Perceptual Decision-Making under Uncertainty in Humans

Maria Ruesseler

A dissertation submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy in the University of Oxford

St John's College **Hilary Term 2021**

Perceptual Decision-Making under Uncertainty in Humans

Maria Ruesseler

Abstract

On a daily basis, humans need to make decisions in a complex uncertain world that requires them to adapt their decision strategies to the environmental volatility. In order to make a decision, humans often need to integrate noisy information and combine it with prior experience they learned through previous similar decision situations. This thesis examines how humans form decisions based on noisy information using electroencephalography (EEG), and how they make use of prior expectation in order to guide their decisions using magnetoencephalography (MEG). For both parts, models that capture the different decision-making processes were used to make predictions of human behaviour and underlying neural activity.

To study how humans integrate noisy evidence a novel continuous random dot motion task was developed in which signal periods had to be detected in a stream of noisy evidence which, in contrast to previous studies, is controlled by the experimenter. Simulations with a leaky accumulator model predicted that participants should discount past evidence to solve the task. Participant behaviour suggests that they adapted their evidence strategy to the frequency and length of signal periods in the noise in a manner that is consistent with leaky accumulation. A contingent negative variation (CNV) signal measured with EEG suggests that the brain adapted its evidence accumulation process in line with the behavioural results. In addition, a P300-like signal encoded the difference between previous and current evidence sample, scaling with the size of the difference and trial frequency.

To study how humans update their prior belief about decision outcomes a different random dot motion task was used in which participants had to integrate evidence within trials but also had to integrate prior expectation over past trial outcomes to inform choice. A Bayesian model was fit to subjects' behavioural data to estimate the prior expectation on each trial that was subsequently used to study human choice behaviour and neural activity in the task. Human choice behaviour depended on both, the within-trial evidence, as well as prior belief. Activity in parietal cortex, measured with MEG, at the time of feedback suggested that the feedback type, i.e., whether a choice was correct or incorrect, is encoded in parietal cortex. There was only a very weak influence of prior belief on the neural signal.

This work was funded by a 4-year studentship from the Wellcome Trust.

Table of Contents

Acknowledgements vi
Declaration vii
Long abstract viii

1 Introduction 1

1.1 Neural basis of perceptual decision-making in non-human primates and humans 1

1.1.1 Evidence accumulation in non-human primates 2

1.1.2 Evidence accumulation in humans 6

1.2 Different roles of parietal and frontal areas in perceptual decision-making 13

1.3 Decision-making models 17

1.3.1 Bayesian models of decision-making and learning 17

1.3.2 The drift-diffusion model 19

1.3.3 The Ornstein-Uhlenbeck process 22

2 A new experimental paradigm to study evidence integration in continuously varying environments 25

2.1 Task design 27

2.1.1 Task properties of a continuous random dot motion task 27

2.1.2 Trial length and frequency should influence the optimal strategy of evidence integration 31

2.1.3 The noise structure of the task 33

2.2 Computational Modelling 35

2.2.1 The Ornstein-Uhlenbeck process can capture the different features of the task 35

2.2.2 The OU process predicts task dependent decay rates and decision thresholds 39

2.2.3 How do we measure the differences in evidence integration for different task conditions reliably? 44

2.3 Conclusion 50

3 Behavioural performance on a continuous evidence integration task 53

3.1 Introduction 53

3.2 Task design and testing protocol 56

3.2.1 Design of the random dot motion stimulus 57

3.2.2 Training protocol 58

3.2.3 Testing protocol 61
3.3 Methods and results 61
3.3.1 Detection rate 62
3.3.2 False alarm rates 63
3.3.3 Reaction times 65
3.3.4 Integration kernels 69
3.3.5 Correlation between integration kernels and reaction times 76
3.3.6 Trial responses 78
3.4 Discussion 79
3.4.1 Evidence for a leaky accumulator model 80
3.4.2 Other models of decision-making might explain the behavioural results 81
3.4.3 Single-subject measures of evidence integration 83
3.4.4 What to look for in the EEG chapter 84

4 Neural signals during a continuous evidence integration task recorded with EEG 85
4.1 Introduction 85
4.2 Results for the conventional ERP analysis 89
4.2.1 ERPs time-locked to response 89
4.2.2 ERPs time-locked to trial start 93
4.3 Method and results for the convolutional GLM analysis 95
4.3.1 Method of the convolutional GLM 96
4.3.2 Replication of findings from conventional analysis 101
4.3.3 Convolutional GLM analysis of intertrial periods 105
4.4 Control experiment 108
4.5 Discussion 111
4.5.1 Neural activity reflects complex adaptation to different task environments 112
4.5.2 Neural markers of decision-formation on the single-subject level 114
4.5.3 Differences between findings in the present study and other studies of evidence integration 115
4.5.4 Conclusion 117
4.6 Methods 117
4.6.1 Task paradigms 117
4.6.2 EEG recording and pre-processing: 118

4.6.3 Analysis 119
4.6.4 Participants 121
4.7 Supplementary Figures 122

5 Influence of prior expectation on neural activity in parietal cortex measured with MEG 125
5.1 Introduction 125
5.2 Study design 128
5.2.1 Task 128
5.2.2 Model details 130
5.2.3 Training and testing regimen 132
5.2.4 Analysis of saccade direction 132
5.3 Behavioural findings 133
5.3.1 Feedback on previous trials affects choice 133
5.3.2 Prior expectation interacts with current evidence to predict choice 135
5.4 MEG methods 138
5.4.1 MEG acquisition 138
5.4.2 MEG pre-processing 138
5.5 MEG analysis and results 139
5.5.1 Effect of trial outcome on feedback processing 139
5.5.2 Effect of expectation strength on feedback processing 141
5.5.3 Interaction of expectation strength with feedback type (correct/incorrect) 144
5.5.4 Congruency of expectation and feedback 145
5.5.5 Single-subject results for prior belief and trial outcomes 147
5.6 Discussion 149
5.6.1 Stimulus certainty influences choice 149
5.6.2 Influence of prior belief on parietal cortex 151
5.7 Supplementary Figure 153

6 Conclusions and future directions 154
6.1 Combining accumulator models and reinforcement learning to understand decision-making 155
6.2 Understanding cognitive processes in mental diseases 158

References 162

Acknowledgements

First and foremost, I would like to thank my supervisors Laurence Hunt and Jill O'Reilly for their invaluable academic guidance throughout my PhD as well as their support and kindness after a very challenging first year and during this challenging final year.

I am also indebted to my colleagues Tom Marshall, Lilian Weber, Cameron Hassall and Yu Takagi for their help with many and various aspects of the work presented herein. I thank Jonathan Hadida, Nils Kolling, Jackson Smith, Anna-Katharina Matke-Bauer, Yvonne Fonken, Rocio Silva Zunino and the entire OHBA crew for their invaluable discussions and advice regarding EEG data collection.

I would also like to thank my family and friends for their support and friendship, especially Jana Rosenau, Alberto Lazari, Sandra Straulino, Sarah Buchaupt, Rita Lang, Hee Kyoung Malkowski, and Anna and Andrew Emberton, which made all the difference.

Declaration

The work in this thesis was performed by the author. All collaboration and assistance are described below:

Chapter 4: The code for eyeblink removal from the EEG was adapted from code provided by Lilian Weber with her permission. Undergraduate Nebojsa Jovanovic adjusted the original task paradigm used in Chapters 1-4 to include vertical motion under my supervision.

Chapter 5: The study was designed by Tom Marshall and Jill O'Reilly. Data collection was jointly carried out by Tom Marshall and the author. Pre-processing of the MEG data to remove artefacts and the analysis of choice direction from eye movement data were performed by Tom Marshall. Jill O'Reilly provided the code for the Bayesian model.

Long Abstract

Chapter 1 introduces the concept of perceptual decision-making and common modelling approaches to study perceptual decision-making which enabled the identification of neural correlates of evidence accumulation. Neural traces of evidence accumulation are found across cortex. Specifically, the similarity in neural activity during evidence accumulation in frontal and parietal cortex leads to the questions why there is redundancy in the neural encoding of evidence accumulation and the different functional roles of those areas. I suggest that parietal activity reflects the integration of prior belief to build a priority map of the visual field while frontal areas integrate current evidence. I introduce the Bayesian model that allows me to make predictions about the influence of prior belief on behavioural choices and neural activity in chapter 5. Further, I describe accumulator models and more specifically *leaky* accumulator models, which I use to make predictions about evidence integration strategies in noisy environments in which evidence constantly fluctuates in Chapter 2. These predictions are then compared to human behaviour and neural activity in chapters 3 and 4.

Chapter 2 introduces a novel continuous random dot task and a modelling approach to study decision-making under uncertainty. The continuous random dot paradigm combines features of evidence integration and signal detection tasks. Observers have to *detect* coherent motion in a stream of noisy randomly moving dots, but they also have to *integrate* the direction of the coherently moving dots. In contrast to previous continuous paradigms, however, the *noise* is manipulated by the experimenter resulting in rapid changes in the evidence observers have to integrate. This allows me to study decision-making when the onset of signal (the onset of coherent motion) is unclear. Using a leaky accumulator model, I am hypothesizing that humans should discount past evidence to detect signal periods. Moreover, the leaky accumulator model makes specific predictions how human observers should adapt their behaviour when the frequency of signal periods and their length varies.

Chapter 3 presents behavioural results suggesting that humans adapt their behaviour in the novel continuous random dot task to the frequency and length of signal periods that is largely consistent with predictions of the leaky accumulator model. Measures of human behaviour indicate that humans adapt their evidence accumulation strategy in a manner that suggests longer integration for long signal periods and a more cautious response behaviour when trials are rare, implying that they adapt their decision-threshold to require more evidence.

Chapter 4 examines neural activity recoded with electroencephalography (EEG) while participants were doing the continuous random dot motion task. Using a convolutional GLM approach I investigate the continuous EEG data without time-locking to specific events. This enables me to examine the influence of rapid changes in evidence introduced by the noise manipulation as well as signals that develop over the presentation of many changes in evidence. A CNV signal measured centrally on the scalp shows significant differences depending on the frequency and length of signal periods. In addition, each new incoming evidence sample is represented by a positive EEG signal that scales with the difference between previous and current evidence sample and the frequency of signal periods. This signal seems to be similar to the P300 or centro-parietal positivity which have been shown to reflect evidence integration and surprise in previous studies.

Chapter 5 presents behavioural and neural results on a different discrete-trial random dot motion task in which participants have to integrate evidence within a trial but also build up a prior expectation over past trial outcomes in order make a correct decision. Behavioural results suggest that participants integrate prior belief over past trial outcomes consistent with a Bayesian model. They use both prior belief and evidence on the current trial to make a decision. Using magnetoencephalography (MEG) I probe whether neural activity in parietal cortex reflects the integration of prior expectation after participants have seen the feedback on the current trial. However, while activity in parietal cortex encodes whether the choice on the trial was correct or incorrect, there is only a weak influence of prior belief on neural activity at the time of feedback.

Chapter 6 brings together the results from both studies represented in this thesis and discusses 2 areas in which continuous evidence integration tasks can advance our understand of decision-making. First, I discuss the advantage of combined models of reinforcement learning and accumulation models together with a continuous paradigm to examine the neural representation of complex decision-making processes requiring both, integration of current evidence and prior expectation. Second, I discuss how continuous task paradigms combined with the convolutional GLM approach allow us to further understand the differences in cognitive processes in health and psychopathologies.

1 Introduction

We are living in a complex environment and often need to base our decisions on insufficient information. Especially in the last year we all had to base our decisions on very uncertain information that sometimes changed daily. At the start of the world-wide coronavirus pandemic in March 2020, scientists constantly provided new insides about the virus. We had to *integrate* this information based on its impact on our daily life to make profound decisions, such as whether it is safe to take the bus to work. Though we are still learning new information about the virus every day, we also *gained experienc*e in the past year how to regulate the number of infections and to protect ourselves from the virus. We learned how effective lockdowns and good testing strategies are and have more information about how the virus spreads that helps us to navigate our daily lives. This raises the question how our brain can make decisions when the information provided is uncertain. We need appropriate neural mechanisms that allow us to integrate noisy evidence and combine it with prior expectation to make a good choice.

In this thesis I will investigate how temporal dependencies influence decision-making in humans. I am using the framework of perceptual decision-making because it is relatively well understood in terms of simple decisions. However, even visual decisions, such as identifying behaviourally relevant objects in the visual field require the integration of current evidence and past experience. In the first part of this thesis, I examine how the brain adapts to fast temporal changes in the environment, in which only the most recent sensory history should be integrated (Chapters 2-4). Then, I will look at how prior belief, built up over several actions, is integrated with current sensory information to form a decision (Chapter 5).

1.1 Neural basis of perceptual decision-making in non-human primates and humans

My research is based on the knowledge of perceptual decision-making processes that have been mostly investigated in monkeys and humans and recently also in rodents. I will briefly describe how electrophysiological studies in monkeys have informed the way we think about perceptual decision-making. I will link this to

studies in humans that build on the neural and behavioural findings in monkeys. Subsequently, I will discuss the role of parietal cortex in perceptual decision-making and contrast this with frontal cortex. Recently, the role of parietal cortex during decision-making has been challenged with some studies suggesting that it represents prior expectations, while frontal cortex integrates evidence. Finally, I will discuss some of the most common models for decision-making, including drift-diffusion models and Bayesian models, which provide the conceptual framework for my thesis.

1.1.1 Evidence accumulation in non-human primates

Electrophysiological studies in monkeys in the second half of the 20th century were able to examine how sensory information and movement planning is represented in the brain. They used specifically developed tasks to test the receptive field properties of neurons in order to examine to which features of a stimulus they would be responsive. For example, single neuron recordings in the middle temporal area (MT) showed that these neurons are selective to specific motion directions (Albright, 1984; Britten et al., 1993, 1996). The understanding of how neurons encode sensory information is important because it allows us to probe how this information is later used for higher cognitive processes.

In order to study motion selectivity in MT, studies used a random dot motion paradigm (RDM). In this paradigm monkeys fixate on a central fixation point while observing a circular aperture with moving dots. The monkeys have to make a judgement on the average motion direction of the dots and saccade to a congruent target to either side of the aperture (Figure 1A). Because the monkey is 'forced' to make a choice between two alternatives, this paradigm is also called a 2-alternative forced choice paradigm (2AFC). The difficulty of RDMs tasks is modified by adjusting the coherence, that is the fraction of dots moving coherently in one direction while the rest of the dots move randomly. If the coherence is small, the task is more difficult. RDM tasks have been used beyond the study of motion direction selectiveness in MT and have become the most prominent task paradigm used to investigate neural processes underlying evidence accumulation.

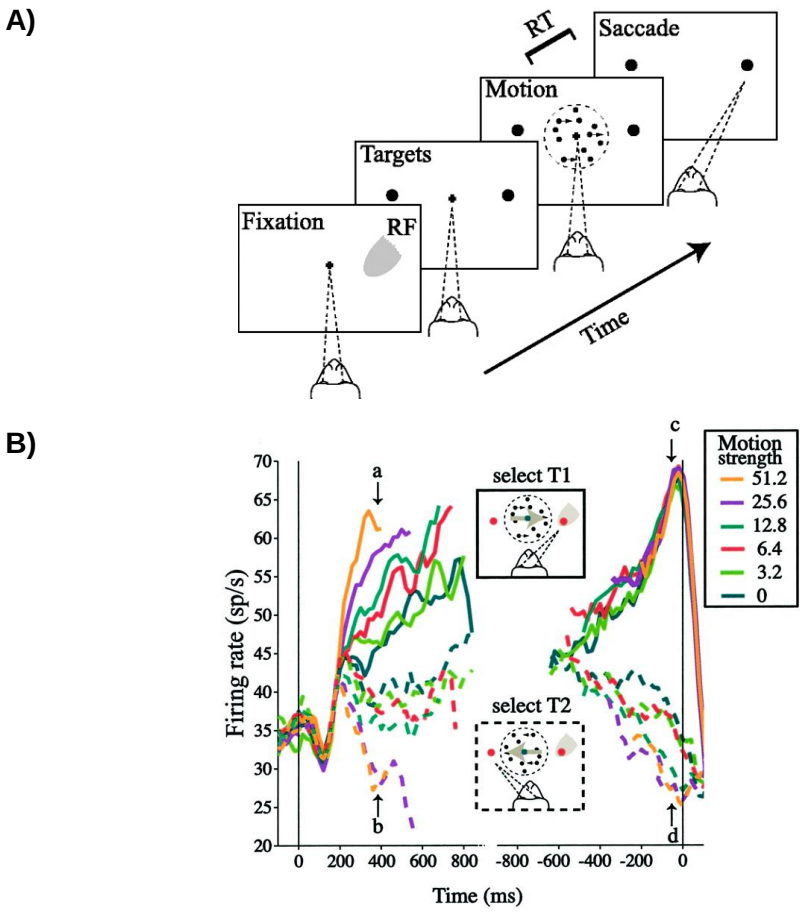


Figure 1: 2AFC RDM task schematic and neural ramping activity of LIP neurons during task performance (adopted from Roitman and Shadlen, 2002). **A)** Each trial starts with the monkey fixating on a central fixation point. The targets appear followed by a cloud of randomly moving dots of which a fraction moves in a given direction to either target. As soon as the monkey has made its decision about the rotation direction of the moving dots, it has to make a saccade to the congruent target. Neural activity is recorded from neurons with their receptive field on either target location. **B)** Time course of LIP activity during the motion viewing period shown. If neural activity is time-locked to trial start neurons show stereotypical ramping activity that is scaled by motion strength if their receptive field is at the chosen target (left side). Independent of stimulus strength this ramping activity reaches a threshold that is visible when time-locked to saccadic response (right side).

Not all representations of sensory information in the brain are important for subsequent decision-making. In order for neural activity to encode sensory evidence that is used for downstream decision processes it should be directly related to the choice behaviour. A good example is neural activity in MT while monkeys perform a RDM task. Neural activity in MT is correlated with the strength

of the random dot stimulus and predictive of the choice the monkey is about to make (Britten et al., 1992, 1996; Salzman et al., 1992). The same is true if monkeys have to judge the rotation direction of a bistable figure. Even though there is no physical evidence for either rotation direction, monkeys will still perceive a specific rotation direction that is predicted by activity in MT neurons (Dodd et al., 2001). This means that MT neurons can weakly but significantly predict the perceptual choices that they are well-suited to inform. This type of observation is suggestive of sensory evidence that that builds the foundation of the decision-making process.

Subsequently, sensory evidence has to be integrated into a decision-variable that determines the choice of the monkey. The representation of the decision-variable for a task, in which the motion direction has to be judged, should be in an area that is connected with MT (Blatt et al., 1990; Lewis and Van Essen, 2000). Further it should have connections with motor areas for a subsequent response. For that reason, the lateral intraparietal area (LIP) was chosen as a candidate area for encoding a decision variable. Based on previous studies, LIP is known for its representation of salience and for the planning of saccadic eye movements (Colby and Goldberg, 1999; Mazzoni et al., 1996; Snyder et al., 2000).

Shadlen and Newsome (2001) recorded from neurons in LIP while trained monkeys performed the RDM task and showed that the activity of these neurons reflects evidence integration. During the stimulus viewing period neurons exhibit stereotypical ramping activity to a threshold when their receptive field covers the target location that the monkey is about to saccade to. The slope of the ramp scales with stimulus strength (coherence) (Figure 1B). Higher coherence leads to a steeper ramping slope. In addition, the ramping activity is also present in trials with a coherence of 0%, i.e., when there is no motion signal in the trial, but the animal still chooses the target in the neuron's receptive field. Collectively this provides evidence that these neurons represent an evidence integration signal.

However, this study does not provide direct evidence linking the ramping activity and the timing of the subsequent decision. In order to separate eye movement and integration signals in LIP, the random dot stimulus was always shown for a set amount of time. After the stimulus presentation, there was a short delay period

before the targets appeared and the monkey was allowed to make a decision. This does not test whether there is a direct relationship between neural signals and when the decision is made, because the time frame for the signal development in LIP is the same on each trial.

For LIP to encode a decision variable the ramping activity should correlate with reaction time (RT) and reach a stereotypical threshold eliciting a response that is independent of the stimulus strength. Roitman and Shadlen (2002) changed the RDM task so that monkeys had to make an eye movement to the target as soon as they decided on the rotation direction of the motion stimulus. Independent of coherence, LIP activity reaches a stereotypical threshold eliciting a response and the slope of the ramp correlates with RT (Figure 1B). Further, causal evidence that LIP is crucial in the decision-making process comes from a study which showed that electrical microstimulation of neurons during a RDM trial increases the proportion of choices toward their receptive field (Hanks et al., 2006).

In addition, for perfect accumulation LIP neurons should weigh evidence by its relevance, i.e., its likelihood that it supports a certain choice. This means that activity in LIP should reflect changes in evidence in a stimulus. Indeed, it has been shown that LIP activity reflects the temporal integration of RDM stimuli in which motion coherence fluctuates within a trial until the accumulated evidence reaches a decision threshold (Huk and Shadlen, 2005; Kiani et al., 2008). Further experiments have been able to directly show that LIP weighs each new incoming evidence sample by its relevance using a task in which monkeys had to integrate discrete visual cues which were assigned different probabilities in support of one of two choices (Kira et al., 2015; Yang and Shadlen, 2007).

Though there is compelling evidence for the role of LIP in perceptual decision-making, only a small subset of its neurons has typically been studied. In most studies, researchers only recorded from neurons in LIP with very specific memory activity during delay periods that are located in the dorsal part of LIP (Gold and Shadlen, 2007; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001). However, more recent studies that recorded from a broader set of neurons showed that LIP neurons reflect diverse dynamics representing decision-relevant and

decision-irrelevant signals (Meister et al., 2013; Park et al., 2014; Zoltowski et al., 2019). It is therefore not possible to pin down the exact role of LIP in decision-making. Further, signals of evidence accumulation have been found across the brain in cortical and subcortical areas, such as the superior colliculus (Ratcliff et al., 2007) and the frontal eye field (FEF) (Ding and Gold, 2012; Kim and Shadlen, 1999), indicating that the process of evidence integration is much more complex. This will be further discussed in section 1.2 (Different roles of parietal and frontal areas in (perceptual) decision-making).

1.1.2 Evidence accumulation in humans

The principles of evidence integration as set out by electrophysiological studies in monkeys have been subsequently used to study decision-making in humans. This includes the identification of brain signals reflecting true sensory evidence and an evidence accumulation signal that reflects the integral of the sensory evidence. It allows for the identification of homolog areas to the brain areas that are activated in animal models during decision-making, such as the identification of a true sensory evidence signal in hMT+ during a RDM task (Serences and Boynton, 2007). Further, non-invasive imaging in humans has the advantage that the neural dynamics across the whole brain can be studied while recordings in animal models are often limited to single neurons across only a few brain areas.

Functional magnetic resonance imaging (fMRI) studies were able to identify brain structures involved in evidence accumulation such as the dorsolateral prefrontal cortex (Heekeren et al., 2004), intraparietal sulcus (Kayser et al., 2010; Liu and Pleskac, 2011) and posterior parietal cortex (Tosoni et al., 2008). However, due to the nature of the blood-oxygen-level dependent (BOLD) signal which is only an indirect measure of brain activity this has been more difficult. Though fMRI allows a high spatial resolution, the temporal resolution of the BOLD signal is very limited and insensitive to fast transient changes in activity. A prediction of how the BOLD signal would behave during evidence accumulation is not possible and depends on assumptions about the behaviour of the underlying neural activity. Indeed, while some reported an increase of the BOLD signal during evidence accumulation

(Heekeren et al., 2004; Pleger et al., 2006), others found a decrease in the BOLD signal (Liu and Pleskac, 2011).

Though fMRI experiments could identify brain areas involved in decision-formation, they are not able to fully explain the underlying neuronal signals. Due to their high temporal resolution electroencephalography (EEG) and magnetoencephalography (MEG) are much better suited to investigate how the decision-formation signals unfold in the brain on the level of milliseconds. For example, a network study using MEG showed that different stages of the decision-making processes show oscillatory power in different frequency bands that can partially overlap in time (Siegel et al., 2011). This highlights the advantage of the high temporal resolution achieved with MEG as it allows to better separate brain signals belonging to different sensory stages.

Donner et al. (2009) were the first to measure an evidence integration signal using MEG. They used a random dot motion stimulus in a signal detection task. Participants had to detect the presence of a weakly coherent motion signal that was imbedded in randomly moving dots on only half of the trials. They showed that the oscillatory power in the beta (15 – 30Hz) and gamma (30 – 80Hz) band in the motor preparation signal increases leading to the response. Moreover, this activity reflected the integral of the gamma band activity in hMT+ that encodes the sensory evidence. More compellingly, the activity predicted the choice and was similar for correctly detected coherent motion as well as for responses in trials without a coherent motion signal (false alarms). This means that this fulfils a further criterion of an evidence integration signal which is that it depends upon the subjects' choice rather than the external evidence. In addition, this signal seems to be modulated by other external and internal factors, such as prior expectation and past choices that affect the activity prior to stimulus onset (de Lange et al., 2013). This is similar to evidence accumulation models which also show variability in the starting point depending on a bias when fitted to behaviour (Mulder et al., 2012; Ratcliff and Smith, 2004).

A problem with a lot of tasks in EEG and MEG used to study decision signals is that the abrupt onset of a stimulus elicits a prominent visual evoked potential (VEP)

overlapping and masking other signals representing higher cognitive processes. In order to distinguish between a purely visual response and other processes, several EEG studies used continuous stimuli that had gradual changes in their properties which have to be detected (Kelly and O'Connell, 2013; O'Connell et al., 2012). For example, in one study participants observed a flickering annulus stimulus for several minutes (O'Connell et al., 2012). They had to detect gradual changes in the contrast of that stimulus that happened sporadically. This allowed the experimenters to measure the steady state visual evoked response (SSVEP) based on the flickering rate of the stimulus and allowed them to separate it from other signals related to the decision process. In addition, they used lateralized button press responses to track the motor preparation signal contralateral to the response movement.

Three signals representing sensory evidence, evidence accumulation and motor preparation can be distinguished using EEG and continuous stimuli with lateralised button press responses. The task using gradual changes in contrast described above showed that the SSVEP represents a true sensory evidence signal because the amplitude of the signal correlated with the stimulus contrast changes and predicted the time point of the impending choice. They also found motor preparation signals showing signs of evidence integration similar to previous research (Donner et al., 2009). These signals gradually decreased (lateralised beta) or increased (lateralised readiness potential) until they reached a fixed threshold irrespective of RTs. They resembled the grand averaged evidence accumulation. However, these signals are independent of the response type following an all-or-nothing relationship, meaning they were not present without response execution such as in tasks in which participants only tracked the evidence without an overt response. This means that they are more closely related to execution signals (O'Connell et al., 2012; Twomey et al., 2016).

The third signal they discovered represents a freely evolving evidence integration signal over centro-parietal electrodes, called the centro-parietal positivity (CPP). The CPP like the motor preparation signals represents integrated evidence and builds up until a fixed threshold that is independent of RTs. The build-up rate is dependent on stimulus strength (Figure 2). In contrast to motor preparation signals

it is present even without an overt response (but is weaker for false alarms). This is similar to the neural activity in monkey LIP that also scales with evidence strength and shows weaker ramping activity in trials with 0% coherence. Further, the CPP is an universal evidence integration signal because it develops independently of the modality of the sensory evidence and is insensitive to a distractor stimulus (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2016). However, an urgency signal is also represented in the motor signal independently of the CPP, that has a top-down influence on the evidence accumulation and activity of the CPP (Kelly et al., 2019). It is therefore argued, that CPP is a processing stage of the decision process between the sensory evidence representation and motor execution signals.

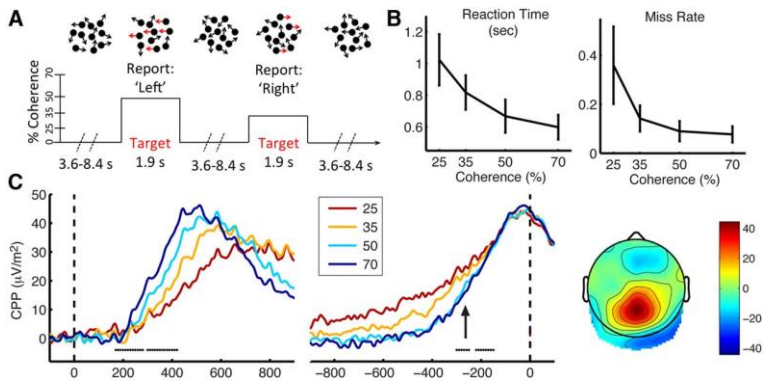


Figure 2: CCP during motion integration shows stereotypical ramping to a threshold (adopted from Kelly and O'Connell, 2013). A) Continuous random dot motion task in which participants have to judge the motion direction in periods with coherent motion. **B)** As in monkeys, low coherences are more difficult leading to increased RTs and misses of coherent motion periods. **C)** The CPP recorded with EEG shows stereotypical ramping activity scaled by motion strength when time-locked to trial start (left panel) that reaches a threshold eliciting response (middle panel). Topographically it is located over centro-parietal cortex (right panel).

The CPP shows a resemblance with the widely studied P300, including scalp topography, peak coincidence with movement activation and sensitivity to task relevance (O'Connell et al., 2012; Twomey et al., 2015). Though the P300 has been linked to many cognitive processes (Polich and Kok, 1995) such as context updating (Donchin and Coles, 1988), uncertainty (Duncan-Johnson and Donchin, 1977; Mars et al., 2008) and confidence in a decision (Hillyard et al., 1971; Squires et al., 1973), its exact function remains unknown. Twomey et al. (2015) used an

odd-ball task to show that the P300 represents evidence accumulation similar to the CPP discovered by O'Connell et al. (2012). Similar to the CPP it scales with the strength of sensory evidence and builds up to a threshold eliciting an action. The rate of the build-up predicts RT as does the build-up rate of the CPP and can be fitted with a common evidence integration model. Further research should revisit its role in other cognitive functions such as surprise with respect to its role in evidence accumulation.

Other external factors can influence the CPP during evidence accumulation that highlight how flexibly the CPP can adapt to different circumstances. In the natural world the source of evidence is not always clear neither its onset. Sometimes, that means we have to track different sources of information simultaneously. Loughnane et al. (2016) showed that the build-up rate and start of the CPP is influenced by target selection signals that are also important for visual search, when more than two sources of information have to be tracked and one of those is a distractor. Other studies have shown that an urgency signal influences the build-up and starting point of the CPP, for example in paradigms that favour speed over accuracy or in which the onset of evidence in the stimulus is uncertain (Devine et al., 2019; Steinemann et al., 2018).

Studies on the CPP have been able to identify decision-making signals that seem to be homolog to those in non-human primates. All of these studies assume that the evidence or stimulus strength remains constant during the trial period. However, to understand the mechanisms that give rise to the CPP we also need to understand how each new piece of incoming evidence affects this signal. This can be achieved by manipulating sensory evidence across time, and studying the impact of new incoming evidence on brain activity (Cheadle et al., 2014; De Lange et al., 2010; Wyart et al., 2012). In one study participants had to judge whether the orientation of a stream of Gabor patches was more diagonal or cardinal while they recorded EEG (Wyart et al., 2012). They found that activity over the occipital lobe encodes a perceptual update of each pattern (Figure 3A). In contrast, activity over parietal cortex represents a decision update, encoding the weight that each pattern contributes to the overall evidence accumulation (Figure 3B). The decision update is underpinned by slow rhythmic fluctuations over parietal cortex. This means that

the weight of each new incoming evidence (observed pattern) that is added to a decision variable depends on the oscillation of the brain activity. Similar to the CPP this activity is not observed for a distractor (Wyart et al., 2015). Further, the weight of the decision update does not only depend on the oscillatory state of the brain but also on previous accumulated evidence. The change in activity after a new evidence sample decreases when more evidence has already been accumulated (De Lange et al., 2010).

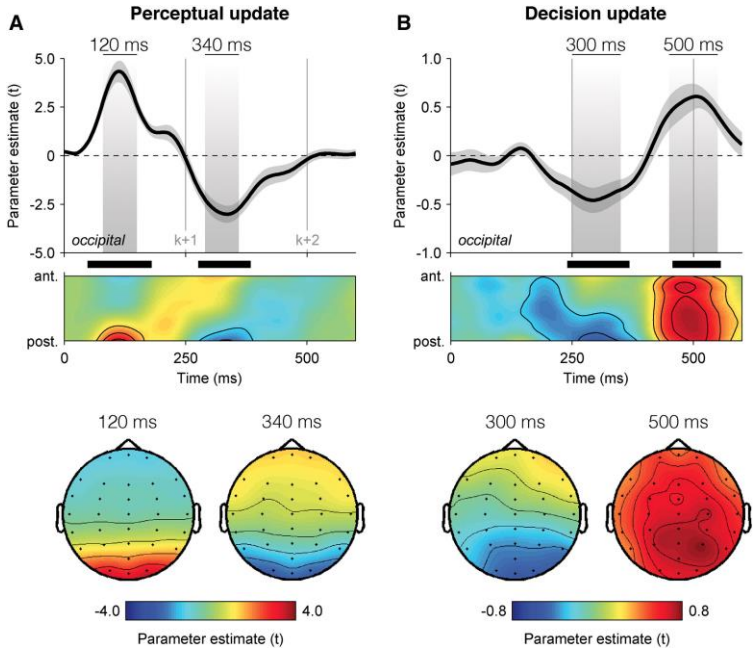


Figure 3: Perceptual and decision updates can be dissociated with EEG. A 2AFC evidence integration task in which evidence is shown as a stream of discrete samples, allows to investigate the influence of each new incoming sample on neural activity. Shown are parameter estimates. **A)** Each new evidence sample leads to a perceptual update which is topographically located over occipital cortex and precedes a decision update. **B)** The decision update is topographically different with a negative deflection that is more over occipital cortex that is followed by a positive deflection over parietal cortex. It represents the update of a decision variable.

For the main part of my thesis I developed a new continuous RDM paradigm similar to Kelly and O'Connell (2013, Figure 2A) but I combine it with constantly changing evidence similar to (De Lange et al., 2010; Wyart et al., 2012). This will allow me to study the influence of discrete new evidence samples on brain activity while at the same time investigating event related potentials (ERPs) that usually come with

typical single trial event locked experiments, such as the CPP or a motor preparation signal. In addition, this paradigm enables me to study the manipulation of evidence when no decision is made because during intertrial periods evidence is still changing constantly but participants should not emit a response.

In order to analyse the EEG data, I will make use of a regression approach because traditional time-locked ERP analysis is suboptimal for this paradigm. It cannot separate overlapping transient EEG signals that relate to the rapidly changing evidence samples as well as separating these from EEG signals that develop over the representation of many evidence samples. However, there are regression methods that can estimate the independent contribution of separate predicator variables to the ERP waveforms (Ehinger and Dimigen, 2019; Smith and Kutas, 2015a, 2015b). Crucially, in order to obtain regression ERPs, I will use the continuous EEG rather than the time-locked data. By selecting predictors for different events such as a change in evidence, trial start or response I can separate the overlapping ERPs.

Together the continuous paradigm and the use of regression ERPs allow me to study how uncertainty influences decision-making. In the past uncertainty in humans has been studied using time-varying evidence. In order to solve these tasks many studies suggest leaky integration in which more weight is given to recent evidence while earlier evidence is discounted (Bronfman et al., 2016; Ganupuru et al., 2019; Glaze et al., 2015; Harun et al., 2020; Levi et al., 2018; Ossmy et al., 2013; Tsetsos et al., 2012). The statistics or uncertainty of changes in the environment determines the integration kernel participants use, i.e. how fast they should discount past evidence (Glaze et al., 2015; Veliz-Cuba et al., 2016). For example, humans can adapt their evidence integration kernels to signal durations using integration kernels with higher decay rates for shorter signal durations (Ganupuru et al., 2019; Harun et al., 2020; Ossmy et al., 2013). In my task design I specifically manipulate the frequency and duration of the sensory evidence in order to test how flexibly participants can adapt their evidence accumulation process to different timescales when evidence is noisy and inconsistent.

1.2 Different roles of parietal and frontal areas in perceptual decision-making

Based on previous studies there is evidence that a network consisting of parietal, frontal, motor and sensory areas is actively involved in the decision-making process (Siegel et al., 2011, 2015). For example, in a dot-motion task the direction of coherently moving dots is first encoded in MT in non-human primates (Albright, 1984; Britten et al., 1993, 1996). Downstream areas LIP (Huk and Shadlen, 2005; Roitman and Shadlen, 2002; Shadlen and Newsome, 2001) and FEF (Ding and Gold, 2012; Kim and Shadlen, 1999) will then use this information to compute a decision and a movement plan that then can be transformed into action in motor areas.

However, there is still an ongoing debate about the different functions of parietal and frontal areas during the decision-making process. Both areas show ramping activity during a decision-making task reflecting the integration of sensory evidence in monkeys (Ding and Gold, 2012; Kim and Shadlen, 1999; Roitman and Shadlen, 2002), rodents (Erlich et al., 2015; Hanks et al., 2015) and humans(Heekeren et al., 2004; Kayser et al., 2010; Liu and Pleskac, 2011; Philiastides et al., 2011; Tosoni et al., 2008). Studies imply that both areas are activated simultaneously during the evidence accumulation process (Hanks et al., 2015; Siegel et al., 2011, 2015). This poses the question what the different contributions of each of these areas are.

Recently it has been called into question whether single neural activity in monkey LIP can be characterised as ramping to a threshold (Latimer et al., 2015). Instead, it was proposed that neurons in LIP follow a step-like change in activity, meaning that the activity of neurons either steps up or down to different states depending on the evidence. This means that LIP might be still involved in the decision-process but might not reflect ongoing evidence accumulation. However there is an ongoing debate about these results with some studies still arguing that the activity of neurons reflects ramping to threshold (Chandrasekaran et al., 2018; Latimer et al., 2017; Shadlen et al., 2016; Zhao and Kording, 2018; Zylberberg and Shadlen, 2016).

Recent electrophysiological evidence from awake, behaving monkeys suggests that parietal cortical area LIP may not have as direct a role in evidence accumulation as previously thought. Another way to probe the function of LIP in monkeys is by simultaneous recordings from MT and LIP during a dot-motion evidence integration task (Yates et al., 2017). This is based on the notion that MT encodes the sensory information, i.e., the motion direction in the stimulus, while downstream area LIP integrates the sensory evidence into a decision. Multiarea recordings from both of these sites showed that there is only a weak motion integration signal in LIP. The activity in LIP can be mainly described by a premotor activity for planning the next saccade. However, in order to plan the next saccade LIP might integrate information from other areas, which might include decision signals. Other studies have shown that neurons in LIP are multiplexed and show responses to task relevant stimuli as well as oculomotor responses (Ipata et al., 2009; Meister et al., 2013; Park et al., 2014). This puts LIP in a perfect place to help plan saccades as it represents all relevant stimuli at a time.

More compelling evidence for a less important role of LIP in evidence accumulation comes from inactivation studies. Inactivation of monkey LIP (Katz et al., 2016) or homolog posterior parietal cortex (PPC) in rats (Erlich et al., 2015) does not impair performance on evidence integration tasks. This is surprising, because microstimulating neurons in LIP increases choices towards the receptive fields of the stimulated neurons while also reducing the RTs of these choices (Hanks et al., 2006). However, LIP's involvement in saccade planning (Colby et al., 1996; Mazzoni et al., 1996) might also account for these results. Perhaps, microstimulation affirms a saccade plan to the target in the receptive field of a stimulated neuron which in turn leads to a faster reaction.

Similarly, in humans there is also conflicting evidence concerning the location of evidence accumulation. Using a discrimination task, a fMRI study has found a correlation between evidence accumulation and the BOLD signal in dorsolateral prefrontal cortex (DLPFC), homolog to the FEF in monkeys (Heekeren et al., 2004). The BOLD signal correlated with trial difficulty, meaning it increased when trials were easy because there was less noise in the stimulus. In addition, the BOLD signal in DLPFC also correlated with the difference in activity in the sensory

areas which represented the stimuli participants had to compare to make a decision. This is in line with a TMS study which found that stimulation of the DLPFC reduces accuracy while prolonging RTs (Philiastides et al., 2011). In contrast, as described in the previous section, some EEG studies suggest that the CPP represents the accumulation of evidence. (Kelly and O'Connell, 2013; O'Connell et al., 2012). The ramping of the CPP depends on stimulus strength and reaches a threshold eliciting the motor execution of the choice. In contrast to the BOLD signal, the CPP can be directly linked to the evidence integration process. However, due to the low spatial resolution of EEG it is not possible to determine the exact source of the signal.

The function of LIP in decision-making might be more complex than previously thought. There is evidence that prior belief can bias the activity of LIP neurons (Hanks et al., 2011). A similar effect has been found in PCC of rats performing an auditory task (Akrami et al., 2018). Rats had to compare two auditory stimuli presented on either side that were separated by a couple of seconds and report the stimulus which was louder. They cannot optimally perform this task because the first stimulus presented in the current trial is biased towards the average loudness of the first stimulus presented in the last trials observed. This means that if the average is closer to the loudness of the second stimulus it becomes harder for the rat to discern which stimulus is louder. Surprisingly, silencing of PPC neurons leads to an improvement of the task. In subsequent analyses, the improvement in performance could only be explained by the reduced influence of the encoded bias on decision-making.

LIP activity does not only reflect decision-making processes but also the allocation of attention (Bisley and Goldberg, 2003; Colby et al., 1996; Gottlieb et al., 1998; Kusunoki et al., 2000). For example, inactivation of LIP impairs behaviour on a visual search task when monkeys have to saccade to a target stimulus that is similar to its surrounding distractors (Wardak et al., 2002). Further, though a pop-out distractor can momentarily increase the activity of neurons with their receptive field at this location, the activity of neurons with their receptive fields on the saccade target remains elevated (Ipata et al., 2006). This supports other research

that not only visual salience influences activity in LIP but also the relevance for future behaviour (Colby et al., 1996; Gottlieb et al., 1998; Kusunoki et al., 2000).

Together with the findings from the multiarea recordings in MT and LIP in monkeys, these results suggest that parietal cortex might be involved in creating a priority map that describes the locations in the visual field that are important for future behaviour and to which a saccade needs to be planned (Bisley and Goldberg, 2010; Ipata et al., 2009; Kusunoki et al., 2000). LIP might use a mixture of information to create this map (Ipata et al., 2009; Meister et al., 2013; Park et al., 2014). This includes past experience about locations in the visual field that previously guided behaviour (Akrami et al., 2018; Hanks et al., 2011).

In contrast, FEF presumably is more important for representing the categorical decision but possibly not the evidence accumulation process per se (Erlich et al., 2015; Hanks et al., 2015; Kim and Shadlen, 1999). A study in rats found that FEF homolog frontal orienting field (FOF) in rats represents a categorical update of the decision process, i.e., the choice that would be made if the trial were to end (Hanks et al., 2015). This means if one would inactivate it early in the trial it does not impair the subsequent decision process. Only inactivation at the end of the trial impairs behaviour (Erlich et al., 2015). Recordings in monkey FEF also point to a role in the formation of a categorical decision for motor preparation (Kim and Shadlen, 1999).

In conclusion, though parietal and frontal cortex have been extensively studied, it is not yet completely understood how they fit into the decision-making network. Both areas seem to carry important information about upcoming decisions. This means for parietal cortex that further studies should look at its role in creating a salience map by integrating over past experience and evidence. The study in Chapter 5 of this thesis examines the representation of prior belief in parietal cortex with a task in which the outcome probability of a trial changes over the course of a session.

1.3 Decision-making models

1.3.1 Bayesian models of decision-making and learning

One of the biggest challenges the brain faces is uncertainty. Uncertainty arises from internal and external factors. Neurons are noisy (Faisal et al., 2008) and the information the brain has to base its decision on is often noisy as well, such as in a random dot motion task in which randomly moving dots make the detection of direction of the coherent ones harder. An optimal model for the brain to deal with this uncertainty rationally is described by a Bayesian model. It has been shown that the behaviour of humans is nearly optimal and approaches the predictions of a Bayesian model under experimental conditions (Behrens et al., 2007; Körding and Wolpert, 2004; O'Reilly et al., 2013).

In its simplest form a Bayesian model can be described as follows:

$$P(H|e) \propto P(e|H) \times P(H)$$

Where *P(H|e)* is the conditional probability that a hypothesis or parameter (*H*) is true given the observed evidence, also called posterior. This can be calculated as the product of the conditional probability that we observe this evidence (*e*) assuming it has been generated by *H*, called likelihood, and a prior probability *P(H)*. For example, in a random dot motion task the likelihood in one trial may be proportional to the estimated coherence of moving dots. Therefore, the ramping activity of neurons during RDM tasks might be interpreted as representation of the likelihood in the brain (Kira et al., 2015; Yang and Shadlen, 2007). In this example, prior belief can be thought of as a bias (for example based on past choice history) that acts as an additional source of information influencing the current choice. Indeed, it has been shown that in a 2AFC RDM task prior probability that captures past trial outcomes influences choice, specifically if it the trial was difficult (Braun et al., 2018; Drugowitsch et al., 2019; Fritsche et al., 2017; Hanks et al., 2011; Urai et al., 2018). In contrast, if coherence is strong, the influence of the prior is very weak. This is also referred to as precision weighting because the Bayesian model allows to give more weight to the more reliable source of information.

Some studies have specifically investigated where prior belief updating is represented in the brain (O'Reilly et al., 2013; Summerfield and De Lange, 2014; Summerfield and Koechlin, 2008; Vilares et al., 2012). It has long been suggested that in visual perception prior expectation plays an important role in how we interpret the world (Seriès and Seitz, 2013). A study in which participants had to make simple perceptual decisions has shown that prior belief is represented in higher visual areas and influences the activity in early visual areas (Summerfield and Koechlin, 2008). There is also evidence that downstream areas in frontal and parietal cortex are involved in updating the belief system when participants had to infer the location of a visual target (O'Reilly et al., 2013; Vilares et al., 2012). In Chapter 5, I set out to test whether human parietal cortex integrates prior belief across trials to build a priority map that informs saccade choice on the current trial. Using a Bayesian model, I can explicitly calculate how participants should update their prior belief to optimally solve the task and test whether this is reflected in the activity of parietal cortex.

An important feature of the brain is its ability to learn, which can also be described as basing current behaviour on past experience to reduce uncertainty. Learning is an ongoing process in the brain because the natural world around us changes constantly and we need to adapt our expectations based on the current state of the environment, weighing recent evidence more than older evidence (Glaze et al., 2015; Veliz-Cuba et al., 2016). For that reason, it would be suboptimal to build our prior belief of the current state of the environment on the entire evidence presented in the past.

In Chapter 5, we design an experiment in which the outcome probabilities in a 2AFC task change over time and participants should only use more recent trials to calculate the prior expectation of which outcome is more likely. There are periods in this task in which choices A and B are equally likely and periods in which either choice A or B have a higher probability to be correct. In order to predict the outcome probability, participants should use the veridical feedback after each trial. They should only use the most recent trials to build up their prior belief about the outcome probability because otherwise they would not be able detect the change points at which the outcome probability changes. The rate of change in the

outcome probability is given by the hazard rate and determines the average number of past trials that should be used for calculating the prior expectation. It has already been shown that the anterior cingulate cortex can learn and represent the volatility of how fast environments change (Behrens et al., 2007). The focus of Chapter 5 is more on the neural processes that represent prior belief in parietal cortex.

A Bayesian approach is a normative model because it tells us how to infer the state of the environment optimally given the evidence. In the next section, I will discuss descriptive models of decision-making that allow us to make predictions about the process of how the brain might encode the representation of the likelihood and prior expectation. For example, these models can explain how prior belief might influence the evidence accumulation process represented in the ramping activity of neurons during decision-making (Hanks et al., 2011).

1.3.2 The drift-diffusion model

Evidence accumulation models such as the drift-diffusion model (DDM) allow us to better understand the behaviour and neural processes that underlie decision-making. 2AFC tasks capture many aspects of decision-making because they force participants to make a choice based on noisy evidence that they have to integrate over time. Evidence accumulation models are capable of optimally solving these tasks (Bogacz et al., 2006). They have the advantage that the noise can be averaged out through the process of evidence integration until a threshold is reached, at which point there is enough accumulated evidence to commit to a choice. In addition, they are able to explain behaviour and neural activity during decision-making (Brown and Heathcote, 2008; Brunton et al., 2013; Ivanoff et al., 2008; Ratcliff and Rouder, 1998; Reddi et al., 2003; Van Veen et al., 2008).

All evidence accumulation models have their roots in signal detection theory (SDT). The goal of SDT is to determine whether a noisy evidence sample (e) stems from one of two states (s_1, s_2). These states are often defined in terms of a stimulus being present or absent, but could also be considered in terms of two competing hypotheses (e.g. motion left vs. motion right) Each of those two possibilities can

be described as a likelihood $P(e|s_1)$ and $P(e|s_2)$. Calculating the log-likelihood ratio of these ($log(P(e/s_1)/P(e/s_2))$) results in a decision variable. This is another way to express the likelihood, as in a Bayesian model. Based on whether the decision variable is below or above a threshold the source of the evidence sample is either assumed to be state 2 or state 1, respectively. If both states are equiprobable, then a log-likelihood ratio above 0 favours s_1 and below 0 favours s_2 (Green and Swets, 1967).

The shortcoming of the SDT is that it cannot capture the different dynamics at play in the process of forming a decision. Essentially, the decision process in SDT is only described by the decision variable. It does not capture how the evidence integration process unfolds over time. In order to understand how the decisions are formed in the brain, we need to be able to interpret the temporal dynamics of neural activity and how they are influenced by other factors such as prior beliefs and reward.

The SDT can be easily extended to encompass the accumulation of evidence over time. Put simply, the stimulus viewing period in the 2AFC random dot motion task can be considered a sequence of discrete samples of evidence of which each presents evidence for a certain choice. Calculating the likelihood ratio over all of these samples results in a decision variable. By using the log-likelihood ratio, we can write this as an addition, in which the likelihood ratio of each new evidence sample is added to a developing decision variable. By setting a positive and negative threshold the running total only continues until one of the two thresholds is crossed. The threshold that is crossed will determine the choice. This model is called a sequential probability ratio test (SPRT) (Wald, 1945). It is optimal because it minimizes the error rate while keeping the RT as short as possible (Bogacz et al., 2006). There is evidence that neurons in LIP compute the decision variable by adding the log likelihood ratio of each new incoming evidence, when evidence is presented as discrete samples (Kira et al., 2015; Yang and Shadlen, 2007). If evidence is displayed continuously as in a RDM task, the SPRT converges to become the DDM in which the decision variable becomes a time dependent integral (Bogacz et al., 2006).

In the DDM, the trajectory of the decision variable diffuses to a decision bound (Ratcliff, 1978; Ratcliff and McKoon, 2008). As in the SPRT model, the decision variable represents the difference between 2 choices. This means that the choice depends on whether the decision variable crosses the positive or negative decision bound. The decision variable is defined by a drift rate overlayed with white noise. The drift rate determines the mean trajectory of the decision variable and is mostly influenced by the stimulus strength. However, other factors such as prior belief and speed pressure can influence its starting point or slope, as well (Forstmann et al., 2008; Hanks et al., 2014, 2011; Mulder et al., 2012). This means, we can think about the DDM as a descriptive model of the decision process that can be optimal because both likelihood and prior are accounted for, allowing optimal inference as described by a normative Bayesian model.

The advantage of the DDM is its flexibility in accounting for different influences on the evidence accumulation process (Brunton et al., 2013; Forstmann et al., 2008; Shinn et al., 2019). In the decision-making literature the most important influence on the evidence accumulation is speed pressure. In the natural world we often have to make decisions under time pressure in which we might not be able to consider all the evidence. This leads to an increased error rate and is called the speed-accuracy trade-off. One way to think about this is by imagining that the distance between starting point of the accumulation process and the decision-bound is altered. If the distance is rather short, the subject has not enough time to accumulate enough evidence for an accurate decision. There is conflicting evidence for whether subjects adjust the starting point of the evidence accumulation process or the decision-bounds in order to adapt to different speed regimes (Forstmann et al., 2008; Hanks et al., 2014; Heitz and Schall, 2012; Ivanoff et al., 2008; Steinemann et al., 2018).

The biggest disadvantage of the DDM is the assumption that evidence is held constant throughout time and that each evidence sample equally contributes to the evidence accumulation process until the bound is reached. As discussed in the section about the Bayes model, in the natural world state changes are often uncertain, and the state has to be inferred from recent evidence history. Therefore,

we need a model that can adapt to situations in which the best strategy is to discount past evidence.

1.3.3 The Ornstein-Uhlenbeck process

The Ornstein-Uhlenbeck process (OU process) is a descriptive model related to the DDM which can adapt the evidence integration kernel to linearly discount past evidence (Bogacz et al., 2006; Busemeyer and Townsend, 1993; Ratcliff and Smith, 2004). In contrast to the DDM the OU process is able to capture behaviour when evidence is uncertain. Consider a signal detection problem, in which you are in the forest and trying to detect the bird song of a very rare bird species. There is uncertainty because of the background noise of other birds but there is also temporal uncertainty because you don't know whether and when you will hear the bird. In this scenario a DDM is not the optimal model because in order to identify the bird song you don't want to consider all the different birds you have heard since entering the forest. Instead, you should only consider the most recent songs over a time window that corresponds with the length of the bird song you want to detect. The OU process has a decay parameter that can be adapted to only consider most recent evidence depending on the temporal contingencies of the environment which makes it a suitable evidence integration model for such problems.

One set of paradigms that particular allows us to study this kind of accumulation processes are change-point detection tasks. In these paradigms there is at least one switch in the environmental state that subjects have to detect. For example, there are auditory click tasks in which subjects would listen to a stream of clicks and they had to detect changes in the click rate (Ganupuru et al., 2019; Harun et al., 2020; Piet et al., 2018). Further, this model is also useful for continuous paradigms in which subjects have to detect reoccurring signal in a stream of noise. In a study by Ossmy et al (2013) subjects had to detect signals of varying lengths in a stream of continuous background noise. They showed that subjects discounted past evidence in order to detect the signal periods by adapting the decay rate to the signal length.

The OU process belongs to the class of leaky accumulator models. These encompass a range of neurocomputational models that can explain neurophysiological recording data during evidence accumulation as well as psychophysical measurements like RTs (Mazurek et al., 2003; Usher and McClelland, 2001; Wang, 2002). They can explain behaviour with 2 and more choices, the influence of reward on decision-making and how subjects should weigh fluctuations in evidence within trials (Bogacz et al., 2007; Bronfman et al., 2016; Gao et al., 2011; Levi et al., 2018; Ossmy et al., 2013; Piet et al., 2018; Roe et al., 2001; Tsetsos et al., 2011, 2012; Usher and McClelland, 2004). All biophysical leaky accumulator models have in common that they consist of two input units which provide noisy evidence that is integrated by two leaky integrator units, which means that the activity of the integrator units decays over time. In addition, all of these models are characterised by mutual inhibition which is implemented either through a pool of inhibitory interneurons (Wang, 2002), direct inhibition between the integrators (Usher and McClelland, 2001) or via the input units (Mazurek et al., 2003). However, these models are computationally equivalent to each other for certain parameter values in a 2AFC task and reduce to the OU process (Bogacz et al., 2006; Ratcliff and Smith, 2004; Usher and McClelland, 2001). For this reason, I am choosing the OU process and I explain in Chapter 2 why it is a suitable model for my task paradigm.

In the OU process the rate of change (dX/dt) depends on the previously accumulated evidence. Previously accumulated evidence is reflected in the value of the decision variable (X) before the current evidence sample is added. When calculating the rate of change, the current value of the decision variable is scaled by an additional parameter λ which defines the behaviour of the OU process. This weighted decision variable of previously accumulated evidence is then added to the current evidence sample. A negative λ means that the decision variable is decreased in the oppositive direction of its current value. In other words, in the absence of any additional evidence, it will decay towards 0 (which is a 'stable attractor' of the dynamical system). In this case the development of the decision variable only depends on the most recent evidence samples that have been observed. In contrast, if λ is positive, then the decision variable will naturally increase in the same direction as the previously accumulated evidence. This

means that the effect of evidence samples observed further back outweigh evidence samples seen more recently in determining the development of the decision variable, leading to a run-away process of the decision variable (exponential growth, or 'unstable attractor'). If λ equals 0 the OU process reduces to a DDM, meaning that every evidence sample observed contributes equally to the evidence accumulation process (Bogacz et al., 2006; Busemeyer and Townsend, 1993; Ratcliff and Smith, 2004). In other words, the OU process behaves as a stable attractor when λ is negative and there is no or minimal noise. This makes the model more robust to noise, because even if the noise pushes the model away from the stable state, it nevertheless always moves back in the direction of it.

The task I developed to study decision-making in uncertain environments consists of constantly changing visual evidence and participants have to detect certain stable periods. This is conceptually similar to detecting the bird song of a rare bird species in a forest full of other birds. My study is novel in that both trial length over which participants have to integrate varies as well as the frequency with which the trial periods appear. This leads to unique predictions of how the decision thresholds and the decay rate (λ) should be adapted. In addition, I am not only looking at behavioural but also neural correlates by simultaneously recording EEG while participants are doing the task. In the next chapter, I am going to explain the task paradigm in detail and discuss why the OU process is suitable to model subject behaviour on this task. This is the foundation on which I will discuss the behavioural and EEG results from human participants in Chapters 3 and 4, respectively.

2 A new experimental paradigm to study evidence integration in continuously varying environments

Unlike in most psychology experiments, in everyday life we do not make decisions in confined trials. Instead, we have to continuously accumulate evidence and decide when we have sufficient evidence to commit to a choice. Depending on the circumstances, we need to adapt the length of the period over which to integrate evidence. For example, the decision to switch lanes on the motorway will require only seconds. By contrast, the decision to sell stocks will require longer deliberation, as you would want to study the underlying stock market for days or months.

In this chapter, I introduce an experimental task and a formal mathematical framework to investigate how people integrate evidence in a dynamic world that changes constantly. The task is based on the 2-alternative-forced-choice (2AFC) random dot motion (RDM) task, which has been extensively studied in the past and has been used to identify the neural mechanisms of evidence integration in perceptual decision-making (Gold and Shadlen, 2007). In the 2AFC RDM task, ramping activity of single neurons (Huk and Shadlen, 2005; Mazurek et al., 2003; Roitman and Shadlen, 2002), populations of neurons (Bollimunta et al., 2012), or scalp electroencephalography (EEG) potentials (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015) signal evidence integration in favour for a choice alternative (for more details, see Chapter 1). This ramping activity depends on the signal-to-noise ratio (SNR) in the stimulus and the duration of a single trial. The neural activity is predictive of the subject's choice behaviour and reaction time (Gold and Shadlen, 2007). However, in humans it has not been well studied how these neural signals might adapt their timescales of evidence integration to different dynamic environments. For this reason, I developed a modified RDM task, to study decision signals in continuous environments with EEG.

The modified RDM task lacks the trial structure of traditional RDM tasks meaning that participants continuously watch a stream of noise that is interleaved with 'trial-like' periods in which signal is added to that noise (Figure 1). Participants then

have to report whenever they believe that they are currently in a trial period with signal. Unlike in the classic RDM task, a drift diffusion model (DDM) will not allow for an optimal solution of this task, because the decision variable would equally depend on the entire history of evidence presented – that is, its value would be influenced by the past noise stream as much as the recent sensory evidence. The optimal strategy in solving this task is instead to adapt the evidence integration kernel, i.e., the weighted average, to the length and frequency of the trial periods. For example, if trial periods typically last 3 seconds, then the decision variable should disregard the evidence presented more than 3 seconds ago.

Leaky accumulator models allow for past evidence to be forgotten and to only consider the most recent evidence (Bogacz et al., 2006; Ratcliff and Smith, 2004; Usher and McClelland, 2001). A leak term controls how past evidence is discounted. This separates the model from a normal DDM, which doesn't have a parameter to control the weighting of past evidence (Bogacz et al., 2006). Previous behavioural studies in rodents and humans have suggested that leaky accumulator models can account for continuous evidence integration in a dynamic environment (Glaze et al., 2015; Ossmy et al., 2013; Piet et al., 2018). However, none of these studies can extract the subjects' integration kernels directly from their behaviour. As I will show in this chapter, the modified RDM task design allows one to read out the integration kernel directly. Once it is possible to measure this integration kernel, one can then ask questions about whether it differs across different task environments. In particular, I hypothesised that these kernels would change as a consequence of altering the *frequency* and *length* of the trial periods embedded in the noise (see also Ossmy et al., 2013). I show that a leaky accumulator model that maximises reward rates in the task adopts different integration kernels for these different task environments.

In addition, previous studies have not studied the neural signals underlying leaky integration in continuous evidence integration tasks. I acquired EEG data while participants completed the task. This allows me to study signals of decision formation in a more natural setting that can also capture the relationship between these signals and behaviour. Previous EEG studies suggested a centro-parietal positivity signal tracks evidence integration (Kelly and O'Connell, 2013; O'Connell

et al., 2012; Twomey et al., 2015). However, these studies altered the noise structure continuously, and had a constant trial length and frequency. Crucially, these studies were not able to investigate how the centro-parietal positivity evolves during 'noise' periods, because they did not place the noise structure under experimental control. I hypothesise that this centro-parietal positivity is influenced by the evidence integration kernel that subjects use in the experiment, and this in turn depends upon the length of the trials and their frequency.

In this chapter I first introduce the modified RDM task paradigm (section 2.1). I then use computational modelling of a leaky integrator to make predictions of the behavioural and neural outcomes (section 2.2). I show that trial length and frequency modulate the key parameters in a leaky integrator model in order to optimise reward. This gives rise to the behavioural and neural predictions that I will test in Chapters 3 and 4.

2.1 Task design

2.1.1 Task properties of a continuous random dot motion task

In the continuous version of the 2AFC RDM task, participants observe a stream of randomly moving dots in a circular aperture (Figure 1). Some of these dots will move coherently to the left or to the right. The coherence is the proportion of dots moving in the same direction, whereas the other dots move randomly. In this study the coherence varied between -1 (all dots move to the left) and 1 (all dots move to the right). At 0 coherence, all dots move randomly.

Unlike in trial-based versions of the task, and even previous RDM tasks where motion is continuous (e.g., Kelly and O'Connell, 2013), during the present task the coherence *changes* constantly. During 'intertrial' periods, the *average* of these constantly changing values remains 0. Within this stream of constantly changing coherence, there are trial periods in which the *average* coherence is either to the left or to the right (see Figure 1). The aim of the subject performing this task is to detect these 'stable' periods of predominantly leftward or rightward motion (trials). Exact details of how this noise structure is generated are given below.

For participants, this means that they should not make a decision as soon as they think they know in which direction the dots are moving coherently, as it is the case for a discrete trial version of the RDM task. Instead, participants have to weigh the recent history of motion directions in order to detect periods where the average motion direction of the dots is consistently going to the left or right.

Participants indicate their decision about the average motion direction by pressing keyboard button 'L' to report a trial period with average motion to the right and 'A' to report a trial period with average motion to the left. Figure 1 shows an example period of a stream of randomly moving dots. During 'intertrial' or 'noise' periods, dots can move randomly either to the left or right with various coherence levels. Once a trial period starts, they move more to one side than the other, on average. This is indicated by circles of moving dots which are examples of possible changes during trial and intertrial periods (Figure 1). Every time a button is pressed a change in colour of the central fixation point indicates the feedback. A correct response is indicated by a green fixation point on the next frame, whereas an incorrect response during the trial is followed by a red fixation point. False alarms, i.e., button presses between trial periods, are indicated by a yellow fixation point. If a trial is missed the fixation point will turn blue after 500ms (a button press within these 500ms is still counted as a successful response to the trial, in order to account for non-decision time and allow participants to integrate over the full length of the trial). Note that following correct or incorrect responses (made during a 'trial period'), the trial is terminated immediately, and the stimulus returns to baseline.

Participants are rewarded for correct responses but lose points for any other response. They receive 3 points for correct responses and are punished with -3 points for incorrect responses. Missed trials and false alarms are punished with -1.5 points. A reward bar is shown beneath the stream of continuously moving dots, to indicate how many points participants have won. As long as participants win more points on average than they lost, their reward increases to the right border of the reward bar. Every time they hit the right border of the reward bar, which equals 15 points, the bar is reset, and they receive £0.50 to take at the end of the experiment. In the rare cases where participants hit the left border of the reward bar (-15 points), participants have £0.50 deducted from their take-home pay. This

reward structure encourages participants to continuously pay attention to the movement direction of the dots and to avoid responses to the wrong direction during a trial.

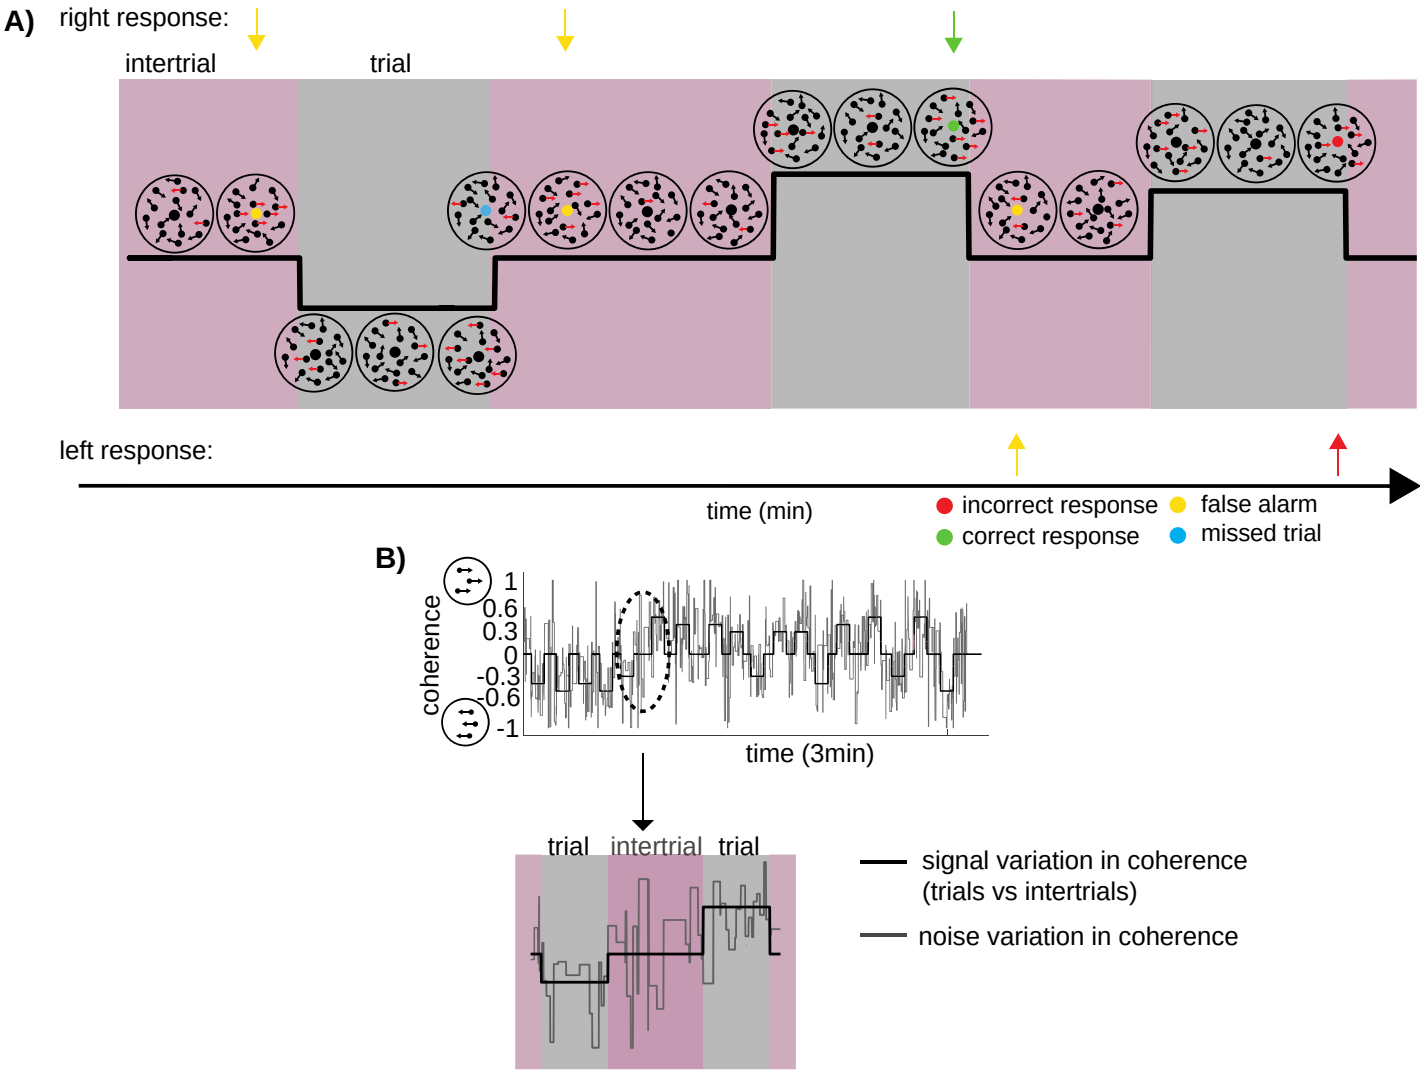

A)
right response:
intertrial
trial
left response:
time (min)
incorrect response
correct response
false alarm
missed trial
B)
coherence
1
0.6
0.3
0
-0.3
-0.6
-1
time (3min)
trial
intertrial
trial
signal variation in coherence
(trials vs intertrials)
noise variation in coherence

Figure 1 (previous page): Task structure of the continuous RDM paradigm. A) Participants observe a stream of randomly moving dots, with interleaved stable trial periods they have to detect. Shown is a snippet of such a stream of moving dots. During intertrial periods (red) dots move equally to the left or right, whereas during trial periods (grey) dots move more often in one direction than the other. Participants can respond with a button press of either L or A on a keyboard if they think that they are either in a leftward or rightward stable trial period. Participants receive feedback about their button presses through a changing fixation point and win or lose points accordingly. Underneath the stream of moving dots a reward bar shows the number of points won. **B)** Structure of the noise in the task paradigm. The noise structure is partly the driving force for participants to adjust their integration kernels to detect trials, and also the means by which to measure behavioural integration kernels and EEG responses during the intertrial period. Participants have no knowledge of the mean coherence (black line) that describes the trial periods (every time the mean is not zero). To infer whether they are in a trial they have to integrate over the noise (grey line), which describes the coherent movement of dots participants observe at any point in time. The noise consists of little steps in coherences (grey line) and the coherences are drawn from a random distribution with a mean that is dictated by the black line. This means that every time the mean coherence is not zero there will be a shift in the noise towards that mean and fluctuate around it. Importantly this noise is under experimental control, allowing us to interrogate the relationship between fluctuations in noise and both behavioural and neural data.

2.1.2 Trial length and frequency should influence the optimal strategy of evidence integration

The constantly varying coherence levels in the task force participants to integrate the motion direction of the dots over an extended time period in order to gather enough information to decide whether they are in a trial. If a certain threshold of evidence is reached in that period participants should respond. Yet in order to maximise their accuracy, we hypothesised that participants should adapt the time period for evidence integration to match the length of the trial periods. This weighting of past evidence might be reflected in an exponentially decaying integration kernel (Figure 2). They should also be more likely to commit to making a response if they know that trial periods are more frequent.

In order to understand whether participants can flexibly adapt their integration kernel, I tested the continuous evidence integration task under different conditions following a within-subject 2x2 design. In different blocks, participants were told that they would either have long (5 seconds) or short (3 seconds) trial periods, and either frequent (intertrial period range 3s-8s) or rare (intertrial period range 5s-40s) trials. Subjects received extensive training [details in Chapter 3] so that these

environmental statistics were well learnt prior to the experimental session. First, I will explain how these different conditions should intuitively influence the integration kernel participants might ideally use. Then, I will test these predictions by simulating the evidence integration process with a leaky accumulator model.

We hypothesised that trial length would have the most profound effect on the integration kernel (Figure 2). If trials are short, participants have less time to think about whether they are currently in a trial, whereas if trials are very long, they can take more time to integrate sensory evidence in order to confirm that they are actually in a trial. In other words, participants should discount past evidence faster when trials are short, leading to a steeper decay of the integration kernel (Figure 2B). Behaviourally, this means that participants should have longer reaction times for long trials and a higher detection rate. Short trials will be easier to miss, leading to a decrease in the detection rate.

In contrast to the trial length, the frequency of trials should influence how cautiously participants should respond. If trials are frequent participants should be less cautious because the probability of being in a trial is high. Accordingly, participants should increase their number of button presses, resulting in a higher total number of correct trials, but also potentially a higher false alarm rate. Conceptually this means that participants should adapt their decision threshold. The decision threshold determines how much evidence needs to be accumulated before the participant commits to a response (Figure 2B). For frequent trials participants should lower their decision threshold to minimise the probability of a missed trial. In contrast, a low decision threshold would lead to a high false alarm rate when trials are rare. Consequently, participants should increase their decision threshold.

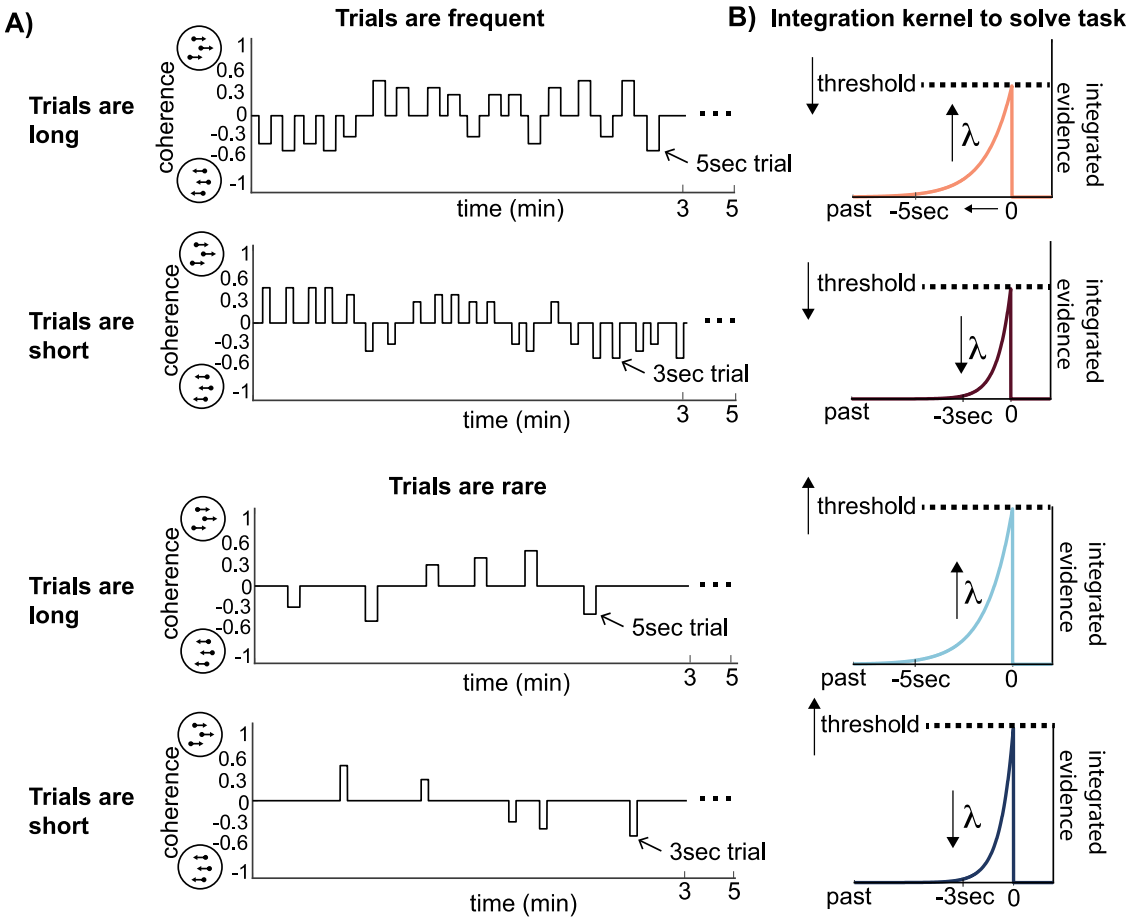


Figure 2: Task features and hypothesised integration kernels. A) To manipulate the integration kernel participants should use to solve the task optimally we introduced 4 different conditions following a 2x2 design. Trials could either be *long* or *short*, and *frequent* or *rare*. **B)** Depending on the condition, participants should either alter their decision threshold and decay (λ) over past evidence. Long trials should lead to a decay that is stretched in time in comparison to short trials. For frequent trials, participants should accumulate evidence to a smaller threshold in order to avoid missing trials than for rare trials.

2.1.3 The noise structure of the task

A very important feature of this task paradigm is its noise structure, which leads to the continuously varying coherence levels as shown in Figure 1B. Importantly, the noise is placed under experimental control rather than being randomly generated, meaning that we can examine how fluctuations in the noise impact upon participants' behavioural and neural data. For example, the noise structure provides us with a direct window on evidence integration. By averaging the noise preceding a behavioural response (in particular for false alarms, which occur when

the motion has zero mean coherence) we can measure the degree to which recent samples of the stimulus trigger a subject's response. Therefore, we can use features of the noise structure to directly infer the integration kernels participants have used in different conditions from the behavioural data. In addition, it enables us to investigate the influence of each new incoming piece of evidence (i.e. each change in the coherence level) on the subsequent EEG signal, using a convolutional general linear model (GLM) analysis strategy (Ehinger and Dimigen, 2019; Lalor et al., 2006; Smith and Kutas, 2015a, 2015b) that I describe in further detail in Chapter 4.

More generally, the noise can be described as a series of steps that vary in duration and coherence (Figure 1B). The step duration was sampled from an exponential distribution with a mean duration of 100ms. This distribution was then truncated, with a minimum duration of 10ms and a maximum duration of 250ms for each step. Each step in the noise is at a different coherence which is sampled randomly from a normal distribution, with a standard deviation of 0.3 for trial periods and 0.5 for intertrial periods. Steps whose resulting coherence were greater than 1 (or less than -1) were set to be 1 (or -1) respectively. The mean of the normal distribution depends on whether the step occurred during an intertrial period or during a trial. During an intertrial period the mean of the normal distribution is 0. That means it is equally likely that negative or positive coherences are drawn. During trial periods, the mean of the normal distribution was sampled uniformly from the set [-0.5, -0.4, -0.3, 0.3, 0.4, 0.5]. This means that during trial periods, the dots move on average more to one side than to another. The further the mean of the normal distribution is away from 0, the stronger the average motion coherence is and so the easier it should be to detect a trial period.

A step-like noise structure was chosen instead of white noise to ensure that participants constantly perceive different percepts of motion direction that are stable for a short time. During piloting, we found that high frequency white noise does not induce percepts of fast changes in motion direction and would have made the detection task too easy. In addition to making the task more difficult and forcing participants to integrate, the step-like noise structure allows us to measure how the recent time-varying evidence influences both subjects' propensity to respond

(during both trial and intertrial periods), and also influences subthreshold decision-signals in the brain. This is possible because we can time-lock to individual steps in the coherence and analyse the EEG signal corresponding to changes and magnitudes of these steps.

2.2 Computational Modelling

2.2.1 The Ornstein-Uhlenbeck process can capture the different features of the task

The continuous RDM task paradigm requires the integration of evidence over long periods to detect trial periods. An agent has to be able to integrate new incoming evidence on a constant basis and make a decision as soon as there is enough evidence to reach an internally set threshold. However, in order to solve the task, the agent should not take into account all past evidence, because this could cause the agent to reach the decision threshold prematurely. Therefore, the best strategy is to only consider the past evidence presented in the number of samples that make up a trial period and to 'forget' previous evidence. There are different generative models that describe the incremental increase or decrease of evidence variable that follows each new sample of evidence.

In previous literature, evidence integration has often been described by means of a drift diffusion model (DDM). The DDM is a generative sequential analysis model in which incoming evidence is sequentially added up (Ratcliff, 1978; Ratcliff and Rouder, 1998). A random walk is superimposed on this linear process to account for noise. In a 2AFC task the resulting decision variable drifts over time to one of two bounds. Once one of the bounds is reached, a decision is made. The change in the decision variable is described as follows:

$$dx = I_t * dt \ + c^2 * dt * W \quad (1)$$

Where *dx* is the change in the decision-variable over time interval *dt*, *I* is the input, i.e. the strength of the evidence which determines the drift rate and *W* is gaussian white noise, scaled by the variance $c^2 * dt$. Because of the white noise the decision

variable diffuses and does not develop to a choice boundary in a straight line. The sign of *dx* indicates to which choice alternative the diffusion process is drifting. After each time step the decision variable is updated by *dx* until a decision-threshold *z* is reached, at which point the model elicits a response. One of the reasons why this model is widely used is because of its ability to predict both behaviour and neural signals.

The DDM makes accurate predictions about accuracy and reaction times under different task regimes of the 2AFC task (Hanks et al., 2006; Huk and Shadlen, 2005; Mazurek et al., 2003; Palmer et al., 2005; Ratcliff and Rouder, 1998; Rorie et al., 2010). Extensions such as varying the starting points of the drift-diffusion process (Rorie et al., 2010) and collapsing bounds (Drugowitsch et al., 2012; Murphy et al., 2016) can account for behavioural variability for example in choice biases and for the speed-accuracy trade-off (Bogacz et al., 2010; Forstmann et al., 2008; Mazurek et al., 2003; Ratcliff and Rouder, 1998; Usher and McClelland, 2001). Furthermore, it is possible to fit the model directly to neural data and to investigate how the different parameters are encoded in the brain (Brown and Heathcote, 2008; Ivanoff et al., 2008; Reddi et al., 2003; Van Veen et al., 2008).

However, one limitation of the DDM is that its decision variable is the summation of the entire past evidence (Equation 1). This is a result of the fact that, in the absence of noise, the DDM model will never decay back to baseline, as *dx* has no dependence upon the current value of *x* (see also Figure 3A below, $\lambda = 0$). This is not quite plausible in real-life scenarios when it would be unfavourable to include information in our decision-making process that is outdated. For example, when we want to buy stocks of a company, we are probably not interested in the performance of the company 5 years ago. Consequently, the DDM model cannot adequately reflect the current evidence in the continuous RDM task paradigm. It is only a useful model for single-trial task paradigms, in which the input to the model (defined by the strength of the coherence) is always relevant to the final decision.

In contrast to the DDM, the Ornstein-Uhlenbeck process (OU process) is able to adapt to different time intervals of evidence integration based on the underlying decision (Equation 2). This is because the derivative of the decision variable in the

OU process not only depends on the input, but also on its current value, discounted by a learning rate which we constrain to be negative (λ), as can be seen below (Busemeyer and Townsend, 1993):

$$dx = (\lambda * x + I) * dt + c^2 * dt * W \quad (2)$$

Where *dx* is the change in the value of the decision-variable. The current value of the decision-variable (*x*) is scaled by λ. *I* defines the drift rate, as in the DDM (Equation 1). When λ<0, the derivative of the OU process can be described as a weighted average of past evidence. Assuming a negative λ (and *c* = 0) the derivative pushes the decision variable always towards the x-axis intersection which is a stable attractor under these circumstances (Figure 3A). If the decision-variable is negative, the derivative is positive and vice versa. More negative values of λ lead to a faster approach to the x-axis intersection, where a stable attractor exists as *dx*=0 (Bogacz et al., 2006). However, if λ = 0, the OU process reduces back to a standard DDM. Note that we do not consider cases where λ > 0, in which past evidence is amplified across time rather than forgotten (Bogacz et al., 2006).

The influence of different values of λ on the integration process of a simple stimulus of the OU process is also illustrated in Figure 3B. For illustration purposes, I normalised all models by dividing them by their maximum value (although note that from Figure 3A, it can be seen that different values of λ do in fact lead to different stable attractor values for *x*). The decision-variable *dx* changes its value in the direction of the input for all models but the temporal dynamics differ. The more negative λ, the faster the OU process adapts to the change in input and decays faster back to baseline, meaning that past evidence is 'forgotten' faster. This would also be true for my task-paradigm, in which the evidence does not step 'on' and 'off' as in Figure 3B, but instead varies continuously over time. For these reasons, I expect that the OU process can optimally solve the different conditions of my task-paradigm minimizing the number of false alarms and missed trials. Based on the hypotheses in the previous section this means that λ should adapt to the trial length. For short trials λ should be more negative than for long trials. In contrast, trial frequency should be more influenced by adding decision thresholds that define

when the OU process has reached a decision. In general, the decision thresholds should be higher when trials are rare, because subjects should be more conservative in their responses in these conditions. How the decision thresholds are implemented is described below.

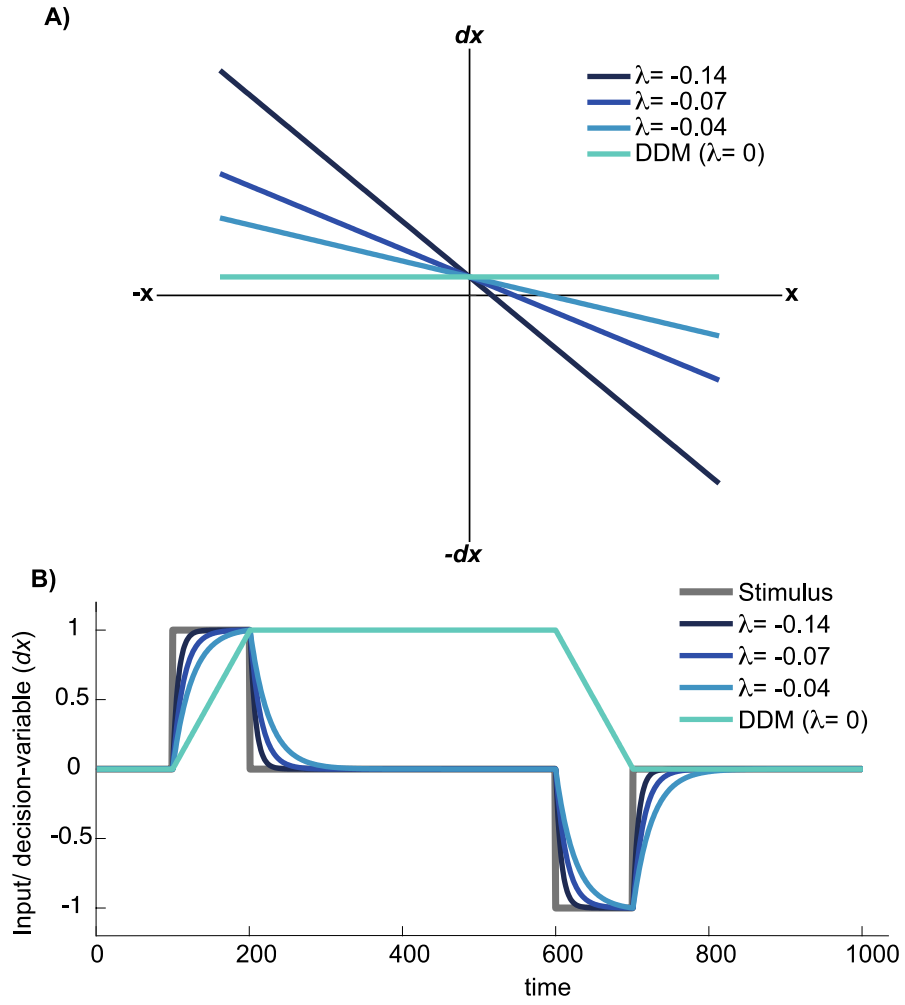


Figure 3: The behaviour of the OU process with different decay rates. **A)** The derivative of the OU process (*dx*) pushes the decision-variable (*x*) towards the intersection with the x-axis under negative λs. If *x* is very negative, the derivative of the OU process is positive pushing the system towards the x-axis intersection, and vice versa for very positive *x*. At the point where the x-axis is intersected, then *dx* equals 0 and *x* is held at a constant value (except for noise). **B)** depicts how the OU process integrates over changing input under different decay rates (λs). Note that this simulation is shown in the absence of any noise. All model instantiations were normalised by dividing each instantiation by its highest value.

2.2.2 The OU process predicts task dependent decay rates and decision thresholds

In order to formally test whether the OU process predicts different integration kernels for different conditions I ran a parameter search to identify the decay rate and threshold that maximised overall reward using Equation 2. I tested a range of thresholds and negative λs and used stimulus streams of the actual experiment as input. As we were not interested in variability between different instantiations of the model, the noise was scaled to a very low value (c = 0.01).

Every time the decision variable reached the decision threshold (+z or -z), the decision variable was reset to 0 and the outcome was either rewarded with points or penalised, as if it were a subject. If the threshold was reached during a trial period, the rest of that period was replaced with noise at baseline, as in the actual paradigm. For each condition I ran 90 different stimulus streams per condition (this corresponds to the number of streams observed by 15 subjects) and repeated each 50 times. I recorded the number of points won for each λ and z and the number of incorrect, correct, missed trials and false alarms. I then calculated the average number of points for each λ and threshold per stimulus stream (Figure 4). The best λ and threshold for each condition are those for which most points were won.

Figure 4A shows the reward obtained by an OU process for an example stimulus stream in each condition. The first column depicts the average number of points for each λ and z pair. This is a function of the total number of correct responses, false alarms, missed responses and incorrect responses, which are shown in the remaining four columns. In the first column, the optimal value (denoted with a black point) corresponding to the most points won shifts slightly to the upper border of the figures for long trials, meaning that long trials have less negative λs for a longer decay back in time. This is confirmed by significant differences in the average best λs across stimuli shown in Figure 4B (long versus short trial conditions: $F(1/356)$ = 28.8, $p < 1e^{-6}$; frequent versus rare trial conditions: $F(1/356)$ = 12.5, $p = 1e^{-4}$; interaction between frequency and length: $F(1/356)$ = 0.69, p = 0.41, tested with a 2-way ANOVA).

As predicted z also varies with task regime. It is higher for rare trials than for frequent (Figure 4C). This is consistent with my hypotheses that the model should be more cautious when trials are rare, and collect more evidence to confirm that it is in a trial period. Nevertheless, the best estimated z is not only dependent on trial frequency but also increased for long trials in comparison to short ones (Figure 4C). There are significant differences in the best z for different trial conditions (trial length: $F(1/356) = 99.7$, $p < 1e^{-21}$; trial frequency: $F(1/356) = 123.2$, $p < 1e^{-24}$; interaction between trial frequency and trial length: $F(1/356) = 9.1$, $p = 0.003$) This indicates that there is a relationship between λ and z. In theory a higher z means that the subject has to integrate for longer. The difference to λ is that λ influences the slope of the kernel while z does not. This means that z and λ are mechanistically quite different. By adjusting both z and λ the model can achieve a compromise between false alarms which increase with more negative λs and the number of missed trials that increase with higher z (Figure 4A).

There is not an optimal solution in which there are neither misses nor false alarms. The number of incorrect responses is negligible for all conditions. A correct detection of the motion direction is necessary at all times to detect the stable trial periods. If participants could not discriminate between the changing motion directions well, they could not infer whether they are in a stable period or not. This is different to the classic 2AFC task in which the evidence integration process consisted of detecting the motion direction of the dots.

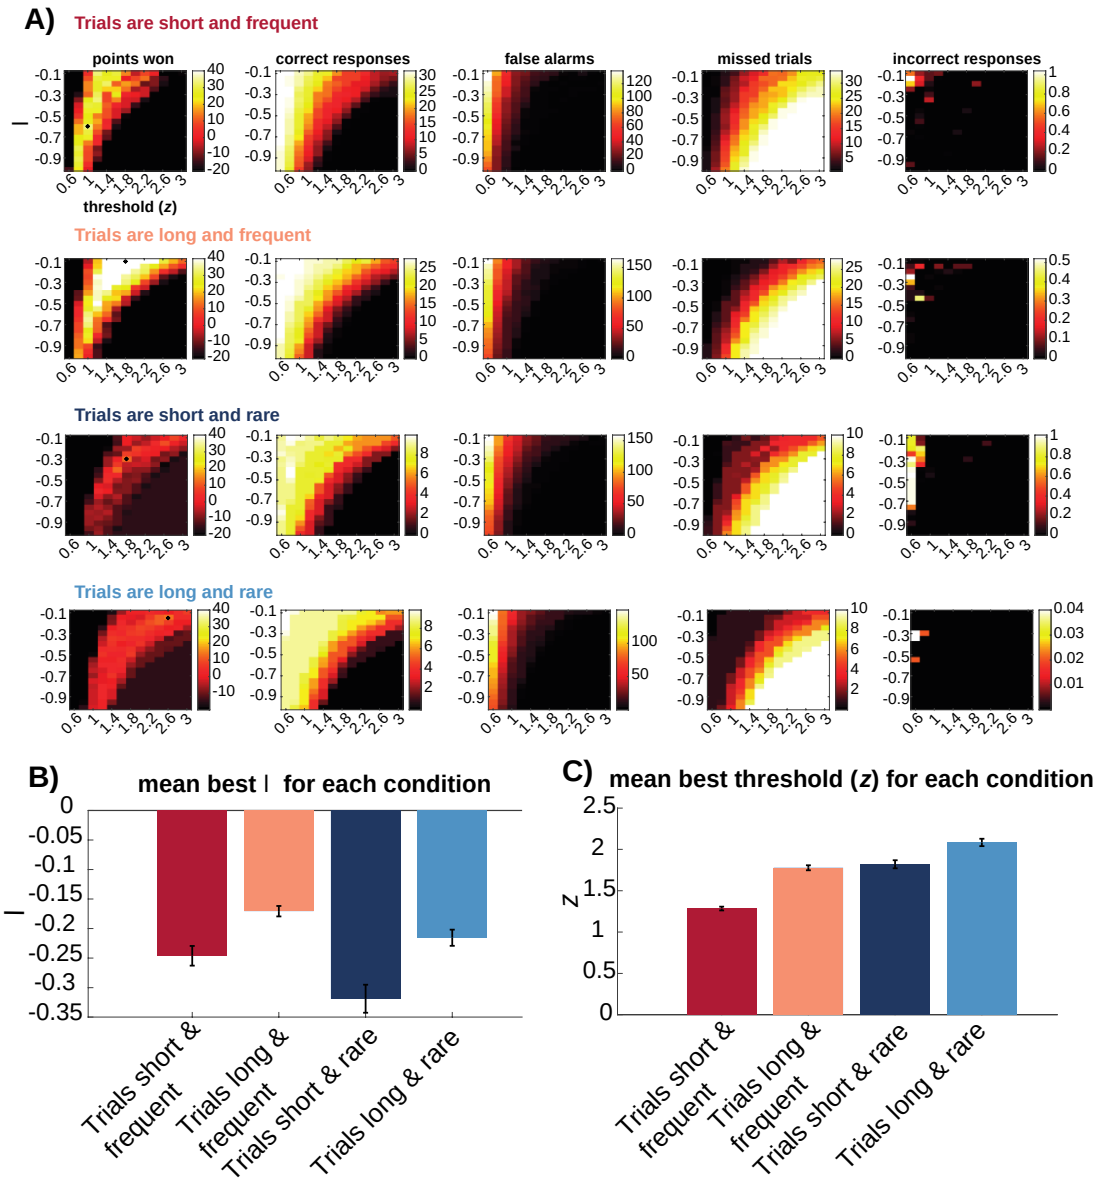


Figure 4: Model behaviour and best parameter estimates for different task conditions. The OU process was run for different λs and z across stimuli. z is the decision threshold at which the OU process is reset to 0 and a response is recorded. The noise was kept low with a scaling of c = 0.1. **A)** Shown are heatmaps for the model behaviour across different λs and zs for different behavioural parameters. Each row shows these outcomes for a single stimulus stream of a different condition. For each λ and threshold the model evaluation on the stimulus stream was repeated 50 times. The first column shows the distribution of the average number of points won. The model was rewarded in the same way humans were. The higher the number of points won the more optimal the model behaved. Best model parameters are indicated by a black point. The other columns show the average number of false alarms, correct trial responses, missed trials and incorrect trial responses. **B)** Mean best λs for each condition with standard error of the mean (s.e.m) **C)** Mean best z for each condition with s.e.m. The model evaluation of the task was repeated on 90

different stimulus streams per condition. For each stimulus stream λ and z that yielded on average the highest number of points were selected and subsequently averaged. A 2-way ANOVA was performed on both λs (trial length: $F(1/356) = 28.8$, $p < 1e^{-6}$; trial frequency: $F(1/356) = 12.5$, $p = 1e^{-4}$; interaction between frequency and length: $F(1/356) = 0.69$, $p = 0.41$) and zs (trial length: $F(1/356) = 99.7$, $p < 1e^{-21}$; trial frequency: $F(1/356) = 123.2$, $p < 1e^{-24}$; interaction between trial frequency and trial length: $F(1/356) = 9.1$, $p = 0.003$).

Under optimal λs and zs for the OU process there are significant differences in behavioural measures between conditions. Figure 5 shows that the false alarm rate, detection rate and median reaction time (RT) vary with trial length and frequency. The false alarm rate decreases with trial length and rare trial conditions, which means that the model is more 'cautious' when trials are rare and trials are long (trial length: $F(1/356) = 155$, $p < 1e^{-29}$, trial frequency: $F(1/356) = 299$, $p < 1e^{-47}$, interaction trial length and frequency: $F(1/356) = 100$, $p < 1e^{-20}$ from a 2-way ANOVA). There is a significant effect of trial frequency ($t = 12.4$, $p < 1e^{-25}$), trial length ($t = -12.9$, $p < 1e^{-26}$) and coherence ($t = 20.6$, $p < 1e^{-40}$) on detection rate tested with a GLM with a constant regressor, coherence regressor [0.3, 0.4, 0.5], trial frequency regressor [frequent: 1, rare: -1] and trial length regressor [short: 1, long: -1]. The detection rate improves for long and frequent trials. Theoretically this can be explained by longer integration kernels for long trials and lower decision-thresholds for short trials. Within each condition performance increases with coherence strength as is typical for random dot motion paradigms. In addition, shorter RTs for short and frequent trial conditions also indicate that the model seems to adjust the decay and the decision threshold according to the condition (2-way ANOVA on RTs: trial length: $F(1/356) = 439$, $p < 1e^{-63}$ and trial frequency: $F(1/356) = 122$, $p < 1e^{-23}$, interaction length and frequency: $F(1/356) = 0.13$, $p = 0.72$).

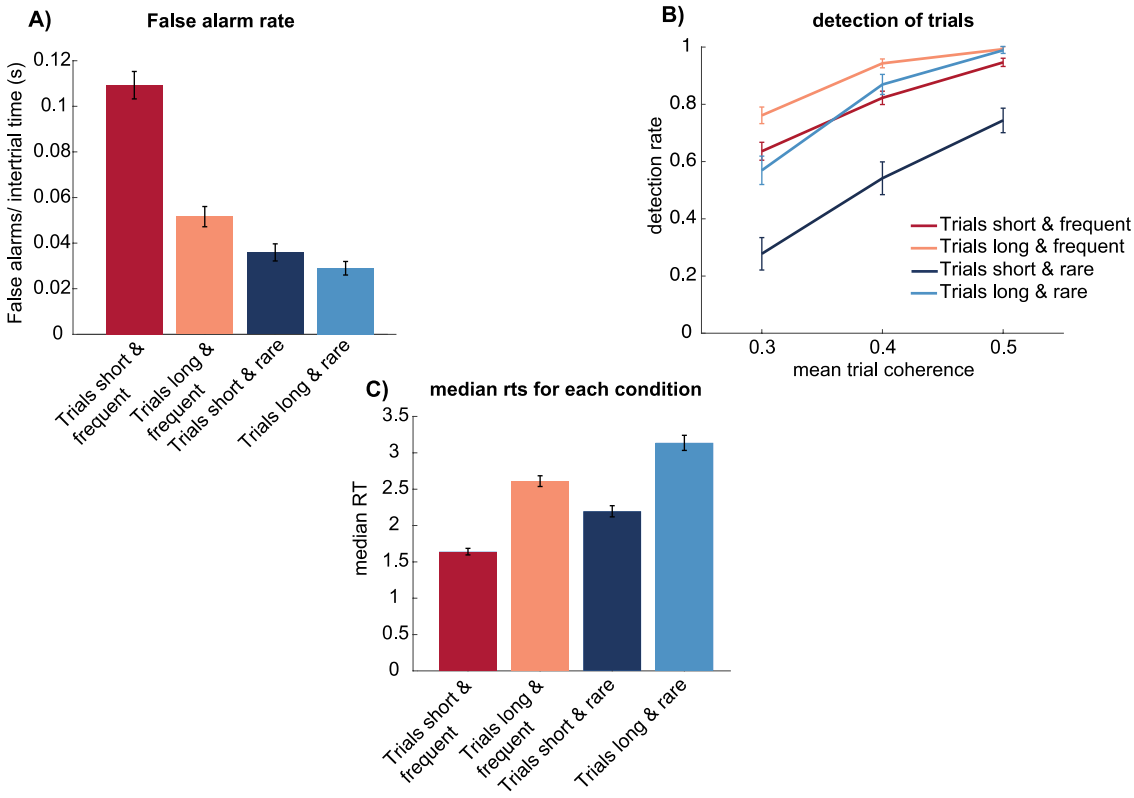


Figure 5: The behaviour of the OU process varies across conditions. A) For each simulated stimulus stream of a condition the ratio of number of false alarms and intertrial time was calculated. Shown are the mean and standard error of the mean (s.e.m.). There is a significant effect of trial frequency and length on false alarm rate (trial length: $F(1/356)$ = 155, $p < 1e^{-29}$, trial frequency: $F(1/356)$ = 299, $p < 1e^{-47}$, interaction trial length and frequency: $F(1/356)$ = 100, $p < 1e^{-20}$ from a 2-way ANOVA). **B)** The detection rate was calculated as ratio of correct responses to trials and total number of trials of a certain coherence for each condition. Shown are the mean + s.e.m. A GLM analysis with regressors for coherence, trial length and trial frequency was performed, revealing significant effects for all regressors (trial frequency: (t = 12.4, $p < 1e^{-25}$, trial length: t = -12.9, $p < 1e^{-26}$ and coherence: t = 20.6, $p < 1e^{-40}$). **C)** Median rt distribution for correct trials across conditions. Shown are the mean + s.e.m. (2-way ANOVA on RTs: trial length: $F(1/356)$ = 439, $p < 1e^{-63}$ and trial frequency: $F(1/356)$ = 122, $p < 1e^{-23}$, interaction length and frequency: $F(1/356)$ = 0.13, p = 0.72).

In conclusion, the OU process seems to be a suitable model for the task, which makes specific predictions about parameters that influence how to integrate information over different timescales. Behavioural measures such as RT, detection rate and false alarm rate, of the model output differ for different conditions according to the hypotheses stated above. This means that these behavioural measures can be compared with participant behaviour to infer whether participants adapted their evidence integration strategy. In the next section, I will use this model

to develop an approach to analyse the noise and the signal of my paradigm's coherent motion stimulus, and how these vary between the different conditions.

2.2.3 How do we measure the differences in evidence integration for different task conditions reliably?

In order to investigate how behaviour and the encoding of decision variables in the brain changes during the continuous RDM task, I need to develop strategies to test the influence of trial structure (including length and frequency) on evidence integration. In this section, I will show why the noise structure is advantageous in extracting evidence integration kernels directly from the behaviour. In particular, I will use psychophysical reverse correlation, a commonly used technique to investigate how fluctuating sensory evidence is weighted and influences behaviour (Ahumada, 1996; Neri and Heeger, 2002; Neri et al., 1999). This technique has been used to show that both humans and non-human primates can adapt their evidence integration process to the stimulus statistics (Ganupuru et al., 2019; Levi et al., 2018).

In order to extract integration kernels from the behavioural data, it is important to understand that the model and the participants have to guess when the mean from which the coherence levels for the noise are drawn is different to zero. This will be reflected in the noise structure preceding the responses of the participant or model. We can easily test this by averaging the noise structure leading up to a response. Figure 6 depicts a cut-out from a stimulus stream with some button press responses during trials (green) and during intertrial periods (yellow). We can time-lock to a specific response type, such as the false alarms during intertrial periods, and take the preceding 5 seconds of the stimulus and average that across all false alarms for each condition. This can be repeated for correct responses during trial periods (green responses, Figure 6).

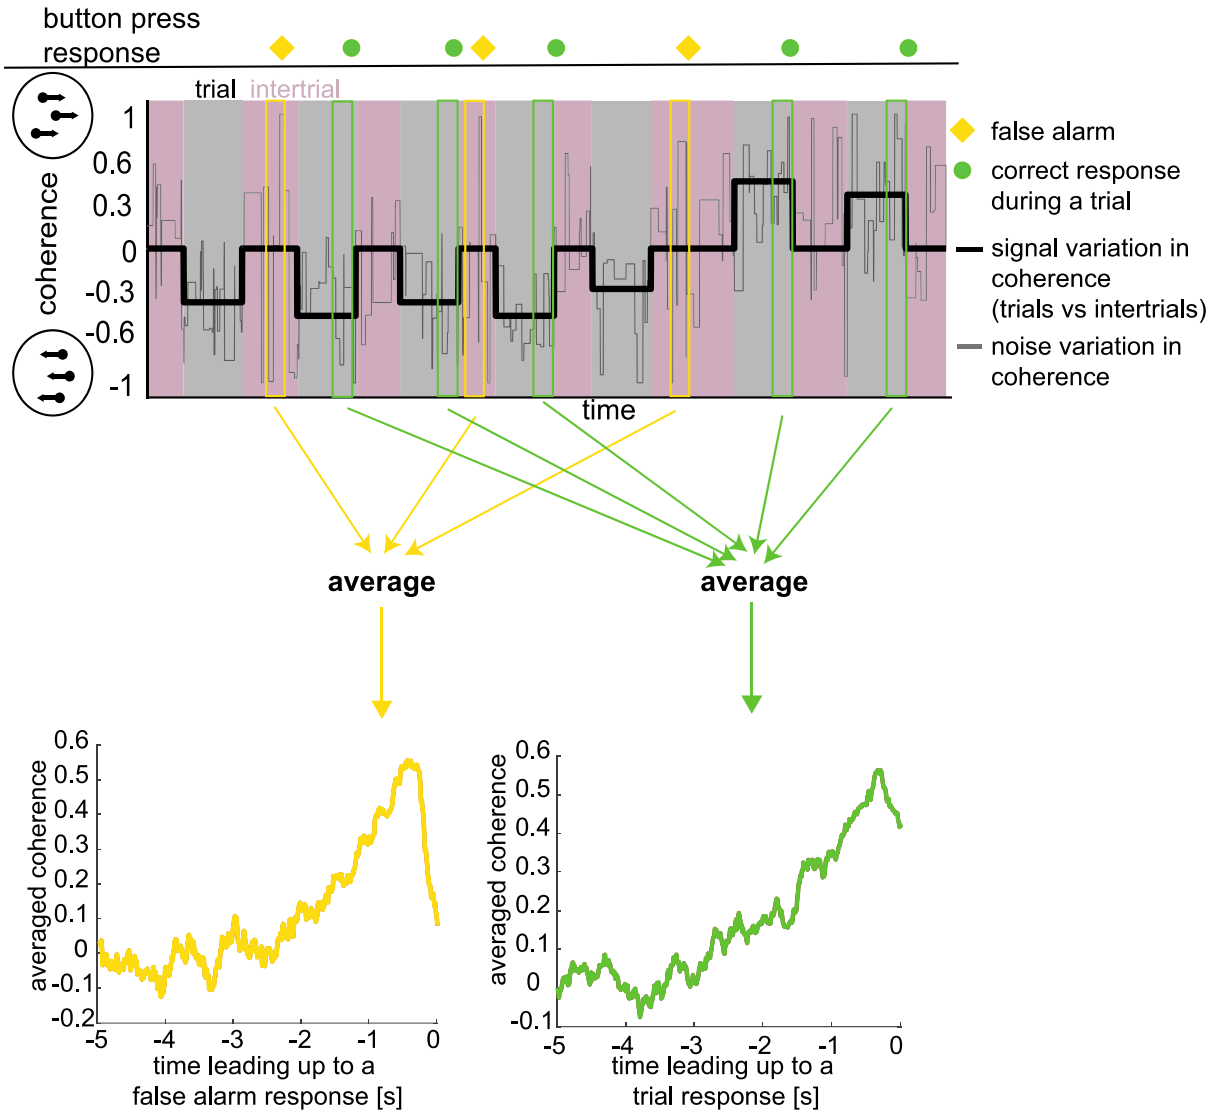


Figure 6: Analysis approach to retrieve behavioural integration kernels. Either participants or models make responses during trial and intertrial periods, as indicated by the button press responses at the top for a cut-out of a stimulus stream. These responses are labelled yellow for false alarms (during intertrial periods) and green for correct responses (during trial periods). For each response, the preceding stimulus, which is the noise variation in coherence, has been extracted and then averaged with all other responses of the same response label. This leads to averaged integration kernels that are shown below. Both kernels show a ramping of average coherences that peak shortly before the button press occurs. If the different conditions (long/short, rare/frequent trials) have an influence on these kernels, then their decay and peaks should vary across conditions.

If participants or the model relied on recent information to make a response, then I would expect to see a rise in the averaged coherences leading up to a response that represents an integration kernel. This provides an empirical estimate of the integration kernel. As stated above, integration kernels for long trials should go further back in time than for short trials and the height of the integration kernel at the time of the response, *z*, should be increased for rare in comparison to frequent trial conditions. In addition, the integration kernels for false alarms and responses to trials should be similar because false alarms result from having mistaken the intertrial baseline noise for signal.

I performed this analysis for the optimal model identified in the last section, and found different evidence integration kernels for false alarms across different frequency conditions. For the analysis I used the false alarm responses and correct trial responses of the of the optimal models. In order to test for significance, I performed a permutation test without replacement to correct for multiple comparisons across time. I used a 2-way ANOVA to search for significant effects of trial length, frequency and their interaction at each timepoint. Then I compared the original F-scores for each timepoint to a distribution of F-scores from 100 simulations with randomly permuted labels (in which I took the maximum F-value across time to enter into the null distribution) and a cut-off of alpha = 0.05.

For false alarms the analysis shows different integration kernels for frequent and rare trial conditions. The average level of coherences required to decide for a response of the model is significantly higher for rare trial conditions as indicated by a cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -4s and 0s, meaning that on average more evidence is needed (Figure 7). In contrast, there is no evidence for a difference between short and long trial conditions, which may be because the differences in best λ between long and short trial conditions are too small to be picked up with this analysis method.

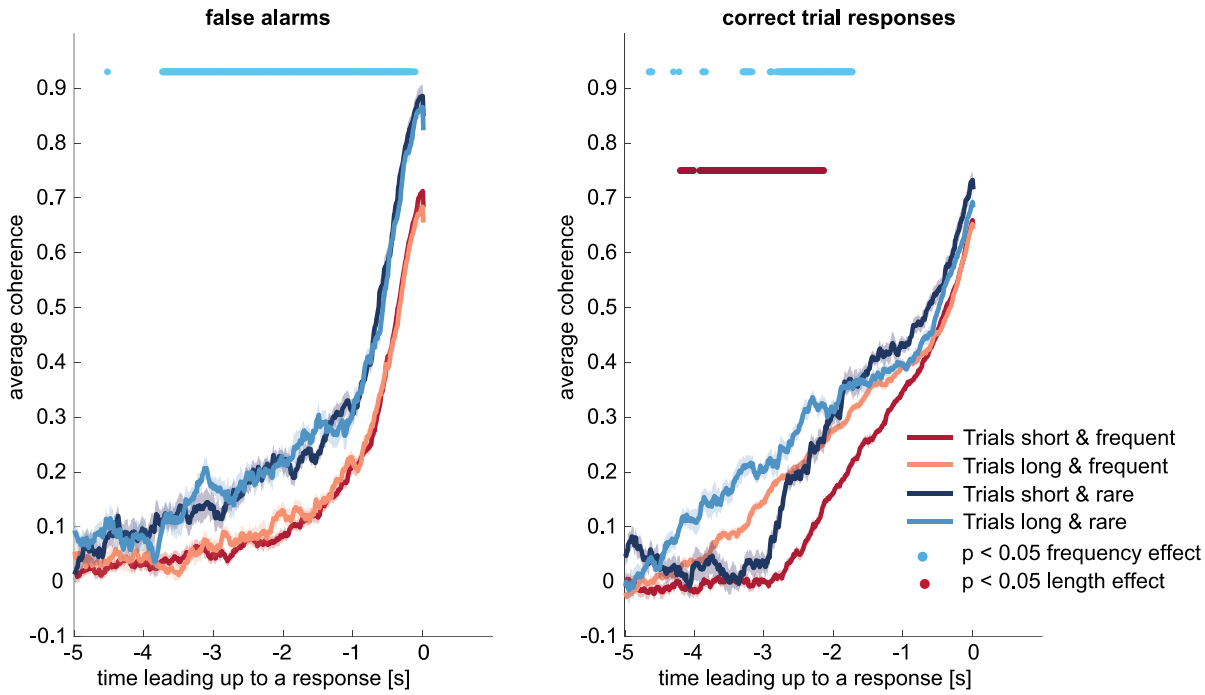


Figure 7: Integration kernels for false alarm and correct responses for the best models. Best estimates of z and λ were chosen to test the model on stimulus streams of different conditions and its responses during intertrial periods (false alarms) and trials recorded. The recorded responses were analysed as explained in Figure 6. The 5 seconds leading up to a response were taken and averaged across all responses and plotted for false alarms on the right and correct responses during trials on the right. To examine whether the effects of trial length and frequency are significant I performed a 2-way ANOVA at each timepoint and used a permutation test without replacement to correct for multiple comparisons. Red and blue dots above the kernels indicate timepoints at which there is a significant effect with an alpha of 0.05.

For correct trial responses the integration kernel was different across all 4 conditions (Figure 7). There is a similar significant effect of trial frequency on the integration kernel with rare trials requiring stronger evidence and a higher average coherence threshold to elicit a model response (cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -3s and -1s). However, the overall threshold of coherences needed to elicit a response is much smaller for trial periods. This supports the theory that false alarms are at least partly triggered by peaks in coherences in the noise structure, whereas responses during trials are primarily driven by a recency-weighted average of past evidence. In contrast to false alarms I do find significant differences in the integration kernels

between long and short trial conditions that are consistent with my predictions (cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -4s and -1.5s).

There are 2 possible reasons for why differences in the response kernel might arise during correct trial responses. It might be, as hoped, that this reflects differences in the parameters of the evidence integration process – or it might be that this is a side effect of the *stimulus statistics* being different for the different trials. This might reflect a change in the underlying mean coherence level rather than a difference in the evidence integration process across conditions. I performed two control analyses to test for possible confounds by the stimulus statistics.

The recency effect I observed during trial periods is at least partially due to decision formation processes. By fixing the parameters of the OU process across conditions (λ = -0.2, c = 0.01 and z = 1.8) I can examine whether the recency effects observed for trial integration kernels reflect evidence integration. Under constant parameters all integration kernels should look the same, if they are not confounded by the stimulus statistics. As shown in Figure 8C the overall difference between *rare* and *frequent* trial conditions has disappeared, indicating that this difference is driven by true differences in the underlying integration kernel rather than by differences in stimulus statistics. This effect should be visible in the integration kernels of participants as well. Nevertheless, a significant difference in the integration kernels between short and long trials remained (as indicated by a cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -4 and -2s). At least to some extant the stimulus statistics are driving this difference, meaning that this is potentially an artefact and does not reflect differences in evidence integration across different conditions.

The second control analysis for trial length implies that the prolonged integration kernels for long trial periods is most likely an artefact of the *mean signal* being different in these long trials. For example, in a response from the long trial condition that is made 4 seconds into the trial, then all of the preceding 4 seconds will have a non-zero mean signal (something that is impossible in the short trial condition,

where the maximum trial length is 3 seconds). I therefore reanalysed the integration kernels only for model responses with a reaction time *smaller* than 3 seconds (so that the maximum trial length was *matched* for the long-trial and short-trial conditions), for the model with the pest parameter estimate and the model with constant parameters (Figure 8B+D). For both models the integration kernels decay now back to baseline within 3 seconds for all conditions. Although notably the optimal model for short/frequent trials still decays back to baseline more quickly than any of the other conditions and a permutation test with a cut-off alpha of 0.05 still shows a small significant effect of trial length between -5s and -2s (Figure 8B). This means that the manipulation of the trial length might not have a measurable impact on the integration kernels. A possible reason could be that the lower variance during the trial periods pushes the evidence faster over the response threshold of the model than if the variance would have been as high as during the intertrial periods. This would mean that the optimal strategy to solve the task is to track changes in the variance and not the length of a trial.

Overall, these results suggest that it is possible to gain some insight on the integration process by making use of the noise structure of the task. Specifically, differences between frequent and rare trial conditions seem to have a measurable effect on the integration kernels, but differences between short and long trials are more challenging to detect. I can apply the same analysis on behavioural data from participants. Moreover, none of these potential confounds apply to the analysis of false alarm response, when the structure of the noise is truly matched across all four conditions. A comparison of the behavioural data from model and participants might therefore allow me to further understand whether participants integrate evidence similar to a leaky accumulator or whether they use a different strategy.

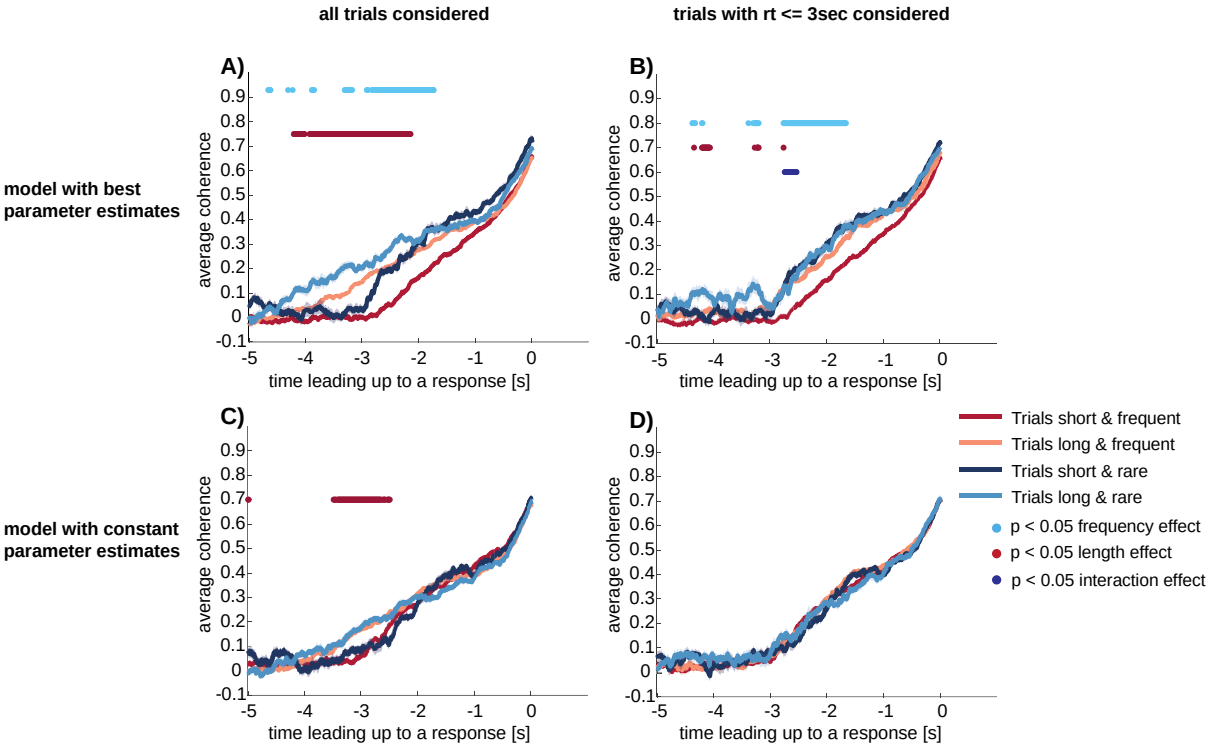


Figure 8: Model comparison for integration kernels for correct trial responses. Responses to trial periods were modelled with either best parameter estimations (**A+B**) or constant parameters (**C+D**). For **A,C** all correct responses were included. For **B,D** only responses with <= 3s reaction times were included for long trial conditions. In order to test for significant effects I performed a 2-way ANOVA at each timepoint and used a permutation test without replacement to correct for multiple comparisons. Dots show significant effects on timepoints at an alpha of 0.05.

2.3 Conclusion

In this chapter I have introduced a novel task paradigm that allows me to measure how flexibly participants can adapt their integration kernels to different task environments. In this paradigm participants watch a stream of moving dots that constantly change direction. Within this stream of randomly moving dots, there are so-called stable trial periods in which the dots move on average more to one side than the other. Participants have to detect the side to which the dots are moving by pressing a button. In addition, exposure to 4 different task environments, in which trial length and frequency have been altered, should enforce an adaption of the task-solving strategy by changing the threshold and decay rate of integration kernels. Short trial lengths should induce integration kernels with shorter decay

than for longer trial lengths. The frequency of trials should have an impact on the threshold of the evidence integration process. For rare trial conditions evidence should be accumulated to a higher threshold than for frequent trials. To investigate whether these manipulations affect the integration process, I investigated how a commonly-used accumulator model, the OU process model, solves the task under different conditions.

I tested whether the parameters of an OU process reflect changes in the evidence integration behaviour of the model. A parameter search revealed that the best model performance was indeed dependent on different parameter values for the threshold and decay of the model. Frequent trial conditions required lower evidence accumulation thresholds (z) and shorter trials required shorter decay rates (λ). Behavioural model results for the different conditions reflect the differences in the best parameter fits. Detection rates are higher for frequent and long trials because participants have longer decay rates and lower thresholds meaning that they are less inhibited to respond. However, low thresholds mean also that the false alarm rate is increased for frequent trials as well as for short trials because the decision is based on less evidence. RTs reflect the different model parameters as well because short decay rates and low thresholds lead to shorter RTs.

It has already been demonstrated that such a model is a good fit to behavioural data in rodent and human studies (Ossmy et al., 2013; Piet et al., 2018). One study (Ossmy et al., 2013) has shown that in an experiment in which participants have to detect luminance signals that vary in length within a noise stream, participants' performance improves with longer signals. They interpreted this result as an indicator that participants adapt their integration kernels according to the trial length and successfully fit a leaky accumulator model to the data (Ossmy et al., 2013). Another study found similar results when rodents had to integrate noisy clicks to form a decision. They have also shown that rodents can adapt their integration kernel based on the signal length (Piet et al., 2018).

However, none of these papers were able to directly infer the integration kernels their subjects used, and had to fit their models to the data. In order to investigate

how subjects adapt their integration kernel we need a task that allows for direct read-out of the decision strategy. The structure of the noise in my task allows for it to be directly empirically estimated using psychophysical reverse correlations, as long as we control for potential differences in the stimulus structure (such as differences in reaction time) across conditions. By averaging over the noisy stimulus leading up to a response, I was able to show that we can retrieve condition dependent integration kernels. This is, therefore, a more direct way in inferring the evidence accumulation strategy participants use.

3 Behavioural performance on a continuous evidence integration task

3.1 Introduction

Decision-making does not happen under static environments and discrete periods, but rather the evidence on which we base our decisions in daily life constantly changes over time. Evidence that was previously relevant for a decision, might be irrelevant for decisions that we are making now. Therefore, we need to be able to adapt our behaviour to the volatility in the evidence in order to make a correct decision. This requires dynamic and flexible processes of decision-making. In Chapter 2, I proposed the OU process as a suitable model to deal with uncertainty in a continuous RDM task. It allowed me to make predictions of how subjects should adapt their behaviour to adapt to the variability in trial frequency and trial lengths in the continuous RDM. In this chapter I investigate whether humans adapt their behaviour in line with those predictions.

In order to study how subjects adapt to changes in evidence we need task paradigms with fluctuating evidence and analyses that allow us to investigate how subjects weigh evidence over time. In most evidence accumulation tasks in the past evidence was typically held constant within a discrete trial (Hanks and Summerfield, 2017). The decision-making behaviour was usually described by a drift-diffusion model in which each evidence sample was equally weighted to form a decision (Gold and Shadlen, 2007). Fewer studies used within-trial variance to investigate how evidence accumulation adapts when evidence fluctuates within a trial (in humans: Cheadle et al., 2014; Ganupuru et al., 2019; Harun et al., 2020; De Lange et al., 2010; Tsetsos et al., 2011, 2012; Wyart et al., 2012, 2015; in monkeys: Huk and Shadlen, 2005; Kiani et al., 2008). For example, a study using fMRI has shown that the brain can modulate the gain associated with each evidence sample. They showed that new evidence, consistent with previous evidences, has smaller influence on the decision variable, if a participant is already relatively certain about the choice (Cheadle et al., 2014).

Depending on the temporal structure of evidence within a trial, subjects discount past evidence (Levi et al., 2018; Tsetsos et al., 2012). The most common technique to study how evidence is weighted in trials with fluctuating evidence are psychophysical reverse correlations (Ganupuru et al., 2019; Harun et al., 2020; Levi et al., 2018; Okazawa et al., 2018; Tsetsos et al., 2012). For this technique evidence leading to a response is averaged across trials. In addition, some of these studies used change-point detection tasks, which show that if the onset of a trial is uncertain or a sudden change in evidence for the opposite choice occurs, discounting past evidence is necessary to detect these state changes (Ganupuru et al., 2019; Harun et al., 2020; Piet et al., 2018; Shinn et al., 2019). This discounting process can be described by leaky accumulator models such as the OU process (Bronfman et al., 2016; Levi et al., 2018; Piet et al., 2018; Shinn et al., 2019; Tsetsos et al., 2011, 2012). Although this provides an important first step in investigating how decision-making might work when evidence changes over time, the discrete trial paradigm does not consider what happens when trials are imbedded in a continuous stream of noisy evidence. In continuous paradigms the onset of the trial needs to be inferred and the length of the trial needs to be estimated in order to solve the task.

Continuous decision-making paradigms in humans have been developed to study evidence accumulation while simultaneously recording EEG (Kelly and O'Connell, 2013; O'Connell et al., 2012). In a continuous decision-making paradigm, subjects observe a stream of noisy evidence, such as randomly moving dots. Subjects have to detect when signal is added to that noise, for example when dots start moving coherently. As such a continuous decision-making task combines elements of a binary decision task (e.g. choose left vs. right) with elements of signal detection (detect signal present vs. absent). For example, a continuous version of the RDM paradigm was among the first to describe signatures of sensory evidence accumulation in the human brain (Kelly and O'Connell, 2013) using a paradigm in which the main source of errors was the failure to detect coherent motion, rather than making an incorrect choice. However, this and similar studies have only analysed the data in correctly identified signal periods similar to studies that used discrete trials with constant evidence (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2016).

Only a few decision-making studies have manipulated the *noise* in a continuous environment (Orsolic et al., 2019; Ossmy et al., 2013). In these studies, signal periods of different length are not confined to discrete trials, and the onset of decision-relevant information can only be inferred by discounting past evidence. Using this approach, humans and mice have shown to be able to discount past evidence in order to detect periods of signal in a noisy environment. A leaky accumulation model can explain the behavioural results, in which the best strategy to solve the task is to discount past evidence depending on the length of the period with signal (Ossmy et al., 2013).

In my experiment, I tested variations in both the length and frequency of the signal period in a continuous RDM task, in which the noise was also continuously varying. As laid out in Chapter 2 (Section 2.1), I used a 2x2 design in which participants either had to detect frequent or rare and long or short trial periods. In this chapter I examined whether human behaviour is consistent with a leaky accumulation model, the OU process (see Chapter 2.2). Simulations of the task with an OU process have shown that in order to optimise the reward, the decay rate (λ) and the decision threshold (z) were adapted to the different conditions (Chapter 2.2.2). In particular, for long trial conditions the decay rate was less negative in contrast to short trials, meaning that past evidence was discounted more slowly. Similarly, the optimal threshold was increased for rare trial conditions, meaning the OU process should require more accumulated evidence to make a response (in order to avoid increased false alarms). The behavioural results of the model presented in Chapter 2 including the detection rate, reaction time (RT) and false alarm distribution (see Chapter 2, Section 2.2.2) as well as behavioural integration kernels measured with psychophysical reverse correlation (see Chapter 2, Section 2.2.3), should allow me to draw comparisons with subject behaviour.

I found the behaviour of the OU process and subject behaviour were quite similar across conditions. Except for false alarm rate, participants adapted their behaviour between different conditions in line with the OU process predictions. Psychophysical reverse correlations showed that participants discounted past evidence. Participants adapted their integration kernels to frequent and rare trial conditions. Integration kernels for both correct responses during trial periods and

for false alarms were consistent with the integration kernels from the OU process. In addition, I found that integration kernels were subject specific, meaning that they were more similar across conditions within a subject than between subjects.

3.2 Task design and testing protocol

In total 32 participants were tested, 12 males. All were below the age of 40 and had normal or corrected-to-normal vision. Four participants were excluded from the analysis. Three subjects were unable to learn the task either because they were not able to identify trial periods in the continuous task or they were at chance level in the first training task (see below). Another participant fell repeatedly asleep during the testing session and was also excluded. In total 28 participants were included for the behavioural data analysis. The study was approved by the local ethics committee (CUREC approval: R60298) and all participants gave written consent prior to taking part in the study.

The task is described in detail in Chapter 2. Every participant completed 2 sessions, a 'training' session and a 'test' session. Both sessions took place in a darkened and sound attenuated booth. In the first session, participants underwent 2h of training in which they gradually learnt the structure of the continuous task. On a separate day from training, participants performed the continuous task while EEG was recorded. The two sessions were never more than a week apart.

Our training protocol was designed to overtrain participants to reach a high level of performance on the task, so that during the main testing session they were not having to learn online how to perform optimally on the task. Our training taught participants about the structure of the long/short and rare/frequent trial periods, and how to discriminate such trial periods from background noise fluctuations. Crucially, it also incentivised participants to perform the task so as to maximise their overall reward rate. This gives validity to our modelling approach in Chapter 2, where we sought to find the parameters of the OU process that would optimise reward rate on the task (rather than optimising other measures of task performance, such as maximising detections or minimising false alarms).

3.2.1 Design of the random dot motion stimulus

I coded the task and stimulus in Psychtoolbox (Brainard, 1997) and parameters were chosen similar to Shadlen and Newsome (2001). Participants were seated 87 cm in front of the screen. The moving dots were displayed in a circular aperture subtending with a radius of 5 visual degrees on a Dell monitor with a refresh rate of 100Hz. Dots had a size of 0.1 visual degrees and were displayed with a density of 2.5 dots per squared visual degree. A fixation point in the centre of the screen had a size of 0.3 visual degrees. All dots were black and displayed on a mid-grey background (rgb: 0.5, 0.5, 0.5).

Originally, random dot motion stimuli were used to study motion perception in middle temporal area (MT) which is highly responsive to global motion signals (Britten et al., 1993; Movshon and Newsome, 1992; Newsome and Pare, 1988; Salzman et al., 1992). This puts an important constraint on the design of the random dot motion task. The trajectory of a single dot should not reveal the motion direction of the entire stimulus. Instead, the task needs to be designed in a way that the entire field of moving dots is necessary to detect the motion direction (Movshon and Newsome, 1992). This feature enables the study of decision-making because the sensitivity of MT neurons to these global motion signals defines the sensory evidence in the evidence accumulation process later studied in lateral intraparietal areal (LIP).

For this study I incorporated coding principles by Shadlen and Newsome (2001) to compute global motion signals. Dots were equally divided in 3 sets. These sets were shown sequentially, meaning only one set per frame was shown. Each time a set reappeared on the screen the coherence on that frame dictated the likelihood of that dot either being displaced randomly or in the direction of the coherence. Randomly displayed dots moved like Brownian motion particles with no particular speed. Dots that moved coherently were displaced according to a speed of 7 degrees per second. This approach means that subjects are forced to observe the entire field of moving dots to figure out the motion direction. A single dot is not reliable because it only reappears on every third frame and does not necessarily move coherently for more than 2 frames.

3.2.2 Training protocol

Task 1

During training, participants were not immediately confronted with the continuous motion task. Participants were introduced to the conventional RDM task based on discrete trials (Shadlen and Newsome, 2001). They completed 70 trials of increasing difficulty, in which motion direction was either leftward or rightward and coherences progressively decreased from 90% to 10% in blocks of ten trials. In all discrete trial blocks (tasks 1 to 5), participants were given feedback by changing the colour of the fixation cross to indicate correct/incorrect responses and by either a growing or shrinking reward bar.

Task 2

Participants completed 70 trials as in task 1, but coherences now varied randomly from trial to trial, at levels [10, 20, 30, 40, 50%].

Task 3

Participants completed 28 trials of an intermediate task version that still consisted of discrete trials, but had features of the continuous task. In each trial some noise was superimposed on the mean motion coherence. The noise had a step-like structure inducing fast changes in the percept of motion direction (see Chapter 2.1, Figure 1). The coherence steps were sampled from a normal distribution with a mean coherence and a standard deviation of 0.3 as in the continuous RDM (see Chapter 2.1). This meant that as in the final continuous task version coherences would fluctuate throughout the trial; participants had to estimate the average motion direction across the entire trial but respond before the trial had completed. Trial length was 5s (trial length of long trial conditions) and the fixation point shrank over the course of the trial, so that participants would have an idea how much time they had left to respond. Mean coherences were fixed at 80, 70, 60 and 0%. Trials with a mean coherence of 0% were included to simulate intertrial periods of the continuous task that participants had to contrast to the other trials. Participants had to suppress a response in these trials.

Task 4

This task was the same as task 3, but trials were increased to 56 and mean coherences reduced to 50, 40, 30 and 0%.

Task 5

This task was the same as task 4. However, the trial length was reduced to 3s to simulate short trials of the continuous task. This served to train participants about the need to make judgments based on longer/shorter durations of evidence based on the length of time needed to respond.

Task 6

In this task participants encountered the continuous version for the first time, in order to introduce them to the idea of rare and frequent trial periods. Participants performed 4 blocks of 3min. The first two blocks were long and frequent trial conditions, the last two long and rare trial conditions. An additional 500ms was allowed following each trial period for subjects to respond to account for non-decision time. In contrast to the final version, mean trial coherences were increased to 60, 50 and 40%, in order to make trials more straightforward to detect in the superimposed noise of step-like changes in the coherence (see Chapter 2.1, Figure 1). Crucially, the fixation point turned white to indicate the onset of each trial period. This manipulation, although not present in the final continuous task, was designed to help participants to learn the frequency of trials in different conditions. In addition to this, both here and in the main experiment, participants were explicitly cued on an instruction screen whether in the current block the trial periods would be ‘long’ or ‘short’, and whether they would be ‘rare’ or ‘frequent’.

To show participants how well they were performing on the continuous task, a ‘reward bar’ was shown below the stimulus throughout the block. The ‘reward bar’ grew in size every time the subjects responded correctly (+3 points), and reduced in size every time the subjects made a false alarm (-1.5 points), missed a trial (-1.5 points) or responded incorrectly (-3 points). Once the reward bar reached a ‘target’ (= 15 points), a bonus was added to the participant payment and the reward bar reset. This was designed to incentivise participants to maximise their overall reward rate in the task, rather than to simply maximise correct detections or

minimise false alarms. It was in addition to the central fixation point changing colour to indicate whether responses were correct/incorrect. The reward bar was left onscreen throughout the continuous training blocks (tasks 6 to 11), but was not presented onscreen during task performance of the main testing blocks (day 2) in order to eliminate eye movements to the reward bar and maintain participants' fixation on the RDK stimulus.

Task 7

This task was the same as task 6 but mean trial coherences were decreased to 50, 40 and 30% - these are the mean coherences of the trials in the final version of the task. In addition, the fixation point did not change to white during trials anymore, so participants had to infer when trial periods were present or absent.

Task 8

This task was the same as task 6, but the first two blocks were of the short and frequent trial condition and the last two blocks of the short and rare condition. The fixation point changed to white during trial periods, allowing participants to learn about trial frequency in a condition with shorter trial length.

Task 9

This task was the same as task 8, but with mean trial coherences 50, 40 and 30% were included and the fixation point did not turn white during trials anymore.

Task 10

Participants practised the full continuous version. In each task participants completed 4 blocks of 5min length, one for each condition in randomised order, with the participant instructed as to which condition they were currently performing. As a further reminder of the current task condition, the fixation point took different shapes and sizes for each of the different conditions. The fixation point did not turn white to indicate the onset of trial periods. Mean coherences for trials were fixed at 50, 40 and 30%.

For the discrete trial tasks performance was checked qualitatively by looking at the psychometric functions and participants had to perform at 80% correct or higher to

move on to the continuous tasks. Performance in the continuous task was only checked qualitatively, by plotting the stimulus stream and responses after each block. That way I could get an impression of the frequency with which participants responded within trials, how often they missed a trial and their false alarm rate. I gave verbal feedback to participants based on their performance here, which also helped participants to improve their behaviour during training. If, in the latter parts of training, participants missed more than half of the trials, they were excluded from the subsequent EEG session (n=3). For those participants who failed to progress from the training session, they were only paid for that session and did not move on to the EEG session. During training, participants earned £10/hour baseline payment, plus £0.15 bonus payment each time the reward bar reached its target.

3.2.3 Testing protocol

After participants completed training successfully, they participated in a second session not more than one week later. In the second session participants first performed a 'reminder' where they repeated training task 10 from the training session; this was performed while the experimenter put on the EEG cap. Then, while EEG data was collected, they completed 5-6 task 'sessions', each lasting 20 minutes. Each session consisted of all 4 conditions, presented for 5 minutes each, in randomised order. The reward bar was only shown after a block had ended; participants now earned £0.50 bonus payment each time the reward bar was filled. At the start of each block, participants were informed which task condition they were about to perform. Further, after each response or missed response the fixation point would change colour to give participants feedback about their performance; it turned green to indicate a correct response, red to indicate an incorrect response, yellow to indicate a false alarm and blue to indicate a missed trial. The present chapter focusses only on behavioural analyses; EEG results will be presented in Chapter 4.

3.3 Methods and results

All data analysis was carried out in MATLAB (Version R2018b, The MathWorks, Inc., Natick, Massachusetts, United States).

3.3.1 Detection rate

First, I investigated whether the detection rate of stimuli, and the false alarm rate, depended on the task conditions (trial length and frequency). I predicted an increase in detection rate for longer trials, as there is more coherent motion to detect, and an increase in false alarm rate for more frequent trials (as the prior probability of a trial is higher).

Methods. First, I calculated detection rates as the proportion of correctly identified trial periods, for each mean coherence level over all trials in a block. Detection rates for blocks were averaged across different sessions in each subject, and are plotted in Figure 1 (mean +/- standard error of the mean (s.e.m.) across subjects). In order to determine whether trial duration, frequency or coherence level effect the detection rate, I chose the following linear model for subject level analysis:

$$DT = \beta_1 \cdot Coh + \beta_2 \cdot TF + \beta_3 \cdot TL$$

where *DT* is the detection rate, *Coh* is a regressor for the mean trial coherence level, *TF* for trial frequency (+1 for frequent, -1 for rare) and *TL* for trial length (+1 for short, -1 for long). Except for the *Coh* regressor all others were binary. The *Coh* regressor was demeaned. Then I performed a group-level analysis for random effects.

Results. As expected, detections rates varied not only with trial coherence but also with trial frequency and length (Figure 1). As shown in previous studies the detection rate increased with the mean coherence level (random effects analysis: $t(27) = 16.08$, $p < 1e^{-13}$) and for long trial conditions (random effects analysis: $t(27) = -15.55$, $p < 1e^{-14}$, Kelly and O'Connell, 2013; Ossmy et al., 2013). In addition, the detection rate was higher for frequent trial conditions than for rare (random effects analysis: $t(27) = 7.26$, p-value $< 1e^{-7}$). This suggests that even for trials consisting of the same amount of sensory evidence, subjects were more cautious in reporting trial periods in rare conditions (because the baseline probability of being in a trial was lower). These findings are qualitatively similar to simulations with the OU process shown in Figure 6 of Chapter 2, supporting the theory that participants

adapted their evidence integration parameters according to the task demands. Note that this conclusion does not depend upon a parameterised fit of the model to the data, but instead upon a parameterisation of the model that maximises reward rate in the task (see Chapter 2).

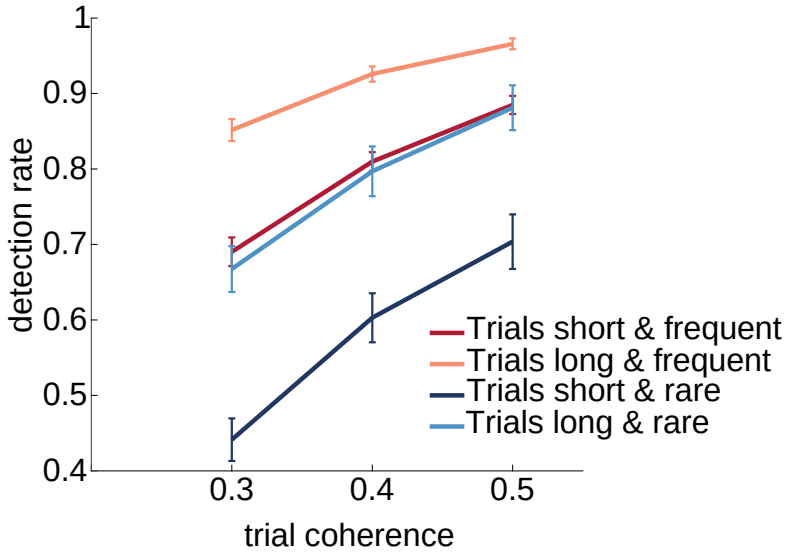


Figure 1: Performance under different conditions. Shown are the average detection rate of trial periods for different mean coherences (collapsed over left and right trials) and for different conditions (see legend). The detection rate was calculated as ratio of correctly identified trial periods of a mean coherence over all trial occurrences of that mean coherence. There are significant differences between trial frequency ($p < 1e^{-7}$), trial length ($p < 1e^{-14}$) and mean coherence levels ($p < 1e^{-13}$) tested with a random effects analysis.

3.3.2 False alarm rates

Methods. In order to calculate the false alarm rate, I counted for each subject and condition the number of responses during intertrial periods across all sessions. Then I calculated the total time spent in intertrial periods in seconds for each condition. This allowed me to calculate the number of false alarms per second of intertrial period. These were averaged across subjects for each condition (Figure 2, shown with individual data points). I implemented a 2-way ANOVA to test for significant differences between the false alarm rates of different conditions. One subject was removed because the false alarm rate was further away than 2 standard deviations from the mean for one condition.

Results. False alarms varied between conditions. Trial frequency did not have a significant influence on false alarm rates ($F(1/104)$ = 0.11, p = 0.73 from 2-way ANOVA), whereas trial length did, with shorter trial conditions leading to much higher false alarm rates ($F(1/104)$ = 12.2, p = 0.007 from 2-way ANOVA). The interaction between length and frequency was also not significant ($F(1/104)$ = 0.3, p = 0.59 from 2-way ANOVA). This suggests that despite rare trials having a lower detection rate than frequent trials, this did not translate into a lower false alarm rate (Figure 2). By contrast, in short trial conditions, participants were prepared to false alarm slightly more frequently (Figure 2), perhaps motivated by the high proportion of missed trials in these conditions (Figure 1).

The overall false alarm rate appeared considerably lower than that of the optimised OU process (compare to Chapter 2, Figure 6). This might indicate that participants used a different strategy to leaky integration, or it might imply that other parameters of the OU process (such as noise) should be further adjusted in order to provide a more appropriate fit to participants' behaviour. However, we would expect that most predictions from the OU process would match, suggesting participants should increase false alarm rates for frequent trial conditions, as they should be less cautious; this is clearly inconsistent with the data (Figure 2).

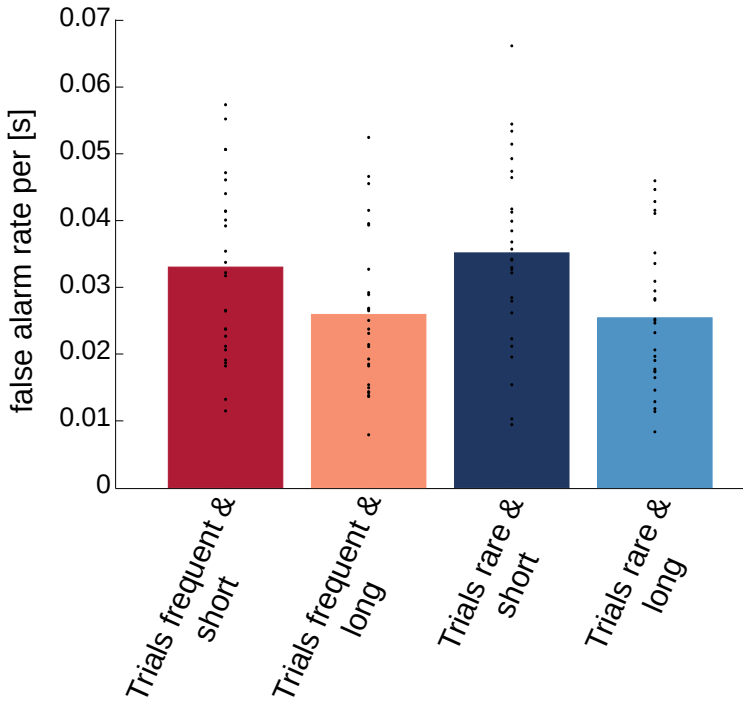


Figure 2: False alarm rate across conditions. For each condition the mean of false alarms rate was calculated across subjects by dividing the number of false alarm responses by the seconds of intertrial time in the stimuli of different conditions. Each black point represents the mean false alarm rate of a subject. There is a significant difference between short and long trial conditions (p = 0.007, 2-way ANOVA).

3.3.3 Reaction times

I hypothesised that RTs would be longer for long-trial conditions, as more information and confidence can be gained by observing the stimulus for longer. Furthermore, predictions from the OU process suggest that RTs should be longer in rare trial conditions – as the prior probability of a trial is lower and more evidence is needed to conclude that a trial is indeed occurring.

Methods. In order to examine the influence of different task conditions, I calculated the median RT per trial coherence and condition across subjects (Figure 2). In order to ensure a fair comparison between short and long trial conditions, I repeated this analysis only including correct trial responses with RTs smaller or equal to 3.5s (the maximum response window for short trials, as an additional

500ms was allowed following each trial period for subjects to respond to account for non-decision time).

It could be argued that the increased RT in the long trial conditions might be an artefact because the response window in short trials was truncated at 3.5s in contrast to long trial conditions in which participants had 5.5s to respond. In other words, participants could have used the same integration kernel in long- and short-trial conditions, but been prevented from responding in a subset of short trials (when the RT would have been over 3.5s), resulting in a difference in mean RT that does *not* reflect a difference in the integration kernel. Therefore, I corrected for the possibility that RT distributions might be longer for long trials simply because there is a longer response window available for these trials.

I performed a random effects analysis to investigate the influence of different trial conditions and coherences on RTs. For each subject, I performed a regression analysis with the same regressors previously used for the detection rate analysis, plus an additional interaction term between trial frequency and trial length. The coherence regressor was demeaned.

$$RT = \beta_1 \cdot Coh + \beta_2 \cdot TF + \beta_3 \cdot TL + \beta_4 \cdot TF \cdot TL$$

I obtained betas of each regressor for each subject. Subsequently, I performed a *t*-test on all regression coefficients for each regressor separately.

Results. Across all coherences, the different trial conditions, trial frequency and trial length, influenced the median RT distributions (Figure 2, upper panel); this remained true even when controlling for response window length by restricting analyses to consider responses made with RTs < 3.5s (Figure 2, lower panel). Longer RTs were observed for long trial conditions in comparison to short trial conditions (Figure 2, lower panel, random effect analysis: $t(27) = 11.25$, $p < 1e^{-10}$). In long trial conditions participants should adapt their integration kernel to be as long as the trial length. Similarly, rare trial conditions led to an increase in RTs, as well (Figure 2, lower panel, random effect analysis: $t(27) = 34.67$ $p < 1e^{-23}$). In rare trial conditions, participants should increase their decision threshold that leads to

longer RTs. This is exactly what I found for the decision threshold parameter (z) and decay parameter (λ) of the OU process. Further, there was an interaction between trial frequency and trial length that led to prolonged RTs for rare and long trial conditions (Figure 2, lower panel, random effects analysis, $t(27) = -7.51$, $p < 1e^{-7}$). As previously shown, participants' RTs decreased with increasing coherence strength (Roitman and Shadlen, 2002, Figure 2, lower panel, random effect analysis: $t(27) = -9.15$, $p < 1e^{-9}$).

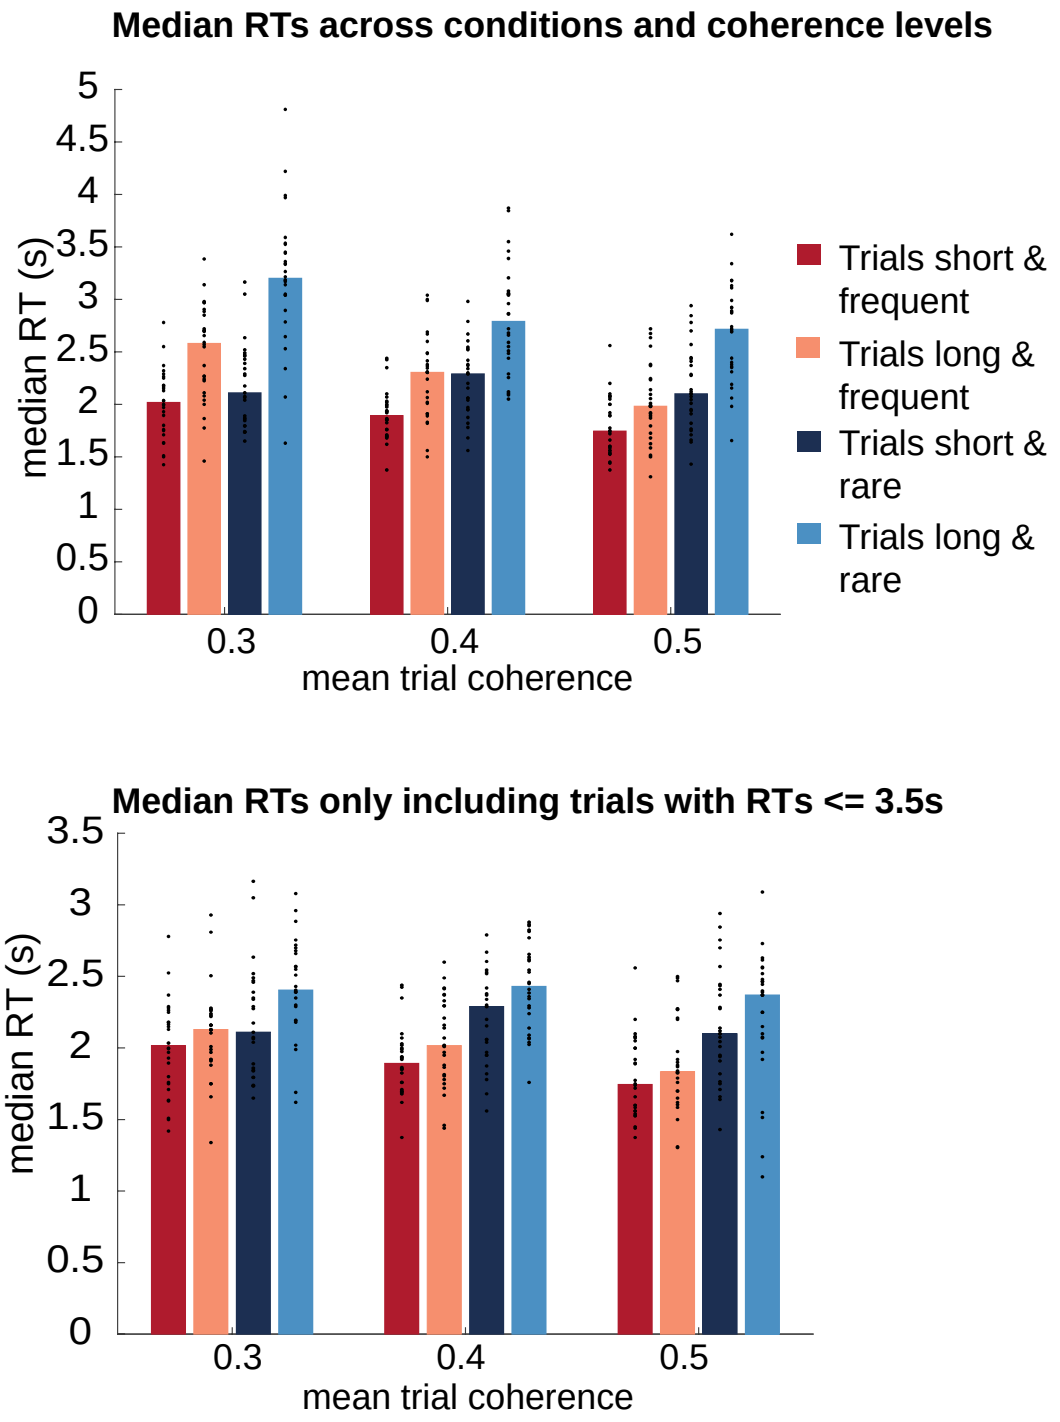


Figure 3: Median RTs across conditions and coherence levels. The upper panel shows median RTs of correct trial responses for different mean coherences and conditions. Bars show the median across subjects. Black points are subject medians. The lower panel depicts the same but only for correct trial responses with RTs that do not exceed the time to respond to a trial period in short trial conditions of 3.5s. A random effect analysis showed that there was a significant influence of trial mean coherences and conditions on RTs.

However, none of these measurements provide conclusive evidence for leaky integration. Other models might also be able to explain the changes in false alarm rate, detection rate and RTs. For example, some studies suggest that a normative Bayesian model with a non-linear transition function that determines how prior belief is discounted might be a better fit in contrast to the linear discount in an OU process if the brain interprets this task as a change-point detection task (Glaze et al., 2015; Murphy et al., 2020; Veliz-Cuba et al., 2016, see Discussion). For this reason, I am using psychophysical reverse correlations in the next section to better understand the evidence integration strategies participants have applied.

3.3.4 Integration kernels

So far, I have only several key comparisons of behaviour to predictions from the OU process modelling (Chapter 2) to investigate how participants adapt their task performance under different conditions. However, an empirical read-out of the integration kernel provides a more direct window into the integration process participants used, and the degree to which more recent evidence was weighed against past evidence (leakiness). As discussed in Chapter 2, the optimal strategy to solve this task is to discount past evidence dependent on the task condition. In general, participants should adapt their integration kernel to the length of a trial and adapt their decision threshold depending on trial frequency. For rare trial conditions, participants should adapt a higher decision threshold to avoid false alarms. If they would not discount evidence and adapt their decision threshold, participants would base their decision on the entire past evidence. Therefore, they would not be able to 'forget' evidence of past intertrial periods. In my task, the noise structure allows me to get a direct measure of the integration kernels participants used to detect trial periods, by using psychophysical reverse correlations (see Chapter 2, Figure 7).

Methods. To get a direct measurement of the integration kernels, I used the psychophysical reverse correlation method for determining spatiotemporal weighting functions of sensory input (Ahumada, 1996; Neri and Heeger, 2002; Neri et al., 1999; Tsetsos et al., 2012). The logic of this approach is that by looking at the time points in the run up to the response, I can determine which timepoints

contributed to that response, as they will on average contain motion compatible with the response (see Chapter 2 for details of the method).

For each subject and condition, I averaged 5s preceding each response, having first separated trials into false alarms or correct trial responses. I then averaged across subjects and plotted the results with s.e.m. (Figure 4). To test for significant differences between the integration kernels of different conditions and to account for multiple comparisons, I implemented a permutation test without replacement. For this test I shuffled the condition labels for the averaged data per subject and calculated the *F*-value for each time point using a 2-way ANOVA across subjects. I repeated this 100 times. To correct for multiple comparisons across time, I selected the highest *F*-value of the distribution across all possible timepoints to build a null-distribution across the 100 permutations. Then I compared the true *F*-value of a 2-way ANOVA analysis using the correct labels, for each data point with the respective null-distribution and selected a cut-off of an alpha of 0.05. I repeated this analysis for correct trial responses, but again excluded responses with RTs exceeding 3.5s (to control for potential influences of the response window on the integration kernel).

In addition, I simulated random responses to 150 stimulus streams per condition to examine the 'auto-correlation' within the stimulus streams and how much this accounts for the results I find in the participants. At each timepoint at which the stimulus exceeded a coherence of 0.5 in the stimulus there was a 20% chance of a response irrespective of trial or intertrial period. I then averaged across the 5s leading up to the simulated responses for each condition.

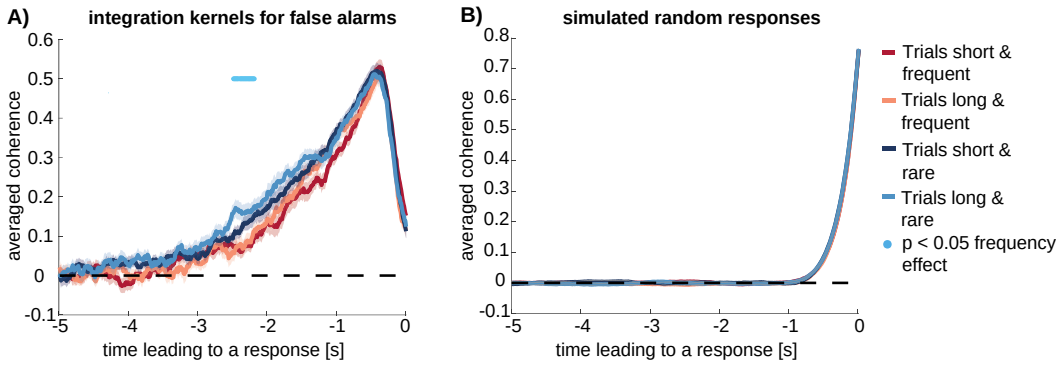


Figure 4: Integration kernels for false alarms. A) Integration kernels for false alarms. Stimulus streams leading to a false alarm were averaged across subjects for each condition (mean + s.e.m. – shaded area), blue dots indicate significant influence of trial frequency (see legend). **B)** Integration kernel for simulated random responses. Responses were simulated with a random chance of 0.2 at any timepoint for a response when the coherence was > 0.5. Then the stimulus streams leading to a response were averaged for each condition.

In order to quantify the fitted integration kernels to false alarms with a single number per condition, I fitted an exponential decay to the averaged integration kernels per condition for each subject. For each integration kernel, I selected the data from the peak to 5s before a response and used an in-built gradient descent algorithm in MATLAB to minimise the least-squares fit between the exponential Equation and the integration kernel. The exponential function is defined as follows:

$$Model = A * e^{\frac{t}{\tau}}$$

Where *A* is the amplitude, *t* = time and τ describes the decay. The loss function was defined as the squared error between model and real data plus an added *L2*-regularisation:

$$loss = \sum (data - model)^2 + \sum params^2 * \lambda$$

Where *params* are the estimated parameters for *A* and τ, while λ is the regularization factor for L2-regularisation, set to 0.01. I only included subjects with more than 20 responses in any condition, to exclude subjects for which a good fit was not possible due to an unreliable reverse correlation analysis. This method

was adapted from Murray et al. (2014). Subjects with parameter fits > 2 standard deviations from the mean were excluded as well.

Results. The integration kernels leading up to false alarms are the most readily interpretable, as the distribution of motion coherence outside trial periods was matched across conditions – therefore any difference in integration kernel leading up to false alarms can only reflect adaptation to the task requirements, rather than any possible differences in the stimulus.

Participants adapted their integration kernels to different trial conditions. There is a significant difference in the decay of the integration kernels between rare and frequent trial conditions as indicated by a cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -3 and -2s (Figure 4). For rare trial conditions participants needed more evidence to make a decision than for frequent trial conditions. In contrast, an effect of trial length on integration kernels was not observed. It did not appear that the peak of the evidence integration kernel differed between frequent and rare trials.

OU process simulations led to qualitatively similar false alarm integration kernels (Chapter 2, Figure 7). Both participants and the OU process showed slower decay rates for rare trial conditions. For both the OU process simulations and actual participant behaviour, there was no measurable influence of trial length on the integration kernels (Figure 4). One possibility is that the difference in trial lengths (3s vs. 5s) was not sufficiently great to cause participants to adapt their integration kernel, but instead they only adapted their rate of responding across these conditions (see Figures 1 and 2). However, in contrast to participants the integration kernels measured for the OU process showed differences in the peak of evidence integration kernels that were higher for rare trial conditions. This suggests that participants may not adapt their decision threshold to these different task conditions, but instead the decay rate that determines how they weigh past evidence.

One potential confound in this analysis is that, due to the structure of the noise in the task (see Chapter 2, Figure 1), the stimulus is itself autocorrelated. This might,

for instance, mean that a decision strategy that only considered a single piece of evidence (Stine et al., 2020) could nevertheless appear to have some degree of ‘leaky integration’. To exclude the possibility that these results are due to the stimulus stream and not due to cognitive processes, I therefore simulated responses at random timepoints, triggered whenever the instantaneous stimulus exceeded a certain threshold (0.5). I then calculated the resulting integration kernels (Figure 4B). For these simulations there was no difference between conditions, and the integration kernels showed a far shorter decay than those observed in behaviour. This means that the integration kernels retrieved from the behavioural responses and stimuli reflect cognitive processes. In addition, these behavioural results are similar to the simulation results of the OU process (Chapter 2), further supporting that participants discounted past evidence differently for frequent and rare trial conditions.

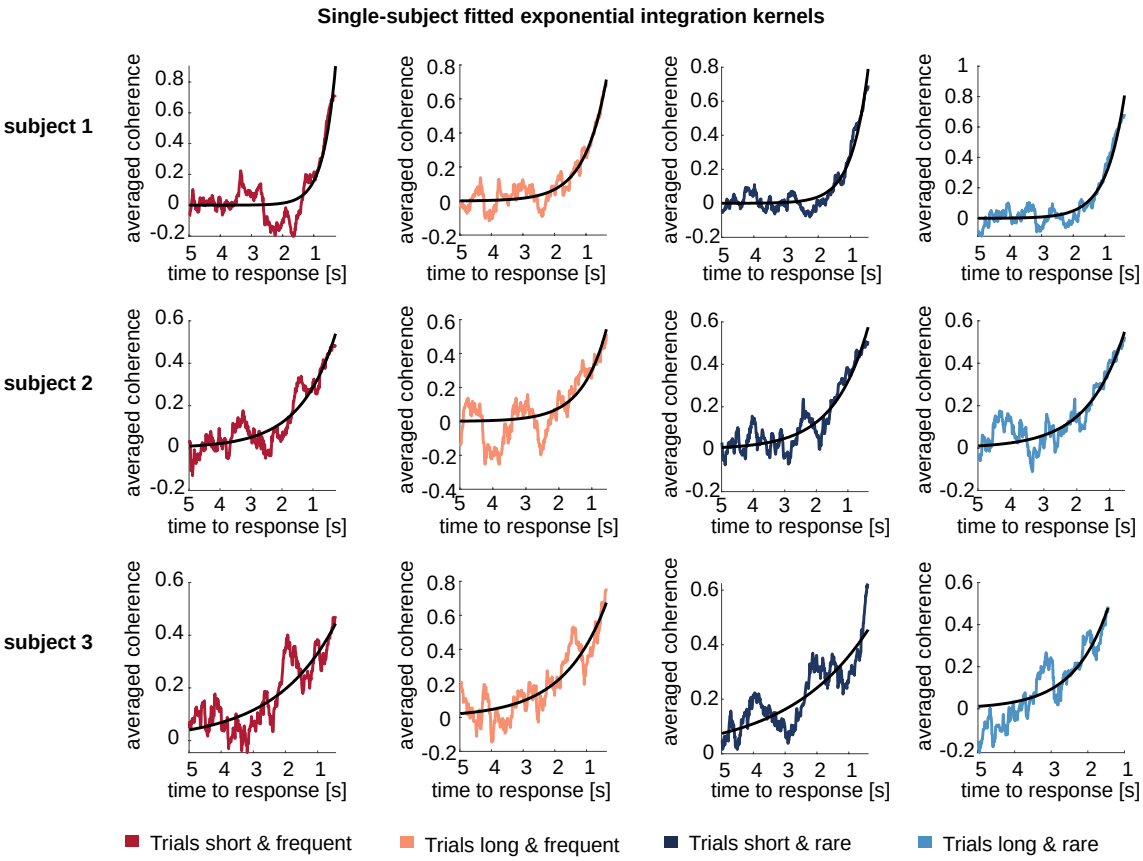


Figure 5: Single-subject kernels for false alarms with fitted exponentials. For each subject I fitted exponentials to averaged integration kernels for false alarms across sessions for each condition. Exemplary integration kernels and fitted exponentials are shown for 3 subjects.

Looking at the single subject fits, there was one further result that was unexpected and notable: integration kernels appeared to be highly subject-specific, with some subjects having fast, slow or intermediate decay rates (Figure 5). In order to quantify the similarity between τ and *A*s across conditions within subjects, I computed the Pearson's correlation (R) coefficient between single subject τs for long versus short and frequent versus rare trial conditions and repeated this for *A*s (Figure 6). I found that both parameters are significantly correlated across conditions within subjects (long vs short: (τ) R = 0.86, $p < 1e^{-5}$, 95% CIs: [0.65 0.95]; (*A*) R =0.76, $p < 1e^{-4}$, 95% CIs: [0.48 0.9]; rare vs frequent: (τ) R = 0.83, $p < 1e^{-4}$, 95% CIs: [0.58 0.94], (*A*) R = 0.81, $p < 1e^{-4}$, 95% CIs: [0.57 0.92]), which further confirms that the integration kernels within subjects are very similar.

The similarity of integration kernels within subjects is important because it might allow me to probe the cross-subject variation in the parameters of the exponential fit to other behavioural measures such as RT. In addition, these subject specific integration kernels could suggest that the underlying brain signals vary between subjects as well, and although not explored in this thesis, one possible avenue would be to explore the relationship in across-subject variation in these parameters with the EEG signals studied in Chapter 4. The high single-subject reliability of these markers of evidence integration also suggests that these could be studied in terms of their variation across the population, or how they are changed as a consequence of neurological or psychiatric disease.

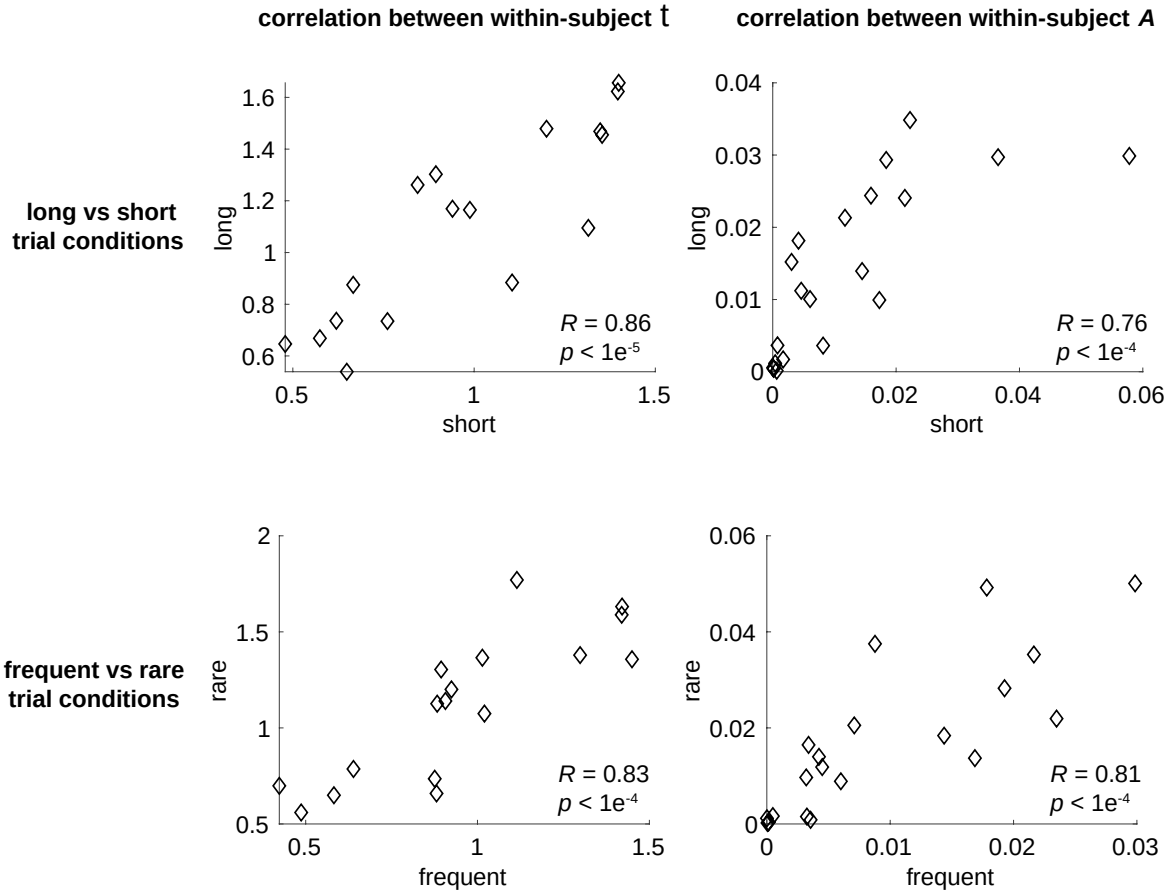


Figure 6: Within subject correlation of fitted τ and *A*s to integration kernels of false alarms. For each subject I fitted exponentials to averaged integration kernels for false alarms across sessions for long, short, frequent and rare trial conditions. Then I correlated τ and *A*s for long versus short and rare versus frequent trial conditions. For each correlation I reported the Pearson's correlation coefficient (*R*) and significance (*p*).

3.3.5 Correlation between integration kernels and reaction times

If the fitted integration kernels do capture individual differences in evidence integration (see the striking within-subject consistency in kernels, Figure 5-6), the parameters of the fitted exponentials may correlate with more elementary measures of behaviour, such as variation in RTs (<= 3.5s) across subjects.

Correlating the parameter estimates of the integration kernels for false alarms with median RTs showed a positive relationship specifically between fitted τ and median RTs within subjects and for each condition (Figure 7, *short and frequent*: (*A*) *R* = 0.46, *p* = 0.07, 95% CIs: [-0.05, 0.77]; (τ) *R* = 0.51, *p* = 0.04, 95% CIs: [0.04, 0.8]; *long and frequent:* (*A*) R = 0.63, *p* = 0.01, 95% CIs: [0.2, 0.86]; (τ) *R* = 0.53, *p* = 0.03, 95% CIs: [0.25, 0.86]; *short and rare:* (*A*) *R* = 0.43, *p* = 0.1, 95% CIs: [-0.09, 0.76]; (τ) *R* = 0.59, *p*-value = 0.01, 95% CIs: [0.16, 0.84]; *long and rare*: (*A*) *R* = 0.53, *p* = 0.03, 95% CIs: [0.05, 0.81]; (τ) *R* = 0.61 *p* = 0.01, 95% CIs: [0.18, 0.84]). This means that the behavioural integration kernels and RTs of participants evidence integration strategy are directly linked to each other. Longer median RTs means that participants actually integrate on average over a longer period of time to form a decision. However, as shown in Figures 5 and 6 this might be due to individual differences within in the integration kernels rather than due to the influence of different conditions.

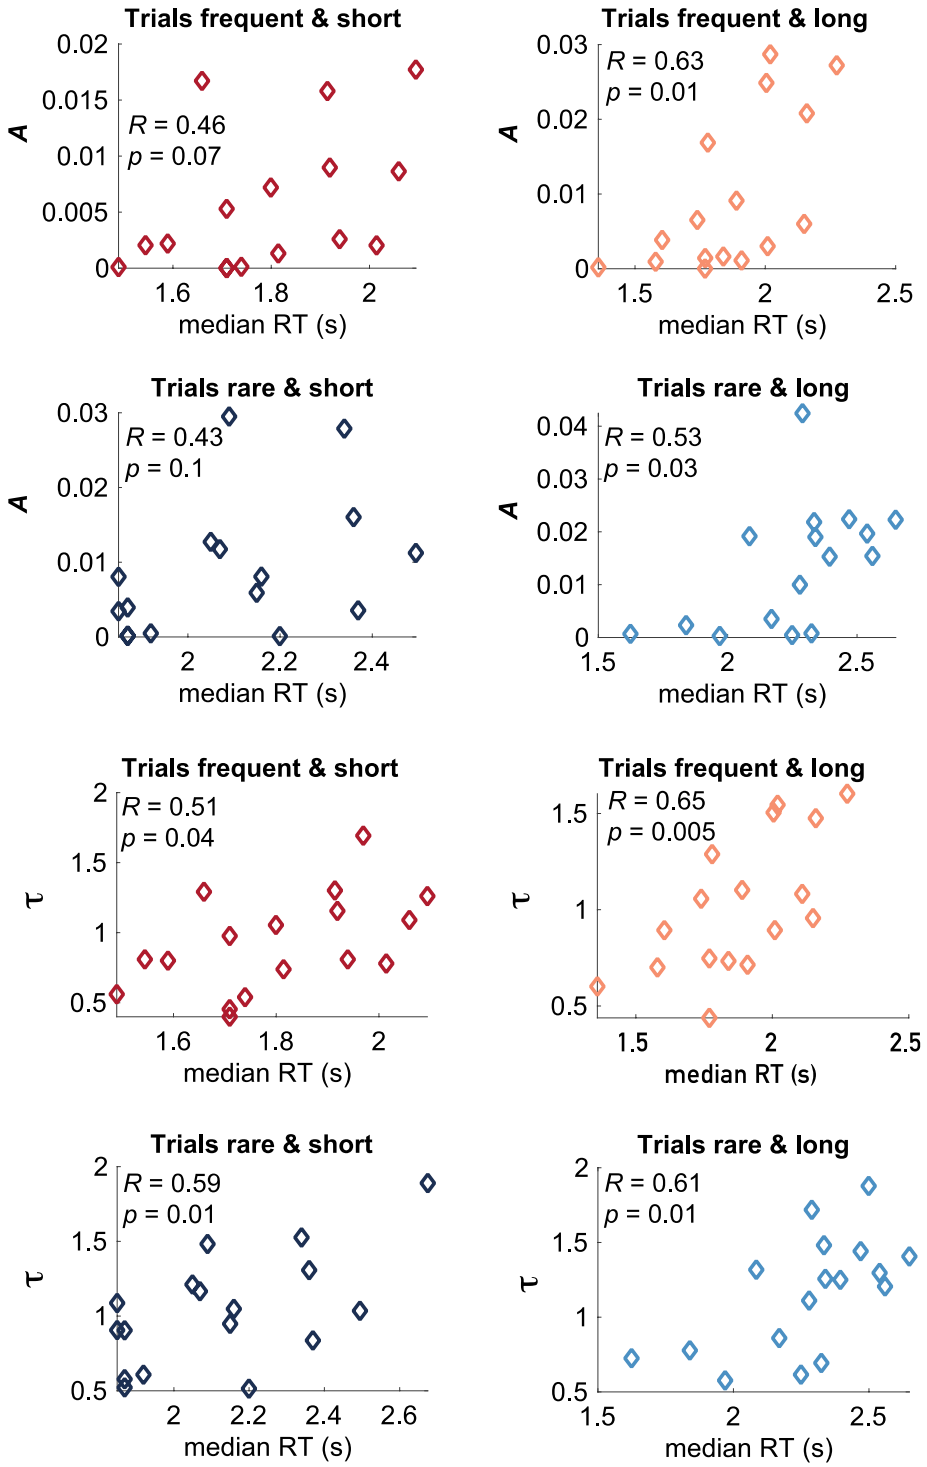


Figure 7: Within subject correlation of median RTs and fitted *A* and τ of false alarms. Correlation between median RT (<= 3.5s) and fitted *A* or fitted τ for each condition. Reported are Pearson's correlation coefficients and p-values. Subjects with less than 20 responses of either false alarms or correct trial responses were excluded, as well as outliers >2standard deviations from the mean.

3.3.6 Trial responses

The integration kernels in the run-up to correct trial responses are also of interest – although there is a shift in the mean coherence during these trial periods, there is also still variation in the stimulus stream across trials, meaning it is possible to perform a psychophysical reverse correlation analysis (Figure 8).

Again, the integration kernels suggest that participants required more evidence in rare trial conditions than in frequent trials. Unlike in the analysis of false alarms, there was also an apparent effect of long versus short trials (Figure 8, upper left panel). However, this effect disappeared when only trials with RTs < 3.5s were included while the effect of trial frequency remained as indicated by a cluster of timepoints that have been detected with a permutation test at an alpha of 0.05 between -4 and 0s (Figure 8, upper right panel). This suggests that the apparent difference in integration kernels for different trial lengths was in fact an artefact of the different permitted response windows in the long- and short- trial conditions, rather than an adjustment of the kernel per se.

It could be argued that there is a sampling bias, in that low-coherence trials were detected less frequently in the rare/short conditions (Figure 1). This would mean that the kernels in these conditions are built from higher coherence trials, and this in turn might affect the shape of the integration kernel. However, If I split the integration kernels by trial coherence (again limiting the analysis to for trials with RTs not exceeding 3.5s), the significant effect of trial frequency on the integration kernels remained (cluster of timepoints between -3s and -1s detected with a permutation test at an alpha of 0.05) while there is still no effect of trial length (Figure 8, bottom panels).

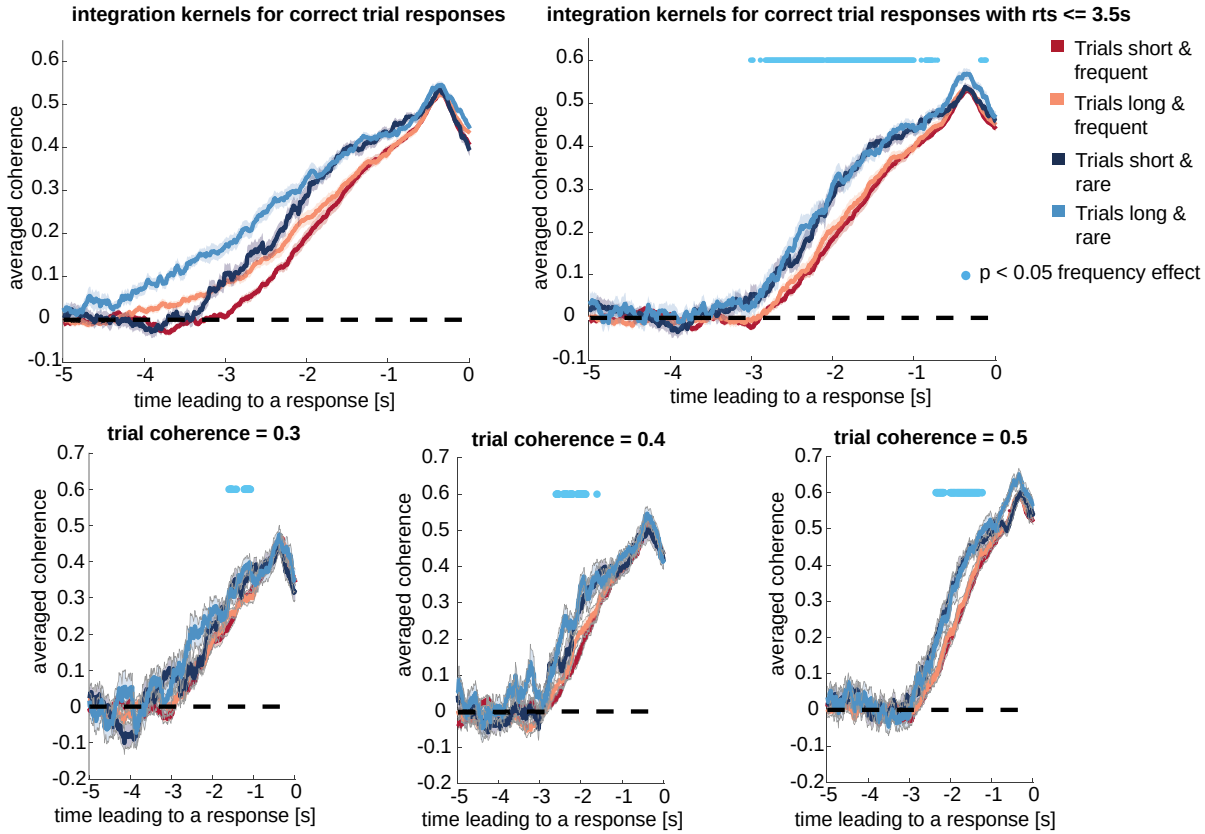


Figure 8: Integration kernels for trial periods. Stimulus streams leading to a correct trial response were averaged across subjects for each condition (mean + s.e.m.). The upper left panel shows integration kernels for all correct trial responses. In the right panel only trials with RTs smaller than 3.5 seconds were considered to make short and long trial conditions more comparable. The lower panel shows integration kernels split for conditions and coherence levels for trials with RTs smaller than 3.5 seconds. I performed a permutation test without replacement to identify significant effects of conditions on integration kernels split for coherence and for integration kernels across conditions for trials with RTs <= 3.5. Circles indicate significant difference between frequent and rare conditions at an alpha of 0.05. There were no significant effects between long and short trial conditions.

3.4 Discussion

I have shown that I can investigate the influence of different task conditions that demand different evidence integration strategies using several different measures of behaviour. Comparing rare versus frequent trials, I found that frequent trials have higher detection rates, shorter RTs and require less evidence for the participant to commit to a response. Comparing short vs. long trials, I found that longer trials have an improved detection rate, longer average RT and lower false alarm rate. There was not, however, any difference in the evidence integration

kernels between short vs. long trials. Surprisingly, I found that the integration kernels for different conditions within subjects are highly correlated.

3.4.1 Evidence for a leaky accumulator model

The behavioural data in this chapter suggests that the training protocol was successful in training participants to solve the task. My behavioural results show that subjects adapted their decision strategies across the different conditions in a manner that is consistent with them maximising the reward on the task. It is unlikely that participants used non-integration strategies such as extrema detection - which can result in similar behaviour (Stine et al., 2020) - because of the difference in the autocorrelation of the stimulus and actual behavioural integration kernels. Therefore, participants have most likely based their decisions on past evidence.

I argue that the behavioural results support a leaky accumulator model as a strategy used to solve the task, because of the similarity of several of these behavioural findings to the OU process simulations. The strongest indication comes from evidence integration kernels that show a decay across time. This fits within the wider literature that has shown flexible adaptation of leaky evidence integration based on within-trial fluctuations of the stimulus and trial length, as well as other cognitive factors such as confidence (Ganupuru et al., 2019; Harun et al., 2020; Ossmy et al., 2013; Piet et al., 2018; Tsetsos et al., 2012). The current study adds to that showing that trial frequency also influences the integration kernel. Therefore, leaky evidence integration might be a common mechanism used by the brain to adapt to evidence integration processes in uncertain environments.

Surprisingly, unlike for trial frequencies, there was no measurable effect of short versus long trials on evidence integration kernels. In contrast, other behavioural measures such as RT distributions show a significant difference between short and long trial conditions (Chapter 3, Figure 3). Additionally, the decay parameter λ for the OU process also supports longer integration kernels for long trial conditions (Chapter 2, Figure 5). Perhaps, the parameter choices for long and short trials might have been too insensitive to pick up differences in the integration kernels between short and long trial conditions. Another possibility is that the different RTs

for long and short trial conditions might be due to other non-decision processes, such as motor preparation or sensory evidence encoding. Fitting a OU process directly to the behavioural data might allow to further investigate whether participants adapt their integration kernel between short and long trial conditions or whether this might be due to other cognitive processes (Ossmy et al., 2013).

3.4.2 Other models of decision-making might explain the behavioural results

Though it looks like the OU process is a strategy that is used by the brain to solve decision-making, it might not be the only possible process that could explain the behavioural results. One possible mechanism that might influence evidence integration could be an urgency signal. An urgency signal is often understood as a function of collapsing decision-bounds that can be of various forms (Drugowitsch et al., 2012; Hanks et al., 2014; Shinn et al., 2019). For example, a study by Shinn et al. (2019) has shown that subjects use leaky integration, but a time-varying urgency signal is important when subjects have to adapt their evidence integration kernels. The adaptation of an urgency signal to different task conditions could explain the differences in behaviour between participants and the OU process. Perhaps, the brain adapts the urgency signal between conditions which might be a reason why there was no difference in the thresholds of the behavioural integration kernels in comparison to the OU process. It has been argued that such an urgency signal is represented in motor preparation signals in the brain (Kelly et al., 2019; Murphy et al., 2016).

If the brain uses multiple mechanisms to control evidence accumulation, it makes sense that a single linear model such as the OU process might not be able to explain the behaviour, specifically when the environment is uncertain. There could be more complex mechanisms involved that influence how past evidence and prior belief shape the evidence integration process. One of these mechanisms might be the build-up of prior expectation about the outcome of a decision. There is evidence that participants use prior belief to guide their decisions and evidence accumulation processes (Braun et al., 2018; Hanks et al., 2011; Urai et al., 2018). For example, it has been shown that the influence of the prior increases during evidence accumulation within a trial (Hanks et al., 2011).

Recent studies have proposed a normative Bayesian model for decision-making under uncertainty in which a non-linear transition function defines the prior belief about the environmental state based on past evidence in order to detect change points between different environmental states (Glaze et al., 2015; Veliz-Cuba et al., 2016). Because of its non-linearity the prior likelihood in the normative model does not increase linearly with increasing posterior belief of the current state of the environment. Instead, prior expectation reaches an absorbing bound based on the volatility, that is the likelihood with which environmental states change. This gives the model an advantage in detecting state changes in contrast to the OU process, in which the volatility of the environment has no influence on the evidence accumulation process.

Therefore, it is plausible that in the current task evidence accumulation is shaped by the expected changes in the environmental states, such as between trial and intertrial periods (Glaze et al., 2015; Murphy et al., 2020; Orsolic et al., 2019; Veliz-Cuba et al., 2016). In these previous studies, subjects had to judge whether they are in one of two states based on noisy evidence. To solve this task participants should balance evidence integration and change detection processes, i.e. the belief that there has been a change in the environmental state. Murphy et al (2020) have shown that a normative Bayesian model can be a better fit to participants performance. However, some studies suggest that under certain assumptions the OU process can be an approximation to this normative model (Glaze et al., 2015; Veliz-Cuba et al., 2016).

My task could in theory be interpreted as a change point task with 3 states. A normative Bayesian model is able to track the posterior likelihood of 3 states, as described by Veliz-Cuba et al (2016). Within this framework, participants have to decide whether they are in an intertrial period, leftward motion period or rightward motion period i.e., the 3 states (see Figure 9A as example, with $H_{1\text{-}3}$ representing the different states). The transition probability (volatility) between states would be described by ϵ. An advantage of the normative model is that these state changes between different states are independent of one another. For example, in my task the transition probability from a leftward to rightward trial period would be 0, whereas the transition probability between intertrial and any trial period would be

defined by the trial frequency. Figure 9B can be understood as describing the likelihood distribution that an evidence sample ξ which is a noisy coherence step in my task has been generated by one of the 3 states, that are each defined by a mean μ. For intertrial periods in my task, the likelihood distribution would be centred around a coherence of 0. Figure 9C-D both show realisations of the log-likelihood ratios for the different states. In my task participants might try to estimate the change point probability. This means they adopt their absorbing bounds of the prior to reflect the change point probability. Maybe, this explains why there was no change in false alarm rates between frequent and rare trial conditions. Perhaps, the brain is able to estimate the transition probabilities between states based on trial frequency but might not be able to adjust other free parameters of the model optimally to minimize the false alarm rate for short trials.

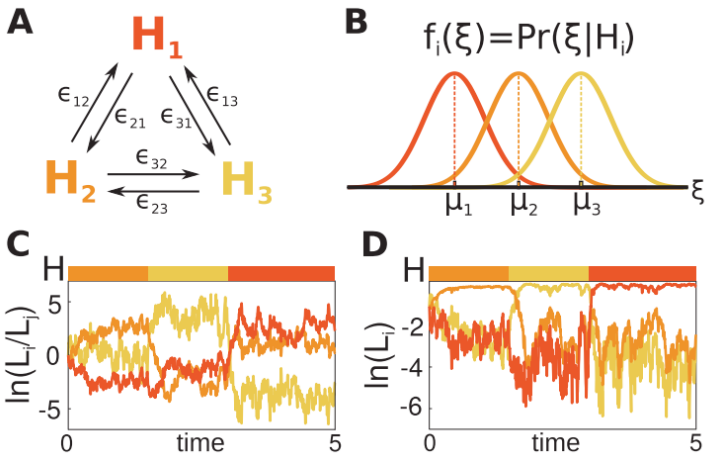


Figure 9: Interpretation of my task as a change point detection task with 3 tasks (adapted from Veliz-Cuba et al., 2016). A) 3 environmental states H_{1-3} that represent the intertrial period, leftward and rightward trial periods that each have a certain transition probability (ϵ) to transition to another state. **B)** Likelihood distributions that describe the likelihood that a certain evidence sample (ξ) stems from a given state. μ represents the average coherence of each state, with a coherence of 0 for the intertrial period. **C)** and **D)** show realisations of the loglikelihood ratios for the 3 different choices.

3.4.3 Single-subject measures of evidence integration

Subject-specific integration kernels might allow for a more fine-grained investigation of evidence accumulation in health and disease. In this chapter, I found that the integration kernels differ for individual subjects but were relatively

stable across conditions within a subject. This indicates that these integration kernels and task design might be useful in the future to draw conclusions about the evidence accumulation process in different populations such as young and old (McGovern et al., 2018) or between healthy people and people with a psychiatric disease such as schizophrenia (Chen et al., 2003).

It has been shown that global motion detection that is necessary to solve my task is abnormal in patients with schizophrenia and that a hypofunction of NMDA receptors might be responsible for abnormal decision-making under uncertainty (Cavanagh et al., 2020; Chen et al., 2003; Salvador et al., 2020). However, to my knowledge, there are no previous studies showing behavioural markers of evidence accumulation at a single-subject level. It will be interesting to investigate how these behavioural markers are related to subject-specific changes in the EEG signal. The task design, together with the convolutional GLM approach (explained in Chapter 4), will allow to further investigate this in the future.

3.4.4 What to look for in the EEG chapter

In the next chapter I will focus on how the different parameters that define my task such as the trial frequency and trial length are represented in the brain. First, I am going to demonstrate that the convolutional GLM approach allows me to analyse the EEG data without time-locking the data to different trial or response events and averaging across them. I will focus on the intertrial period, and make use of the noise structure to investigate how every single jump influences the EEG signal and how that differs between different task conditions. With respect to the behavioural findings, I particularly expect to find differences in the EEG signal between rare and frequent trials based on the integration kernel results. I would predict differences in the ramp-to-threshold between different conditions. I will then discuss how these results relate to behaviour.

4 Neural signals during a continuous evidence integration task recorded with EEG

4.1 Introduction

We understand the neural underpinnings of perceptual decision-making quite well in the context of discrete trials. However, to understand how humans form decisions in a world in which information constantly fluctuates, we need to devise novel task designs and analysis methods. In the past 2 chapters I have introduced a new continuous RDM paradigm to study decision-making under uncertainty and discussed how participants should adapt their behaviour to the environmental statistics. In this chapter, I am investigating the neural signals that underlie the decision-making processes in this paradigm. I am applying a convolutional GLM approach that allows me to examine how the brain processes each new incoming evidence sample but also how it accumulates evidence over time to form a decision.

The gold standard in magnetoencephalography and electroencephalography (M/EEG) analysis is to time-lock the signal to events within each trial (for example, the onset of a stimulus, or the response of the participant), and average it for each kind of event across many trials. This results in an event-related potential (ERP) for each event type that can be compared between different task conditions (Luck, 2014). Yet, there are problems with this approach. First of all, if two events are too close to each other in time, it is possible that the resulting ERPs are mixtures that contain signal from different event types (Ehinger and Dimigen, 2019). Explicitly modelling this overlap can help tease signals apart that belong to different events. A second issue with a discrete trial approach is the amount of recording time needed per subject to achieve a sufficient signal-to-noise ratio (SNR). In a typical EEG task there is usually a lot of ‘dead’ time between trials, also called intertrial periods, which do not contribute to the analysis at a later stage. The fraction of data that comprises the trials is rather small.

In my continuous version of the random dot motion task I can overcome both of these issues. I am using a convolutional GLM approach (Ehinger and Dimigen,

2019; Gonçalves et al., 2014; Lalor et al., 2006; Smith and Kutas, 2015a, 2015b) which does not require me to chop my data up into trials. Instead, I am making use of every sample of the data for estimating the response – thus allowing me to tease apart signals that belong to different events. This approach enables me to investigate the influence of different events presented in rapid succession while increasing the SNR substantially.

The analysis in this chapter focuses on EEG signals measured over (centro-) parietal cortex given their established role in evidence accumulation. (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015). Sequential sampling models, most notably the drift-diffusion model (DDM), have been shown to describe evidence accumulation of neurons in monkey parietal cortex (Gold and Shadlen, 2007). Signals in human parietal cortex are thought to be homologous to the signal measured in monkey LIP (FitzGerald et al., 2015) and the activity varies according to the DDM (Forstmann et al., 2008; Hunt et al., 2012; Kelly and O'Connell, 2013). Most importantly, EEG studies suggest that the P300 component further described as centro-parietal positivity (CPP) reflects a signal that represents the accumulation of evidence independent of sensory modality (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015). It has been shown that the rate of the build-up of the CPP depends on signal strength and correlates with reaction times (RTs). Moreover, the CPP is even present in the absence of a motor preparation signal (Twomey et al., 2016). This suggests that the CPP reflects an intermediate process that purely reflects the accumulated evidence, which is terminated when the motor preparation signal reaches an execution threshold (Kelly and O'Connell, 2013; Kelly et al., 2019; Twomey et al., 2016). In addition, there is evidence that external factors such as urgency and task environments such as speed pressure can influence the CPP build-up and the motor preparation signal (Brosnan et al., 2020; Kelly et al., 2019; Murphy et al., 2016; Steinemann et al., 2018). For this reason, I would expect that in my task the build-up of the CPP varies with different trial conditions and mean coherences in trial periods.

In a dynamic world, the extent to which a given piece of information is relevant for a decision depends on the speed with which things are changing. For example,

selling stocks requires the observation of the stock market for longer periods of time, in order to be certain about nascent trends. In contrast, riding a bike on a busy street trying to switch lanes requires evidence over only the last few seconds. In order to study how the brain deals with this uncertainty, it is possible to examine how beliefs are updated in response to single pieces of information. Studies in monkeys have shown that lateral intraparietal area (LIP) encodes the log-likelihood ratio of a decision outcome with each new discrete evidence sample in a evidence integration task. In contrast to the random dot motion task in which evidence is constant during a trial, in this task evidence slightly fluctuates across samples (Kira et al., 2015; Yang and Shadlen, 2007). Based on these results studies in M/EEG in humans have shown that we find similar signals in human parietal and frontal cortex that represent the size of the update of each incoming evidence sample (Cheadle et al., 2014; De Lange et al., 2010). They found that the size of the update depends on whether a new sample supported the target that is supported by the evidence accumulated so far. Critically, each new piece of information was important to judge which decision was more likely correct. However, this does not explain how the brain deals with the uncertainty of whether a piece of new information is informative or belongs to some random noise.

In recent years, some studies have looked at what model the brain might use to solve decision-making under uncertainty (Glaze et al., 2015; Murphy et al., 2020; Ossmy et al., 2013; Piet et al., 2018). Some of them have shown that leaky accumulation can explain the behavioural adaptation (Ossmy et al., 2013; Piet et al., 2018). However, others propose a Bayesian normative model with a non-linear transition function that is a better fit to behaviour. For example, it has been shown in humans and mice that mostly motor preparation signals encode these normative decision variables, while lower visual areas seem to encode more evidence related variables (Murphy et al., 2020; Orsolic et al., 2019). However, these studies have not looked at how new pieces of information are updated in the brain when their relevance for an upcoming decision is not known.

In this chapter, I analysed the EEG signal recorded while participants were performing the continuous RDM. The behavioural results presented in the previous chapter imply that participants seem to adapt their strategy specifically between

frequent and rare trial conditions. Behavioural integration kernels show that participants needed more evidence to commit to a decision in rare trial conditions than frequent. The reduced detection rates for rare trial conditions reflected this, as well. However, there is little evidence for an adaptation of the integration kernel between long and short trial conditions. In this chapter, I investigated how these changes in strategy between different conditions are implemented in the brain. In addition, my paradigm allowed me to examine the decision-process in real time by making use of the high frequency of the changes in the stimulus. This enabled me to examine the EEG response to each new stimulus.

First, I investigated how EEG signals which precede a response differ between different conditions by performing a conventional ERP analysis on my data, time-locking to trial start and responses. I found that rare trial conditions show a significantly earlier onset of a negative deflection in signal than frequent trial conditions leading to a response. Next, I implemented a convolutional GLM analysis and verified it through comparison with the results from the conventional ERP analysis. Convolutional GLMs were then used to investigate how the brain updates new incoming evidence during intertrial periods. This analysis revealed that participants have different prior expectations about the stimulus variation in rare and frequent trial conditions. Using a control experiment with a decision irrelevant motion direction (vertical) that has the same statistics as the decision relevant motion direction (horizontal) I showed that the decision irrelevant motion direction led to a diminished brain response compared with the motion relevant direction. This means that the EEG signals measured during the intertrial period actually reflected processes of evidence integration and are not merely reflecting a form of surprise to big changes in the stimulus or just reflecting the visual input.

4.2 Results for the conventional ERP analysis

4.2.1 ERPs time-locked to response

I hypothesise that there is a difference between frequent and rare trial conditions in EEG signals measured during trial periods. There should, if at all, only be a minimal difference between short and long trial conditions. Behavioural results show that participants needed significantly more evidence to commit to a decision in rare trial conditions in comparison to frequent. However, there were only weak differences between integration kernels for short and long trial conditions (see Chapter 3). Previous studies investigating evidence accumulation in humans using EEG have reported a CPP signal that ramps up over the time of a trial. Conventional ERP analysis time-locked to either trial start or response reveals that this ramping activity is scaled by the strength of evidence in the trial with stronger evidence leading to increased ramping activity (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015). For this reason, I expect that the ERPs measured in this study show ramping activity that is modulated by trial frequency and mean trial coherence in my paradigm. In this section I investigated the influence of different conditions on ERPs time-locked to trial start and response.

ERP signal was time-locked to correct responses of each participant. We observed a strong negative deflection in the ERP signal at centro-parietal electrodes preceding the button press response that resembles a contingent negative variation (CNV) signal (Figure 1E,F). CNV signals are slow negative ERPs that have been studied with respect to temporal expectation of an behaviourally relevant event or response (Macar and Vidal, 2004; Nobre et al., 2007; Pfeuty et al., 2005; van Rijn et al., 2011; Trillenberg et al., 2000). They are typically located in areas that are related to motor preparation (Nobre et al., 2007; Rektor, 2002). The CNV in my study was localised near the midline but extending slightly more anterior (Figure 1E).

Surprisingly, I did not initially find a CPP-signal which I would expect to show a positive deflection (Kelly and O'Connell, 2013; O'Connell et al., 2012). Perhaps,

the strong CNV signal masked this signal. For that reason I performed a current source density (CSD) transform, a method that has been used previously to separate the source of different EEG signals underlying different cognitive processes (Kelly and O'Connell, 2013). The analysis of the CSD transformed data shows that there was indeed a positive ramp across parietal cortex that is similar to the CPP signal in other studies (Supplementary Figure 1). The continuous design of my task allows me to study evidence integration in the absence of motor preparation during intertrial periods. In addition, using convolutional GLMs I can separate overlapping EEG signals without a CSD transform. I will first discuss traditional time-locked ERP results without CSD transform in order to compare them to the results from the convolutional GLM analysis. Then I will apply this method to investigate how the EEG signal adapts to new evidence samples during the intertrial period.

There was a significant effect of trial frequency on the CNV signal which suggests that participants adapted their temporal expectation of an upcoming trial (Figure 1 A,C,G,I). For correct responses to rare trials, the negative deflection in the ERP started significantly earlier and reached a significantly more negative threshold in comparison to frequent trials (cluster-based 2-tailed permutation test, with 3 significant clusters: (1) p = 0.005, (2) p = 0.004, (3) p=0.01 at a priori selected centro-parietal electrodes at an alpha of 0.05, depicted by white triangles in Figure 1E, in the 2-second window before the button press, Figure 1A). A difference waveform between frequent and rare trial conditions averaged over subjects shows that the greater negativity for rare trials was consistent across subjects (Figure 1C). The same trend was visible for false alarms, though there is only one significant cluster (p = 0.007) that indicates an earlier negative deflection in the ERP for rare trials but the difference in thresholds 0.5s before the button press was smaller (Figure 1G+I). This indicates that participants are less committed to respond when trials are rare.

In contrast, trial length had a weak influence on response-related EEG signals. For both false alarms and correct trial responses, there was an earlier onset of a negative deflection in the EEG signal starting about 2s before a response for long trials (Figure 1B,D,H,J). For false alarms this difference was significant (p = 0.01)

but not for correct trial responses. Further, for short trials brain activity seemed to reach a slightly more negative threshold, which might indicate that participants needed stronger evidence for short trials over a shorter time period to commit to a decision (Figure 1B,H). These effects of trial length were also apparent in the difference waveforms that show quite consistent positive difference for the onset of integration at about 2s before a response and a negative difference for the threshold reached (Figure 1D,J). However, the difference in threshold was not significant.

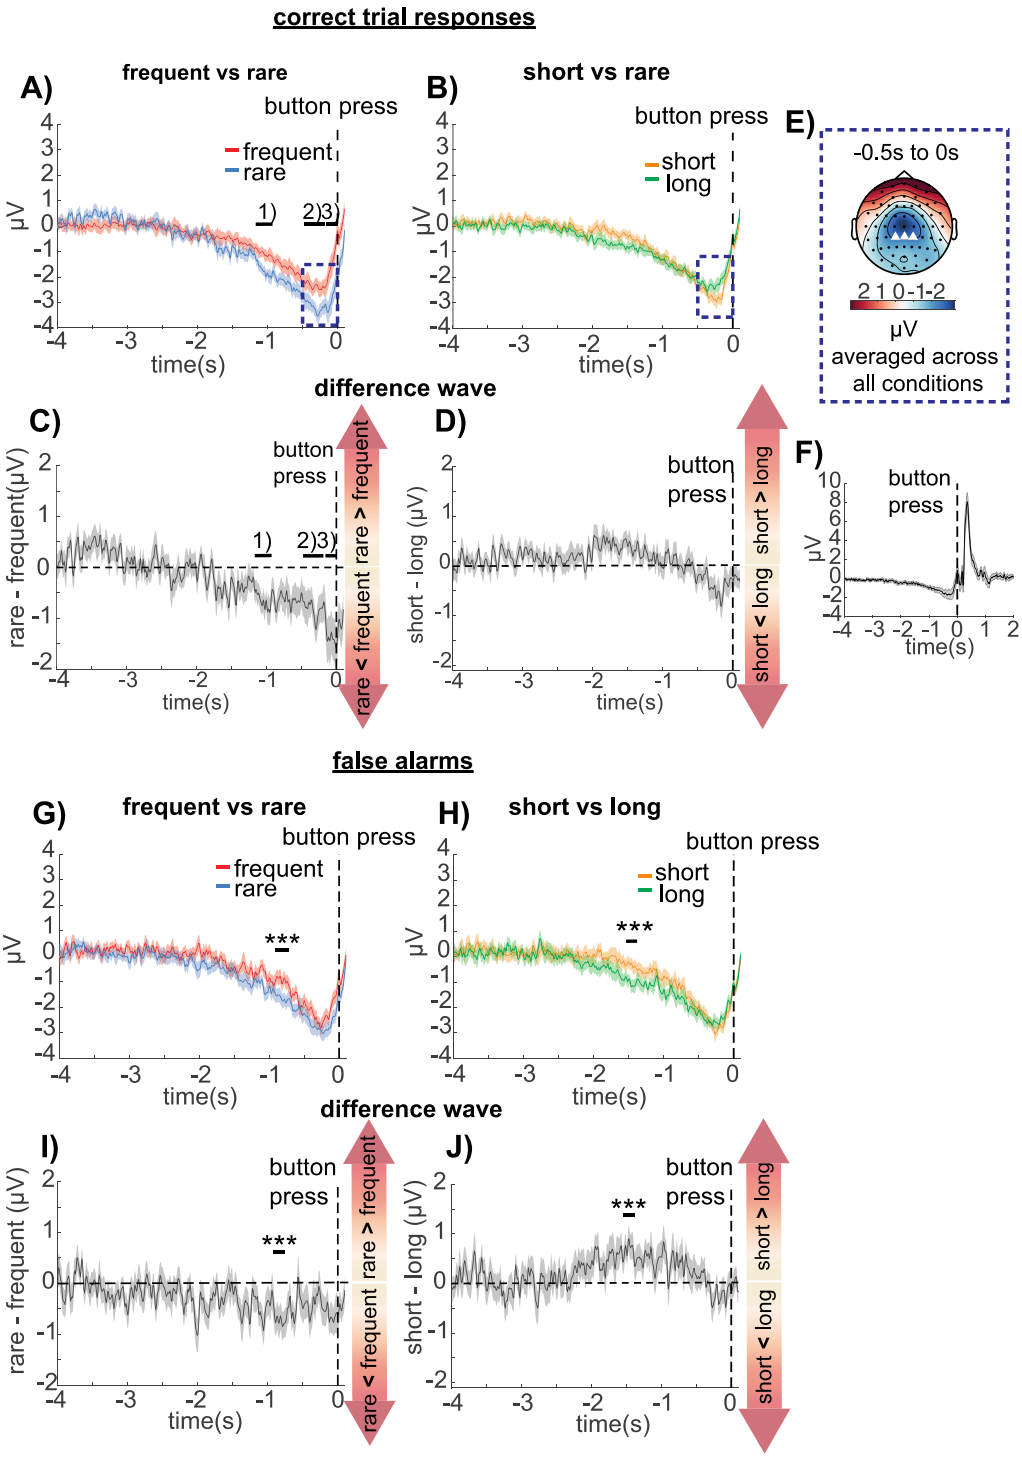


Figure 1: Response ERPs. This figure shows response ERPs on the group-level for correct trial responses (upper panels (**A-D**) and false alarms (lower panels **G-J**) averaged over centro-parietal electrodes CP1, CP2 and CPz with standard error (s.e.) of the difference between frequent and rare. **A)** shows averaged ERPs for correct responses in frequent (red) and rare (blue) trial conditions and 3 significant clusters (1) p = 0.005, 2) p = 0.004 3) p = 0.01, 2-tailed cluster-based

permutation test with a priori selected centro-parietal region and time window) **B)** shows ERPs for correct responses for short (yellow) and long (green) trial conditions. **C)** and **D)** show difference waves for rare versus frequent as well as short versus long that were calculated separately for each subject and then averaged. **E)** is a topography plot averaged across all conditions for the time window framed by the blue squares in A) and B) which is from 0.5s before a button press until the time of a response. The 3 triangles indicate the centro-parietal electrodes (CP1,CP2, CPz) for which I plotted the timeseries. **F)** Shows the average ERP for a correct response over a longer time window after the response. **G)** and **H)** are the same as A) and B) respectively just for false alarms with a significant cluster of $p = 0.007$ for G) and $p = 0.01$ for H). **I)** and **J)** depict the difference waves similar to C) and D)

4.2.2 ERPs time-locked to trial start

ERPs time-locked to trial starts showed weak differences, that were not significant (tested with a cluster-based permutation test), between different conditions and the overall signal had a smaller scale than the ERP time-locked to responses (Figure 2). Only trials with correct trial responses were included in this analysis. Two sources of signal were observed over the time course of 3-4s after a trial start. At the start there was a central negativity which was followed by a negative signal in the frontal part of the brain and a positive signal in parietal areas. The positive signal might be the centro-parietal positivity signal observed before (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015) , while the negative signal in the front has been described as a form of an urgency signal (Brosnan et al., 2020; Kelly et al., 2019; Murphy et al., 2016). Figure 2A,C shows that there were only small differences between frequent and rare as well as short and long trial conditions at parietal electrodes. After the initial dip in activity, the EEG signal for rare trial conditions seemed to increase and reach a positive peak over a longer period than for frequent trials which could reflect the need of more evidence for rare trial conditions across participants. The effect of trial length (Figure 2B,D) showed only a very small difference in a higher peak for short trials, that might reflect less evidence needed for a decision. However, none of these results were significant.

Time-locking the EEG signal to trial starts might not be fully appropriate for this paradigm, as it is not apparent for participants when a trial starts, and they could misjudge it either believing a trial starts earlier or later than the actual starting point in the stimulus. Moreover, it is also not possible to completely unmix signals locked

to responses from those locked to trial start. Both of them will overlap to some extent, and the amount of overlap depends on how participants judge the start of trial periods – which we do not know – and respond to it. For this reason, it might be useful to use a different analysis technique to investigate how much of the EEG signal during trial periods is due to the preparation of a response and what part of the EEG signal might reflect a more general evidence integration process. In the next part I will show how convolutional GLMs allow me to separate these signals and verify them as a useful method to analyse continuous EEG data.

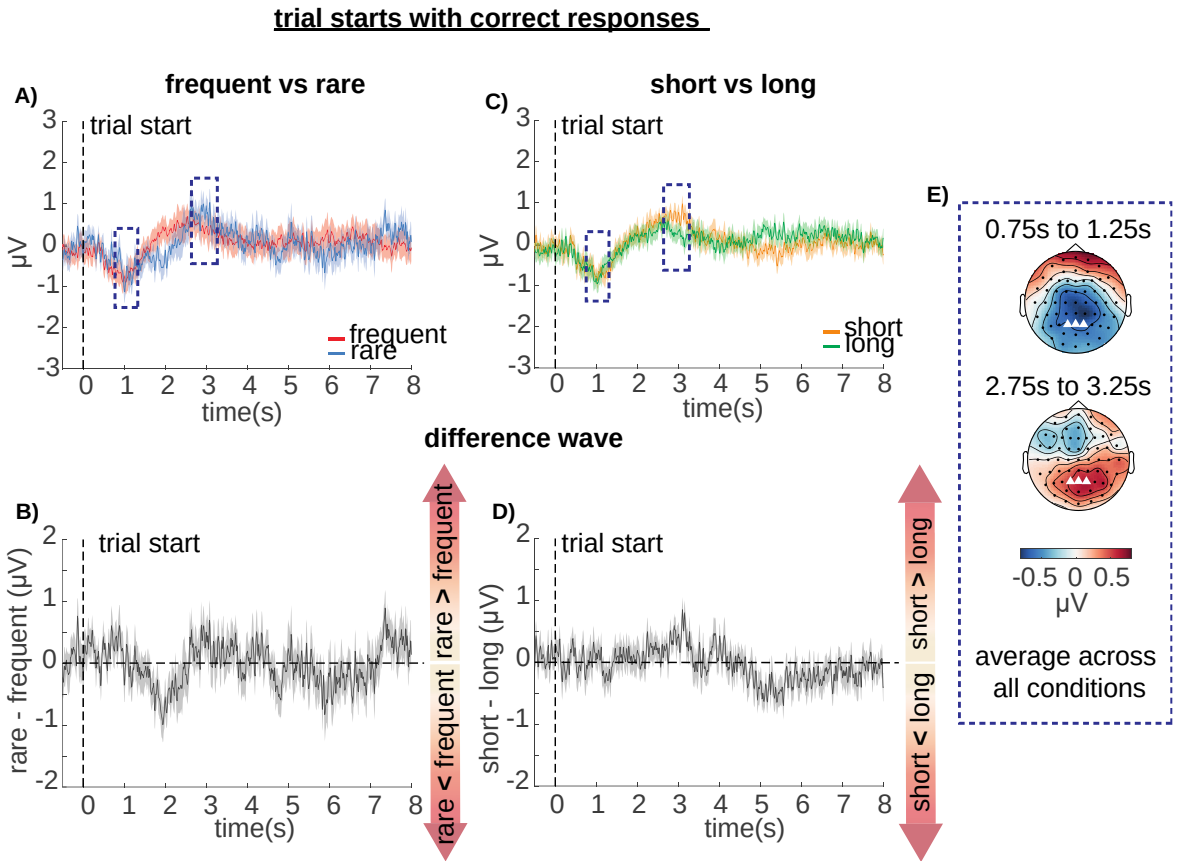


Figure 2: ERPs time-locked to trial starts. **A)** depicts ERPs time-locked to trial starts for frequent (red) and rare (blue) trial conditions averaged across subjects with s.e. of the difference between frequent and rare on the group-level at parietal electrodes (P1, P2, Pz). **B)** shows the same as A) but for short (yellow) vs long (green) trial conditions. **C)** and **D)** are the difference waves for rare versus frequent and short versus long calculated for each subject and then averaged across subjects. **E)** Topographies for blue squares in A) and B) averaged across all conditions.

4.3 Method and results for the convolutional GLM analysis

Convolutional GLMs allow for the investigation of the influence of events on EEG without the need to take windows of data around certain events. Others have used this approach to separate and measure visual evoked potentials, and to remove confounds from the EEG signal. They called this technique visual evoked spread spectrum (Lalor et al., 2006), regression based ERPs (Smith and Kutas, 2015a, 2015b) or linear deconvolution (Ehinger and Dimigen, 2019). This approach is different to the traditional use of GLMs in ERP analyses in which a single regressor for a condition would be used and applied across different chunks of data (epochs or trials) and separately for a given EEG channel. In contrast, in the convolutional GLM, for each event such as a trial start there will be a set of regressors, one for each time step. This means that there is a regressor to look at the influence of trial start on the EEG signal at t_0, the next regressor will look at timepoint t_1 and will be repeated until a time at which we expect that there is no influence of the event on brain activity anymore.

If events are sufficiently far apart in time, this approach is mathematically equivalent to (and exactly returns) the conventional event-related potential (Smith and Kutas, 2015a). Crucially, however, if events appear rapidly after each other leading to overlapping signals in the EEG data, the convolutional GLM can be used to separate different ERPs for different events. In my continuous random dot paradigm, the coherence of the dots changes every hundreds of milliseconds (see Chapter 2, Figure 1). Participants have to integrate over these ‘jumps’ in order to determine whether they are in a trial period. However, in this case these jumps appear so rapidly that the influence of one jump on the EEG signal will overlap with the effect of previous and future jumps that follow directly after the current jump. This means that a standard ERP analysis will be confounded. For this reason, employing a convolutional GLM is crucial, as it allows me to explicitly model the overlap of signals to different events. In turn, this enables me to investigate how participants integrate evidence between trials and how the brain updates incoming evidence in a continuous and noisy environment.

4.3.1 Method of the convolutional GLM

For the convolutional GLM analysis, regressors for different events and timesteps in continuous EEG recordings are built. This means that the EEG is not epoched, but instead we analyse each 5-minute block as a whole. Figure 3 shows a simulated continuous EEG trace. The goal of the convolutional GLM is to explicitly model the evoked response of multiple overlapping events, e.g. condition A and condition B in Figure 3 (Ehinger and Dimigen, 2019). The design matrix is built in a way such that each regressor has as many entries as there are timepoints in the EEG recording which might be influenced by any given event type. In its simplest form we can use binary regressors to investigate the evoked response of different overlapping events. This means that for each event we have a number of regressors to investigate the influence of the event on the EEG signal at different timepoints. The first regressor for an event, for example condition A, will have a ‘1’ at each timepoint the event occurs. Given that the event is likely to influence subsequent timepoints as well, the next regressor will have a ‘1’ one sample later than the previous one. This is continued for as many timesteps as one expects an influence of the even on the EEG signal. A next regressor set, equally built, is used to model condition B.

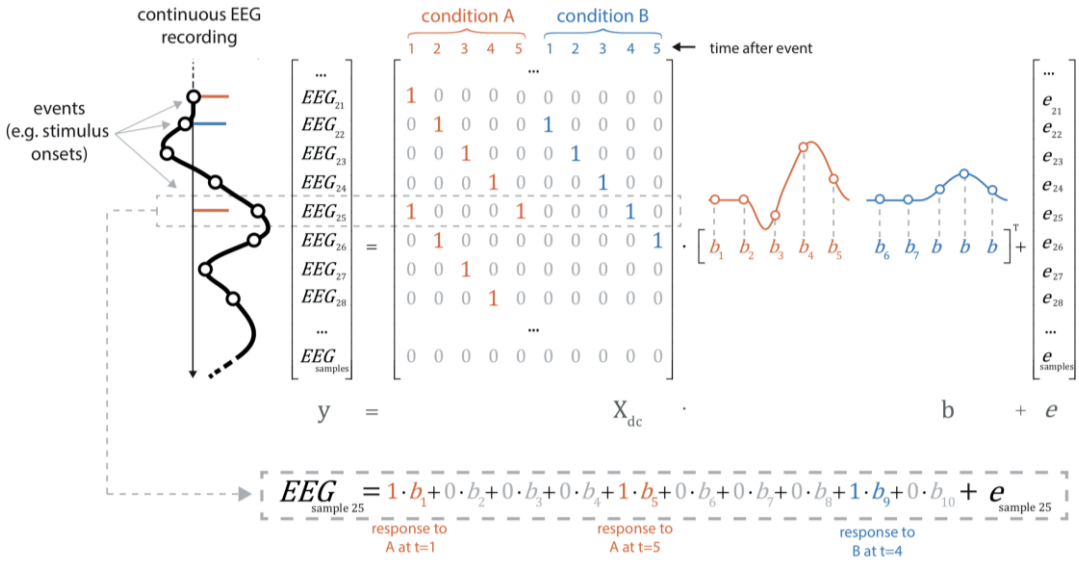


Figure 3: Schematic of a convolutional GLM. On the left is a simulated continuous EEG recording with 2 different events. These two events, condition A and condition B, are very close together in time. In order to differentiate between the ERPs for both of them, a convolutional GLM approach is shown. For each timestep there is a different regressor for both condition A and condition B. Plotting the effect size for each regressor reveals the ERPs for both conditions. The Equation below the figure shows the equation for the convolutional GLM for EEG sample 25 (Figure taken from Ehinger & Dimigen 2019).

For my analysis I implemented the same approach. First, I investigated the regression ERPs for trial start and responses. I used binary regressors for trial starts that were followed by a correct response, and for correct responses as well as false alarms. Further, I included regressors to account for differences between left and right button presses [1 right response, -1 left response] for correct responses and false alarms (these are not shown in the results). This analysis was conducted to verify my analysis pipeline, and compare the findings from the convolutional GLM with those obtained via conventional event-related averaging as in Figures 1 and 2. In addition, I repeated the convolutional GLM analysis adding a regressor to investigate the influence of the mean trial coherences at 30, 40, 50% on correct responses. Previous studies have focused on the influence of coherences on evidence integration, but, in contrast to my study, they held the coherence constant during a trial.

The main reason for the convolutional GLM analysis is to examine brain processes during intertrial periods. This is particularly useful because the signal during this period is not contaminated by motor activity or trial onsets. It allows me to purely look at the decision update of each new sample. In addition, this approach increases the SNR massively. Within each 5 minute block there are about 3000 jumps in the noise. Roughly 750 of these jumps will fall within trial periods, leaving over 2000 jump events that I can investigate in only one out of 6 blocks per condition. This is in stark contrast to typical trial-based designs which have a lot fewer events.

I set up a convolutional GLM with different regressors for changes in coherences during intertrial periods. The first regressor, called jump regressor, identifies the main effect of a change in coherence during intertrial periods, irrespectively of the size of the change in coherence (Figure 4, upper panel, a). This means that this regressor is binary. Further, I have implemented a regressor to examine the influence of the size of the change between two jumps in coherences. I call this regressor 'prediction error' in the future (Figure 4, upper panel, c). It describes the absolute difference in coherence between a jump at t_0 and a jump at t_1. Another regressor I included describes the new absolute coherence level after a jump occurred (Figure 4, upper panel, b). This regressor tracks the amount of evidence presented after a jump occurred and is called jump level regressor. Both prediction error and jump level regressor were demeaned. An additional regressor was added to track the overall evidence presented at each timepoint (Figure 4, subpanel 5). This is the 'absolute stimulus' regressor. If the brain implements an OU process to integrate evidence, then I would expect there to be a signal that reflects an integration kernel based on past evidence as discussed in Chapter 2. To ensure that none of these regressors share too much variance with each other I calculated the correlation (R^2 value) for each regressor pair (Table 1), showing that the regressors measure different things. Once all regressors were defined I created a design matrix with lagged regressors for the convolutional GLM. Figure 4 shows a schematic of how a design matrix is created for different regressors. All regressors are shown at t_0 but also for the two following two timesteps to depict how the same regressor is used but just horizontally shifted to extract the effect of a jump over time. Additionally, I included confound regressors for false alarms, correct

responses and trial starts in this GLM. Table 2 describes the size of design matrix including how many rows each regressor occupied. Then I calculated the effect sizes for each channel and block of a session separately, before I averaged across conditions and subjects. Difference waves were calculated in the same manner as for the conventional ERP analysis for correct responses and false alarms. For effect sizes of trial start only an average across all conditions was calculated.

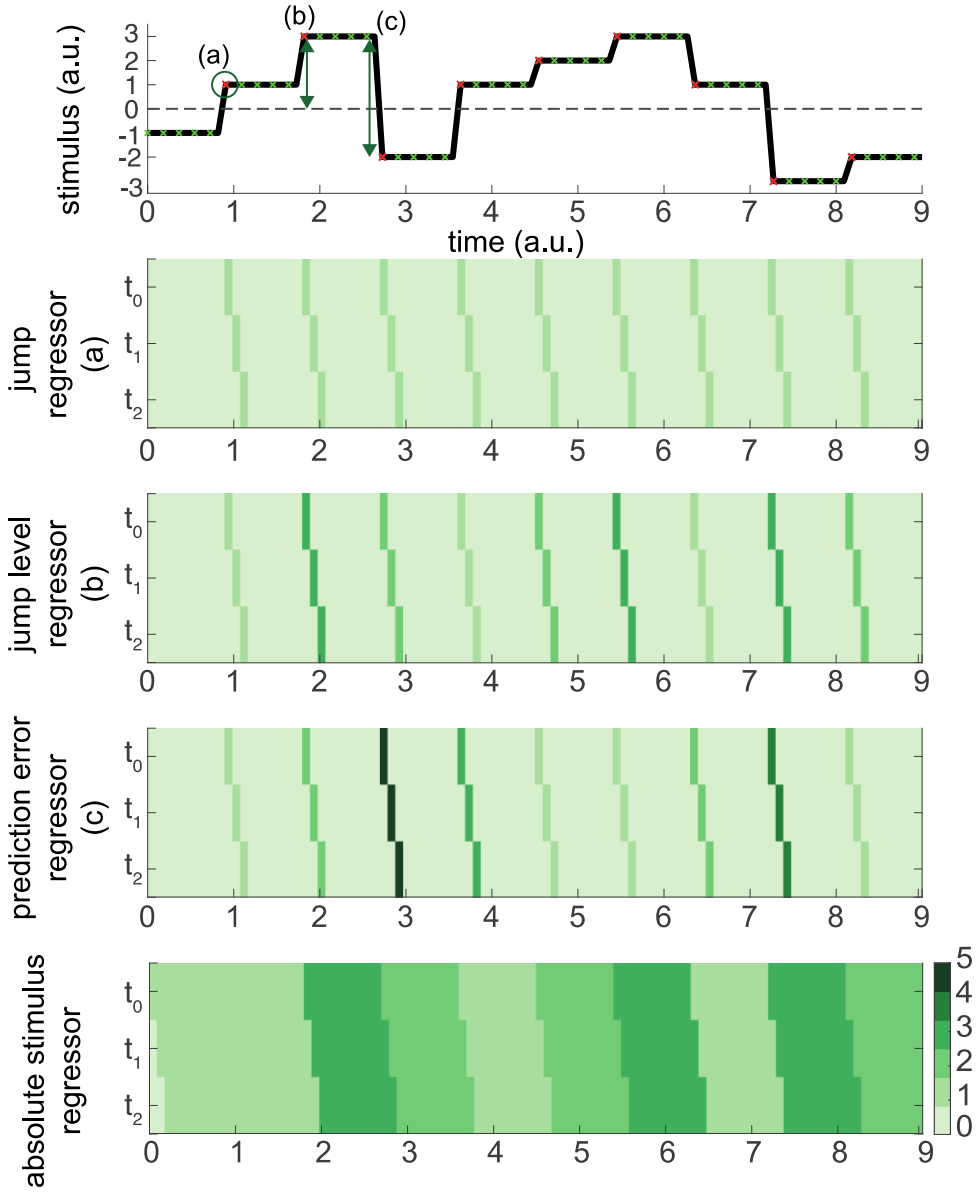


Figure 4: Schematic of regressors for investigating intertrial periods. The first subplot shows a simulated stimulus stream. The jump regressor in the second subplot is

a binary regressor which is 1 every time there is a jump. For this and the following regressors there are 3 example regressors for different lags including the regressor at time t_0, t_1 and t_2. The jump level regressor is the absolute level of coherence after a jump (3rd subplot). In the 4th subplot is an example for the prediction error regressor which is the absolute difference in coherence between the previous and current jump level. The final regressor in subplot 5 is the absolute stimulus regressor which is the entire stimulus stream absoluted.

Table 1: Explained variance between regressors of the convolutional GLM for intertrial times. Columns and rows are the different regressors used to investigate intertrial periods. Between each regressor pair the explained variance was calculated to make sure that the regressors were not correlated with each other.

R^2	Jump	Jump level	Prediction error	Absolute stimulus
Jump	1	0	0.01	0
Jump level	0	1	0.18	0.03
Prediction error	0.01	0.18	1	0.01
Absolute stimulus	0	0.03	0.01	1

Table 2: Design matrix. This table describes the size of the design matrix assuming a sampling frequency of the EEG signals of 100Hz. For each regressor I stated the number of lags pre- and post event and the total number of rows this regressor covers in the design matrix.

Regressor	Pre-event lags	Post-event lags	Total rows in the design matrix
Jump	100	150	250
Prediction Error	150	150	300
Jump level	100	150	250
Absolute stimulus	150	150	300
False alarms	500	350	850
Correct responses	500	350	850
Trial starts	50	800	850

4.3.2 Replication of findings from conventional analysis

In order to establish the convolutional GLM approach for examining evidence integration for non-trial periods, I verified this approach by reanalysing the ERPs for responses and trial starts. I should be able to get very similar results to what I have shown before with the conventional analysis in section 4.2. In addition, the convolutional GLM approach should be able to separate the variance explained by trial starts and motor responses. This was not possible with the conventional ERP analyses, as these average across epochs of data around either trial start or response, and signals related to response and trial start will overlap partially. For the convolutional GLM approach I included binary regressors for successful trial starts, correct responses and false alarms.

I found that I could replicate the results of the conventional ERP analysis with the convolutional GLM approach (Figure 5). For both correct responses (1-tailed cluster-based permutation test and an alpha of 0.05, $p = 0.03$ with a priori selected centro-parietal region and time window of 2s before a response) and false alarms (2-tailed cluster-based permutation test with 2 significant clusters 1) $p = 0.02$, 2) $p = 0.006$) rare trials showed an earlier negative deflection in the ERP. There were still small effects of trial length on false alarms (1-tailed cluster-based permutation test and an alpha of 0.05, $p = 0.04$ with a prior selected centro-parietal region and time window of 2s before a response) and correct responses (1-tailed cluster-based permutation test and an alpha of 0.05, $p = 0.04$). However, the earlier onset of a negative deflection in the ERP for long trials before correct responses disappeared. These differences are likely due to regressors for trial start which implicitly models the effect of RTs. In the ERP analysis I don't control for the fact that different conditions have different RT distributions (see Chapter 3, Figure 3). Mathematically, the binary regressors for false alarms and correct responses should be the same as taking the average over epoched EEG data. The only other reason for slightly different results, specifically for correct responses, might be the improved distinction from signals evoked by trial starts.

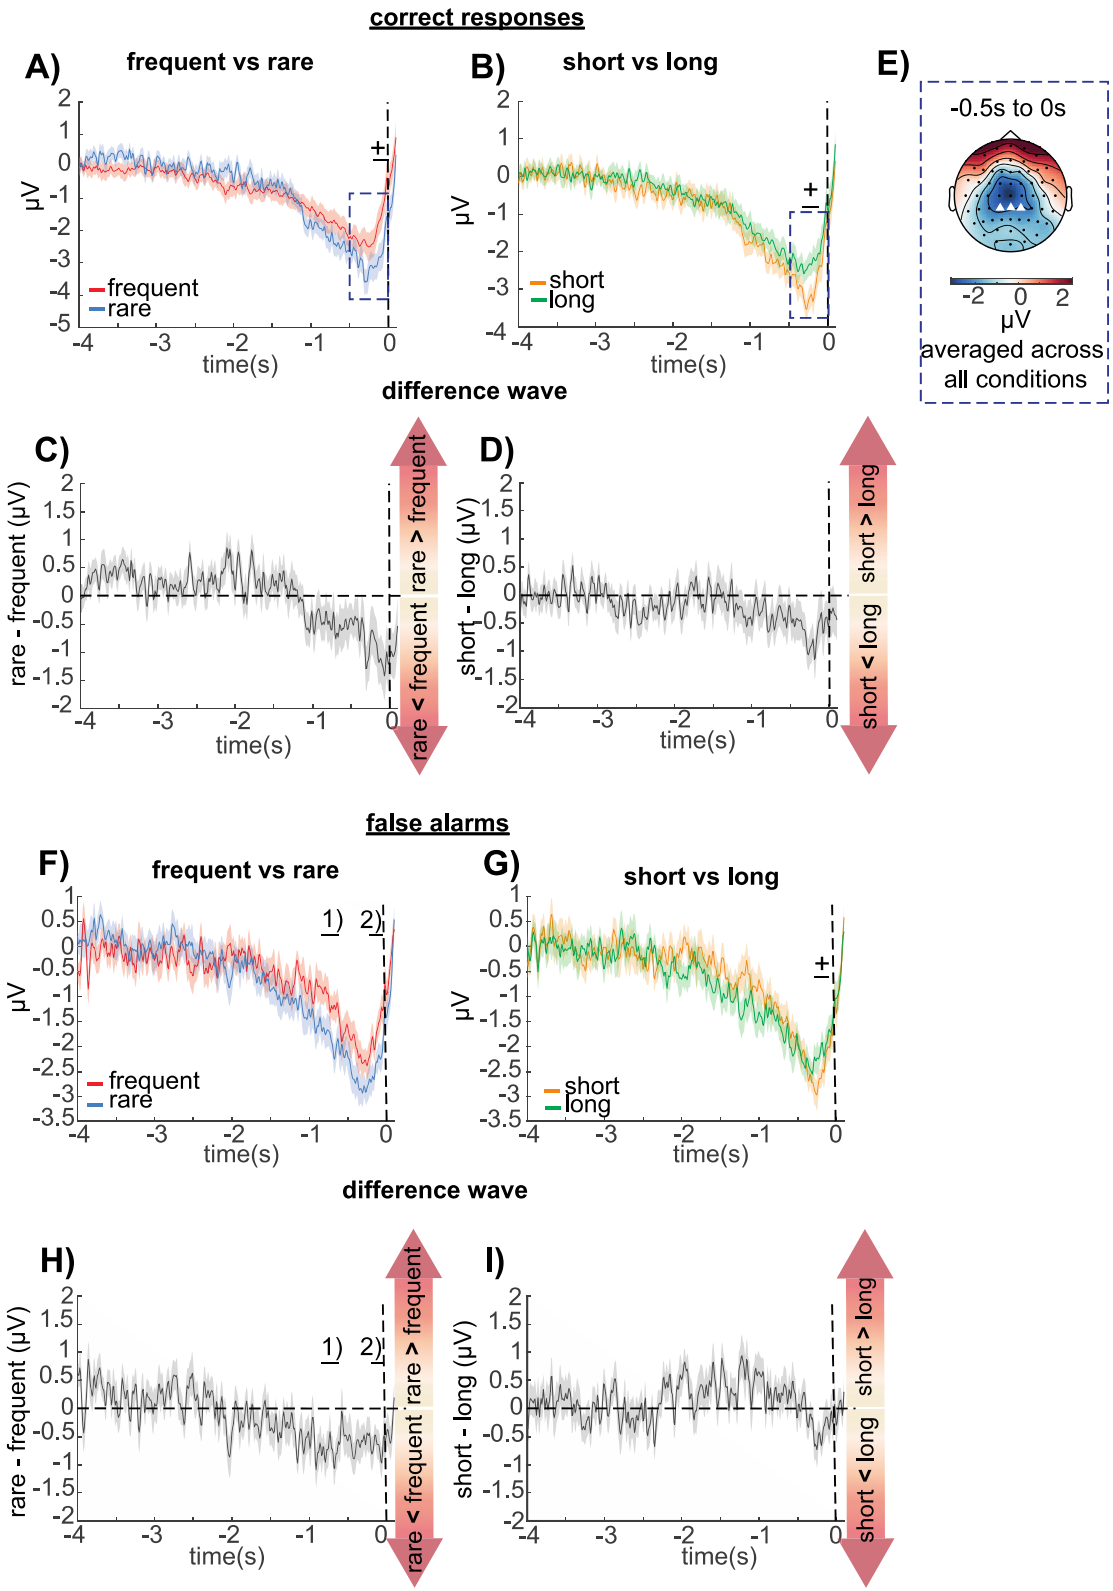

correct responses
A)
frequent vs rare
B)
short vs long
E)
-0.5s to 0s
μV
frequent
rare
short
long
time(s)
-2 0 2
μV
averaged across
all conditions
difference wave
C)
D)
rare - frequent (μV)
short - long (μV)
rare < frequent
rare > frequent
short < long
short > long
false alarms
F)
frequent vs rare
G)
short vs long
1)
2)
difference wave
H)
I)

Figure 5 (previous page): Regression ERPs for correct responses and false alarms. **A)** Regression ERPs for correct responses and frequent and rare trial conditions with s.e. of the difference between frequent and rare on the group-level at centro-parietal electrodes CP1, CP2 and CPz and a significant cluster with a 1-tailed test between -0.5s and 0s. **B)** Same as in A) but for short and long trial conditions. Significant cluster with a 1-tailed test between -0.5s and 0s. **C)** and **D)** difference waves for rare versus frequent and long versus short, respectively. **E)** Topography for time window -0.5 to 0s averaged across all conditions. White triangles indicate electrodes used for plots in A) -D). **F)** - **I)** same as A) - D) but for false alarms. 2 significant clusters with a 2-tailed test in **F)** between -1s and 0s. Significant cluster with a 1-tailed test for **G)** between -0.5s and 0s.)

The convolutional GLM approach revealed a smaller influence of trial starts on the EEG signal than shown with the conventional ERP analysis conducted in 4.2.2 (Figure 6). This means that most of the variance in EEG signal seen at the time of the trial start (Figure 2) is actually explained by the CNV signal. As stated above, this is not very surprising because participants cannot be sure about the start time of a trial period due to the noise structure of the stimulus. Participants might perceive the start of a trial period as earlier or later than the actual timepoint. This is quite different to results reported by other studies that show a distinct CPP time-locked to trial start in continuous paradigms (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015). However, their intertrial and trial periods were not overlayed with additional noise, meaning that a trial start is much easier to detect and might therefore lead to different results.

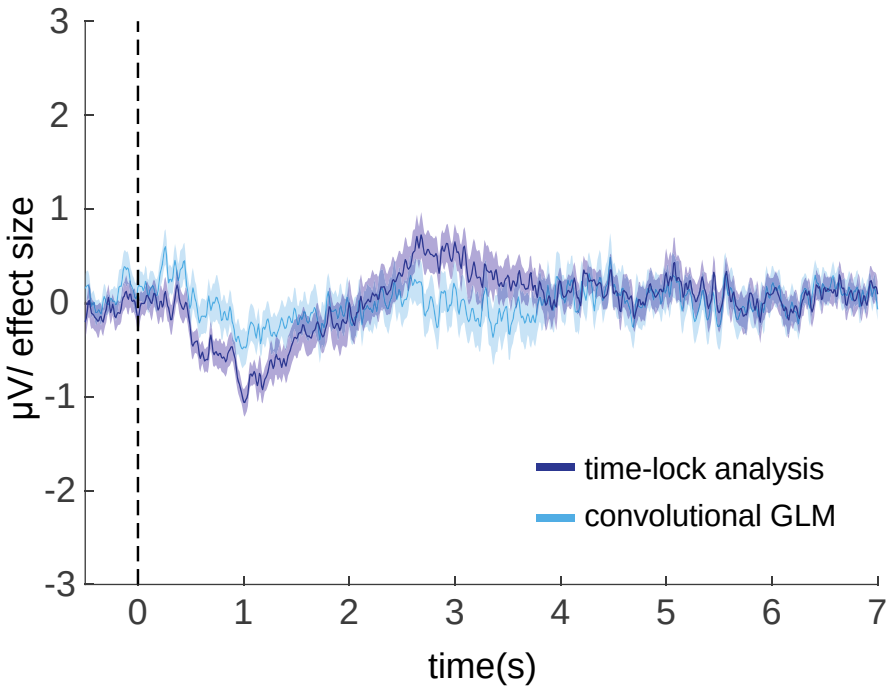


Figure 6: Comparison of time-locked and regression ERPs for trial start. Plotted is the ERP on the group-level and standard error of the mean (s.e.m) for the conventional ERP analysis averaged across conditions (dark blue) and the regression ERP on the group-level from the convolutional GLM analysis with s.e.m. (light blue). Both signals are plotted at parietal electrodes P1, P2 and Pz.

Mean trial coherences did not have an impact on response-locked EEG signals (Figure 7). It is well known that strength of evidence controls the speed of the ramping signal observed over centro-parietal electrodes (Kelly and O'Connell, 2013). Therefore, I expected that the mean trial coherence in the continuous task would influence the negative ramping activity leading to a response. I repeated the convolutional GLM analysis and added a regressor controlling for the mean coherence of a trial for correct responses. Surprisingly, I did not find any effect of the mean coherence on the ramping activity. A possible explanation for this is that participants paid more attention to the noisy jumps than the average of coherence at any time. In previous studies, coherences were held constant during a trial making it more important to just track in which direction motion is going and not for how long (Kelly and O'Connell, 2013). In this case, coherence strength is more important and higher coherence facilitates faster detection of the correct motion direction. In my continuous task the occurrence of single high coherence jumps alone are not enough to decide whether a participant is in a trial period, and the

regressor that I used here also does not capture any of the frequent changes in motion coherence that occur even during a trial.

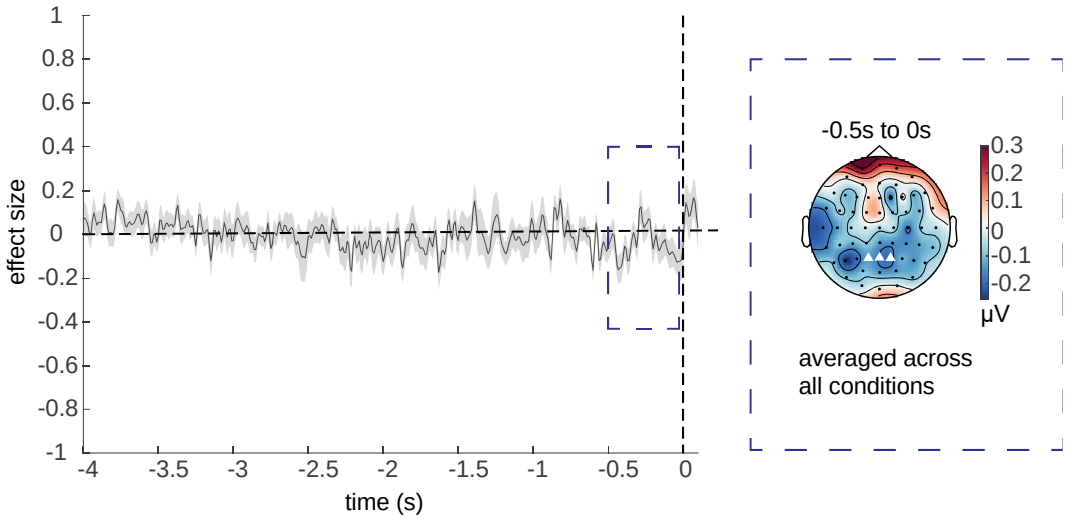


Figure 7: Regression ERP for coherence regressor for correct responses. Plotted is the average effect size of a regressor on the group-level with s.e.m. in the convolutional GLM for different mean coherences of trial periods that were followed by a correct response. Topography plot averaged across conditions for a time window of 0.5s before a response. White triangles depict electrodes plotted in A).

4.3.3 Convolutional GLM analysis of intertrial periods

My continuous paradigm enables me to study how the brain encodes continuous noisy stream of evidence to form a decision. Participants have to integrate over the entire stimulus presentation including intertrial periods. This means that I can study decision-processes in time windows in which the participant does not know whether the presented information supports any response and separate it from motor preparation signals. In addition, this method increases the SNR because the number of events I can analyse is much higher than in typical EEG experiments.

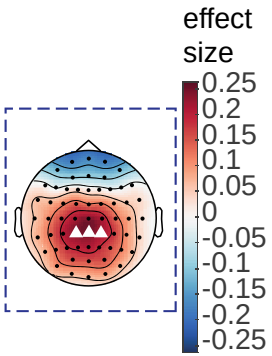
In my continuous paradigm I expected that the brain updates evidence for being in a trial period every time the coherence changes. The main effect of jumps during intertrial periods was similar across conditions with a positive peak at about 300ms after a jump occurred (a cluster-based permutation test did not identify significant differences, Figure 8). It was localized across central and centro-parietal areas. Its scalp topography looked similar to the P300 that has been claimed to track surprise

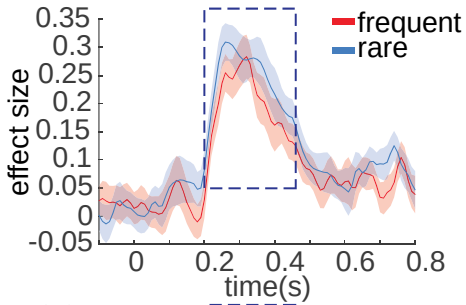
(Mars et al., 2008; Twomey et al., 2015) and has been related to the centro-parietal positivity signal measured in other evidence integration studies (Kelly and O'Connell, 2013; O'Connell et al., 2012; Twomey et al., 2015).

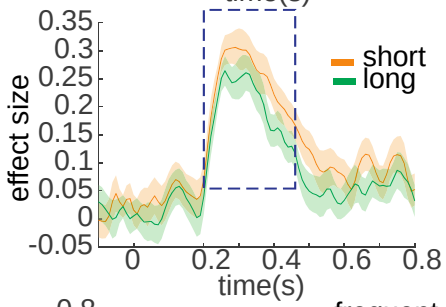
Trial frequency had a significant effect on the update of the prediction error measured at centro-parietal electrodes (Figure 8). Large differences in the coherence level of previous and current jump led to a significant increase in signal for rare trials (p = 0.006, 2-tailed cluster-based permutation test at a prior selected electrodes CPz, CP1, CP2 and a time window of 0s – 0.8s). The time course and topography were similar to the main effect of a jump in coherence but the prediction error led to overall higher effect sizes. This suggests that large differences in coherence between previous and current jumps are more surprising in rare trial conditions. In other words, participants have different expectations of jump sizes during rare and frequent trial conditions. They expect large jumps more often in frequent trials, because they overall spend more time in trials. In contrast, there was not a significant difference in updating of new evidence between short and long trials (Figure 8).

Both the absolute stimulus and jump level showed no difference between conditions (tested with a cluster-based permutation test) and a signal reflecting leaky accumulation could not be identified for the absolute stimulus regressor (Figure 8). Both the absolute stimulus and jump level regressors track absolute coherence levels but on different timescales. Despite this similarity, the scalp topography for both regressors was quite different. The jump level regressor showed a negative deflection at about 400ms after a jump with a central negative topography. In contrast, the absolute stimulus regressor showed 2 negative deflections at about 150ms and 400ms located over parietal areas and a positive peak in between at about 320ms over occipital areas. This suggests that the EEG signal related to the absolute stimulus regressor could reflect a mixture of visual and cognitive processes. The time course and scalp topography of these findings appear similar to results by Wyart et al. (2012), where a perceptual update in occipital areas was followed by a decision update in parietal regions (Chapter 1, Figure 3).

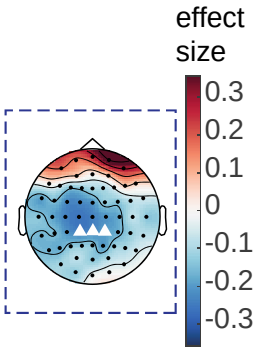
jump regressor

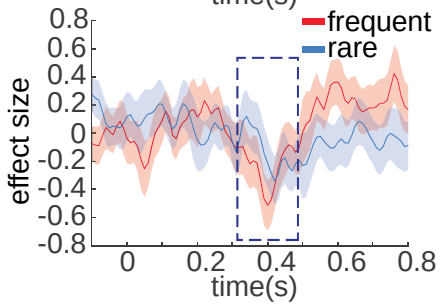


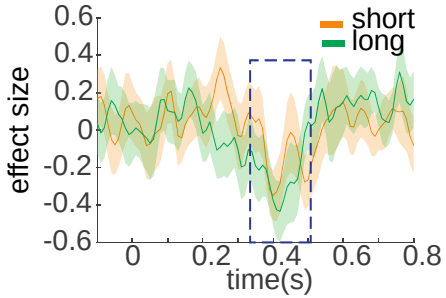




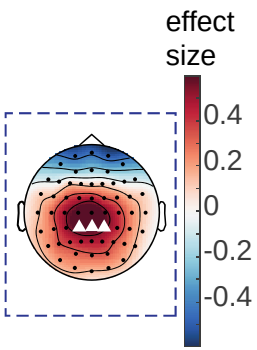
jump level regressor

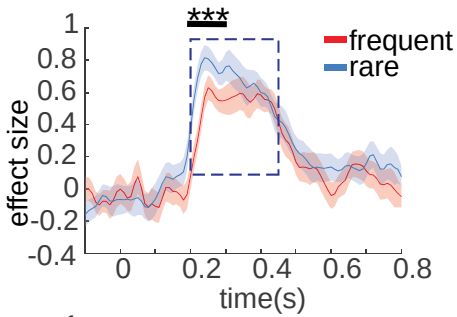


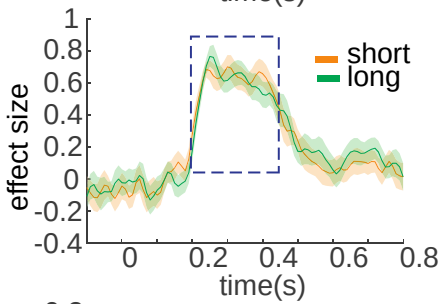




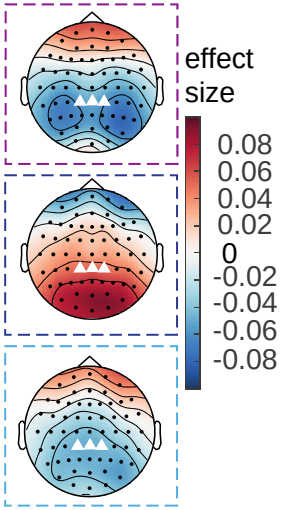
prediction error regressor

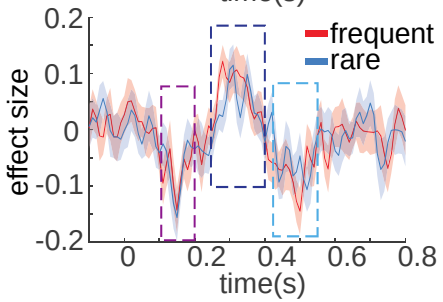






absolute stimulus regressor





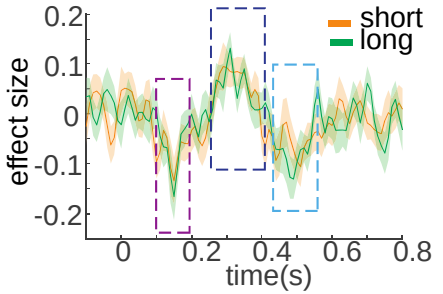


Figure 8 (previous page): Regression ERPs for convolutional GLM analysis of intertrial periods. Shown are the regression ERPs on group-level with s.e. for the difference for frequent versus rare and short versus long for different regressors defining jumps during intertrial periods and for the absolute stimulus. All ERPs were plotted for centro-parietal electrodes CP1, CP2 and CPz as indicated by the white triangles in the topography plots. Topography plots were averaged across conditions for a regressor and encompass the time window indicated by the blue square. A 2-tailed cluster-based permutation test for a priori selected region of centro-parietal electrodes and a time window of 0s to 0.8s revealed a significant difference between rare and frequent trials (p = 0.006) for the prediction error regressor

4.4 Control experiment

Based on the current task design it is not possible to determine whether the evoked responses I observed for changes in the stimulus, including prediction error and jump regressors, are due to low-level visual surprise (i.e., detecting changes in the stimulus) or are reflecting a process of evidence integration. Therefore, I conducted a small control study with 6 participants, in which stimuli contained both horizontal (decision-relevant) and vertical (decision-irrelevant) motion. The vertically moving dots behaved statistically exactly as the horizontally moving dots and also included the same trial periods but participants were instructed to only respond to trials in horizontal motion direction. Participants still had to judge whether they were in a trial period in which dots move either more to the right or more to the left. However, at the same time some dots moved vertically as a distractor. I implemented a convolutional GLM with jump, jump level, prediction error and absolute stimulus regressors for both the horizontal and vertical motion as well as regressors for correct responses, false alarms and trial starts which are not shown.

EEG signals influenced by jumps in coherences were not only reflecting the surprise of a big jump in coherences but actually tracked decision-relevant evidence (Figure 9). I compared the overall EEG signals between vertical and horizontal conditions. Specifically, the effects of jump and prediction error regressors were a lot smaller for irrelevant jumps in vertical motion. This means that the signal I measured represented only decision-relevant information, as has been shown before (O'Connell et al., 2012; Wyart et al., 2015). For both the jump level and absolute stimulus regressor, the differences between decision-relevant and irrelevant motion direction were a lot smaller. This control suggests that the

modulation of EEG signals measured for the coherence jump and prediction error regressors reflects cognitive processes in evidence integration as opposed to encoding the motion reflected by the influence of jump level and absolute stimulus regressors on the EEG signal.

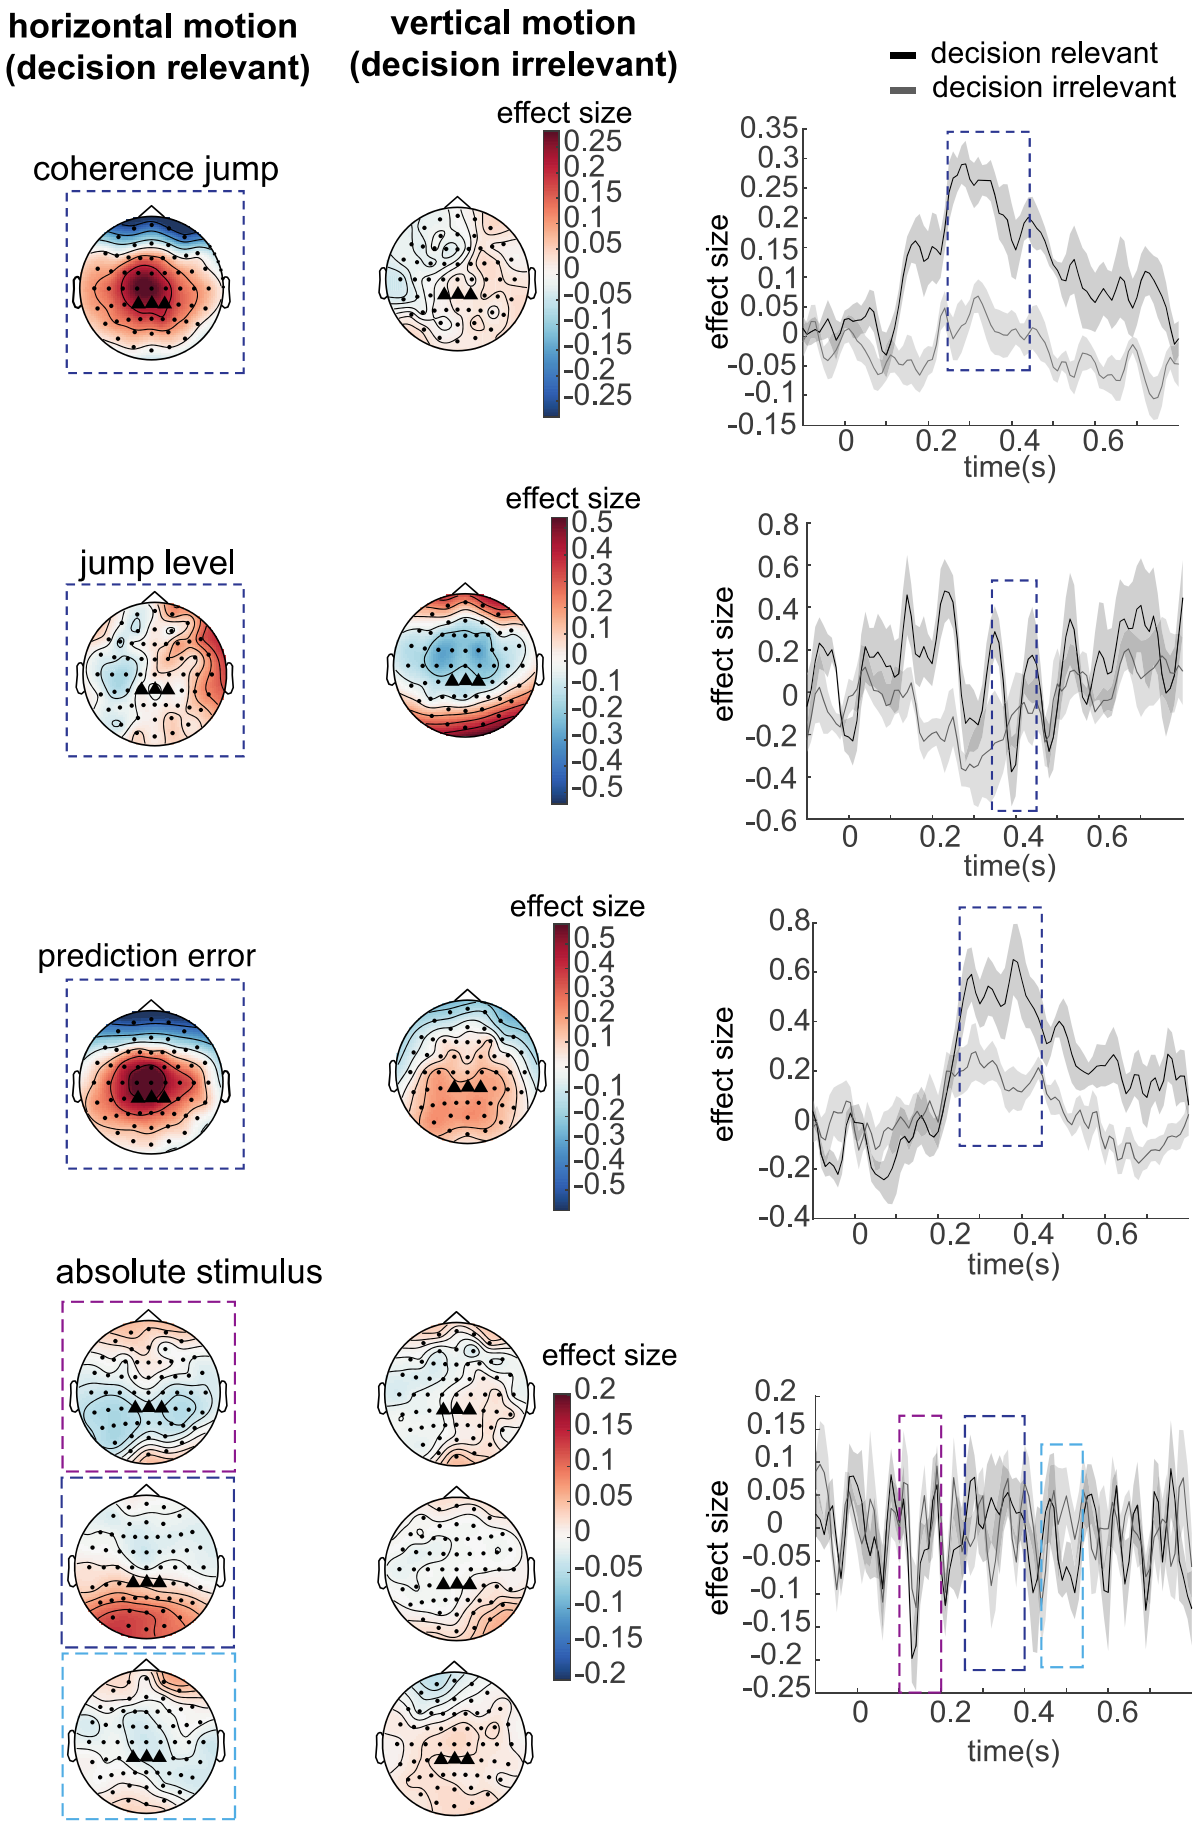

horizontal motion
(decision relevant)
vertical motion
(decision irrelevant)
decision relevant
decision irrelevant
coherence jump
effect size
0.25
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
effect size
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
0
0.2
0.4
0.6
time(s)
jump level
effect size
0.5
0.4
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
effect size
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
0
0.2
0.4
0.6
time(s)
prediction error
effect size
0.5
0.4
0.3
0.2
0.1
0
-0.1
-0.2
-0.3
-0.4
-0.5
effect size
0.8
0.6
0.4
0.2
0
-0.2
-0.4
0
0.2
0.4
0.6
time(s)
absolute stimulus
effect size
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
effect size
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
0
0.2
0.4
0.6
time(s)

Figure 9 (previous page): Comparison of regression ERPs for decision relevant and decision irrelevant information. In a control study, vertical motion was included as a distractor. Participants had to track the horizonal motion only. The left column shows topography plots averaged across conditions for horizontal (decision relevant) motion in comparison to topography plots in the 2nd column for vertical motion (decision irrelevant motion). Black triangles depict electrodes picked for the timeseries plots in the right column. The 3rd column shows regression ERPs on the group-level with s.e.m. for decision relevant (horizontal motion, dark grey) and decision irrelevant information (vertical motion, light grey). All regression ERPs were plotted for centro-parietal electrodes CP1, CP2, CPz).

4.5 Discussion

In this chapter I have shown that the brain adapts to different conditions of the task paradigm, which leads to changes in the CNV and the processing of new evidence samples across frequent and rare trial conditions. This is consistent with the broader conclusion that the brain adapts the processing of incoming sensory evidence in accordance with the statistics of the sensory environment (Murphy et al., 2020; Orsolic et al., 2019; Ossmy et al., 2013). The convolutional GLM approach has been shown to be a suitable method to analyse data of a continuous task with sequential events very close in time. This analysis revealed that most of the variance in the signal during a trial period was best explained by a CNV signal reflecting the temporal expectation or commitment to a trial, rather than trial onset. There was a difference in response preparation between long and short trial conditions (Figure 1), but most prominent was the difference between frequent and rare trial conditions (Figure 1,5): both rare and long trial conditions required longer response preparation. I only found weak evidence for an EEG signal build-up similar to the CPP time-locked to trial start (Figure 6); this is probably due to the overlap with the CNV signal which seems to mask the CPP which is visible when the data is CSD transformed (Kelly and O'Connell, 2013; O'Connell et al., 2012)). In addition, new samples of evidence were updated differently between frequent and rare trial conditions. Larger jumps in the noise are less expected in the rare trial condition, and so this would explain why they were associated with a larger update in the EEG signal (Figure 8). This signal was decision-related and not an artefact of other processes such as low-level statistics of the visual stimulus, as it showed a strongly reduced signal for distractor motion in the stimulus (Figure 9). Convolving the past stimulus stream with the EEG signal did not lead, as expected,

to a signal representing an integration kernel, but rather a waveform that might potentially reflect the bottom-up and top-down interactions between visual and parietal areas (Figure 8).

4.5.1 Neural activity reflects complex adaptation to different task environments

Updates in EEG signal following a jump in the noise fit with behaviour and support an approach in which the brain not only updates evidence but also its belief of which decision is correct. In rare trial conditions, higher prediction errors led to higher amplitudes of the EEG signal than for frequent trials. There are fewer trials in the rare condition, which means that large differences between consecutive jumps are less likely. This means that a large prediction error might count more towards a developing decision variable or accumulated evidence in the rare condition because it is more likely that there was a trial onset. More evidence is needed to form a decision, as shown in the behavioural integration kernels (Chapter 3, Figure 4). In order to reach the decision-threshold in the rare trial condition it might therefore be necessary to weigh 'outliers' e.g., big jumps in the noise more than in the frequent trial condition. Topographically and temporally the ERPs for a jump in the noise and the prediction error are similar to the P300, which has been linked to reflecting prior expectation of behaviourally relevant events (Mars et al., 2008). The differences in amplitude for the prediction error for different trial frequencies might also explain why the detection rate of rare trials is reduced (Chapter 3, Figure 1) and why RTs are prolonged for rare trial conditions (Chapter 3, Figure 3). This finding is important because it indicates that the brain can flexibly adapt the gain of each evidence sample based on the uncertainty in the environment. This is an extension to studies that have shown how the gain of new incoming evidence changes depending on how much evidence has already been accumulated (Cheadle et al., 2014; De Lange et al., 2010).

I hypothesise that participants try to solve the task as a change-point detection problem, rather than an evidence accumulation problem. Other studies that examine similar change-detection paradigms have shown that a leaky integration approach neither explains behaviour nor the brain signal (Murphy et al., 2020; Orsolic et al., 2019). Rather, participants combined evidence accumulation with

the update of a prior belief of what the correct decision might be and whether a change-point has occurred from an intertrial to a trial period. They showed that while parietal areas continued to represent accumulated evidence, motor related signals represented the actual belief update (Murphy et al., 2020). This appears consistent with the findings in the present study, in which signals over parietal cortex reflected the update of each new evidence sample, and a CNV signal located centrally possibly represented the commitment to a choice.

The CNV could represent an urgency signal that has a top-down effect on motor preparation. Multiple studies have shown that neural signals in motor-related areas represent expectation of an event (Macar and Vidal, 2004; Nobre et al., 2007; Pfeuty et al., 2005; van Rijn et al., 2011; Trillenberg et al., 2000). For example, an electrophysiology study in monkey LIP has shown that these neuros encode elapsed time and the probability of a behaviourally relevant cue occurring to which the monkey has to saccade to (Janssen and Shadlen, 2005). Moreover, another study has found a direct impact of the CNV and slow theta oscillatory power on the beta power that reflects motor preparation (Cravo et al., 2011). Therefore, the CNV might have an influence as an urgency signal on neural signals representing evidence accumulation in a decision-making task. A previous study found that urgency is important for the adaptation of evidence integration to different timescales (Shinn et al., 2019). Other studies have shown that this urgency might be encoded in motor related signals (Kelly et al., 2019; Murphy et al., 2016). This seems to be consistent with the present study because there were significant differences in the CNV between frequent and rare trials with an earlier onset of this signal for rare trial conditions. It may reflect an adaption of the urgency to commit to a response in which less evidence is needed for frequent trials. In that sense, the CNV signal for rare trial conditions indicates that participants built up a prior expectation over a longer period of time (Figure 1A,C,G,I), which is consistent with behavioural integration kernels discussed in Chapter 3.

In addition to the difference in the EEG signal for prediction errors between frequent and rare conditions, the absence of an exponential integration kernel for the stimulus regressor may also imply another strategy than the OU process. In simulations, I found that applying a deconvolutional GLM approach to the output

(the decision-variable) of the OU process does indeed retrieve the integration kernel (Supplementary Figure 2). Under the assumption that the CPP represents the decision variable, I would have therefore predicted that I can also retrieve an integration kernel by applying the same analysis approach to the EEG signal measured over (centro-) parietal cortex. Instead of an integration kernel, this analysis of the EEG signal showed a negative peak followed by a positive and ending with a negative peak (Figure 8). The corresponding topography plots showed that negative peaks were located over occipital areas and the positive more over parietal areas. This could represent the gating of sensory evidence before it reaches motor areas. The first negative peak might represent the pure sensory evidence in visual areas and looks like the topography of the steady-state visual evoked response found for the stimulus representation in O'Connell et al. (2012).This is followed by an update of the decision variable (Wyart et al., 2012, Chapter 1, Figure 3) which might then project back a prediction error to visual areas. Indeed, adopting a similar deconvolutional approach to EEG data acquired during randomly varying visual input produces long-lasting alpha-band oscillations (referred to as 'perceptual echoes' in the previous literature) that can be explained as a signature of predictive coding (Alamia and VanRullen, 2019). Thus, the neural signal might help to weigh each new piece of incoming evidence in upstream areas based on previous input.

4.5.2 Neural markers of decision-formation on the single-subject level

A convolutional GLM approach opens possibilities for new experiments in EEG that allow for analysis of neural markers on the single-subject level. The hallmark of EEG analysis has always been its averaging approach over EEG signals time-locked to certain events (Luck, 2014). The caveat of this method is that a lot of data is needed to sufficiently increase the SNR. In addition, events have to be sufficiently distant from each other to ensure that EEG signals for different events do not overlap. This would falsify the resulting ERPs because they would be partly due to the change in signal for other events that are nearby. The convolutional GLM successfully overcomes both of these challenges. With the GLM approach it possible to space events much closer in time without running into the problem of overlapping ERPs. This means that it is possible to reduce the recording time

because there is no need for intertrial or delay periods to space trials and events sufficiently apart. At the same time, this method allows for an increase in the SNR relative to the recording time and allows to investigate neural markers of decision formation on the single subject level at timescales of clinical relevance (Gonçalves et al., 2014; Lalor et al., 2006). Therefore, I intend to use this method in the future to analyse whether I can get robust single-subject measures of decision-making from this experiment, and correlate them with the findings in the behavioural section. Early analyses of this approach suggest that markers of the decision process may indeed be robust at the single-subject level (Supplementary Figure 3). Specifically, it would be interesting to correlate EEG signal parameters with the parameters of the subject-specific behavioural integration kernels measured for false alarms. This would allow me to draw conclusions about how robustly single-subject behavioural markers are related to processes in the brain. More importantly, it might be possible to identify new neural and behavioural markers on the single-subject level in patient populations with psychiatric diseases. Novel continuous task designs with a convolutional GLM approach could increase the use of EEG in patient populations because recording times can be reduced but at the same time allow for robust subject specific results.

4.5.3 Differences between findings in the present study and other studies of evidence integration

I did not find a prominent CPP, as reported by previous studies of continuous evidence accumulation, which might be due to the task design (although it should be noted that a CPP-like signal is visible when the CSD transformation is applied). When time-locked to trial start, I only found a small increase in signal, possibly reflecting evidence integration. However, given that evidence integration and motor preparation will overlap during a trial, this signal is difficult to interpret. Typically, I would expect to see a build-up in signal to a threshold at the time of response for both time-locking to trial start and response that defines a CPP (Kelly and O'Connell, 2013; O'Connell et al., 2012). In contrast, I found a decrease in the EEG signal over parietal areas, when the data is response-locked, that looks like a CNV signal. Another possible reason why I did not find a prominent CPP is because of the noise structure in my task. Participants might not clearly detect the start of the trial. They might think that the start is either earlier or later, meaning

that time-locking to trial start might not work. Though other studies have also used noisy intertrial periods, their trial onsets were more obvious because they were followed by a constant signal that was less perturbed by noise than in my paradigm (Kelly and O'Connell, 2013; O'Connell et al., 2012).

The advantage of the structure of the noise in the continuous paradigm together with the convolutional GLM is that it enabled the analysis of evidence integration signals in the absence of motor preparation signals. One possible way to separate different signal sources is a CSD transform, that has previously been shown to separate motor preparation and evidence integration signals in a continuous random dot motion task (Kelly and O'Connell, 2013). Applying a CSD-transform to the response-locked data of the present study resulted in a positive signal over parietal areas that possibly reflects integration of evidence reflecting the CPP, and a negative signal more central representing the CNV signal. A convolutional GLM applied to the EEG signal without a CSD transform, showed that there is only a minimal effect of trial start on the EEG signal, meaning that most of the variance is explained by the CNV signal. This indicates that the convolutional GLM cannot separate the evidence integration and CNV signals when the EEG data is not CSD transformed. In the future it might therefore be useful to repeat the convolutional GLM analysis on CSD transformed EEG data.

Unlike the manipulation of trial frequency, the manipulation of trial length was only partially effective. There was only weak evidence in the response-locked EEG signal. One reason for that are the selected trial lengths. Behavioural integration kernels for false alarms indicate that participants did not need 5 seconds of trial length but rather integrated only for the length of short trials of 3 seconds (see Chapter 3, Figure 4). This manipulation might have worked better if the trial lengths would have been shorter for both long and short trial conditions. An increase in difficulty of the mean trial coherence might have also forced participants to integrate for longer. In addition, plotting CSD transformed EEG data sorted for different RTs (O'Connell et al., 2012) might allow me to investigate which signals might sometimes lead to responses of participants longer into a trial than 3s.

4.5.4 Conclusion

During decision-making, brain activity can flexibly adapt to uncertainty in incoming evidence. In this chapter, we showed that an update of new evidence samples may reflect the expectation the brain has about new incoming evidence. Moreover, a dominant CNV signal might reflect urgency signals that track the overall belief of which choice to make. In addition, the complex activity pattern across visual and parietal cortex in response to the absolute stimulus indicates that there are complex interactions between these brain areas. Perhaps these determine how the expectation about new evidence is formed. In the future the high SNR that can be achieved with this continuous task design and the convolutional GLM might allow for the investigation of subject-specific markers of decision-making.

4.6 Methods

4.6.1 Task paradigms

The task paradigm and training of participants is described in detail in Chapter 2 Section 2.1 and Chapter 3 Section 3.2. Participants responded with button presses on an English keyboard. For rightward motion participants pressed ‘L’ with the right index finger and ‘A’ for leftward motion with their left index finger. In addition, to investigate the difference in processing of decision-relevant and irrelevant information, I conducted a small control study with 6 participants. In order to investigate whether the regression ERPs found for intertrial periods are merely due to surprise or actually representing processing of decision-relevant information, I analysed the EEG with horizontal and vertical motion. In this paradigm vertical motion is a distractor and represents irrelevant information. To implement this in practice, I split the population of dots in half, meaning 50% of dots were assigned as potentially having vertical motion, and 50% of dots were assigned as potentially having vertical motion. Mean coherence levels for trial periods were held constant, meaning that at 50% mean coherence on average all horizontal dots would move in one direction. Statistically both horizontally and vertically moving dots behaved the same and I included trial periods for the vertical motion as well. However, participants were only instructed to respond to the horizontal motion. Participants should only track horizontal motion to decide whether they are in a trial. Speed,

size and density of dots as well as the size of the aperture were not changed. The code for the vertical motion control task was implemented by Nebojsa Jovanovic as a summer project under my supervision.

4.6.2 EEG recording and pre-processing:

EEG data was collected at a sampling rate of 1000 Hz with Synamps amplifiers and Neuroscan data acquisition software (Compumedics) and 61 scalp electrodes following the 10-20 layout. Additionally, I placed bipolar electrodes below and above the right eye and on the temples to measure eyeblinks as well as horizontal and vertical eye movements (HEOG and VEOG channels). A ground electrode was attached to the left elbow bone. The EEG signal was referenced to the left mastoid but later re-referenced to the average of left and right mastoids. Impedances of electrodes were kept below 15 kΩ.

Data was pre-processed using spm12 (http://www.fil.ion.ucl.ac.uk/spm/), the FieldTrip toolbox for EEG/MEG-analysis (http://fieldtriptoolbox.org) and MATLAB (Version R2018b, The MathWorks, Inc., Natick, Massachusetts, United States). Each session for each participant was pre-processed as continuous data. First, each session was downsampled to 100Hz. Then, I referenced the data to the average of left and right mastoid electrodes and bandpass filtered the data between 0.1Hz and 30Hz. In a next step, I used customized scripts from post-doc Lilian Weber to remove eye movement artefacts. The method includes the signal space projection eye correction methods in SPM. The bipolarized VEOG channel was used to build a spatial confound topography of eye blinks that helps to delineate ocular source components. Subsequently, these components were regressed out of the EEG data. For this, segments of 1000ms around eye blink events in the VEOG channel were generated. Principal component analysis was used to define the noise subspace of eyeblinks across all channels. The first component was subsequently regressed out (Nolte and Hämäläinen, 2001). For each participant and session the spatial confound map of the first component was visually checked to ensure it showed a typical eye blink topography before the regression was applied. Further, I thresholded the EEG data to remove artefacts that were >= 100µV in a single channel by labelling a 500ms window around the

peak of the artefact. In a last step, I used current source density (CSD) to transform the data in FieldTrip.

4.6.3 Analysis

Conventional EEG analysis:

I used FieldTrip to take epochs of the data for a conventional event-related potential analysis. Labelled artefacts from the pre-processing step were replaced with nan. For each session and subject, I time-locked the data to the start of correctly identified trials and response-locked the data to correct response and false alarms. Then, I averaged across sessions and subjects. For topography plots I averaged across all conditions and, for timeseries analysis, I averaged the data separated for different trial frequencies (frequent and rare) or trial lengths (short or long). Timeseries analysis was usually reduced to the average of a selected group of 3 electrodes that is always indicated in the topography plots. I did not apply baseline correction because events could easily overlap in time which makes it difficult to find a baseline period that I could use. Difference waves were calculated for each subject separately for rare vs frequent and short versus long trial conditions and then averaged across subjects.

Cluster-based permutation test for conventional ERP analysis

I implemented a cluster-based permutation test within FieldTrip to test for differences between frequent and rare trial conditions as well as long and short trial conditions. For the main analysis, permutation tests were applied to a-priori selected time and region of interest in the data. For response locked analyses I selected centro-parietal electrodes (CP1,CP2,CPz) because I expected to measure a similar centro-parietal positivity signal as reported by O'Connell et al. (2012). I selected a time window of -2s to 0s for the permutation test. I did not increase the time window beyond the response as the spiky EEG signal in that time period still might be contaminated by eye movements. For the analysis of epochs time-locked to trial start, I selected parietal electrodes (P1,P2,Pz) as the topography clearly revealed a positive peak slightly posterior to centroparietal electrodes. Central parietal electrodes seemed to be right at the border between a negative frontal and positive posterior signal. A time window from 0s to 6.5s was

selected for the permutation test. 1000 permutations were generated with the Monte Carlo method and clusters were selected based on an alpha of 0.05 for uncorrected and corrected clusters for a two-tailed test.

Convolutional GLM analysis

During the EEG recordings I sent triggers between the computer presenting the stimulus and the EEG computer. Given the high density of triggers not all triggers get recorded in the EEG file. For this reason, I matched the EEG triggers and a copy of all triggers sent on the stimulus computer for each block of a session separately. First, I made sure that all the triggers in the EEG file were represented in the copy of all triggers sent. If that was the case, I slid both trigger streams against each other to find the optimal position in which most of the triggers between EEG and stimulus match. This was necessary as there was sometimes specifically towards the end of a 5min block a small delay of a couple of milliseconds between EEG triggers and the triggers saved with the stimulus computer. In some subjects there were 1-2 blocks (of 24 blocks in total, 4 blocks per session, see Chapter 2.1) that cannot be matched so far as there are some difficulties with incorrect triggers sent to the EEG computer. For a couple of subjects 5 blocks in different sessions had to be excluded but those were outliers and the convolutional GLM analysis still led to usable results for those subjects. There were 2 subjects in total in which there were problems with the connection between the EEG and stimulus computers in which the EEG data was not usable at all.

Permutation test for convolutional GLM analysis

I implemented a cluster-based permutation to test whether there is a difference between frequent versus rare and short versus long trial conditions. This was done the same way as described above for the conventional ERP analysis. However, instead of using the actual averaged EEG traces, I used coefficients from the convolutional GLM. For the first convolutional GLM analysis that verified the results from the conventional ERP analysis, I used the same a-priori selected region of interests and time windows. In contrast, for the convolutional GLM analysis of intertrial periods, I always restrained the cluster-based permutation test to centro-parietal electrodes (CP1,CP2,CPz) and a time window from 0s to 0.8s.

4.6.4 Participants

I tested 32 participants, 12 male for the main study. Of those 3 were not able to learn the task. One was at chance level at the training of single trial random dot task. The other two were not able to identify trial periods in the continuous task version. One person fell asleep repeatedly during the EEG session. Another 5 were excluded from the EEG analysis because there were issues with matching the continuous EEG with the stimulus stream or it was not yet possible to remove eyeblink confounds due to bad channels. This means 23 subjects were included in the analysis. For the control study 6 participants were tested, 4 male. All participants were below the age of 40, had normal or corrected-to-normal vision and gave written consent prior to taking part in the study. The study was approved by the local ethics committee (CUREC R60298).

4.7 Supplementary Figures

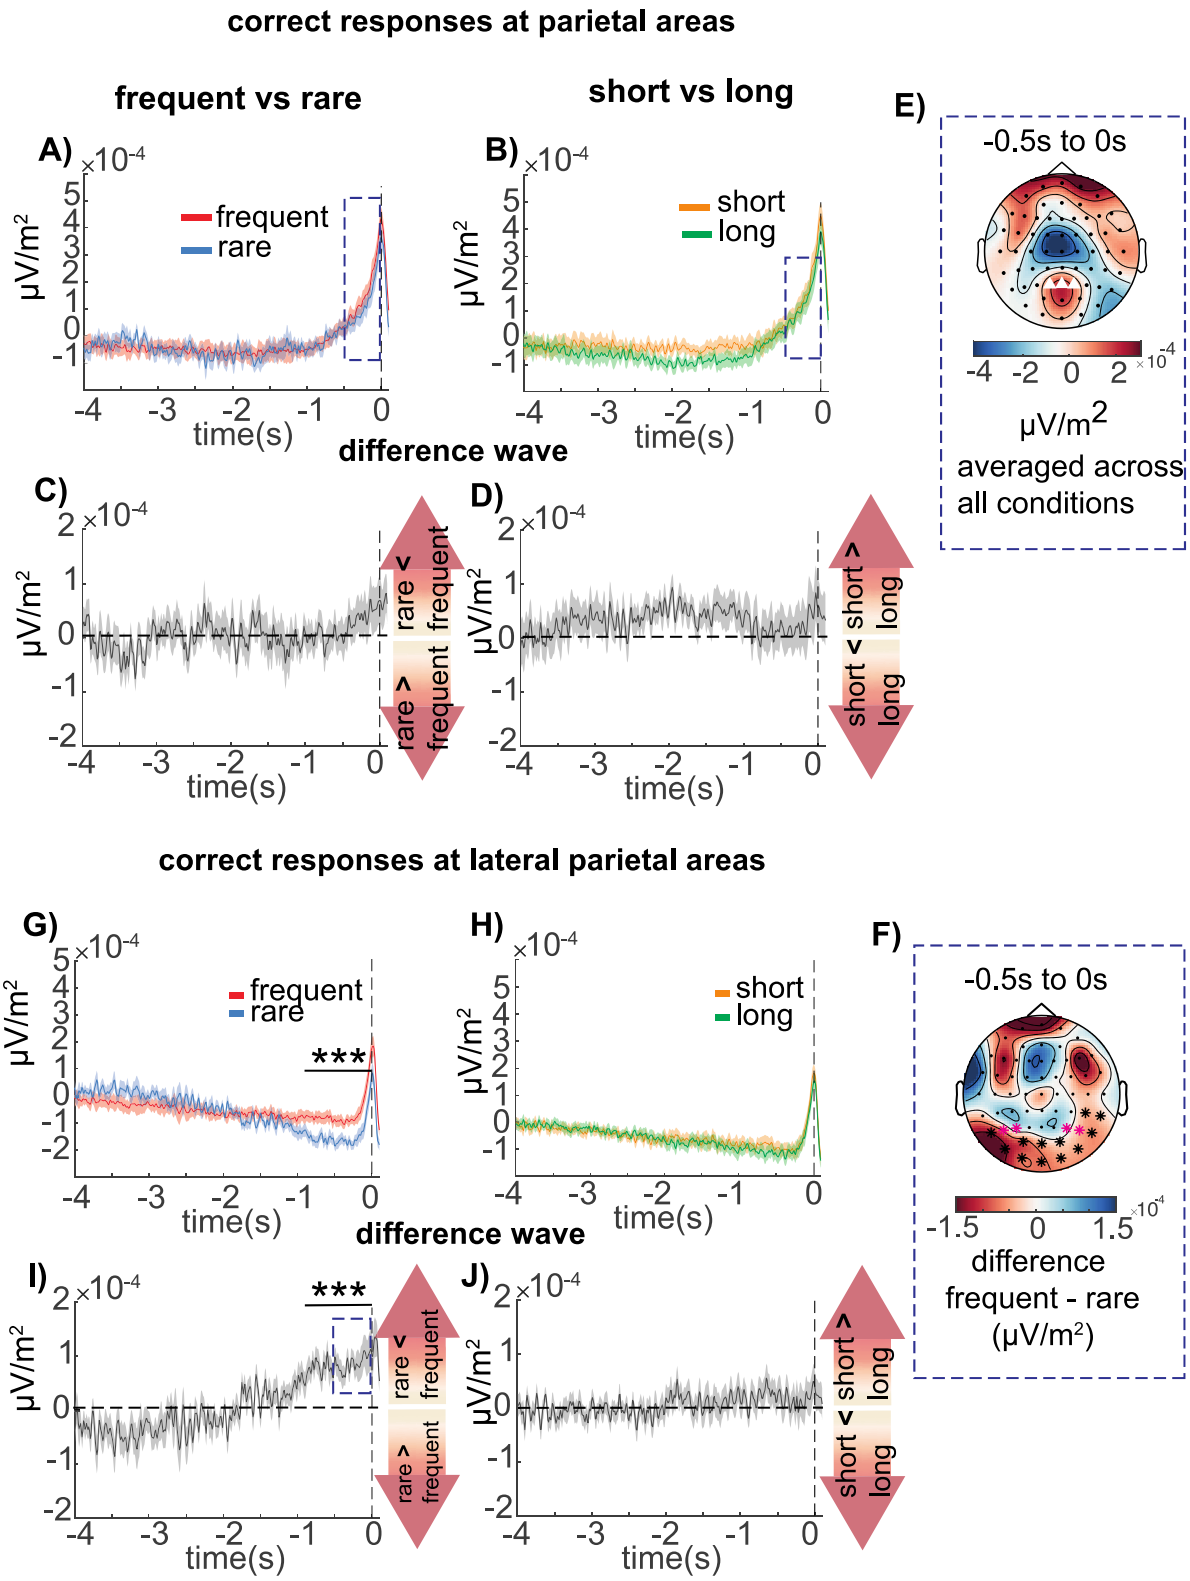


Figure 1 (previous page): CSD transformed ERPs for correct responses. A) CSD transformed ERPs (with standard error of the difference between frequent and rare) for frequent and rare trial conditions do not show a difference between conditions at parietal electrodes (P1,P2,Pz). **B)** same as in A) but for short and long trial conditions. **C)** and **D)** difference waves calculated for each participant and then averaged across participants with standard error of the difference. **E**) Topography plot, averaged across all conditions, shows two separate sources in frontal and parietal areas. White triangles indicate the electrodes plotted in A) - D). **F)** is a topography plot showing the difference between frequent and rare trials across the brain. Asteriks indicate significant differences with H1 hypothesis that there are differences between rare and frequent trial conditions in the brain at the a priori selected time window of -2 to 0s. Pink asterisks indicate the electrodes plotted in **G)** and **H)**. **G)** shows ERPs for rare and frequent trial conditions at P3, P4, P5, and P6 with standard errors of the difference and a significant difference at these a priori defined electrodes and time window of -2 to 0s (p = 0.0001, alpha = 0.05). **H)** the same as in G) but for short and long trial conditions. **I)** and **J)** are the difference waves averaged across subjects with the standard error of the difference.

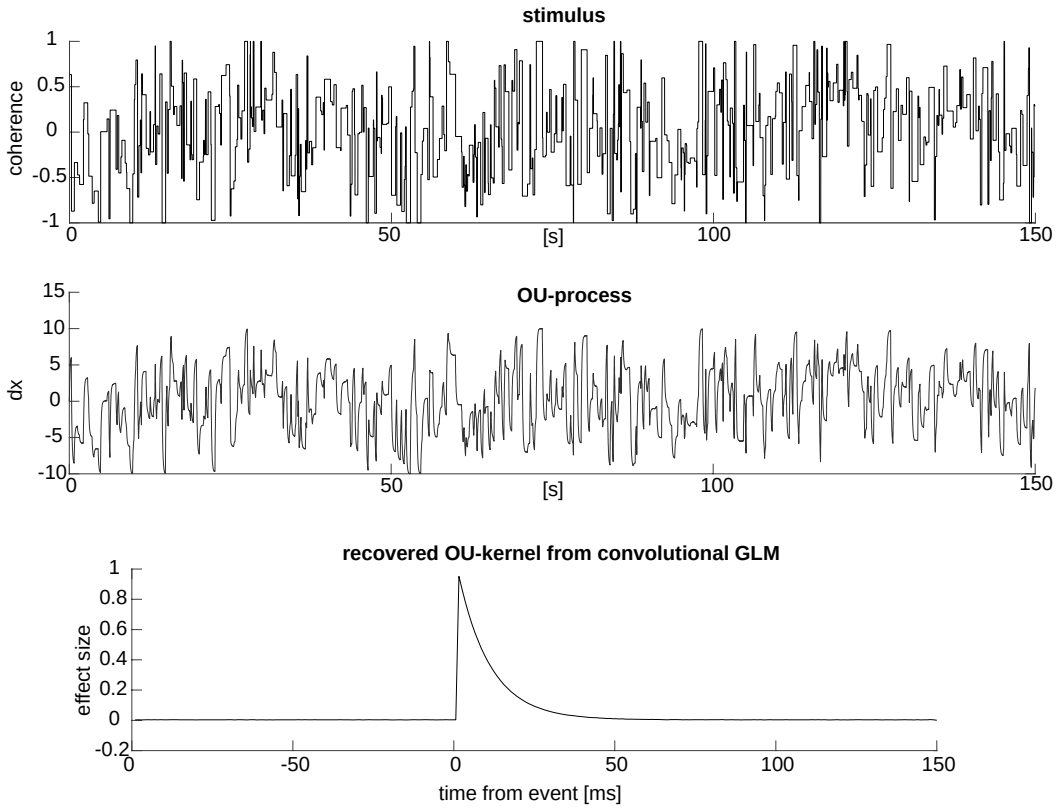


Figure 2: Recovery of the OU-kernel with the convolutional GLM. The first row shows a cut-out of a stimulus stream. In the second row the decision variable (dx) computed with an OU process with λ = -0.1 and minimal noise (c = 0.01) is depicted. The third row shows the GLM coefficients for a convolutional GLM with the stimulus regressed onto the decision variable.

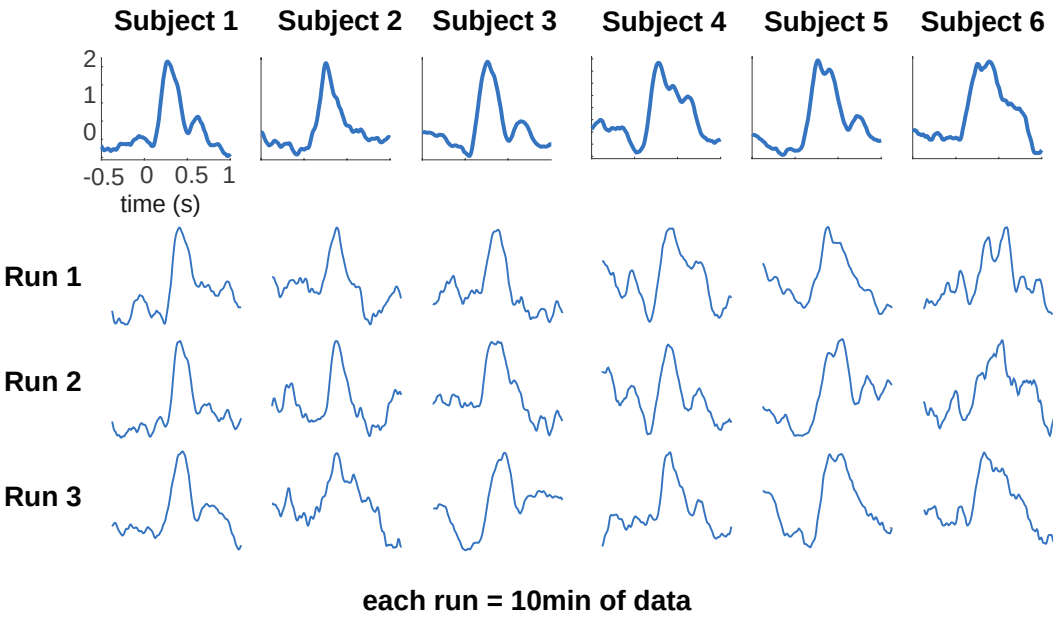


Figure 3: Robust single-subject neural signatures. Shown are 3 pilot subjects who did the continuous RDM task while I recorded EEG. The first row shows the prediction error across all 6 sessions for a subject. Each run shows the prediction error calculated over 10min of data. In the future, I would like to repeat this analysis on the subjects that are included in the present study.

5 Influence of prior expectation on behaviour and neural activity in parietal cortex measured with MEG

5.1 Introduction

The past chapters investigated how humans integrate evidence over the timescale of hundreds of milliseconds to detect a stimulus amongst noise; one aspect of this question was how the temporal extent of the evidence integration kernel may adapt to the temporal features of the stimulus stream.

This chapter investigates evidence accumulation on a longer timescale, spanning tens of seconds, and importantly, spanning multiple decisions and actions. Rather than looking at how coherent motion is accumulated over time from a single moving dot stimulus, this chapter investigates how people build up an expectation over multiple such stimuli about the likely characteristics of future stimuli.

In this chapter I report behavioural and MEG data from a project done jointly with Tom Marshall; although Tom had already started this project when I got involved, I collected some of the data (some with Tom and some independently) and the analyses presented here are my own work unless explicitly labelled otherwise.

Background

The underlying circuits of evidence accumulation in a perceptual decision-making task are not yet fully understood. Neural traces of evidence accumulation are found across cortex. The frontal eye fields (FEF) and lateral intraparietal area (LIP) show similar activity and are co-activated across humans and non-human primates (Ding and Gold, 2012; Gold and Shadlen, 2007; Kim and Shadlen, 1999; Nobre et al., 2000; Roitman and Shadlen, 2002; Siegel et al., 2015; Ungerleider, 2000). In both FEF and LIP, neurons are more active if the next saccade is to a location in the receptive field of a neuron. This means that during a trial of a random dot motion task the activity of neurons with their receptive field at the target location increases. However, this poses the question what the functional difference between FEF and LIP might be.

It has been suggested that FEF might be more important for integrating current evidence for different upcoming saccadic targets, while parietal cortex integrates over previous saccades and their behavioural relevance. Studies in both monkeys and rodents have found that LIP, whose homologue is posterior parietal cortex (PPC) in rodents (Erlich et al., 2015; Hanks et al., 2015; Katz et al., 2016), although correlated with evidence accumulation, is not causally necessary for the performance in a random dot motion task on a trial-by-trial basis, in contrast to FEF, whose homologue is frontal orienting field (FOF) in rodents (Erlich et al., 2015). What is the purpose of evidence-correlated activity in LIP?

In the natural world we do not only rely on the current evidence that is presented to us but we also base our behaviour on past experience. Past experiences help determine which features in the visual field become important for future behaviour. Monkey area LIP has been shown to represent the salience of objects that are important for future behaviour in the visual field by integrating visual, motor and cognitive signals (Colby et al., 1996; Gottlieb et al., 1998; Ipata et al., 2006, 2009; Kusunoki et al., 2000). Hence, it is possible that LIP also integrates past experiences to build a priority map of visual features that will be important for future behaviour.

In typical perceptual decision tasks, each trial is independent of the others, meaning that only the evidence presented with the current trial is needed to make a decision. In contrast, in a memory task PPC seems to be crucial for integrating sensory history to solve the task (Akrami et al., 2018).

In the original project studied in this chapter, Tom Marshall set out to investigate the distinct functional roles of homologues of FEF and LIP in human cortex using MEG. He designed a novel random dot motion task based on two hypotheses:

1) FEF reflects evidence accumulation *within* a trial. Competition between different neural populations for the saccadic targets is reflected in the typical ramping activity observed during random-dot-motion tasks. A winner-takes-all behaviour of these networks defines which target is chosen.

2) Parietal cortex integrates trial outcome *across* trials to build a graded prior belief over saccadic targets that represents the likelihood of a target being correct.

To facilitate testing of both of these hypotheses, two changes were applied to the typical random dot motion tasks. Firstly, to directly modulate the competition in FEF independently of the coherence level or signal-to-noise, the stimulus included two competing sets of randomly moving dots (both leftwards and rightwards dots). Of each set, some dots would move randomly while the others would move to the left and to the right depending on the set. Participants had to make a judgement on each trial in which direction most dots were moving.

Secondly, to investigate the integration of outcomes across multiple trials, the task was modified so that the underlying prior probability of rightwards (vs leftwards) motion being dominant was not 50% (as in typical random dots tasks), but either 80%, 50% or 20%. These prior probabilities changed approximately every 25 trials without explicit cueing. As participants were not informed about the direction and timepoint of changes in the prior probability of the correct target they had to build a prior belief based on the past veridical feedback.

My analysis focusses on the timepoint at which feedback is received and the prior expectation should be updated for subsequent trials. I analysed behavioural evidence for the use of the prior, and analysed MEG signals from central parietal cortex for evidence of feedback- and update-related processing. Data concerning neural signals prior to and during evidence accumulation in a trial were analysed separately by Tom Marshall (Marshall, Ruesseler, Hunt and O'Reilly, *submitted*).

Behavioural results show that participants based their choice on a combination of veridical feedback and accumulated evidence. For the neural signals I found only a very weak effect of the prior. Only whether the choice is correct or incorrect is tracked by parietal cortex at the time of feedback.

5.2 Study design

5.2.1 Task

This task is based on a random dot motion paradigm (Figure 1). Participants had to integrate randomly and coherently moving dots to decide on each trial whether more dots move to the right or to the left by making a congruent saccade to a target. The aperture of moving dots spanned 4.2 visual degrees and was placed centrally on the screen with a fixation point in its centre. 100 moving dots were displayed in the aperture. Circular saccade targets with a diameter of 4 visual degrees were placed 0.8 visual degrees below the midline and 6.8 visual degrees next to the aperture of moving dots. After each trial participants received feedback in form of a semicircle at the location of the fix dot pointing always to the correct target. The task was coded by Tom Marshall using Psychtoolbox (Brainard, 1997).

Each trial started with a preparation period of 1s. Then the aperture appeared, and all dots started moving randomly. After 1s some of the dots started to move coherently for 2.5s. This means that participants saw moving dots for 3.5s in total. Once the dots disappeared participants were instructed to make a saccade to one of the targets within 1s. Subsequently the feedback cue was shown at fixation for 1s (Figure 1F).

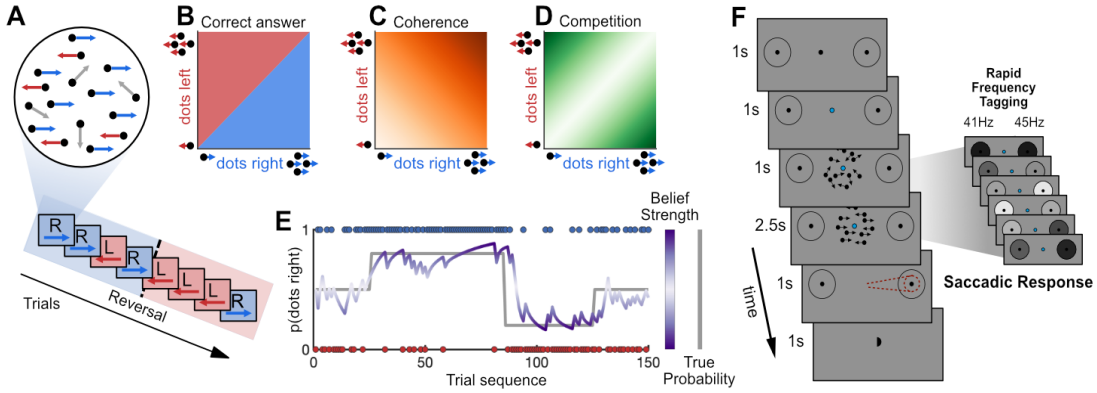


Figure 1: Task description. **A)** Participants had to judge whether more dots move to the left or to the right by making a congruent saccade. Within each trial a fraction of dots would move coherently to the left and right, while the rest moves randomly. **B)** In each trial participants had to judge the dominant motion direction. **C)** The fraction of dots moving coherently to the left and right is termed coherence and can be visualised in a 2D map. **D)** Competition between left and rightward moving dots. If dots move equally to the left and right (white area in the middle) competition is high**. E)** Across trials the outcome probability would vary between equal outcome probability (50% of trials are to the right) and a probability either shifted to the left target (20% right) or right target (80% right). About every 25 trials there is a reversal in outcome probability (grey line). Based on a Bayesian model participants should develop their belief state (purple line) by observing the veridical feedback displayed at the end of each trial. The belief state, also called prior, becomes more stable the longer a participant encounters the same outcome probability and is more unstable after a reversal. **F)** shows the structure of a single trial. The trial starts with a prepatority cue indicating that the random dot motion is about to start. For the 1st second dot motion is completely random which is succeeded by a 2.5s of coherent motion. Once the dots disappeared participants were instructed to saccade to the target in the direction of the domination motion direction. At the end of the trial participants received feedback about the correct target (reproduced from Marshall et al, *submitted*).

There are 2 modifications that were implemented which make this task unique in comparison to other random dot motion tasks. First, the outcome probability (prior probability that the dominant motion direction was left or right on any given trial) was not 50%, as in most RDM tasks, but instead varied so that there were blocks of trials in which the probability of dots moving rightwards was 20, 50 or 80% (Figure 1E). This prior probability changed periodically with a hazard rate of 0.04, so the outcome probability switched approximately every 25 trials in a session. This means that participants had to integrate over past trial outcomes in order to build a prior distribution of which target is most likely correct on the current trial. As stated in the introduction, I hypothesize that the build-up of a prior across trials is integrated over parietal cortex.

Second, the random dot motion stimulus was defined not only by the coherence, i.e., the total fraction of dots moving coherently but also by competition. The overall coherence (proportion of dots moving non-randomly) took values of 10%, 30%, 50%, 70% and 90% (Figure 1C). The competition is defined as the fraction of all coherently moving dots moving to the right, which could take values of 10%, 30%, 50%, 70% or 90% of the dots (Figure 1D). The rest of the coherently moving dots would move to the left. At 50% competition equal number of dots of all coherently moving dots would move to the left and right, while at competition below 50% more dots move to the left than to the right and vice versa for values above 50%. A signal reflecting the competition in integrating evidence for the two targets should be located in the FEF. This has been investigated by Marshall et al (*submitted*).

Intuitively, participants should rely more on evidence from the current trial than on a prior if coherence is high, i.e., most dots move coherently and if competition is low, i.e., most of the dots move in one direction. In contrast, when trials are difficult, meaning low coherence and a lot of competition, that is similar numbers of dots move to the left and right, participants should rely more on the prior. Given that the feedback depends on the evidence observed in a trial, a Bayesian model that depends on the veridical feedback should be able to solve the task optimally.

5.2.2 Model details

The task behaviour is modelled using a Bayesian model that has been designed by my supervisor Jill O'Reilly:

On each trial, the prior probability that the dominant motion direction would be 'right' followed a Bernoulli distribution with parameter q_t, i.e., the prior probability that the correct answer would be 'right' on trial t was q_t.

The prior distribution over q_t was initiated as a uniform on the range (0,1) on the first trial, and thereafter obtained from the posterior over q_{t-1}, based on the outcomes $x_{1:t-1}$, combined with a uniform 'leak'; the posterior and 'leak' distributions were weighted by a factor H representing the true hazard rate:

$$p(Q_t = q|X_{1:t-1}) = (1 - H) \cdot p(Q_{t-1} = q|X_{1:t-1}) + H \cdot U(0,1)$$

Eq. 1

Where:

$$H = p(q_t \neq q_{t-1}) = {}^{1}/_{25}$$

Eq. 2

i.e., it was assumed that participants know approximately the true value of H following extensive pre-training.

The posterior $p(Q_t = q|x_{1:t})$ was obtained iteratively from the prior $p(Q_t = q|x_{1:t-1})$ and the likelihood:

$$p(Q_t = q|X_t) = p(X_t|x{\sim}B(q))$$

Eq. 3

… using Bayes' theorem:

$$p(Q_t = q|X_{1:t}) \propto p(X_t|x{\sim}B(q)) \cdot p(Q_t = q|X_{1:t-1})$$

Eq. 4

where the posterior over q_t was normalized to integrate to 1.
Where a scalar value for the 'prior' is used in data analysis, this is the expected value of q_t based on the prior distribution

$$E(q_t) = \int p(Q_t = q|X_{1:t-1}) \cdot q \, dq$$

Eq. 5

Subsequently, this model has been fitted to the behavioural data based on the maximum log-likelihood (Marshall et al, *submitted*). For my analysis I am using the fitted prior (*p (Qt = q)*) to investigate its influence on behaviour and brain signal. The prior (Figure 1, E purple line) changes throughout the session based on whether *p(dots right)* is tilted to the right, left or not in favour of any target (Figure

1, E grey line). This is reflected in the belief strength in Figure 1, E. If the prior is close to 1 then the models' prior belief favours strongly a rightward choice. In contrast, if the prior is close to 0, then it is more likely that a left choice is correct. At a prior of 0.5 the belief is not in favour for any choice. The belief strength follows the changes in *p(right)* with some delay, because in the Bayesian model the update of the prior depends on several previous trial outcomes as defined by the hazard rate.

5.2.3 Training and testing regimen

34 participants were trained on 300 trials on a separate day before they were tested in the MEG. Due to low accuracy 5 participants were excluded after training. A total of 29 participants were tested in the MEG. In the MEG participants completed 4 sessions of 150 trials each. Each session lasted about 20min. Another 4 participants were excluded after the MEG acquisition due to low performance during testing and too many saccades during fixation periods. This means that the data of 25 participants was analysed in this thesis. Approval was granted through a local ethics committee: CUREC R53476. All participants gave written consent.

5.2.4 Analysis of saccade direction

Saccade direction was extracted from eye tracker data using custom MATLAB scripts by Tom Marshall. For one participant this was not possible as there was a technical problem with the eye tracker. In that case the saccades were reconstructed from the horizontal EOG.

5.3 Behavioural findings

5.3.1 Feedback on previous trials affects choice

Methods. To determine whether the outcomes of previous trials, and hence the prior, influenced perceptual decision-making, I analysed the influence of past feedback and evidence on choice in 25 participants using logistic regression. In the first regression I examined how feedback over the past 3 trials, coherence and competition on the current trial and the interaction of competition and coherence (called 'evidence') predict choice. Except for the coherence regressor all others are signed with positive being in favour of right and negative for being in favour of a left response. The dependent variable is choice on the current trial – a binary regressor with '1' being a rightward choice.

In this regression I included a regressor for the evidence on the previous trial (t_{-1}) as well. Although participants could learn the prior probability of a leftwards- or rightwards trial from the feedback alone, I was curious as to whether a previous trial in which strong sensory evidence for rightwards motion (say) had been presented would lead to a stronger prior for rightwards motion on the subsequent trial. A constant regressor was also included, which captures any average bias subjects have towards choosing the left or right option. Reported are the *t*-statistics on the group-level. All analysis was carried out in MATLAB (Version R2018b, The MathWorks, Inc., Natick, Massachusetts, United States).

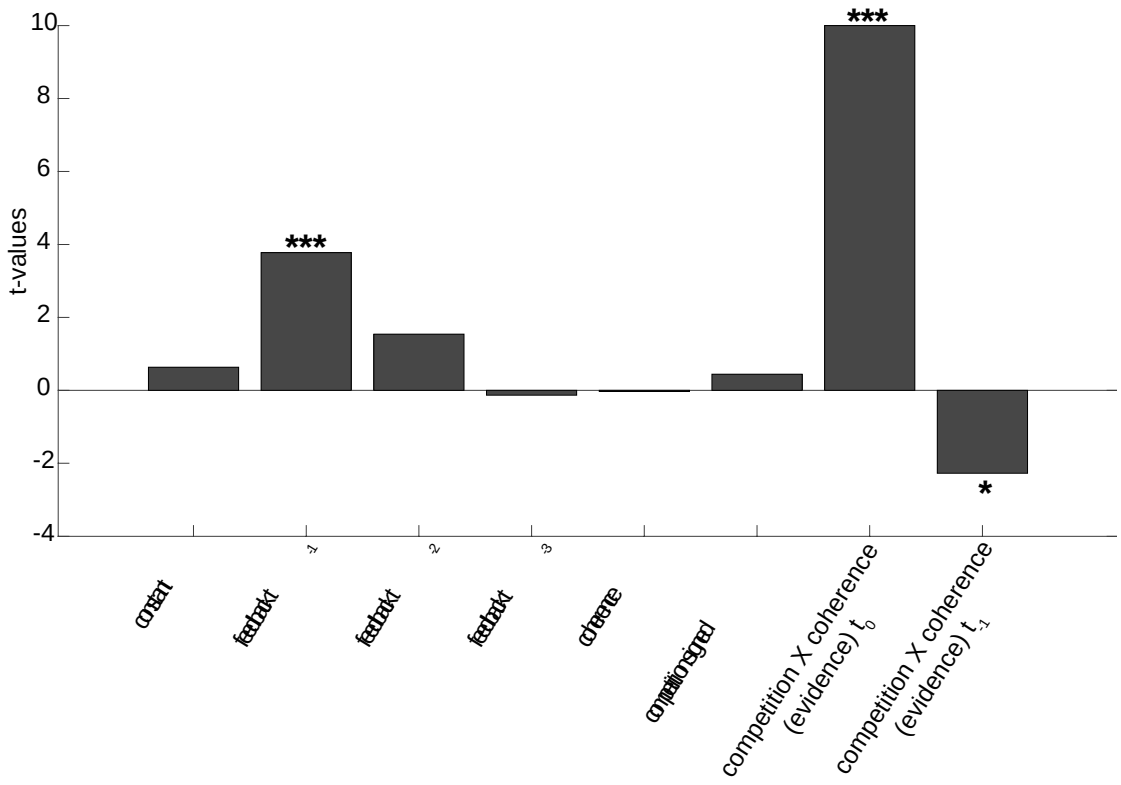


Figure 2: Logistic regression of feedback and evidence on choice. Represented are the t-values for a one-sided group test for each regressor. The first 3 regressors were for the influence of feedback of the past 3 trials (t) on current choice. Feedback was signed, positive for a rightward choice. Only the feedback on the last trial influenced current choice significantly (one-sided t-test, *t(24)* = 3.78 *p* = 0.0009). Coherence (not signed) and competition (signed positive for rightward choice) on current trial did not influence choice. Only the interaction between these two regressors, the evidence, had a significant impact (*t(24)* = 10, *p* < 1e-9). In addition, the evidence presented on the last trial also had a significant influence on current choice (final regressor on the right, *t(24)* = -2.27, *p* = 0.03).

Results. Past feedback impacted choice behaviour on the current trial (Figure 2). A logistic regression shows that participants specifically used the feedback from the last trial to guide their current choice (one-sample t-test of beta coefficients against zero, feedback t_{-1}: t(24) = 3.78, *p* = 0.0009). For the penultimate feedback and feedback further back in time the influence decreased steadily (one-sample t-test of beta coefficients against zero, feedback t_{-2}*: t(24)* = 1.54, *p* = 0.14; feedback t_{-3}: *t(24)* = -0.9, *p* = 0.13). This was expected because the outcome probabilities vary throughout a session. Therefore, participants should not have based their decision on all previous trials in a session.

Interestingly, the effect of evidence strength on the previous trial (t_{-1}) on the prior was negative – in other words strong evidence for rightwards motion on trial t_{-1} made participants less likely to choose 'right' on trial t (one-sample t-test of beta coefficients against zero, $t(24) = -2.27$, $p = 0.03$). This is possibly the result of collinearity between evidence and feedback on trial t_{-1}.

As expected participants based their decision mostly on the evidence in the current trial (one-sample t-test of beta coefficients against zero, $t(24) = 10$, $p < 1e^{-9}$). The effect of past feedback and evidence should have guided participants behaviour particularly when evidence is weak. There was no significant effect of coherence (one-sample t-test of beta coefficients against zero, $t(24) = -0.04$, $p = 0.97$), competition (one-sample t-test of beta coefficients against zero, $t(24) = 0.44$, $p = 0.66$) or the constant (one-sample t-test of beta coefficients against zero, $t(24) = 0.63$, $p = 0.63$).

5.3.2 Prior expectation interacts with current evidence to predict choice

Methods. Having established that feedback on at least one previous trial affects choice (Figure 2), I went on to use the prior *p(right)* from the Bayesian model to summarize the collective effect of recent trial outcomes in a single variable. I could then investigate the interaction of prior belief with evidence on the current trial. This regression included a regressor for a prior, the interaction between prior and coherence, prior and competition, as well as evidence. All regressors were signed positive if rightward leaning, except for the coherence which is not signed. This regression included a constant term, as well. For all regressors I am reporting the t-values for a one-sided t-test on the group-level.

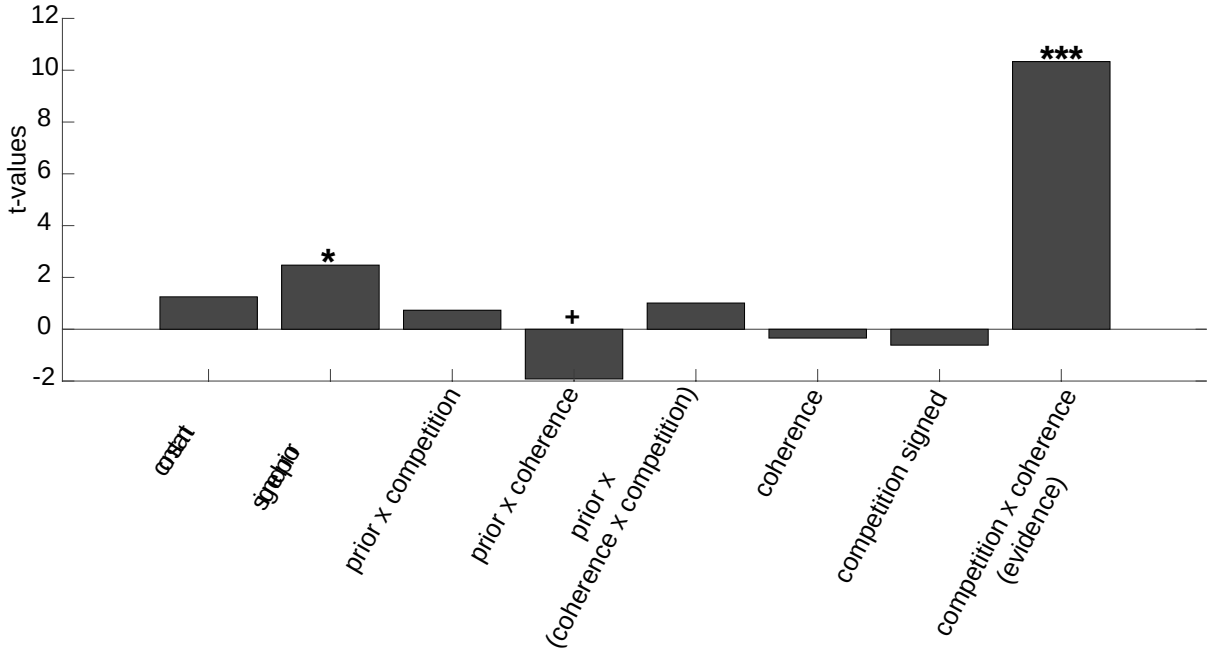


Figure 3: Logistic regression of prior belief on choice. T-values are shown for a one-sided group test. The prior for each trial was calculated by fitting a Bayesian model to each participants' behavioural data and had a significant influence on current choice ($t(24$= 2.47, p=0.02). In addition, the impact of the interaction between prior and coherence, as well as prior and competition and prior and evidence are depicted. Each regressor was signed positive for a rightward choice. Only the interaction between prior and coherence showed a trend in influencing current choice ($t(24)$ = -1.92, p = 0.07). The other regressors were the same as in Figure 2. As for Figure 2 there was a significant influence of evidence on current choice ($t(24)$ = 10.33, p = $1e^{-9}$). Each regressor except for the constant and coherence regressor were signed (positive = rightward choice).

Results. I found that both the prior (now summarized using *p(right)* from the Bayesian model), and the evidence, drove choice (one-sample t-test of beta coefficients against zero, prior: $t(24)$ = 2.47, p = 0.02; evidence: $t(24)$ = 10.33, p = $1e^{-9}$, Figure 3). However the combined effect of prior and coherence was sub-additive as evidenced by the negative effect of the interaction term (prior x coherence, one-sample t-test of beta coefficients against zero, $t(24)$ = -1.92, p = 0.07). In other words, the negative interaction between coherence and prior belief implies that participants relied more heavily on a prior when coherence was low, because then there was less evidence participants could integrate (Figure 4C). There was no significant effect of the other regressors on choice (one-sample t-test of beta coefficients against zero, constant: $t(24)$ = 1.25, p = 0.22;

priorXcompetition: *t(24)* = 0.73, *p* = 0.47; priorXcompetitionXcoherence: *t(24)* = 1, *p* = 0.32, coherence: *t(24)* = -0.34, *p* = 0.73; competition: *t(24)* = -0.61, *p* = 0.54).

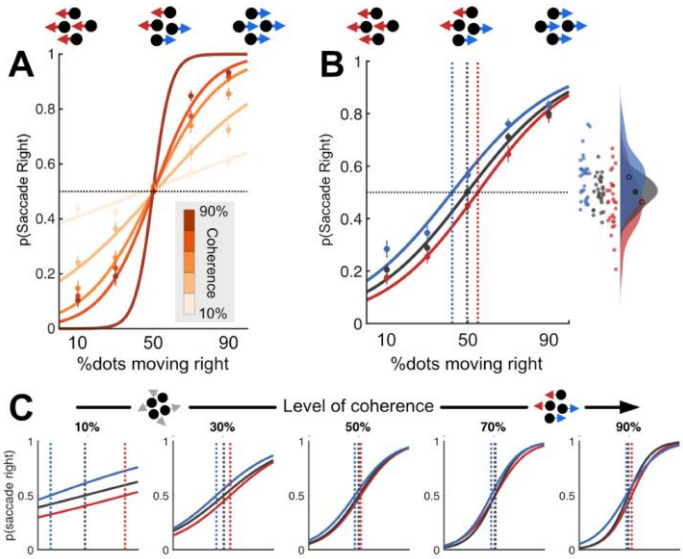


Figure 4: Influence of within-trial parameters and across-trial parameters on choice. **A)** Psychometric functions on the group-level for different coherences. Higher coherence led to better performance. **B)** Psychometric functions for different outcome probabilities. Adaptation of the prior belief shifted the point of subjective equality in favour of the outcome probability. **C)** shows that the influence of the prior belief was stronger for low coherence levels, i.e. when there was less evidence to base the decision on (reproduced from Marshall et al (*submitted*)).

These behavioural results suggest that

a) At the time of feedback, processing should occur relating to the updating of the prior;
b) The processing of the veridical feedback may be affected by the strength of evidence presented during the trial.

We went on to test these hypotheses in MEG data.

5.4 MEG methods

5.4.1 MEG acquisition

Data was acquired with an Elekta Neuromag VectorView 306 MEG system at 1kHz. Before participants were seated in the MEG scanner, their head position was identified using 4 head position indicator coils that were placed on the head. This allowed us to determine the head position of the participant in the scanner at the start of each session. In order to get a full model of the head, head landmarks including the pre-auricular points, nasion and 200 additional points on scalp, nose and face were acquired with a Polhemus Isotrack system. Further EEG electrodes were placed below and above the left eye and on each temple to record eye movements for artifact rejection later on. Another electrode was placed on the wrist to record ECG. Additional to the MEG acquisition, the eye position was recorded with an eyetracker (Eyelink 2000).

5.4.2 MEG pre-processing

For the pre-processing, customized MATLAB scripts were written using fieldtrip and OSL by Tom Marshall. First, data was maxfiltered. Then data was downsampled to 200Hz and bandpassed filtered between 1 and 80Hz. To remove the line-noise an additional bandstopfilter was applied around 50Hz. In order to remove trials with outliers and other artifacts including eyeblinks and ECG the function osl_detect_artefacts and independent component analysis was used.

While the main project used source-reconstruction with additional acquired MRI scans, my analysis is limited to the gradiometers over central parietal cortex. We chose to analyse planar gradiometers rather than magnetometers as these are maximally sensitive to sources directly underneath the sensor location. I took the clean data and first combined the gradiometers based on the ft_combineplanar function, which calculates the root mean squared (RMS) across the two planar gradiometers at a given location. Then, I z-scored the data and averaged over all combined channels for each participant for subsequent trial based GLM analysis. I averaged across all sensors because I don't have strong topographic predictions about the effect, other than that I should be able to measure it at parietal sensors.

5.5 MEG analysis and results

The aim of this chapter is to investigate how brain signals over parietal areas are influenced by evidence and prior belief. The longstanding theory that parietal cortex tracks the accumulation of evidence has recently been challenged by studies showing that inactivation of parietal cortex does not affect performance on these tasks (Erlich et al., 2015; Hanks et al., 2015). One suggestion is that parietal cortex accumulates evidence but is not necessary for action selection on a trial to trial basis; rather parietal cortex accumulates information across many trials to build a prior belief (Akrami et al., 2018). From behavioural results I know that both prior and evidence affect the choice behaviour of participants. I would expect that I find a modulation of the signal measured over parietal areas by prior belief and possibly evidence (the interaction of absolute and relative coherence) at the time of feedback. Only at the time of feedback the participant has seen the evidence and has a new estimate of the prior based on the feedback she just received.

5.5.1 Effect of trial outcome on feedback processing

Methods. If the MEG signal over parietal cortex reflects feedback processing, then we might expect a difference in the event related field (ERF) following 'correct' and 'incorrect' feedback events. To investigate this, I first plotted the cross-subjects mean ERF for correct and incorrect feedbacks, broken down by leftwards and rightwards saccades (Figure 5A). Seeing no effect of saccade direction, I collapsed across leftwards and rightwards trials and calculated the difference wave between correct and incorrect feedback events (Figure 5B,C).

Results. In Figure 5 the mean effect of correct and incorrect trials for left and right button presses across the trial is depicted. Only after the participant had observed the feedback a difference between correct and incorrect trials emerged. The first deflection at about 3.9s is possibly a saccade artefact, linked to the saccade to the feedback. This small peak is followed by a larger peak with a higher amplitude for incorrect trials, irrespective of a left or right choice. This difference is consistent across participants as shown by the difference wave (Figure 5B,C). Incorrect trials were more surprising than correct trials and therefore might have required a larger update of the prior.

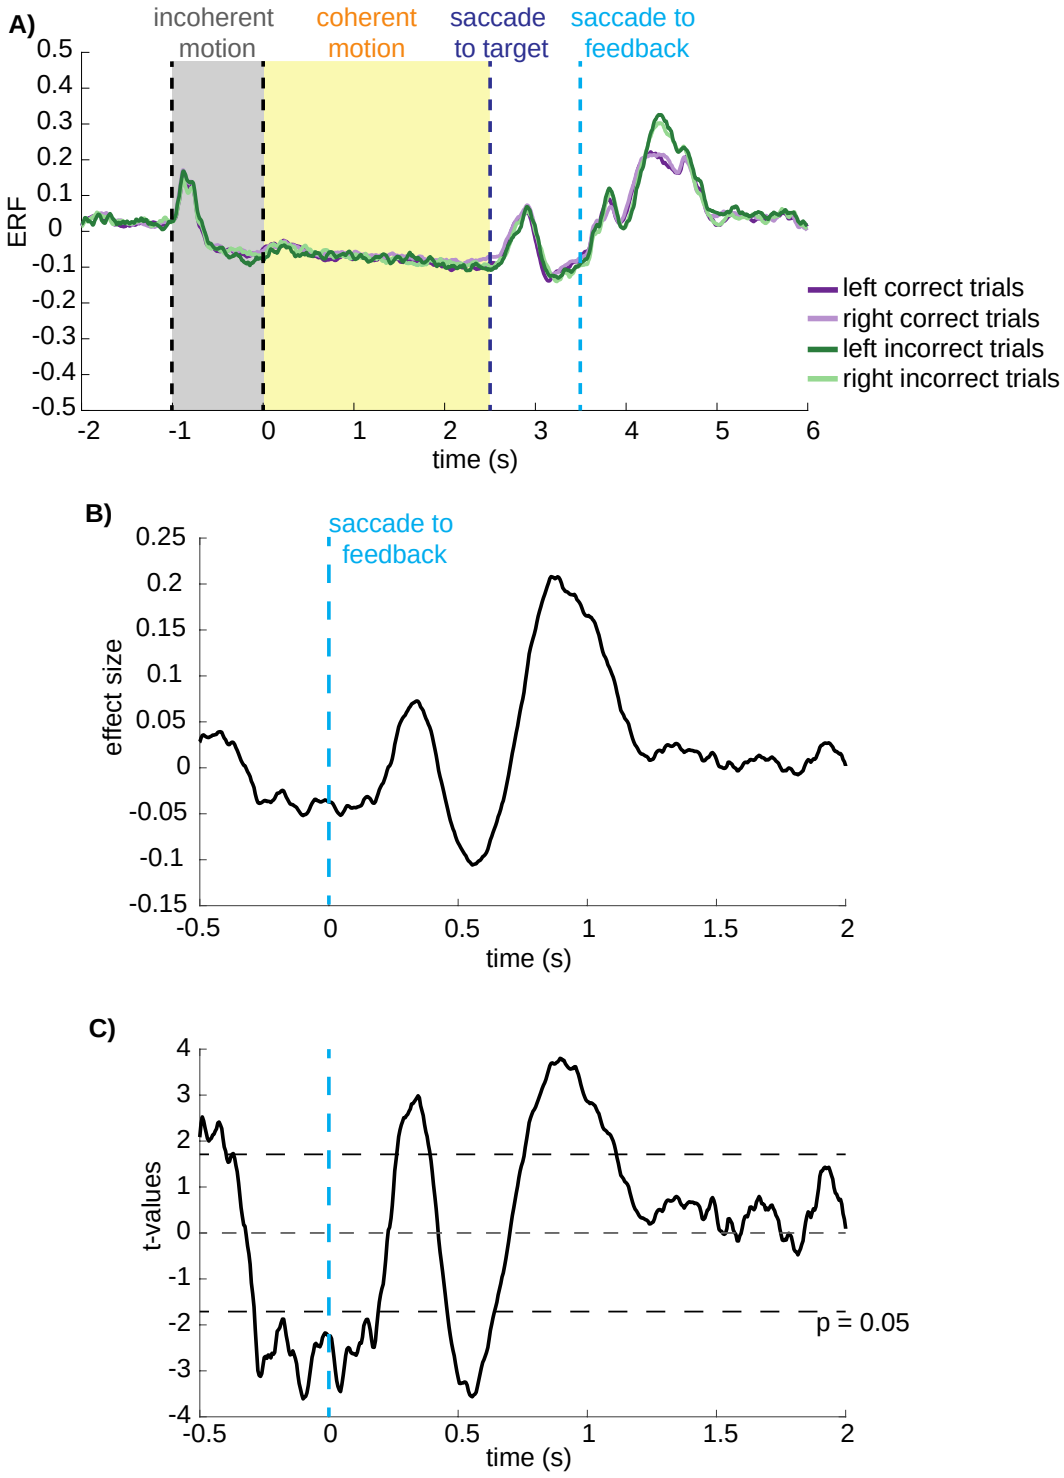


Figure 5 (previous page): Incorrect trials lead to a higher update of the feedback signal in MEG. **A)** Averaged MEG signal across trials at the group-level separated for correct and incorrect trials as well as choice on current trial. Only after the feedback has been shown there was a difference between incorrect and

correct trials, which did not depend on choice. The first peak after the feedback at about 4s is most likely an artefact of the saccade from the target back to the feedback. However, at about 4.5s there was a significant higher update for incorrect trials than for correct trials. **B,C)** The difference between incorrect and correct in the MEG signal was present across participants. The difference wave was calculated for each subject and then averaged. T-values and effect sizes correspond to the group-level. Note that zero on the time axis is now the presentation of feedback.

5.5.2 Effect of expectation strength on feedback processing

Methods. If the ERF following feedback reflects feedback processing, then we might expect it to be affected by the strength of the prior, as feedback is more informative when the prior is weak (when there is high uncertainty about the expected dots direction). Similarly, feedback is more informative when evidence was week (i.e. the correct dots direction was hard to discern from the perceptual evidence).

Figure 6 depicts the results of a trial based GLM analysis in which I included regressors for the difference between correct and incorrect trials, the absolute size of the prior (based on the Bayesian model fit) and the absolute strength of the interaction between coherence and competition, which is the total amount of evidence in the trial. A constant regressor is also added but not shown. In order to get a sense whether there is a significant effect of any regressor on the data I plotted a dotted line that depicts a t-stat that corresponds to a p-value of 0.05 (Figure 6, bottom column). If a cluster of t-values is below or above this border, the effect should be significant. In addition, I averaged the subject-specific effect sizes for the time window in which the t-value for the regressor describing the difference between correct and incorrect trials is consistently above 2. This is at about 0.8 to 1s (Figure 5A). Then I performed a t-test across subjects to test for significance. The time window was chosen based on this regressor because of its significant influence on the MEG signal and because I assumed that the timepoint at which this regressor has the largest influence, the other regressors will also have the largest influence.

Results. In addition to the effect of correct and incorrect trials, the prior also seemed to have a very small negative influence on the ERF at about 0.5s.

However, this was not detected by the test for significance with the averaged betas because of the chosen time window ($t(24) = 0.24$, $p = 0.81$). (Note that although the prior does cross the p=0.05 line in Figure 6, this is uncorrected for multiple comparisons across time). The effect of the trial was slightly early at about 0.5s than the effect of correct and incorrect trials. The evidence in the trial did not seem to have a significant impact ($t(24) = -1.02$, $p = 0.32$). It is notable though that the effects of correct and incorrect trials and the evidence regressor are in opposite directions (Figure 6, at about 1s). Perhaps, the regressor for the difference in correct and incorrect trials reflects how surprising the feedback is. It might be less surprising to be incorrect when the evidence was weak.

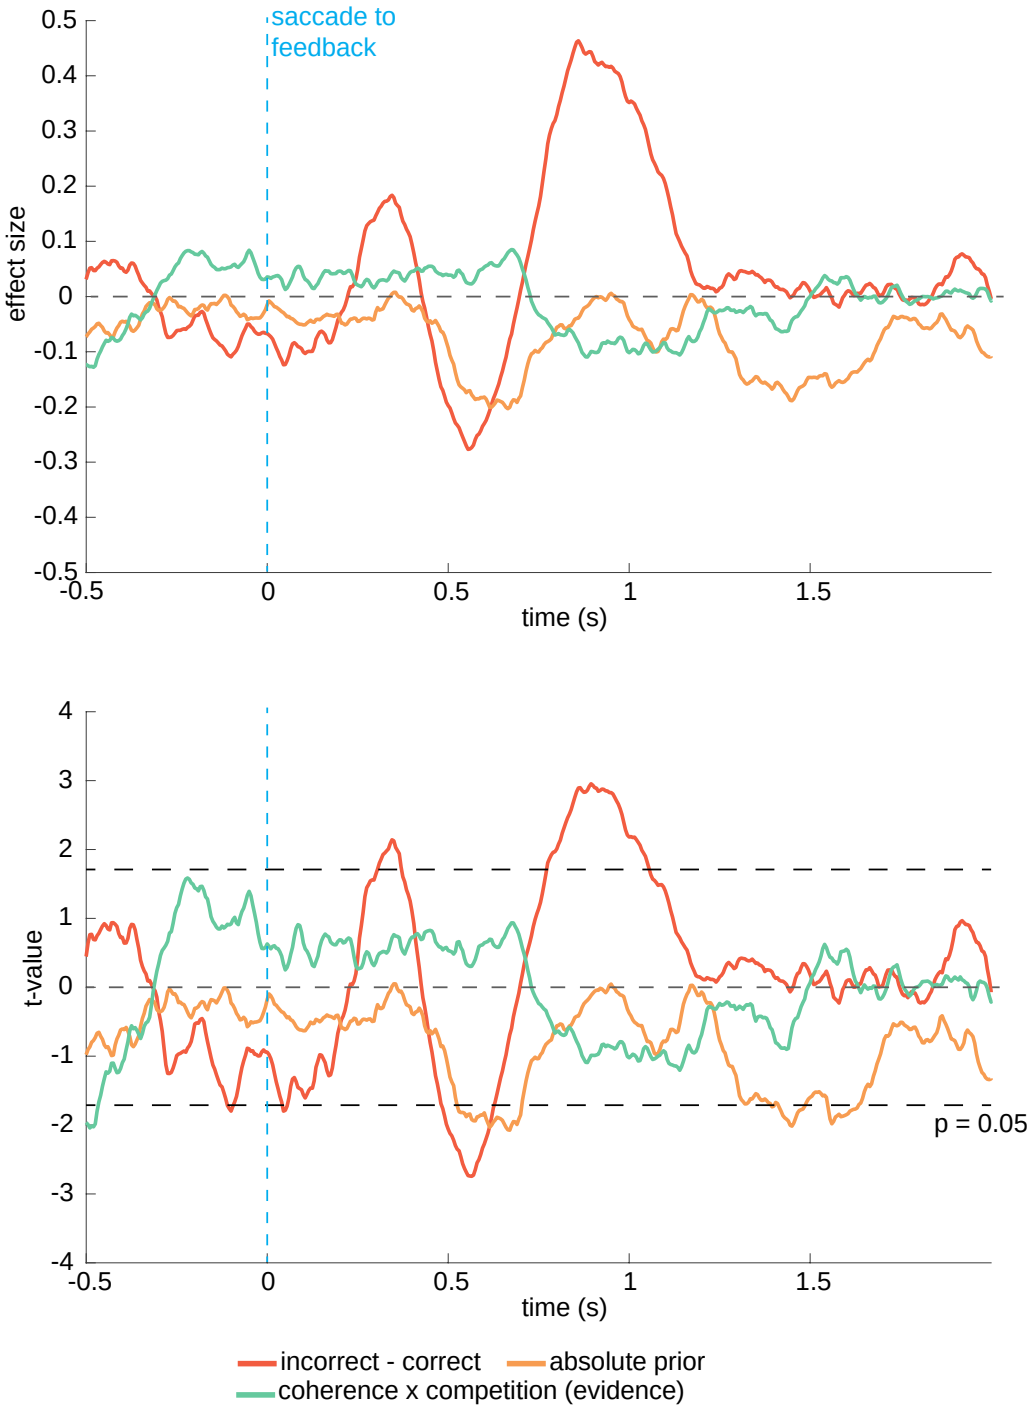


Figure 6: The prior showed a weak influence on the MEG signal at the time of feedback. Group-level regression analysis with a signed regressor for the difference between correct and incorrect trials, a regressor for the absolute prior scaled for strength but not for target direction, and a regressor for the interaction between coherence and competition (evidence) on the current trial. Constant regressor not shown. Upper panel shows effect sizes and lower panel the uncorrected t-values.

5.5.3 Interaction of expectation strength with feedback type (correct/incorrect)

Methods. Intuitively, feedback processing should be enhanced when participants chose the incorrect target, because the prediction of the correct target based on the previous trials was incorrect. We previously observed that feedback-linked ERFs were larger for 'incorrect' than 'correct' feedback. If this difference reflects processing of the feedback in relation to one's prior beliefs, then there should be an interaction between the feedback type (correct/incorrect) and the prior strength. Therefore, I repeated the above analysis separately for correct and incorrect trials. Then I calculated the difference wave between the effect sizes for priors on correct and incorrect trials as well as evidence on correct and incorrect trials (Figure 7).

Results. However, the influence of the prior on the MEG signal did not depend on whether trials were correct or incorrect. Nor was there a difference for the effect of evidence between correct and incorrect trials.

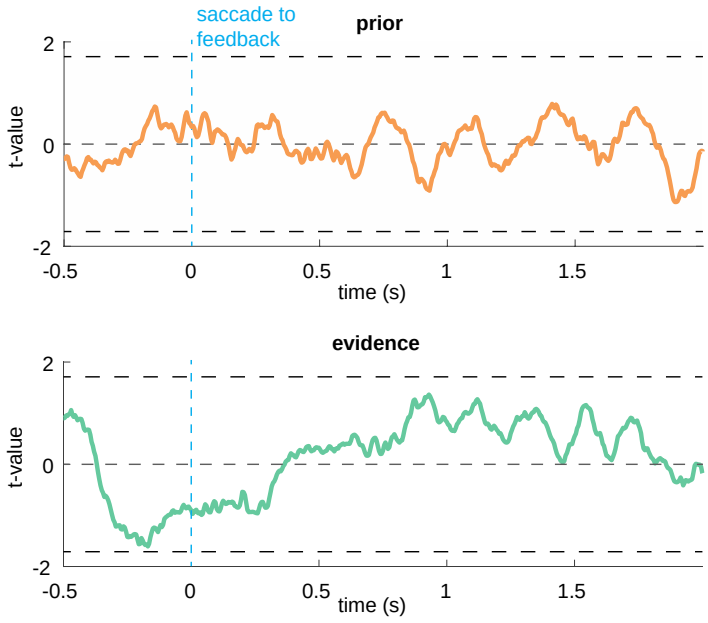


Figure 7: Updating of the prior does not depend on correct or incorrect trials. Group-Level GLM with a constant regressor (not shown), absolute prior and interaction between coherence and competition (evidence). Upper panel shows the t-values for the difference wave between the prior in correct and incorrect trials. Lower panel shows the t-values for the difference wave between the evidence in correct and incorrect trials.

5.5.4 Congruency of expectation and feedback

Methods. Updating of the prior did not depend directly on the feedback type (correct/incorrect), but rather the congruency between expectation and feedback. The feedback was in the form of an arrow indicating the correct dominant motion direction (left/right). It is possible, for example, that a participant had a prior favouring rightwards motion, for some reason (perhaps based on the perceptual evidence) responded leftwards, and then obtained an 'incorrect' feedback which actually confirmed their prior belief.

In order to investigate whether the update of the prior depended on whether it aligned with feedback or choice on a trial, I implemented two GLMs (Figure 8). Both GLMs are similar to the one in Figure 6. I included a constant, which is not shown and regressors for evidence and prior signed for choice or feedback, respectively. This means, that both prior and evidence regressors were signed positive if they favour the same direction as the subsequent choice (or feedback) and negative otherwise. I did not include a regressor for the difference between correct and incorrect trials if evidence was signed for choice because this regressor is correlated with the regressor for the difference between correct and incorrect trials (Supplementary Figure 1). As for Figure 6 I performed a t-test on the averaged subject-specific betas for all regressors in the window at which the t-values for the regressor for the difference between incorrect and correct trials is above 2 (~0.8 – 1s).

Results. This analysis showed that a negative effect of absolute evidence is strongly signed for choice at the time of the peak at 1s (Figure 8, left column, $t(24) = -3.57$, $p = 0.002$). In other words, the ERF relating to feedback processing was reduced if the participant made a choice *against* strong perceptual evidence for correct and incorrect trials. Perhaps, the peak for correct and incorrect trials at 1s represents an update of the prior based on what the feedback told the subject, over and above what she already learned from seeing the trial itself, or perhaps trials on which the choice was counter to strong perceptual evidence were those on which the prior was particularly strong and less prone to updating. There was no effect of the prior signed for choice ($t(24) = -1.27$, $p = 0.22$). The update in the MEG

signal did not seem to depend whether prior and evidence were aligned with feedback (prior: *t(24)* = 1.21, *p* = 0.24; evidence: *t(24)* = -1, *p* = 0.33), perhaps favouring the second explanation (Figure 8, right column - note in this GLM the evidence regressor is the same as in Figure 6, because evidence is always aligned with feedback as designed by the experimenter).

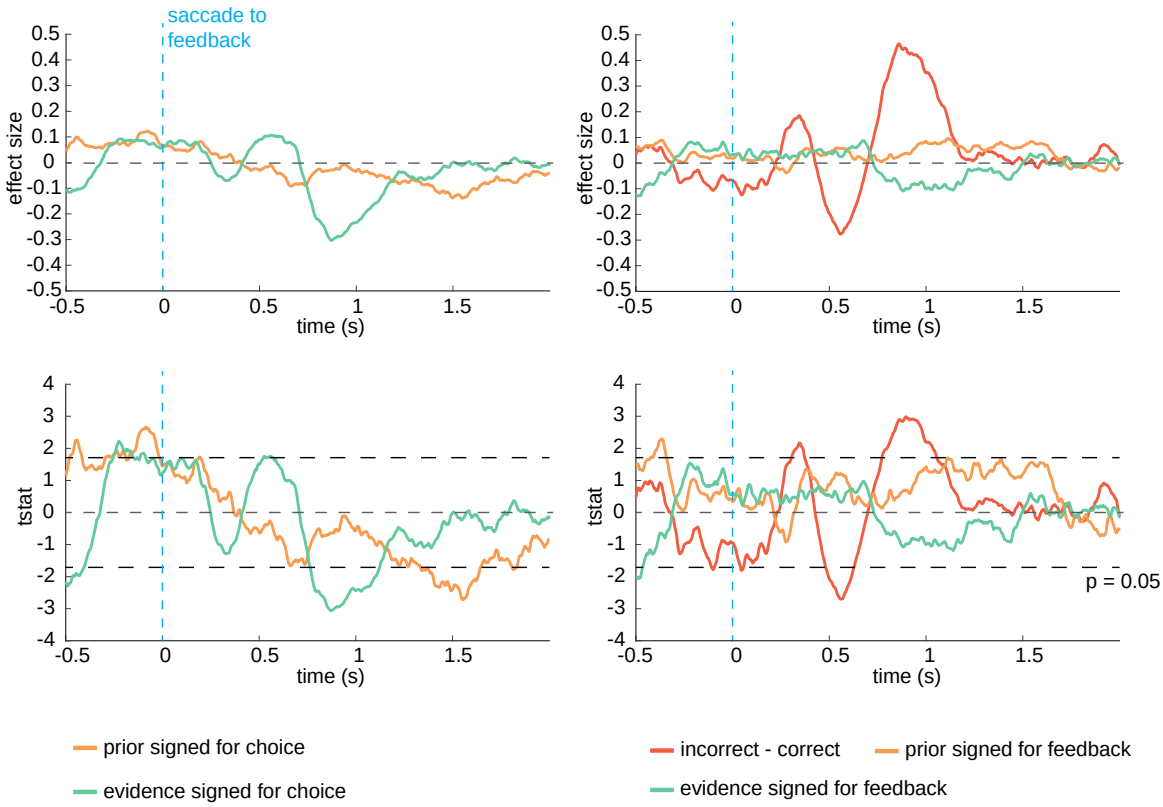
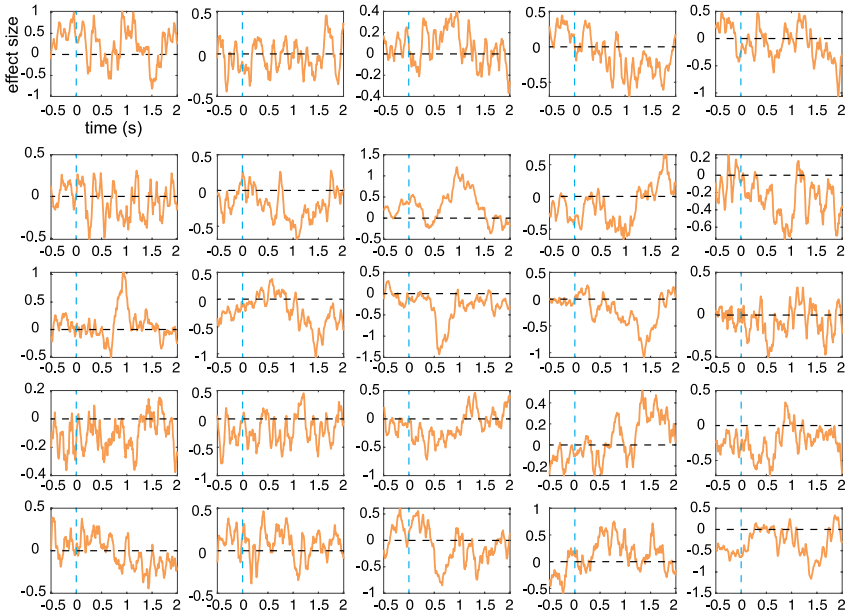


Figure 8: Strong evidence signed for choice had a negative effect on the MEG signal. Group-level GLMs for prior and evidence signed for choice (right column) and signed for feedback (left column). Both GLMs included a constant regressor not shown. Evidence and absolute prior signed either for choice or feedback were positive if they aligned with choice and feedback, negative otherwise. While a regressor for the difference in incorrect and correct trials was included when evidence and absolute prior were signed for feedback, it has been omitted for the GLM signed for choice, because evidence signed for choice and the regressor for the difference in incorrect and correct trials were correlated (Supplementary Figure 1). The regressor for evidence signed for feedback was the same as the regressor for the absolute evidence in Figure 6 because by experimental design the feedback always pointed in the direction dictated by the competition. Upper panel shows effect sizes in the lower panel uncorrected t-values are plotted.

5.5.5 Single-subject results for prior belief and trial outcomes

The influence of the absolute prior (Figure 6) was quite weak. This could have 2 reasons. Either, only a few participants integrated the prior probability of which target is correct into their choice or for all participants there was only a small effect of prior on the MEG signal at the time of feedback. I plotted the absolute prior from the group-level GLM in Figure 6 for each subject separately (Figure 9A). There was a notable effect of the prior in only some of the participants and the timing and direction of the effect varied quite a lot. In general, the results on the single subject level were similar for the effect of the difference between correct and incorrect trials. Though the effect was overall stronger in most participants it was still quite noisy (Figure 9B). This signal was sometimes in the same direction as the prior but sometimes also in the opposite direction. Either the two signals have different underlying generators or these effects are due to a single neural population that operates differently across participants.

A)



B)

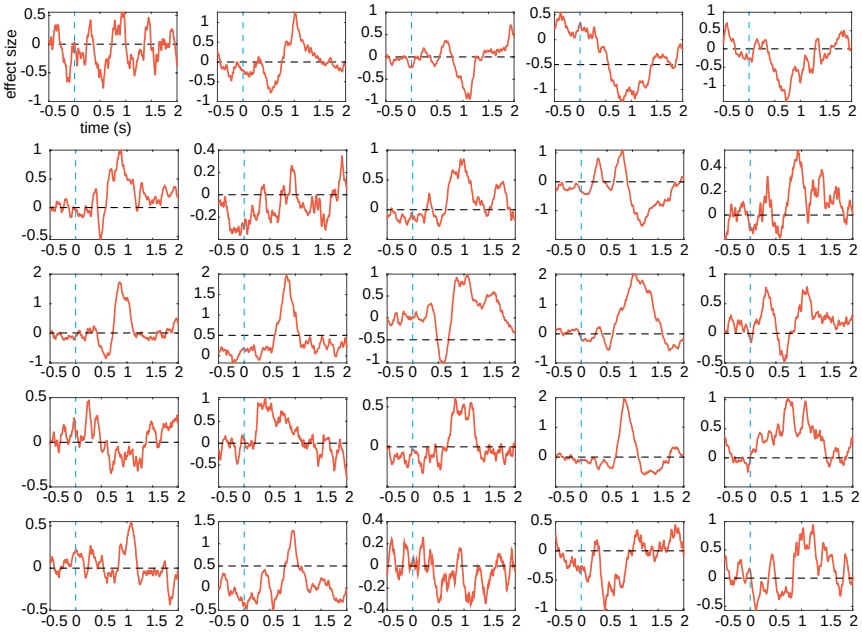


Figure 9: Single subject effects for prior belief were noisy, and more variable across subjects, than the effect of the difference between correct and incorrect trials. Plotted are the results on the single subject level for prior belief (**A**) and the difference between correct and incorrect trials (**B**) of the group level GLM in Figure 6. The group level GLM included a regressor for the evidence and a constant, as well. Blue dotted line is the time of the saccade to the feedback.

5.6 Discussion

A combination of evidence and prior belief integrated over the veridical feedback of past trials guided choice. The influence of the prior on the current choice was smaller if current evidence was strong. If the past trial was difficult the prior had a stronger impact on the current choice. Neural signals recorded over parietal cortex showed that at the time of feedback, that there is a stronger update in the signal for incorrect trials. However, surprisingly, there is only a small impact of the prior that is independent of choice and feedback. Evidence has only a minor effect at the time of feedback. If evidence aligns with choice, it reduces the effect on the update related to the outcome of a trial (whether a trial is correct). Possibly, this represents what participants learn from the trial over and above what they saw during the dot-motion.

5.6.1 Stimulus certainty influences choice

In theory, a Bayesian model that only takes into account the veridical feedback should be sufficient to solve the task optimally. I expected that participants learned the hazard rate that determines when the outcome probability switches. From fMRI we know that the brain can track the task volatility and adapt its learning rate accordingly (Behrens et al., 2007). Therefore, it is no surprise that participants seem to use prior probabilities to solve the task (Figure 3).

However, the decision is not only influenced by the integration of past feedback, but also by the strength of the evidence. If evidence is strong, trials are easy and participants do not need to rely on their estimated prior as much. In other words, their certainty about the stimulus is small and their decision confidence high. In this case there is not much participants can learn from the trial, therefore they don't need to update their prior. This behaviour has also been found in rats (Lak et al., 2017, 2020a).

In a similar manner, participants might use evidence of the previous trial to guide choice. If the last trial was easy, they did not learn much and it might have a weaker impact on the current choice (Figure 2). In a perceptual choice task rats specifically tend to repeat the last choice if the previous trial was difficult. In that case decision

confidence was low on the previous trial and therefore the prior probability might have driven the decision more strongly (Lak et al., 2020b). This is possibly represented by the negative effect of past evidence on current choice (Figure 2). It is similar to a prediction error from the previous trial. The prediction error is defined as the difference between expectation and evidence. In this case the expectation is represented by the influence of past feedback, which has a positive sign, while the evidence on that trial has a negative sign. This could mean that if evidence is strong, it reduces the influence of the past feedback, but if it is weak, the influence of the previous feedback is strong. For this reason a reinforcement learning model might be better to explain the behaviour, as it allows to incorporate decision-confidence (Braun et al., 2018; Lak et al., 2020b).

These results are in contrast to a study by Zylberberg et al. (2018) who used a traditional random dot motion task. Participants were required to complete several blocks of trials. Each block was either biased to the left or right. After each trial participants reported the confidence in their choice and their belief about the bias in the block. They found that high confidence on the previous trial, i.e., the last trial had a high motion coherence and was easy, had a positive impact on the updating of the prior belief of the bias in a block. However, when previous trials were easy in the task in this thesis, the influence of the feedback on trial $t\text{-}_1$ was reduced on the current one, meaning that the update of the prior is negatively correlated with the confidence on the last trial. One reason for this difference might be the design of the tasks. Crucially, in the study by Zylberberg et al. (2018), participants did not receive veridical feedback after each trial and used their confidence instead to update their prior believe over the outcome probability. In this case, I would assume that participants update their prior more strongly when their confidence is high about their choice. However, in the task in this study this is not necessary as participants receive veridical feedback after each trial. In this case participants might not update their prior as much when confidence is high on the previous trial, i.e. evidence is strong because they were certain about the correctness of their choice and might have paid less attention to the feedback. That is also reflected in the MEG signal over parietal cortex which shows a higher peak when trials are incorrect. Participants are less likely to be correct when evidence and confidence

are low. Another possibility for the negative effect could also be collinearity between coherence and feedback.

5.6.2 Influence of prior belief on parietal cortex

I expected that I would find brain signals recorded over parietal cortex at the time of feedback reflecting the update of the prior based on what we know from rodent and monkey studies (Akrami et al., 2018; Erlich et al., 2015; Hanks et al., 2015; Katz et al., 2016). However, there is only very weak evidence for an update of the prior at the time of feedback. Only whether the choice is correct is processed by parietal cortex at the time of feedback but this does not seem to influence the prior belief updating (Figures 6-7).

In the study by Marshall et al (*submitted*), a neural signal encoding the prior belief emerges before the coherent dot-motion period biasing the evidence integration process during the period the coherently moving dots are presented. However, while the brain integrates the evidence presented in the coherently moving dots, the prior belief gets also updated (Figure 10). This could mean that there is only little to be learned in terms of the prior belief once participants receive their feedback. Modelling of behaviour suggests that prior expectation is added during the accumulation process for example by lowering time-dependently the decision bounds (Drugowitsch et al., 2019) or by influencing the build-up rate of the accumulation process (Hanks et al., 2011; Urai et al., 2018).

These results suggest that human parietal cortex is involved in the online integration of prior belief as has been suggested in rodents (Akrami et al., 2018). It fits to the assumption that at least part of parietal cortex represents a priority or salience map, which sorts behavioural targets based on their importance for future behaviour (Bisley and Goldberg, 2010; Ipata et al., 2009). The homolog area LIP in monkeys has connections to upstream visual areas as well as downstream areas such as FEF and anterior cingulate cortex (ACC) (Blatt et al., 1990; Lewis and Van Essen, 2000). This puts the area in perfect position to integrate lower visual information about stimuli as well as motor intention possibly mediated by

FEF but also information about the volatility of a stimulus and prior probability that might be integrated in the ACC (Behrens et al., 2007; O'Reilly et al., 2013).

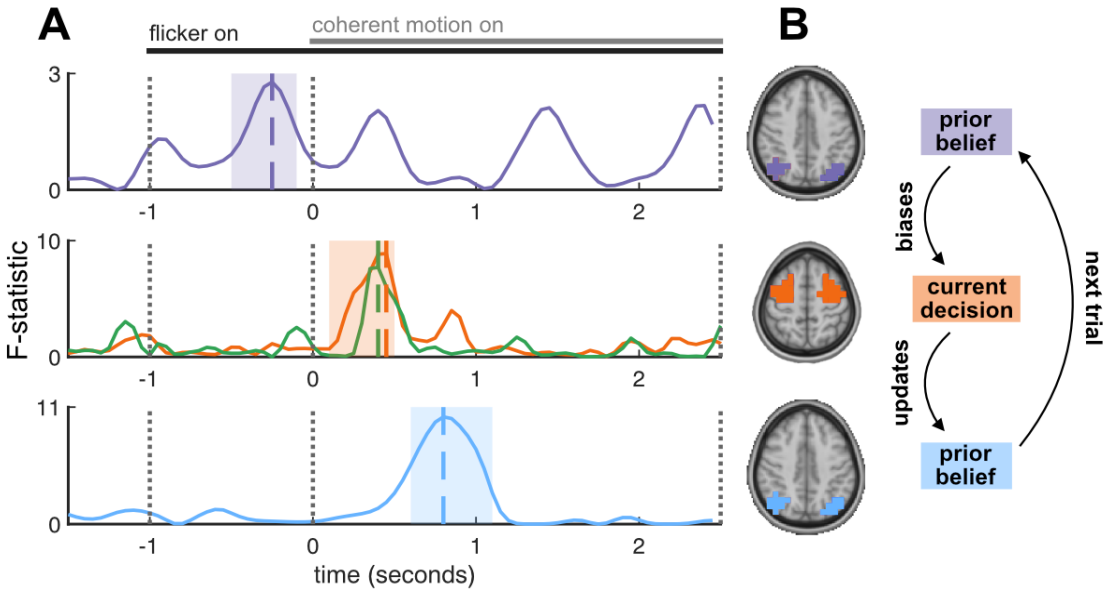


Figure 10: Schematic of proposed interaction between parietal and frontal cortex during perceptual decision-making. A) Time-course of prior belief and current decision signals (evidence) during dot motion viewing in the task described in this chapter. Purple = prior belief during incoherent motion, blue = prior belief during coherent motion, orange = coherence, green = competition **B)** Proposed interplay between parietal cortex and FEF during dot motion viewing. (reproduced from Marshall et al, *submitted*)

5.7 Supplementary Figure

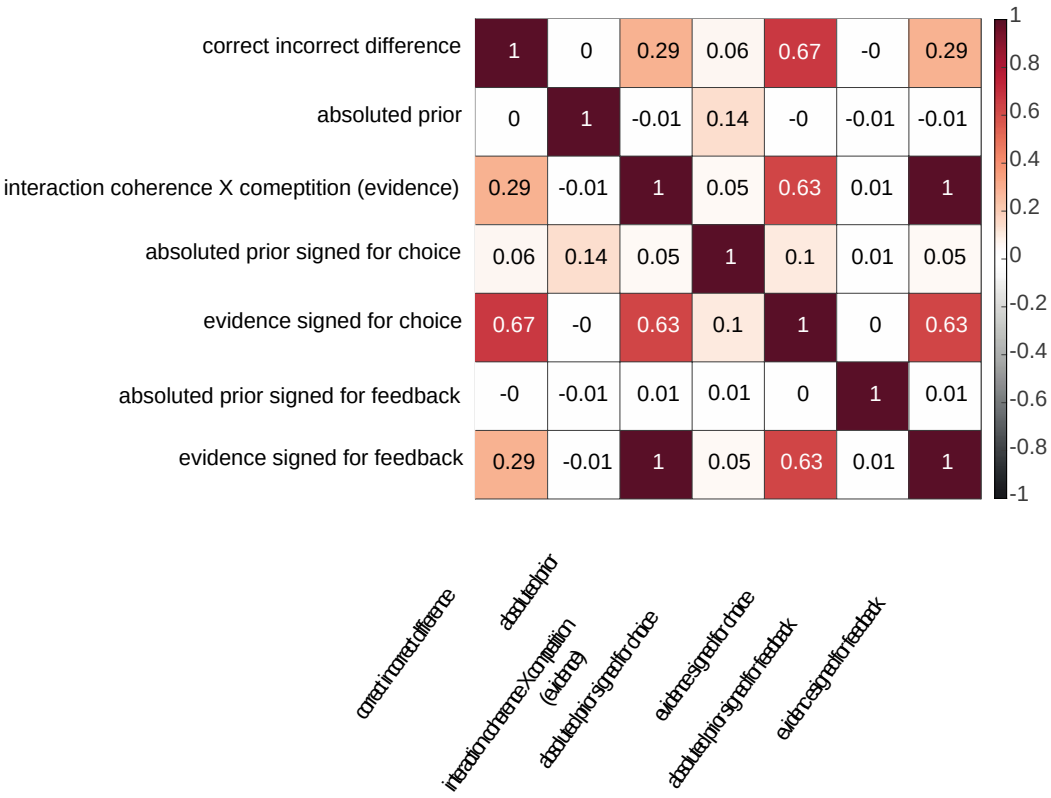


Figure 1: Correlation matrix of different regressors for MEG analysis Regressors for each participant were correlated with each other and averaged across participants. This correlation matrix informed the subsequent regression analysis in Figures 6 to 8.

6 Conclusions and future directions

This thesis has aimed to further our knowledge of evidence accumulation during decision-making under uncertainty in humans. In particular, in Chapters 2-4 I investigated how human behaviour and neural activity adapted to a novel continuous evidence integration task. Subjects had to constantly integrate changing evidence in order to decide whether they were in a trial period. This means that trial period onsets had to be inferred from noisy evidence. Using a 2x2 design I tested how participants adapted to the frequency of trial onset and trial length. I showed that a leaky accumulator model, the OU process, can solve this task by adapting decay rate (λ) and threshold (z) to maximise reward (Chapter 2). The behaviour of participants and the OU process was quite similar, suggesting that participants discounted past evidence as well in order to solve the task (Chapter 3). There is also evidence that neural signals recorded with EEG adapted strongly to different trial frequencies but less to trial length (Chapter 4). A CNV signal showed an earlier negative deflection that reached a lower threshold for rare trial conditions, perhaps reflecting the need for more evidence to commit to a response as reflected by the behavioural integration kernels (Chapter 3). A positive potential located centrally on the scalp was found for the update of each new incoming evidence sample. The peak of this signal was higher for rare trial conditions, implying that a large change in coherence between previous and current evidence sample is more surprising when trials are rare. However, I could not find a signal in the brain that directly reflected a discounting process of past evidence as predicted by the OU process. This might be explained by the fact that participants used a mixture of accumulation and Bayesian inference processes to detect change points at which intertrial periods switched to trial periods (Glaze et al., 2015; Murphy et al., 2020).

In contrast, in Chapter 5 participants had to integrate static evidence over several seconds within a trial but they also had to integrate the veridical feedback over many trials to make a decision. Behavioural analyses showed that participants relied on past feedback and current evidence to make a choice. I hypothesized that after participants received feedback, I would find a signal in parietal cortex reflecting the update of the prior belief. However, I only found a strong update of

whether the choice of a participant was correct and only a very weak update of the prior. Instead, analyses carried out by Marshall et al (*submitted*) showed that there already is an update of prior belief during the evidence viewing period before participants make a choice. This could indicate that prior belief has an influence on the evidence accumulation process as suggested by other studies that showed an influence of prior belief on the drift rate of a DDM.

6.1 Combining accumulator models and reinforcement learning to understand decision-making

The two studies in this thesis looked at decision-making processes that develop over very different timescales. More specifically, there are at least 3 layers to many naturally occurring decision processes: 1) Observers have to integrate immediate evidence, which happens on a short timescale, such as the noisy evidence in the continuous RDM. 2) They have to assess at each timepoint how likely it is that a change in the state of environment occurs, such as change from intertrial period to trial period in the continuous RDM (Chapter 2) or a change in the prior in the RDM task in Chapter 5. Together with the first step this determines how likely a response is. 3) In addition, observers need to learn how likely a change in the environmental state is. For the continuous RDM study this is the trial frequency subjects learned through training. In the RDM task in Chapter 5 it is the hazard rate that determines the average number of trials before a change in prior occurs. Therefore, we need to ask how the brain integrates these different processes that develop over very different timescales and allows for adaptation if the environment evolves. For this we need models that can take into account these different layers of the decision-making process.

The models used in this thesis were both designed to capture decision-processes on a single level, but not well suited to capture decision-processes on different layers simultaneously. I assumed a mechanistic OU process can explain the behaviour in the continuous RDM, in which the decay rate and decision-threshold determine the adaptation to trial frequency and trial length. However, the OU process is agnostic as to how the brain adapts to different trial frequencies and lengths in the first place. We instead assumed that participants learned these

statistics, and adopted appropriate OU parameters, during training. In contrast, the Bayesian model in Chapter 5 only takes into account the prior belief updating across trials. It does not consider how current evidence within a trial might influence this belief updating – it only considers the veridical feedback. A model is needed to explain how the brain might integrate both the immediate evidence and possible changes in the environmental state that have to be learned. Recent studies that have combined reinforcement learning (RL) models or Bayesian models with accumulator models such as the DDM (Miletić et al., 2020) may provide a promising way forward in this direction.

Reinforcement learning (RL) as well as Bayesian models both describe how subjects should adapt to changes in the environment over time in order to maximise reward. Traditional RL models were originally used to investigate value-based decision-making and consist of two steps (Lee et al., 2012). The first step is to calculate the reward associated with a choice. It can be understood as the sum of the current estimated value of this choice option and the prediction error which is the difference between expected and actual reward on the last trial. The influence of the prediction error on the expected value of a choice is determined by a learning rate (α) describing how much past experience influences the current choice. The second step is described by a soft-max function to calculate the probability to choose an option. This function contains a parameter, called inverse temperature (β) that describes the sensitivity to choose the option with the highest reward or to explore other options. In Bayesian RL models, as described in Chapter 5, the update from the feedback on the last trial alters the likelihood that a choice is correct on the next trial while taking into account the *uncertainty* associated with this likelihood, which effectively alters the learning rate that is associated with the feedback (O'Reilly et al., 2012).

Some studies have shown that it is possible to combine RL with a DDM by linking the soft-max function with the drift-diffusion process. Conceptually, β is the equivalent to the ratio of the decision-threshold scaled by noise, while the drift-rate in the DDM can be thought of as the difference between the value expectations of choice options (Fontanesi et al., 2019a, 2019b; Miletić et al., 2020; Milosavljevic et

al., 2010; Pedersen et al., 2017; Tuerlinckx and De Boeck, 2005). Others have combined Bayesian models with accumulator models to investigate the influence of prior belief on evidence integration by using different link-functions that describe the influence of prior belief on the accumulation process (Glaze et al., 2015; Veliz-Cuba et al., 2016; Yu and Cohen, 2008; Zhang et al., 2014). Such a model has been described in the discussion in Chapter 3 (Glaze et al., 2015; Veliz-Cuba et al., 2016).

The advantage is that, theoretically, different RL or Bayesian algorithms can be integrated with different accumulator models including an OU process. Therefore, a framework of combined RL/Bayesian and accumulator models is very flexible and can be applied to a variety of different decision paradigms (Miletić et al., 2020). They can give better explanations of behavioural measures such as reaction times and accuracy in value-based decision-making tasks. (Fontanesi et al., 2019a; Miletić et al., 2020; Pedersen et al., 2017). In perceptual decision-making tasks they can account for trial-to-trial variability and make explicit hypotheses about the drift-rate in accumulator models (Fontanesi et al., 2019a, 2019b; Frank et al., 2015; Miletić et al., 2020).

Using a combined model of RL/Bayesian models and evidence accumulation we can make explicit hypotheses about how neural populations might integrate different processes that require learning, prior expectation and accumulation of evidence. There is already some evidence for an overlap in the encoding of these different decision-processes. For example, one study in humans has shown that attractor models originally used as a biological explanation of evidence accumulation also predicted value based decision-signals in parietal and prefrontal cortex (Hunt et al., 2012). Marshall et al (*submitted*) has also shown that attractor networks can account for activity representing prior expectation in parietal cortex. These results suggest that the same neural processes underly the representation of current evidence and prior expectation and that they might be found in the same neural populations. However, only by modelling the encoding of prior expectation and evidence accumulation simultaneously can we explicitly probe how prior expectation and evidence accumulation are encoded in the brain.

Thus far, most studies applied this combined model framework to well-studied paradigms in decision-making. However, in the future we should combine this new set of models with novel task paradigms such as developed in this thesis to understand how the brain deals with uncertainty on different levels of the decision-formation processes. Specifically, the continuous RDM paradigm developed in this thesis might be of use because it combines uncertainty on the level immediate evidence integration but also on the level of expected changes in the environmental state. In addition, we can modulate the reward expectation by altering the points participants gain or lose by making a choice and investigate how this influences the model predictions and hence brain activity.

6.2 Understanding cognitive processes in mental diseases

In order to treat neuropsychiatric disorders and their symptoms we need to understand their underlying cognitive causes. Many psychopathologies show similar impairments in goal-directed behaviour such as decision-making (McTeague et al., 2016). However, we have limited understanding of the neural substrates causing these impairments and how they differ between different mental diseases. One issue is that in clinical psychology tests for impairment of cognitive function are often quite unspecific and require different cognitive processes (Snyder et al., 2015). This means, that if patients show similar impaired behaviour, this might result from a number of different deficits in neural processes underlying cognition. Therefore, we need precise tasks that test for different cognitive impairments as well as models that allow us to probe the underlying neural processes.

A good example for a combined approach of specific tasks with models in order to understand cognitive deficits comes from research in schizophrenia. On the synaptic level it has long been suggested that N-methyl-D-aspartate receptor (NMDA) hypofunction is a common characteristic in patients with schizophrenia (Kehrer, 2008; Nakazawa et al., 2012). However, only recently has there been research on how NMDA hypofunction maps onto circuit function using attractor networks that can explain deficits in cognitive functions such as working memory and decision-making (Cavanagh et al., 2020; Lam et al., 2017; Murray et al.,

2014b). Yet, there is further research needed to understand the neural underpinnings of cognitive dysfunction and how it affects different brain areas.

The improvement in our ability to measure and understand neural signals underlying complex decision-making tasks in humans can help us to characterise how they differ in patient populations. It has been shown that sub-anaesthetic Ketamine mimics impairments found in patients with schizophrenia and the dysfunction of NMDA receptors (Murray et al., 2014b; Starc et al., 2017). Salvador et al. (2020) investigated how sub-anaesthetic Ketamine influenced neural decision signals in a previously established decision-making task using EEG. They found abnormal encoding of evidence in frontal and parietal cortex and premature motor preparation signals leading to premature responses to uncertain beliefs. Similarly, other established task paradigms in cognitive neuroscience which we understand on the neural and computational level can be used to characterise deficits of cognition in schizophrenia as well as other mental diseases.

I anticipate that the novel continuous task paradigm developed in this thesis will allow our group and others to further explore the differences in cognition between healthy and different populations of psychopathologies. It has several advantages to other task paradigms:

1) It resembles an uncertain environment in which complex cognitive processes are necessary to solve it. We can apply computational models to understand the different cognitive processes involved in solving the task. This means we can measure how subjects deal with environmental volatility. The ability to process uncertainty is impaired in psychopathologies such as psychotic disorders (Cole et al., 2020). Therefore, it allows us to further characterise and understand the differences in normal and abnormal decision-making.

2) This paradigm is very flexible and can be adapted to fit the pathology of different mental diseases. For example, patients with depression often have problems to disengage from negative stimuli

(Grahek et al., 2018). In the past, tasks to probe this rumination of negative stimuli and thoughts included the identification of different facial expressions or the memorisation of words associated with more or less negative stimuli (Mor and Daches, 2015). For example, one could adapt my task in order to show a stream of facial expressions or words in which participants need to find more stable periods of positive or negative stimuli.

3) The continuous paradigm allows for a much higher SNR in EEG data which means that shorter sessions are required to acquire enough data and that we might be able to acquire single-subject specific behavioural and neural markers. This is specifically relevant for patient populations that might not be able to concentrate for longer periods of time. In Chapter 3 I have shown that behavioural integration kernels are subject specific (see Chapter 3, Figures 5-6), while preliminary analysis of regression ERPs in Chapter 4 suggest that subject-specific prediction error ERPs look strikingly similar across different runs consisting of 10min of data (see Chapter 4, Supplementary Figure 3). This means that this continuous paradigm might allow me to extract behavioural markers of complex cognitive processes. So far, the reproducibility of EEG markers has mostly been tested with resting state EEG (Corsi-Cabrera et al., 2007; Williams et al., 2005) and only a few cognitive tasks, such as an auditory oddball and Sternberg task (Cassidy et al., 2012; McEvoy et al., 2000; Williams et al., 2005). In a next step, the test-retest reliability has to be established for the continuous RDM paradigm. If it is possible to reproduce these subject-specific markers across sessions on different days, they could be used to track the development of cognitive processes over time. For example, one could compare the difference in these markers when patients are off or on medication, or one could test the dose-response to different levels of a drug.

In summary, the studies presented in this thesis have been able to further characterise decision-making under uncertainty in human behaviour and in neural signals. Both RDM tasks described in this thesis might be useful in the future to further understand how the brain integrates different components of decision-making processes such as evidence integration and belief updating. However, they should be combined with novel models that can capture these different components such as a model combining RL and evidence accumulation. In addition, the continuous decision paradigm I developed in Chapter 2 could be a promising diagnostic clinical application to characterise psychopathologies.

References

Ahumada, A.J. (1996). Perceptual Classification Images from Vernier Acuity Masked by Noise. Perception *25*, 2–2.

Akrami, A., Kopec, C.D., Diamond, M.E., and Brody, C.D. (2018). Posterior parietal cortex represents sensory history and mediates its effects on behaviour. Nature *554*, 368–372.

Alamia, A., and VanRullen, R. (2019). Alpha oscillations and traveling waves: Signatures of predictive coding? PLoS Biol. *17*, e3000487.

Albright, T. (1984). Direction and orientation selectivity of neurons in visual area MT of the macaque. J. Neurophysiol. *52*, 1106–1130.

Behrens, T.E.J., Woolrich, M.W., Walton, M.E., and Rushworth, M.F.S. (2007). Learning the value of information in an uncertain world. Nat. Neurosci. *10*, 1214–1221.

Bisley, J.W., and Goldberg, M.E. (2003). Neuronat activity in the lateral intraparietal area and spatial attention. Science *299*, 81–86.

Bisley, J.W., and Goldberg, M.E. (2010). Attention, intention, and priority in the parietal lobe. Annu. Rev. Neurosci. *33*, 1–21.

Blatt, G.J., Andersen, R.A., and Stoner, G.R. (1990). Visual receptive field organization and cortico-cortical connections of the lateral intraparietal area (area LIP) in the macaque. J. Comp. Neurol. *299*, 421–445.

Bogacz, R., Brown, E., Moehlis, J., Holmes, P., and Cohen, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. Psychol. Rev. *113*, 700–765.

Bogacz, R., Usher, M., Zhang, J., and McClelland, J.L. (2007). Extending a biologically inspired model of choice: Multi-alternatives, nonlinearity and value-based multidimensional choice. Philos. Trans. R. Soc. B Biol. Sci. *362*, 1655–1670.

Bogacz, R., Wagenmakers, E.J., Forstmann, B.U., and Nieuwenhuis, S. (2010). The neural basis of the speed-accuracy tradeoff. Trends Neurosci. *33*, 10–16.

Bollimunta, A., Totten, D., and Ditterich, J. (2012). Neural dynamics of choice: single-trial analysis of decision-related activity in parietal cortex. J Neurosci *32*, 12684–12701.

Brainard, D.H. (1997). The Psychophysics Toolbox. Spat. Vis. *10*, 433–436.

Braun, A., Urai, A.E., and Donner, T.H. (2018). Adaptive history biases result from confidence-weighted accumulation of past choices. J. Neurosci. *38*, 2418–2429.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1992). The analysis of visual motion: A comparison of neuronal and psychophysical performance. J. Neurosci. *12*, 4745–4765.

Britten, K.H., Shadlen, M.N., Newsome, W.T., and Movshon, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. Vis. Neurosci. *10*, 1157–1169.

Britten, K.H., Newsome, W.T., Shadlen, M.N., Movshon, J.A., and Celebrini, S. (1996). A relationship between behavioral choice and the visual responses of neurons in macaque MT. Vis. Neurosci. *13*, 87–100.

Bronfman, Z.Z., Brezis, N., and Usher, M. (2016). Non-monotonic Temporal-Weighting Indicates a Dynamically Modulated Evidence-Integration Mechanism. PLoS Comput. Biol. *12(2)*, e1004667.

Brosnan, M.B., Sabaroedin, K., Silk, T., Genc, S., Newman, D.P., Loughnane, G.M., Fornito, A., O'Connell, R.G., and Bellgrove, M.A. (2020). Evidence accumulation during perceptual decisions in humans varies as a function of dorsal frontoparietal organization. Nat. Hum. Behav. *4*, 844–855

Brown, S.D., and Heathcote, A. (2008). The simplest complete model of choice response time: Linear ballistic accumulation. Cogn. Psychol. *57*, 153–178.

Brunton, B.W., Botvinick, M.M., and Brody, C.D. (2013). Rats and Humans Can Optimally Accumulate Evidence for Decision-Making. Science *340*, 95–98.

Busemeyer, J.R., and Townsend, J.T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. Psychol. Rev. *100*, 432–459.

Cassidy, S.M., Robertson, I.H., and O'Connell, R.G. (2012). Retest reliability of event-related potentials: Evidence from a variety of paradigms. Psychophysiology *49*, 659–664.

Cavanagh, S.E., Lam, N.H., Murray, J.D., Hunt, L.T., and Kennerley, S.W. (2020). A circuit mechanism for decision-making biases and NMDA receptor hypofunction. Elife *9*: e53664

Chandrasekaran, C., Soldado-Magraner, J., Peixoto, D., Newsome, W.T., Shenoy, K. V., and Sahani, M. (2018). Brittleness in model selection analysis of single neuron firing rates. BioRxiv 430710.

Cheadle, S., Wyart, V., Tsetsos, K., Myers, N., deGardelle, V., HerceCastañón, S., and Summerfield, C. (2014). Adaptive gain control during human perceptual choice. Neuron *81*, 1429–1441.

Chen, Y., Nakayama, K., Levy, D., Matthysse, S., and Holzman, P. (2003). Processing of global, but not local, motion direction is deficient in schizophrenia. Schizophr. Res. *61*, 215–227.

Colby, C.L., and Goldberg, M.E. (1999). Space and attention in parietal cortex. Annu. Rev. Neurosci. *22*, 319–349.

Colby, C.L., Duhamel, J.R., and Goldberg, M.E. (1996). Visual, presaccadic, and cognitive activation of single neurons in monkey lateral intraparietal area. J. Neurophysiol. *76*, 2841–2852.

Cole, D.M., Diaconescu, A.O., Pfeiffer, U.J., Brodersen, K.H., Mathys, C.D., Julkowski, D., Ruhrmann, S., Schilbach, L., Tittgemeyer, M., Vogeley, K., et al. (2020). Atypical processing of uncertainty in individuals at risk for psychosis. NeuroImage Clin. *26*:102239

Corsi-Cabrera, M., Galindo-Vilchis, L., del-Río-Portilla, Y., Arce, C., and Ramos-Loyo, J. (2007). Within-subject reliability and inter-session stability of EEG power and coherent activity in women evaluated monthly over nine months. Clin. Neurophysiol. *118*, 9–21.

Cravo, A.M., Rohenkohl, G., Wyart, V., and Nobre, A.C. (2011). Endogenous modulation of low frequency oscillations by temporal expectations. J. Neurophysiol. *106*, 2964–2972.

Devine, C.A., Gaffney, C., Loughnane, G., Kelly, S.P., and O'Connell, R.G. (2019). The role of premature evidence accumulation in making difficult perceptual decisions under temporal uncertainty. Elife *8*: e48526

Ding, L., and Gold, J.I. (2012). Neural correlates of perceptual decision making before, during, and after decision commitment in monkey frontal eye field. Cereb. Cortex *22*, 1052–1067.

Dodd, J. V, Krug, K., Cumming, B.G., and Parker, A.J. (2001). Perceptually bistable three-dimensional figures evoke high choice probabilities in cortical area MT. J. Neurosci. *21*, 4809–4821.

Donchin, E., and Coles, M.G.H. (1988). Is the P300 component a manifestation of context updating? Behav. Brain Sci. *11*, 357–374.

Donner, T.H., Siegel, M., Fries, P., and Engel, A.K. (2009). Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. Curr. Biol. *19*, 1581–1585.

Drugowitsch, J., Moreno-Bote, R.N., Churchland, A.K., Shadlen, M.N., and Pouget, A. (2012). The cost of accumulating evidence in perceptual decision making. J. Neurosci. *32*, 3612–3628.

Drugowitsch, J., Mendonça, A.G., Mainen, Z.F., and Pouget, A. (2019). Learning optimal decisions with confidence. Proc. Natl. Acad. Sci. U. S. A. *116*, 24872–24880.

Duncan-Johnson, C.C., and Donchin, E. (1977). On Quantifying Surprise: The Variation of Event-Related Potentials With Subjective Probability. Psychophysiology *14*, 456–467.

Ehinger, B. V., and Dimigen, O. (2019). Unfold: An integrated toolbox for overlap correction, non-linear modeling, and regression-based EEG analysis. PeerJ *2019*, e7838.

Erlich, J.C., Brunton, B.W., Duan, C.A., Hanks, T.D., and Brody, C.D. (2015). Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. Elife *4*: e05457

Faisal, A.A., Selen, L.P.J., and Wolpert, D.M. (2008). Noise in the nervous system. Nat. Rev. Neurosci. *9*, 292–303.

FitzGerald, T.H.B., Moran, R.J., Friston, K.J., and Dolan, R.J. (2015). Precision and neuronal dynamics in the human posterior parietal cortex during evidence accumulation. Neuroimage *107*, 219–228.

Fontanesi, L., Gluth, S., Spektor, M.S., and Rieskamp, J. (2019a). A reinforcement learning diffusion decision model for value-based decisions. Psychon. Bull. Rev. *26*, 1099–1121.

Fontanesi, L., Palminteri, S., and Lebreton, M. (2019b). Decomposing the effects of context valence and feedback information on speed and accuracy during reinforcement learning: a meta-analytical approach using diffusion decision modeling. Cogn. Affect. Behav. Neurosci. *19*, 490–502.

Forstmann, B.U., Dutilh, G., Brown, S., Neumann, J., von Cramon, D.Y., Ridderinkhof, K.R., and Wagenmakers, E.-J. (2008). Striatum and pre-SMA facilitate decision-making under time pressure. Proc. Natl. Acad. Sci. *105*, 17538–17542.

Frank, M.J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J.F., and Badre, D. (2015). FMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. J. Neurosci. *35*, 485–494.

Fritsche, M., Mostert, P., and de Lange, F.P. (2017). Opposite Effects of Recent History on Perception and Decision. Curr. Biol. *27*, 590–595.

Ganupuru, P., Goldring, A.B., Harun, R., and Hanks, T.D. (2019). Flexibility of Timescales of Evidence Evaluation for Decision Making. Curr. Biol. *29*, 2091-2097.e4.

Gao, J., Tortell, R., and McClelland, J.L. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. PLoS One *6*, 16749.

Glaze, C.M., Kable, J.W., and Gold, J.I. (2015). Normative evidence accumulation in unpredictable environments. Elife *4*:e08825

Glaze, C.M., Filipowicz, A.L.S., Kable, J.W., Balasubramanian, V., and Gold, J.I. (2018). A bias-variance trade-off governs individual differences in on-line learning in an unpredictable environment. Nat. Hum. Behav. *2*, 213–224.

Gold, J.I., and Shadlen, M.N. (2007). The neural basis of decision making. Annu. Rev. Neurosci. *30*, 535–574.

Gonçalves, N.R., Whelan, R., Foxe, J.J., and Lalor, E.C. (2014). Towards obtaining spatiotemporally precise responses to continuous sensory stimuli in humans: A general linear modeling approach to EEG. Neuroimage *97*, 196–205.

Gottlieb, J.P., Kusunoki, M., and Goldberg, M.E. (1998). The representation of visual salience in monkey parietal cortex. Nature *391*, 481–484.

Grahek, I., Everaert, J., Krebs, R.M., and Koster, E.H.W. (2018). Cognitive Control in Depression: Toward Clinical Models Informed by Cognitive Neuroscience. Clin. Psychol. Sci. *6*, 464–480.

Green, D., and Swets, J. (1966). Signal detection theory and psychophysics. Wiley.

Hanks, T.D., and Summerfield, C. (2017). Perceptual Decision Making in Rodents, Monkeys, and Humans. Neuron *93*, 15–31.

Hanks, T., Kiani, R., and Shadlen, M.N. (2014). A neural mechanism of speed-accuracy tradeoff in macaque area LIP. Elife *3*:e02260

Hanks, T.D., Ditterich, J., and Shadlen, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. Nat. Neurosci. *9*, 682–689.

Hanks, T.D., Mazurek, M.E., Kiani, R., Hopp, E., and Shadlen, M.N. (2011). Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task. J. Neurosci. *31*, 6339–6352.

Hanks, T.D., Kopec, C.D., Brunton, B.W., Duan, C.A., Erlich, J.C., and Brody, C.D. (2015). Distinct relationships of parietal and prefrontal cortices to evidence accumulation. Nature *520*, 220–223.

Harun, R., Jun, E., Park, H.H., Ganupuru, P., Goldring, A.B., and Hanks, T.D. (2020). Timescales of Evidence Evaluation for Decision Making and Associated Confidence Judgments Are Adapted to Task Demands. Front. Neurosci. *14*. 826

Heekeren, H.R., Marrett, S., Bandettini, P.A., and Ungerleider, L.G. (2004). A general mechanism for perceptual decision-making in the human brain. Nature *431*, 859–862.

Heitz, R.P., and Schall, J.D. (2012). Neural Mechanisms of Speed-Accuracy Tradeoff. Neuron *76*, 616–628.

Hillyard, S.A., Squires, K.C., Bauer, J.W., and Lindsay, P.H. (1971). Evoked potential correlates of auditory signal detection. Science *172*, 1357–1360.

Huk, A.C., and Shadlen, M.N. (2005). Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. J. Neurosci. *25*, 10420–10436.

Hunt, L.T., Kolling, N., Soltani, A., Woolrich, M.W., Rushworth, M.F.S., and Behrens, T.E.J. (2012). Mechanisms underlying cortical activity during value-guided choice. Nat. Neurosci. *15*, 470–476.

Ipata, A.E., Gee, A.L., Gottlieb, J., Bisley, J.W., and Goldberg, M.E. (2006). LIP responses to a popout stimulus are reduced if it is overtly ignored. Nat. Neurosci. *9*, 1071–1076.

Ipata, A.E., Gee, A.L., Bisley, J.W., and Goldberg, M.E. (2009). Neurons in the lateral intraparietal area create a priority map by the combination of disparate signals. Exp. Brain Res. *192*, 479–488.

Ivanoff, J., Branning, P., and Marois, R. (2008). fMRI Evidence for a Dual Process Account of the Speed-Accuracy Tradeoff in Decision-Making. PLoS One *3*, e2635.

Janssen, P., and Shadlen, M.N. (2005). A representation of the hazard rate of elapsed time in macaque area LIP. Nat. Neurosci. *8*, 234–241.

Katz, L.N., Yates, J.L., Pillow, J.W., and Huk, A.C. (2016). Dissociated functional significance of decision-related activity in the primate dorsal stream. Nature *535*, 285–288.

Kayser, A.S., Buchsbaum, B.R., Erickson, D.T., and D'Esposito, M. (2010). The Functional Anatomy of a Perceptual Decision in the Human Brain. J. Neurophysiol. *103*, 1179–1194.

Kehrer, C. (2008). Altered excitatory-inhibitory balance in the NMDA-hypofunction model of schizophrenia. Front. Mol. Neurosci. *1*. doi: 10.3389/neuro.02.006.2008

Kelly, S.P., and O'Connell, R.G. (2013). Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. J. Neurosci. *33*, 19434–19441.

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

Kiani, R., Hanks, T.D., and Shadlen, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the

environment. J Neurosci *28*, 3017–3029.

Kim, J.N., and Shadlen, M.N. (1999). Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. Nat. Neurosci. *2*, 176–185.

Kira, S., Yang, T., and Shadlen, M.N. (2015). A Neural Implementation of Wald's Sequential Probability Ratio Test. Neuron *85*, 861–873.

Körding, K.P., and Wolpert, D.M. (2004). Bayesian integration in sensorimotor learning. Nature *427*, 244–247.

Kusunoki, M., Gottlieb, J., and Goldberg, M.E. (2000). The lateral intraparietal area as a salience map: The representation of abrupt onset, stimulus motion, and task relevance. Vision Res. *40,* 10–12.

Lak, A., Nomoto, K., Keramati, M., Sakagami, M., and Kepecs, A. (2017). Midbrain Dopamine Neurons Signal Belief in Choice Accuracy during a Perceptual Decision. Curr. Biol. *27*, 821–832.

Lak, A., Okun, M., Moss, M.M., Gurnani, H., Farrell, K., Wells, M.J., Reddy, C.B., Kepecs, A., Harris, K.D., and Carandini, M. (2020a). Dopaminergic and Prefrontal Basis of Learning from Sensory Confidence and Reward Value. Neuron *105*, 700-711

Lak, A., Hueske, E., Hirokawa, J., Masset, P., Ott, T., Urai, A.E., Donner, T.H., Carandini, M., Tonegawa, S., Uchida, N., et al. (2020b). Reinforcement biases subsequent perceptual decisions when confidence is low: A widespread behavioral phenomenon. Elife *9*:e49834

Lalor, E.C., Pearlmutter, B.A., Reilly, R.B., McDarby, G., and Foxe, J.J. (2006). The VESPA: A method for the rapid estimation of a visual evoked potential. Neuroimage *32*, 1549–1561.

Lam, N.H., Borduqui, T., Hallak, J., Roque, A.C., Anticevic, A., Krystal, J.H., Wang, X., and Murray, J.D. (2017). Effects of Altered Excitation-Inhibition Balance on Decision Making in a Cortical Circuit Model. BioRxiv. 100347

de Lange, F.P., Rahnev, D.A., Donner, T.H., and Lau, H. (2013). Prestimulus oscillatory activity over motor cortex reflects perceptual expectations. J. Neurosci. *33*, 1400–1410.

De Lange, F.P., Jensen, O., and Dehaene, S. (2010). Accumulation of evidence during sequential decision making: The importance of top-down factors. J. Neurosci. *30*, 731–738.

Latimer, K.W., Yates, J.L., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2015). Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. Science *349*, 184–187.

Latimer, K.W., Huk, A.C., and Pillow, J.W. (2017). No cause for pause: New analyses of ramping and stepping dynamics in LIP (Rebuttal to Response to Reply to Comment on Latimer et al 2015). BioRxiv 160994.

Lee, D., Seo, H., and Jung, M.W. (2012). Neural Basis of Reinforcement Learning and Decision Making. Annu. Rev. Neurosci. *35*, 287–308.

Levi, A.J., Yates, J.L., Huk, A.C., and Katz, L.N. (2018). Strategic and dynamic temporal weighting for perceptual decisions in humans and macaques. ENeuro *5(5)*:ENEURO.0169-18.2018

Lewis, J.W., and Van Essen, D.C. (2000). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. J. Comp. Neurol. *428*, 112–137.

Liu, T., and Pleskac, T.J. (2011). Neural correlates of evidence accumulation in a perceptual decision task. J. Neurophysiol. *106*, 2383–2398.

Loughnane, G.M., Newman, D.P., Bellgrove, M.A., Lalor, E.C., Kelly, S.P., and O'Connell, R.G. (2016). Target selection signals influence perceptual decisions by modulating the onset and rate of evidence accumulation. Curr. Biol. *26*, 496–502.

Luck, S.J. (2014). An Introduction to the Event-Related Potential Technique, 2nd Edition.

Macar, F., and Vidal, F. (2004). Event-Related Potentials as Indices of Time Processing: A Review. J. Psychophysiol. *18*, 89–104.

Mars, R.B., Debener, S., Gladwin, T.E., Harrison, L.M., Haggard, P., Rothwell, J.C., and Bestmann, S. (2008). Trial-by-trial fluctuations in the event-related electroencephalogram reflect dynamic changes in the degree of surprise. J. Neurosci. *28*, 12539–12545.

Mazurek, M.E., Roitman, J.D., Ditterich, J., and Shadlen, M.N. (2003). A Role for Neural Integrators in Perceptual Decision Making. Cereb. Cortex *13*, 1257–1269.

Mazzoni, P., Bracewell, R.M., Barash, S., and Andersen, R.A. (1996). Motor intention activity in the macaque's lateral intraparietal area. I. Dissociation of motor plan from sensory memory. J. Neurophysiol. *76(3)*, 1439–56

McEvoy, L.K., Smith, M.E., and Gevins, A. (2000). Test-retest reliability of cognitive EEG. Clin. Neurophysiol. *111*, 457–463.

McGovern, D.P., Hayes, A., Kelly, S.P., and O'Connell, R.G. (2018). Reconciling age-related changes in behavioural and neural indices of human perceptual decision-making. Nat. Hum. Behav. *2*, 955–966.

McTeague, L.M., Goodkind, M.S., and Etkin, A. (2016). Transdiagnostic impairment of cognitive control in mental illness. J. Psychiatr. Res. *83*, 37–46.

Meister, M.L.R., Hennig, J.A., and Huk, A.C. (2013). Signal Multiplexing and Single-Neuron Computations in Lateral Intraparietal Area During Decision-Making. J. Neurosci. *33*, 2254–2267.

Miletić, S., Boag, R.J., and Forstmann, B.U. (2020). Mutual benefits: Combining reinforcement learning with sequential sampling models. Neuropsychologia *136*, 107261.

Milosavljevic, M., Malmaud, J., Huth, A., Koch, C., and Rangel, A. (2010). The Drift Diffusion Model can account for the accuracy and reaction time of value-based choices under high and low time pressure. Judgm. Decis. Mak. *5*, 437–449.

Mor, N., and Daches, S. (2015). Ruminative thinking: Lessons learned from cognitive training. Clin. Psychol. Sci. *3*, 574–592.

Movshon, J.A., and Newsome, W.T. (1992). Neural Foundations of Visual Motion Perception. Curr. Dir. Psychol. Sci. *1*, 35–39.

Mulder, M.J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., and Forstmann, B.U.

(2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. J. Neurosci. *32*, 2335–2343.

Murphy, P.R., Boonstra, E., and Nieuwenhuis, S. (2016). Global gain modulation generates time-dependent urgency during perceptual choice in humans. Nat. Commun. *7*, 13526

Murphy, P.R., Wilming, N., Hernandez-Bocanegra, D.C., Ortega, G.P., and Donner, T.H. (2020). Normative circuit dynamics across human cortex during evidence accumulation in changing environments. BioRxiv 2020.01.29.924795

Murray, J.D., Bernacchia, A., Freedman, D.J., Romo, R., Wallis, J.D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., et al. (2014a). A hierarchy of intrinsic timescales across primate cortex. Nat. Neurosci. *17*, 1661–1663.

Murray, J.D., Anticevic, A., Gancsos, M., Ichinose, M., Corlett, P.R., Krystal, J.H., and Wang, X.J. (2014b). Linking microcircuit dysfunction to cognitive impairment: Effects of disinhibition associated with schizophrenia in a cortical working memory model. Cereb. Cortex *24*, 859–872.

Nakazawa, K., Zsiros, V., Jiang, Z., Nakao, K., Kolata, S., Zhang, S., and Belforte, J.E. (2012). GABAergic interneuron origin of schizophrenia pathophysiology. Neuropharmacology *62*, 1574–1583.

Neri, P., and Heeger, D.J. (2002). Spatiotemporal mechanisms for detecting and identifying image features in human vision. Nat. Neurosci. *5*, 812–816.

Neri, P., Parker, A.J., and Blakemore, C. (1999). Probing the human stereoscopic system with reverse correlation. Nature *401*, 695–698.

Newsome, W.T., and Pare, E.B. (1988). A Selective Impairment of Motion Perception Following Lesions of the Middle Temporal Visual Area (MT). J Neurosci. 8(6), 2201-11

Nobre, A., Correa, A., and Coull, J. (2007). The hazards of time. Curr. Opin. Neurobiol. *17*, 465–470.

Nobre, A.C., Gitelman, D.R., Dias, E.C., and Mesulam, M.M. (2000). Covert visual spatial orienting and saccades: Overlapping neural systems. Neuroimage *11*, 210–216.

Nolte, G., and Hämäläinen, M.S. (2001). Partial signal space projection for artefact removal in MEG measurements: A theoretical analysis. Phys. Med. Biol. *46*, 2873–2887.

O'Connell, R.G., Dockree, P.M., and Kelly, S.P. (2012). A supramodal accumulation-to-bound signal that determines perceptual decisions in humans. Nat. Neurosci. *15*, 1729.

O'Reilly, J.X., Jbabdi, S., and Behrens, T.E.J. (2012). How can a Bayesian approach inform neuroscience? Eur. J. Neurosci. *35*, 1169–1179.

O'Reilly, J.X., Schüffelgen, U., Cuell, S.F., Behrens, T.E.J., Mars, R.B., and Rushworth, M.F.S. (2013). Dissociable effects of surprise and model update in parietal and anterior cingulate cortex. Proc. Natl. Acad. Sci. U. S. A. *110*(38): E3660-9

Okazawa, G., Sha, L., Purcell, B.A., and Kiani, R. (2018). Psychophysical reverse correlation reflects both sensory and decision-making processes. Nat.

Commun. *9*. 3479

Orsolic, I., Rio, M., Mrsic-Flogel, T.D., and Znamenskiy, P. (2019). Mesoscale cortical dynamics reflect the interaction of sensory evidence and temporal expectation during perceptual decision-making. BioRxiv.552026

Ossmy, O., Moran, R., Pfeffer, T., Tsetsos, K., Usher, M., and Donner, T.H. (2013). The timescale of perceptual evidence integration can be adapted to the environment. Curr. Biol. *23*, 981–986.

Palmer, J., Huk, A.C., and Shadlen, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. J. Vis. *5*, 376–404.

Park, I.M., Meister, M.L.R., Huk, A.C., and Pillow, J.W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. Nat. Neurosci. *17*, 1395–1403.

Pedersen, M.L., Frank, M.J., and Biele, G. (2017). The drift diffusion model as the choice rule in reinforcement learning. Psychon. Bull. Rev. *24*, 1234–1251.

Pfeuty, M., Ragot, R., and Pouthas, V. (2005). Relationship between CNV and timing of an upcoming event. Neurosci. Lett. *382*, 106–111.

Philiastides, M.G., Auksztulewicz, R., Heekeren, H.R., and Blankenburg, F. (2011). Causal role of dorsolateral prefrontal cortex in human perceptual decision making. Curr. Biol. *21*, 980–983.

Piet, A.T., El Hady, A., and Brody, C.D. (2018). Rats adopt the optimal timescale for evidence integration in a dynamic environment. Nat. Commun. *9*. 4265

Pleger, B., Ruff, C.C., Blankenburg, F., Bestmann, S., Wiech, K., Stephan, K.E., Capilla, A., Friston, K.J., and Dolan, R.J. (2006). Neural coding of tactile decisions in the human prefrontal cortex. J. Neurosci. *26*, 12596–12601.

Polich, J., and Kok, A. (1995). Cognitive and biological determinants of P300: an integrative review. Biol. Psychol. *41*, 103–146.

Ratcliff, R. (1978). A Theory of Memory Retrieval. Psychol. Rev. *85*, 59–108.

Ratcliff, R., and McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. Neural Comput. *20*, 873–922.

Ratcliff, R., and Rouder, J.N. (1998). Modeling Response Times for Two-Choice Decisions. Psychol. Sci. *9*, 347–356.

Ratcliff, R., and Smith, P.L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. Psychol. Rev. *111*, 333–367.

Ratcliff, R., Hasegawa, Y.T., Hasegawa, R.P., Smith, P.L., and Segraves, M.A. (2007). Dual diffusion model for single-cell recording data from the superior colliculus in a brightness-discrimination task. J. Neurophysiol. *97*, 1756–1774.

Reddi, B.A.J., Asrress, K.N., and Carpenter, R.H.S. (2003). Accuracy, Information, and Response Time in a Saccadic Decision Task. J. Neurophysiol. *90*, 3538–3546.

Rektor, I. (2002). Scalp-recorded Bereitschaftspotential is the result of the activity of cortical and subcortical generators - A hypothesis. Clin. Neurophysiol. *113*, 1998–2005.

van Rijn, H., Kononowicz, T.W., Meck, W.H., Ng, K.K., and Penney, T.B. (2011).

Contingent negative variation and its relation to time estimation: A theoretical evaluation. Front. Integr. Neurosci. *5*, 91.

Roe, R.M., Busemeyer, J.R., and Townsend, J.T. (2001). Multialternative decision field theory: A dynamic connectionist model of decision making. Psychol. Rev. *108*, 370–392.

Roitman, J.D., and Shadlen, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. J Neurosci *22(21)*, 9475–9989.

Rorie, A.E., Gao, J., McClelland, J.L., and 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 19,*5*(2):e9308

Salvador, A., Arnal, L., Vinckier, F., Domenech, P., Gaillard, R., and Wyart, V. (2020). Premature commitment to uncertain beliefs during human NMDA receptor hypofunction. BioRxiv 2020.06.17.156539.

Salzman, C.D., Murasugi, C.M., Britten, K.H., and Newsome, W.T. (1992). Microstimulation in Visual Area MT: Effects on Direction Discrimination Performance. J Neurosci. 12(6), 2331-55.

Serences, J.T., and Boynton, G.M. (2007). The representation of behavioral choice for motion in human visual cortex. J. Neurosci. *27*, 12893–12899.

Seriès, P., and Seitz, A.R. (2013). Learning what to expect (in visual perception). Front. Hum. Neurosci. *7*, 668.

Shadlen, M., and Newsome, W. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. J. Neurophysiol. 86(4), 1916–1936.

Shadlen, M.N., Kiani, R., Newsome, W.T., Gold, J.I., Wolpert, D.M., Zylberberg, A., Ditterich, J., De Lafuente, V., Yang, T., and Roitman, J. (2016). Comment on “single-trial spike trains in parietal cortex reveal discrete steps during decision-making.” Science (80-.). *351*, 1406–1406.

Shinn, M., Ehrlich, D., Lee, D., Murray, J.D., and Seo, H. (2019). Confluence of timing and reward biases in perceptual decision-making dynamics. BioRxiv 865501

Siegel, M., Engel, A.K., and Donner, T.H. (2011). Cortical network dynamics of perceptual decision-making in the human brain. Front. Hum. Neurosci. *5*, 1–12.

Siegel, M., Buschman, T.J., and Miller, E.K. (2015). Cortical information flow during flexible sensorimotor decisions. Science *348*, 1352–1355.

Smith, N.J., and Kutas, M. (2015a). Regression-based estimation of ERP waveforms: I. The rERP framework. Psychophysiology *52*, 157–168.

Smith, N.J., and Kutas, M. (2015b). Regression-based estimation of ERP waveforms: II. Nonlinear effects, overlap correction, and practical considerations. Psychophysiology *52*, 169–181.

Snyder, H.R., Miyake, A., and Hankin, B.L. (2015). Advancing understanding of executive function impairments and psychopathology: Bridging the gap between clinical and cognitive approaches. Front. Psychol. *6*, 328.

Snyder, L.H., Batista, A.P., and Andersen, R.A. (2000). Intention-related activity

in the posterior parietal cortex: A review. In Vision Res. *40(10-12)* 1433–1441.

Squires, K.C., Hillyard, S.A., and Lindsay, P.H. (1973). Vertex potentials evoked during auditory signal detection: Relation to decision criteria. Percept. Psychophys. *14*, 265–272.

Starc, M., Murray, J.D., Santamauro, N., Savic, A., Diehl, C., Cho, Y.T., Srihari, V., Morgan, P.T., Krystal, J.H., Wang, X.J., et al. (2017). Schizophrenia is associated with a pattern of spatial working memory deficits consistent with cortical disinhibition. Schizophr. Res. *181*, 107–116.

Steinemann, N.A., O'Connell, R.G., and Kelly, S.P. (2018). Decisions are expedited through multiple neural adjustments spanning the sensorimotor hierarchy. Nat. Commun. *9,* 3627.

Stine, G.M., Zylberberg, A., Ditterich, J., and Shadlen, M.N. (2020). Differentiating between integration and non-integration strategies in perceptual decision making. BioRxiv 2020.01.24.918169.

Summerfield, C., and Koechlin, E. (2008). A Neural Representation of Prior Information during Perceptual Inference. Neuron *59*, 336–347.

Summerfield, C., and De Lange, F.P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. Nat. Rev. Neurosci. *15*, 745–756.

Tosoni, A., Galati, G., Romani, G.L., and Corbetta, M. (2008). Sensory-motor mechanisms in human parietal cortex underlie arbitrary visual decisions. Nat. Neurosci. *11*, 1446–1453.

Trillenberg, P., Verleger, R., Wascher, E., Wauschkuhn, B., and Wessel, K. (2000). CNV and temporal uncertainty with "ageing" and "non-ageing" S1-S2 intervals. Clin. Neurophysiol. *111*, 1216–1226.

Tsetsos, K., Usher, M., and McClelland, J.L. (2011). Testing Multi-Alternative Decision Models with Non-Stationary Evidence. Front. Neurosci. *5*, 63.

Tsetsos, K., Gao, J., McClelland, J.L., and Usher, M. (2012). Using Time-Varying Evidence to Test Models of Decision Dynamics: Bounded Diffusion vs. the Leaky Competing Accumulator Model. Front. Neurosci. *6*, 79.

Tuerlinckx, F., and De Boeck, P. (2005). Two interpretations of the discrimination parameter. Psychometrika *70*, 629–650.

Twomey, D.M., Murphy, P.R., Kelly, S.P., and O'Connell, R.G. (2015). The classic P300 encodes a build-to-threshold decision variable. Eur. J. Neurosci. *42*, 1636–1643.

Twomey, D.M., Kelly, S.P., and O'Connell, R.G. (2016). Abstract and effector-selective decision signals exhibit qualitatively distinct dynamics before delayed perceptual reports. J. Neurosci. *36*, 7346–7352.

Ungerleider, S.K. and L.G. (2000). Mechanisms of Visual Attention in the Human Cortex. Annu. Rev. Neurosci. *23*, 315–341.

Urai, A.E., de Gee, J.W., Tsetsos, K., and Donner, T.H. (2018). Choice History Biases Subsequent Evidence Accumulation. BioRxiv 251595.

Usher, M., and McClelland, J.L. (2001). The Time Course of Perceptual Choice: The Leaky, Competing Accumulator Model. Psychol. Rev. *108*, 552–592.

Usher, M., and McClelland, J.L. (2004). Loss aversion and inhibition in dynamical models of multialternative choice. Psychol. Rev. *111*, 757–769.

Van Veen, V., Krug, M.K., and Carter, C.S. (2008). The neural and computational basis of controlled speed-accuracy tradeoff during task performance. J. Cogn. Neurosci. *20*, 1952–1965.

Veliz-Cuba, A., Kilpatrick, Z.P., and Josić, K. (2016). Stochastic models of evidence accumulation in changing environments. SIAM Rev. *58*, 264–289.

Vilares, I., Howard, J.D., Fernandes, H.L., Gottfried, J.A., and Kording, K.P. (2012). Differential representations of prior and likelihood uncertainty in the human brain. Curr. Biol. *22*, 1641–1648.

Wald, A. (1945). Sequential Tests of Statistical Hypotheses. Ann. Math. Stat. *2*, 117–186.

Wang, X.-J. (2002). Probabilistic decision making by slow reverrberation in cortical circuits. Neuron *36*, 955–968.

Wardak, C., Olivier, E., and Duhamel, J.-R. (2002). Saccadic Target Selection Deficits after Lateral Intraparietal Area Inactivation in Monkeys. J. Neurosci. *22*. 9877–9884

Williams, L.M., Simms, E., Clark, C.R., Paul, R.H., Rowe, D., and Gordon, E. (2005). The test-retest reliability of a standardized neurocognitive and neurophysiological test battery: “Neuromarker.” Int. J. Neurosci. *115*, 1605–1630.

Wyart, V., de Gardelle, V., Scholl, J., and Summerfield, C. (2012). Rhythmic Fluctuations in Evidence Accumulation during Decision Making in the Human Brain. Neuron *76*, 847–858.

Wyart, V., Myers, N.E., and Summerfield, C. (2015). Neural mechanisms of human perceptual choice under focused and divided attention. J. Neurosci. *35*, 3485–3498.

Yang, T., and Shadlen, M.N. (2007). Probabilistic reasoning by neurons. Nature *447*, 1075–1080.

Yates, J.L., Park, I.M., Katz, L.N., Pillow, J.W., and Huk, A.C. (2017). Functional dissection of signal and noise in MT and LIP during decision-making. Nat Neurosci. *20(9),* 1285–1292

Yu, A.J., and Cohen, J.D. (2008). Sequential effects: Superstition or rational behavior? Adv. Neural Inf. Process. Syst. *21*, 1873–1880.

Zhang, S., Huang, C.H., and Yu, A.J. (2014). Sequential effects : A Bayesian analysis of prior bias on reaction time and behavioral choice. Proc. Cogn. Sci. Soc. Conf. 36, 1844–1849.

Zhao, X., and Kording, K.P. (2018). Rate fluctuations not steps dominate LIP activity during decision-making. BioRxiv. 249672

Zoltowski, D.M., Latimer, K.W., Yates, J.L., Huk, A.C., and Pillow, J.W. (2019). Discrete Stepping and Nonlinear Ramping Dynamics Underlie Spiking Responses of LIP Neurons during Decision-Making. Neuron *102*, 1249-1258.e10.

Zylberberg, A., and Shadlen, M. (2016). Cause for pause before leaping to conclusions about stepping. BioRxiv 085886.

Zylberberg, A., Wolpert, D.M., and Shadlen, M.N. (2018). Counterfactual

Reasoning Underlies the Learning of Priors in Decision Making. Neuron *99*, 1083-1097.e6.