while maintaining the within-trial and between-group dynamics. Importantly it is not the raw amplitude of the signal that is of interest to neurally-informed modelling, but rather its pre-evidence and emergent dynamics across a trial. This normalisation is important, as raw signal amplitude may be taken as sufficient evidence for differences between groups on some underlying cognitive construct, such as pre-evidence starting point adjustments. However, when we correct for these differences, as was done in Chapter Two, we see mu/beta levels which are far more in line with model predictions. This lends further support for the value of applying normalisation to signals where possible, as they allow us to extract the important features of the neural signals, without inferring any information from potentially decision-irrelevant sources.

As outlined in Chapter One (section 1.4.1), the same cannot be achieved for signals such as the CPP. As some signals may not accumulate to a consistent bound across coherence or reaction time (Steinemann et al., 2018), they cannot be normalised by their boundary at response and instead require the exploration of other means for reducing decision-irrelevant contributions. An example of the importance of this is seen when we compare older and younger adults' P3b/CPP which shows lower parietal amplitudes in older adults and more frontal positivity (see van Dinteren et al., 2014b; and Chapter Four for detailed discussion). However, ERP signals such as the P3b/CPP may be affected by overlapping signals which emerge differently across age groups, thus potentially influencing the observed amplitude and build-up differences of the P3b/CPP in older adults (Anderer et al., 1996, 1998; Gajewski et al., 2018; Bertoli et al., 2005; Golob et al., 2001; Kuruvilla-Mathew et al., 2022; Polich, 1997a; Kolev et al., 2002; O'Connell, Balsters, et al., 2012). Furthermore, physiological changes in skull

conductance may affect how the signal appears on the scalp (Antonakakis et al., 2020), but not the underlying process at hand. Collectively these studies indicate a potential range of non-decision-related differences in older adults which may affect the older adult P3b/CPP, and thus the interpretations we make when comparing signals across groups.

Without a reliable means for removing decision-irrelevant group differences, it is difficult to ascertain the true effects of ageing on the processes of decision-making represented in the CPP. In Chapter Four, the application of a CSD transform reduced age differences in amplitude, but maintained a frontalisation effect, suggesting that some degree of differences in signal amplitude between older and younger adults may be driven by overlapping signals. This is an important thing to consider for future research, as previous work has indicated significant group differences in stimulus-locked peak amplitude without the application of these spatial filters. In this way, the effects of ageing on P3b/CPP amplitude may have been overstated to some degree.

This study then selected a functionally derived CPP component via independent component analysis (ICA) based on a set of predefined evidence and RT-defined difference in the build-up. However, Functional Source Separation (FSS) as applied by Porcaro and colleagues (2008, 2011, 2019) allows for a potentially more appropriate method. Here, the function constrains the isolation of components, meaning independent components can be extracted based on temporal independence and similarity on a specified metric, rather than grouping similar ICs after they have been isolated, as was done here. Thus far, only stimulus-locked P3b/CPP amplitude in a given time window has been used as a constraint, however, utilising a technique such as FFS to isolate a CPP component based on their response-locked functional differences in coherence and reaction time may have provided a more refined neural metric.

As the CNV has been shown to be a potentially useful marker of dynamic urgency (Devine, 2019), it presents a key tool for future work. However, as outlined in Chapter Four, it may suffer from the effects of overlapping signals in much the same way as the CPP and may have to some extent potentially influenced the

frontalisation effects observed in older adults. Extraction of a CNV-IC in the same way as was done for the CPP was not achievable for this study, as there was no manipulation which could be hypothesised to reasonably cause characteristic changes to the CNV signal. This would require a task designed to elicit distinct effects on the CNV such as a speed-accuracy trade-off manipulation. While beyond the scope of this study, the same approach could be applied to the data from Chapter Two, which investigated a speed-accuracy tradeoff in older adults. This study had the two requirements needed to elicit changes in both the CPP and the CNV: namely an evidence strength manipulation and a speed-accuracy manipulation. Application of the same ICA technique used in Chapter Four would allow for the extraction of both a CNV and CPP-IC profile. This would allow us to more conclusively investigate potential topographical shifts of the CNV, alongside allowing us to derive a spatially independent CNV component for use in neurally-constrained modelling. While beyond the scope of this work, this is an exciting next step which will be pursued in the future.

5.4: Limitations

While this body of work was successful in demonstrating the inherent value of neurally-informed modelling, it was not without its limitations. Primarily, the studies in this thesis may have been limited by a lack of statistical power. Chapters Two and Four had relatively large sample sizes (N >= 49 for both), with power analyses indicating they were sufficiently powered to detect small-to-medium effect sizes, Chapter Three only analysed data from 14 participants. Although this was a repeated-measures design, and as a result more robust and statistically powerful than a between-groups design, a post-hoc analysis indicated it may have only been sufficiently powered to detect medium-to-large effect sizes (f > 0.36). For example, a repeated-measures ANOVA of the effects of session on pre-evidence mu/beta amplitude while non-significant gave an effect size of f = 0.25. To detect an effect of this size a sample of 28 participants would have been needed. As such the study may have benefitted from a larger sample size. This may have been of particular importance for the model constraints. Additional data may have engendered more concrete and statistically significant differences in pre-evidence mu/beta and internal noise reductions, which may have provided more precise means for constraining the data. As the sample was smaller, the study may have been incapable of extracting some of these more subtle differences.

The studies in this thesis were also limited by the low number of possible constraints employed. While the models benefited from neural information which encouraged the inclusion of urgency and a period of early accumulation, the only constraints employed were pre-evidence mu/beta as a marker of starting-point adjustments, and motor time constraints via motor potentials. Additional neural constraints would potentially have furnished new results. For example, recent work has indicated the utility of visual N1s latencies, measured over the occipital cortex, as means for constraining the visual-encoding component of non-decision time (Ghaderi-Kangavari et al., 2022; Nunez et al., 2019). In Chapter Three, behavioural estimates of internal noise were introduced as a potentially viable constraint for within-trial noise. However, this required participants to complete an additional task. The shortened version of the EQN that was used (150 trials)

may also have resulted in less reliable estimates of internal noise. Future work may use a larger number of trials to increase reliability or explore neural indices of the CPP, such as signal entropy (Sleimen-Malkoun et al., 2015), or another metric of ongoing within-trial signal variability as a means of estimating within-trial noise parameters. As outlined in Chapter Four, ICA-derived CPP and CNV components may be a useful avenue for refining neural constraints, potentially providing more precise means of estimating neural constraints of urgency, drift rates or within-trial noise. The richer a roster of potential neural constraints, the more freedom we have in allowing other parameters to vary by condition, allowing us to investigate more nuanced and complex models without any cost to model comparison scores.

5.5: Concluding Remarks

The nascent field of neurally-informed modelling has a huge amount to offer the world of perceptual decision-making. As has been demonstrated throughout this thesis, neurally-informed modelling allows us to explore with more confidence novel parameters such as urgency and early accumulation. Furthermore, utilising neural constraints enables us to explore increasingly complex models, supported by neural data, to give new insights into perceptual decision-making. As a whole, this body of work not only presents some of these potential new insights but also presents some novel techniques for isolating the neural signatures of decision-making on which neurally-informed modelling rely. This field is burgeoning with possibility, and this work presents a significant advancement in our understanding of the potential applications and benefits of neurally-informed modelling. Through this, we can bring forth an increasingly nuanced understanding of how decisions are formed in the human brain.

References

Acerbi, L., & Ma, W. J. (2017). *Practical Bayesian Optimization for Model Fitting with Bayesian Adaptive Direct Search* (arXiv:1705.04405). arXiv. https://doi.org/10.48550/arXiv.1705.04405

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

Alba, G., Pereda, E., Mañas, S., Méndez, L. D., Duque, M. R., González, A., & González, J. J. (2016). The variability of EEG functional connectivity of young ADHD subjects in different resting states. *Clinical Neurophysiology*, *127*(2), 1321–1330. https://doi.org/10.1016/j.clinph.2015.09.134

Anderer, P., Pascual-Marqui, R. D., Semlitsch, H. V., & Saletu, B. (1998). Differential effects of normal aging on sources of standard N1, target N1 and target P300 auditory event-related brain potentials revealed by low resolution electromagnetic tomography (LORETA). *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *108*(2), 160–174. https://doi.org/10.1016/S0168-5597(97)00080-4

Anderer, P., Semlitsch, H. V., & Saletu, B. (1996). Multichannel auditory event-related brain potentials: Effects of normal aging on the scalp distribution of N1, P2, N2 and P300 latencies and amplitudes. *Electroencephalography and Clinical Neurophysiology*, *99*(5), 458–472. https://doi.org/10.1016/S0013-4694(96)96518-9

Antonakakis, M., Schrader, S., Aydin, Ü., Khan, A., Gross, J., Zervakis, M., Rampp, S., & Wolters, C. H. (2020). Inter-Subject Variability of Skull Conductivity and Thickness in Calibrated Realistic Head Models. *NeuroImage*, *223*, 117353. https://doi.org/10.1016/j.neuroimage.2020.117353

Awwad Shiekh Hasan, B., Joosten, E., & Neri, P. (2012). Estimation of internal noise using double passes: Does it matter how the second pass is delivered? *Vision Research*, *69*, 1–9. https://doi.org/10.1016/j.visres.2012.06.014

Bertoli, S., Smurzynski, J., & Probst, R. (2005). Effects of Age, Age-Related Hearing Loss, and Contralateral Cafeteria Noise on the Discrimination of Small Frequency Changes: Psychoacoustic and Electrophysiological Measures. *Journal of the Association for Research in Otolaryngology*, *6*(3), 207–222. https://doi.org/10.1007/s10162-005-5029-6

Boehm, U., van Maanen, L., Forstmann, B., & van Rijn, H. (2014). Trial-by-trial fluctuations in CNV amplitude reflect anticipatory adjustment of response caution. *NeuroImage*, *96*, 95–105. https://doi.org/10.1016/J.NEUROIMAGE.2014.03.063

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., & Cohen, J. D. (2006). The physics of optimal decision making: A formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, *113*(4), 700–765. https://doi.org/10.1037/0033-295X.113.4.700

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

Cabeza, R. (2002). Hemispheric asymmetry reduction in older adults: The HAROLD model. *Psychology and Aging*, *17*, 85–100. https://doi.org/10.1037/0882-7974.17.1.85

Chen, Y., Norton, D., & McBain, R. (2014). Effects of Domain-Specific Noise on Visual Motion Processing in Schizophrenia. *PLOS ONE*, *9*(6), e99031. https://doi.org/10.1371/journal.pone.0099031

Corbett, E., Martinez-Rodriguez, A., Judd, C., O'Connell, R.G., Kelly, S.P. (in press) Multiphasic Value Biases in Fast-Paced. *eLife*

David, N., Schneider, T. R., Peiker, I., Al-Jawahiri, R., Engel, A. K., & Milne, E. (2016). Variability of cortical oscillation patterns: A possible endophenotype in autism spectrum disorders? *Neuroscience and Biobehavioral Reviews*, *71*, 590–600. https://doi.org/10.1016/j.neubiorev.2016.09.031

Davis, S. W., Dennis, N. A., Daselaar, S. M., Fleck, M. S., & Cabeza, R. (2008). Que PASA? The posterior-anterior shift in aging. *Cerebral Cortex (New York, N.Y.: 1991)*, *18*(5), 1201–1209. https://doi.org/10.1093/cercor/bhm155

Debener, S., Makeig, S., Delorme, A., & Engel, A. K. (2005). What is novel in the novelty oddball paradigm? Functional significance of the novelty P3 event-related potential as revealed by independent component analysis. *Cognitive Brain Research*, *22*(3), 309–321. https://doi.org/10.1016/j.cogbrainres.2004.09.006

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

Desender, K., Boldt, A., Verguts, T., & Donner, T. H. (2019). Confidence predicts speed-accuracy tradeoff for subsequent decisions. *ELife*, *8*, e43499. https://doi.org/10.7554/eLife.43499

Devine, C. A. (2019). *An Electrophysiological Investigation into Temporal Factors in Human Perceptual Decision-Making* [Thesis, Trinity College Dublin. School of Psychology. Discipline of Psychology]. http://www.tara.tcd.ie/handle/2262/88733

Devine, C. A., Gaffney, C., Loughnane, G. M., Kelly, S. P., & O'Connell, R. G. (2019). The role of premature evidence accumulation in making difficult perceptual decisions under temporal uncertainty. *ELife*, *8*, e48526. https://doi.org/10.7554/eLife.48526

Diaz, J. A., Queirazza, F., & Philiastides, M. G. (2017). Perceptual learning alters post-sensory processing in human decision-making. *Nature Human Behaviour*, *1*(2), 1–9. https://doi.org/10.1038/s41562-016-0035

Diaz, M. T., Barrett, K. T., & Hogstrom, L. J. (2011). The influence of sentence novelty and figurativeness on brain activity. *Neuropsychologia*, *49*(3), 320–330. https://doi.org/10.1016/j.neuropsychologia.2010.12.004

Dinstein, I., Heeger, D. J., & Behrmann, M. (2015). Neural variability: Friend or foe? *Trends in Cognitive Sciences*, *19*(6), 322–328. https://doi.org/10.1016/j.tics.2015.04.005

Dockree, P. M., Kelly, S. P., Foxe, J. J., Reilly, R. B., & Robertson, I. H. (2007). Optimal sustained attention is linked to the spectral content of background EEG activity: Greater ongoing tonic alpha (~10 Hz) power supports successful phasic goal activation. *European Journal of*

Neuroscience, *25*(3), 900–907. https://doi.org/10.1111/j.1460-9568.2007.05324.x

Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: Rethinking the scaling problem. *Psychonomic Bulletin & Review*, *16*(6), 1129–1135. https://doi.org/10.3758/PBR.16.6.1129

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

Dosher, B. A., Jeter, P., Liu, J., & Lu, Z.-L. (2013). An integrated reweighting theory of perceptual learning. *Proceedings of the National Academy of Sciences*, *110*(33), 13678–13683. https://doi.org/10.1073/pnas.1312552110

Dosher, B. A., & Lu, Z. L. (1998). Perceptual learning reflects external noise filtering and internal noise reduction through channel reweighting. *Proceedings of the National Academy of Sciences of the United States of America*, *95*(23), 13988–13993. https://doi.org/10.1073/pnas.95.23.13988

Dosher, B. A., & Lu, Z.-L. (1999). Mechanisms of perceptual learning. *Vision Research*, *39*(19), 3197–3221. https://doi.org/10.1016/S0042-6989(99)00059-0

Dosher, B. A., & Lu, Z.-L. (2006). Level and mechanisms of perceptual learning: Learning first-order luminance and second-order texture objects. *Vision Research*, *46*(12), 1996–2007.

https://doi.org/10.1016/j.visres.2005.11.025

Dosher, B. A., & Lu, Z.-L. (2017). Visual Perceptual Learning and Models. *Annual Review of Vision Science*, *3*, 343–363. https://doi.org/10.1146/annurev-vision-102016-061249

Dully, J. (2020). *The Impact of Natural Aging on Perceptual Decision Making* [Thesis, Trinity College Dublin. School of Psychology. Discipline of Psychology]. http://www.tara.tcd.ie/handle/2262/94134

Dully, J., McGovern, D. P., & O'Connell, R. G. (2018). The impact of natural aging on computational and neural indices of perceptual decision making: A review. *Behavioural Brain Research*, *355*, 48–55. https://doi.org/10.1016/j.bbr.2018.02.001

ElShafei, H. A., Masson, R., Fakche, C., Fornoni, L., Moulin, A., Caclin, A., & Bidet-Caulet, A. (2022). Age-related differences in bottom-up and top-down attention: Insights from EEG and MEG. *European Journal of Neuroscience*, *55*(5), 1215–1231. https://doi.org/10.1111/ejn.15617

Fabiani, M., Friedman, D., & Cheng, J. C. (1998). Individual differences in P3 scalp distribution in older adults, and their relationship to frontal lobe function. *Psychophysiology*, *35*(6), 698–708.

Fahle, M. (2005). Perceptual learning: Specificity versus generalization. *Current Opinion in Neurobiology*, *15*(2), 154–160. https://doi.org/10.1016/j.conb.2005.03.010

Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, *9*(4), Article 4. https://doi.org/10.1038/nrn2258

Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible

statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, *39*, 175-191.

Finkel, D., Reynolds, C. a, McArdle, J. J., & Pedersen, N. L. (2007). Age changes in processing speed as a leading indicator of cognitive aging. *Psychology and Aging*, *22*(3), 558–568. https://doi.org/10.1037/0882-7974.22.3.558

Fischer, A. G., Nigbur, R., Klein, T. A., Danielmeier, C., & Ullsperger, M. (2018). Cortical beta power reflects decision dynamics and uncovers multiple facets of post-error adaptation. *Nature Communications*, *9*(1), Article 1. https://doi.org/10.1038/s41467-018-07456-8

Fjell, A. M., & Walhovd, K. B. (2004). Life-span changes in P3a. *Psychophysiology*, *41*(4), 575–583. https://doi.org/10.1111/j.1469-8986.2004.00177.x

Forstmann, B. U., Anwander, A., Schäfer, A., Neumann, J., Brown, S., Wagenmakers, E.-J., Bogacz, R., & Turner, R. (2010). Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(36), 15916–15920. https://doi.org/10.1073/pnas.1004932107

Forstmann, B. U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D. Y., Ridderinkhof, K. R., & Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. *Proceedings of the National Academy of Sciences*, *105*(45), 17538–17542. https://doi.org/10.1073/pnas.0805903105

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, B. U., Tittgemeyer, M., Wagenmakers, E.-J., Derrfuss, J., Imperati, D., & Brown, S. (2011). The Speed-Accuracy Tradeoff in the Elderly Brain: A Structural Model-Based Approach. *Journal of Neuroscience*, *31*(47), 17242–17249. https://doi.org/10.1523/JNEUROSCI.0309-11.2011

Frazier, P., & Yu, A. J. (2007). Sequential Hypothesis Testing under Stochastic Deadlines. *Advances in Neural Information Processing Systems*, *20*. https://proceedings.neurips.cc/paper/2007/hash/9c82c7143c102b71c593d98d96093fde-Abstract.html

Gajewski, P. D., Ferdinand, N. K., Kray, J., & Falkenstein, M. (2018). Understanding sources of adult age differences in task switching: Evidence from behavioral and ERP studies. *Neuroscience & Biobehavioral Reviews*, *92*, 255–275. https://doi.org/10.1016/j.neubiorev.2018.05.029

Ghaderi-Kangavari, A., Rad, J. A., Parand, K., & Nunez, M. D. (2022). Neuro-cognitive models of single-trial EEG measures describe latent effects of spatial attention during perceptual decision making. *Journal of Mathematical Psychology*, *111*, 102725. https://doi.org/10.1016/j.jmp.2022.102725

Gilbert, C. D., Sigman, M., & Crist, R. E. (2001). The Neural Basis of Perceptual Learning. *Neuron*, *31*(5), 681–697. https://doi.org/10.1016/S0896-6273(01)00424-X

Gold, J., Bennett, P. J., & Sekuler, A. B. (1999). Signal but not noise changes with perceptual learning. *Nature*, *402*(6758), Article 6758. https://doi.org/10.1038/46027

Golob, E. J., Miranda, G. G., Johnson, J. K., & Starr, A. (2001). Sensory cortical interactions in aging, mild cognitive impairment, and Alzheimer's disease. *Neurobiology of Aging*, *22*(5), 755–763. https://doi.org/10.1016/S0197-4580(01)00244-5

Gonen-Yaacovi, G., Arazi, A., Shahar, N., Karmon, A., Haar, S., Meiran, N., & Dinstein, I. (2016). Increased ongoing neural variability in ADHD. *Cortex*, *81*, 50–63. https://doi.org/10.1016/j.cortex.2016.04.010

Grady, C. (2012). The cognitive neuroscience of ageing. *Nature Reviews Neuroscience*, *13*(7), Article 7. https://doi.org/10.1038/nrn3256

Grady, C. L., Maisog, J. M., Horwitz, B., Ungerleider, L. G., Mentis, M. J., Salerno, J. A., Pietrini, P., Wagner, E., & Haxby, J. V. (1994). Age-related changes in cortical blood flow activation during visual processing of faces and location. *Journal of Neuroscience*, *14*(3), 1450–1462. https://doi.org/10.1523/JNEUROSCI.14-03-01450.1994

Groppe, D. M., Urbach, T. P., & Kutas, M. (2011). Mass univariate analysis of event-related brain potentials/fields I: A critical tutorial review. *Psychophysiology*, *48*(12), 1711–1725. https://doi.org/10.1111/j.1469-8986.2011.01273.x

Hanes, D. P., & Schall, J. D. (1996). Neural control of voluntary movement initiation. *Science (New York, N.Y.)*, *274*(5286), 427–430. https://doi.org/10.1126/science.274.5286.427

Hanks, T., Kiani, R., & Shadlen, M. N. (2014). A neural mechanism of speed-

accuracy tradeoff in macaque area LIP. *ELife*, *3*, e02260. https://doi.org/10.7554/eLife.02260

Hawkins, G. E., Forstmann, B. U., Wagenmakers, E.-J., Ratcliff, R., & Brown, S. D. (2015). Revisiting the Evidence for Collapsing Boundaries and Urgency Signals in Perceptual Decision-Making. *Journal of Neuroscience*, *35*(6), 2476–2484. https://doi.org/10.1523/JNEUROSCI.2410-14.2015

Hecker, L., Wilson, M., Tebartz van Elst, L., & Kornmeier, J. (2022). Altered EEG variability on different time scales in participants with autism spectrum disorder: An exploratory study. *Scientific Reports*, *12*(1), Article 1. https://doi.org/10.1038/s41598-022-17304-x

Heitz, R. P. (2014). The speed-accuracy tradeoff: History, physiology, methodology, and behavior. *Frontiers in Neuroscience*, *8*. https://www.frontiersin.org/articles/10.3389/fnins.2014.00150

Heitz, R. P., & Schall, J. D. (2012). Neural Mechanisms of Speed-Accuracy Tradeoff. *Neuron*, *76*(3), 616–628. https://doi.org/10.1016/j.neuron.2012.08.030

Huettel, S. A., & McCarthy, G. (2004). What is odd in the oddball task?: Prefrontal cortex is activated by dynamic changes in response strategy. *Neuropsychologia*, *42*(3), 379–386. https://doi.org/10.1016/j.neuropsychologia.2003.07.009

Huk, A. C., & Shadlen, M. N. (2005). Neural activity in macaque parietal cortex reflects temporal integration of visual motion signals during perceptual decision making. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *25*(45), 10420–10436.

https://doi.org/10.1523/JNEUROSCI.4684-04.2005

Huk, A., & Meister, M. (2012). Neural correlates and neural computations in posterior parietal cortex during perceptual decision-making. *Frontiers in Integrative Neuroscience*, *6*. https://www.frontiersin.org/articles/10.3389/fnint.2012.00086

Ivanov, V., Manenti, G., Sudmann, S. S., Kagan, I., & Schwiedrzik, C. M. (2022). *Decision-making processes in perceptual learning depend on effectors* (p. 2022.06.29.498152). bioRxiv. https://doi.org/10.1101/2022.06.29.498152

Jervis, B. W., Bigan, C., & Besleaga, M. (2020). New-Onset Alzheimer's Disease and Normal Subjects 100% Separated Statistically by P300 and ICA. *American Journal of Alzheimer's Disease & Other Dementias®*, *35*, 1533317520935675. https://doi.org/10.1177/1533317520935675

Jia, K., Xue, X., Lee, J.-H., Fang, F., Zhang, J., & Li, S. (2018). Visual perceptual learning modulates decision network in the human brain: The evidence from psychophysics, modeling, and functional magnetic resonance imaging. *Journal of Vision*, *18*(12), 9. https://doi.org/10.1167/18.12.9

Jung, T.-P., Humphries, C., Lee, T.-W., Makeig, S., McKeown, M., Iragui, V., & Sejnowski, T. J. (1997). Extended ICA Removes Artifacts from Electroencephalographic Recordings. *Advances in Neural Information Processing Systems*, *10*. https://proceedings.neurips.cc/paper/1997/hash/674bfc5f6b72706fb769f5e93667bd23-Abstract.html

Karalunas, S. L., Geurts, H. M., Konrad, K., Bender, S., & Nigg, J. T. (2014).

Reaction time variability in ADHD and autism spectrum disorders: Measurement and mechanisms of a proposed trans-diagnostic phenotype. *Journal of Child Psychology and Psychiatry*, *55*(6), 685–710. https://doi.org/10.1111/jcpp.12217

Karalunas, S. L., & Huang-Pollock, C. L. (2013). Integrating impairments in reaction time and executive function using a diffusion model framework. *Journal of Abnormal Child Psychology*, *41*(5), 837–850. https://doi.org/10.1007/s10802-013-9715-2

Kayser, J., & Tenke, C. E. (2015). On the benefits of using surface Laplacian (Current Source Density) methodology in electrophysiology. *International Journal of Psychophysiology : Official Journal of the International Organization of Psychophysiology*, *97*(3), 171–173. https://doi.org/10.1016/j.ijpsycho.2015.06.001

Kelly, S. P., Corbett, E. A., & O'Connell, R. G. (2019). *Multifaceted adaptation of the neural decision process with prior knowledge of time constraints and stimulus probability* (p. 715318). bioRxiv. https://doi.org/10.1101/715318

Kelly, S. P., Corbett, E. A., & O'Connell, R. G. (2021). Neurocomputational mechanisms of prior-informed perceptual decision-making in humans. *Nature Human Behaviour*, *5*(4), Article 4. https://doi.org/10.1038/s41562-020-00967-9

Kelly, S. P., & O'Connell, R. G. (2013). Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. *Journal of Neuroscience*, *33*(50), 19434–19441. https://doi.org/10.1523/JNEUROSCI.3355-13.2013

Kim, M., Lee, T. H., Kim, J.-H., Hong, H., Lee, T. Y., Lee, Y., Salisbury, D. F., & Kwon, J. S. (2018). Decomposing P300 into correlates of genetic risk and current symptoms in schizophrenia: An inter-trial variability analysis. *Schizophrenia Research*, *192*, 232–239. https://doi.org/10.1016/j.schres.2017.04.001

Kolev, V., Yordanova, J., Basar-Eroglu, C., & Basar, E. (2002). Age effects on visual EEG responses reveal distinct frontal alpha networks. *Clinical Neurophysiology*, *113*(6), 901–910. https://doi.org/10.1016/S1388-2457(02)00106-2

Kononowicz, T. W., & Penney, T. B. (2016). The contingent negative variation (CNV): Timing isn't everything. *Current Opinion in Behavioral Sciences*, *8*, 231–237. https://doi.org/10.1016/j.cobeha.2016.02.022

Krishnaswamy, P., Obregon-Henao, G., Ahveninen, J., Khan, S., Babadi, B., Iglesias, J. E., Hämäläinen, M. S., & Purdon, P. L. (2017). Sparsity enables estimation of both subcortical and cortical activity from MEG and EEG. *Proceedings of the National Academy of Sciences*, *114*(48), E10465–E10474. https://doi.org/10.1073/pnas.1705414114

Kumral, D., Şansal, F., Cesnaite, E., Mahjoory, K., Al, E., Gaebler, M., Nikulin, V. V., & Villringer, A. (2020). BOLD and EEG signal variability at rest differently relate to aging in the human brain. *NeuroImage*, *207*, 116373. https://doi.org/10.1016/j.neuroimage.2019.116373

Kuruvilla-Mathew, A., Thorne, P. R., & Purdy, S. C. (2022). Effects of aging on neural processing during an active listening task. *PLOS ONE*, *17*(9), e0273304. https://doi.org/10.1371/journal.pone.0273304

Lachaux, J. P., Rudrauf, D., & Kahane, P. (2003). Intracranial EEG and human

brain mapping. *Journal of Physiology-Paris*, *97*(4), 613–628. https://doi.org/10.1016/j.jphysparis.2004.01.018

Law, C.-T., & Gold, J. I. (2008). Neural correlates of perceptual learning in a sensory-motor, but not a sensory, cortical area. *Nature Neuroscience*, *11*(4), Article 4. https://doi.org/10.1038/nn2070

Levi, D. M., Klein, S. A., & Chen, I. (2005). What is the signal in noise? *Vision Research*, *45*(14), 1835–1846. https://doi.org/10.1016/j.visres.2005.01.020

Liu, C. C., & Watanabe, T. (2012). Accounting for speed–accuracy tradeoff in perceptual learning. *Vision Research*, *61*, 107–114. https://doi.org/10.1016/j.visres.2011.09.007

Lu, Z.-L., & Dosher, B. A. (2004). Perceptual learning retunes the perceptual template in foveal orientation identification. *Journal of Vision*, *4*(1), 5. https://doi.org/10.1167/4.1.5

Lu, Z.-L., & Dosher, B. A. (2008). Characterizing observers using external noise and observer models: Assessing internal representations with external noise. *Psychological Review*, *115*, 44–82. https://doi.org/10.1037/0033-295X.115.1.44

Lynn, S. K., & Barrett, L. F. (2014). “Utilizing” Signal Detection Theory. *Psychological Science*, *25*(9), 1663–1673. https://doi.org/10.1177/0956797614541991

Makeig, S., Westerfield, M., Jung, T.-P., Covington, J., Townsend, J., Sejnowski, T. J., & Courchesne, E. (1999). Functionally Independent Components of the Late Positive Event-Related Potential during Visual Spatial Attention. *Journal of Neuroscience*, *19*(7), 2665–2680.

https://doi.org/10.1523/JNEUROSCI.19-07-02665.1999

Malhotra, G., Leslie, D. S., Ludwig, C. J. H., & Bogacz, R. (2018). Time-varying decision boundaries: Insights from optimality analysis. *Psychonomic Bulletin & Review*, *25*(3), 971–996. https://doi.org/10.3758/s13423-017-1340-6

Maniglia, M., & Seitz, A. R. (2018). Towards a whole brain model of Perceptual Learning. *Current Opinion in Behavioral Sciences*, *20*, 47–55. https://doi.org/10.1016/j.cobeha.2017.10.004

Manning, C., Tibber, M. S., Charman, T., Dakin, S. C., & Pellicano, E. (2015). Enhanced Integration of Motion Information in Children With Autism. *Journal of Neuroscience*, *35*(18), 6979–6986. https://doi.org/10.1523/JNEUROSCI.4645-14.2015

Marriott Haresign, I., Phillips, E., Whitehorn, M., Noreika, V., Jones, E. J. H., Leong, V., & Wass, S. V. (2021). Automatic classification of ICA components from infant EEG using MARA. *Developmental Cognitive Neuroscience*, *52*, 101024. https://doi.org/10.1016/j.dcn.2021.101024

McCarthy, G., & Donchin, E. (1981). A Metric for Thought: A Comparison of P300 Latency and Reaction Time. *Science*, *211*(4477), 77–80.

McCarthy, P., Benuskova, L., & Franz, E. A. (2014). The age-related posterior-anterior shift as revealed by voxelwise analysis of functional brain networks. *Frontiers in Aging Neuroscience*, *6*, 301. https://doi.org/10.3389/fnagi.2014.00301

McGovern, D. P., Astle, A. T., Clavin, S. L., & Newell, F. N. (2016). Task-specific transfer of perceptual learning across sensory modalities. *Current Biology*, *26*(1), R20–R21.

https://doi.org/10.1016/j.cub.2015.11.048

McGovern, D. P., Hayes, A., Kelly, S. P., & O'Connell, R. G. (2018). Reconciling age-related changes in behavioural and neural indices of human perceptual decision-making. *Nature Human Behaviour*, *2*(12), 955–966. https://doi.org/10.1038/s41562-018-0465-6

Milne, E. (2011). Increased intra-participant variability in children with autistic spectrum disorders: Evidence from single-trial analysis of evoked EEG. *Frontiers in Psychology*, *2*, 51. https://doi.org/10.3389/fpsyg.2011.00051

Morcom, A. M., & Henson, R. N. A. (2018). Increased Prefrontal Activity with Aging Reflects Nonspecific Neural Responses Rather than Compensation. *Journal of Neuroscience*, *38*(33), 7303–7313. https://doi.org/10.1523/JNEUROSCI.1701-17.2018

Murphy, P. R., Boonstra, E., & Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. *Nature Communications*, *7*(1), Article 1. https://doi.org/10.1038/ncomms13526

Muthukumaraswamy, S. D., Myers, J. F. M., Wilson, S. J., Nutt, D. J., Lingford-Hughes, A., Singh, K. D., & Hamandi, K. (2013). The effects of elevated endogenous GABA levels on movement-related network oscillations. *NeuroImage*, *66*, 36–41. https://doi.org/10.1016/j.neuroimage.2012.10.054

Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus—Norepinephrine system. *Psychological Bulletin*, *131*(4), 510–532. https://doi.org/10.1037/0033-2909.131.4.510

Nunez, M. D., Gosai, A., Vandekerckhove, J., & Srinivasan, R. (2019). The

latency of a visual evoked potential tracks the onset of decision making. *NeuroImage*, *197*, 93–108. https://doi.org/10.1016/j.neuroimage.2019.04.052

O'Connell, R. G., Balsters, J. H., Kilcullen, S. M., Campbell, W., Bokde, A. W., Lai, R., Upton, N., & Robertson, I. H. (2012). A simultaneous ERP/fMRI investigation of the P300 aging effect. *Neurobiology of Aging*, *33*(10), 2448–2461. https://doi.org/10.1016/j.neurobiolaging.2011.12.021

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

O'Connell, R. G., & Kelly, S. P. (2021). Neurophysiology of Human Perceptual Decision-Making. *Annual Review of Neuroscience*, *44*(1), 495–516. https://doi.org/10.1146/annurev-neuro-092019-100200

O'Connell, R. G., Shadlen, M. N., Wong-Lin, K. F., & Kelly, S. P. (2018). Bridging Neural and Computational Viewpoints on Perceptual Decision-Making. *Trends in Neurosciences*, *41*(11), 838–852. https://doi.org/10.1016/j.tins.2018.06.005

Onton, J., Westerfield, M., Townsend, J., & Makeig, S. (2006). Imaging human EEG dynamics using independent component analysis. *Neuroscience & Biobehavioral Reviews*, *30*(6), 808–822. https://doi.org/10.1016/j.neubiorev.2006.06.007

Oostenveld, R., Fries, P., Maris, E., & Schoffelen, J.-M. (2010). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. *Computational Intelligence and Neuroscience*,

2011, e156869. https://doi.org/10.1155/2011/156869

Pardhan, S., Gilchrist, J., Elliott, D. B., & Beh, G. K. (1996). A comparison of sampling efficiency and internal noise level in young and old subjects. *Vision Research*, *36*(11), 1641–1648. https://doi.org/10.1016/0042-6989(95)00214-6

Park, W. J., Schauder, K. B., Zhang, R., Bennetto, L., & Tadin, D. (2017). High internal noise and poor external noise filtering characterize perception in autism spectrum disorder. *Scientific Reports*, *7*(1), 17584. https://doi.org/10.1038/s41598-017-17676-5

Petrov, A. A., Van Horn, N. M., & Ratcliff, R. (2011). Dissociable perceptual-learning mechanisms revealed by diffusion-model analysis. *Psychonomic Bulletin & Review*, *18*(3), 490–497. https://doi.org/10.3758/s13423-011-0079-8

Pion-Tonachini, L., Kreutz-Delgado, K., & Makeig, S. (2019). ICLabel: An automated electroencephalographic independent component classifier, dataset, and website. *NeuroImage*, *198*, 181–197. https://doi.org/10.1016/j.neuroimage.2019.05.026

Pirrone, A., Dickinson, A., Gomez, R., Stafford, T., & Milne, E. (2017). Understanding perceptual judgment in autism spectrum disorder using the drift diffusion model. *Neuropsychology*, *31*, 173–180. https://doi.org/10.1037/neu0000320

Polich, J. (1997a). EEG and ERP assessment of normal aging. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *104*(3), 244–256. https://doi.org/10.1016/S0168-5597(97)96139-6

Polich, J. (1997b). On the relationship between EEG and P300: Individual differences, aging, and ultradian rhythms. *International Journal of Psychophysiology*, *26*(1), 299–317. https://doi.org/10.1016/S0167-8760(97)00772-1

Polich, J. (2007). Updating P300: An Integrative Theory of P3a and P3b. *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *118*(10), 2128–2148. https://doi.org/10.1016/j.clinph.2007.04.019

Porcaro, C., Balsters, J. H., Mantini, D., Robertson, I. H., & Wenderoth, N. (2019). P3b amplitude as a signature of cognitive decline in the older population: An EEG study enhanced by Functional Source Separation. *NeuroImage*, *184*, 535–546. https://doi.org/10.1016/j.neuroimage.2018.09.057

Ratcliff, R., Cherian, A., & Segraves, M. (2003). A comparison of macaque behavior and superior colliculus neuronal activity to predictions from models of two-choice decisions. *Journal of Neurophysiology*, *90*(3), 1392–1407. https://doi.org/10.1152/jn.01049.2002

Ratcliff, R., Hasegawa, Y. T., Hasegawa, R. P., Smith, P. L., & Segraves, M. A. (2007). Dual Diffusion Model for Single-Cell Recording Data From the Superior Colliculus in a Brightness-Discrimination Task. *Journal of Neurophysiology*, *97*(2), 1756–1774. https://doi.org/10.1152/jn.00393.2006

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., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, *111*(2), 333–367. https://doi.org/10.1037/0033-295X.111.2.333

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., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, *9*(3), 438–481. https://doi.org/10.3758/BF03196302

Ratcliff, R., Voskuilen, C., & McKoon, G. (2018). Internal and external sources of variability in perceptual decision-making. *Psychological Review*, *125*, 33–46. https://doi.org/10.1037/rev0000080

Reuter-Lorenz, P. A., & Cappell, K. A. (2008). Neurocognitive Aging and the Compensation Hypothesis. *Current Directions in Psychological Science*, *17*(3), 177–182. https://doi.org/10.1111/j.1467-8721.2008.00570.x

Ribeiro, M.J., & Castelo-Branco, M. (2022). Slow fluctuations in ongoing brain activity decrease in amplitude with ageing yet their impact on task-related evoked responses is dissociable from behavior. *ELife*, *11*, e75722. https://doi.org/10.7554/eLife.75722

Ribeiro, M. J., Paiva, J. S., & Castelo-Branco, M. (2016). Spontaneous Fluctuations in Sensory Processing Predict Within-Subject Reaction Time Variability. *Frontiers in Human Neuroscience*, *10*, 200. https://doi.org/10.3389/fnhum.2016.00200

Richardson, C., Bucks, R. S., & Hogan, A. M. (2011). Effects of aging on

habituation to novelty: An ERP study. *International Journal of Psychophysiology*, *79*(2), 97–105. https://doi.org/10.1016/j.ijpsycho.2010.09.007

Rogge, J., Jocham, G., & Ullsperger, M. (2022). Motor cortical signals reflecting decision making and action preparation. *NeuroImage*, *263*, 119667. https://doi.org/10.1016/j.neuroimage.2022.119667

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), e9308. https://doi.org/10.1371/journal.pone.0009308

Rossiter, H. E., Davis, E. M., Clark, E. V., Boudrias, M. H., & Ward, N. S. (2014). Beta oscillations reflect changes in motor cortex inhibition in healthy ageing. *NeuroImage*, *91*, 360–365. https://doi.org/10.1016/J.NEUROIMAGE.2014.01.012

Salthouse, T. (1996). The processing-speed theory of adult age differences in cognition. *Psychological Review*, *103*(3), 403–428. https://doi.org/10.1037/0033-295X.103.3.403

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

Shadlen, M. N., & Newsome, W. T. (1996). Motion perception: Seeing and deciding. *Proceedings of the National Academy of Sciences*, *93*(2), 628–633. https://doi.org/10.1073/pnas.93.2.628

Sleimen-Malkoun, R., Perdikis, D., Müller, V., Blanc, J.-L., Huys, R., Temprado, J.-J., & Jirsa, V. K. (2015). Brain Dynamics of Aging: Multiscale Variability of EEG Signals at Rest and during an Auditory Oddball Task,,.

ENeuro, *2*(3), ENEURO.0067-14.2015. https://doi.org/10.1523/ENEURO.0067-14.2015

Song, Y., Ding, Y., Fan, S., Qu, Z., Xu, L., Lu, C., & Peng, D. (2005). Neural substrates of visual perceptual learning of simple and complex stimuli. *Clinical Neurophysiology*, *116*(3), 632–639. https://doi.org/10.1016/j.clinph.2004.09.019

Stacey, J. E., Crook-Rumsey, M., Sumich, A., Howard, C. J., Crawford, T., Livne, K., Lenzoni, S., & Badham, S. (2021). Age differences in resting state EEG and their relation to eye movements and cognitive performance. *Neuropsychologia*, *157*, 107887. https://doi.org/10.1016/j.neuropsychologia.2021.107887

Starns, J. J., & Ratcliff, R. (2010). The effects of aging on the speed-accuracy compromise: Boundary optimality in the diffusion model. *Psychology and Aging*, *25*(2), 377–390. https://doi.org/10.1037/a0018022

Starns, J. J., & Ratcliff, R. (2012). Age-related differences in diffusion model boundary optimality with both trial-limited and time-limited tasks. *Psychonomic Bulletin and Review*, *19*(1), 139–145. https://doi.org/10.3758/s13423-011-0189-3

Steinemann, N. A., O'Connell, R. G., & Kelly, S. P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. *Nature Communications*, *9*(1), 3627. https://doi.org/10.1038/s41467-018-06117-0

Swick, D., Ashley, V., & Turken, A. U. (2008). Left inferior frontal gyrus is critical for response inhibition. *BMC Neuroscience*, *9*(1), 102. https://doi.org/10.1186/1471-2202-9-102

Tenke, C. E., & Kayser, J. (2005). Reference-free quantification of EEG spectra: Combining current source density (CSD) and frequency principal components analysis (fPCA). *Clinical Neurophysiology*, *116*(12), 2826–2846. https://doi.org/10.1016/j.clinph.2005.08.007

Tibber, M. S., Anderson, E. J., Bobin, T., Carlin, P., Shergill, S. S., & Dakin, S. C. (2015). Local and global limits on visual processing in schizophrenia. *PLoS ONE*, *10*(2), e0117951. https://doi.org/10.1371/journal.pone.0117951

Tibber, M. S., Kelly, M. G., Jansari, A., Dakin, S. C., & Shepherd, A. J. (2014). An Inability to Exclude Visual Noise in Migraine. *Investigative Ophthalmology & Visual Science*, *55*(4), 2539–2546. https://doi.org/10.1167/iovs.14-13877

Tsolaki, A., Kosmidou, V., Hadjileontiadis, L., Kompatsiaris, I. (Yiannis), & Tsolaki, M. (2015). Brain source localization of MMN, P300 and N400: Aging and gender differences. *Brain Research*, *1603*, 32–49. https://doi.org/10.1016/j.brainres.2014.10.004

Turner, B. M., Forstmann, B. U., Wagenmakers, E.-J., Brown, S. D., Sederberg, P. B., & Steyvers, M. (2013). A Bayesian framework for simultaneously modeling neural and behavioral data. *NeuroImage*, *72*, 193–206. https://doi.org/10.1016/j.neuroimage.2013.01.048

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

Van de Vondervoort, J. W., & Hamlin, J. K. (2016). Evidence for Intuitive

Morality: Preverbal Infants Make Sociomoral Evaluations. *Child Development Perspectives*, *10*(3), 143–148. https://doi.org/10.1111/cdep.12175

van den Broek, S. P., Reinders, F., Donderwinkel, M., & Peters, M. J. (1998). Volume conduction effects in EEG and MEG. *Electroencephalography and Clinical Neurophysiology*, *106*(6), 522–534. https://doi.org/10.1016/S0013-4694(97)00147-8

van Dinteren, R., Arns, M., Jongsma, M. L. A., & Kessels, R. P. C. (2014a). Combined frontal and parietal P300 amplitudes indicate compensated cognitive processing across the lifespan. *Frontiers in Aging Neuroscience*, *0*. https://doi.org/10.3389/fnagi.2014.00294

van Dinteren, R., Arns, M., Jongsma, M. L. A., & Kessels, R. P. C. (2014b). P300 Development across the Lifespan: A Systematic Review and Meta-Analysis. *PLOS ONE*, *9*(2), e87347. https://doi.org/10.1371/journal.pone.0087347

van Dinteren, R., Huster, R. J., Jongsma, M. L. A., Kessels, R. P. C., & Arns, M. (2018). Differences in Cortical Sources of the Event-Related P3 Potential Between Young and Old Participants Indicate Frontal Compensation. *Brain Topography*, *31*(1), 35–46. https://doi.org/10.1007/S10548-016-0542-Y/FIGURES/7

Verleger, R. (2020). Effects of relevance and response frequency on P3b amplitudes: Review of findings and comparison of hypotheses about the process reflected by P3b. *Psychophysiology*, *57*(7), e13542. https://doi.org/10.1111/psyp.13542

Vidal, F., Burle, B., Spieser, L., Carbonnell, L., Meckler, C., Casini, L., &

Hasbroucq, T. (2015). Linking EEG signals, brain functions and mental operations: Advantages of the Laplacian transformation. *International Journal of Psychophysiology*, *97*(3), 221–232. https://doi.org/10.1016/j.ijpsycho.2015.04.022

Vilidaite, G., & Baker, D. H. (2017). Individual differences in internal noise are consistent across two measurement techniques. *Vision Research*, *141*, 30–39. https://doi.org/10.1016/j.visres.2016.10.008

Voss, A., Lerche, V., Mertens, U., & Voss, J. (2019). Sequential sampling models with variable boundaries and non-normal noise: A comparison of six models. *Psychonomic Bulletin & Review*, 1–20. https://doi.org/10.3758/s13423-018-1560-4

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory and Cognition*, *32*(7), 1206–1220. https://doi.org/10.3758/BF03196893

Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods*, *39*(4), 767–775. https://doi.org/10.3758/BF03192967

Voytek, B., Kramer, M. A., Case, J., Lepage, K. Q., Tempesta, Z. R., Knight, R. T., & Gazzaley, A. (2015). Age-Related Changes in 1/f Neural Electrophysiological Noise. *Journal of Neuroscience*, *35*(38), 13257–13265. https://doi.org/10.1523/JNEUROSCI.2332-14.2015

Vrieze, S. I. (2012). Model selection and psychological theory: A discussion of the differences between the Akaike information criterion (AIC) and the Bayesian information criterion (BIC). *Psychological Methods*, *17*(2), 228–243. https://doi.org/10.1037/a0027127

Wagenmakers, E.-J., & Farrell, S. (2004). AIC model selection using Akaike weights. *Psychonomic Bulletin & Review*, *11*(1), 192–196. https://doi.org/10.3758/BF03206482

Walhovd, K. B., & Fjell, A. M. (2001). Two- and three-stimuli auditory oddball ERP tasks and neuropsychological measures in aging. *NeuroReport*, *12*(14), 3149–3153.

West, R., Schwarb, H., & Johnson, B. N. (2010). The influence of age and individual differences in executive function on stimulus processing in the oddball task. *Cortex*, *46*(4), 550–563. https://doi.org/10.1016/j.cortex.2009.08.001

Wiecki, T. V., Sofer, I., & Frank, M. J. (2013). HDDM: Hierarchical bayesian estimation of the drift-diffusion model in Python. *Frontiers in Neuroinformatics*, *7*(JULY 2013). https://doi.org/10.3389/fninf.2013.00014

Woltering, S., Jung, J., Liu, Z., & Tannock, R. (2012). Resting state EEG oscillatory power differences in ADHD college students and their peers. *Behavioral and Brain Functions*, *8*(1), 60. https://doi.org/10.1186/1744-9081-8-60

Wright, B. A., & Sabin, A. T. (2007). Perceptual learning: How much daily training is enough? *Experimental Brain Research*, *180*(4), 727–736. https://doi.org/10.1007/s00221-007-0898-z

Yamaguchi, S., Hale, L. A., D'Esposito, M., & Knight, R. T. (2004). Rapid Prefrontal-Hippocampal Habituation to Novel Events. *Journal of Neuroscience*, *24*(23), 5356–5363. https://doi.org/10.1523/JNEUROSCI.4587-03.2004

Yang, G. J., Murray, J. D., Repovs, G., Cole, M. W., Savic, A., Glasser, M. F., Pittenger, C., Krystal, J. H., Wang, X.-J., Pearlson, G. D., Glahn, D. C., & Anticevic, A. (2014). Altered global brain signal in schizophrenia. *Proceedings of the National Academy of Sciences*, *111*(20), 7438–7443. https://doi.org/10.1073/pnas.1405289111

Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, *8*(8 APR), 69. https://doi.org/10.3389/fnins.2014.00069

Appendices

Chapter Four: Additional Statistics

Contrast Change Detection

Stimulus-Locked Peak Amplitude and Latency

Analysis of the effect of localisation on stimulus-locked peak amplitude showed no significant effect of Age (p = 0.118), Location (p = 0.205), or Age by Location interaction (p = 0.126). Analysis of the effect of Age and Location on the latency of these peaks showed a significant effect of Location ($F_{(2, 76)}$ = 4.616, p = 0.011), but not of Age (p = 0.252) or any Interaction (p = 0.124), suggesting no difference in latency across age groups at any location. Trends indicated a general linear trend, with faster latencies at frontal electrodes and central sites, and slower latencies over the parietal cortex.

Non-CSD'd Data Results

Analysis of this data indicated a significant main effect of Location ($F_{(1.64,76)}$ = 86.42, p < 0.001), Age ($F_{(1,76)}$ = 9.003, p = 0.003) on response-locked amplitude. There was a significant Age by Location interaction ($F_{(1.64,64)}$ = 20.41, p < 0.001). Post-hoc analyses indicated significant age-differences trends at all locations, with a significantly less positive in younger adults over frontal electrodes (t = -4.288, p <0.001), but a more positive amplitude over central (t = 4.332, p<0.001), and parietal electrodes (t = 5.61, p <0.001).

Analysis of response-locked slopes indicated a significant main effect of Location ($F_{(1.44,76)}$ = 23, p < 0.001), Age ($F_{(1,76)}$ = 6.08, p < 0.001) and Age by Location interaction ($F_{(1.44,76)}$ = 22.31, p < 0.001). Post-hoc analyses indicated significant age-differences trends at all locations, with a significantly less positive slope in younger adults over frontal electrodes (t = -4.685, p = 0.007), but a more positive slope over central (t = 3.232, p = 0.01), and parietal electrodes (t = 4.9, p < 0.001).

Random-Dot Motion

Stimulus-Locked Peak Amplitude and Latency:

Analysis of the effect of Age and Location on the stimulus-locked peak amplitude of these peaks showed a significant effect of Location ($F_{(2, 66)}$ = 46.82, p < 0.001), Age ($F_{(2, 66)}$ = 33.84, p < 0.001) and Coherence ($F_{(2, 66)}$ = 85.19, p < 0.001). There was no Coherence by Age interaction (p = 0.196) or Location-Coherence Age interaction (p = 0.324). There was a significant Location by Coherence Interaction ($F_{(2, 66)}$ = 9.797, p < 0.001). Post-hoc tests indicated that peak amplitudes were significantly lower for low coherence over frontal electrodes (t = -5.266, p < 0.001), and parietal electrodes (t = -7.375, p < 0.001), but not significantly different over central electrodes (p = 0.479). There was a significant Location by Age Interaction ($F_{(2, 66)}$ = 15.11, p < 0.001). Post-hoc tests indicated that peak amplitudes were significantly higher for younger adults over parietal electrodes (t = 7.193, p < 0.001), and central electrodes (t = 3.442, p < 0.006), but not significantly different over frontal electrodes (p = -0.592, p = 1).

Analysis of the effect of localisation on stimulus-locked peak latency showed no significant effect of Age (p = 0.383), or Coherence (p = 0.83), or Age-Coherence Interaction (p = 0.63). There was a significant effect of Location ($F_{(1.61\ 64)}$ = 48.12, p < 0.001), but no Location-Coherence interaction (p = 0.397). The Location-Age interaction was significant ($F_{(1.61\ 64)}$ = 44.095, p < 0.001), with a significant Location-Age-Coherence interaction ($F_{(1.61\ 64)}$ = 4.99, p = 0.012). The interaction term was likely driven by a complex effect. Young adults showed a linearly increasing latency from frontal to parietal electrodes, with a potential but non-significantly faster latency for low coherence trials over central electrodes (Mean Difference = -38.5, p = 0.981), but slower latencies for low-coherence trials over parietal regions (Mean Difference = 46.03, p = 0.422). Older participants showed no significant coherence effects at any location (all p = 1), but slower latencies over central electrodes than for either parietal or frontal regions.

Non-CSD'd Data Results:

Analysis of non-CSD transformed data indicated a significant main effect of Location ($F_{(1.76,64)}$ = 83.93, p < 0.001), Age ($F_{(1,64)}$ = 16.08, p < 0.001) and

Coherence ($F_{(1,64)}$ = 30.19, p < 0.001) on response-locked amplitude. There was a significant Age by Location interaction ($F_{(1.76,64)}$ = 29.17, p < 0.001), no significant Location by Coherence interaction (p = 0.095), and no significant Coherence by Age interaction (p = 0.135) or three-way interaction (p = 0.397). Post-hoc analyses indicated significant age-difference trends at all locations, with a significantly less positive in younger adults over frontal electrodes (t = -4.531, p <0.001), but a more positive amplitude over central (t = 5.039, p<0.001), and parietal electrodes (t = 5.61, p <0.001).

Analysis of this response-locked slopes indicated a significant main effect of Location ($F_{(1.33,64)}$ = 12.17, p < 0.001), Age ($F_{(1,64)}$ = 20.63, p < 0.001) and Coherence ($F_{(1,64)}$ = 41.46, p < 0.001). There was a significant Age by Location interaction ($F_{(1.38,64)}$ = 20.04, p < 0.001), no significant Location by Coherence interaction (p = 0.08), no significant Coherence by Age interaction (p = 0.169) or three-way interaction (p = 0.632). Post-hoc analyses indicated significant age-difference trends at all locations, with a significantly less positive in younger adults over frontal electrodes (t = -3.382, p = 0.007), but a more positive amplitude over central (t = 4.352, p < 0.001), and parietal electrodes (t = 5.49, p < 0.001).