The Dynamic Interplay of External and Internal Attention





Daniela Gresch

Worcester College

University of Oxford

A thesis submitted in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Trinity Term 2024

Abstract

During natural behaviour, our attention is in constant flux, seamlessly transitioning between information available in the external environment and internal representations stored in working memory. However, as past research has primarily investigated external and internal attention in isolation, relatively little is known regarding the dynamic interplay between these two attentional domains. In this doctoral thesis, I explore scenarios where individuals encounter both external and internal information in quick succession, necessitating rapid shifts between perception and working memory. The experimental work in this thesis can be divided into two main branches. The first branch explores cross-domain attentional modulations; that is, how attention within one domain influences attention within the alternative domain. The second branch takes a deeper dive into the intricacies of shifting attention between domains. To provide an overview of previous contributions, Chapter 1 reviews past studies on external and internal attention, both as independent facets of attention and as interdependent phenomena. Building on this, Chapter 2 investigates the behavioural consequences of attentional shifts and examines how these can be modulated. Chapter 3 explores between-domain shifts by employing neural measures to examine the timing of reactivating internal representations following engagement in an external task. To understand the overarching nature of between-domain shifts, Chapters 4 and 5 introduce a novel, combined perception and working-memory task that allows within- and between-domain shifts to be contrasted in each respective domain. While Chapter 4 focusses on the behavioural correlates of between-domain shifts, Chapter 5 investigates the neural signatures associated with these transitions. Finally, in Chapter 6, I discuss the implications of my results and suggest potential avenues for future research. The findings of my doctoral research reveal that the dynamic interplay between attentional domains imposes behavioural costs; however, these costs are not immutable and can be influenced by various modulatory factors from multiple sources. Further, I demonstrate that prompt (but not always complete) shifts between attentional domains can be triggered by various events. Taken together, this thesis advocates for a holistic approach to examining external and internal attention.

Acknowledgements

With the submission of this thesis, my time in Oxford also slowly comes to an end. Reflecting on the last three and one-half years, I feel incredibly fortunate to have been surrounded by such a remarkable network of individuals. Their presence and support have been invaluable, transforming what could have been a daunting experience into one filled with growth and memorable moments. Although I am sure words cannot fully capture my gratitude, I hope to express how deeply thankful I am to those who have left an indelible mark on both my professional path and personal life.

First and foremost, I want to express my gratitude to my supervisor team: Kia Nobre, Freek van Ede, and Sage Boettcher. Kia, through your guidance, I have learned to embrace confusion as an opportunity for growth, to question the status quo, and to understand that a DPhil journey is about daring to take risks. Each meeting with you ignites a fresh spark of inspiration within me, and I feel immensely honoured to have the opportunity to work with and continue learning from you beyond the completion of my doctorate. Besides your undeniable scientific brilliance, I truly value the vibrant atmosphere you have cultivated in the lab, which is anything but dull and impersonal. Freek, despite you not being in Oxford, it always felt like you were just around the corner, ever ready to tackle any questions I had. I am truly impressed by your efficiency and ability to swiftly decipher the story hidden within the data (and your enthusiasm for jumping into any swimming pool at the slightest opportunity). I am very grateful for the chance to have been part of the Proactive Brain Lab, even from afar. Thank you for providing me the opportunity to learn from you the art of “nailing it”. Sage, you have been a constant inspiration since the day we first met nearly 10 years ago. Little did I know then that our initial friendship as master’s and bachelor’s students would evolve into a supervisor-DPhil-student relationship (and other adventures, such as being trapped in the Laotian jungle for 24 hours after a bus breakdown). Your scientific advice and positivity have helped me navigate through moments of uncertainty, bringing me back on track. You have always made time to sit down with me, debrief after meetings, guide me through analyses, or brainstorm about the next 200 experiments we will definitely run together. With you, science becomes playful and effortless, something you start looking forward to doing on a Sunday afternoon. I owe you so much and can say without a doubt that I would not be here without guidance throughout the years. Thank you for being not only exceptional role model but also a caring friend.

In addition to my supervisors, I would like to extend my gratitude to past and present B&C members who have shaped my time in Oxford. I am especially thankful for the lucky "coincidence" that brought me to join a lab that also serves as the workplace for two of my close friends. Thank you, Oli, for your companionship since our master's program in Munich began, for the thought-provoking conversations, and for the thesis-writing-break runs that made the past few months so much more enjoyable. Thank you, Dejan, for being a fantastic mentor and an incredibly generous friend. The internship under your guidance during my first year of bachelor's studies in Frankfurt has opened the door to incredible opportunities, for which I will always be immensely grateful. I would also like to thank my colleagues and friends Irene, Gwen, Levi, Giedre, Jasper, Lauren, Nahid, Katharina, Sarah, and Anna from the B&C Lab. Working at OHBA and NRH has been an enjoyable experience, thanks to you. I have truly appreciated your willingness to share your skills, insights, and helpful suggestions, all of which have been incredibly valuable to me. Moreover, I would like to express my thanks to Rose from the Proactive Brain Lab, who despite physical distance has always been a supportive colleague and has become a great friend. To cut a long story short, working with all of you has shown me how well research and friendship can go together.

I have been fortunate to make some lovely friends in Oxford. Thank you, Helen, Claire, Ryan, and Amy, for all the fun times. From home, I want to express my gratitude to Aylin. For discovering together that research is actually quite cool, for the intense climbing sessions, and for always being there for me. Even though we do not see each other often these days, whenever we do, it feels as if nothing has changed.

I feel deeply indebted to my family. Thanks to my mother Heike, father Thomas, sister Melanie, grandmother Edith, and godfather Dietmar for your unstoppable support and belief in me. I am extremely grateful for the freedom you have given me to explore my own path – without you, I would not be where I am today. Thank you.

Finally, I owe an immeasurable thanks to my best friend and partner Gordon (a.k.a. Friendo), who has been my consistent pillar of support, both personally and academically. I never thought I would enjoy being with a fellow scientist, let alone marrying one, but you provided some compelling reasons to convince me otherwise. Thank you for singing your silly little songs to me when I feel stressed, for introducing me to the world of road cycling and videogaming, and for being my personal encyclopaedia. I love you.

Statement of Contribution

The work presented within this doctoral thesis has greatly benefitted from the guidance provided by my DPhil supervisors: Sage Boettcher, Freek van Ede, and Kia Nobre. The experiments presented in each of the empirical chapters were designed in collaboration with my supervisory team. I collected the data for the experiments presented in Chapter 2 and 4. The EEG data for Chapter 3 were collected together with Larissa Behnke, while the MEG data for Chapter 5 were collected with the assistance of Anna Camera and Sven Braeutigam. I analysed the data presented in all empirical chapters. Some of the work has already been published or submitted to a peer-reviewed journal. I have generated the figures and authored the text for these manuscripts, with guidance from my supervisory team.

Chapters relating to a previous publication or preprint can be found here:

- Chapter 2.3:
  Gresch et al. (2022) – https://doi.org/10.1016/j.cognition.2022.105156

- Chapter 2.4:
  Gresch et al. (2021) – https://doi.org/10.1016/j.cognition.2021.104915

- Chapter 4.3 and 4.4:
  Gresch et al. (2024b) – https://doi.org/10.1016/j.cognition.2024.105731

- Chapter 5:
  Gresch et al. (2024a) – https://doi.org/10.1101/2024.02.14.580280

In the thesis, I provide transparent acknowledgement of any previously published or submitted materials, clearly indicating the corresponding chapters where applicable.

My co-authors (i.e., DPhil supervisors) have confirmed that I have permission to include published/submitted materials in this doctoral thesis.



Sage Boettcher

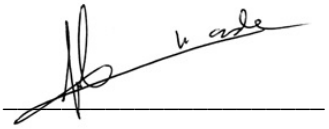

Freek van Ede



Kia Nobre

Table of Contents

List of Abbreviations xi

1 General Introduction 1

1.1 Towards a Dynamic Framework of External and Internal Attention 3

1.1.1 External Attention 4

1.1.2 Internal Attention as Working Memory 6

1.1.3 Similar but Not the Same 8

1.1.4 From a Comparative to a Dynamic Study 13

1.2 Cross-Domain Attentional Modulations 15

1.2.1 Internal-to-External Attentional Modulations 15

1.2.2 External-to-Internal Attentional Modulations 18

1.2.3 Protecting From vs. Mitigating External-to-Internal Modulations 22

1.2.3.1 Protecting Internal Representations During Maintenance 23

1.2.3.2 Mitigating External Interference During Maintenance 26

1.3 Shifting Between Attentional Domains 29

1.3.1 An Updated Framework and the Big Questions 29

1.3.2 The Little We Know From Behaviour 32

1.3.3 The Very Little We Know From the Brain 39

1.3.4 Recommendations for Studying Between-Domain Shifts 45

1.4 Overview and Aims of the Thesis 47

2 Cross-Domain Modulations via Temporal Expectations 50

2.1 Abstract 50

2.2 Introduction 51

2.3 Experiment 1: How Temporally Guided Internal Attention Modulates External Attention 54

2.3.1 Methods 55

2.3.1.1 Participants 55

2.3.1.2 Task and Procedure 56

2.3.1.3 Analysis 60

2.3.2 Results 63

2.3.2.1 Working-Memory Performance 63

2.3.2.2 External-Task Performance 64

2.3.2.3 External-Task Performance: Sequential Effects 65

2.3.3 Interim Discussion 67

2.3.3.1 Cross-Domain Consequences of Predictable Temporal Structures 67

2.3.3.2 Short-Term Cross-Domain Consequences....69
2.3.3.3 The Contribution of Task Relevance....70
2.3.3.4 Temporal-Expectations Benefit in Working Memory....71
2.3.3.5 Conclusion....74
2.4 Experiment 2: How Temporally Guided External Attention Modulates Internal Attention....75
2.4.1 Methods....78
2.4.1.1 Participants....78
2.4.1.2 Task and Procedure....78
2.4.1.3 Analysis....84
2.4.2 Results....86
2.4.2.1 Working-Memory Performance: No Interference vs. Interference....86
2.4.2.2 Three Possible Scenarios....87
2.4.2.3 Working-Memory Performance: Fixed vs. Variable....88
2.4.2.4 Working-Memory Performance: Suppression vs. Shielding....89
2.4.2.5 Working-Memory Performance: Shielding vs. Reduced Interference-Encoding....90
2.4.2.6 External-Task Performance....92
2.4.3 Interim Discussion....94
2.4.3.1 The Detrimental Effects of Interference....94
2.4.3.2 Working-Memory Protection via Shielding....96
2.4.3.3 Possible Neural Mechanisms Subserving Working-Memory Shielding....98
2.4.3.4 Ruling Out Alternative Interpretations....99
2.4.3.5 Conclusion....100
2.5 Discussion....100
3 Reactivating Internal Contents During External-to-Internal Shifts....103
3.1 Abstract....103
3.2 Introduction....104
3.3 Methods....106
3.3.1 Participants....106
3.3.2 Task and Procedure....106
3.3.3 Apparatus and Data Acquisition....111
3.3.4 Behavioural Analysis....112
3.3.5 EEG Pre-Processing....113
3.3.6 Time-Frequency Analysis....113
3.3.7 Relationship Between EEG and Behaviour....114
3.3.8 Statistical Analysis....115
3.4 Results....115

3.4.1 Behaviour: Working-Memory Performance 115
3.4.2 Behaviour: External-Task Performance 117
3.4.3 EEG Time-Frequency: Internal Selection During No-External-Task Trials 118
3.4.4 Two Possible Scenarios of Internal Reselection 120
3.4.5 EEG Time-Frequency: Internal (Re)Selection During External-Task Trials 122
3.4.6 EEG Time-Frequency: Internal Reselection After External-Task Responses 124
3.4.7 EEG Time-Frequency: Timing of Visual Reselection Is Behaviourally Relevant 126
3.5 Discussion 127
3.5.1 Prioritising Visual and Motor Contents 128
3.5.2 Immediate Reselection of Visual and Motor Contents 130
3.5.3 The Influence of Temporal Predictability (and the Lack Thereof) 131
3.5.4 Conclusion 132

4 Behavioural Correlates of Shifting Between Attentional Domains 133

4.1 Abstract 133
4.2 Introduction 134
4.3 Experiment 1: Costs of Shifting Between Attentional Domains 139
4.3.1 Methods 139
4.3.1.1 Participants 139
4.3.1.2 Apparatus and Data Acquisition 140
4.3.1.3 Task and Procedure 140
4.3.1.4 Behavioural Analysis 145
4.3.1.5 Eye-Tracking Analysis 146
4.3.2 Results 148
4.3.2.1 RTs and Errors: Staying vs. Shifting 148
4.3.2.2 Gaze: Staying vs. Shifting 150
4.3.2.3 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts 152
4.3.2.4 Gaze: Within-Domain Shifts vs. Between-Domain Shifts 154
4.3.3 Interim Summary 156
4.4 Experiment 2: Temporal Dynamics of Shifting Between Attentional Domains 156
4.4.1 Methods 157
4.4.1.1 Participants 157
4.4.1.2 Apparatus and Data Acquisition 157
4.4.1.3 Task and Procedure 158
4.4.1.4 Analysis 159
4.4.2 Results 160
4.4.2.1 RTs and Errors: 0 ms vs. 250 ms vs. 750 ms 160

4.4.2.2 Swaps and Guesses: Within-Domain Shifts vs. Between-Domain Shifts 162
4.4.3 Interim Summary 164
4.5 Experiment 3: Spatial Dynamics of Shifting Between Attentional Domains.....165
4.5.1 Methods.....165
4.5.1.1 Participants.....165
4.5.1.2 Apparatus and Data Acquisition 166
4.5.1.3 Task and Procedure 166
4.5.1.4 Analysis.....169
4.5.2 Results.....170
4.5.2.1 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts.....170
4.5.2.2 RTs and Errors: Spatial Shift vs. Non-Spatial Shift.....172
4.5.2.3 Swaps and Guesses: Within-Domain Shifts vs. Between-Domain Shifts 173
4.5.3 Interim Summary 175
4.6 Discussion 176
4.6.1 Validation of a Novel Task Design 177
4.6.2 Costs of Between-Domain Shifts 178
4.6.3 Behaviour-Gaze Dissociations When Shifting Within vs. Between Domains.....181
4.6.4 Spatial Reorienting During Between-Domain Shifts 182
4.6.5 Behavioural Differences in Reporting External and Internal Targets.....184
4.6.6 Conclusion.....186
5 Neural Dynamics of Shifting Between Attentional Domains..... 187
5.1 Abstract.....187
5.2 Introduction 188
5.3 Methods.....189
5.3.1 Participants 189
5.3.2 Task and Procedure 190
5.3.3 Apparatus and Data Acquisition 194
5.3.4 Behavioural Analysis.....196
5.3.5 MEG Pre-Processing 196
5.3.6 Time-Resolved Multivariate Pattern Analysis 197
5.3.7 Relationship Between Decoding and Behaviour.....199
5.3.8 Control Analysis: Eye Movements 199
5.3.9 Time-Frequency Analysis.....200
5.3.10 Statistical Analysis 201
5.4 Results.....202
5.4.1 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts 202

5.4.2 MEG Decoding: Within-Domain Shifts vs. Between-Domain Shifts 203
5.4.3 MEG Decoding and RTs: Within-Domain Shifts vs. Between-Domain Shifts 205
5.4.4 MEG Decoding: External Attention vs. Internal Attention 206
5.4.5 MEG Decoding: Shift Type and Attentional Domain 209
5.4.6 MEG Time-Frequency: External Attention vs. Internal Attention.....211
5.4.7 MEG Time-Frequency: Within-Domain Shifts vs. Between-Domain Shifts 214
5.5 Discussion 214
5.5.1 Distinct Neural Patterns of Within- and Between-Domain Shifts 215
5.5.2 Distinct Neural Patterns of External and Internal Attention 217
5.5.3 Relationship Between Shift-Type and Attentional-Domain Decoding.....218
5.5.4 Limitations and Future Directions 219
5.5.5 Conclusion.....219
6 General Discussion 220
6.1 The Dynamic Interplay Is Behaviourally Costly 222
6.2 The Dynamic Interplay Is Subject to (Some) Modulations.....224
6.3 The Dynamic Interplay Is Triggered by Various Events.....227
6.4 The Dynamic Interplay Is Prompt but Leaky 231
6.5 Final Conclusions.....233
7 References..... 234
8 Appendix 280
8.1 Supplementary Information for Chapter 2 280
8.2 Supplementary Information for Chapter 3 293
8.3 Supplementary Information for Chapter 4 300
8.4 Supplementary Information for Chapter 5 307

List of Abbreviations

2AFC	Two-alternative forced-choice
ANOVA	Analysis of variance
BOLD	Blood oxygen level dependent
CDA	Contralateral delay activity
DFPN	Dorsal frontoparietal network
DMN	Default mode network
ECG	Electrocardiography
EEG	Electroencephalography
EOG	Electrooculography
ERP	Event-related potential
ESD	Extreme studentised deviate
fMRI	Functional magnetic resonance imaging
HPI	Head position indicator
ICA	Independent component analysis
IDEA	Internal dominance over external attention
ISI	Inter-stimulus interval
ITC	Inferior temporal cortex
ITI	Inter-trial interval
LDA	Linear discriminant analysis
LFPN	Lateral frontoparietal network
MEG	Magnetoencephalography
MRI	Magnetic resonance imaging
MTP	Multiple trace theory of temporal preparation
PCA	Principal component analysis
P_D	Distractor positivity
PFC	Prefrontal cortex
PPC	Posterior parietal cortex
RBF	Radial basis function
ROC-AUC	Area under the receiver operating characteristic curve
RSC	Retrosplenial cortex
RT	Reaction time

SAG	Supervisory attentional gateway
SEM	Standard error of the mean
SN	Salience network
SVM	Support vector machine
TMS	Transcranial magnetic stimulation

1 General Introduction

Nearly every behaviour we exhibit is a product of the dynamic interplay between external stimulation originating in the world around us and internal contents stored in our minds. When cycling to work, for instance, one must remain focussed on various external events, such as the changing traffic lights, pedestrians crossing the street, or a sudden halt by a vehicle ahead. Yet, beyond these external events lie internal factors, such as memorised directions to the office, awareness of the vehicles behind us, and the recollection of upcoming obstacles. Though all such demands present independent challenges in and of themselves, our overall behaviour can only be most efficient when both external and internal information are operated upon in tandem. When the vehicle ahead abruptly stops, we need to react swiftly to our external surroundings without losing track of the memorised directions. When noticing that a pedestrian crossing the street happens to be a colleague from work, we find ourselves comparing external information to internal representations. When keeping track of the number of vehicles behind us as they pass, we integrate external stimulation into our mental representations. At times, external input may even prompt us to overwrite our internal contents, for example when replacing the memory of one initial route with another, due to the presence of new roadworks. Relying solely on either external or internal information in any of these scenarios might lead to inefficiencies if not potentially dangerous errors. Instead, to guide adaptive behaviour, we constantly need to shift our focus between the sensory environment and mental representations while also safeguarding our goal-relevant information processing from irrelevant interference in either domain – be it the flashing billboard on the other side of the road, or (perhaps more critically) the ingrained directions to our favourite ice cream shop.

Selective attention (hereafter referred to simply as “attention”) provides an excellent framework for investigating how the brain seamlessly draws upon both sensory and internal content to guide adaptive behaviour. Attention comprises the functions for anticipating, orienting

toward, selecting, prioritising, and transforming relevant signals within a given task, thereby subserving the fundamental requirements for mental experiences.

At their inception, experimental psychology and cognitive neuroscience were primarily focussed on understanding how external attention to sensory signals drive behaviours such as action, decision making, and learning. As these fields advanced, methods for quantifying internal content gradually emerged, allowing researchers to investigate cognitive operations based on internal information, including processes such as retrieving, relating, and constructing contents from prior experience. However, exploration within these two realms has unfolded relatively independently, with less consideration given to the interaction between external and internal contents. Consequently, the investigation of behaviours resulting from the dynamic interplay between sensory stimulation and internal states remains in its early stages.

The primary objective of my doctoral research has been to bridge the gap between the study of attention directed to the external environment and attention focussed on internal contents. Therefore, the empirical investigation within this doctoral thesis explores scenarios where individuals encounter both external and internal information in quick succession, necessitating rapid shifts in attention between these two domains. Specifically, my experimental work can be divided into two main branches. The first branch exclusively employs behavioural measures to explore cross-domain attentional modulations in behaviour; that is, how attention within one domain influences attention within the alternative domain (Chapter 2). The second branch takes a deep dive into the intricacies of shifting attention between domains by leveraging neurophysiological measures as well as other biomarkers, including fixational gaze behaviour. At the core of this branch lies an exploration of the behavioural consequences (Chapter 4) and neural dynamics (Chapter 5) associated with attentional shifts between external and internal sources of information, elucidating how these compare to shifts between contents within the same domain. Chapter 3 bridges the first and second branch by investigating the neural dynamics of reinstating internal representations after shifting to the external domain.

Before reporting the empirical portion of my work, I will start with a review of the concepts of external and internal attention. This will include an overview of how external and internal attention can be defined, as well as how these two attentional domains have been independently and interdependently studied in the past. Regarding the latter, I will introduce relevant literature on cross-domain modulations and between-domain shifts of attention. I will conclude this introductory chapter by outlining the aims of the current thesis.

1.1 Towards a Dynamic Framework of External and Internal Attention

Although it may appear cliché to introduce the concept of attention with a well-worn William James quote, the indispensable role it plays in the current thesis compelled me to incorporate it, nonetheless. In 1890, William James asserted:

> *It [Attention] is the taking possession by the mind, in clear and vivid form, of one out of what appears several simultaneously possible objects or trains of thoughts. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others […] (pp. 403-404)*

While there are a lot of insights to unpack from these 48 words, I would like to highlight two aspects. First, the quote depicts attention as a process of selectively focusing on relevant information while ignoring irrelevant contents in order to maintain efficient goal-directed behaviour. This original definition still resonates with many contemporary proposals of attention (Nobre & Kastner, 2014). In its most fundamental essence, the definition of attention has been boiled down to "*the prioritisation of processing information that is relevant to current task goals*" (Nobre & Kastner, 2014, p. 1204).

Second, James' quote demonstrates the multiplicity of attention. It can select physically present items in our external environment, but also representations stored in our minds, thus, operating on different domains of information. Later in the "The Principles of Psychology", James

(1890) makes this subdivision more explicit and states "*Attention may be divided into kinds in various ways. It is either to a) Objects of sense (sensorial attention); or to b) Ideal or represented objects (intellectual attention)*" (p. 416). What James originally labelled as sensorial and intellectual attention is now commonly referred to as external and internal attention, respectively (Chun et al., 2011; Dixon et al., 2014; Narhi-Martinez et al., 2023). The taxonomy of attention in regard to its domain (also sometimes referred to as targets of attention (Chun et al., 2011) or objects of attention (Oberauer, 2019)) is not the sole framework. Several additional categorisations delineate the intricacies of attention. For instance, James (1890) made further distinctions between voluntary versus non-voluntary attention as well as immediate versus derived attention (for extensive overviews see: B. Anderson, 2011; Nobre, 2018). However, most of these additional classifications operate within the broader context of external or internal attention, thus, coexisting in diverse combinations within each domain rather than being mutually exclusive (Chun et al., 2011).

As the experimental work of this thesis centres on visual external and internal attention, the subsequent findings reviewed in the following sections are predominantly rooted in the visual-attention literature. However, I believe that many of the lessons may generalise to other modalities and multisensory situations.

1.1.1 External Attention

External attention (a.k.a., perceptual attention, attention to perception, stimulus-oriented attention) refers to the selection and prioritisation of sensory information, that is, to information with its source outside of one's own body and mind (Burgess et al., 2007; Chun et al., 2011; Dixon et al., 2014; Narhi-Martinez et al., 2023). While external attention can occur in all senses, it is particularly prone to occur during stimulation from more distal senses, such as vision and audition, which draw attention outward to the external environment (Dixon et al., 2014).

The neurophysiological basis for external attention as a selective process arises from the severe capacity limitations of our sensory system (Broadbent, 1958; Carrasco, 2011; Neisser, 1967; W. Schneider & Shiffrin, 1977), likely due to the finite amount of overall energy available to the brain, and by the high energy cost of the neuronal activity involved in cortical computations (Lennie, 2003). Additional requirements, stemming from constraints within the brain's functional architecture – for example, the need to integrate activity among local neuronal populations or to link corresponding populations across sensory areas within and across modalities – may impose further demands that necessitate a selective process to manage these complexities effectively (Nobre, 2018). From an evolutionary perspective, these neurophysiological limitations may have evolved purposefully, as selectively focusing on sensory information relevant to adaptive behaviour and ultimately survival may be more critical than comprehensively representing the entire external environment (Nobre, 2018).

To avoid swamping the capacity-limited sensory system with irrelevant information and to withstand additional biological limitations, external attention selectively processes relevant signals from the unfolding stream of information. According to the biased competition model, external attention accomplishes this task by resolving the inherent competition for neuronal representations when multiple stimuli are presented at the same time within a common receptive field (Desimone, 1998; Desimone & Duncan, 1995; Luck et al., 1997; Moran & Desimone, 1985). When two visual stimuli fall within the same receptive field, competition for neuronal representations is biased in favour of the stimulus that is currently attended, while responses to the unattended stimulus are reduced. Thus, overcoming competition among stimuli within neuronal receptive fields could be a major contribution of attention in surmounting a significant limitation of sensory systems.

However, what prompts external attention to resolve competition among multiple stimuli? Traditional models propose two sources of attention that control the processing of external stimuli (Beck & Kastner, 2009; Corbetta & Shulman, 2002; Theeuwes, 2010). First,

stimulus-driven sources bias processing towards sensory inputs that are (intrinsically) salient. Second, goal-directed sources allow us to prioritise information consistent with our current behavioural goals. Accordingly, the focus of external attention dynamically changes based on the current priority given to bottom-up stimulus salience and top-down goal relevance. Beyond this dichotomy, numerous other factors have been proposed to operate as additional sources of attentional control. To date, there is broad consensus that attention also guides perception by drawing on explicit and implicit memories from multiple timescales, highlighting the experience-driven sources of attentional control (B. A. Anderson et al., 2021; Awh et al., 2012; Failing & Theeuwes, 2018; Narhi-Martinez et al., 2023; Nobre & Stokes, 2019).

The selective prioritisation of relevant external information plays a significant role in situations where multiple stimuli occur at once, for instance when searching for the matching sock in drawer of other "distracting" socks. However, external attention also operates in the absence of competition on isolated stimuli (Cohen & Maunsell, 2014; Maunsell & Treue, 2006). In fact, concurrency with stimulus presentation is by no means a mandatory prerequisite for attention to operate. Via goal- and experience-driven sources of control, attention can also proactively orient perception in anticipation of upcoming relevant events (Battistoni et al., 2017; Chelazzi et al., 1993; Hopfinger et al., 2000; Kastner et al., 1999; Luck et al., 1997; Posner, 1980; Stokes et al., 2009; Worden et al., 2000). Consequently, external attention can occur in the presence and absence of stimuli.

1.1.2 Internal Attention as Working Memory

The definition of internal attention requires careful formulation due to the absence of an agreed-upon taxonomy of internal mental states. Even though compelling efforts within the field of philosophy of neuroscience have been dedicated to crafting a precise definition of stimulus-independent representations and states (Morgan, 2014; Thomson & Piccinini, 2018), navigating

through the ongoing heated debate often proves more perplexing than beneficial. Hence, it may be prudent to begin with the more obvious aspects.

Undeniably, internal mental states represent a broad spectrum of experiences occurring within the body and mind, pertaining to phenomena not immediately accessible to our senses. The degree to which internal states correspond to a particular state of the external environment can vary. For instance, internal states can clearly emerge from external signals such as memories of different types and lifetimes. These can include perceptual traces; working-memory contents and operations; and sensory, motor, semantic, and episodic long-term memories. Internal states can also build more loosely on sensory content and previous experience, such as planning, problem solving, prospection, imagination, mind wandering, meta-awareness and dreaming (Christoff et al., 2016). In addition, internal states can be classified based on their effectiveness in guiding action and behaviour, with ruminations, hallucinations, and suicidal thoughts positioned on the maladaptive end of the spectrum. Certainly, there are numerous alternative ways to categorise internal mental states. However, what all internal mental states have in common is that their processing is largely decoupled from the external perceptual environment, regardless of the presence or absence of external stimuli (Dixon et al., 2014; Schooler et al., 2011; Smallwood, 2013).

At times, the term internal attention has been used inconsistently, being equated with a multitude of different internal mental states (Gjorgieva et al., 2022; Kam et al., 2018; Kritzman et al., 2022; Verschooren & Egner, 2023; Villena-González et al., 2017; Wulf et al., 1998; Ziegler et al., 2018). This complicates navigation through the scientific literature and increases the risk of overgeneralising findings. To this end, I will adopt the terms of "internally directed cognition" and "internally directed mental activity" to refer to any of these aforementioned mental states (Dixon et al., 2014). In line with a substantial body of research, I will reserve the concept of "internal attention" to denote the selective processing of past sensory contents, particularly working-memory contents (Chun, 2011; Kiyonaga & Egner, 2013; van Ede & Nobre, 2023).

But this leads to the next question – how exactly do we define working memory? William James was not exclusively ahead of the game in the field of attention; as it turns out, he also demonstrated considerable acumen in the area of working memory. James introduced the concept of "primary memory", referring to the initial, temporary storage of events that have recently occurred, appearing to us as part of the present moment even though they belong to the immediate past (referred to as "specious present"; James, 1890). From this early notion, the concept of working memory evolved, first coined in 1960 as a "quick-access repository" used to maintain and carry out behaviourally relevant plans (G. A. Miller et al., 1960). However, formal investigations of working memory did not commence until the 1970s, when Baddeley adopted the term for the famous multicomponent model (Baddeley & Hitch, 1974). In this thesis, I would like to refrain from implying any theoretical constructs of working memory, defining it as the capacity-limited temporal maintenance or manipulation of internal contents to guide future behaviour, as many researchers have commonly done (Baddeley, 2012; D'Esposito & Postle, 2015; E. K. Miller et al., 2018; Nobre, 2018; van Ede, 2023).

1.1.3 Similar but Not the Same

Throughout the era of cognitive research, various theoretical dichotomies have framed the study of attention. While invigorating copious debates, this dogmatic approach has inadvertently hindered exploration, often forcing examinations of attention to adhere to one binary division or the other, rather than delving into the intricate dynamics of their interaction (B. Anderson, 2011). This also holds true for the investigation of external and internal attention. The sequential evolution of external- and internal-attention research likely contributed to the prevalent practice of studying the two processes in isolation. External attention, akin to the elder sibling, has been extensively examined, while the exploration of internal attention only gained momentum in the late 20th century. When investigating the relationship between these two processes, much emphasis has been placed on identifying cognitive and mechanistic commonalities between

attention in perception and working memory. Consequently, as of today, a substantial degree of parallels has been drawn between attention in the external environment and in working memory. Nevertheless, it is important to refrain from assuming that attention operates in the same way when directed externally and internally. The following section aims to provide a brief overview of the comparative study of external and internal attention in terms of their sources of control, their processing stages, their targeted attributes, and their modulatory outcomes, highlighting commonalities and differences between domains.

Both external and internal attention involve a set of mechanisms modulating contents on multiple stages of processing (Myers et al., 2017; Nobre & Kastner, 2014; Nobre & van Ede, 2023; van Ede & Nobre, 2023). Within the sensory and working-memory domain, a considerable amount of research has examined the initial attentional orienting stage. Orienting towards external information has been investigated in tasks using predictive pre-cues, proactively providing information concerning a target-defining attribute prior to a target array (Posner, 1980). Conversely, tasks using informative retro-cues have proven effective for investigating internal orienting of attention to selective contents within working memory (Griffin & Nobre, 2003; Landman et al., 2003). Past research has demonstrated that external attention can be oriented towards space (Kastner et al., 1999; Luck et al., 1997; Posner, 1980), objects (Duncan, 1984; Egly et al., 1994; Esterman & Yantis, 2010; Gayet & Peelen, 2022; Peelen & Kastner, 2011; Puri et al., 2009; Stokes et al., 2009), features (Chawla et al., 1999; Giesbrecht et al., 2003; Kok et al., 2017; T. Liu et al., 2007; Maunsell & Treue, 2006; Shibata et al., 2008; Snyder & Foxe, 2010), temporal structures (Correa et al., 2005; Coull & Nobre, 1998; Denison et al., 2017; Miniussi et al., 1999; Rohenkohl et al., 2014), sematic attributes (Cristescu & Nobre, 2008; Neely, 1976; Posner & Snyder, 1975), and more (e.g., Miniussi et al., 2005). Likewise, internal attention can be directed towards these attributes once they have been encoded into working memory (space: (Astle et al., 2012; Griffin & Nobre, 2003; Landman et al., 2003); objects: (Lin et al., 2021; Peters et al., 2015); features (Hajonides et al., 2020; Serences et al., 2009); time: (Heuer & Rolfs, 2022; Jin et al., 2020;

van Ede et al., 2017), semantic attributes: (Backer et al., 2020)). Therefore, both external and internal attention can each be further subdivided based on identical dimensions, which is a logical relationship considering that content in working memory arises from previous external stimulation.

Comparing behavioural and neural implications of external and internal attention reveals additional commonalities. At a basic level, both pre- and retro-cues exhibit similar behavioural outcomes, wherein directing attention to task-relevant perceptual or mnemonic contents results in enhanced performance characterised by increased speed and accuracy (Griffin & Nobre, 2003; Lepsien & Nobre, 2006; Wallis et al., 2015). Further behavioural markers, such as fixational gaze behaviour, have been shown to accompany both orienting attention to external locations (Engbert & Kliegl, 2003; Hafed et al., 2011) as well as directing attention to internal representations held within the spatial framework of working memory (Draschkow et al., 2022; van Ede, Chekroud, & Nobre, 2019; van Ede et al., 2021). The neural systems responsible for directing attention in perception and working memory demonstrate analogous patterns, as evidenced by the involvement of the areas in the dorsal frontoparietal network (DFPN; Nee & Jonides, 2009; Nobre et al., 2004; Tamber-Rosenau et al., 2011; Wallis et al., 2015). In line with this, similar patterns of activity have been observed in human visual and parietal cortices during attentional orienting in perception and working memory, suggesting that orienting to sensory and mnemonic items share common mechanisms (Y. Zhou et al., 2022). Moreover, neurophysiological markers of attention, such as relative desynchronisation of alpha-band oscillations in the hemisphere contralateral to the location of the pre- and retro-cued items, appear to be roughly comparable between external attention and internal attention (Mok et al., 2016; Myers et al., 2015; Wallis et al., 2015), highlighting an overlap in the neural modulations in both domains.

Informative pre- and retro-cues provide goal-directed (i.e., endogenous, voluntary) sources of attentional control, allowing us to prioritise information consistent with our current behavioural objectives. As mentioned earlier, external attention can be additionally governed by

stimulus- and experience-driven sources of control. While internal attention to working-memory contents has been regarded as being goal- or experience-driven, recent studies have additionally provided evidence for an involuntary, stimulus-driven capture of internal attention within working memory (Ester & Nouri, 2023; van Ede, Board, et al., 2020). Consequently, similar sources of control (i.e., goal-directed, stimulus-driven, and experience-drive factors) seem to compete in perception as well as in working memory to guide attention to sensory or internal signals, respectively. Nevertheless, so far, the comparison of goal-directed external and internal attention has most commonly been studied.

Some researchers have suggested that goal-directed attention in working memory can simply be thought of as sustained internal attention (i.e., equating to preparatory external attention; Chun, 2011), and that external attention and working memory share a single, capacity-limited resource (Kiyonaga & Egner, 2013; Verschooren & Egner, 2023; Y. Zhou et al., 2022). According to this view, the same cognitive and neural mechanisms are deployed in each case, with the main difference being the domain on which they operate. However, it is crucial to acknowledge that mechanisms underlying selective external and internal attention are not identical. Differences between external and internal attention become especially apparent when taking a closer look at their affordances across various attentional stages.

For instance, while external orienting encompasses anticipating the appearance of perceptual events, orienting within working memory occurs in the context of pre-existing contents. This implies that within working memory, information becomes available for immediate selection and prioritisation upon orienting to the relevant memorised content. By contrast, during attention to external events, selection cannot take place until the cued event occurs, and further requires the identification of cued objects and their associated features among all perceptual input. Accordingly, an electroencephalography (EEG) study revealed that pre-cues induced a pronounced lateralisation in late cue-evoked potentials linked to anticipating a target (Myers et al., 2015). In contrast, retro-cues elicited a distinct pattern of event-related potentials (ERPs)

indicative of an early selection mechanism rather than anticipation of a stimulus. Further support for the presence of selection processes during internal attention comes from neuroimaging investigations, which have demonstrated the additional involvement of medial and lateral prefrontal areas (Lepsien et al., 2005; Lepsien & Nobre, 2006; Nee & Jonides, 2009; Nelissen et al., 2013; Nobre et al., 2004; Tamber-Rosenau et al., 2011; Wallis et al., 2015) along with motor regions (Chatham & Badre, 2015; Tamber-Rosenau et al., 2011) following orienting attention within working memory. Conversely, these regions exhibit less consistent involvement during proactive orienting of external attention based on pre-cues, likely because selection mechanisms are not activated until target onset (van Ede & Nobre, 2023).

The selective recruitment of these additional brain areas might represent the transformation of the prioritised internal content into a prospective, action-oriented representation during internal attention (Myers et al., 2017; Olivers & Roelfsema, 2020). Consistent with this idea are recent findings from a study in non-human primates, revealing dynamic transformations of internal representations into task-relevant codes that were shifted away from non-relevant attributes such as the spatial location of the memorised item (Panichello & Buschman, 2021). Moreover, EEG studies have highlighted that motor plans are represented in working memory to prepare for upcoming behaviour (Boettcher et al., 2021; Ester & Weese, 2023; Nasrawi et al., 2023; Nasrawi & van Ede, 2022; van Ede, Chekroud, Stokes, et al., 2019). The transformation into motor codes and selection thereof cannot occur during preparatory external attention since the relevant information has not yet been presented.

In summary, this section has demonstrated substantial overlaps between external and internal attention regarding their targeted attributes and sources of control. However, it is important to keep in mind that the temporal dynamics of the attentional stages and the representational format of the contents within each of these stages might drastically differ across domains. Hence, internal attention is not entirely analogous to external attention. While not the

primary focus of this thesis, Chapters 4 and 5 will delve further into uncovering shared and unique traits of external and internal attention.

1.1.4 From a Comparative to a Dynamic Study

It is tempting to concentrate on a comparative investigation of external and internal attention. Given that external attention has a longer history of study, there is a natural inclination to extend the same research questions that were explored decades ago in the realm of external attention to the study of internal attention. While this perspective undoubtedly enhances our understanding of internal attention, it overlooks a crucial aspect – the dynamic interplay between domains. To advance in the field, our scientific inquiry must move beyond comparison and strive to understand the interactions between external and internal attention.

Advocating for a dynamic framework that focusses on the interactions between external and internal attention is compelling due to the interconnected nature of the two domains in everyday situations. Rather than operating in isolation, they often collaborate synergistically, with information from the environment and our working memory harmoniously contributing to foster adaptive behaviours (van Ede & Nobre, 2023). For example, when assembling a new piece of furniture, our initial step involves carefully attending to the provided assembly instructions. As we proceed, it is important to maintain a visual representation of the necessary component in working memory before searching for it among the other pieces. This adaptive, temporally extended behaviour is possible because attentional contents transcend their individual domain boundaries, allowing for cooperation between perception and working memory (van Ede & Nobre, 2023). In particular, sensory information selected by external attention can enter working memory, be stored as internal representations, and subsequently committed to long-term memory (i.e., operating outside-in). In turn, stored representations (which either just entered working memory from perception or were retrieved from long-term memory) can be accessed by internal attention to guide anticipated perception and action in the external world (i.e., operating

inside-out). This highlights the central role that working memory occupies in the framework of external and internal attention, acting as a pivotal exchange point supporting the bidirectional flow of information between domains (Chun et al., 2011; Kiyonaga & Egner, 2013; Myers et al., 2017; Narhi-Martinez et al., 2023; van Ede & Nobre, 2023).

However, the intimate dance of external and internal attention does not always guarantee optimal behaviour. In some instances, irrelevant information may vie for priority, potentially hindering behavioural efficiency. For example, the television playing in the background might capture our external attention, while trying to hold in working memory the specific type of screw we need to find in the toolbox. At other times, we may find ourselves navigating multiple external and internal tasks with distinct behavioural objectives within the same time frame. In such scenarios, safeguarding task-relevant information within domains from unwelcome, competing influences becomes crucial. However, achieving this is not always possible, thus resulting in unwanted cross-domain modulations (Lorenc et al., 2021; Olivers et al., 2011).

Regardless of whether external and internal sources of information serve the same or different behavioural goals, a prevailing premise in the field asserts that we cannot attentively monitor both the external environment and internal mental states at the same time (Awh et al., 1998; Brooks, 1968; Dixon et al., 2014; M. D. Fox et al., 2005; Handy & Kam, 2015; Kiyonaga & Egner, 2013; Macedo-Pascual et al., 2023; Servais et al., 2023; Unsworth & Robison, 2016). That is, whenever internal attention is actively engaged, active processing via external attention is likely suppressed and vice versa. Therefore, to allow behaviour to be guided by both external and internal signals, we need to shift attention from relevant information in one domain to relevant information in the other domain, while avoiding transitions to task-irrelevant information. Only recently, an integrative theory of external and internal attention has been proposed, acknowledging the fact that we are constantly balancing between external and internal influences (Narhi-Martinez et al., 2023). Apart from this, the dynamic interplay between perception and working memory has rarely been discussed within frameworks of attention.

The empirical chapters of this thesis examine the dynamic interplay between external and internal attention from two perspectives. First, I will explore the consequences that arise from the intimate interactions between external and internal attention, specifically focussing on attentional modulations across domains. In essence, I will investigate how external attention influences internal attention and vice versa. Next, I will focus on the exploration of attentional shifting between domains, an area that has largely remained uncharted territory despite its fundamental importance in shaping behaviours.

1.2 Cross-Domain Attentional Modulations

In the context of the dynamic interplay between external and internal attention, how attending to external signals shapes internal attention and, conversely, how attending to internal signals impacts external attention, has probably garnered the most interest. In this thesis, I will term these influences from one domain to the other as "cross-domain attentional modulations". Previous research has demonstrated that external and internal attention bidirectionally impact one another, whether or not internal representations and external information are consistent with current behavioural goals (Dudukovic & Wagner, 2006; Teng et al., 2023). In the following sections, I will highlight relevant findings regarding internal-to-external and external-to-internal modulations, before diving into literature examining the mechanisms mitigating cross-domain influences. This will be relevant for Chapter 2, in which I will introduce two projects. The first project explores internal-to-external attentional modulations, while the second project investigates external-to-internal attentional modulations. Additionally, I will discuss potential neural mechanisms that may elucidate how we safeguard our internal contents in the latter case.

1.2.1 Internal-to-External Attentional Modulations

Working-memory contents influence what we attend to in the external world. This idea resonates across multiple theories of external attention, particularly in the context of visual search. For

example, the popular guided-search framework has argued that the relevant target of attention is represented in working memory (i.e., search template) and primes subsequent deployments of external attention during search (Wolfe, 1994, 2021; Yu et al., 2023). Consistent with this idea, research has demonstrated that internal representations maintained in working memory bias neural activity towards information that is relevant to the current perceptual task (Bundesen et al., 2005; Carlisle et al., 2011; Chelazzi et al., 1993; Desimone & Duncan, 1995; Dodwell et al., 2024).

In visual-search scenarios, working-memory content plays a crucial role in guiding our attention towards task-relevant external events. However, we frequently encounter situations where tasks necessitate external attention, alongside the retention of other internal information that is not directly linked to the current task, but nevertheless held for potential future use. Even in these situations, in which internally maintained representations are irrelevant to the current external task, it has repeatedly been observed that the contents of working memory can affect the allocation of external attention. One of the foundational supports for this assumption stems from studies conducted by Awh et al. (1998) which demonstrated that maintaining the spatial location of an object in working memory enhances the processing efficiency of other external stimuli presented in that same location during the retention interval. Similar results have been obtained for internally represented features and objects (Downing, 2000; Y. Huang et al., 2019; Kiyonaga et al., 2012; Olivers, 2009; Olivers et al., 2006; Sasin & Fougnie, 2020; Soto et al., 2005, 2008). For instance, when participants are asked to maintain a colour or shape of an item in working memory and perform a visual-search task during the retention interval, search performance is improved for memory-matching targets, however, negatively influenced when external distractor items matched items held in working memory. Retro-cue studies have shown that this phenomenon, known as memory-driven attentional capture, is largely restricted to a single prioritised internal item (Mallett & Lewis-Peacock, 2018; Olivers et al., 2011; van Moorselaar, Battistoni, et al., 2015; van Moorselaar et al., 2014). Consequently, an internally

maintained representation appears to be beneficial when it aligns with the current external-task goal but becomes disadvantageous when it lacks relevance for external attention.

Internal representations stored in working memory do not only bias attentional selection in the external world. The contents in working memory can even influence more basic levels of processing, distorting whether and how visual features are perceived (Kang et al., 2011; Macedo-Pascual et al., 2023; Teng & Kravitz, 2019) via alterations of neural response profiles in visual areas (Bahmani et al., 2018; Gayet et al., 2017; Merrikhi et al., 2017). However, as the strength of internal representations decays over time, their influence on perceptual processing becomes weaker (Teng et al., 2023). These findings have been often taken as evidence that visual areas involved in ongoing perception are also recruited during (visual) working-memory maintenance (Harrison & Tong, 2009; Pasternak & Greenlee, 2005; Postle, 2006; Rademaker et al., 2019; Serences et al., 2009; Sreenivasan, Vytlacil, et al., 2014).

The mentioned studies on internal-to-external attentional modulations suggest that mnemonic information can significantly infiltrate external attention, regardless of whether internal representations are consistent with the external perceptual goals. While previous studies have focussed on how spatial (Awh et al., 1998; Downing, 2000), feature- (Y. Huang et al., 2019; Kiyonaga et al., 2012; Olivers, 2009; Olivers et al., 2006; Sasin & Fougnie, 2020; Soto et al., 2005, 2008), or (real-world) object-based internal representations (Jung et al., 2018) influence external attention, another prominent guiding attribute of attention has been rather neglected. Predictable temporal structures can guide attention to the timing of relevant events in the external world (Coull & Nobre, 1998; Denison et al., 2017; Rohenkohl et al., 2014) and have been demonstrated to dynamically prioritise internal representations when they become most relevant for behaviour (Jin et al., 2020; van Ede et al., 2017). In the first part of Chapter 2, I will investigate how predictable temporal structures regarding internal representations maintained in working memory influence the processing of external stimuli.

1.2.2 External-to-Internal Attentional Modulations

External attention can directly influence multiple stages of internal attention (Liesefeld et al., 2020; Lorenc et al., 2021; McNab & Dolan, 2014). For example during encoding, the selection and prioritisation of perceptual events shapes which information enters working memory and the level of detail with which it is retained (Awh et al., 2006; Dube et al., 2017; Emrich et al., 2017; Feldmann-Wüstefeld & Vogel, 2019; McNab & Klingberg, 2008; A. M. Murray et al., 2011; Myers et al., 2014). In this regard, external attention functions as a crucial gatekeeper, prioritising information that aligns with the current and future processing goals, while downweighing irrelevant external signals to prevent them from entering the severely capacity-limited storage (Gazzaley & Nobre, 2012; Liesefeld et al., 2020).

Once internal representations have been stored in working memory, they exhibit a remarkable resilience against external interference (Lorenc et al., 2021). However, unwanted external-to-internal attentional modulations still occur to some extent. On a behavioural level, interference can lead to reduced precision of internal representations (Bae & Luck, 2019; Magnussen & Greenlee, 1992; Rademaker et al., 2015; S. Z. Sun et al., 2017), increased responses to non-target items (Zhang & Lewis-Peacock, 2023), as well as attractive (Gupta & Sridharan, 2023; Mallett et al., 2020; Rademaker et al., 2015; Teng et al., 2023; Teng & Kravitz, 2019; van der Stigchel et al., 2007; Zhang & Lewis-Peacock, 2022, 2023) and sometimes repulsive biases (Bae & Luck, 2019; Gupta & Sridharan, 2023) towards or away from features of the external item, respectively.

Various factors contribute to the degree to which unwanted external events modulate internal representations. These factors encompass the type of interference, the resemblance between interfering external stimuli and memory contents, the ability to prioritise internal representations, and to anticipate interference. In terms of interference types, it has been shown that working memory can be disrupted both by passive exposure to irrelevant distractors (Nemes et al., 2012; Rademaker et al., 2015; D. Schneider, Barth, Getzmann, et al., 2017; S. Z. Sun et al.,

2017) as well as by active visual processing of task-relevant interrupters requiring a response (Fukuda et al., 2022; Lively et al., 2021; Teng et al., 2023; Teng & Kravitz, 2019; Ülkü et al., 2024). Particularly, interruptions prove detrimental to working-memory performance when compared to distractions (Bae & Luck, 2019; Berry et al., 2009; Fukuda et al., 2022; Mishra et al., 2013; B. Wang, Theeuwes, et al., 2019; Zickerick, Kobald, et al., 2021). This is likely due to a higher probability of external stimuli to enter working memory when interference is task-relevant (Hakim et al., 2021). Additionally, the need for reactivation of internal contents after interruption is necessary, a process likely not required during mere perceptual distraction (Clapp et al., 2010; Sakai et al., 2002; Zickerick, Kobald, et al., 2021).

Previous studies have demonstrated that the severity of unwanted external-to-internal modulations further hinges on the similarity between memoranda and interference. For instance, external stimuli tend to have a greater impact on working-memory contents when they belong to the same object category (Jha et al., 2004; Sreenivasan & Jha, 2007; Yoon et al., 2006), are perceptually similar in feature space (Rademaker et al., 2015; Sun et al., 2017), and/or are subjectively perceived as similar (Fukuda et al., 2022). This similarity-induced effect likely varies depending on the time between working-memory encoding and perceptual interference onset: as maintained representations decay over time, their likelihood of being impacted by incoming perceptual signals increases (Hautekiet et al., 2024; Teng et al., 2023).

The finding that a high overlap between interference and internal contents impairs working-memory performance has been interpreted as support for the sensory recruitment hypothesis (Harrison & Tong, 2009; Pasternak & Greenlee, 2005; Postle, 2006; Rademaker et al., 2019; Serences et al., 2009; Sreenivasan, Vytlacil, et al., 2014). This hypothesis, originally mentioned by William James over a century ago (James, 1890), suggests that both perception and working memory recruit sensory regions such as early visual areas for perceptual processing and mnemonic retention, respectively. Accordingly, competition over such a shared substrate and, consequently, behavioural interference, would increase as the overlap increases between

perceived information and maintained representations. Indeed, functional Magnetic Resonance Imaging (fMRI) studies have provided support for this hypothesis by showing that behavioural decrements in working-memory performance due to interference can go in hand with disruptions of mnemonic information in early visual cortices (Hallenbeck et al., 2021; Lorenc et al., 2018; Rademaker et al., 2019).

Another factor influencing the impact of external interference on working memory is the current priority status of the internal representation. Priority can be manipulated experimentally with retro-cues presented during the retention interval of a working-memory task indicating the likelihood of an item being subsequently probed (Griffin & Nobre, 2003; Landman et al., 2003; Souza & Oberauer, 2016). Retro-cue prioritisation has been shown to shield internal representations from perceptual interference (Barth & Schneider, 2018; Makovski et al., 2008; Makovski & Jiang, 2007; D. Schneider, Barth, Getzmann, et al., 2017; van Moorselaar, Gunseli, et al., 2015; Vergauwe et al., 2023) by enhancing protection through top-down influences (Souza & Oberauer, 2016). For instance, when a single item is retroactively cued, working-memory performance can remain consistent regardless of the presence of distractors (D. Schneider, Barth, Getzmann, et al., 2017). In comparison, when presented with a neutral cue devoid of information about the item to be tested, working-memory performance significantly deteriorates in the presence of distractors.

These findings stand in contrast with previously mentioned findings regarding internal-to-external modulations, suggesting that prioritised internal representations are more prone to interact with perception (Mallett & Lewis-Peacock, 2018; Olivers et al., 2011; van Moorselaar, Battistoni, et al., 2015; van Moorselaar et al., 2014). Likewise, in the field of external-to-internal modulations, some researchers have postulated that external events may exert a more pronounced influence on prioritised memory representations than on unprioritised ones. This could be attributed to higher priority yielding enhanced activation in sensory areas and thus more interactions with ongoing perception (Lorenc et al., 2021). Evidence supporting this notion arises

from studies demonstrating that prioritised and unprioritised memory items are either equally susceptible to external interference (Hautekiet et al., 2024; Zhang & Lewis-Peacock, 2022) or, in some cases, even more strongly impacted (Mallett & Lewis-Peacock, 2019). The susceptibility of retro-cued memories is further heightened when visual interference emerges during active manipulation of internal content, suggesting that retroactive reinforcement is less effective at safeguarding items that are currently in use (Lout et al., 2023). Along similar lines, Zokaei et al. (2014) demonstrated that transcranial magnetic stimulation (TMS) to the visual cortex disrupted the precision of working-memory items only after they had been cued. These findings give rise to speculation that prioritisation might exert a dual influence through opposing factors of protection and vulnerability. Seeking clarity amidst this complexity, a recent study revealed that prioritisation shielded internal representations from being completely outcompeted by interfering items, yet rendered them more susceptible to attractive biases towards task-irrelevant interference (Zhang & Lewis-Peacock, 2023).

Lastly, learned regularities and expectations regarding upcoming perceptual interference could additionally influence external-to-internal modulations (Lorenc et al., 2021). In an EEG study, Hakim et al. (2020) observed an enhanced amplitude of the contralateral delay activity (CDA) following perceptual interference in blocks with a high (i.e., 75%) versus low (i.e., 25%) likelihood of interference occurrence during the retention interval. This amplified activity, attributed to active working-memory maintenance, might suggest that participants were able to safeguard their working-memory representations when anticipating the distracting perceptual input with higher certainty. However, on a behavioural level, no discernible advantage was observed between trials with higher versus lower probabilities of interference occurrence (Hakim et al., 2020). Nevertheless, in certain circumstances, expectations about interference might enhance working-memory performance. For instance, in a recent study, working-memory performance of older participants improved in trials where the occurrence of an interfering task during the retention interval was 100% predictable, as compared to those trials where

interference could not be anticipated (Ülkü et al., 2024). Moreover, Wang et al. (2022) provided evidence that expectations regarding external events might protect working-memory contents. Participants were tasked with memorising both the colour and orientation of a presented grating. In colour-interference blocks, most trials involved an intervening colour task, while the minority of trials involved an intervening orientation task during the retention interval (and vice versa for orientation-interference blocks). Behavioural findings revealed that memory of a specific feature (i.e., colour or orientation) suffered more from interference defined within the same feature-dimension than from interference defined within the other feature-dimension, however, only when such interference was unanticipated. This was interpreted as evidence that the to-be-remembered features for which interference was expected were protected, whereas to-be-remembered features for which no interference was expected were left vulnerable. fMRI decoding results corroborated this notion, showing that preceding interference, the degree to which a particular feature of the memorised content was represented in the early visual cortex relied on the expected interference feature (i.e., higher decoding accuracy of anticipated features; B. Wang et al., 2022). In the second part of Chapter 2, I will delve deeper into this topic, investigating whether expectations regarding the temporal onset of external interference during the retention interval can serve to protect internal contents from undesired external-to-internal modulations.

1.2.3 Protecting From vs. Mitigating External-to-Internal Modulations

The preceding section has underscored several factors that influence the degree to which external interference can impact internal representations. To understand the underlying mechanisms supporting interference-robust working-memory contents, I would like to highlight two areas of research. Firstly, investigations into the neural substrates of internal representations may offer valuable insights into how they are protected from unwanted external interference. I will discuss persistent delay activity, parallel coding, activity-silent storage, and dynamic coding as potential

mechanisms for safeguarding internal representations. Secondly, research on inhibitory processes involved in limiting perceptual processing during maintenance may offer additional insights about working-memory resistance. In the second part of Chapter 2, I will present a behavioural experiment that examines the contribution of protective and inhibitory mechanisms in situations where the temporal onset of an interfering event can be anticipated during working-memory maintenance.

1.2.3.1 Protecting Internal Representations During Maintenance

Early findings suggested that persistent, stimulus-specific delay activity of neurons in the prefrontal cortex (PFC) plays a key role in working-memory maintenance (Courtney et al., 1998; D'Esposito et al., 1995; Funahashi et al., 1989; Fuster & Alexander, 1971; Goldman-Rakic, 1995). This activity was assumed to support the resilience of internal representations against distractors (Compte et al., 2000). Evidence stems, for example, from neurophysiological studies in non-human primates demonstrating that delay activity in PFC neurons is more robust against distractors as compared to areas closer to primary sensory areas such as the inferior temporal cortex (ITC; E. K. Miller et al., 1993, 1996) or the posterior parietal cortex (PPC; Constantinidis & Steinmetz, 1996).

However, the involvement of the PFC in the content-specific storage of internal representations has encountered scepticism from proponents of the sensory recruitment hypothesis, who argue instead that neurons in the sensory areas are responsible for visual working-memory storage (Harrison & Tong, 2009; Pasternak & Greenlee, 2005; Postle, 2006; Rademaker et al., 2019; Serences et al., 2009; Sreenivasan, Curtis, et al., 2014). This perspective claims that the delay activity in the PFC lacks a high degree of feature selectivity, a critical factor for accurately retaining subtle distinctions between items stored in working memory (Lara & Wallis, 2014; Lebedev et al., 2004; Pasternak et al., 2015; Riggall & Postle, 2012). Instead, the role of the PFC might rather lie in regulating processes that support the active information storage in

visual cortices (Awh & Jonides, 2001; Curtis & D'Esposito, 2003; D'Esposito & Postle, 2015). Despite concerns that storing information in sensory regions may render memory representations vulnerable to interference (Bettencourt & Xu, 2016), evidence suggests that simultaneous representations of both working memories and sensory inputs can be observed in the early visual cortex during perceptual interference (Hallenbeck et al., 2021; Rademaker et al., 2019). Such sensory representations resilient to distractions have been shown to rely heavily on top-down support from the PFC (Feredoes et al., 2011; T. G. Lee & D'Esposito, 2012; B. T. Miller et al., 2011; Zanto et al., 2011).

In addition to the involvement of the PFC and sensory regions, working-memory storage and protection from external interference may be further supported by parallel coding of internal contents in multiple brain regions (Christophel et al., 2017; Hallenbeck et al., 2021). Persistent neural activity that is informative about a currently memorised stimulus has be found in sensory, parietal, and prefrontal brain regions (Christophel et al., 2017), although these representations may vary in format and quality (Rademaker et al., 2019). Representations in the early visual cortex may be analogue and more precise (but see: Kwak & Curtis, 2022), while representations in parietal and prefrontal regions may be coarser, that is, more abstract (S.-H. Lee et al., 2013) or transformed into action-based codes (Boettcher et al., 2021; Myers et al., 2017; Panichello & Buschman, 2021). Parallel coding in different brain regions could hold significant importance in enhancing resistance to distractions (Bettencourt & Xu, 2016). For instance, an fMRI study revealed that working-memory content was redundantly maintained in early visual areas and in the intraparietal sulcus (IPS). While internal representation in early visual areas showed susceptibility to external interference, the parietal representations demonstrated greater resilience, thus, potentially minimising behavioural biases induced by interference (Lorenc et al., 2018).

Many studies have failed to demonstrate activity of neurons or neuronal populations in the PFC to persist throughout the working-memory maintenance period (Barak et al., 2010; Brody

et al., 2003; Shafi et al., 2007). As a consequence, there is an ongoing debate regarding whether alternative neural mechanisms may offer a more accurate understanding of how content-specific representations are maintained in working memory. Within this discourse, the activity-silent working-memory framework has garnered considerable traction (Mongillo et al., 2008; Stokes, 2015; Stokes et al., 2013). According to this framework, the maintenance of internal information in an "activity-silent" state, devoid of concurrent neural activity, may rely on transient functional changes in synapses through short-term synaptic plasticity. These synaptic changes may temporarily link neural populations coding for mnemonic representations. Critically, this model suggests that non-specific stimulations can reactivate memories stored in silent synaptic traces, thus reinstating the original firing patterns. Thus, short-term synaptic changes may effectively support the maintenance of internal representations without persistent neural activity while simultaneously enabling robust readouts several seconds later.

Evidence for an activity-silent working-memory storage has come from human fMRI, EEG, and magnetoencephalography (MEG) studies using multivariate pattern analyses. In particular, these studies have illustrated that when participants are selectively attending to a subset of stored items, the decoding accuracy of unattended items often decreases to baseline levels (LaRocque et al., 2013; Lewis-Peacock et al., 2012; Muhle-Karbe et al., 2021; Rose et al., 2016; Sprague et al., 2016; Wolff et al., 2017). When attention shifts back to the previously unattended items, the neural traces associated with those items are reactivated. Remarkably, this reactivation even occurs following the presentation of external interference during the working-memory delay period (Lewis-Peacock et al., 2012). This suggests that unattended but still relevant mnemonic representations were likely kept in activity-silent (also sometimes referred to as "latent") states. Importantly, latent internal information can be reinstated by uninformative visual or neural impulses, so called "pings" (Rose et al., 2016; Wolff et al., 2017). The perturbation caused by the pinging mechanism might function as distractors, potentially initiating a top-down protection that leads to the reactivation of latent internal representations (Lorenc et al., 2021). Furthermore,

Lorenc et al. (2021) propose that internal representations maintained in a latent state may possess greater resilience against external interference. This resilience may arise from the fact that latent memories minimally interfere with ongoing behaviour (Kreither et al., 2022; Mallett & Lewis-Peacock, 2019; Muhle-Karbe et al., 2021) and cannot be manipulated without prior reactivation (Trübutschek et al., 2019). Consequently, if information remains unalterable while in a latent state, it should be less likely to be unintentionally disrupted by incoming perceptual input (Lorenc et al., 2021).

In addition to the activity-silent framework, the dynamic coding model of working memory has been put forward, positing that internal representations are coded by non-stationary spatiotemporal patterns of neural activity (Meyers, 2018; Meyers et al., 2008; Sreenivasan, Curtis, et al., 2014; Sreenivasan, Vytlacil, et al., 2014). Despite dynamic neural activity, it has been suggested that the neuronal population can still represent stable working-memory representations (J. D. Murray et al., 2017; Spaak et al., 2017; Sreenivasan, Curtis, et al., 2014). For example, a study on non-human primates demonstrated that dynamic neuronal populations in the PFC remained resilient to external distractions while preserving remembered information within a temporally stable subspace (Parthasarathy et al., 2019). Thus, interference resistance in working memory could be provided by a dynamic working-memory code that dynamically diverges from the initial sensory-evoked pattern. This divergence could render mnemonic contents more distinguishable from the activity induced by incoming perceptual signals, all while still retaining internal representations in a temporally stable subspace (Lorenc et al., 2021; Serences, 2016).

1.2.3.2 Mitigating External Interference During Maintenance

Alongside neural substrates and mechanisms dedicated to protecting working-memory contents, control processes aimed at mitigating the processing of perceptual interference might further increase working-memory resistance. The majority of research investigating inhibition of

perceptual interference has been conducted within the realm of external attention tasks, such as visual search or spatial orienting (Geng & Witkowski, 2019; Noonan et al., 2018; D. Schneider et al., 2022; van Moorselaar & Slagter, 2020; Wöstmann et al., 2022), or during working-memory encoding (Liesefeld et al., 2020). Nevertheless, accumulating evidence suggests that interference inhibition is also crucial for bolstering working-memory resistance during maintenance. Neuroimaging studies have reported enhancement of activity in visual areas associated with stimuli required to be maintained, and suppression of activity associated with stimuli that should be ignored, relative to a stimulus-matched passively viewed baseline condition (Gazzaley, 2011; Gazzaley et al., 2005). Furthermore, in a recent EEG study, the amplitude of the distractor positivity (P_D), an ERP component indicative of distractor suppression (Gaspar & McDonald, 2014; Gaspelin et al., 2023; Hickey et al., 2009; Sawaki & Luck, 2010), was enhanced when irrelevant as compared to relevant stimuli were presented during the working-memory retention interval, indicating that sensory processing of interference was reduced (Hakim et al., 2021). In cases where working-memory prioritisation via retro-cues was possible, other EEG studies have shown that sensory processing and encoding of interference were reduced (Barth & Schneider, 2018; D. Schneider, Barth, Getzmann, et al., 2017). This underscores that retro-cue prioritisation might not only stabilise internal representations (Hakim et al., 2020; Souza & Oberauer, 2016) but also influence interference handling.

Research on the neurophysiological basis of distractor suppression during the maintenance of internal contents remains sparse, and a clear picture has yet to emerge. There is some evidence suggesting that the PFC may not only play a role in enhancing or protecting relevant mnemonic representations during the retention interval, but also in down-regulating the gain of sensory processing of subsequent intervening stimuli (Postle, 2005). Using EEG, Chao and Knight (1998) demonstrated that patients with PFC lesions exhibit enhanced sensory processing of distracting stimuli occurring during the delay period of a working-memory task in comparison to control participants. This finding was interpreted as evidence that PFC damage impairs the

ability to inhibit distracting information from overwriting internal representations stored in sensory areas (Chao & Knight, 1998). Moreover, an invasive working-memory study conducted on non-human primates reported that the processing of a perceptual distractor presented during the delay period was more reduced in neurons in the PFC as compared to PPC, implying that the suppression of interference is supported by the PFC (Suzuki & Gottlieb, 2013). In line with the earlier finding by Chao and Knight (1998), reversible inactivation of the PFC neurons produced greater distractibility during the retention period than did inactivation of neurons in the parietal cortex (Suzuki & Gottlieb, 2013). This underscores the causal involvement of the PFC in inhibiting external interference during working-memory maintenance (but see: Jacob et al., 2016; Jacob & Nieder, 2014).

In addition to reduced interference-evoked responses, Suzuki and Gottlieb (2013) also found evidence for anticipatory suppression in the PFC prior to interference onset. The circumstances under which distractor suppression operates proactively (i.e., by pre-emptively preventing attention to interference) or reactively (i.e., by disengaging attention from interference) are highly debated topics in the literature on external attention (Geng, 2014; Noonan et al., 2018; Theeuwes et al., 2022; van Moorselaar & Slagter, 2020). Within the working-memory literature, it has been claimed that proactive interference suppression may be achieved through coordinated modulation of frequency-specific oscillations. Particularly, oscillations in the alpha band (8-12 Hz) might play a special role in setting the preparatory state for selective inhibition (Jensen & Mazaheri, 2010). For example, when the timing of visual interference can be expected during working-memory retention, occipital alpha activity increased before the anticipated distractor onset, and this increase was enhanced in anticipation of stronger versus weaker distractors, thus, suggesting that alpha activity is involved in proactive distractor suppression (Bonnefond & Jensen, 2012; Payne et al., 2013). However, other working-memory studies have only observed enhanced alpha oscillations as indicative of reactive inhibition (Liljefors et al., 2023; Y. J. Zhou et al., 2023), while some research observed reduced alpha power in

response to interference occurring during retention (Schroeder et al., 2018). Consequently, more work is necessary to determine the functional significance of alpha-band oscillatory activity in inhibiting interference proactively and reactively during working-memory maintenance. Taken together, despite the relatively limited research on interference inhibition during working memory and its underlying mechanisms, past findings suggest that inhibitory processes are essential for achieving optimal memory-guided behaviour.

1.3 Shifting Between Attentional Domains

The previously described internal-to-external and external-to-internal attentional modulations share a significant characteristic beyond their focus on both external and internal information. For a modulation to occur from one domain to the other, a shift between attentional domains is necessary. Despite extensive research on the consequences of such shifts, surprisingly little is known about *how* the brain transitions between external and internal information. Therefore, a crucial aspect of understanding the dynamic interplay between external and internal attention lies in exploring the mechanisms underlying shifts between domains. In this section, I will provide an overview of the current state of research on between-domain shifts. Given the limited literature addressing shifts between perception and working memory (and vice versa), I will draw upon findings from related fields that examine shifts between externally and internally directed cognition.

1.3.1 An Updated Framework and the Big Questions

While the widely recognised taxonomy of external and internal attention (Chun et al. 2011) has offered valuable insights regarding how the two attentional domains compare to each other, it adopts a relatively static perspective, overlooking the dynamic nature of everyday situations that often necessitate shifting between these domains. Only recently, an updated framework of external and internal attention has been proposed, acknowledging that our focus dynamically

fluctuates within and between domains (Narhi-Martinez et al., 2023). According to Narhi-Martinez et al. (2023), attention can be conceptualised as a multi-level system of weights and balances. Within both external and internal attention, a mechanism of weighting and balancing determines which competing sensory and mnemonic information is given the highest priority, respectively. Within each attentional domain, this process is likely influenced by stimulus-, goal-, and/or experience-driven sources of attentional control. On a more global level, this mechanism determines which domain holds stronger priority at any given moment, contingent upon whether external or internal signals carry greater weight. Moreover, relative attentional weightings within one attentional domain have been theorised to bias attention within the other domain, thereby also providing a theoretical framework for cross-domain attentional modulations. An attentional framework of weights and balances operating dynamically both within and across external and internal information is laudable for its ability to encapsulate the broader gamut of attention functions and further provides a promising theoretical foundation for investigating shifts between domains.

The proposal of an attentional system of weights and balances invites questions about the mechanistic implementation of transitions between sensory signals and internal representations. Indeed, there are many open questions we can ask: Do shifts between external and internal attention occur automatically through a competitive winner-takes-all mechanism of weights resetting? Alternatively, is there an additional higher-order arbitrator involved, responsible for evaluating the competition between external and internal signals and determining which should be prioritised, thereby moderating the shift? Furthermore, are there additional obstacles to overcome when shifting between domains as compared to within domains? If so, at which attentional stage do these obstacles arise, and do they impose costs? Such obstacles might hinder the attentional orienting function or leave it unimpeded but delay subsequent stages such as selecting, prioritising, or output gating. Finally, is an attentional shift from perception to memory more or less efficient than the reverse? In summary, these questions ultimately revolve

around identifying the cognitive and neural signatures of shifts between domains and determining whether they differ from shifts within a single domain.

Considering the purpose of shifts, one might argue that they represent a mere biological epiphenomenon resulting from capacity limitations in perception and working memory. This perspective would suggest that between-domain shifts serve no discernible purpose whatsoever. However, the theoretical and computational literatures have elucidated different reasons why between-domain shifting is a fundamental requirement for adaptive behaviour (Honey et al., 2017; Weilnhammer et al., 2023). One crucial reason might be the necessity to distinguish external reality from internal representations, especially within overlapping networks subserving both external and internal states (Honey et al., 2017). Furthermore, transitioning between domains may be required for driving neural plasticity and facilitate learning by internal states informing external states and vice versa. In fact, it is theorised that between-domain transitions play a pivotal role in facilitating learning across various timescales, ranging from tens of milliseconds to the circadian sleep-wake cycle (Honey et al., 2017; Weilnhammer et al., 2023). Given the rapid pace at which shifting between domains may occur, we can imagine attention swiftly and automatically oscillating between recurrent intervals of external and internal information (Honey et al., 2017). These rapid intrinsic transitions could be nested within more gradual shifts between domains, guided by current priorities and bottom-up environmental cues. A system of weights and balances, as proposed by Narhi-Martinez et al. (2023), could provide a suitable framework to account for the latter.

As a first step towards understanding how attention operates between the external and internal domains, I will consider intentional between-domain shifts, which will serve as the central theme throughout Chapters 3, 4, and 5. In particular, Chapter 3 explores how we shift back to internal contents after engaging in an external task, while Chapter 4 and 5 focusses on how attentional shifts between domains might differ from those within domains.

1.3.2 The Little We Know From Behaviour

Initial empirical evidence for the costs associated with between-domain shifts come from list-completion studies conducted in the late 1980s and early 1990s (Carlson et al., 1993; Dark, 1990; Weber et al., 1986). Weber et al. (1986) developed a task wherein participants were presented with two separate displays, each containing either a series of letters or asterisks. In the perception condition, the presentation of three asterisks in the first display (e.g., * * *) was followed by the presentation of three letters in the second display (e.g., L H F), which participants simply had to read out as quickly as possible. Conversely, in the memory condition, the first display contained three letters to be memorised (e.g., D T R), while the second display only comprised of three asterisks (* * *). Upon the onset of the second display, participants were required to name the memorised characters as quickly as possible. Lastly, in the alternate condition, participants were presented with letters in the first (e.g., D T R) and the second display (e.g., L H F). In this condition, participants had to alternate as quickly as possible between letters in perception and memory once the second display was presented, always starting with the perceptual item (e.g. L D H T F R). After adjusting for the difference in list length between pure and mixed lists, Weber et al. (1986) observed shorter list-completion times when the second display required only perception or memory compared to the between-domain shifting condition. The authors suggested that the increased cost of shifting between domains indicates the additional engagement of a processing-intensive control mechanism responsible for selectively accessing contents when alternating between domains.

Dark (1990) refined the experimental setup introduced by Weber et al. (1986) by incorporating separate conditions for shifting from memory to perception and from perception to memory. Additionally, the task included conditions where only a single shift was necessary between domains, in contrast to Weber et al.'s (1986) study where the alternate condition consistently entailed five between-domain shifts. The results indicated that the time required for switching may vary depending on the direction of the shift, with shifts from perception to memory

taking longer than shifts from memory to perception. Strikingly, the external-to-internal switch cost disappeared when participants were cued about the to-be-relevant memory item 750 ms in advance and thus could select the appropriate domain before being required to respond to it. Based on these findings, it was argued that the between-domain switch cost was not due to attentional transitioning between external and internal tasks, but instead was reflecting memory-retrieval processes triggered by the internal task (Dark, 1990).

A study aiming to resolve the contentions between these interpretations revealed that the true effects of between-domain shifts may have been obscured by additional factors (Carlson et al., 1993). This was revealed by utilising a revised version of Weber et al.'s (1986) original list-completion task, which necessitated a manual response after each digit in a list, as opposed to a verbal response at the end of the trial. By employing this approach, the authors demonstrated that reaction times (RTs) were contingent upon the serial position of a digit within a list, rather than on the attentional domain being shifted from or to. Specifically, participants seemed to chunk item responses into pairs, resulting in elevated RTs for a current digit to be succeeded by diminished RTs for the subsequent digit. This fluctuating pattern of RTs throughout the list was observed regardless of whether perception or memory items were reported first in mixed-domain lists, thus, contradicting Dark's (1990) earlier findings. Rather than attributing performance decreases in mixed-domain lists to processes that manage between-domain shifts (Weber et al., 1986) or to memory retrieval (Dark, 1990), the observed results were interpreted as being due to control processes that organise information into digestible chunks, thus facilitating easier processing and production (Carlson et al., 1993). Information chunking was feasible since all required information for the between-domain shift was perceptually and mnemonically accessible by the time the second display appeared. This implies participants might have prepared their responses in advance, rendering it uncertain whether the shift between perception and memory occurred between two manual responses or before any response was initiated. Taken together, list-completion tasks may not be the most effective way for studying between-domain shifts, since

they likely evoke additional processing that is challenging to distinguish from the costs of shifting between perception and working memory.

The research field examining shifts between perception and memory appears to have gone dormant until recently. Verschooren, Schindler, et al. (2019) have adopted an interesting approach to revive research in this field, utilising task-switching designs to examine differences in performance when participants alternately engage in perceptual and memory tasks. In the original task design, participants performed a primary task (i.e., baseline task) requiring external attention and were intermittently interrupted by a secondary task requiring either external or internal attention. This experimental design was intended to enable comparisons between performance on the trials of the baseline task that immediately followed versus preceded the secondary tasks. At the beginning of each block, participants had to memorise four digits presented in a 2×2 grid. The baseline task involved a simple visual discrimination task where participants had to identify whether the largest opening of a geometric shape was located at the top or bottom. In the interrupting external task, participants were presented with a 2×2 grid containing four digits and were instructed to determine whether a subset of these digits met predefined criteria (i.e., being bold or italicised). Conversely, the interrupting internal task required determining whether displayed digits matched the locations of the memorised digits. The authors identified a shift cost between domains comparable in magnitude to the cost of shifting within domains. However, this task design only allowed for comparison of internal-to-external (i.e., internal-to-baseline) and external-to-external (i.e., external-to-baseline) shifts, making it challenging to generalise these findings to the overall effects of within- versus between-domain shifts. Furthermore, the interrupting external and internal tasks differed in their task demands which may have introduced additional behavioural variabilities.

To address these limitations, Verschooren, Liefooghe, et al. (2019) developed a new experimental protocol, drawing inspiration from the cued task-switching literature. In their experiment, participants alternated between trials predominantly relying on perception (i.e.,

external stimuli) or memory (i.e., internal representations) to perform a matching task. In external trials, participants had to compare a target item to one out of four items appearing in 2×2 grid. In internal trials, participants were instructed to compare a target item to one of the four memorised items. These memorised items and their corresponding positions in the grid had been learned prior to the start of each experimental block. On each trial, a cue was displayed, informing participants to use perception or memory to perform the matching task as well as the position of the to-be-compared item. In this task design, four distinct types of task switches could occur between trial *n*-1 and trial *n*: external-to-external, internal-to-internal, external-to-internal, and internal-to-external. The first two are categorised as within-domain switches, while the latter two are considered between-domain switches. Compared to previous list-completion methods, this task design provided a more controlled approach for examining between-domain switches. Each switch in Verschooren, Liefooghe, et al.'s (2019) experiment necessitated a manual response, and the unpredictable order of trials restricted the participants' ability to prepare in advance for these switches. Between-domain switch costs were quantified by comparing RTs and error rates for trials preceded by a trial in the same domain (i.e., within-domain switch) with those preceded by a trial in the alternative domain (i.e., between-domain switch). The findings revealed that switching between domains incurred higher costs than switching within domains. However, this between-domain switch cost was asymmetric. Higher costs were observed for switches from external-to-internal versus internal-to-internal as compared to switches from internal-to-external versus external-to-external. These results mirror the findings previously reported by Dark (1990). The asymmetric between-domain switch cost was validated through several follow-up experiments, by employing variations of the original task design (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019).

The associative interference account, rooted in the literature on task switching (Mayr et al., 2014), was proposed as the most suitable explanation for the asymmetric between-domain switch cost (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019).

This learning-based theory posits that associations are learned between task stimuli and attentional states resulting in interference when participants are required to switch within and between external and internal tasks. Accordingly, when transitioning between two tasks, costs are typically higher when switching to the more automatic, dominant task, compared to switching to the less automatic, non-dominant task (Kiesel et al., 2010). Because switching between versus within domains incurred a larger cost when shifting to the internal domain than to the external domain (i.e., higher cost for external-to-internal versus internal-to-internal than internal-to-external versus external-to-external), the researchers interpreted internal attention as being dominant over external attention (Verschooren, Liefooghe, et al., 2019; Verschooren & Egner, 2023). The dominant (i.e., internal) task is efficiently shielded during within-switch trials as the memory traces from the non-dominant (i.e., external) task are not strong enough to interfere. However, this is not the case for the non-dominant task, which faces a continuous stream of interference from the dominant task during within-domain switch trials. When a between-domain switch takes place, the current attentional state must be updated, thereby increasing its susceptibility to interference from the alternate domain (Dreisbach & Haider, 2008; Dreisbach & Wenke, 2011). Consequently, during between-domain switch trials, the dominant task suddenly experiences interference originating from the non-dominant task, resulting in a large performance difference when comparing within- and between-domain switch trials. In contrast, the non-dominant task experiences some level of interference from the dominant task continuously during within-domain switch trials, resulting in a smaller cost when comparing within- and between-domain switch trials. This reasoning has led to the conclusion that an internal-attention task can be shielded more efficiently than an external-attention task (at least in repetition trials). A broader hypothesis has been derived from these findings and considerations, proposing a pervasive default bias for attention (and cognition in general) to be directed to internal versus external information – the Internal Dominance over External Attention (IDEA) hypothesis (Verschooren & Egner, 2023).

Predictions derived from the associative interference account and the IDEA hypothesis exhibit some compatibility with the simpler task-set inertia account (Allport et al., 1994; Wylie & Allport, 2000; Yeung & Monsell, 2003). This account suggests that switch costs stem from residual activity or interference from the task performed in a previous trial, which can have varying effects depending on whether the previous and current tasks are dominant or non-dominant. Specifically, this account posits that switching to a strong, well-practiced task from a weaker, less-practiced task is more challenging than vice versa, thereby also offering an explanation for the asymmetric between-domain switch cost.

Beyond studies utilising adaptations of the task design introduced by Verschooren, Liefooghe, et al. (2019), two other notable investigations have examined between-domain shifts. One significant study involves an fMRI investigation where participants alternated between phases of external attention and mental imagery (Gilbert et al., 2005). This study, which will be discussed in more detail in the following section, observed higher behavioural costs for switching between domains, particularly when switching from internal-to-external rather than external-to-internal (i.e., reverse pattern of results compared to Verschooren, Liefooghe, et al. (2019)). While this line of research has been instrumental in uncovering the neural mechanisms behind these shifts, a limitation lies in the experimental design, which was not optimally controlled for interpreting behavioural effects with high precision. Another study used a task-switching approach to explore switches between external attention and self-referential thoughts (i.e., to current physical sensations or personality traits; Calzolari et al., 2022). Although the internal state investigated in this study was different from those in previous research (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Schindler, et al., 2019, 2019), the authors observed a greater cost associated for between- than within-domain shifts, which was higher when switching from external-to-internal versus internal-to-internal as compared to internal-to-external versus external-to-external (i.e., similar pattern of results as in Verschooren, Liefooghe, et al. (2019)).

The burgeoning interest in the study of between-domain shifts is exciting. However, it is important to acknowledge certain limitations in both the design of these tasks and the conclusions that can be drawn from them. Arguably, the most critical consideration is that all prior studies have employed some variation of task-switching paradigms to assess the costs of switching within versus between domains. Hence, these studies have focussed on switches between external and internal tasks that each necessitate a distinct response. However, in everyday situations, attention also frequently shifts between external and internal stimuli that do not demand immediate responses. Consequently, this leaves uncertain the mechanisms involved in scenarios requiring transitions between *contents* of perception and working memory. Furthermore, a task-switching approach may introduce additional confounding factors that are not exclusively associated with the process of shifting between domains itself. For instance, task completion entails responses based on specific stimuli, potentially leaving additional motor traces that contribute to the between-domain shift cost. Furthermore, evaluative functions related to task performance, such as error monitoring and internal feedback, may also obscured the results. Finally, in task-switching paradigms, the cognitive processes involved in shifting between domains may also be influenced by the need to learn and recall various types of stimulus-response mappings, such as associating certain external and internal items with specific motor actions.

When reviewing Verschooren et al.'s studies, which are arguably the most well-controlled between-domain-switch studies to date, a few issues still emerge. First, although the authors aimed to explore transitions of attention between perception and working memory, the internal items were likely stored in long-term memory, considering they were learned in advance and remained consistent across an experimental block (Verschooren, Schindler, et al., 2019). Second, while the experimental design may have been effective in highlighting differences in RTs, the two-alternative forced-choice (2AFC) format of the responses potentially lacked sufficient sensitivity in detecting variations in accuracy (Hautekiet et al., 2023; Verschooren et al., 2020).

Furthermore, interpreting asymmetric costs of between-domain switches through the lens of the IDEA hypothesis and the associative interference account may not be infallible. For instance, while the authors argue that their results were primarily explained by associative interference, they cannot completely discount the possibility that other factors, such as memory retrieval, may also have contributed (Hautekiet et al., 2023; Verschooren et al., 2020). Moreover, while the IDEA hypothesis suggests that internal attention is more intrusive and therefore a more potent influence on behaviour (Verschooren & Egner, 2023), compelling empirical evidence is lacking to conclusively affirm that internal attention dominates over external attention. The dominance of either external or internal attention may largely hinge on the specific goals within a current context. Lastly, the capacity of the associative interference to explain between-domain switch costs may be limited to results derived from task-switching designs. Often, we encounter scenarios where transitions between external and internal information are not tied to completing distinct tasks. Hence, it remains uncertain whether between-domain shift costs would be observed in these situations, and if so, how these costs can be explained. Taken together, Verschooren et al.'s endeavours pave the way for exploring a vast array of additional questions using alternative task designs.

1.3.3 The Very Little We Know From the Brain

Compared to the scarcity of behavioural research on attentional shifts between domains, the exploration of their neural correlates is even more limited. As far as I am aware, only a single EEG study has directly investigated the temporal dynamics of attentional shifts between perception and memory. This EEG study, carried out by Verschooren et al. (2021), adopted the same task-switching design as their initial behavioural experiment (Verschooren, Schindler, et al., 2019). To reiterate, participants were tasked with an external discrimination (baseline) task, which was sporadically and unpredictably interrupted by trials from an external- or internal-attention task. The design precluded testing between-domain shifts in both directions, given that the baseline

task consistently involved external attention. Although this task design necessitates careful interpretation of the results, it may still yield some noteworthy insights.

The start of each trial was signalled by the presentation of a cue presented in the centre of the screen. The amplitude of the visual P1 component elicited by this cue was found to be significantly diminished in external baseline trials following an internal-task trial as compared to those following an external-task trial. Prior research has attributed enhanced amplitudes of the P1 to increased early sensory processing due to top-down deployment of attention (Hillyard & Anllo-Vento, 1998; Mangun & Hillyard, 1991; Pratt, 2011). Consequently, the authors proposed that the observed reduction in the P1 component following a shift from the internal domain might indicate early competition between domains, which in turn hindered the deployment of attention. This competition might be resolved by the involvement of additional top-down control processes that might not be engaged when shifting from the interrupting external task to the external baseline task (Verschooren et al., 2021).

Evidence suggesting the involvement of a top-down control process in regulating shifts between domains can be inferred from a previous fMRI study that utilised three distinct task-switching paradigms (Gilbert et al., 2005). In all three experiments, participants alternated between phases relying on external attention and phases relying on mental imagery. These phases were analogous in terms of the visual stimuli, whether perceived or imagined, and the nature of the required responses. For instance, during the external phase of one of these experiments, participants were tasked with navigating clockwise around the edge of a visually presented shape, by indicating whether the upcoming corner necessitated a left or right turn. In the internal phase of the same experiment, the actual shape was substituted with a thought bubble. This signalled that participants should mentally visualise the shape from the previous external phase and proceed their response from the last known position. Across all three experiments, Gilbert et al. (2005) observed a transient activation of the lateral rostral PFC during shifts between domains, as opposed to within-domain shifts. Further analyses showed that this activation was consistent

during transitions both from external to internal phases and from internal to external phases, suggesting that this activation pattern was independent of the direction of the shift. Consequently, these findings led to the conclusion that the lateral rostral PFC plays a critical role in goal-directed shifts between domains, but this area does not seem to be involved in shifts that occur within a single domain. Moreover, considering the shared neural activity patterns between mental imagery and working memory (Albers et al., 2013), it is conceivable that similar mechanisms might be implicated when shifting attention between external stimuli and the contents of working memory.

Gilbert et al.'s (2005) findings laid the foundation for the formulation of the gateway hypothesis (Burgess et al., 2006, 2007). This hypothesis posits that the lateral rostral PFC functions as a supervisory attentional gateway (SAG), orchestrating the coordination of external and internal information in service of goal-directed behaviours. Burgess et al. (2006) compared the role of the SAG to a railway switch point, responsible for steering the brain's information flow towards central representations. In this metaphor, information in the brain is symbolised by a train, while the tracks represent the pathways along which this information is conveyed. The switch point does not influence the train (i.e., the information) but rather directs its route. One track, controlled by this switch point, may lead to prioritised internal information processing, while another track facilitates the information flow from and to the external environment via central representations. Burgess et al. (2006, 2007) noted that in many scenarios, access to central representations is determined through continuous competition between external and internal contents, frequently bypassing the need for the SAG. However, the SAG becomes essential when there is a need to prioritise one source of information over another, when rapidly transitioning between them, or when navigating unfamiliar situations. In such cases, the SAG acts as a bridge, mediating between internal mental processes and those that involve engagement with the external environment.

Although Burgess et al. (2006, 2007) did not explicitly mention a behavioural between-domain shift cost linked to the SAG mechanism, their findings could indirectly shed light on the potential existence of such a cost. Between-domain shift trials likely impose greater processing

demands, possibly attributable to the necessary involvement of the SAG, as attention must be redirected from previously attended external or internal information. Given that the activation of the SAG likely entails additional neural processing, the gateway hypothesis provides a plausible explanation for the previously observed costs in between-domain shift trials.

A promising direction for exploration lies in bridging these findings, which propose that the rostral PFC may operate as control mechanism during between-domain shifts, with research on the dynamics of large-scale brain networks. Although the network literature is not specifically focussed on external and internal attention per se, it has provided significant insights into the broader context of externally and internally directed cognition. For example, it has been shown that externally directed cognition is linked with heightened activity in the DFPN and a reduction in activity in the default mode network (DMN), while internally directed cognition shows an inverse pattern, characterised by elevated activity in the DMN and diminished activity in the DFPN. The DMN, initially labelled as task-negative network (Raichle et al., 2001; Spreng, 2012), has been linked to a variety of internally directed cognitive processes, encompassing mind wandering, day dreaming, self-reflection, memory retrieval, and future-event simulation – all of which are guided by information other than the immediate sensory input (Andrews-Hanna et al., 2014; Christoff et al., 2016; Denny et al., 2012; K. C. R. Fox et al., 2015; Konishi et al., 2015; Spreng et al., 2009). The antagonistic interaction between the DFPN and DMN has been observed in across various task paradigms (Greenwood et al., 2018; Kelly et al., 2008; Kim et al., 2010; Shulman et al., 1997; Weissman et al., 2006), as well as in intrinsic fluctuations of the blood-oxygen-level-dependent (BOLD) signal during rest (M. D. Fox et al., 2005). Several lines of research have suggested that the anticorrelated relationship between the DMN and the DFPN may be modulated by the lateral frontoparietal network (LFPN). Specifically, it is proposed that the LFPN may flexibly couple with either network, supporting externally or internally directed cognition, respectively (Christoff et al., 2016; Dixon et al., 2017, 2018; Spreng et al., 2010, 2013; Vincent et al., 2008). For instance, Spreng et al. (2010) observed that the LFPN exhibited significant functional connectivity with the DFPN

when performing a task requiring externally directed visuospatial attention, but coupled more strongly with the DMN to support internally focussed autobiographical planning. Furthermore, it has been proposed that the rostral PFC may be part of the LFPN (Dixon et al., 2017, 2018; Spreng et al., 2013, 2013; Zabelina & Andrews-Hanna, 2016). Building on the previously mentioned findings from Gilbert et al. (2005), this further supports the potentially significant role of the LFPN in orchestrating transitions between externally and internally directed cognition. However, it remains to be tested in future studies whether these findings also extend to shifts between external and internal attention.

Additional lines of investigation have implicated the engagement of other brain networks and regions in shifts between external and internal information. Again, it is crucial to note, that unlike Gilbert et al. (2005), many of these studies often lack task designs specifically tailored to directly target shifts between domains. Instead, these studies frequently compare conditions that necessitate externally or internally directed cognition, using these comparisons to infer the potential importance of specific brain areas and networks in domain-shifting processes. In this context, it has been suggested that the salience network (SN), especially the right fronto-insular cortex, might coordinate transitions between externally and internally directed cognition (Christoff et al., 2016; Goulden et al., 2014; Servais et al., 2023; Sridharan et al., 2008; Uddin, 2015; Zabelina & Andrews-Hanna, 2016). Moreover, a recent fMRI study demonstrated that the retrosplenial cortex (RSC) adjusted its coupling with other networks based on whether cognitive processing was externally or internally oriented. As such, this area has been proposed to operate as a pivotal hub facilitating transitions between the two domains (Y. P. Li et al., 2023).

Another interesting vein of work highlights the involvement of the hippocampus in shifts between external and internal information. Traditionally, the hippocampus has been recognised to facilitate both memory encoding and retrieval (Honey et al., 2017; Marr, 1971; Servais et al., 2023). Given that the encoding of novel experiences (i.e., external signals) and the retrieving of relevant memories (i.e., internal signals) impose contrasting demands on the hippocampus

(Sherman & Turk-Browne, 2020), it has been suggested that the hippocampus prioritises external or internal information by transitioning between distinct states (Honey et al., 2017). Only recently, the hippocampus has been more directly linked to attentional processes (Aly & Turk-Browne, 2016a, 2016b). In an fMRI study, Poskanzer and Aly (2023) investigated how the hippocampus prioritises sensory signals from the external world versus internal signals related to memories. Participants performed a search task in two different conditions; one condition primarily demanded external attention (i.e., explicitly instructed condition) while the other additionally required internal attention (i.e., memory-guided condition). The authors observed stronger functional connectivity between the hippocampus and the cholinergic basal forebrain during explicitly instructed versus memory-guided attention. In contrast, the hippocampus exhibited stronger functional connectivity with the DFPN when the task demanded higher versus lower levels of internally directed cognition. These findings have been interpreted to suggest that the basal forebrain and the DFPN may bias the hippocampus, prompting it to transition between processing external and internal sources of information (Poskanzer & Aly, 2023). An alternative interpretation was reached from a recent cued-recall study integrating intracranial EEG with high-density scalp EEG (Treder et al., 2021). These researchers suggested that the hippocampus drives the transition from perception to memory by triggering a recall cascade in different areas of DMN. However, this study only investigated shifts from perception to memory but not vice versa.

Summing up this section, the existing body of literature concerning the neural correlates of transitions between domains remains sparse. There is some initial evidence for the involvement of different neural dynamics and systems when shifting between- versus within-domains. Notably, however, Gilbert et al. (2005) and, to some extent, Verschooren et al. (2021) are among the few researchers who have directly employed between-domain shift tasks in their studies. In contrast, other investigations have implied the involvement of specific neural regions and networks indirectly, either by comparing external and internal conditions or by examining between-domain shifts in only one direction.

1.3.4 Recommendations for Studying Between-Domain Shifts

The preceding literature review underscores the challenging nature of examining attentional shifts between domains. In addition to these challenges, a fundamental given when studying between-domain shifts is that any external- or internal-attention condition involves some extent of processing within the opposite domain. Consequently, neither condition can be deemed entirely pure – both external and internal attention are inevitably influenced by both external and internal factors. For instance, directing a participant to attend to the internal domain without presenting an explicit external cue is challenging. Therefore, the investigation of shifts between external and internal attention presents a considerable complexity. I propose that meeting a set of essential prerequisites is crucial for the controlled study of between-domain shifts. Recommendations for the systematic investigation include:

1. Implementing opposing experimental conditions requiring inversely disproportional amounts of external and internal attentional demands (i.e., a higher proportion of internal attention and a lower proportion of external attention in one condition, and vice versa in another condition), thereby allowing both domains to be studied.
2. Equating the stimulus materials and motor demands of the external and internal conditions to avoid merely comparing patterns of activity associated with differential processing of sensory attributes or motor codes.
3. Including a robust within-domain control condition to draw conclusions about the potential engagement of distinct processes during between-domain shifts. Establishing such a condition necessitates encompassing both external-to-external and internal-to-internal shifts.
4. Measuring shifts in both directions – that is, from external-to-internal and from internal-to-external attention – to allow for the formulation of generalisable statements concerning between-domain shifts.

5. Ensuring balanced behavioural performance across conditions, especially in neural studies, to prevent confounding factors stemming from variations in task difficulty.
6. Avoiding discrepancies in nuisance variables such as levels of reward, motivation, and arousal associated with different domains, to avoid different levels of engagement between conditions.

Recent studies investigating between-domain switches have presented notable methodological advancements (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019; Verschooren, Schindler, et al., 2019) as compared to earlier endeavours from the late 1980s and early 1990s (Carlson et al., 1993; Dark, 1990; Weber et al., 1986). However, the most well-controlled series of experiments are those based on task-switching paradigms, wherein each trial requires a response. As mentioned in Chapter 1.3.2, this approach might not fully capture the essence of between-domain shifts, as additional confounding factors could be introduced through the utilisation of task-switching paradigms. Furthermore, real-life situations often require attentional shifts between domains without an immediate need to respond to every stimulus. Thus, while studies inspired by task-switching paradigms provide valuable contributions, they may not represent the optimal or sole pathway for future research in this area.

In the past, pre- and retro-cues have served as valuable tools for examining the (neuro)cognitive signatures of external and internal attention, respectively (Carrasco, 2011; Griffin & Nobre, 2003; Landman et al., 2003; Posner, 1980; van Ede & Nobre, 2023). Employing both pre- and retro-cues within a single trial could present a lucrative methodology for exploring between-domain shifts, while also adhering to the six recommendations outlined above.

In prior attention literature, researchers have utilised task designs with two or more pre- or retro-cues to examine potential costs associated with reorienting attention within perception or memory, respectively. For instance, Dowd and Golomb (2019) demonstrated that precision in

reporting an item remains similar between single and double pre-cue trials. In double retro-cue task designs, an initial cue indicates the item to be prioritised, but is occasionally superseded by a second cue directing attention to another item in working memory that was previously uncued. Behavioural findings suggest that reorienting attention to an initially uncued representation in double-cue trials does not diminish performance compared to a no-cue condition (van Moorselaar, Olivers, et al., 2015). Some studies even observed retro-cue benefits in double-cue trials (Guo et al., 2024), which, under certain conditions, reached an equivalent magnitude to those observed in single-cue trials (Landman et al., 2003; Q. Li & Saiki, 2014; Rerko & Oberauer, 2013). Hence, considering that information regarding the previously uncued (i.e., second-cued) working-memory item remains preserved, the implementation of double retro-cue trials might offer an effective condition to study internal within-domain shifts.

To date, studies have employed pre- or retro-cues in different task conditions to examine and compare the behavioural and neural consequences of external and internal attention, respectively (Myers et al., 2015; Panichello & Buschman, 2021; Y. Zhou et al., 2022). However, double-cueing paradigms incorporating both pre- and retro-cues have yet to be developed. In Chapters 4 and 5, I will introduce a novel task design that incorporates both pre- and retro-cues within a single trial to examine the consequences of shifting attention between external and internal contents.

1.4 Overview and Aims of the Thesis

In my doctoral research, I investigate the dynamic interplay between external and internal attention. I explore scenarios where individuals encounter both external and internal information in quick succession, necessitating rapid shifts in attention between these two domains. Specifically, I approach this topic from two distinct perspectives. First, I explore the phenomenon of cross-domain modulations. Second, I analyse the underlying processes of between-domain shifting.

The first empirical chapter of this thesis explores cross-domain attentional modulations (Chapter 2). Here, I examine how manipulations of one domain can influence the other and vice versa. My approach involves leveraging temporal expectations to modulate the attentional processing of internal representations while monitoring for potential consequences within the external domain. Through this inquiry, I also aim to shed light on whether temporally anticipating an external event offers protection for internal representations compared to scenarios devoid of prior expectations. Furthermore, I seek to discern whether potential benefits of temporal expectations emerge from mitigating external interference or safeguarding internal representations.

In Chapter 3, my focus turns to between-domain transitions, with a specific zoom-in on scenarios involving shifts from external to internal events. I examine the temporal dynamics of reactivations of internal representations following external interruptions, capturing the natural flow of shifting attention from one domain to another. Utilising EEG, I aim to pinpoint the moments when various facets of working-memory contents resurface, providing insights into the reactivation process during external-to-internal shifts.

In Chapter 4, I introduce a novel experimental approach featuring double-cueing and undertake a systematic exploration of between-domain attentional shifts. Through the integration of behavioural and gaze measures, my aim is to discern whether transitions between the contents of different domains incur greater costs as compared to shifts within a single domain. Additionally, I investigate whether potential between-domain shift costs are modulated by the time given to transition between domains as well as by spatial-reallocation demands.

In Chapter 5, my research focusses on identifying the temporal dynamics of between-domain shifts by utilising MEG. I employ state-of-the-art decoding methods to discern differences in neural dynamics between within- and between-domain shifts. Furthermore, I assess from which attentional stage these differences might originate, enhancing our understanding of attentional processes underlying such shifts.

Lastly, Chapter 6 provides an integrative overview of the entire body of work, highlighting the novel insights into the dynamic interplay between external and internal attention presented in this thesis. Moreover, this final chapter serves as a platform for discussing the future trajectory of this field of research.

2 Cross-Domain Modulations via Temporal Expectations

2.1 Abstract

In this chapter, I explore attentional effects across the external and internal domains. Specifically, I utilised temporal expectations to modulate attentional processing within one domain while monitoring for potential consequences within the opposing domain. In the experiments detailed in this chapter, participants were required to maintain information in working memory while encountering an external event during the retention interval. The critical manipulation concerned the attentional domain to which the temporal expectations were applied. Chapter 2.3 presents an experiment investigating whether temporally guided internal attention influences performance in an interfering task requiring external attention. The results indicate that anticipating when internal contents will be required for an upcoming working-memory task can not only improve working-memory guided behaviour but also affect performance in an external task occurring during the period of anticipation. In Chapter 2.4, I adopted the opposite approach to investigate whether predicting the timing of an external event could alleviate its negative impact on subsequent working-memory guided behaviour. The findings reveal that anticipating the timing of external interference can effectively mitigate detrimental effects thereof on the quality of internal representations. Furthermore, the results suggest that temporal expectations enable a protective mechanism which shields internal content from external interference, rather than suppressing such interference. Collectively, the empirical work presented in this chapter demonstrates that temporal expectations have far-reaching consequences, affecting not just the domain to which they are directly applied but also extending their influence on the opposite domain. This suggests a significant cross-domain modulation attributable to temporal expectations.

The findings presented in Chapter 2.3 have been published as Gresch et al. (2022).
The findings presented in Chapter 2.4 have been published as Gresch et al. (2021).

2.2 Introduction

Sensory input in the external world and internal representations in working memory can acquire relevance in rapid succession, causing us to continually shift between the two domains of attention. During this process, the contents of external and internal attention often engage in a reciprocal exchange, collaborating towards achieving a shared task objective. Revisiting a previous example, when following furniture assembly instruction, external information is absorbed and stored in working memory. These internalised representations then steer our search through the various furniture pieces, aiding us in locating the exact part needed for the next stage in the assembly process.

However, the flow of information between the external and internal domain in pursuit of a common objective does not always follow a straightforward, linear process. Instead, we frequently engage in multiple subtasks within a common time frame, where attentional shifts between external and internal inputs serve broader, but not necessarily directly linked, purposes. For instance, while preparing a Sunday roast, numerous subtasks demand our attention, such as making the sauce, roasting the meat, and preparing side dishes. To ensure all these elements are ready simultaneously, we must adeptly alternate between focusing on external and internal contents within each subtask, while also smoothly transitioning from one subtask to another.

Previous literature has shown that we are susceptible to cross-domain modulations in these multi-task scenarios. Among various findings, past research has indicated that perceptual processing (Kang et al., 2011; Macedo-Pascual et al., 2023; Teng & Kravitz, 2019) as well as the rhythmic allocation of attention to external stimuli (Balestrieri et al., 2022) can be modulated by items held in working memory. Moreover, studies on memory-guided attentional capture have shown that objects of a certain colour or shape held in working memory draw our attention towards similar-coloured or-shaped stimuli in a subsequent, unrelated visual-search task (Y. Huang et al., 2019; Kiyonaga et al., 2012; Olivers, 2009; Olivers et al., 2006; Sasin & Fougnie, 2020; Soto et al., 2005, 2008). Relating this to the roast example, when keeping in mind that we need to add

salt to the sauce, the internal content might inadvertently steer our attention towards the saltshaker while we are actually looking for a knife to chop vegetables. In addition to internal-to-external modulations, a vast literature has demonstrated that the stimuli we perceive or interact with while maintaining contents in working memory can lead to distortions and biases in internal representations (Lorenc et al., 2021).

This empirical chapter aims to uncover attentional effects across domains in multi-task scenarios; specifically, how attending to internal contents can modulate external attention, and vice versa. The work presented here diverges from previous studies investigating whether and how external or internal contents impact each other, by introducing an additional dimension to this dynamic interplay. Specifically, I employ predictable temporal structures as a tool to modulate the attentional processing of external sensory input (Chapter 2.4) and internal representations (Chapter 2.3) while monitoring for potential consequences within the respective opposing domain.

Temporal expectations facilitate upcoming perception and action by orienting attention to relevant information occurring at predictable moments in time (Denison, 2024; Nobre, 2001; Nobre & van Ede, 2018, 2023). Such expectations can be formed based on different types of temporal structures that exhibit predictability. For instance, the effectiveness of temporal associations in guiding attention has been demonstrated through temporal-cueing tasks (Correa et al., 2004; Coull & Nobre, 1998; Denison et al., 2017; Griffin, 2002). Drawing inspiration from Posner's spatial orienting task (Posner, 1980) and building upon the extensive body of foreperiod literature (Los, 2010; Niemi & Näätänen, 1981), these tasks utilise symbolic cues to forecast at which time point a target is likely to appear. When targets appear at intervals accurately predicted by these cues, a significant enhancement in performance is typically observed. Notably, temporal expectations can enhance perception not only when predictions are explicitly instructed, but also through implicitly learned associations (Cravo et al., 2017; Los et al., 2017; Olson & Chun, 2001). Moreover, temporal expectations can arise from rhythmic occurrences at fixed intervals (Jones,

2019; Mathewson et al., 2010; Rohenkohl et al., 2012) or from reoccurring sequences of temporal intervals (Heideman et al., 2016; Heideman, van Ede, et al., 2018; O'Reilly et al., 2008). Lastly, temporal expectations can also be based on the probability of an event occurring at a specific time, given that it has not occurred yet (Luce, 1991). These expectations are behaviourally evident in quicker RTs to stimuli appearing at highly predictable moments (Coull, 2009; Grabenhorst et al., 2019; Vangkilde et al., 2012).

Beyond biasing attention during perception, temporal expectations also operate in working memory, prioritising internal representations at times when they are anticipated to be most relevant for guiding behaviour (Boettcher et al., 2021; Heuer & Rolfs, 2021; Jin et al., 2020; Olmos-Solis et al., 2017; van Ede et al., 2017; Zokaei et al., 2019). Supporting evidence can be drawn from a change-detection experiment in which the interval between encoding and probe displays was systematically manipulated across blocks (Jin et al., 2020). During expect-early blocks, the majority of trials were characterised by a brief memory delay, while the remaining trials featured a longer delay. This ratio was reversed in expect-late blocks. Both accuracy and RTs showed improvement during shorter memory delays in expect-early versus expect-late blocks, mirroring the performance enhancements commonly seen in perceptual tasks (Nobre & van Ede, 2018). Furthermore, research has indicated that temporal expectations can become associated with individual items held in working memory. This association enables dynamic shifts in prioritisation, moving from enhanced accuracy and accessibility for items expected to be relevant early to those anticipated later as time progresses (van Ede et al., 2017; Zokaei et al., 2019). In line with these findings, a recent eye-tracking study revealed that in a dynamic visual-search task, temporal expectations facilitate attentional shifts between independent search templates by flexibly prioritising features of the most imminently anticipated external target (Williams et al., 2023).

To date, research has primarily investigated the consequences of temporal expectations on either perception or working memory, with these attentional domains being examined in

isolation (but see: Capizzi et al., 2012, 2013). This fails to capture the dynamic interplay of external and internal attention faced during everyday activities, in which we are often required to juggle external and internal demands from multiple tasks across a shared time period. In these scenarios, we may be able to anticipate the timing of an upcoming event in one domain, while concurrently having to engage in another task in the other domain. In the context of the Sunday-roast example, our ability to anticipate the precise moment to add specific ingredients to the sauce might influence how we attend to other external and internal events.

In this empirical chapter, I aim to investigate whether predictable temporal structures within one domain influence performance in the other domain. The first experiment investigates the impact of temporally guided internal attention on external attention, whereas the second experiment explores the impact of temporally guided external attention on internal representations.

2.3 Experiment 1: How Temporally Guided Internal Attention Modulates External Attention

In this experiment, I investigate whether performance in a working-memory task still benefits from temporal expectations when an external intervening task occurs within the period of anticipation. I additionally explore whether the predictable temporal structure of the working-memory task also influences performance in the external intervening task, thereby affecting overall performance when juggling contemporaneous tasks.

To address these questions, I inserted an external-attention task within the retention period of a visual working-memory task. Between blocks, the duration of the retention interval was either predictable (i.e., 100% certainty to be short or long) or unpredictable (i.e., 50% short, 50% long). Critically, the temporal-expectation manipulation only applied to the time when the items in working memory would be probed for report. In contrast, the intervening task always occurred at the same time after working-memory encoding, irrespective of the expected memory-

probe time. Accordingly, any effects on the performance in the external task would reflect a between-task consequence of temporal expectations. My primary aim in the current study was to demonstrate proof-of-concept as to whether any such between-task consequences of predictable temporal structures exist. In the upcoming sections of Chapter 2.3, the terms "between-task consequences" and "cross-domain modulations" will be used interchangeably, as the two tasks within a trial vary with respect to the targeted attentional domain.

2.3.1 Methods

2.3.1.1 Participants

An initial power analysis conducted in G*Power (Faul et al., 2007) aimed at detecting medium effects (d = 0.5, α = 0.05, 1-β = 0.95) suggested a sample size of *n* = 54 To yield the targeted number of participants, I collected data from 75 online participants in total. Data from 20 participants were excluded following an a-priori trial-removal procedure before splitting data by conditions. One additional participant was also removed for using a non-memory-based strategy to complete the task (see 2.3.1.3 Analysis for details). Out of 54 participants (age range: 19 to 40; mean age: 28.37; 15 females; 38 males, 1 non-binary), 45 self-reported to be right-handed (9 left-handed).

The study was approved by the Central University Research Ethics Committee of the University of Oxford. Participants were recruited via Prolific Academic (https://www.prolific.co/) and pre-screened based on demographic criteria (age range 18 to 40, fluent in English), general health (normal or corrected-to-normal vision, no history of mental illnesses), and previous participation history on Prolific Academic (participated in at least 10 studies with a study approval rate above 90%). All participants provided informed consent prior to participation and were paid £7.50 per hour. An additional monetary reward of up to £5 could be earned depending on participants' task performance in the memory task. Specifically, performance above 80% received

a bonus payment scaling from £0.01 at 80% to £5 at 100%, with an average bonus payment of £0.67 (*SD* = 0.85) across all participants.

2.3.1.2 Task and Procedure

Participants performed a web-based visual working-memory task requiring them first to memorise the angles of two oriented bars and then, when probed after a retention interval, to reproduce the exact angle of one of those memory items. In addition, participants had to perform an intervening task between encoding the memory array and retrieving the probed item (**Figure 2.1A**). Two critical experimental manipulations concerned the time at which the memory probe appeared. First, the probe could occur after either a short (1250 ms) or long (2500 ms) retention interval (delay condition: early vs. late). Second, in fixed (predictable) blocks, the memory probe would always occur at the same time, whereas in variable (unpredictable) blocks, it would occur unpredictably (block type: fixed vs. variable). This resulted in four possible trial conditions: an early-fixed, early-variable, late-fixed, or late-variable onset of the memory probe. Critically, the intervening event always occurred at the same time after the disappearance of the initial encoding display (1000 ms), independent from the memory-probe time.

Participants completed the experiment in a web browser on their personal computers. The recommended internet browsers were Mozilla Firefox and Google Chrome. Participating via mobile phone or tablet was not allowed. Prior to the experiment, participants' individual screen resolution was estimated by asking them to adjust an image of a credit card such that it matched the size of a physical credit card placed on the screen. In this manner, I could calculate the ratio between the card-image width in pixels and the actual card width in centimetres to obtain a measure of pixel density (pixel per cm). Together with the instructed viewing distance of approximately 60 cm (one arm's length away from the monitor), this allowed stimuli to be presented in degrees of visual angle, regardless of monitor size (Q. Li et al., 2020). The

experimental script was generated in PsychoPy (Peirce et al., 2019) and hosted online using Pavlovia (https://pavlovia.org/).

At the start of each trial, two tilted bars were simultaneously presented against a grey background (RGB value: [128,128,128]) for 250 ms. One bar was always positioned to the left and the other to the right of the central fixation cross. Independent of location, one of the bars was tilted to the left (anticlockwise) and the other to the right (clockwise). To avoid angles too close to vertical or horizontal meridians, the items' angles were randomly drawn in increments of ±5° between ±5° and ±85°. Across trials, a leftwards or rightwards oriented bar was equally likely to appear in the left or right position of the screen. The stimuli were approximately 0.8° in width and 5.7° in length and were centred at a viewing distance of approximately 5.7° visual angle from fixation. At encoding, both lateralised items were equally likely to be probed, rendering them equally relevant.

Encoding array disappearance was followed by a delay period of 1000 ms, during which the fixation cross remained on the screen. After this delay, the intervening item (a tilted bar) was presented in the centre of the screen for 250 ms. Participants were required to respond to the intervening item by pressing the F key with their left index finger if the bar was tilted to the left, or the J key with their right index finger if the bar was tilted to the right.

Within blocks, the tilt of the intervening items was pseudo-randomised such that a leftwards or rightwards tilt occurred equally often. The intervening item was always presented in a different colour than the two memory items that preceded it. The colour of the memory items and the intervening item were always drawn from a set of three highly distinguishable colours (RGB values: blue [0,225,228], orange [254,163,0], pink [253,142,253]). The colours used for the memory items and the intervening item varied randomly across trials. The intervening item had the same size as the memory items and its angle was also drawn between ±5° and ±85° in increments of ±5°. Across trials, the orientation of the intervening item could be either leftwards or rightwards, with both orientations having an equal likelihood.

The offset of the intervening item was followed by a second delay period (1250 ms or 2500 ms) before the fixation cross changed to match the colour of one of the two memory items, to indicate which memorised tilt should be reported. Importantly, the colour of the probe would never match that of the intervening item. In fixed blocks, the second delay interval had a fixed duration of 1250 ms (early blocks) or 2500 ms (late blocks) across all trials. In variable blocks, half of the trials had a retention interval of 1250 ms, whereas the other half had a memory delay of 2500 ms. The order of trials within variable blocks was randomised.

For the working-memory task, participants had to try to reproduce the exact angle of the probed memory item. As such, in contrast to the intervening task requiring simple discrimination (leftwards or rightwards), the memory task demanded a precision response. Following the appearance of the probe, participants had unlimited time to decide on their response. After response initiation, a visual response dial was displayed on the screen, always starting in the vertical position. The response dial included markers along a circle that corresponded to the ends of a bar and always appeared surrounding the fixation.

To report a leftwards (rightwards) angle, participants were (as for the intervening task described above) asked to press the F or J key on the keyboard using their left or right index finger, respectively. The dial rotated leftwards when pressing F and rightwards when pressing J (either by holding the key down or pressing the key repeatedly; always in increments of ±5°). The dial could only be rotated in the direction that was initially indicated by the participant. For example, if a participant started pressing the F key after the probe, the dial would only move leftwards, and it would therefore not be possible to move the dial towards the right with the J key. Since the response dial always started in the vertical position and because it could not be rotated beyond ±90°, a leftwards (or rightwards) oriented bar could only be correctly reported with the left (or right) key. Consequently, the hand required for responding was directly linked to the angle of the bar that was probed. This builds on previous tasks (Boettcher et al., 2021; van Ede, Chekroud, Stokes, et al., 2019), though the specifics of this response implementation were not essential to

the current study. Once participants started rotating the dial, they had limited time (4000 ms) to complete the angle reproduction. This was intended to encourage participants to recall the exact orientation before starting to move the dial. When the dial aligned with the remembered tilt of the item, participants pressed the space bar to confirm their response and continue with the task.

At the end of each trial, participants received feedback regarding their working-memory performance and, when relevant, also about their intervening-task performance. The working-memory feedback provided information on how well participants reproduced the probed item. Feedback was presented for 500 ms in the form of a number ranging from 0 to 100, with 100 indicating a perfect report and 0 indicating that the adjusted orientation was perpendicular to the angle of the probed item. However, if the time to adjust the angle ran out, the message "Too slow" was presented instead for 750 ms. Additional feedback could also appear to indicate if participants responded with the wrong key or did not respond at all to the intervening item. To incentivise fast responses in the intervening task, participants also received a feedback message when their RTs were slower than 1000 ms. This feedback message was combined with an image reminding participants to press F (or J) when the intervening item was tilted to the left (or right). Feedback was presented for a minimum of 750 ms and remained on screen until the space key was pressed, encouraging participants to read the feedback message before proceeding to the next trial. Trials were separated by an inter-trial interval (ITI) randomly drawn between 500 and 800 ms.

The experiment consisted of 384 trials divided across 12 blocks (each including 32 trials). In six blocks, the delay between intervening-item offset and probe onset was fixed (predictable) – with the probe occurring early in three of the blocks, and occurring late in the other three. In the remaining six blocks, probe onset was variable (unpredictable; pseudo-randomly varying between early and late within each block). Therefore, the total number of trials in which the probe appeared at any given delay interval after the offset of the intervening item (either early or late) remained equal across both the fixed and variable blocks. The order of blocks was pseudo-

randomised in groups of four containing two variable blocks, one fixed-early, and one fixed-late block.

To become familiarised with the procedure of the experiment, participants performed 32 practice trials each with an unpredictable delay period. Participants were informed that they would never have to reproduce the tilt of the intervening item. However, they were not informed about the block type (fixed vs. variable) or the two possible probe-onset times (early vs. late). The instructions also stressed that for both the intervening and working-memory task, participants should respond as quickly and accurately as possible. At the end of the experiment, participants were redirected to the survey website Qualtrics (http://www.qualtrics.com/), where they were asked about their comprehension of the instructions, potential strategies used to complete the task, and whether or not their data should be analysed. The experiment lasted approximately 50 minutes in total.

2.3.1.3 Analysis

Data were analysed in R Studio (Posit team, 2022). During pre-processing, trials were removed when RTs in the working-memory task (calculated from probe onset to response onset) were below 200 ms or exceeded 5000 ms. Next, I removed trials for which the remaining RTs were 2.5 *SD* above the individual mean across all conditions or if participants took longer than 4000 ms to reproduce the probed angle after response initiation. Regarding the intervening task, I excluded trials if participants either did not respond at all or if they did not respond within a time window ranging from 200 ms to 1500 ms after the intervening-task onset. Datasets with more than 25% of trials rejected during these pre-processing steps or with average reproduction errors higher than 40° in the working-memory task (across all conditions) were removed from further analysis (n = 20). Additionally, one dataset was also removed in which the participant self-reported as having employed an explicit non-memory-based strategy to maintain the encoding display (i.e., physically aligning their fingers with the memory items at encoding). After this exclusion step,

datasets from the remaining 54 participants (in which an average of 95.18% [*SD* = 2.78] of trials were retained) were entered into the main analysis.

For the working-memory task, I examined the average RTs for the conditions fixed-early, variable-early, fixed-late, and variable-late. Moreover, I also evaluated reproduction errors by averaging the absolute difference between the original angle of the target (probed) item and the reported angle. For the intervening task, I analysed the average RTs for the fixed-early, fixed-late, and the variable condition. I did not split the variable condition by early and late trials, as at time of intervening-task onset, it was unknown whether the working-memory probe would occur early or late. For the same conditions, I also calculated the average error rates. Participants committed an error when using the wrong key to respond to the intervening item. Since I expected error rates for this simple discrimination task to be quite low, RTs were deemed the more sensitive dependent variable for the intervening-task performance. To examine potential sequential effects in variable blocks, I analysed RTs and error rates as a function of the delay condition associated with the working-memory task in the previous trial (previous probe early vs. previous probe late). Unlike classic sequential effects that are considered within single-task situations (Los, 2010), I here investigated potential sequential effects of the preceding delay of the working-memory task on performance in the intervening task that always occurred at the same time after memory encoding.

When comparing more than two means, I applied a repeated-measures analysis of variance (ANOVA) and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples *t*-test and report Cohen's *d* as a measure of effect size. For post-hoc *t*-tests, I report Bonferroni-corrected *p* values that I denote as $p_{Bonferroni}$. I used the ggplot2 package (version 3.3.3; Wickham, 2009) for plotting results. Where relevant, the within-subject standard error of the mean was calculated from normalised data using the approach from Morey (2008).

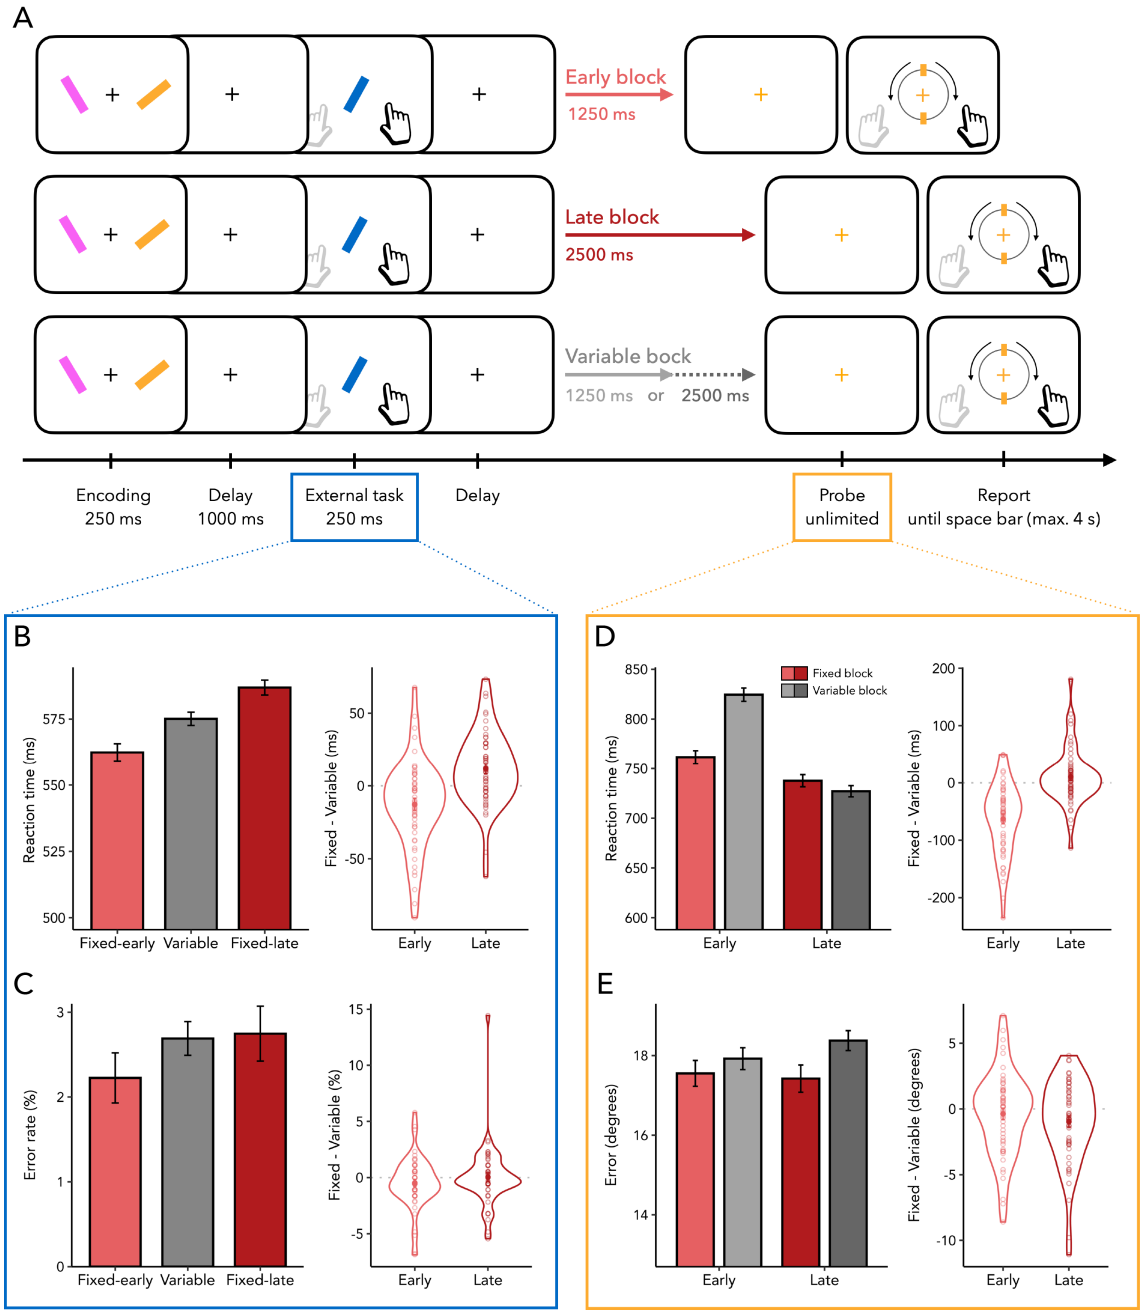

Figure 2.1. Task schematic and results. **(A)** Two lateralised oriented bars were presented for 250 ms. For the working-memory task, participants had to remember the angle of both bars, of which one had to be reported at the end of the trial. Following a delay of 1000 ms, an intervening task occurred wherein another tilted bar was presented centrally for 250 ms. For this external intervening task, participants had to indicate as quickly as possible whether this bar was tilted to the left or to the right. In early blocks, the intervening item offset was followed by a delay of 1250 ms, whereas in late blocks, it was followed by a delay of 2500 ms. In variable blocks, the delay between the offset of the intervening item and the onset of the memory probe was equally often either 1250 ms or 2500 ms. After this delay, a colour change of the central fixation cross indicated which of the two bars from the encoding display would have to be reproduced from memory. **(B)** External-task performance. Left panel: RTs in the external task significantly increased from fixed-early to variable and from variable to fixed-late blocks. Right panel: Cross-domain modulations effect of temporal expectation in fixed-early and fixed-late blocks, relative to variable blocks. **(C)** Left panel: Error rates did not differ between conditions. Right panel: Temporal expectations did not influence error rates. **(D)** Working-memory performance. Left panel: RTs to early probes were faster when the probe could be temporally predicted as compared to when it was temporally unpredictable. Right panel: Temporal expectations reduced RTs in the early condition. Violin plots for each delay condition were calculated by taking the difference in RTs between fixed variable blocks (fixed minus variable). **(E)** Left panel: Reproduction errors were significantly higher in variable versus fixed blocks. Right panel: Temporal expectations decreased reproduction errors for both delays. Dots represent individual participants. Error bars indicate ±1 *SEM*.

2.3.2 Results

2.3.2.1 Working-Memory Performance

I first confirmed that the manipulation of temporal predictions for the working-memory task was effective despite an external intervening task occurring in the period of anticipation. For this, I evaluated RTs, which I defined as the time between the onset of the memory probe and response initiation. RTs served as a proxy for the time it took participants to access the relevant memory information before starting to reproduce the probed angle. I found a significant main effect of delay condition and block type (**Figure 2.1D**). RTs to the probe were faster in late as compared to early trials ($F_{(1,53)}$ = 62.517, p < 0.001, η^2_G = 0.024) and when probe onset was fixed compared to variable ($F_{(1,53)}$ = 22.491, p < 0.001, η^2_G = 0.005). These two main effects were paired with a significant interaction between the delay condition and block type ($F_{(1,53)}$ = 48.396, p < 0.001, η^2_G = 0.009). This interaction showed that temporal predictions conferred a significant benefit (i.e., led to faster response initiation) for early probes ($t_{(53)}$ =-7.437, $p_{Bonferroni}$ < 0.001, d = 1.012), but not for late probes ($t_{(53)}$ = 1.568, $p_{Bonferroni}$ = 0.491, d = 0.213). This finding is typical for studies of temporal expectations (as reviewed in Nobre & van Ede (2018)) and can be attributed to the fact that, once the early interval passes, participants can always anticipate that working memory will be probed at the later interval, irrespective of whether the current block is fixed or variable. Moreover, pairwise comparisons revealed that participants responded more slowly to early as compared to late memory probes in variable blocks ($t_{(53)}$ = 10.633, $p_{Bonferroni}$ < 0.001, d = 1.447), whereas the difference in RTs in early-fixed versus late-fixed blocks did not reach significance ($t_{(53)}$ = 2.531, $p_{Bonferroni}$ = 0.058, d = 0.344).

It should be noted that the response window for the intervening task (relative to memory-probe onset) differed in trials with early versus late memory probes. Although I only analysed trials in which participants responded within 1500 ms after intervening-task onset, it is possible that RTs to the working-memory task were affected by the recovery time available between the intervening- and memory-task responses in early as compared to late trials. To ensure that the

findings could not be attributed to this difference, I reran the pre-processing and analysis while excluding slower responses to the intervening task (i.e., RTs > 1000 ms rejected, with 8.42 ± 8.80% trials removed per participant). The observed benefit of temporal prediction remained the same even when ensuring a minimum recovery time of 500 ms (**Supplementary Table 8.1**, **Supplementary Figure 8.1**).

Next to RTs, I additionally considered the influence of temporal predictions on the quality of working-memory reports. To this end, I analysed the average reproduction error (i.e., the absolute deviation between the reported angle and the true angle of the probed item), for which lower values indicate better performance. I observed a significant effect of block type on reproduction errors (**Figure 2.1E**; $F_{(1,53)} = 4.854$, $p = 0.032$, $\eta^2_G = 0.002$), with smaller errors when the memory probe occurred in a fixed versus a variable block. In contrast to RTs, however, for errors I found no systematic difference between early and late probes ($F_{(1,53)} = 0.313$, $p = 0.578$, $\eta^2_G < 0.001$) nor an interaction between block type and delay condition ($F_{(1,53)} = 0.865$, $p = 0.357$, $\eta^2_G < 0.001$).

These data reveal that temporal expectations were used to anticipate upcoming memory-guided behaviour, building on prior demonstrations of similar effects in the absence of intervening tasks (Jin et al., 2020; van Ede et al., 2017; Zokaei et al., 2019). The current results show that benefits of temporal expectation on working-memory guided behaviour occur even when an external intervening task must be completed during the retention interval.

2.3.2.2 External-Task Performance

Next, I asked whether temporal expectations regarding the subsequent working-memory probe influenced performance of the intervening task. Critically, the intervening task always occurred at the same time after encoding, so any potential performative differences in this task should reflect a cross-domain consequence of temporal expectation regarding the working-memory task. Note that I did not divide variable blocks by subsequent delay intervals because participants could not

have known whether the memory items would be probed early or late at the time of the intervening task.

As shown in Figure 2.1B, responses to the intervening task were fastest when the memory probe was expected early, slowest when the memory probe was expected late, and intermediate when the memory probe occurred unpredictably. This pattern of results was supported by a significant main effect ($F_{(2,106)}$ = 18.282, p < 0.001, η^2_G = 0.007). Pairwise comparisons showed significantly faster RTs in fixed-early as compared to both the variable ($t_{(53)}$ =-3.067, $p_{Bonferroni}$ = 0.010, d = 0.417) and the fixed-late blocks ($t_{(53)}$ =-5.427, $p_{Bonferroni}$ < 0.001, d = 0.739), as well as significantly faster RTs to variable as compared to fixed-late blocks ($t_{(53)}$ =-3.443, $p_{Bonferroni}$ = 0.003, d = 0.468). These results show that temporal expectations do not merely influence performance in the task to which they apply (here, the working-memory task); they also affect performance in another task that occurs during the period of temporal anticipation – even when this task itself always falls within the exact same timeframe after memory encoding.

Consistent with the RTs, error rates in the intervening task (Figure 2.1C) were numerically smallest in fixed-early and largest in fixed-late blocks. However, for error rates I did not find a significant difference between conditions ($F_{(2,106)}$ = 1.060, p = 0.350, η^2_G = 0.004). This may reflect intervening-task performance being very close to ceiling (i.e., error rates below 3%). Hence, I refrained from further analysing and interpreting accuracy of the intervening task.

2.3.2.3 External-Task Performance: Sequential Effects

After having demonstrated that temporal expectations affected intervening-task performance in fixed blocks, wherein participants are persistently exposed to the same temporal structure, I asked whether temporal associations over a short timescale can also guide performance. To this end, I analysed sequential effects in variable blocks by comparing RTs and error rates to the intervening task as a function of the working-memory delay in the immediately preceding trial (n-1). In contrast to classical sequential effects examined in simple RT tasks (Los, 2010), I here tested for

potential sequential effects of the previous working-memory delay on performance in the intervening task.

As shown in **Figure 2.2A**, RTs to the intervening task were faster when the working-memory probe in trial *n*-1 occurred early as compared to when it occurred late ($t_{(53)}$ =-2.937, p = 0.005, d = 0.400). The temporal structure associated with the working-memory delay of the previous trial thus had an immediate effect on RTs in the intervening task of the current trial – consequently, demonstrating a sequential-effect influence on the intervening task. Error rates showed a similar numerical pattern (**Figure 2.2B**), though pairwise comparisons did not reach significance ($t_{(53)}$ =-0.580, p = 0.564, d = 0.079).

Thus, the temporal interval of the working-memory task in the preceding trial significantly affected performance of the intervening task in the current trial. Nevertheless, the effects of blocked temporal predictions were stronger (**Supplementary Figure 8.2**). The difference in RTs between early versus late trials in fixed blocks was significantly larger than the difference in RTs following trials with early versus late working-memory probes in variable blocks ($t_{(53)}$ =-3.260, p = 0.002, d = 0.444).

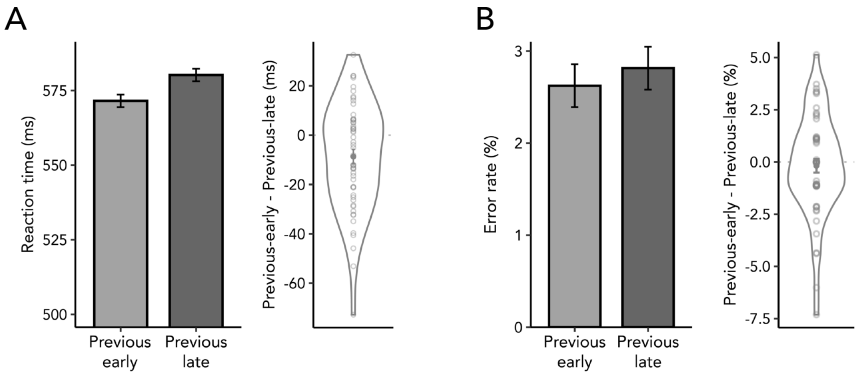

Figure 2.2. Sequential effects of temporal intervals on external-task performance (A) Left panel: RTs in variable blocks of the external task were significantly faster when the immediately preceding working-memory probe occurred early versus late. Right panel: Violin plots showing the difference in RTs between trials in which the working-memory delay of the preceding trial was early versus late (previous-early minus previous-late) **(B)** Left panel: Error rates in variable blocks did not differ significantly depending on the working-memory delay interval of the previous trial. Right panel: As in the right panel in (A), with the difference in errors rates instead of RTs. Dots represent individual participants. Error bars indicate ±1 *SEM*.

2.3.3 Interim Discussion

By manipulating the predictability of temporal structures within a multi-task context involving both external and internal attentional demands, I made two relevant observations. First, temporal expectations about when to utilise working-memory contents confer significant benefits to memory performance even when other external intervening tasks must be performed in the interim. Second, temporal expectations about the later working-memory task also had significant consequences for performance in the external intervening task, even though the temporal-prediction manipulation did not apply to this task. Responses to an intervening task were expedited when the requirement for memory-guided behaviour was expected to occur early and slowed when the working-memory task was expected late, as compared to when the timing of the working-memory task was unpredictable. Thus, this experiment unmasked cross-domain consequences (i.e., between-task consequences) of predictable temporal structures.

2.3.3.1 Cross-Domain Consequences of Predictable Temporal Structures

This work clearly demonstrates the presence of a cross-domain modulations via predictable temporal structures in a multi-task setting. It invites consideration of what mechanisms could account for the effects, and this remains an important direction for future research. At this juncture, I can only speculate regarding the putative mechanisms.

First, it is possible that the between-task effect may result from more efficient scheduling of both the external and internal tasks when the timing of either of them is predictable. In line with this, Kushleyeva et al. (2005) proposed that, in multi-task scenarios, temporal aspects of each task are used to flexibly schedule when to prioritise one task or the other. Given there was consistently less time between the intervening item and the memory probe in the fixed-early blocks, participants may have opted to complete the intervening task as soon as possible to refocus on the working-memory task in time for the probe. In contrast, when the working-memory

probe was consistently expected late – yielding less time pressure – participants may have allowed themselves more time for the intervening task, rendering their responses slower.

Another possible interpretation is that an elevated state of preparedness or alertness may have accompanied temporal expectations for an early working-memory probe, whereas a diminished state of alertness may have prevailed when the probe was expected late (Shalev & Nobre, 2022; Weinbach & Henik, 2012). The differential overarching state of preparation or alertness may have effectively spilled over to the external intervening task. An alternative spill-over interpretation relates to a suggestion in the time-perception literature that cross-contamination between temporal intervals occurs when more than a single interval is estimated simultaneously (Moon & Anderson, 2013; Taatgen & Van Rijn, 2011). Specifically, representations of shorter or longer intervals can shift the representation of an intermediate interval in those respective directions. Such a contamination effect may also have contributed to the observed between-task effect of predictable temporal structures. For example, in fixed-late trials, the representation of time used for guiding performance in the intervening task may have expanded and thereby slowed responses. Likewise, faster RTs in fixed-early trials may have been mediated through compression of time representations. Future studies should take a more granular, and complementary physiological approach to understand exactly how these factors contribute to the between-task consequence of temporal expectations that I exposed here.

In addition to the scenarios above, it may have been easier to anticipate the intervening task when the working-memory task itself was also temporally predictable. However, this is unlikely to account for the current results, as this would have led to faster RTs to the intervening task regardless of whether the working-memory probe occurred early or late. In contrast, I found that it was not the temporal predictability per se that affected performance of the intervening task, but rather the time at which the subsequent probe was expected in the fixed blocks (early vs. late). Yet, it remains an interesting question whether I would observe similar results if the intervening task itself had occurred at random moments during the anticipatory period.

2.3.3.2 Short-Term Cross-Domain Consequences

Aside from the between-task consequences of predictable temporal structure between fixed and variable blocks, I also found between-task short-term effects of temporal structure in variable blocks. RTs to the intervening task in variable blocks were faster when the previous working-memory probe occurred early as opposed to late. The present effects in this multi-task context provide an interesting extension to observations during variable-foreperiod tasks, in which the interval (foreperiod) between a warning signal and an imperative target varies on a trial-by-trial basis (Drazin, 1961; Karlin, 1959; Los, 1996, 2010; Niemi & Näätänen, 1981; Steinborn et al., 2008; Vallesi & Shallice, 2007; Woodrow, 1914). A typical finding in these variable-foreperiod tasks is the asymmetric sequential effect: RTs for any current foreperiod are relatively insensitive to preceding shorter foreperiods, however increase if the previous foreperiod was longer than the current one (Los, 2010). Similarly, in this study the speed of responses to the intervening task may have been scaled relative to the memory-probe delay of the trial directly prior. However, direct comparisons between the current results and single-task sequential effects are difficult to derive and will require further experimentation. Unlike most single perceptual-motor tasks evaluating temporal preparation, the present study comprised a multi-task design engaging both external and internal attention. Moreover, each trial in the experiment contained three important intervals – the interval between encoding offset and intervening-task onset (always 1000 ms), the interval between intervening-item offset and the memory probe (early probe: 1250 ms, late probe: 2500 ms), and the ITI (randomly drawn between 500 ms and 800 ms), rendering it neither a purely fixed nor variable design. Nevertheless, it is intriguing that I observed a sequential effect on RTs in the intervening task (which always occurred at the same point in time) arising from the delay condition of the working-memory task of the previous trial.

At the same time, I note that the existence of sequential effects does not fully account for the performance benefits related to temporal predictability in fixed versus variable blocks. The results show a higher difference in RTs between early and late trials in fixed blocks as opposed to

previous-early versus previous-late trials in variable blocks. To the extent that implicit temporal associations underpin the effects of temporal expectation or preparation in fixed blocks, the current results suggest that strengthening of evidence and learning over longer timescales leads to stronger impacts on performance. The multiple trace theory of temporal preparation (MTP) provides a good candidate mechanism, capturing the ability of temporal associations over multiple timeframes to modulate behaviour (Los et al., 2014, 2017; Salet et al., 2022). In the future, it will be interesting to utilise neural measures to characterise exactly how temporal regularities over the different timescales shape signals guiding temporal expectation and preparation that can operate across tasks and domains.

2.3.3.3 The Contribution of Task Relevance

At the first glance, the beneficial consequence of predictable temporal structures in one task on performance in another task may seem to contradict findings related to the selective nature of temporal expectations. For example, previous research has demonstrated that directing attention to a target occurring at one point in time leads to performance benefits at that specific time but simultaneously to impairments earlier and later (Denison, 2024; Denison et al., 2017, 2021). Instead, the current results show that temporal expectations can facilitate performance of the working-memory task to which they are applied, as well as the external intervening task. In reconciling these findings, it is important to note several key differences which may help resolve this apparent discrepancy. In Denison et al.'s (2017) study, participants performed only one external task. Stimuli within that task had overlapping sensory properties and action associations, competing for priority in guiding performance. As one of the stimuli was more likely to be irrelevant, a trade-off between prioritising the likely-relevant stimulus and ignoring the likely-irrelevant stimulus was therefore an effective strategy. By contrast, a strategic trade-off would have been counterproductive in the present study, as participants completed two separate tasks

that were equally relevant. This provides for a scenario wherein both tasks potentially benefitting from temporal expectations would be much more advantageous.

Accordingly, the current pattern of results is likely driven by the requirement of having to perform two relevant tasks. An additional intriguing question is whether the cross-task benefits of predictable temporal structure are also contingent on the specific parameters and demands of the two tasks. In this study, a choice-reaction task was embedded within a working-memory task, to which participants had to return upon completion of the intervening task. Participants therefore had to maintain internal representations to perform the task adequately. Would similar results occur if another perceptual-motor choice-reaction task replaced the working-memory task? In scenarios where there is no requirement to maintain mental contents, and thus the demands of different tasks do not overlap in time, such a finding would suggest that performance benefits across tasks could solely be attributed to regularities in temporal structure. Therefore, in future studies, it will be interesting to examine the extent to which benefits of temporal structure can occur between two independent tasks rather than superimposed ones. Independently manipulating the perceptual similarity and motor requirements between tasks would provide further insights into the role of overlapping demands in orchestrating behaviour across successive tasks.

2.3.3.4 Temporal-Expectations Benefit in Working Memory

Beyond providing evidence for a between-task consequence of temporal predictability, the current study also yields new insights into the dynamic and prospective nature of working memory, extending the growing literature on the effects of temporal expectations on working-memory performance (Boettcher et al., 2021; Jin et al., 2020; Olmos-Solis et al., 2017; van Ede et al., 2017; Wilsch et al., 2015, 2018; Zokaei et al., 2019). The results uniquely demonstrate that performance benefits of temporal expectations during working memory remain robust, even when having to perform an intervening task in the period of probe anticipation.

The time-perception literature provides abundant evidence that the ability to explicitly track time can be biased when performed in complex multitask situations (Block et al., 2010; S. W. Brown, 1997, 2006; S. W. Brown et al., 2013; Fortin et al., 2007; Hemmes et al., 2004; Polti et al., 2018). Potentially, having to engage in the intervening task could have impeded the utilisation of temporal regularities regarding the working-memory task. Indeed, prior studies have shown that temporal-expectation effects can be substantially reduced by performing a concurrent demanding task (Capizzi et al., 2012, 2013; van der Mijn & van Rijn, 2021; but see: Capizzi et al., 2012 [Experiment 1]). Yet, a carefully controlled study using dual-task designs has shown that foreperiod effects survive in the face of a simultaneously performed task (van Lambalgen & Los, 2008). In a similar vein, recent work has revealed that learned (spatio)temporal regularities regarding external events can guide visual-search performance even when interacting with temporally unpredictable intervening events (Boettcher et al., 2022). Likewise, temporal expectations affect performance even during dynamic streams that require inhibition of temporally competing distractors (Chauvin et al., 2016; Davranche et al., 2011; Zokaei et al., 2021). These results suggest that it is possible to maintain and utilise representations of temporal regularities across other intervening events. The findings of the present study extend these previous results, by demonstrating that temporal regularities can also be highly effective for facilitating working-memory-guided behaviour during multi-task situations.

Nevertheless, our ability to keep track of time might be dependent on the specific demands of the intervening tasks, whereby more cognitively taxing intervening tasks may be more detrimental to the formation of temporal expectations. Moreover, the utilisation of temporal expectations for memory-guided behaviour may be further affected by the overlap between intervening and memory items. Prior research has found evidence for working-memory performance to be most impaired when perceptual distractors are similar to the memory content (Clapp et al., 2010; Hermann et al., 2021; Jha et al., 2004; Sreenivasan & Jha, 2007; Yoon et al., 2006). Therefore, a high degree of similarity in the associated objects between tasks may

modulate the benefit of temporal expectations. Within this study, additional analyses did not support this idea (Supplementary Results 8.1, Supplementary Table 8.2, Supplementary Table 8.3, Supplementary Figure 8.3). Sensory similarity between intervening and memory items did not affect RTs, and reproduction errors were in fact smaller when the intervening and memory item were more similar than when the angular difference between both items was high. I also did not observe any interactions between sensory similarity and either block type or delay condition, suggesting the effects of temporal expectation were not dependent on the sensory overlap. However, it is important not note that these results could be further modulated by an overlap in motor demands between the intervening and working-memory task (i.e., same vs. different response hand). Future studies may wish to explore more systematically the possible contribution of sensory overlap, motor overlap, and task difficulty on temporally guided internal attention in multi-task contexts.

Temporal expectations have been reported to benefit performance particularly at short as opposed to long delays. This effect is usually attributed to the hazard rate: once the potential time of the early probe passes, expectations can be updated, and attention can be reoriented towards the longer interval (Coull, 2009; Coull et al., 2000; Griffin, 2002; Jin et al., 2020; Miniussi et al., 1999; Nobre, 2001). Thus, even in variable blocks, the late memory probe becomes predictable as soon as the early interval has passed (see also MTP for a non-hazard explanation of this effect based on memory traces of past timing experiences; Los et al., 2014, 2017; Salet et al., 2022). This pattern of RTs has been noted in the foreperiod literature, with shorter RTs occurring with increasing delay duration in variable-foreperiod designs. In contrast, RTs in fixed-foreperiod situations increase with increasing delay between warning and target stimulus (Los et al., 2001, 2014; Niemi & Näätänen, 1981). In the variable blocks of the current study, RTs in the working-memory task reflected this typical pattern of results – responses to the memory probe were faster when it appeared unpredictably late as compared to unpredictably early. However, I did not find an effect of delay condition on RTs in fixed blocks as would have been predicted by

the foreperiod literature. This was the case even after excluding slow responses to the intervening task, thus, giving time to recover before performing the memory task. Whether this effect is specific to the multi-task context of the present study is a compelling avenue for future research.

Interestingly, memory representations were more accurate in fixed as compared to variable trials, that is, even when the memory probe appeared late (see also Chauvin et al., 2016; Cravo et al., 2017). This is distinct from a previous study which also examined the role of temporal expectation in working memory (Jin et al., 2020). Critically, however, in this prior study, no intervening task occurred during the working-memory delay. The intervening task in the present study may thus be critical for the here-observed effect of temporal expectations on working-memory performance in the late trials.

2.3.3.5 Conclusion

In the present experiment, I focussed on the intersection of several lines of research often studied in isolation, including external attention, internal attention, temporal expectations, and multitasking. I provide evidence that implicitly learned temporal expectations regarding a working-memory task can be utilised even when having to engage in an external intervening task during the period of anticipation. Moreover, the current findings demonstrate that temporal expectations regarding a working-memory task can influence the performance of the external intervening task. Taken together, this study extends previous findings on internal-to-external attentional modulations and temporal expectations by demonstrating that the temporally guided internal attention can affect and improve performance of an external task within a multi-task situation.

2.4 Experiment 2: How Temporally Guided External Attention Modulates Internal Attention

While Experiment 1 revealed an impact of temporally guided internal attention on external attention, in Experiment 2, I took the complementary approach, asking whether temporal expectations regarding the occurrence of an external event influence performance of a working-memory task. In particular, I aimed to uncover whether temporal expectations concerning external events appearing during the working-memory retention interval could help mitigate their detrimental impact on internal representations. To determine if the potential protective effects of temporal expectations rely on distractor suppression or involve shielding (i.e., protection) of internal representations, I compared effects after irrelevant perceptual distractors that could be ignored versus interrupters that required a response.

It has become well established that external distractors do not only interfere with the encoding (Feldmann-Wüstefeld & Vogel, 2019; McNab & Klingberg, 2008; Vogel et al., 2005; for a review see: Liesefeld et al., 2020) but also with the retention of sensory information in working memory (for a review see: Lorenc et al., 2021). Accordingly, memory reports for both low-level features such as colour (Nemes et al., 2012; S. Z. Sun et al., 2017), location (Marini et al., 2017; van der Stigchel et al., 2007), orientation (Barth & Schneider, 2018; Lorenc et al., 2018; Rademaker et al., 2015; D. Schneider, Barth, Getzmann, et al., 2017), or motion (Berry et al., 2009; Pasternak & Zaksas, 2003), as well as high-level stimuli such as faces (Clapp et al., 2010; Mallett et al., 2020; Yoon et al., 2006) become more prone to errors when a distractor is presented during the maintenance period. At the same time, working-memory content has also been reported to be quite robust to various forms of interference (Bettencourt & Xu, 2016; Rademaker et al., 2019; Zickerick et al., 2020; for a review see: Xu, 2017, 2020), alluding to dynamics that may support such resilience.

To understand how the impact of interference on working memory can be reduced, I first turn to findings from the domain of perception. Studies on distraction during perceptual tasks

have shown that learned regularities can help dampen the effects of distractors through proactive suppression (as reviewed in Geng, 2014; Geng et al., 2019; Noonan et al., 2018; Theeuwes et al., 2022; van Moorselaar & Slagter, 2020). For instance, salient distractors cause less interference when occurring at locations where they are more frequently presented (Failing et al., 2019; Ferrante et al., 2018; Goschy et al., 2014; C. Huang et al., 2021; Leber et al., 2016; Narhi-Martinez et al., 2024; Sauter et al., 2018; Stilwell et al., 2019; B. Wang, van Driel, et al., 2019; B. Wang & Theeuwes, 2018b, 2018a). Moreover, when targets appear at locations that frequently contain a salient distractor, detection of targets is slowed, suggesting a general suppression for the distractor location (Failing et al., 2019; Ferrante et al., 2018; Reder et al., 2003; B. Wang, van Driel, et al., 2019; B. Wang & Theeuwes, 2018b, 2018a). Extending these findings, efficient reduction of distractions can also rely on feature-based regularities. Salient distractors can be more efficiently rejected when they are more likely presented in one colour than any other colour (Failing et al., 2019; Feldmann-Wüstefeld & Schubö, 2016; Gaspelin & Luck, 2018; Stilwell et al., 2019; Vatterott & Vecera, 2012).

Similar to the perceptual domain, working memory may be more resilient if interference can be successfully predicted. There is some evidence that expectations about the occurrence of distractors can help counter their negative consequences during working memory. Findings of a recent study suggest that when interference with specific to-be-remembered features was anticipated, these features were reinforced in working memory. On the other hand, features for which no interference was expected were often left more susceptible (B. Wang et al., 2022). Nevertheless, it remains unclear whether temporal expectations can also mitigate interference during working memory, or through what processes such mitigation might be mediated.

To date, temporal expectations to contents in perception (Coull & Nobre, 1998; van Ede, Rohenkohl, et al., 2020; Vangkilde et al., 2012) and working-memory (Heuer & Rolfs, 2021; Jin et al., 2020; Olmos-Solis et al., 2017; van Ede et al., 2017; Zokaei et al., 2019) have been only investigated in regard to target processing. There is some evidence that temporal expectations

regarding the onset of a perceptual target can improve performance when target stimuli are preceded by (Lamy, 2005), followed by (van Ede et al., 2018), paired with (Menceloglu et al., 2017), or embedded in perceptual distractors (Shalev et al., 2019). Furthermore, research has demonstrated that anticipating the appearance of a perceptual target can still improve performance, even in the presence of concurrent demands on working memory (Capizzi et al., 2012 [Experiment 1]). However, it remains unaddressed whether temporal expectations can help diminish interference from external events during the maintenance of internal representations, and thereby facilitate working-memory performance.

At least two mechanisms could contribute to a potential benefit of temporal expectations in mitigating interference: First, sensory processing of task-irrelevant distractors may be suppressed (Bonnefond & Jensen, 2012; de Vries et al., 2019; Getzmann et al., 2018; Payne et al., 2013), without necessarily affecting internal representations. Alternatively (or additionally), anticipating interference may act on memory contents directly (i.e., shielding the memory items themselves) independently of affecting the processing of external interference. To test for effects of memory shielding, it is necessary to introduce a type of interference that cannot simply be suppressed, such as a sensory input that requires a response during the working-memory retention period. I therefore tested the putative benefits of temporal expectations when faced with two types of interference: to-be-ignored perceptual distractors and to-be-responded-to interrupters (i.e., external intervening tasks) – the latter typically exerting a more detrimental influence over working memory (Bae & Luck, 2019; Berry et al., 2009; Clapp et al., 2010; Hakim et al., 2021; Mishra et al., 2013; B. Wang, Theeuwes, et al., 2019; Zickerick, Kobald, et al., 2021). While suppression can be used to mitigate the effects of task-irrelevant perceptual distractors, it would be counterproductive to inhibit sensory information that is relevant for an external intervening task. Consequently, if I can demonstrate that proactive anticipation of an external intervening task also improves memory performance, this would suggest a contribution from internal shielding beyond any potential influence from external distractor suppression.

2.4.1 Methods

2.4.1.1 Participants

Participants were recruited via Prolific Academic (https://www.prolific.co/). They were pre-screened based on demographic criteria (i.e., age range 18 to 40, fluent in English), general health (i.e., normal or corrected-to-normal vision, no history of mental illnesses) and previous participation history on Prolific Academic (i.e., participated in at least 10 studies, with a study approval rate above 90%). All participants provided informed consent prior to participation and were paid £7.50 per hour. An additional monetary reward of up to £5 could be earned depending on participants' task performance in the experiment. Specifically, performance above 80% received a bonus payment scaling from £0.01 at 80% to £5 at 100%, with an average bonus payment of £0.76 (*SD* = 0.82) across all participants.

As in Experiment 1, sample size was set to *n* = 54. To reach the desired sample size, data were collected from 79 online participants. Data from 22 participants were excluded following an a-priori trial-removal procedure (before splitting data by conditions) and three more participants were removed because they reported utilising an explicit non-memory-based strategy to complete the task (see 2.4.1.3 Analysis for details). This yielded the desired final sample of 54 participants (age range: 18 to 38 years; mean age: 27.17 years; 28 females, 48 right-handed). The online study was approved by the Central University Research Ethics Committee of the University of Oxford.

2.4.1.2 Task and Procedure

In the present study, participants performed a web-based visual working-memory task requiring the reproduction of the exact angle of one out of two tilted bars at the end of a memory delay (Figure 2.3). Two main manipulations of this task were (1) that the interference either appeared at a fixed (i.e., predictable) or variable (i.e., unpredictable) point in time, and (2) that it was either an entirely irrelevant stimulus that should be ignored (i.e., distractor) or a stimulus requiring a

response (i.e., interrupter, external intervening task). These manipulations allowed for the assessment of whether and how temporal expectations mitigate external interference in working memory. I return to these task features at the relevant instances below.

Participants completed the experiment in a web browser on their personal computers. The recommended internet browsers were Mozilla Firefox and Google Chrome; participating via mobile phone or tablet was not allowed. Prior to the experiment, participants' individual screen resolution was estimated by asking them to adjust an image of a credit card such that it matched the size of a physical credit card. In this manner, I could calculate the ratio between the card-image width in pixels and the actual card width in centimetres to obtain a measure of pixel density (i.e., pixel per cm). Together with the instructed viewing distance of approximately 60 cm (i.e., one arm's length away from the monitor), this allowed stimuli to be presented in degrees of visual angle, regardless of monitor size (Li et al., 2020). The experimental script was generated in PsychoPy (Peirce et al., 2019) and hosted online using Pavlovia (http://www.pavlovia.org/).

At the start of each trial, two tilted bars were simultaneously presented against a grey (RGB value: [128,128,128]) background for 250 ms. One bar was always positioned to the left and the other to the right of the central fixation cross. Independent of location, one of the bars was tilted to the left (anticlockwise) and the other to the right (clockwise). To avoid angles too close to vertical and horizontal meridians, the items' angles were randomly drawn in increments of ±5° between ±5° and ±85°. Across trials, a leftwards or rightwards oriented bar was equally likely to appear in the left and right position. The stimuli subtended approximately 0.8° in width and 5.7° in length and were centred at a viewing distance of 5.7° visual angle from fixation. At encoding, both lateralised items were equally likely to be probed, rendering them equally relevant.

Visual encoding was followed by a memory delay of 2750 ms, during which the fixation cross remained on the screen. In 75% of trials within a block, a tilted bar (referred to as "interfering item") was presented during the memory delay in the centre of the screen for 250 ms. The total length of the memory delay was 2750 ms, regardless of whether the interference appeared in

that trial or not. Depending on the type of block, this interfering item was either presented at a variable or fixed time point (temporal predictability: fixed onset vs. variable onset), and participants either had to ignore or respond to it (interference type: distraction vs. interruption).

In fixed-onset blocks, the interfering item always appeared at a specific point in time during the retention interval (interference onset: 250 ms or 1000 ms or 1750 ms after encoding offset depending on the block). In variable-onset blocks the interfering item was equally likely to occur at any of these three time points. That is, within the interference trials of a given variable-onset block, one third of the trials presented the interfering item at 250 ms, one third at 1000 ms, and one third at 1750 ms after encoding offset.

In distraction blocks, participants were instructed to ignore the interfering item. In contrast, in interruption blocks, participants were required to respond to the item; if the bar was tilted to the left (right), participants pressed the F (J) key on the keyboard with their left (right) index finger as fast as possible. Within blocks, leftwards and rightwards angles of the interfering items were counterbalanced. The interfering item was always presented in a different colour than the two memory items that preceded it. The colour of the memory items and the interfering item were always drawn from a set of three highly distinguishable colours (RGB values: blue [0,225,228], orange [254,163,0], pink [253,142,253]). The colours used for the memory items and the interfering item varied randomly across trials. The interfering item had the same size as the memory items and its angle was also randomly drawn in increments of ±5° between ±5° and ±85°.

Directly following the memory delay, the fixation cross changed colour (referred to as “probe”), indicating which memory item's tilt should be reported. Participants were never probed about the interfering item. Following the appearance of the probe, participants had unlimited time to decide on their response. To report a leftwards (rightwards) angle, participants were asked to press the F (J) key on the keyboard using their left (right) index finger. After response initiation, a visual response dial was displayed on the screen, always starting in vertical position. The

response dial had the same diameter as the length of the bars (5.7°) and always appeared surrounding the fixation cross.

The dial rotated leftwards when pressing F and rightwards when pressing J (either by holding the key down or pressing the key repeatedly; always in increments of ±5°). Critically, the dial could only be rotated in the direction that was initially indicated by the participant. For example, if a participant started pressing the F key after the probe, the dial would only move leftwards, and it would therefore not be possible to move the dial towards the right with the J key. Since the response dial always started in a vertical position and because it could not be rotated beyond ±90°, a leftwards (rightwards) oriented bar could only be correctly reported with a left (right) key. As a consequence, the hand required for responding was directly linked to the angle of the bar that was probed. This builds on previous tasks (Boettcher et al., 2021; van Ede, Chekroud, Stokes, et al., 2019), though the specifics of this response implementation were not critical to the current study. Once participants started rotating the dial, they were given only limited time (4000 ms) to complete the angle reproduction. This was intended to encourage participants to recall the exact orientation before moving the dial. When the dial aligned with the remembered tilt of the item, participants pressed the space bar to verify their response and continue with the task.

Next, participants received feedback in the form of a number ranging from 0 to 100, with 100 indicating a perfect report and 0 indicating that the adjusted orientation was perpendicular to the angle of the probed item. Feedback was presented for 500 ms. However, if time to adjust the angle ran out, the message “Too slow” was presented instead for 750 ms. In interruption blocks this was followed by a second feedback message if participants responded with the wrong key (i.e., “Wrong key! Use the correct key to respond to the distractor!”) or did not respond at all to the interfering item (i.e., “Respond to the distractor!”). To incentivise fast responses to the interrupter, participants also received a feedback message when their RTs to the interrupter were slower than 750 ms (i.e., “Too slow! Respond faster to the distractor!”). In distraction blocks, in

which participants had to refrain from responding to the distractor, a feedback message was displayed if participants responded to the interfering item (i.e., “Don’t respond to the distractor!”). The distractor- and interrupter-specific feedback message was combined with an image reminding participants to press F (J) when the interfering item was tilted to the left (right) or to withhold their response, respectively. Feedback was presented for a minimum of 750 ms and remained on screen until the space key was pressed, encouraging participants to read the feedback message before proceeding to the next trial. Trials were separated by an ITI randomly drawn between 500 and 800 ms.

The experiment consisted of 384 trials divided across 12 blocks (each including 32 trials). The blocks were split according to the type of interference. Six blocks included interrupters, while the other six included distractors. These were further subdivided depending on the temporal predictability of the interfering event, which had a fixed onset in three blocks (one block each of 250 ms, 1000 ms, or 1750 ms), and a variable onset in the other three blocks (pseudo-randomised to occur equiprobably at 250 ms, 1000 ms, or 1750 ms). As such, the total number of trials where the interfering event would appear at any one delay interval (e.g., 250 ms) was equal between the corresponding fixed-onset block and across the three variable-onset blocks.

The order of blocks was pseudo-randomised with the two possible interference types nested within block pairs of the same temporal predictability. For example, a fixed-onset block with one type of interference (e.g., distractors) was always followed by another fixed-onset block with the other type of interference (e.g., interrupters). This would then be followed by a pair of variable-onset blocks (with the order of fixed-onset and variable-onset pairs being counterbalanced across participants). The order of the interference types was randomised within each temporal-predictability pair. The order of the fixed-onset blocks with interference at 250, 1000, or 1750 ms after encoding offset was randomised across participants.

The interference type (i.e., distractor vs. interrupter) was made explicit before the start of each block by presenting participants an image of the trial sequence and a reminder to either

ignore or respond to the interfering item. For the sake of simplicity, distractor blocks were referred to as "Ignore blocks" and interrupter blocks as "React blocks". Participants were informed that they would never have to report the angle of the interfering item. However, they were not informed about the temporal predictability (i.e., fixed vs. variable) or about the three possible interference onsets (i.e., 250 ms vs. 1000 ms vs. 1750 ms). In order to become familiarised with the procedure of the experiment, participants performed 16 practice trials of the interruption block and 16 trials of the distraction block, both with variable interference onset. At the end of the experiment, participants were redirected to the survey website Qualtrics (http://www.qualtrics.com/) where they were asked about comprehension of the instructions, potential strategy used to complete the task, and whether they thought their data should be analysed. The whole experiment lasted approximately 50 minutes.

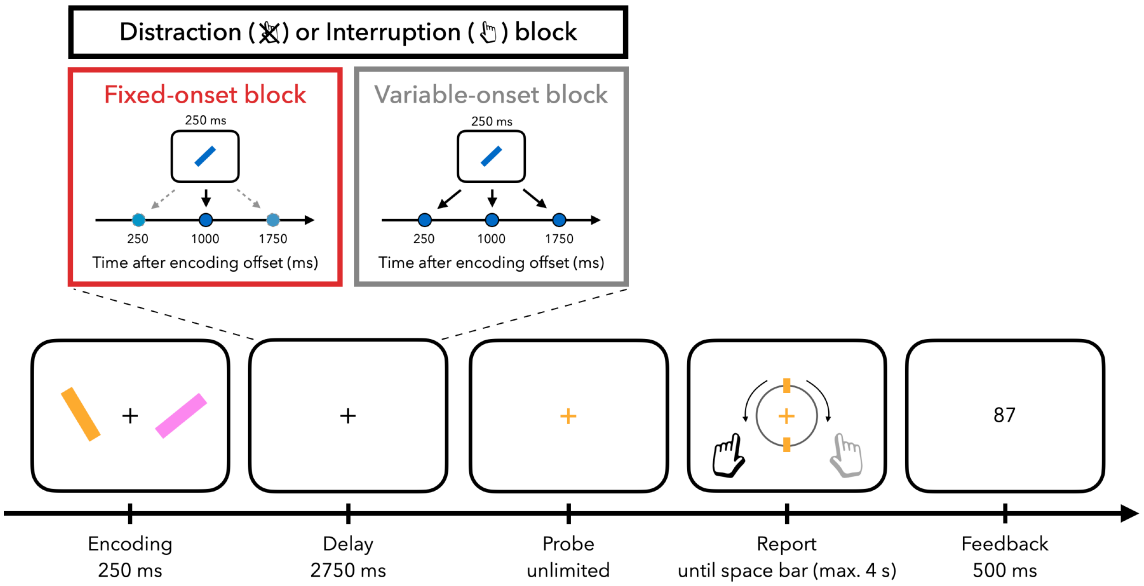

Figure 2.3. Task schematic. Trials started with an encoding display consisting of two, lateralised tilted bars. Participants' task was to remember the angle of both bars, of which one had to be reported at the end of the trial. On 75% of the trials within a block, interference occurred for 250 ms during the memory delay. In distraction blocks, participants could ignore the perceptual task-irrelevant distractor item. In interruption blocks, participants indicated whether the interfering item was tilted to the left or to the right. In fixed-onset blocks, interference occurred at a fixed time point within a block (at 250, 1000, or 1750 ms after encoding offset; counterbalanced across blocks). In variable-onset blocks, interference occurred equally likely at each of the three possible time points. The delay between encoding offset and probe onset was always 2750 ms long. After the delay, a colour change of the central fixation cross indicated which bar's angle had to be reported. Participants were given unlimited time to retrieve the item from working memory and to decide what to report. However, once they started pressing a key, they were given only 4000 ms to complete their report. Following the report, participants received feedback in form of a number ranging from 0 to 100.

2.4.1.3 Analysis

Data were analysed in R Studio (Posit team, 2022). During data pre-processing, trials were removed if RTs (i.e., from probe onset to the first key press) exceeded 5000 ms, were 2.5 *SD* above the individual mean across trials of all conditions, or if participants did not reproduce the probed angle within 4000 ms. I also excluded trials in distractor blocks if participants responded to the interfering item, as well as trials in interrupter blocks if participants did not respond, responded with the wrong key, or did not respond within 1000 ms to the interfering item. Twenty-two datasets where more than 10% of trials were rejected during these pre-processing steps were removed from further analysis. Additionally, three datasets were also removed where participants self-reported to have employed explicit non-memory-based strategies to maintain the encoding display (i.e., aligning their fingers with the memory items). After this exclusion step, the data of 54 participants with an average of 94.00% (*SD* = 1.82) retained trials entered the main analysis.

Reproduction errors were calculated by averaging the absolute difference between the original angle of the target (i.e., probed) item and the reported angle across all trials and within each condition. I also examined RTs to the intervening task during interruption blocks. Moreover, I additionally, analysed the angular deviation between the angle of the interfering item and the reported angle.

To analyse the proportion of trials in which the observer incorrectly reported the angle of the interfering item (i.e., swaps, non-target responses), I fitted a mixture model to the data (Bays et al., 2009). The relative proportion of the non-target (i.e., interfering item) distribution was estimated separately for each participant and condition (i.e., interference type and temporal predictability). The mixture-modelling analysis was performed using the mixtur package (version 1.0.0; Grange & Moore, 2022).

As in previous studies (Lorenc et al., 2018; Mallett et al., 2020; Nemes et al., 2012; Rademaker et al., 2015; S. Z. Sun et al., 2017; van der Stigchel et al., 2007), I also quantified whether the angle reported during recall of the memory item was biased either towards (i.e.,

attractive bias) or away from (i.e., repulsive bias) the angle of the interfering item. To this end, I first subtracted each participant's mean reproduction error calculated across all conditions from their reproduction error in each individual trial to eliminate general response biases. Next, the reproduction errors were sorted ordinally and then binned (i.e., moving-window approach, step size = 5°, bin width = 45°) according to the relative orientation of the target angle with respect to the interference angle. The average response bias was then calculated for each bin, yielding a response-bias curve as a function of the relative target-interference angular difference. This resulted in a response-bias value representing the degree to which the reported angle deviated from the target angle, where a negative value would indicate the report to be shifted clockwise, and a positive value would indicate the report to be shifted anticlockwise. Similarly, when the interference was oriented more clockwise or anticlockwise relative to the target, the target-interference angular difference would be negative or positive, respectively. As such, an attractive bias was indicated when the sign of the response bias matched the sign of the target-interference angular difference, whereas an opposing sign of the response bias relative to the relative target-interference angular difference indicated a repulsive bias. To test whether the interference biased memory reports, the area under the bias curve was integrated – separately for negative and positive target-to-interference angles – and subsequently compared. As a last step, I examined whether temporal predictability or interference type might influence the magnitude of the response bias. Therefore, I equated the average response biases for the negative and positive interference-target angular differences of each participant and condition, then compared the resulting equated biases between conditions.

When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. I computed 2×2 repeated-measures ANOVA with the factors interference type (i.e., distraction vs. interruption) and interference presentation (i.e., no interference vs. interference) to test whether interference has a detrimental impact on working-memory performance. Moreover, I performed a 2×2×3 repeated-measures ANOVA with the

factors temporal predictability (i.e., fixed vs. variable), interference type (i.e., distraction vs. interruption), and interference onset (i.e., 250 ms vs. 1000 ms vs. 1750 ms) to compare differences in reproduction errors. A 2×2 repeated-measures ANOVA with the factors temporal predictability and interference onset was performed to compare differences in RTs to the interrupter. Moreover, I computed a 2×2 repeated-measures ANOVA with the factors temporal predictability and interference type to test for differences in swaps and response biases. When evaluating only two means, I applied paired samples *t*-test and report Cohen's *d* as a measure of effect size. For post-hoc *t*-tests, I report Bonferroni-corrected *p* values that I denote as $p_{\text{Bonferroni}}$. The ggplot2 package (version 3.3.3; Wickham, 2009) was used for plotting results. Where relevant, the within-subject standard error of the mean was calculated from normalised data using the approach from Morey (2008).

2.4.2 Results

2.4.2.1 Working-Memory Performance: No Interference vs. Interference

I first wanted to confirm the negative impact of both types of interference on the accuracy of reports in the working-memory task. To this end, I evaluated the average reproduction error (i.e., the absolute deviation from the probed orientation, for which lower levels indicate better performance) and compared trials with and without interference during the memory delay for distraction and interruption blocks. As depicted in **Figure 2.4** (left panel), errors were systematically larger when the memory delay was disrupted by interference in comparison to trials without interference ($F_{(1,53)} = 53.868$, $p < 0.001$, $\eta^2_G = 0.038$). This was true for nearly every participant in the experiment (**Figure 2.4**, right panel). I also found a significant effect of interference type ($F_{(1,53)} = 73.667$, $p < 0.001$, $\eta^2_G = 0.083$), indicating that working-memory performance was generally worse in interruption as compared to distraction blocks. Further, although the interaction between interference presentation and interference type was significant ($F_{(1,53)} = 11.044$, $p = 0.002$, $\eta^2_G = 0.009$), post-hoc pairwise comparisons revealed that in both

distraction ($t_{(53)}$ =-3.248, $p_{Bonferroni}$ = 0.004, d = 0.442) and interruption blocks ($t_{(53)}$ =-6.588, $p_{Bonferroni}$ < 0.001, d = 0.897), working-memory performance was significantly decreased when interference was present. Although reproduction errors were of primary interest, for the sake of completeness, RTs (i.e., the interval between probe onset and the initial response) are reported in Supplementary Figure 8.4.

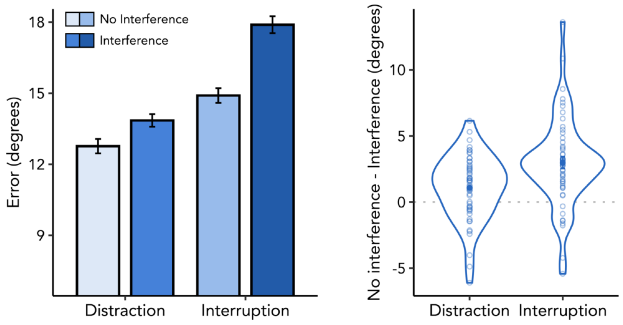

Figure 2.4. Both types of interference impair working memory. Left panel: Participants were less accurate in trials with as opposed to without interference. Right panel: In both distraction and interruption blocks, working-memory performance decreased when interference was present. The interference cost of each interference type was calculated by taking the difference in reproduction errors between trials with and without interference (no interference minus interference). Dots represent individual participants. Error bars indicate ±1 *SEM*.

2.4.2.2 Three Possible Scenarios

Having confirmed that interference negatively affected task performance, I next turned to my main question: Can proactive temporal expectations about the interference onset help overcome the detrimental effects of interference on working-memory performance? I considered three possible scenarios regarding how temporal expectations might influence working memory. First, temporal expectations may be unable to mitigate interference, resulting in reproduction errors being similar between fixed and variable interference onsets (**Figure 2.5**, Scenario 1, red vs. grey bars). Conversely, temporal expectations may help mitigate the effects of interference. Mitigation could occur through two mechanisms, reflected in two possible patterns of results. If temporal expectations work through suppressing external sources of interference, they should only reduce working-memory errors in distractor trials but be ineffective in interruption trials (**Figure 2.5**, Scenario 2). Alternatively, if temporal expectations can directly shield internal representations

from external sources of interference, errors should also be reduced for temporally predictable interrupters (**Figure 2.5**, Scenario 3).

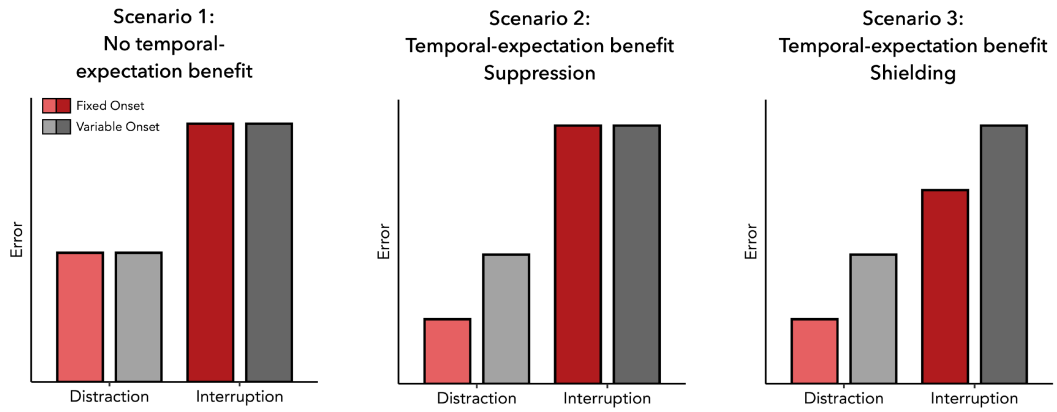

Figure 2.5. Possible patterns of errors in the working-memory task as a function of temporal predictability and interference type. Scenario 1: No temporal-expectation benefit. Scenario 2: temporal-expectation benefit occurs exclusively through distractor suppression. Scenario 3: temporal-expectation benefit occurs through memory shielding.

2.4.2.3 Working-Memory Performance: Fixed vs. Variable

If participants can leverage temporal expectations to mitigate interference, then working-memory performance should be better (i.e., smaller reproduction errors) in blocks where interference occurred at a fixed (i.e., temporally predictable) as compared to a variable (i.e., temporally unpredictable) point in time during the memory delay. In support of this hypothesis, I found significantly smaller errors when interference could be temporally predicted (**Figure 2.6A**; $F_{(1,53)}$ = 8.774, p = 0.005, η^2_G = 0.004), ruling out the first of the hypothetical scenarios (**Figure 2.5**, Scenario 1). In addition, and as expected, I also found a main effect of interference type, showing that participants were overall worse when they were required to respond to the interfering item ($F_{(1,53)}$ = 75.289, p < 0.001, η^2_G = 0.096).

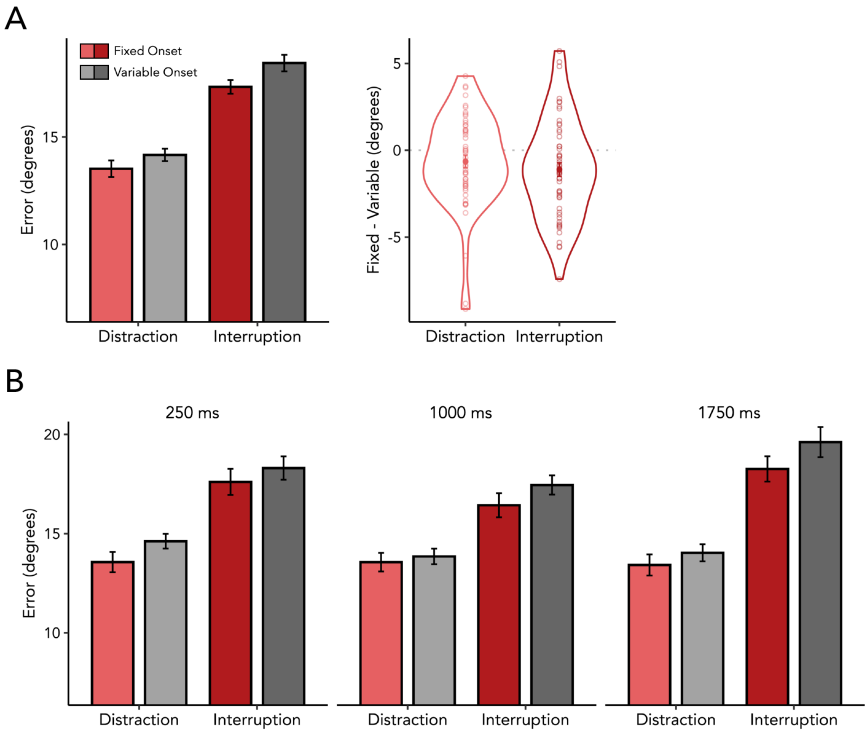

Figure 2.6. Temporal expectations about the onset of interfering events support the robustness of visual working memory. (A) Left panel: Reproduction errors in the visual working-memory task were smaller when participants could temporally predict interference. Right panel: Temporal expectations increased working-memory performance for both interference types. The temporal-expectation benefit of each interference type was calculated by taking the difference in reproduction errors between fixed-onset blocks and variable-onset blocks (fixed minus variable). **(B)** The temporal-expectation effect was similar across all possible interference onsets. Dots represent individual participants. Error bars indicate ±1 *SEM*.

2.4.2.4 Working-Memory Performance: Suppression vs. Shielding

If participants benefited from temporal expectations exclusively by suppressing the interfering item – as in the second possible scenario (**Figure 2.5**, Scenario 2) – then the effect of temporal expectations should be present only in distractor blocks, but not in interrupter blocks where participants had to attend and respond to the interfering item. Alternatively, if proactive anticipation in time instead (or additionally) shields working-memory contents, then I should find smaller errors for fixed compared to variable onsets for both types of interference (**Figure 2.5**, Scenario 3). In line with the third scenario, the effects of temporal expectations were similar in distractor and interrupter blocks (**Figure 2.6A**), without a significant interaction between temporal predictability and interference type ($F_{(1,53)}$ = 0.502, p = 0.482, η^2_G < 0.001). If anything, the temporal-expectation benefit was numerically even larger, albeit not significantly, in trials with

interrupters than distractors. This suggests that temporal expectations help overcome interference not only when the source of interference can be ignored, but also when an external intervening task must be completed during the period of memory retention. This pattern of results was obtained across all three tested interference onsets (Figure 2.6B), indicating that internal representations can be protected against distractors as well as interrupters, regardless of whether interference occurs at an early, intermediate, or late time point after encoding (see Supplementary Table 8.4 for full set of inferential statistics).

Supplementary Figure 8.5 depicts the corresponding analysis for RTs relative to the probe onset (see Supplementary Table 8.5 for full set of inferential statistics). Interestingly, temporal expectations did not significantly modulate participants' RTs in the present experiment.

2.4.2.5 Working-Memory Performance: Shielding vs. Reduced Interference-Encoding

Although the results demonstrate that temporal expectations help mitigate the detrimental effects of interference on working memory even when interfering events cannot be suppressed, it is not immediately possible to conclude unequivocally that shielding alone can account for the effects. An alternative possibility is that, while allowing for sensorimotor processing, temporal expectations may have modulated the degree to which interfering items became encoded into working memory. Such a mechanism would mitigate the impact of external information instead of enhancing protection of internal representations. To consider this option, I firstly tested whether the interference angle entered working memory at all. Although participants were explicitly informed that they will never be probed about the angle of the interfering item, it might be possible that this information was still sometimes incidentally encoded. Using a pairwise comparison, I tested whether the angular deviation between interference and working-memory report differed from a chance level of 45 degrees. As depicted in Figure 2.7A, the actual angular deviation between interference and report (M = 44.121, SD = 1.045) significantly differed from

chance ($t_{(53)}$ = 4.372, p < 0.001, d = 0.595), suggesting that some information regarding the interfering item was carried over into the working-memory reporting stage.

I considered two possible explanations for the above-chance reporting of the interfering item. First, on a subset of trials participants may be erroneously reporting the orientation of the interfering item rather than reporting the memorised item orientation, here termed a "swap". Second, participants may be generally biased (or pulled) towards the orientation of the interfering item in their report of the memory item. If these effects are modulated by temporal expectations, then we can assume that the temporal-expectation benefit is due at least in part to diminished encoding of the interfering item.

Using a mixture-model approach according to Bays et al. (2009), I quantified the proportion of swaps with the interfering item as a function of temporal predictability and interference type (**Figure 2.7B**). Although overall there was a low proportion of swap errors (< 5%), interrupters as opposed to distractors did significantly increase the probability of swaps ($F_{(1,53)}$ = 6.512, p = 0.014, η^2_G = 0.026). Crucially, however, temporal predictability did not affect the swap rate ($F_{(1,53)}$ = 0.761, p = 0.387, η^2_G = 0.002), and the interaction between the two factors did not reach significance ($F_{(1,53)}$ = 0.811, p = 0.372, η^2_G = 0.003).

These findings are complemented by the results of the response-bias analysis (**Figure 2.7C**). Generally, the reproduction of the target angle was biased towards the interference angle ($t_{(53)}$ = 3.779, p < 0.001, d = 0.514). That is, when the interference orientation was anticlockwise (clockwise) to the target orientation there was an anticlockwise (clockwise) shift in participants' orientation reports of the target. Comparing this pull effect between conditions revealed no systematic difference between interrupters and distractors (F(1,53) = 1.308, p = 0.258, η^2_G = 0.007) or between fixed and variable onsets ($F_{(1,53)}$ = 0.707, p = 0.404, η^2_G = 0.003). Moreover, there was no significant interaction between interference type and temporal predictability ($F_{(1,53)}$ = 0.117, p = 0.734, η^2_G < 0.001).

Thus, even though interference entered working memory to some extent, being able to anticipate the timing of an interfering event did not lead to a difference in the memory trace of the distracting or interrupting event. As such, temporal expectations did not modulate the degree to which external information enters working memory. Instead, it is more plausible that temporal expectations affected how well internal representations are shielded against external disturbances.

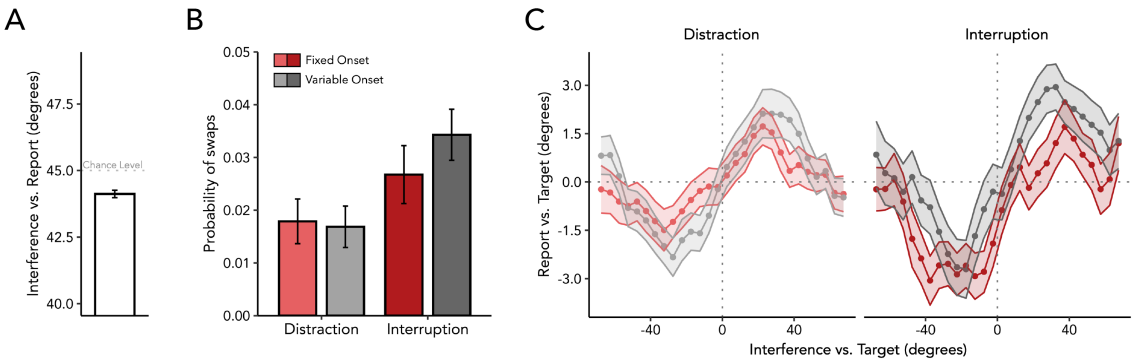

Figure 2.7. No evidence for stronger encoding of temporally predictable interference into working memory. (A) There is a significant difference between the angular deviation between interference and report and the chance level of 45° (depicted as dashed line). **(B)** Interrupters as compared to distractors increased the swap rate, however temporal expectations did not affect the probability of reporting the interfering item. **(C)** shows the average response bias relative to the angular difference between target and interfering item. A negative value along the y-axis would indicate reports biased clockwise from the angle of the target item, while a negative value along the x-axis would indicate that the interfering item's angle was clockwise from the target item's orientation. The memory reports were therefore shown to be biased towards the interference (e.g., an attractive bias). However, neither interference type nor temporal predictability modulated this attraction bias. Shaded areas and error bars indicate ±1 *SEM*.

2.4.2.6 External-Task Performance

To rule out the possibility that participants experienced less interference simply because they chose to ignore the temporally predictable interrupters, I also tested for differences in RTs to the interrupters themselves, when these occurred at predictable versus unpredictable times. Participants responded faster to the interrupter when it occurred at a predictable time in the memory delay compared to a variable onset (**Figure 2.8A**; $F_{(1,53)} = 32.037$, $p < 0.001$, $\eta^2_G = 0.023$). Thus, temporal expectations did not induce a trade-off between performance on the working-memory and the external intervening task, but instead improved performance on both.

For performance to the interrupter, I also found a main effect of interference onset (Figure 2.8C; $F_{(2,106)}$ = 15.655, p < 0.001, η^2_G = 0.037). Overall, RTs occurred faster when the interrupter was presented at 1000 ms ($t_{(53)}$ = 5.123, $p_{Bonferroni}$ < 0.001, d = 0.697) and 1750 ms ($t_{(53)}$ = 4.089, $p_{Bonferroni}$ < 0.001, d = 0.556) as compared to 250 ms; however, there was no difference in RTs between trials with 1000-ms and 1750-ms interrupter onsets ($t_{(53)}$ =-0.368, $p_{Bonferroni}$ = 1.000, d = 0.050).

Interestingly, in contrast to what I observed for the influence of temporal expectations on working-memory performance (Figure 2.6B), I found a significant interaction between temporal predictability and interference onset in the interrupter task (Figure 2.8B; $F_{(2,106)}$ = 15.785, p < 0.001, η^2_G = 0.026). Pairwise comparisons revealed faster RTs to the interrupter when it was temporally predictable compared to unpredictable at 250 ms ($t_{(53)}$ = 6.065, $p_{Bonferroni}$ < 0.001, d = 0.825), but not at 1000 ms ($t_{(53)}$ = 0.187, $p_{Bonferroni}$ = 1.000, d = 0.025) or 1750 ms ($t_{(53)}$ = 1.709, $p_{Bonferroni}$ = 0.280, d = 0.233). The effect of faster RTs to interrupters after short delays but not after long delays is in agreement with prior studies of temporal expectations in simple perception and action tasks (Coull et al., 2000; Cravo et al., 2017; Miniussi et al., 1999; for a review see: Nobre, 2001; Nobre & van Ede, 2018, 2023). Interestingly, however, I did not find a similar onset dependence for the protective effect of temporal expectations on working-memory performance (Figure 2.6B). This may suggest that temporal expectations exert distinct influences on working-memory protection and external-task facilitation – a possibility that remains interesting to address in future research but for which further discussion is beyond the current scope.

Although trials with incorrect interrupter responses were removed in the main analysis, I conducted a separate analysis to examine error rates in the interrupter task. Supplementary Figure 8.6 shows that error rates tended to be lower for fixed versus variable interrupter onsets (see Supplementary Table 8.6 for full set of inferential statistics). However, similar to Experiment 1, error rates were very low (M = 4.131, SEM = 0.404), suggesting a constrained variability within this measure.

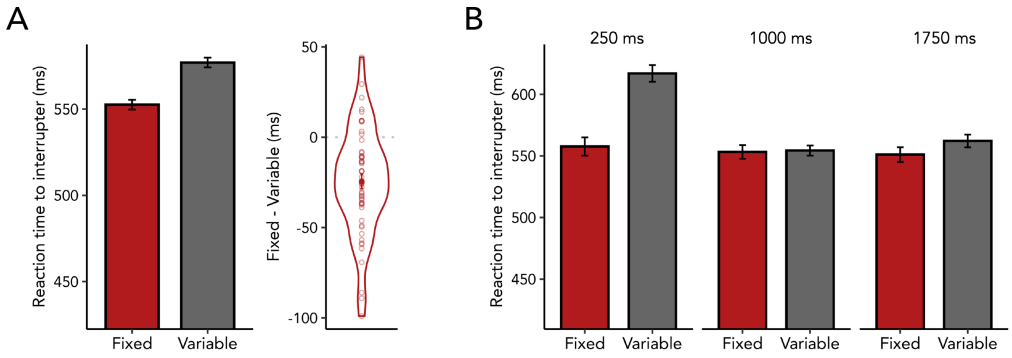

Figure 2.8. Temporal-expectation benefit in the interrupter (external intervening) task. (A) Left panel: RTs to the interrupter were faster for fixed as compared to variable onsets. Right panel shows the difference in RTs to the interrupter between trials where the interrupter occurred at a fixed versus a variable point in time during memory retention (fixed minus variable). **(B)** The temporal-expectation effect was most pronounced when the interrupter occurred at 250 ms compared to at 1000 ms or 1750 ms after encoding offset. Dots represent individual participants. Error bars indicate ±1 *SEM*.

2.4.3 Interim Discussion

To ensure efficient goal-directed behaviour, internal representations must be protected from irrelevant perceptual distractors as well as external intervening tasks. Here, I provide evidence that temporal expectations help overcome the detrimental impact of both types of interference on visual working memory. The current results demonstrate that temporal expectations improved working-memory performance irrespective of the type of interference – that is, even when interference required a response. Because temporal expectations also improved working-memory performance in interrupter trials, it can be concluded that this benefit is unlikely driven solely by increased suppression of the external sensations. Instead, the present findings suggest that shielding of internal contents may provide a potent source to mitigate interference during working-memory retention. This is further supported by the finding that temporal expectations protected working memory even when enhancing performance to the external intervening task.

2.4.3.1 The Detrimental Effects of Interference

In addition to these main insights, I replicated previous findings demonstrating working memory to be substantially more impaired following interruptions as opposed to distractions (Bae & Luck, 2019; Berry et al., 2009; Clapp et al., 2010; Hakim et al., 2021; Mishra et al., 2013; B. Wang, van

Driel, et al., 2019; Zickerick, Kobald, et al., 2021). Besides calling on attentional control processes and mental workspaces that may have been concurrently active for working memory (e.g., Bae & Luck, 2019; Barrouillet et al., 2007; Souza & Oberauer, 2017), interrupters – which required a manual response – may also have directly affected the preparation of memory-guided actions. In the present task design, items were associated with specific actions which likely allowed action plans to be coactivated together with visual representations (Boettcher et al., 2021; van Ede, Chekroud, Stokes, et al., 2019). As such, interrupters may have additionally interfered with action plans in working memory, yielding a more detrimental effect on performance than was elicited by visual distractors not requiring any manual response. In agreement with this, Zickerick, Kobald, et al. (2021) showed that interrupters, but not distractors, were detrimental to the modulation of EEG mu-alpha activity – a neural signature linked to action preparation (McFarland, 2000; Neuper et al., 2006; Salmelin & Hari, 1994). Thus, in addition to the requirement to reactivate sensory representations after interruption (Clapp et al., 2010; Sakai et al., 2002), the impeded retrieval of action plans following interrupters may further account for the greater memory loss – a possibility to be more thoroughly tested in future research.

Besides interrupters having a detrimental influence on working memory, perceptual distractors also impaired memory-guided performance relative to no interference. Interestingly, previous research has yielded mixed evidence regarding the ability of perceptual distractors to impair working-memory performance (Bettencourt & Xu, 2016; Rademaker et al., 2019 [Experiment 1]; Zickerick et al., 2020; for a review see: Xu, 2017, 2020). However, in studies showing working-memory detriments, a high categorical overlap between memory target and distractor is suggested to play an important role (for a review see: Lorenc et al., 2021). In the present work, the interfering item and the working-memory content were defined in the same dimension (i.e., orientation), which likely increased the disruptiveness of perceptual distractors. Future studies will be informative to characterise more comprehensively how sensory and task-demand variables affect internal representations stored in working memory.

2.4.3.2 Working-Memory Protection via Shielding

It has previously been shown that working-memory performance can improve when interfering items gained less attention during retention (Bonnefond & Jensen, 2012; Clapp et al., 2010). Building on these results, the current findings demonstrated that working memory can also improve through anticipation of interference, even when the source of the interference itself cannot be suppressed, as it was the case for interrupters. In fact, RTs were even faster for predictable early interrupters (i.e., 250 ms after encoding offset), indicating that increased attention to the interfering item can co-occur with better working-memory performance. This might potentially be mediated by active allocation of attention to expected interference, as recently demonstrated in a related working-memory task (Makovski, 2019). Thus, the present results argue for a second route by which interference anticipation can facilitate working-memory performance – by shielding of the internal representations, rather than suppressing the external inputs. Critically, I do not intend to suggest that shielding abolishes the detrimental effects of interfering events altogether, given that some information about the external event entered working memory. Instead, as this experiment has demonstrated, working memory performance is better protected from sources of interference when these sources come at expected moments in time. Hence, while temporal expectations versus temporal uncertainty may not follow an all-or-none shielding rule, the key finding is that the ability to predict interference in time diminishes its impact on visual working memory.

While the present findings support the concept of memory shielding, I do not wish to discount the importance of distraction suppression as another mechanism for handling interference. Previous work has demonstrated that perceptual distractors are suppressed while ongoing memory content is maintained (Bonnefond & Jensen, 2012; de Vries et al., 2019; Getzmann et al., 2018; Payne et al., 2013). More specifically, Bonnefond & Jensen (2012) found evidence for suppression through phase shifts within the alpha band prior to the onset of a

temporally predictable distractor. Moreover, task-irrelevant distractors – but not interrupters which required attention – elicited neural signatures reflecting suppression (Hakim et al., 2021).

Nonetheless, suppression alone cannot account for all of the current findings. If the effects were solely due to suppression, I would not have observed improvements in working-memory performance when interference necessitated attentional processing and manual responses. Likewise, task performance to the interrupter would have been unlikely to improve. I showed that shielding (and not suppression) operates in interruption blocks, where suppression is not a viable option (**Figure 2.9**, left). However, based on the present experimental design, I cannot determine whether suppression or shielding, alone or together, operated in distraction blocks (**Figure 2.9**, right). Thus, in future work, it will be of interest to use neural measures to adjudicate between the two mechanisms in different task scenarios. For instance, multivariate decoding of EEG or MEG signals with high temporal resolution may expose differential neural signatures linked to the handling of each interference type. This could inform us whether there is a default mechanism by which temporal expectations are enacted to protect working memory, or whether the mechanism utilised depends on the task and source of interference at hand.

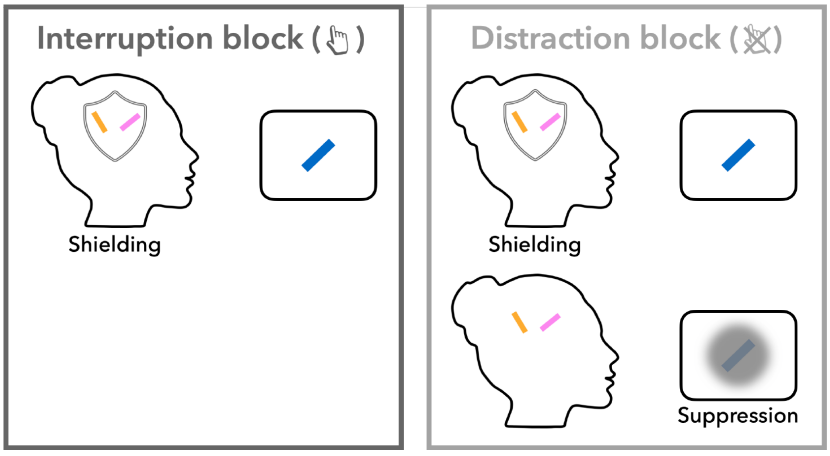

Figure 2.9. Shielding and suppression might operate independently based on current task demands. The experiment revealed that suppression alone cannot account for all the findings, as indicated by the temporal-expectation benefit in interruption blocks. However, future studies need to determine whether interference is suppressed, or working-memory representations also shielded against perceptual distractors.

2.4.3.3 Possible Neural Mechanisms Subserving Working-Memory Shielding

This experiment raises an interesting question regarding how the two mechanisms – suppression of external input and shielding of internal representations – could operate on a neural level. Suppression of anticipated interference during working-memory retention has been linked to modulations of oscillatory activity, such as increases in midfrontal theta (de Vries et al., 2019) and posterior alpha power (Bonnefond & Jensen, 2012; de Vries et al., 2019; Payne et al., 2013), presumably reflecting an inhibition of sensory areas involved in processing of distractors. However, there is less consensus as to how memory contents themselves can be protected against interference. Over the past decade, different types of evidence have emerged regarding the neural mechanisms of how items are held in working memory each offering unique possibilities as to how internal representations could be shielded from interference (Kamiński & Rutishauser, 2020; Serences, 2016). For example, the traditional assumption that working memory relies on sustained neural firing (Funahashi et al., 1989; Fuster & Alexander, 1971; Goldman-Rakic, 1995; E. K. Miller et al., 1996) has recently been challenged. Activity-silent mechanisms have been proposed to underlie working-memory retention through synaptic weight changes (Masse et al., 2020; Mongillo et al., 2008; Stokes, 2015). Neural decoding analyses suggest that memory content may transition to a latent state while other distractors or working-memory items are processed, before re-emerging into an active state when becoming task-relevant (LaRocque et al., 2016; Lewis-Peacock et al., 2012; Sprague et al., 2016). Furthermore, as these silent memories do not seem to bias perception (Mallett & Lewis-Peacock, 2018) or be manipulated without prior reactivation (Trübutschek et al., 2019), it is conceivable that information stored in a latent state might be less susceptible to disruption by interference (Lorenc et al., 2021) – an exciting possibility that remains to be carefully investigated.

In the current work, the time of the working-memory task itself was always fully predictable. Because of this, the momentary task relevance of the memory contents could be modulated by the passage of time (as in Jin et al., 2020; Olmos-Solis et al., 2017; van Ede et al.,

2017; Zokaei et al., 2019). Foreknowledge of when the memory contents become relevant may as well play a role in the ability to shield internal representations from external interference, and it may allow internal representations to be momentarily deprioritised and facilitate interference handling. For example, temporal expectations regarding the onset of interference may indicate when representations should be transformed into a format that is more resistant to interference, such as into an activity-silent code in non-sensory areas. However, it should be noted that this assertion remains purely speculative and remains to be followed-up with dedicated neuroscientific investigations.

2.4.3.4 Ruling Out Alternative Interpretations

As a final point, I would like to consider possible competing interpretations that could account for the shielding effect. First, it is possible that temporal expectations modulate the extent to which interference is encoded into working memory by altering gating mechanisms. When interference is temporally unpredictable, the likelihood that the interfering item will mistakenly enter working memory could be increased, thereby potentially displacing one of the relevant memory representations. Although analyses indicated that interference was encoded to some extent or in some trials, there was no difference between temporally predictable and unpredictable interference in terms of the probability of swaps or the response bias. This finding makes the explanation of interference encoding less likely. However, some caution may be pertinent when interpreting these null results.

Alternatively, working memory might have been conserved by saving attentional resources when the interrupter task was temporally predictable. Yet, when splitting the data into the three interference onsets, I only found a significant effect of temporal expectation for interrupter RTs at 250 ms after encoding offset. In contrast, the results demonstrated a temporal-expectation benefit in the working-memory performance following interrupters occurring at 250 ms, 1000 ms, and 1750 ms after encoding offset. The fact that I observed an effect of temporal

expectations on working memory without effects on the interrupter task (i.e., for 1000 ms and 1750 ms) argues against the possibility that temporally predictable interrupters were simply less interfering because they saved attentional resources by enabling faster responses to the interfering item.

Lastly, the necessity to divide and/or switch attention between internal (i.e., memory content) and external (i.e., interference) space might be reduced based on the temporal predictability of interference. Being able to anticipate when interference will occur might lower the need to divide or shift attention between the two attentional domains. Nevertheless, this alternative interpretation of our results does not discount the idea of memory shielding – with reduced demands to be aware of the external event, internal attention can stay focussed on memory contents, thereby protecting them from anticipated interference.

2.4.3.5 Conclusion

In conclusion, the present study shows firstly that temporal expectations help mitigate interference during visual working memory, and additionally that the influence of proactive temporal anticipation of interference engages processes of memory shielding. In future studies, it will be interesting to reveal the exact (neural) mechanisms that support the handling of these various sources of interference during working-memory maintenance.

2.5 Discussion

Experiment 1 and 2 examined how cross-domain modulations are influenced by temporal expectations – that is, how temporal regularities regarding one domain affect the other domain. Experiment 1 demonstrated that temporally guided internal attention modulates external-attention performance. Experiment 2 took the inverse approach and demonstrated that internal attention can be protected from external interference when there is a temporal regularity associated with external events occurring during the working-memory retention interval.

Collectively, these two experiments showed that in situations requiring both external and internal attention, significant crosstalk occurs between the two. Critically, this interaction can be moderated by applying temporal expectations to either domain.

Beyond considering the role of temporal expectations within the realm of cross-domain modulations, it is crucial to acknowledge that these expectations themselves provide a natural mechanism for attentional shifts. Previous studies have revealed that attentional transitions within domains can be dynamically steered by the incidental, implicit learning of temporal structures, prioritising the processing of external or internal contents at moments when they become most relevant (van Ede et al., 2017; Williams et al., 2023; Zokaei et al., 2019). Furthermore, as demonstrated in this chapter, temporal expectations may also play a crucial role in supporting attentional shifts between different domains. In Experiment 1, anticipating when working-memory would be probed likely helped participants to transition more efficiently from the external intervening task to internal contents, in preparation for memory-guided behaviour. Similarly, in Experiment 2, temporal expectations regarding the interrupter onset may have supported shifts from the internal to external domain (i.e., from memory encoding and maintenance to the external intervening task). With increased certainty about the timing of external events, internal attention may have remained focussed on mnemonic contents, shifting to interrupters only at the most opportune moment. Conversely, in the presence of perceptual distractors, temporal expectations may have signalled the need to maintain internal focus throughout the retention interval, thereby safeguarding memory contents from anticipated disruptions by preventing unnecessary shifts.

So far, the experiments presented in this thesis have demonstrated cross-domain behavioural consequences arising from task contexts in which individuals encounter both external and internal information within a common time frame, necessitating shifts in attention (or the avoidance thereof) between the two domains. In Chapter 3, I will employ a similar task design and investigate the temporal dynamics of reactivations of internal representations following external

interruptions, capturing the natural flow of shifting attention from one domain to another. By employing EEG, my goal is to precisely identify the moments when different working-memory contents re-emerge, offering insights into the reactivation process that occurs during the shift from external to internal attention.

3 Reactivating Internal Contents During External-to-Internal Shifts

3.1 Abstract

During natural behaviour, we must often maintain internal representations in working memory while concurrently engaging with external stimuli. Though previous research has demonstrated working memory to be susceptible to external interference, our internal representations are generally able to withstand such challenges without experiencing complete failure. This study aimed to investigate the temporal dynamics of reactivations of internal representations following external-task engagements, capturing the natural flow of shifting attention from external to internal attention. To this end, I designed a visual-motor working-memory task in which participants were retrospectively cued about a target item during the retention interval. In the majority of trials, following the presentation of the retro-cue, working memory was interrupted by an interfering item that required a response. I employed EEG time-frequency analyses to track the selection of visual information (i.e., item location) and action planning (i.e., prospective response hand) throughout the retention interval. The findings revealed a transient selection of visual and motor information following the retro-cue in external-task trials. Internal contents were revisited immediately after the completion of the external task, rather than being selected just-in-time when internal information was required for memory-guided behaviour. Strikingly, the revisiting of the internally stored visual information was behaviourally relevant. In sum, this study yields new insights into how internal representations are maintained and dynamically shifted to when engaging in ongoing external behaviour.

3.2 Introduction

A key challenge for visual working memory is to maintain task-relevant sensory information, while concurrently handling other external tasks that demand our attention. For example, when cycling to a recently opened ice cream shop, we must retain a memorised set of directions while also responding to external tasks, such as reacting to abruptly stopping cars and navigating around potholes in our path. While previous research has demonstrated that working memory can be disrupted by external-task interference, there typically remains residual internal information that we can return to once the external task is completed (Lorenc et al., 2021). To understand adaptive memory-guided behaviour in situations that additionally require external attention, it is important to investigate the processes involved in reprioritising internal representations after engaging in an external task.

The reprioritisation of internal representations likely extends beyond visual properties. There has been increasing recognition that the purpose of working memory surpasses the mere storage and recollection of past sensory information. Instead, detailed sensory information in working memory serves to guide future behaviour (Chatham & Badre, 2015; Myers et al., 2017; Nobre & Stokes, 2019; Olivers & Roelfsema, 2020; van Ede, 2020; van Ede & Nobre, 2023). Consistent with this notion, recent studies have demonstrated that the encoding and maintenance of visual information in working memory often occurs concurrently with action planning for upcoming behaviour (Boettcher et al., 2021; Ester & Weese, 2023; Nasrawi et al., 2023; Nasrawi & van Ede, 2022; D. Schneider, Barth, & Wascher, 2017; van Ede, Chekroud, Stokes, et al., 2019; Zickerick, Kobald, et al., 2021). Hence, it appears plausible that both visual and action representations should be reprioritised following engagement in an external task. For instance, one study has investigated motor signatures during a working-memory delay wherein an external intervening task was embedded (Boettcher et al., 2021). This study utilised a novel working-memory task in which memorised sensory information was directly linked to prospective actions (see also: Nasrawi et al., 2023, 2024; van Ede, Chekroud, Stokes, et al., 2019). The findings

revealed that action encoding into working memory occurred prior to the onset of the external task, and that motor signatures were reinstated thereafter in the service of memory-guided behaviour. Although this work helped to clarify *whether* action representations of an internal representation are reinstated following an external task, it did not answer *when* this reinstatement takes place. Specifically, this experiment did not allow to disambiguate whether the reselection of motor signatures occurred immediately following the external task or in anticipation of the working-memory task. Examining the temporal dynamics of motor reinstatement is essential for gaining insight into the time course of shifts between the external world and our internal representations.

With regards to visual reselection, past work has demonstrated that visuo-spatial information reactivates immediately following an interrupting task, and sometimes even before its completion (Zickerick, Rösner, et al., 2021). However, the interrupting and working-memory tasks in this previous study were fundamentally different. The working-memory task required participants to memorise and reproduce a tilted bar, whereas the interrupting task involved either a numerical or arithmetic comparison. Hence, transitioning between the interrupting task and the working-memory task also involved shifting from verbal to visual content. Given that working memory is believed to comprise of independent visual and verbal subsystems (Baddeley, 1992; Brooks, 1968; Logie et al., 1990), it is reasonable to assume that a shift from verbal to visual content may encounter less disruption than shifting between visual contents.

While both studies revealed that participants shifted back to internal contents (Boettcher et al., 2021; Zickerick, Rösner, et al., 2021), it remains unclear when exactly visual and motor representations become reselected after engaging in an external task that shares stimulus and response properties with the internal contents. For instance, internal reselection may emerge just-in-time, that is, only at the moment mnemonic representations are needed to perform a memory-guided action. Alternatively, reselection may occur immediately following the external task. To address this question, I used EEG to track the neural dynamics of visual information and

prospective manual actions in a task that systematically manipulated the onset of an external task during the working-memory delay. This task additionally manipulated the temporal predictability of the external task, thus enabling an assessment of whether working-memory performance was enhanced when external-task onset was predictable versus unpredictable.

To preview the results, I show a partial replication of the behavioural benefits of temporally predictable external interference (as observed in Chapter 2.4), that the reactivation of visual and motor signatures occurs immediately following the external-task response, and that variability in the onset of memory-guided behaviour is associated with corresponding changes in the latency of visual reselection after the external task.

3.3 Methods

3.3.1 Participants

The study was approved by the Central University Research Ethics Committee of the University of Oxford. Thirty-four volunteers participated. Two participants were excluded following an a-priori behavioural trial-removal procedure (see 3.3.4 Behavioural Analysis). Furthermore, two datasets were excluded due to the absence of eye tracking data, which made it impossible to verify that participants kept their gaze fixed during the task. All remaining 30 participants (age range: 20 to 40; mean age: 26.2; gender: 16 female, 14 male, 3 non-binary; handedness: 25 right-handed, 5 left-handed) had normal or corrected-to-normal vision. Individuals provided written informed consent before participating in the study and were compensated at a rate of £20 per hour.

3.3.2 Task and Procedure

This experiment adapted a recently developed visual-motor task (see Boettcher et al., 2021; Nasrawi et al., 2023; van Ede, Chekroud, Stokes, et al., 2019) and resembles the task designs employed in Chapter 2. Participants performed a working-memory task (**Figure 3.1**). On each trial,

they had to memorise the orientations of two differently coloured bars. At the end of each trial, participants were asked to reproduce the orientation of one of these bars, indicated by a correspondingly coloured retro-cue which appeared during a memory delay. A key feature of this task was that each bar's orientation was directly linked to the response hand required to reproduce its tilt. Since the prospective response hand and cued item's location and were manipulated orthogonally, this task design allowed me to disambiguate EEG markers of motor selection (i.e., required response hand) from those of visual selection (i.e., location of memorised item) following the retro-cue onset. Moreover, to investigate the reactivation of internal visual and motor traces during external-to-internal shifts, an external task was included in the majority of blocks following the presentation of the retro-cue (block type: no-external-task block vs. external-task block). This external task could occur at one of three potential time points: either 600 ms, 1400 ms, or 2400 ms after the onset of the retro-cue (external-task onset: early vs. medium vs. late). For trials within a fixed block, the external task occurred at only one of these three time points; whereas for trials within a variable block, the external task could occur early, intermediate, or late with equal likelihood (temporal predictability: fixed block vs. variable block).

At the start of each trial, two tilted bars were simultaneously presented against a grey (RGB value: [102, 102, 102]) background for 200 ms. One bar was always positioned to the left and the other to the right of a central, white fixation cross (RGB value: [255, 255, 255]). Independent of the location, one of the bars was tilted to the left (anticlockwise) and the other to the right (clockwise). To avoid angles being too close to the vertical and horizontal meridians, the items' angles were randomly drawn in increments of ±1° between ±10° and ±80°. Within each block, a bar oriented leftwards or rightwards was equally likely to appear in either the left or right position. The stimuli were each approximately 0.7° in width and 5° in length, centred at a viewing distance of 5.5° visual angle left or right from fixation. On each trial, the colours of the memory items were randomly drawn from a set of three highly distinguishable colours (RGB values: blue

[0, 178, 255], orange [239, 169, 33], pink [255, 0, 246]). At encoding, both lateralised items were equally likely to be probed, rendering them equally relevant.

Visual encoding was followed by a memory delay of 800 ms, during which the fixation cross remained on the screen. After this delay, the central fixation cross changed colour for 200 ms to match one of the previously presented bars. This retro-cue was 100% valid, therefore always indicated the to-be-reported item. Retro-cue presentation was followed by a delay interval of 3000 ms, during which the white fixation cross was presented.

During this delay period, in two thirds of blocks, two additional tilted bars were simultaneously presented for 100 ms, appearing in the same size and locations as those in the memory display (referred to as the "external task"). Both bars in this external task had the same orientation and colour. Their angle was pseudo-randomly drawn in increments of ±1° between ±10° and ±80°, with leftwards and rightwards tilts being counterbalanced within each block. Their colour always differed from those of the two preceding memory items. If the bars were tilted to the left (right), participants pressed the F (J) key on the keyboard with their left (right) index finger as quickly as possible.

Depending on the block type, the external task was either presented at a fixed or variable time point (temporal predictability: fixed onset vs. variable onset). In fixed-onset blocks, the external task always appeared at a specific point in time during retention (interference onset: 600 ms or 1400 ms or 2400 ms after retro-cue onset). In variable-onset blocks, the external-task items were equally likely to occur at any of these three time points. That is, in each variable-onset block, one third of the trials presented the external items at 600 ms, one third at 1400 ms, and one third at 2400 ms after retro-cue onset. The total length of the memory delay was 3000 ms, regardless of when and whether the external task appeared.

Following the memory delay, the fixation cross transitioned to black (referred to as the "probe"), indicating that the previously cued item should now be reported. The probe was deliberately presented in a non-target colour to ensure participants utilised the retro-cue when

performing the task. Participants were never probed about the external-task items. Following the appearance of the probe, participants had unlimited time to decide on their response. To report a leftwards (rightwards) angle, participants were asked to press the F (J) key on the keyboard using their left (right) index finger. After response initiation, a visual response dial was displayed on the screen, always starting in vertical position. The response dial had the same diameter as the length of the bars (5°) and always appeared surrounding the fixation.

The dial rotated leftwards when the F key was pressed and rightwards when the J key was pressed (either by holding the key down or pressing the key repeatedly) at a speed of 1/10° per millisecond. Critically, the dial could only be rotated in the direction that was initially indicated by the participant. For example, if a participant started pressing the F key after the probe, the dial would only move leftwards, and it would therefore not be possible to move the dial back towards the right with the J key. Since the response dial always started in a vertical position and could not be rotated beyond ±90°, a leftwards (rightwards) oriented bar could only be correctly reported with a left (right) key. Consequentially, the hand required for responding was directly linked to the angle of the bar that was probed. Critically, because the cued item's tilt and location were pseudo-randomised with respect to the tilt of the external-task items across trials (e.g., in a given trial, a left-positioned memory bar was equally likely to be paired with a left- or right-hand response, in both the external task and the working-memory task), it was possible to independently characterise neural activity reflecting the target item's memorised location and the prospective response hand associated with the target item's tilt. Once the response dial aligned with the memorised tilt of the memory item, participants pressed the space bar to verify their response.

After a 100-ms delay in which only the fixation cross stayed on the screen, participants received a 500-ms feedback regarding their working-memory and external-task performance. Working-memory performance was indicated by number ranging from 0 to 100, with 100 indicating a perfect report and 0 indicating that the adjusted orientation was perpendicular to the angle of the target item. If working-memory performance was below the number 75, the number

was presented in red font (RGB value: [255, 0, 0]), otherwise appearing in green (RGB value: [0, 240, 15]). This was to motivate participants to respond accurately to the working-memory item. When participants took longer than 800 ms to respond to the external task, responded using the wrong key, or did not respond at all, the word “REACT” was displayed in red font, otherwise, the word “REACT” appeared green font. This was to motivate participants to respond accurately and quickly to the external task. In no-external-task blocks, this latter feedback was not presented. Trials were separated by an ITI randomly drawn between 750 and 1000 ms. Between blocks, participants received additional feedback about their average accuracy in the external and working-memory task of the previous block.

The experiment consisted of 648 trials divided across 27 blocks, each including 24 trials. Of the 27 blocks, 9 were no-external task blocks while the remaining 18 included an external task during the memory delay. External-task blocks were further subdivided depending on the temporal predictability of the external task, having a fixed onset in 9 blocks (3 blocks each of 600 ms, 1400 ms, or 2400 ms after retro-cue onset, respectively), and a variable onset in the other 9 blocks (pseudo-randomised to equiprobably occur at 600 ms, 1400 ms, or 2400 ms after retro-cue onset). As such, the total number of trials where the external task would appear at any of the three time points was equal between the corresponding fixed-onset blocks and across the variable-onset blocks.

The experiment was divided into three consecutive 30-minute sessions, with a 5-minute break between each session. Unbeknownst by the participants, the sessions where further split into triplets, each containing one no-external-task block, one fixed external-task block, and one variable external-task block. The order of blocks within these triplets as well as the order of triplets was randomised.

Before the start of each block, participants were informed whether they would have to respond to an external task. However, they were not informed about the temporal predictability (i.e., fixed vs. variable) or about the three possible interference onsets (i.e., 600 ms, 1400 ms,

2400 ms). To familiarise themselves with the task, participants performed variable-onset external-task blocks until they reached an accuracy higher than 85% in the working-memory task and 90% in the external task. Participants completed between 2 and 5 practice blocks to achieve this threshold.

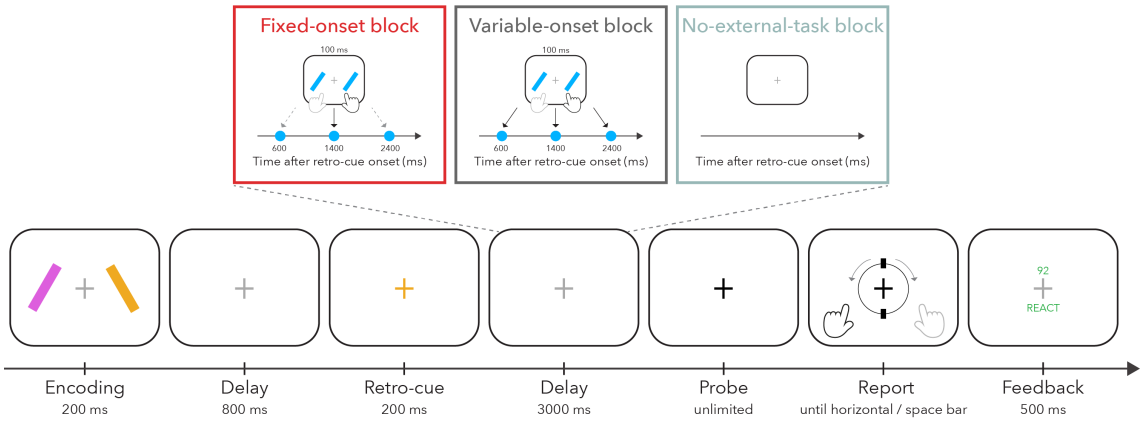

Figure 3.1. Task schematic. Trials started with an encoding display consisting of two, lateralised tilted bars. Participants' task was to remember the angle of both bars, of which one had to be reported at the end of the trial. After a delay, a retro-cue appeared indicating the to-be-reported item. In most of the blocks, external-task items occurred for 100 ms during the memory delay following the retro-cue. In this external task, participants had to indicate whether items were tilted to the left or to the right. In fixed-onset blocks, the external task occurred at a fixed time point within a block (i.e., at 600 ms, 1400 ms, or 2400 ms after retro-cue onset). In variable-onset blocks, the external task occurred equally likely at each of the three possible time points. The delay between retro-cue offset and probe onset was always 3000 ms long. After the delay, the central fixation cross turned black, indicating that the previously retro-cued item should now be reported. Participants reported the memorised tilt using a response dial and were provided feedback at the end of the trial.

3.3.3 Apparatus and Data Acquisition

The experimental script was generated in Python using the PsychoPy package (version 2021.2.3; Peirce et al., 2019). Data acquisition took place in a dimly lit testing booth. Participants sat in front of a monitor (27-inch Acer XB270H; resolution 1920×1080 pixels; refresh rate 100 Hz; screen width 60 cm) at a viewing distance of 100 cm. An eye tracker (EyeLink 1000, SR Research) was positioned on the table approximately 15 cm in front of the monitor. Before data acquisition and between blocks, eye position was calibrated using a 5-point calibration and validation method. Gaze was continuously tracked binocularly (or monocularly if binocular tracking was not possible)

at a sampling rate of 1000 Hz, and participants were instructed to maintain central fixation throughout the trial.

EEG data were acquired using an SynAmps amplifier and Neuroscan data acquisition software (Compumedics Neuroscan, Abootsford, Australia). Sixty-one electrodes were positioned across the scalp using the international 10-10 positioning system. An additional electrode was placed on each mastoid (i.e., behind the left and right ear), and a ground electrode was positioned on the forehead. Vertical and horizontal electrooculography (EOG) were recorded using a bipolar montage, which involved placing electrodes beside the canthi of each eye as well as above and below the left eye. During acquisition, data were low-pass filtered at 500 Hz and digitised at sampling rate of 1000 Hz. Electrode impedances at least below 10 kΩ and ideally below 5 kΩ were targeted. Data were referenced to the left mastoid during data acquisition.

3.3.4 Behavioural Analysis

Data were analysed in R Studio (Posit team, 2022). During data pre-processing, trials were removed if RTs (i.e., from probe onset to the first key press) exceeded 5000 ms or were 2.5 *SD* above the individual mean across the trials of all conditions. I also excluded trials in external-task blocks if participants did not respond, responded with the wrong key, or did not respond within 800 ms to the external task. Datasets with more than 15% of trials rejected during these pre-processing steps were removed from further analysis (n = 2). Moreover, an additional two participants were excluded due to missing eye-tracking data. After these exclusion steps, datasets from 30 participants remained in the main behavioural analysis, with an average of 92.61% (*SD* = 2.67) of trials retained.

The conditions of interest were block type (no-external-task block vs. external-task block), temporal predictability (fixed vs. variable), and external-task onset (early vs. medium vs. late). When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples *t*-test

and reported Cohen's *d* as a measure of effect size. For post-hoc *t*-tests, I reported Bonferroni-corrected *p*-values that I denote as $p_{Bonferroni}$. I used the ggplot2 package (version 3.3.6; Wickham, 2009) for plotting results. The within-subject standard error of the mean was calculated from normalised data using the approach from Morey (2008).

3.3.5 EEG Pre-Processing

EEG data were analysed in Python using MNE-Python (version 1.4.2; Gramfort et al., 2013). The data were re-referenced to the common average, downsampled to 250 Hz, and filtered between 0.1 Hz and 50 Hz. Bad sensors were manually identified and subsequently interpolated. On average, 0.87 ± 0.82 (*M* ± *SD*) bad electrodes were identified and interpolated per participant. I applied an independent component analysis (ICA) to regress out eye-movement-related activity. Data were then epoched from-200 to 4200 ms relative to the retro-cue onset and from-400 to 400 ms relative to the external-task response. No baseline correction was applied. I employed a generalised extreme studentised deviate (ESD) trial rejection approach to detect and exclude trials with high variance. In addition, I removed trials of external-task blocks in which participants did not respond to the external task. On average, 29.07 ± 21.52 trials were removed per participant, but never more than 15% of all trials. Finally, I applied a surface Laplacian transform (Perrin et al., 1989) to increase the spatial resolution of spectral modulations under consideration (as also done in: Boettcher et al., 2021; Nasrawi et al., 2023; van Ede et al., 2019).

3.3.6 Time-Frequency Analysis

I calculated time-frequency representations of power by convolving the data with Morlet wavelets across the frequency range of 3 to 40 Hz. For each frequency, I used a fixed 300-ms time window such that the number of cycles changed with each frequency.

To obtain a marker of visual selection, activity in predefined posterior visual electrodes (PO7/PO8) was contrasted between trials in which the memorised location of the cued item was

contralateral versus ipsilateral to the electrode of interest. I expressed this contrast as a normalised difference (i.e., [(contra-ipsi)/(contra + ipsi)] × 100) and averaged it across the left and right electrodes. To derive time courses of lateralised visual activity, I averaged the contralateral versus ipsilateral response within the predefined alpha-frequency band (8-12 Hz). Topographies of lateralised visual activity were generated by contrasting trials in which cues indicated visual content in the right versus left hemifield, expressed as a normalised difference between right and left trials for each electrode (i.e., right minus left).

In addition to analysing lateralised visual selection, I also investigated lateralised signatures of motor selection. The analysis of motor selection mirrors that of visual selection, with the following distinctions. To examine motor signatures, I contrasted activity in predefined central motor electrodes (C3/C4) between trials in which the prospective response hand was contralateral versus ipsilateral to the electrode of interest. Moreover, I computed time courses of motor lateralisation from the time-frequency contrasts by averaging across the predefined beta-frequency band (13-30 Hz). By independently manipulating item location and required response hand, I ensured that the EEG signatures of visual and motor selection were orthogonal in the trial average.

3.3.7 Relationship Between EEG and Behaviour

To investigate the relationship between neural activity after responding to the external task and subsequent memory-guided behaviour, I split trials by their RTs to working-memory probe for each participant. After sorting the data by RTs, I calculated the same contralateral-versus-ipsilateral time-frequency contrast of visual and motor selection as described above, separately for trials with fast responses (< median RTs) and slow responses (> median RTs) following the working-memory probe.

3.3.8 Statistical Analysis

Statistical evaluation of EEG data was carried out using a cluster-based non-parametric permutation approach, employing 1024 permutations and a cluster alpha of 0.05. This approach offers a solution to the challenges related to multiple comparisons in the statistical analysis of EEG data, especially when dealing with a large number of time-frequency comparisons (Maris & Oostenveld, 2007).

In addition, I utilised a jack-knife approach to identify latency differences between EEG time courses (J. Miller et al., 1998; see also: Nasrawi et al., 2024; van Ede, Chekroud, Stokes, et al., 2019; van Ede et al., 2018). For the jack-knife approach, I computed EEG time-course averages by excluding each of the *n* datasets once, which allowed me to obtain *n* leave-one-out grand averages. Next, I identified the time point at which the time course reached 50% of its largest negative deviation from zero (i.e., minimum). I used this approach to compare the timing of the visual alpha-band and motor beta-band time courses for fast versus slow working-memory trials. Jack-knifed standard errors, *t*-values, and *p*-values were adjusted according to J. Miller et al. (1998).

3.4 Results

3.4.1 Behaviour: Working-Memory Performance

Although it was not the primary objective of the research in this chapter, I aimed to ascertain whether the behavioural findings from Chapter 2.4 could be replicated in the current task design. In comparison to the experiment presented in Chapter 2.4, the current task design utilised slightly different onset times of the external task and included a retro-cue during the delay interval. Nevertheless, as in the previous chapter, I explore whether the shift from an internal working-memory representation to an external task may lead to a behavioural cost and whether some of this cost can be mitigated by temporal expectations of the external task.

First, I aimed to test whether external-task demands during the working-memory delay adversely affected working-memory performance. Previous research has shown that retro-cue prioritisation can protect internal representations from task-irrelevant distractors (Barth & Schneider, 2018; Makovski et al., 2008; Makovski & Jiang, 2007; D. Schneider, Barth, Getzmann, et al., 2017; van Moorselaar, Gunseli, et al., 2015; Vergauwe et al., 2023) by enhancing protection through top-down influences (Souza & Oberauer, 2016). However, the present findings revealed that prioritising a single item in memory could not fully safeguard the memorised contents from interference caused by an external task. As depicted in **Figure 3.2**, reproduction errors were systematically larger when participants had to shift to an external task during the maintenance interval in comparison to trials without external-task demands (**Figure 3.2A**; $t_{(29)} = 4.605$, $p < 0.001$, $d = 0.841$). A similar pattern emerged for RTs, with slower RTs in trials of the external-task blocks compared with trials of the no-external-task blocks (**Figure 3.2B**; $t_{(29)} = 2.423$, $p = 0.022$, $d = 0.442$).

Next, I compared working-memory performance across trials, distinguishing between those with a variable onset of the external task and those where the external task occurred at a fixed point during the delay. While I did not observe any difference in reproduction errors (**Figure 3.2C**; $t_{(29)} = 1.353$, $p = 0.186$, $d = 0.247$), RTs to the probe were delayed in variable as compared to fixed-onset blocks (**Figure 3.2D**; $t_{(29)} = 2.837$, $p = 0.008$, $d = 0.518$).

Therefore, working-memory performance appears to be mostly consistent with the results presented in Chapter 2.4, in that external-task demands negatively affect working-memory performance and these costs can be somewhat mitigated by temporal expectations. It is noteworthy that the temporal-predictability effect observed in RTs warrants careful interpretation, as it was influenced by the timing of the external task and was largely driven by late-onset trials (for details see: **Supplementary Table 8.7**, **Supplementary Figure 8.7**).

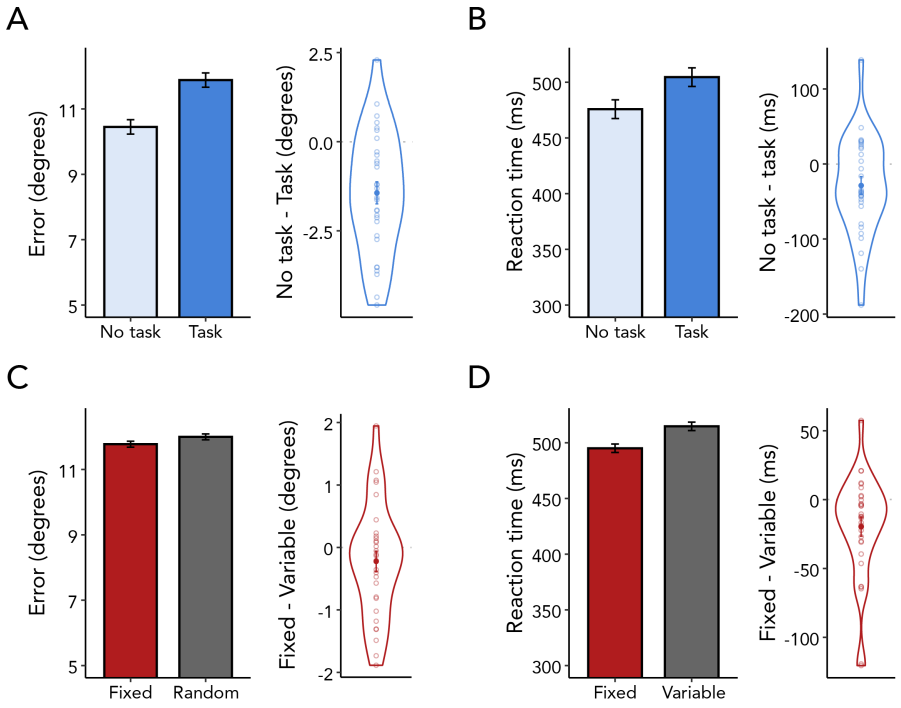

Figure 3.2. Working-memory performance as a function of block type and temporal predictability. **(A)** Reproduction errors were higher in trials with as opposed to without external task. **(B)** RTs were higher in trials with as opposed to without external task. **(C)** There was no difference in reproduction errors between trials with fixed versus variable external-task onsets. **(D)** RTs were higher in trials fixed as opposed to variable external-task onsets. Dots represent individual participants. Error bars indicate ±1 *SEM*.

3.4.2 Behaviour: External-Task Performance

Next, I examined whether RTs to the external task differed when it occurred at fixed times compared to variable times. I mainly focussed on RTs given the extremely high accuracy in completing the external task (error rates in %: M = 2.795, SEM = 0.336; **Supplementary Figure 8.8**), mirroring observations from Chapter 2.3 and Chapter 2.4. Participants responded more slowly to the external task when it occurred with a variable onset in the memory delay compared to a fixed onset (**Figure 3.3**; $t_{(29)} = 3.60$, $p < 0.001$, $d = 0.668$). This result is similar to the external-task performance observed in Chapter 2.4. However, as for the temporal predictability effect in RTs to the working-memory task, this finding has to be interpreted with caution since it was modulated by the exact onset of the external task (for details see: **Supplementary Table 8.8**, **Supplementary Figure 8.9**). In particular, the temporal-predictability effect in the external task was mostly driven by early-onset trials. The following EEG analyses will move away from examining

the temporal-predictability manipulation and instead focus on exploring the processes behind selecting and reselecting internal contents in working memory.

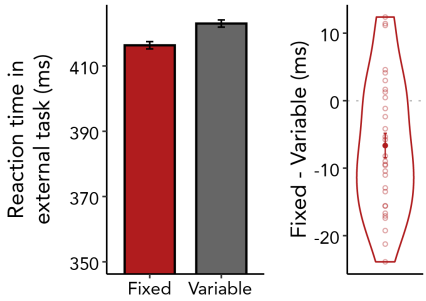

Figure 3.3. RTs in the external task as a function of temporal predictability. RTs in external task were slower for variable as compared to fixed onsets. Dots represent individual participants. Error bars indicate ±1 *SEM*.

3.4.3 EEG Time-Frequency: Internal Selection During No-External-Task Trials

First, I employed an EEG time-frequency analysis to investigate the dynamics of visual and motor selection after retro-cue onset. A well-established neural signature of visual selection is the lateralisation 8-12 Hz alpha-band activity in posterior electrodes contra- versus ipsilateral to the location of an attended item, both in perception (Gould et al., 2011; Myers et al., 2015; Thut et al., 2006; Wallis et al., 2015; Worden et al., 2000) as well as in working memory (e.g., Boettcher et al., 2021; Myers et al., 2015; van Ede, 2018; Wallis et al., 2015). In regard to motor selection, a well-established neural marker is the attenuation of 13-30 Hz beta-band activity in central electrodes contra- versus ipsilateral to the response hand associated with a to-be-reported item (Boettcher et al., 2021; Ester & Weese, 2023; Nasrawi et al., 2023; van Ede, Chekroud, Stokes, et al., 2019; Zickerick, Kobald, et al., 2021).

Prior research has demonstrated the encoding and subsequent preparation of prospective actions in working memory (Boettcher et al., 2021; Nasrawi et al., 2023; Nasrawi & van Ede, 2022). These findings have been complemented by other EEG studies revealing selection of motor signatures during the response stage of a working-memory tasks (Nasrawi et al., 2024; van Ede, Chekroud, Stokes, et al., 2019). However, while visual selection following retro-cues has

been frequently observed (Myers et al., 2015; Wallis et al., 2015), the investigation of whether motor selection similarly occurs during the prioritisation of memorised information – but prior to the requirement to act on it – remains comparatively less explored (but see: Ester & Weese, 2023). To examine whether action selection occurred alongside visual selection following the retro-cues, I focussed on trials without external-task demands during working-memory maintenance.

In line with previous studies, lateralisation of neural activity relative to the visual location of the cued memory item showed a clear relative attenuation of alpha-band activity in contra- versus ipsilateral electrodes PO7/PO8 (Figure 3.4A, cluster $p < 0.001$). The time course of the averaged 8-12 Hz alpha-band lateralisation in the posterior electrodes depicts that this effect arose early after retro-cue onset (Figure 3.4C, cluster $p < 0.001$). Topographies show the difference between trials where the cued item occurred on the right versus left side of the display, and further confirm that the attenuated alpha-band activity was particularly prominent in the left and right posterior electrodes, consistent with reflecting visual selection (Figure 3.4E).

Critically, the present task design also allowed for the investigation of neural lateralisation relative to the prospective response hand that was associated with the cued item. I observed clear relative attenuation of beta-band activity (13-30 Hz) in electrodes C3/C4 contra- versus ipsilateral to the required response hand associated with the cued item. (Figure 3.4B and D, time-frequency representation cluster: cluster $p < 0.001$; time-course cluster: cluster $p < 0.001$). In contrast to the transient selection of visual information, action selection was sustained following retro-cue onset, persisting until after probe presentation. Moreover, topographies confirmed that the beta-band modulations (13-30Hz) between trials requiring a right- versus left-hand response were particularly prominent in the left and right central electrodes, consistent with representing action selection (Figure 3.4F).

These data thus reveal that response-hand specific action selection occurred immediately following the retro-cue alongside the selection of visual information. While visual selection was transient, action-selection signatures remained sustained throughout the maintenance period.

This confirms that both visual and motor information can be prioritised promptly, even when memory-guided behaviour is not immediately needed.

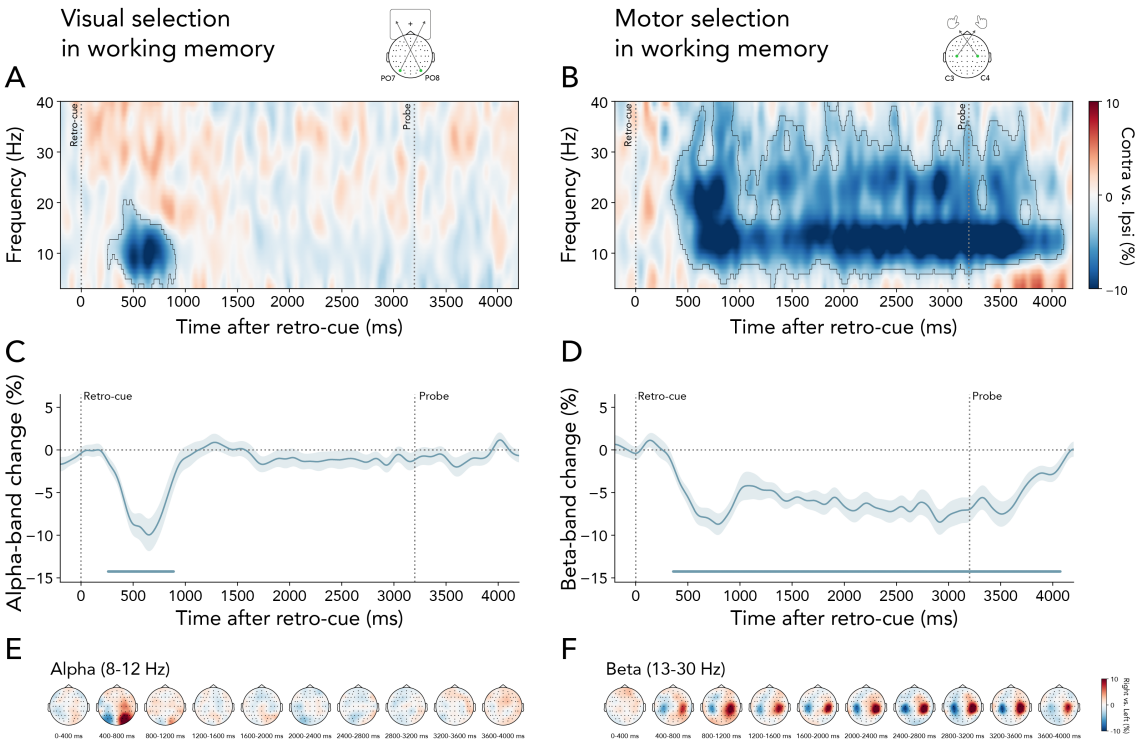

Figure 3.4. Visual and motor selection following retro-cue onset in the no-external-task condition. (A) Lateralised time-frequency spectrum contra- versus ipsilateral relative to the cued item's location (left/right), calculated in electrodes PO7/PO8. The black outline indicates a significant cluster. **(B)** Same as (A) but relative to the response hand required to reproduce the cued item's orientation (left/right), calculated in electrodes C3/C4. **(C)** Time course of lateralised visual activity at 8-12 Hz. The time course shows *M ± SEM* across participants. The horizontal line indicates a significant cluster. **(D)** Same as (C) but for motor activity at 13-30 Hz. **(E)** Topographies of the difference in 8-12 Hz alpha power between trials in which the retro-cue indicated the selection of the right or left item in memory (right minus left), averaged in steps of 400 ms after retro-cue onset. **(F)** Same as (E) but showing the difference in 13-30 Hz beta power between trials in which the retro-cue indicated the right or left response hand.

3.4.4 Two Possible Scenarios of Internal Reselection

Having confirmed visual and motor selection following the retro-cue, I next turned to my main question: when is internal information regarding the visual location and the associated response hand reactivated following an external task? I investigated this question by systematically manipulating the onset of the external task, along with orthogonally manipulating the cued item's location, the cued item's associated responses hand and the response hand required for the external task. This enabled me to disentangle two possible scenarios regarding when we recover

visual and motor signatures associated with working-memory contents. Scenario 1 suggests that visual and motor signatures could be reinstated just-in-time, that is, in anticipation of when working memory is needed for utilisation (**Figure 3.5A**). Importantly, in all three external-task onset conditions, the probe consistently occurred at the same time following the retro-cue. If participants prepared the internal information in a just-in-time manner, I would expect to see enhanced visual and motor activity at the time of the probe, which should be indistinguishable across trials with early, medium, and late external-task onset. Alternatively, in Scenario 2, if internal contents are reactivated immediately after the completion of an external task, visual and motor signatures should appear graded, depending on the timing of the external task (**Figure 3.5B**). In other words, the reactivation of working-memory representations should happen earliest in early-onset trials and latest in late-onset trials.

These two scenarios focus exclusively on the timing of internal-content reactivation and do not aim to illustrate potential variations between visual and motor signatures. Nevertheless, considering the insights from previous section (Chapter 3.4.3) and prior research (Boettcher et al., 2021; Ester & Weese, 2023), we can infer that visual signatures are likely to exhibit a brief reactivation phase, in contrast to motor signatures, which tend to demonstrate a more prolonged reactivation.

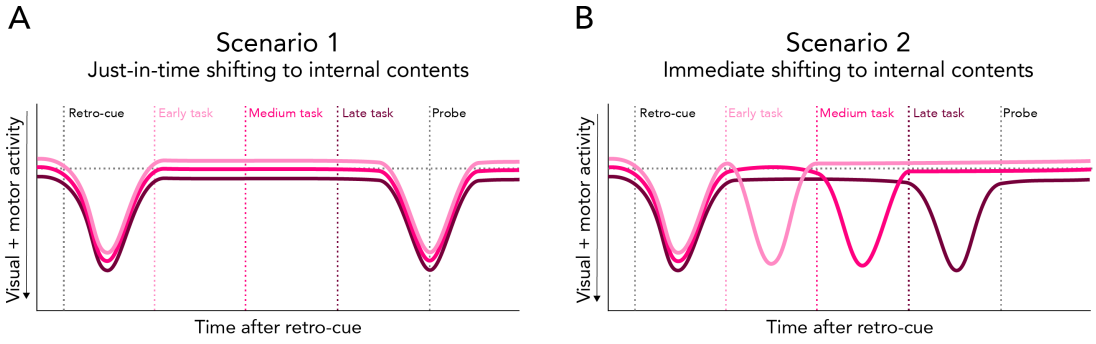

Figure 3.5. Two possible scenarios regarding the reactivation of working-memory contents. (A) Scenario 1: visual and motor signatures could become reactivated just-in-time when required to respond to the working-memory task. **(B)** Scenario 2: visual and motor signatures could become reactivated immediately following the external task. The magnitude differences shown in this figure (i.e., following the retro-cue) are intended solely to enhance the visibility of the three distinct time courses.

3.4.5 EEG Time-Frequency: Internal (Re)Selection During External-Task Trials

To address when reselection of visual and motor contents occurred, I employed time-frequency analyses separately for each of the three external-task onsets. I combined the data from both the fixed- and variable-onset trials, as there were no significant differences in lateralised visual and motor activity across the temporal-predictability conditions for early, medium, and late external-task onsets (**Supplementary Figure 8.10**).

Following the retro-cue, and in line with the no-external-task trials, I observed relative alpha-band attenuation in electrodes PO7/PO8 contra- versus ipsilateral to the cued item's location for each of three onset conditions (**Figure 3.6A**; early onset: 1st cluster $p < 0.001$, medium onset: 1st cluster $p < 0.001$, late onset: 1st cluster $p < 0.001$; upper, middle, and lower panel, respectively). This was succeeded by a second significant cluster of alpha-band attenuation in early- and medium-onset trials (early onset: 2nd cluster $p = 0.003$; medium onset: 2nd cluster $p = 0.033$). These results were complemented by time courses of averaged 8-12-Hz alpha-band lateralisation highlighting that significant attenuation occurred twice during the trial for all onset conditions (**Figure 3.6C**). While the initial significant cluster appeared shortly after the retro-cue in all onset conditions, the second cluster emerged in a graded manner: first in the early-onset, followed by the medium-onset, and finally in the late-onset condition (early onset: 1st cluster $p < 0.001$, 2nd cluster $p = 0.001$; medium onset: 1st cluster $p < 0.001$, 2nd cluster $p = 0.004$; late onset: 1st cluster $p < 0.001$, 2nd cluster $p = 0.005$). Topographies of lateralised visual activity for each of the three onset conditions are shown in **Supplementary Figure 8.11A**.

Motor signatures immediately following the retro-cue showed a similar pattern as the visual selection. For all three external-task onsets, time-frequency representations revealed a relative attenuation of beta power in electrodes C3/C4 contra- versus ipsilateral to the response hand required to reproduce the tilt of the cued item (**Figure 3.6B**; early onset: 1st cluster $p = 0.021$, medium onset: 1st cluster $p = 0.049$, late onset: 1st cluster $p = 0.011$; upper, middle, and lower panel, respectively). This transient lateralisation was followed by a more sustained lateralisation,

which appeared to vary temporally in accordance with the onset of the external task (early onset: 2nd cluster $p < 0.001$; medium onset: 2nd cluster $p < 0.001$, late onset: 2nd cluster $p < 0.001$). The time courses of averaged lateralisation in the 13-30 Hz beta-band confirmed a sequential pattern in the reactivation of motor signatures, occurring earliest in early-onset trials and latest in those with late onsets (**Figure 3.6D**; early onset: 1st cluster $p = 0.021$, 2nd cluster $p < 0.001$; medium onset: cluster $p < 0.001$; late onset: 1st cluster $p = 0.011$, 2nd cluster $p < 0.001$). Topographies of lateralised motor activity for each of the three onset conditions are shown in **Supplementary Figure 8.11B**.

Together, these results yield three important insights. First, they provide compelling evidence for Scenario 2, indicating an immediate reactivation of both visual and motor information following an external task. Second, even though motor activity was sustained throughout the maintenance interval in the no-external-task condition (**Figure 3.4D**), the initial selection of motor representations was transient across all external-task onset conditions, mirroring the selection of visual contents following the retro-cue. Lastly, while the reactivation of visual contents was short-lived, the reselection of motor information persisted for the remainder of the maintenance interval.

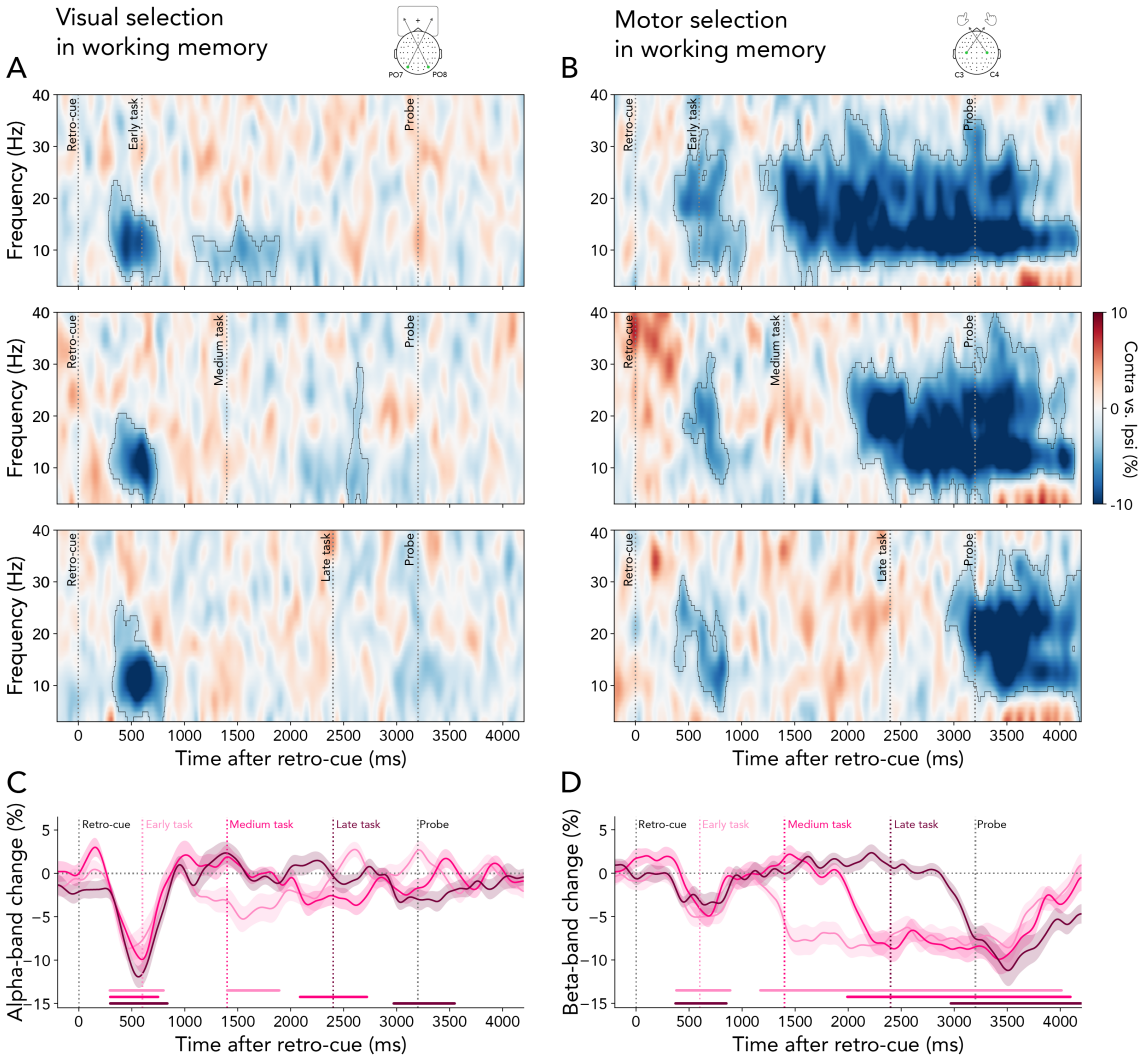

Figure 3.6. Visual and motor reselection following the external task. (A) Lateralised time-frequency representations contra- versus ipsilateral relative to the cued item's location (left/right), calculated in electrodes PO7/PO8. Top panel: early external task, middle panel: medium external task, bottom panel: late external task. Black outlines indicate significant clusters. **(B)** Same as (A) but relative to the response hand required to reproduce the cued item's orientation (left/right), calculated in electrodes C3/C4. **(C)** Time courses of lateralised visual alpha-band activity (8-12Hz) for early, medium, and late external-task trials. Time courses show *M ± SEM* across participants. Coloured horizontal lines indicate significant clusters. **(D)** Same as (C) but for motor beta-band activity (13-30 Hz).

3.4.6 EEG Time-Frequency: Internal Reselection After External-Task Responses

In line with Scenario 2, the results highlight that reactivation of visual and motor signatures occurred after the external task, rather than immediately before the onset of the probe. This prompts the question of how quickly participants shifted from the external task to the location and associated actions of internal contents. To address this, I combined all external-task trials and epoched the EEG data from -400 ms to 400 ms relative the external-task response. This epoching window was selected to avoid contamination of the results with the onset of the probe and the

following working-memory response. Given that RTs to the external task in the late-onset condition were on average 401.54 ± 2.23 ms (*M* ± *SEM*), this window effectively isolates the neural activity linked to the reactivation of internal contents without interference from later events.

Immediately after the external-task response, I observed a clear attenuation of alpha-band power in visual electrodes contra- versus ipsilateral to the location of the cued item (**Figure 3.7A** and **C**; time-frequency representation cluster $p < 0.001$; time-course cluster $p < 0.001$). Analysis of beta-band lateralisation relative to the required response hand yielded similar signatures of motor selection in central electrodes (**Figure 3.7B** and **D**; time-frequency representation cluster $p < 0.001$; time-course cluster $p < 0.001$). Visual and motor reselection had lateralised posterior and central topographies, respectively, within 0 to 400 ms after the external-task response (**Figure 3.7C** and **D**). Collectively, these findings imply that visual and motor reselection are closely linked to the moment participants responded to the external task.

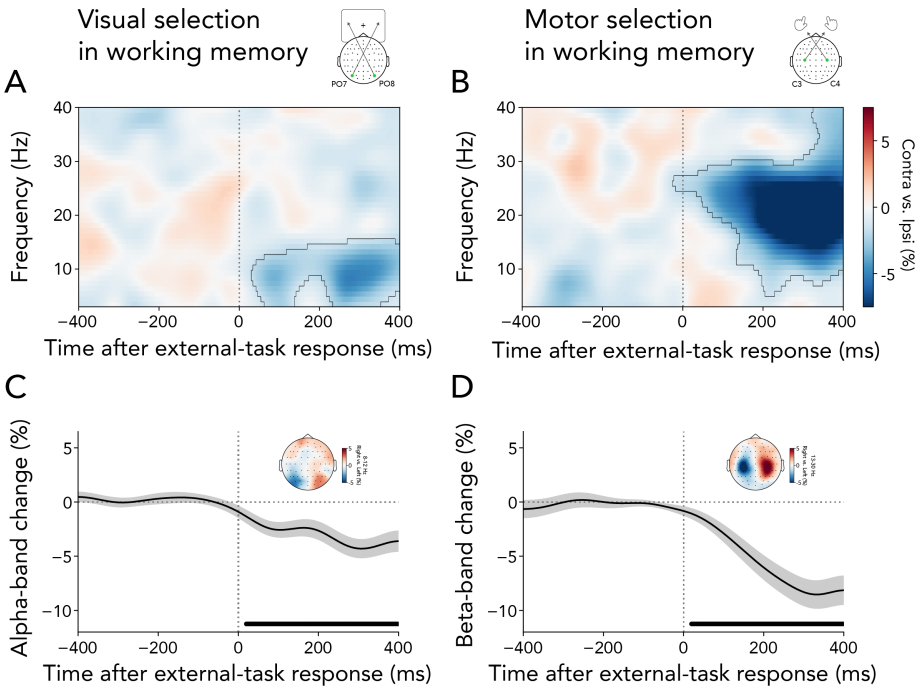

Figure 3.7. Visual and motor reselection immediately after external-task completion. (A) Lateralised time-frequency spectrum contra- versus ipsilateral relative to the cued item's location (left/right), calculated in electrodes PO7/PO8. The black outline indicates a significant cluster. **(B)** Same as (A) but relative to the response hand required to reproduce the cued item's orientation (left/right), calculated in electrodes C3/C4. **(C)** Time course of lateralised visual activity at 8-12 Hz. Time course shows *M* ± *SEM* across participants. Black horizontal line indicates significant cluster. The topography shows associated visual lateralisation between 0 and 400 ms. **(D)** Same as (C) but for motor activity at 13-30 Hz.

3.4.7 EEG Time-Frequency: Timing of Visual Reselection Is Behaviourally Relevant

Lastly, I investigated whether the reselection of visual and motor information following the external task could enhance behavioural performance in the working-memory task, particularly in terms of RTs. To this end, I split the EEG data into trials with fast versus slow responses to the working-memory probe. Next, I compared the lateralised alpha- and beta-band time courses for fast versus slow trials. This analysis can provide valuable indications into whether the speed of memory-guided behaviour is influenced primarily by visual reselection, motor reselection (i.e., the preparedness to act), both factors, or neither.

As shown in Figure 3.8A and B, I observed significant visual and motor reselection for both fast and slow trials (visual reselection: fast-trial cluster $p < 0.001$, slow-trial cluster $p = 0.013$; motor reselection: fast-trial cluster $p < 0.001$, slow-trial cluster $p < 0.001$). To quantify statistically the timing of reselection, I compared fast versus slow visual and motor time courses using a jack-knife approach (J. Miller et al., 1998; see also: Nasrawi et al., 2024; van Ede, Chekroud, Stokes, et al., 2019; van Ede et al., 2018). I identified the point at which the time courses first reached 50% of their most negative deflection, marking the midpoint of their minimum value. I observed significant differences in the latency of visual reselection, with alpha-band lateralisation occurring notably earlier in fast trials as compared to slow trials (Figure 3.8A; fast trials: 19.07 ± 31.80 ms [*M* ± *SEM*]; slow trials: 204.67 ± 40.14 ms; $t_{(29)} = -4.029$, $p < 0.001$). In contrast, there were no significant difference in the latency of beta-band lateralisation between fast and slow trials (Figure 3.8B; fast trials: 129.60 ± 34.55 ms; slow trials: 160.40 ± 22.48 ms; $t_{(29)} = -0.852$, $p = 0.401$).

In sum, these findings demonstrate that the timing of reactivating internal contents has significant consequences for subsequent working-memory performance. In particular, variations in the timing of visual reselection predicted the initiation of memory-guided behaviour.

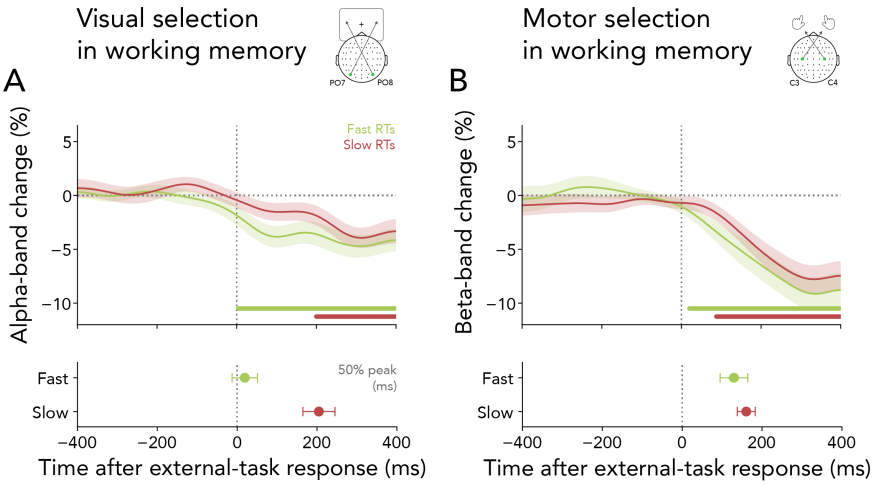

Figure 3.8. Variability in visual reselection predicts later memory-guided behaviour. (A) Time course of lateralised visual activity at 8-12 Hz, separated by fast and slow trials. Time courses show *M* ± *SEM* across participants. Coloured horizontal lines indicate significant clusters. Narrow panels below the time courses show averaged jack-knife estimates of the time the visual time courses first reach 50% of their minimum. Error bars indicate ±1 *SEM*. **(B)** Same as (A) but for motor activity at 13-30 Hz.

3.5 Discussion

Working memory maintains past sensations to guide immediate and future behaviours. Acknowledging this fact, recent frameworks of working memory have shifted focus to underscore not only its capability to encode and retain visual attributes but also its inherently action-oriented nature (Chatham & Badre, 2015; Myers et al., 2017; Nobre & Stokes, 2019; Olivers & Roelfsema, 2020; van Ede, 2020; van Ede & Nobre, 2023). Here, I built on these recent advancements and investigated when visual and motor signatures become reselected after engaging in an external task. The present results revealed that both visual and motor signatures reoccurred immediately following the external-task response, and that variability in the onset of memory-guided behaviour was associated with corresponding changes in the latency of visual (but not motor) reselection after the external task. Consistent with the findings outlined in Chapter 2.4, the introduction of an external task during the retention interval resulted in diminished working-memory performance compared to no-external-task trials. Furthermore, I observed a partial replication of the temporal-predictability effects on working-memory performance when comparing trials where the external task occurred predictably versus unpredictably.

3.5.1 Prioritising Visual and Motor Contents

The present task design represents one of the first attempts to track both visual and motor selection in response to a retro-cue. While past work has investigated visual and motor signals during encoding (Boettcher et al., 2021; Nasrawi et al., 2023; Nasrawi & van Ede, 2022) and during the response stage of working-memory tasks (Nasrawi et al., 2024; van Ede, Chekroud, Stokes, et al., 2019), less work has focussed on scenarios where internal selection is possible but responses are not immediately required. When no external task occurred during the working-memory retention interval, visual selection was associated with a robust but transient decrease in 8-12 Hz alpha power over posterior electrodes contra- versus ipsilateral to the location of the cued item. Previous research has linked the transient nature of alpha lateralisation following retro-cues to more efficient working-memory selection processes compared to more sustained alpha lateralisation (Mok et al., 2016). In contrast, motor selection was associated with a robust and sustained decrease in 13-30 Hz beta power over central electrodes contra- versus ipsilateral to the prospective response hand. The dissociation between transient alpha and sustained beta lateralisation replicates recent findings from Ester and Weese (2023) who employed a comparable retro-cueing visual-motor task.

One potential explanation for the differing durations of visual and motor selection could be that the spatial location of the cued item, indexed by alpha lateralisation, was not strictly necessary for memory-guided behaviour. Participants may have briefly redirected their attention to the hemifield where the cued item was previously presented, followed by recollection processes focussed on the detailed sensory information (i.e., the precise orientation of the item). These recollection processes may not necessarily be traceable through alpha lateralisation. Furthermore, the appearance of the dial in the centre of the screen may have prompted participants to redirect their visuospatial attention towards the middle, potentially contributing to the transient signatures of visual selection. In future analyses of these data, it will be intriguing

to examine the selection of task-relevant visual features associated with the cued item, such as its orientation, to determine whether these aspects exhibit transient or sustained characteristics.

During external-task trials, visual selection following the retro-cue closely resembled that of no-external-task trials. Interestingly, however, when compared to no-external-task trials, motor selection was much more transient in external-task trials, regardless of the onset and temporal predictability of the external task (Supplementary Figure 8.10). This transience of motor signatures prior to an external task resembles previous findings regarding the encoding of prospective actions into working memory (Boettcher et al., 2021). By systematically manipulating the delay between a memory and probe display, Boettcher et al. (2021) demonstrated that both the visual information and the associated motor planning were encoded simultaneously, regardless of the time at which memory was expected to be utilised. Action encoding was then followed by gradual motor preparation, adapted to when memory contents were required to guide behaviour (Experiment 1). Moreover, when introducing an external task during the working-memory delay, this first stage of action encoding was fully segregated from the subsequent stage of action preparation by the requirement to perform a separate motor response to the external task (Experiment 2).

In the current experiment, when visually inspecting the time course of no-external-task trials, it seems that motor signatures formed therein may also exhibit such a biphasic profile, despite a significant motor signal being persistently sustained until after probe onset. This included an initial peak of beta lateralisation between approximately 400 to 1000 ms, which was followed by a minor decline before gradually returning to peak magnitudes by the time of the probe. In external-task trials, these early action-selection and subsequent-preparation phases were separated by the occurrence of an external task (similar to the action encoding and preparation observed in Boettcher et al. (2021), Experiment 2). Since participants knew that they would have to engage in the external task, they may have immediately activated or prioritised the action plan following the retro-cue, then returned to baseline in anticipation of the external task.

Additionally, or alternatively, opting to promptly safeguard motor signatures by placing them into a latent state might be the most prudent approach, rather than deactivating them shortly before the onset of an external task irrespective of whether temporal expectations have been learned or not. In the future, it will be of interest to explore systematically whether motor signatures following retro-cues indeed exhibit a two-stage profile, consisting of time-invariant action selection followed by preparatory reselection.

3.5.2 Immediate Reselection of Visual and Motor Contents

The main focus of this study was to determine when visual and motor signatures become reselected following engagement in an external task – essentially pinpointing the moment when participants shift from external to internal attention. Previous research has indicated an immediate return of alpha lateralisation following the completion of an interrupting arithmetic task embedded within a working-memory delay (Zickerick, Rösner, et al., 2021). The present results build on these findings and reveal that transient visual and sustained motor reselection take place immediately after the completion of the external task, rather than occurring just-in-time before the onset of the probe. In addition, and in contrast to Zickerick, Rösner, et al. (2021), I show that this holds true even when the external and working-memory task share similar stimulus-response mappings.

The present results also demonstrate that the reactivation of internal contents following external-task engagement predicts later memory-guided behaviour. In particular, I observed earlier visual reselection, but not motor reselection, in fast as compared to slow working-memory trials. To date, other research has yielded mixed results regarding the relationship between visual and motor signatures and memory-guided behaviour. While Ester & Weese (2023) found no link between visual or motor selection after a retro-cue and subsequent working-memory performance, Nasrawi et al. (2024) reported variability in the onset of memory-guided behaviour to covary with the timing of motor but not visual selection. My findings complement this previous

line of research by examining the specific time point when participants reselect internal contents after engagement in an external task. It appears that during external-to-internal shifts, the primary determinant of the speed of memory-guided behaviour is the visual accessibility of the cued memory item, rather than the readiness to act on it.

A higher variability in visual but not motor reselection might be explained by the continued engagement of the motor system in the external task at the time of its completion. This engagement may limit the time to shift back to the action plan associated with the retro-cued item until the task is fully completed. In contrast, participants may be more flexible in redirecting their attention to the location of the cued item, since the response to the external task does not depend on visuospatial processing at a specific location. In future research, it would be intriguing to utilise a task design that incorporates task-irrelevant distractors to investigate (potential) visual and motor reselection in scenarios where external events do not require a response.

3.5.3 The Influence of Temporal Predictability (and the Lack Thereof)

The current task design additionally manipulated the temporal predictability of the external task, similar to the task design presented in Chapter 2.4. When aggregating trials across various external-task onsets (i.e., early, medium, and late) and comparing working-memory performance between the fixed and variable condition, I observed a significant effect of temporal predictability in RTs, but not in reproduction errors. One possible explanation is that performance in the working-memory task may have been at ceiling, thereby limiting the variability between conditions. Indeed, reproduction errors in the working-memory task were very low in external-task trials (11.89 ± 0.22 degrees [*M* ± *SEM*]) as compared to those in Chapter 2.4 (external-task trials: 17.89 ± 0.36 degrees). The improved accuracy, in contrast to Chapter 2.4, could be attributed to the introduction of a retro-cue during the working-memory delay. In particular, participants might have relied on the explicit top-down guidance of retro-cues rather than utilising implicitly learned temporal expectations. Moreover, in comparison to the online experiment

described in Chapter 2.4, participants underwent a more extensive practice session (i.e., until they reached an accuracy higher than 85% in the working-memory task and higher than 90% in the external task). Given that the practice blocks exclusively consisted of variable-onset trials, participants may have adapted to the temporal unpredictability of the external task, thus masking any potential differences in temporal predictability in the main experiment.

Regarding temporal-expectation effects on the external task, I found faster RTs for fixed versus variable onsets. Splitting the data by the three external-task onsets revealed that this effect was mostly driven by comparisons in early-onset trials which mirrors the external-task (i.e., interrupter) results of Chapter 2.4. Furthermore, the effect of faster RTs to early but not to later external tasks is in agreement with prior studies of temporal expectations in simple perception and action tasks and can be attributed to the hazard rate (Coull et al., 2000; Cravo et al., 2017; Miniussi et al., 1999; for a review see: Nobre, 2001; Nobre & van Ede, 2018, 2023).

3.5.4 Conclusion

In this chapter, I investigated the temporal dynamics of reactivations of internal representations following external-task engagement, capturing the natural flow of shifting attention from one domain to another. Utilising EEG, I was able to track when visual and motor contents resurface, providing insights into the reactivation process during external-to-internal shifts. Taken together, the findings of this chapter revealed that (1) the results of Chapter 2.4 partially replicated, (2) that the reactivation of visual and motor signatures occurred immediately following the external-task response, and (3) that variability in the onset of memory-guided behaviour was associated with corresponding changes in the latency of visual reselection after the external task. While this chapter offered critical observations into the natural progression of attention from the external to internal domain, the following two chapters will directly explore the cognitive and neural processes associated with shifts between domains (i.e., external-to-internal, internal-to-external) compared to shifts within a single domain (i.e., external-to-external, internal-to-internal).

4 Behavioural Correlates of Shifting Between Attentional Domains

4.1 Abstract

Most everyday tasks require shifting the focus of attention between sensory signals in the external environment and internal contents in working memory. To date, shifts of attention have been investigated within each domain, but shifts between the external and internal domain remain poorly understood. I developed a combined perception and working-memory task to investigate and compare the consequences of shifting attention within and between domains in service of a common orientation-reproduction task. Participants were sequentially cued to attend to items either in working memory or to an upcoming sensory stimulation. Stay trials provided a baseline condition, while shift trials required participants to shift their attention to another item within the same or different domain. Validating the experimental approach, I found evidence that participants shifted attention effectively in either domain (Experiment 1). In addition, I observed greater costs when transitioning attention between as compared to within domains (Experiments 1, 2, and 3). Strikingly, these costs persisted even when participants were given more time to complete the attentional shift (Experiment 2) and when the necessity to spatially shift attention was removed (Experiment 3). Biases in fixational gaze behaviour tracked attentional orienting in both domains but revealed no latency or magnitude difference for within- versus between-domain shifts (Experiment 1). The results suggest that shifting between attentional domains might be regulated by one or more unique control functions. The current findings open new ground for exploring the ubiquitous act of shifting attention between perception and working memory to guide adaptive behaviour in everyday cognition.

The findings presented in Chapter 4.3 and 4.4 have been published as Gresch et al. (2024b)

4.2 Introduction

Efficient goal-directed behaviour necessitates attention to be dynamically and selectively directed towards task-relevant contents available in the external world as well as those maintained in mind as internal representations (Narhi-Martinez et al., 2023; van Ede & Nobre, 2023). Accordingly, attention can be classified based on the domain in which it operates – either perception or memory (i.e., external attention or internal attention; Chun et al., 2011).

In laboratory tasks, pre-cues and retro-cues have emerged as valuable tools for studying attentional orienting mechanisms in relation to both perceptual and mnemonic contents, respectively. Orienting towards external information has been extensively investigated in tasks using predictive pre-cues that proactively provide information concerning a target-defining attribute prior to a target array (Posner, 1980). Conversely, tasks using retro-cues have proven effective for investigating internal orienting of attention to selective contents within working memory. Predictive retro-cues presented during the retention interval retroactively guide attention according to a target-defining attribute, such as its encoded location (Griffin & Nobre, 2003; Landman et al., 2003). These lines of research have provided important insights into the limits, mechanisms, and associated neural systems of attentional orienting to perception and working memory (Carrasco, 2011; Chun et al., 2011; Corbetta & Shulman, 2002; Gazzaley & Nobre, 2012; Kiyonaga & Egner, 2013; Lepsien & Nobre, 2006; Myers et al., 2017; Nobre et al., 2004; Panichello & Buschman, 2021; Souza & Oberauer, 2016; Wallis et al., 2015; Y. Zhou et al., 2022). For instance, pre-cue and retro-cue studies have consistently demonstrated that selectively attending to perceptual or mnemonic task-relevant information results in improved behavioural outcomes and enhanced neural processing for the prioritised information.

Yet, surprisingly, the focus of these previous works has been almost entirely on how we orient and reorient attention *within* either perception or working memory. Outside of the laboratory, we often shift attention dynamically *between* contents in the environment and in our mind. For instance, a cyclist entering a busy intersection needs to attend to sensory contents, such

as the changing traffic lights and pedestrians crossing the street, as well as focus on mnemonic contents, such as memorised directions. Thus, the cyclist's focus of attention must be in constant flux, rapidly transitioning between information available in the external environment and internal traces stored in memory. In this manner, adaptive behaviour within an ever-changing world calls for our ability to seamlessly shift attention between contents in perception and memory.

To date, a limited number of studies have investigated the flexible transitions of attentional selection between perception and memory. Existing work has emphasised that shifting between external and internal attentional domains may incur higher costs than reorienting within either domain. These studies have employed variations of task-switching paradigms, in which participants alternate between phases that require responding to sensory stimuli and phases that necessitate responding to encoded memoranda. Early behavioural studies used list-reproduction tasks requiring verbal or manual reports of digits or letter series which were either perceptually available, stored in memory, or required integration across both domains (Carlson et al., 1993; Dark, 1990; Weber et al., 1986). After a relative hiatus, the question of shifting attention between sensory and internally generated information has attracted renewed interest (Calzolari et al., 2022; Gilbert et al., 2005; Hautekiet et al., 2023; Honey et al., 2017; Poskanzer & Aly, 2023; Servais et al., 2023; Steel et al., 2024; Treder et al., 2021; Verschooren et al., 2020, 2021; Verschooren, Liefooghe, et al., 2019; Verschooren, Schindler, et al., 2019; Weilnhammer et al., 2023). For instance, researchers have replicated earlier work within the visual domain using tasks that require participants to shift between perception and memory on a trial-by-trial basis (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019). In these studies, the external trial required comparing stimuli that were both shown on the screen simultaneously, while the internal trial involved retrieving a stimulus from memory and comparing it to a target on the screen. This line of research takes an important step towards understanding attentional flexibility. However, to fully capture the dynamic nature of attentional shifting that occurs in our daily lives, we must go further.

Two central issues remain to be addressed. The first issue is the need to capture the everyday situation of smoothly shifting between the contents available in the external world and our minds to perform a single task, without necessarily being required to respond to each of these contents. Thus far, paradigms inspired by task-switching approaches have been used to calculate the costs of shifting between domains. For example, recent research has employed match-to-sample tasks in which participants responded either to externally presented or internally represented targets, with all other task requirements held constant across both attentional domains (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019). In these studies, participants responded according to either sensory information or memory representations on a per-trial basis. The cost of shifting between domains was revealed by reduced performance in subsequent trials of the alternate domain, as compared to those in the same domain. In contrast, in our daily lives, it is common to transition between external and internal attention in service of a unified goal-directed action. Returning to the example of the cyclist; taking a left turn at the traffic light involves multiple shifts between perception and memory (i.e., from externally attending to the traffic light, to internally attending to the memorised locations, to externally attending to the children crossing the street, etc.) before performing the single intended manual action (i.e., taking a turn).

The second important issue to address is how attention shifts between visual perception and visual working memory. Working memory is the core internal buffer for guiding adaptive behaviour (Funahashi et al., 1989; Fuster & Alexander, 1971; Myers et al., 2017; Nobre & Stokes, 2019; van Ede & Nobre, 2023). Yet, most previous studies have focussed on attentional shifting between external and other internal domains, such as imagination (Gilbert et al., 2005) or self-oriented cognition (Calzolari et al., 2022). Even when internal attention was operationalised as memory, it was often limited to the verbal modality (Carlson et al., 1993; Dark, 1990; Weber et al., 1986) or closely linked to long-term memory representations (Hautekiet et al., 2023;

Verschooren et al., 2020, 2021; Verschooren, Liefooghe, et al., 2019; Verschooren, Schindler, et al., 2019).

The present three experiments build on a promising foundation from the task-switching and the attentional (re-)orienting literature to investigate how we shift attention within and between contents in visual perception and visual working memory to guide performance towards a unified action. I developed a double-cueing task using both pre-cues and retro-cues, in which participants had to orient and then reorient the focus of attention to external and internal stimuli within trials. Importantly, participants were presented with the same stimulus material and response demands across both attentional domains. On a behavioural level, the task allowed for testing all four possible directions of shifts: external-to-external, internal-to-internal, internal-to-external, and external-to-internal. The first two engage within-domain shifts, while the latter two engage between-domain shifts.

In the first experiment, I recorded gaze position as an implicit continuous proxy of spatial attention shifts as they unfolded, to supplement accuracy and response-time measures obtained at the end of the trials. Recent research has shown that small directional eye movements track the orienting of attention to memorised visual locations (Draschkow et al., 2022; van Ede, Chekroud, & Nobre, 2019; van Ede et al., 2021), similar to orienting to external locations (Engbert & Kliegl, 2003; Hafed et al., 2011). By leveraging this sensitive and time-resolved metric, I could compare the dynamics of attentional orienting to visual sensory and working-memory locations and of shifting within and between domains. In the second experiment, I included multiple intervals between the final attention-orienting cue and the response probe to gain further evidence regarding the time courses of shifting attention within versus between domains. The third experiment examined a task scenario devoid of spatial attention shifts, yielding insights into the spatial dynamics involved in shifting between domains.

I predicted additional costs for shifting between as opposed to within domains in line with the previous literature (Calzolari et al., 2022; Carlson et al., 1993; Dark, 1990; Gilbert et al., 2005;

Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019; Weber et al., 1986). However, I made no directional prediction concerning whether shifts would be faster or incur greater costs when moving from one specific domain to the other (i.e., internal-to-external vs. external-to-internal). Previous work has highlighted that shifting from the internal to the external domain may incur less costs than vice versa (Calzolari et al., 2022; Dark, 1990; Hautekiet et al., 2023; Verschooren et al., 2020). However, given the differences in the current experimental design – specifically investigating attentional shifts towards a single response – I chose to remain neutral to the possible outcomes. Inclusion of gaze biases was intended to reveal any differences in temporal dynamics between shift types and to shed light on whether the timing of spatial shifts ultimately determined or impacted the final behavioural performance measure. I was mindful that performance differences could arise depending on whether participants responded to external or internal targets. Therefore, my analysis only compared responses to targets within a particular domain, albeit as a function of whether attention was initially oriented within that domain or to information in the alternate domain.

As a preview of the current results, Experiments 1, 2, and 3 demonstrate significantly higher response-time costs after shifting attention between as compared to within domains. Moreover, all experiments show increased accuracy costs for between- as compared to within-domain shifts when reporting an internal target, whereas no differences are detected when reporting an external target. Experiment 2 demonstrates that between-domain shift costs remain consistent even when additional time is provided to shift attention after the final cue. Moreover, Experiment 3 indicates that between-domain shifts do not differ whether or not spatial shifts of attention are required. Interestingly, the gaze results of Experiment 1 show similar timings for shifting attention within and between domains. This divergence between fixational gaze behaviour and RTs provokes consideration of additional bottlenecks that may occur when shifting attention between domains, beyond just the early attentional orienting phase.

4.3 Experiment 1: Costs of Shifting Between Attentional Domains

Experiment 1 launches my investigation into the dynamic shifts of attention within versus between contents in perception and working memory. The primary objective of this study was to investigate whether shifting attention between different domains results in a greater behavioural cost compared to shifts within the same domain. Additionally, this experiment allowed me to determine whether potential costs associated with between-domain shifts are particularly pronounced in one specific domain. Furthermore, by tracking fixational gaze behaviour the study aimed to assess whether spatial orienting processes are differentially modulated during within- versus between-domain attentional shifts.

4.3.1 Methods

4.3.1.1 Participants

The study was approved by the Central University Research Ethics Committee of the University of Oxford. The sample size was set to n = 25, based on previous working-memory studies with similar outcome measures (i.e., RTs, reproduction errors, gaze biases; van Ede et al., 2019, 2021). Additionally, I conducted a power analysis based on preliminary pilot data (n = 7). Employing the mixedpower package (Kumle et al., 2021) in RStudio (Posit team, 2022), I computed power simulations with 1000 iterations based on a model including both shift type and target domain as fixed effects. Given that the primary focus was on comparing different levels of shift type, I directed my attention towards this factor during the power-analysis interpretation. To account for the inherent uncertainty in effect-size estimates from pilot data, I systematically varied the mean differences between conditions in order to perform a more conservative power assessment. As illustrated in Supplementary Figure 8.12A, with a sample size of n = 25, a 50% smaller effect in RTs would still provide power greater than 80%.

To yield the targeted number of participants, I collected data from 32 participants in total. Four participants were excluded following an a-priori behavioural trial-removal procedure (see 4.3.1.4 Behavioural Analysis). In addition, three datasets were removed due to poor eye tracking (see 4.3.1.5 Eye-Tracking Analysis for further details). All remaining 25 participants (age range: 18 to 37; mean age: 25.48; gender: 12 female, 10 male, 3 non-binary; handedness: 21 right-handed, 4 left-handed) had normal or corrected-to-normal vision. Individuals provided written informed consent before participating in the study and were paid £15 per hour.

4.3.1.2 Apparatus and Data Acquisition

Participants sat in front of a monitor (27-inch Acer XB270H; resolution 1920×1080 pixels; refresh rate 100 Hz; screen width 60 cm) at a viewing distance of 100 cm. An eye tracker (EyeLink 1000, SR Research) was positioned on the table approximately 15 cm in front of the monitor. Before acquisition, I calibrated the eye tracker using a 5-point calibration and validation method. Gaze was continuously tracked binocularly at a sampling rate of 1000 Hz, and participants were required to maintain fixation throughout the trial. The experimental script was generated in Python using the PsychoPy package (version 2021.2.3; Peirce et al., 2019).

4.3.1.3 Task and Procedure

Participants performed a double-cueing task requiring the reproduction of the exact angle of one out of four tilted bars (**Figure 4.1A**). Two arrays containing laterally presented tilted bars appeared before and after two spatially informative cues. At the end of each trial, participants reproduced the orientation of the last item to be cued. Importantly, items were presented across both an internal working-memory array, encoded before the cues, and an external perceptual array, appearing after the cues. As such, cues indicating the location of a previously encoded item shifted attention internally (i.e., retro-cues, internal cues), instructing which item in working memory would need to be reported. In contrast, cues indicating a previously unoccupied location shifted

attention externally (i.e., pre-cues, external cues), instructing which item in the upcoming perceptual array would need to be reported.

The two main manipulations of this task were (1) that the first and second cue could either indicate the same or different items (i.e., stay vs. shift) and (2) that cued items could be upcoming sensory information or internal representations (i.e., external vs. internal). For shift trials, this meant that, at the time of the second cue, attention needed to be reoriented to either an item of the same domain (i.e., within-domain shift, Figure 4.1B, left panel) or an item in the other domain (i.e., between-domain shift, Figure 4.1B, right panel). This resulted in four possible shift conditions: external-to-external, internal-to-internal, internal-to-external, and external-to-internal, allowing me to compare the effect of shifting within and between attentional domains. In stay trials, the second-cued item was the same as the first-cued item. That is, both cues indicated either the same memory item or the same perceptual location. The second cue was 100% instructive, and thus, always indicated the target for later report. To ensure participants used the first cue, I made the first cue highly predictive for the later report by having disproportionally more stay trials (70%) relative to shift trials (30%). Figure 4.1C illustrates the ratio of conditions and number of trials across the experiment.

At the start of each trial, a white central fixation cross (RGB value: [255, 255, 255]) surrounded by four coloured circles (i.e., placeholders) was presented against a black background (RGB value: [0, 0, 0]) for 750 ms. The colours of the placeholders in each trial were always randomly drawn from a set of four highly distinguishable colours (RGB values: blue [0, 159, 183], orange [217, 120, 0], green [47, 167, 0], pink [247, 35, 255]. The colour-location mapping was randomly determined on each trial. Placeholders subtended 3.5° visual angle (~195 pixels) and were centred in the top-left, top-right, bottom-left, and bottom-right at a distance of 4.5° visual angle (~251 pixels) from fixation. Placeholders stayed on the screen until the response was required at the end of the trial.

Next, two tilted grey bars (RGB value: [50, 50, 50]) subtending approximately 0.5×3.4° (~28×190 pixels) appeared for 250 ms inside two of these placeholders (referred to as "internal display"). The orientation of the bars was randomly determined between 0 and 180°. Participants were instructed to memorise the orientation of these two bars. There were four possible configurations of the internal display, which occurred equally often throughout the experiment: the bars could appear in the placeholders located in the (1) top-left and top-right, (2) bottom-left and bottom-right, (3) top-left and bottom-right, and (4) bottom-left and top-right. As such, at the time of the internal display, there was always one bar presented in each hemifield.

The internal display was followed by a delay of 750 ms, in which only the fixation cross and the placeholders remained on the screen. The central fixation cross then changed colour to match that of one of the placeholders for 150 ms. The colour change acted as the first cue, indicating the most likely to-be-reported item. This first cue could indicate a placeholder that previously contained an item during the internal display (i.e., retro-cue) or a placeholder in which no item had been presented yet (i.e., an item of the upcoming external display; pre-cue). The cued location and attentional domain were pseudo-randomised such that each of the four placeholder locations and each of the two domains were equally often indicated by the cue.

Following another delay of 750 ms, a second cue was presented by another colour change of the central fixation cross for 150 ms. This cue instructed the target item for report with a certainty of 100%, thus overruling the first cue. Accordingly, the target item was always identical to the second-cued item. In the stay condition, the colour of the second cue was identical to the colour of the first cue, therefore indicating the same item. I included 70% of stay trials to encourage participants to use the first cue.

On 30% of the trials, the colour of the first and second cue differed, requiring a shift of attention. The shift condition was further divided into two sub-conditions: within-domain and between-domain shifts (**Figure 4.1B**). In within-domain shift trials (i.e., external-to-external or internal-to-internal), the second cue indicated the item appearing within the opposing hemifield

of the same domain as indicated by the first cue. In between-domain shift trials (i.e., internal-to-external or external-to-internal), the second cue indicated the item in the alternative domain occurring within the other hemifield. Thus, in within- and between-domain shift trials, attention always needed to be reoriented from one hemifield to the other. The second-cued location and domain were pseudo-randomised such that each of the four placeholder locations and each of the two domains were equally often indicated.

Immediately following the second cue, a second set of randomly tilted bars appeared in the perceptual display (referred to as "external display"). The new items appeared in the yet-unoccupied placeholders for 100 ms. For example, if the bars in the internal display appeared in the top-left and bottom-right placeholders, bars in the external display occurred the bottom-left and top-right placeholders. To make the perceptual discrimination challenging, items in the external array were masked. The masking array consisted of four overlayed titled bars presented in each of the placeholders for 100 ms. At each mask location the four tilted bars differed in 45° from the neighbouring bars. The overall orientation of each mask was randomly drawn at 0 ms and 50 ms, thus creating the impression of a dynamic display. The offset of the dynamic mask was followed by a 100-ms delay, after which a visual response dial was displayed, always starting in a random position. The response dial surrounding the central fixation cross was the same size as the placeholders and included markers along a circle that corresponded to the ends of a bar. The dial rotated leftwards when the F key was pressed and rightwards when the J key was pressed (either by holding the key down or pressing the key repeatedly) at a speed of 1/10° per millisecond. Participants reproduced the orientation of the last-cued item, with unlimited time to complete the orientation report. When the dial aligned with the orientation of the target item, participants pressed the space bar to confirm their response and continue with the task. The early onset of the external array, mask, and dial after the second cue was set to capture the costs of incomplete shifts of attention. The masking parameters were set based on piloting efforts to equate the difficulty of external and internal reports.

After a 100-ms delay in which only the fixation cross stayed on the screen, participants received feedback in the form of a number ranging from 0 to 100, with 100 indicating a perfect report and 0 indicating that the adjusted orientation was perpendicular to the angle of the target item. Feedback was presented 0.7° (~39 pixels) above the central fixation cross for 300 ms. Trials were separated by an ITI randomly drawn between 750 and 1000 ms. Between blocks participants were presented with their average reproduction accuracy in the previous block.

The experiment consisted of 640 trials divided across 16 blocks, each including 40 trials. Figure 4.1C displays the proportion of conditions across the experiment. Of the total trials, 70% (448) were stay trials while the remaining 30% (192) were shift trials. Participants were equally likely to report an external or internal item in stay trials (224 each). Shift trials varied in the type of shift performed, that is, 50% (96) required a within-domain shift and the other 50% (96) a between-domain shift. Within-domain and between-domain shift trials were further split by the target domain. Ergo, the within-domain shift trials included 50% (48) external-to-external and 50% (48) internal-to-internal shifts, while the between-domain shift trials included 50% (48) internal-to-external and 50% (48) external-to-internal shifts. Each block consisted of 70% (28) stay and 30% (12) shift trials.

To become familiarised with the procedure of the experiment, participants performed four practice blocks before the data collection started. Practice blocks had 10 trials each, consisting of 70% stay and 30% shift trials. The whole experiment lasted approximately 60 minutes.

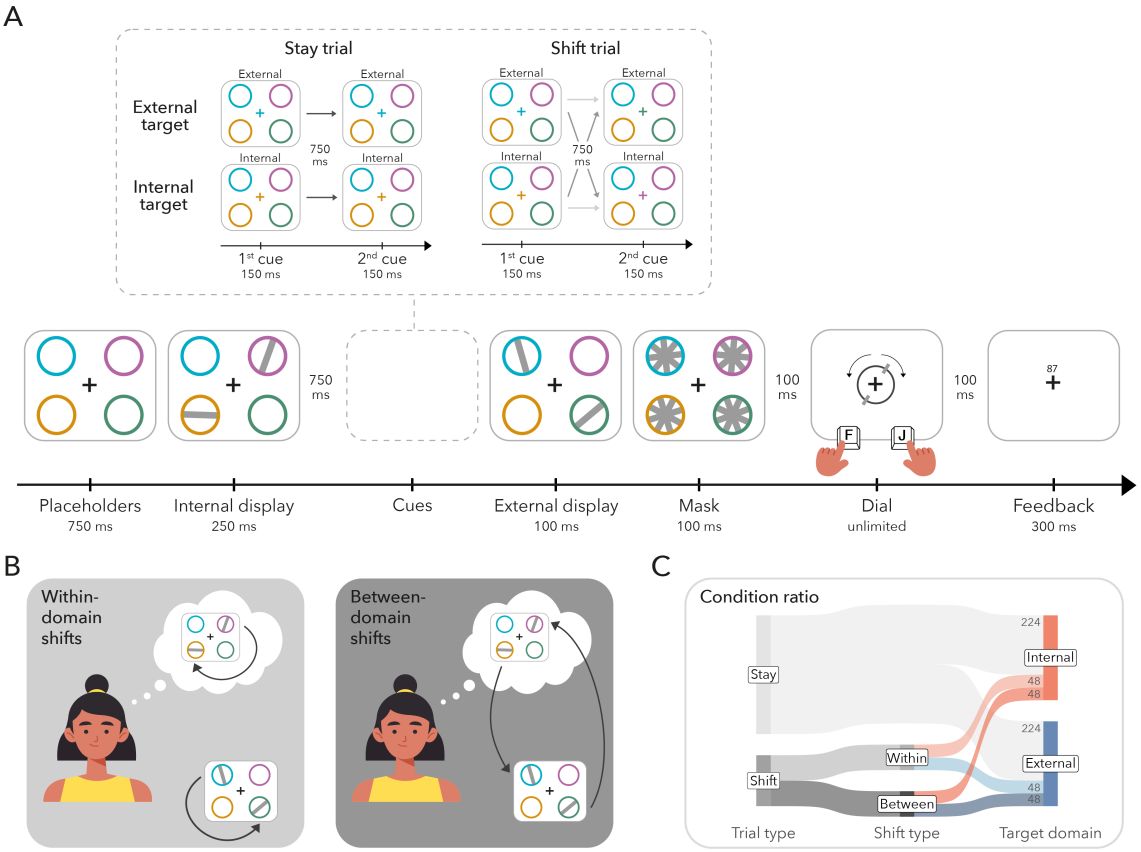

Figure 4.1. Task schematic and condition details of Experiment 1. (A) Illustration of stay and shift trials. In stay trials, the second cue indicated the same item as the first cue. In shift trials, the first and the second cue indicated different items. Light-grey arrows depict within-domain shifts and dark-grey arrows represent between-domain shifts. **(B)** In within-domain shift trials, participants either had to shift attention external-to-external or internal-to-internal. In between-domain shift trials, participants either had to shift attention external-to-internal or internal-to-external. **(C)** Ratio of conditions and number of trials across the experiment.

4.3.1.4 Behavioural Analysis

Data were analysed in R Studio (Posit team, 2022). During pre-processing (similar to Chapters 2 and 3), trials were removed when RTs exceeded 5000 ms. RTs refer to the interval between the appearance of the dial and the moment at which participants first started manipulating the dial. The choice of this RT measure aimed to minimise the variance related to participants' speed and strategy when finalising and confirming their precise responses, thereby providing a useful approximation of decision time (van Ede et al., 2017). Next, I removed trials for which the remaining RTs were 2.5 *SD* above the individual mean across all conditions. Datasets with more than 10% of trials rejected during these pre-processing steps were removed from further analyses (n = 1). Moreover, I removed datasets with average reproduction errors (calculated by averaging

the absolute difference between the angle of the target item and the reported angle) equal or higher than 45° in any of the conditions (n = 3). To analyse the same datasets in the behavioural and eye-tracking analysis, I removed an additional three participants due to bad eye-tracking data (see 4.3.1.5 Eye-Tracking Analysis). After these exclusion steps, datasets from 25 participants remained in the behavioural analysis, retaining 96.88% (SD = 0.69) of trials on average.

The conditions of interest were target domain (external vs. internal), trial type (stay vs. shift), and shift type (within vs. between). When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples t-test and report Cohen's d as a measure of effect size. For post-hoc t-tests, I report Bonferroni-corrected p-values that I denote as $p_{Bonferroni}$. In the case of significant interactions, I exclusively employed post-hoc t-tests to identify differences between the levels of trial type or shift type within each target domain, refraining from extending these comparisons across target domains. This approach was chosen to handle potential baseline differences between external- and internal-report trials. I used the ggplot2 package (version 3.3.6; Wickham, 2009) for plotting results. The within-subject standard error of the mean was calculated from normalised data using the approach from Morey (2008).

4.3.1.5 Eye-Tracking Analysis

Pre-processing and analysis procedures followed previous work (Draschkow et al., 2022; van Ede, Chekroud, & Nobre, 2019; van Ede et al., 2021). Eye-tracking data were first converted from edf to asc format, and subsequently read into R Studio (Posit team, 2022). Eye blinks were detected and interpolated ±100ms around identified blinks. Data from the left and right eyes were averaged, yielding a single time course of horizontal gaze position (x-position) and vertical gaze position (y-position). Next, data were downsampled to 250 Hz (i.e., one sample every 4 ms), then epoched from 200 ms before to 1400 ms after the relative onsets of the first and second cue. To focus on biases in fixational gaze behaviour, trials were removed in which vertical or horizontal

gaze positions exceeded ±2.25° (~126 pixels) from fixation (i.e., further than half the distance to either of the four placeholders, as in Draschkow et al., 2022; van Ede et al., 2021). I smoothed all time-course data over 15 samples (i.e., 60-ms smoothing window average) and I applied baseline correction trial-wise data. For the first-cue epochs, I subtracted the mean value between-200 and 0 ms before the onset of the first cue from all subsequent time points. Similarly, for the second-cue epochs, I subtracted the mean value between -200 and 0 ms before the onset of the second cue from all following time points in the trial. Eye-tracking data from 25 participants remained in the gaze analysis, retaining 93.64% (*SD* = 5.89) of trials on average.

For simplicity, only lateral gaze biases were analysed. To compare gaze biases across trials, I first compared leftwards and rightwards gaze-position time courses within each domain relative to the first-cued item. In addition, leftwards and rightwards gaze-position time courses were also compared across shift types, within each target domain and relative to the second cue. Time-course data were plotted in a 1600-ms window beginning 200 ms prior to the relative cue (i.e., first or second cue). To increase sensitivity and interpretability, I calculated the mean of the lateral gaze bias according to the location of the cued item (i.e., towardness; as in Draschkow et al., 2022; van Ede et al., 2019, 2021). Towardness for each time step was given by the trial-average horizontal position in right-item trials minus the trial-average horizontal position in left-item trials (where position values left of fixation were negative) divided by two. A positive towardness indicated a gaze bias in the direction of the cued item. I used the towardness variable to determine the significance of gaze biases compared to zero, as well as to test for differences between conditions. I ran cluster-based permutation testing with the permuco package (Frossard & Renaud, 2021), using cluster-mass permutation tests with 5000 permutations, which controls for multiple comparisons across time points while retaining high sensitivity.

Moreover, I computed cross-correlations to test for potential latency differences between towardness time courses. Calculating cross-correlations involves shifting one signal relative to the other and computing the correlation coefficient of the two signals at each shift (i.e., lag). The

resulting cross-correlation function indicates how well the two signals match at each shift, that is, a high cross-correlation value at a specific time lag indicates a strong relationship between the two signals at that time lag. To accurately determine the lag at which the cross-correlation reached its maximum, I employed two methods. Firstly, I identified individual peaks in the cross-correlation. Additionally, I utilised a jack-knife approach to enhance the analysis. For the jack-knife approach, I computed cross-correlation averages by excluding each of the *n* datasets once, allowing me to obtain *n* leave-one-out grand averages. Next, I identified the lags at which the cross-correlation coefficient reached its maximum. When possible, jack-knifed standard errors, *t*-values, and *p*-values were adjusted according to J. Miller et al. (1998).

For visualisation purposes, I constructed heat maps of gaze positions between -400 and 1400 ms relative to the first cue. For these analyses, I did not remove trials with gaze values exceeding more than half the distance of the placeholder centres. Two-dimensional kernel density estimations were obtained at a 125×125-pixel spacing on a 500×500-pixel square grid separately for cues indicating right and left placeholders. Lastly, I subtracted density maps of left- from right-cued items which allowed me to illustrate the magnitude of the observed gaze bias.

4.3.2 Results

4.3.2.1 RTs and Errors: Staying vs. Shifting

As I developed a task for quantifying the costs of unfolding shifts of attention, the initial analyses focussed on confirming the success of the experimental design in this regard. Specifically, I tested whether participants shifted their attention according to the initial cue before reorienting according to the second cue. Since the second cue was always 100% instructive for the later report, it was necessary to confirm that participants also engaged with the first cue. If participants made use of the highly predictable first cue, performance should be better in stay trials wherein the cued information was repeated across the first and second cue and therefore required no subsequent shift of attention.

I evaluated RTs (i.e., interval between dial onset and initial dialling response) and compared stay versus shift trials for both external- and internal-target reports (**Figure 4.2A**). I observed a significant main effect of trial type ($F_{(1,24)}$ = 66.840, p < 0.001, η^2_G = 0.130). Trials that required participants to shift their attention to a different item based on the second cue resulted in slower RTs compared to trials where the second cue was redundant, allowing participants to maintain their attention on the same item. This suggests that participants relied on the information provided by the first cue. I also found a significant a main effect of target domain, indicating faster reports for internal compared to external targets ($F_{(1,24)}$ = 8.104, p = 0.009, η^2_G = 0.005). The interaction between trial type and target domain was not significant ($F_{(1,24)}$ = 1.516, p = 0.230, η^2_G < 0.001). I observed the same statistical pattern when analysing response-completion times, that is, RTs based on the interval between the appearance of the dial and participants' confirmation of their responses by pressing space bar (**Supplementary Figure 8.13B**).

I also considered the influence of trial type and target domain on reproduction errors (**Figure 4.2B**). Errors were systematically larger when participants were required to shift attention ($F_{(1,24)}$ = 147.301, p < 0.001, η^2_G = 0.332), again consistent with an initial use of the first cue. This was true for all participants in this experiment (**Figure 4.2B**, right panel), as it was also the case for RTs (**Figure 4.2A**, right panel). There was also a significant effect of target domain ($F_{(1,24)}$ = 8.590, p = 0.007, η^2_G = 0.071), whereby internal items were reported with higher errors than external items. In contrast to RTs, trial type and target domain interacted for errors ($F_{(1,24)}$ = 6.997, p = 0.014, η^2_G = 0.015). Pairwise comparisons revealed that the difference between stay and shift trials was more pronounced when reporting the orientation of internal targets as compared to external targets ($t_{(24)}$ = 2.645, $p_{Bonferroni}$ = 0.043, d = 0.529). However, shift costs for errors were significant for both external-target ($t_{(24)}$ = 7.742, $p_{Bonferroni}$ < 0.001, d = 1.548) and internal-target trials ($t_{(24)}$ = 11.169, $p_{Bonferroni}$ < 0.001, d = 2.234).

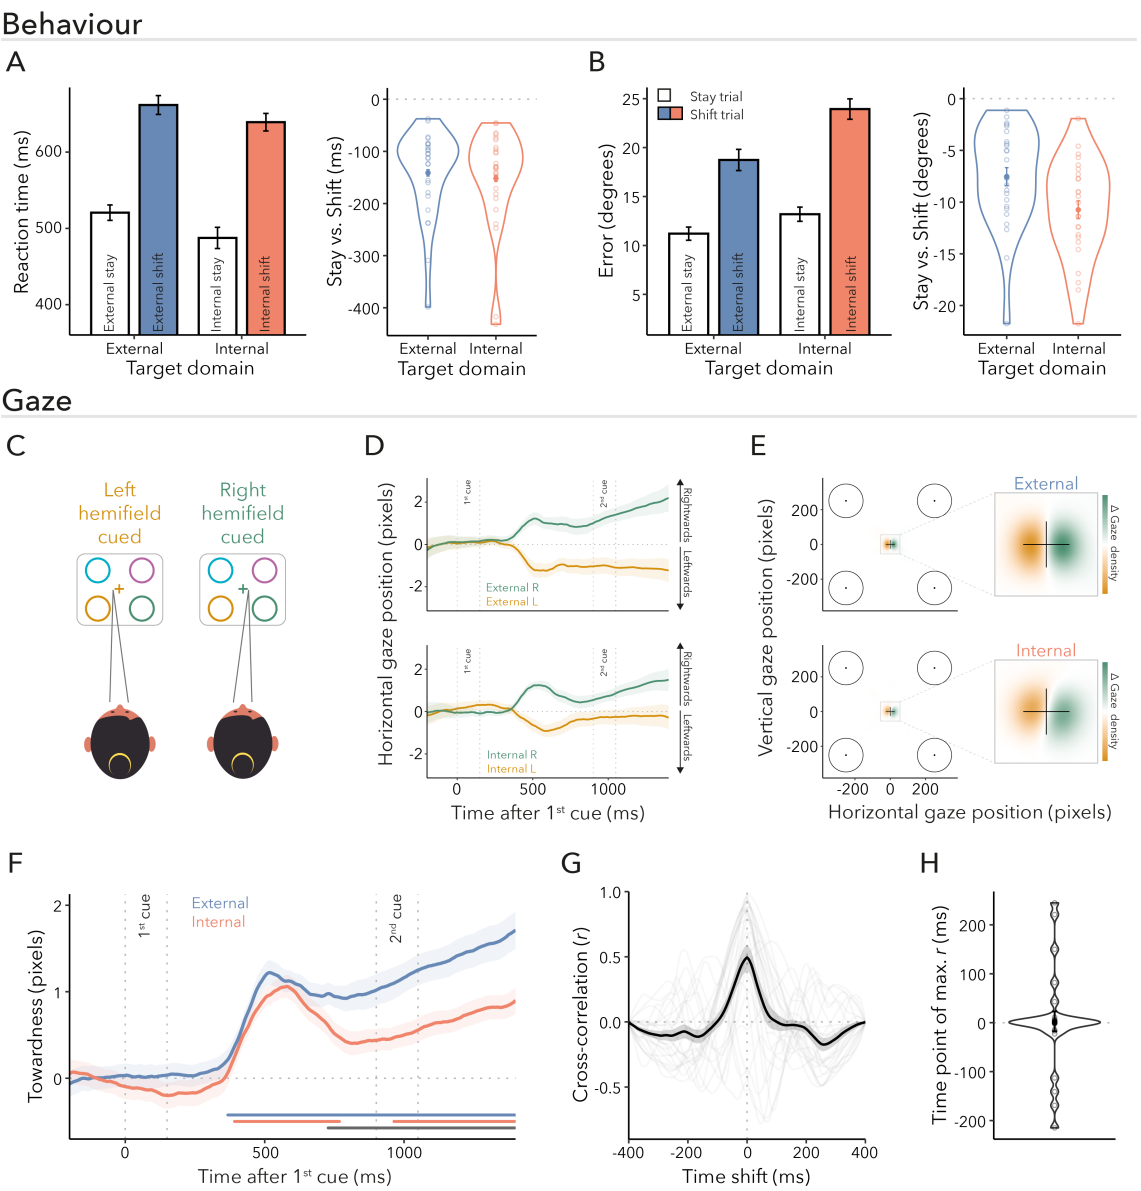

Figure 4.2. Behavioural and gaze effects of orienting attention. (A) RTs were slower for shift than for stay trials and for external- than for internal-target trials. **(B)** Reproduction errors were higher for shift than for stay trials, whereby the effect was larger for internal- than for external-target trials. **(C)** Schematic depicting predicted gaze biases following left and right cues. **(D)** Average gaze position (in the horizontal plane) after the onset of an external (top panel) and internal (bottom panel) first cue as a function of time and the location of the first-cued item. **(E)** Difference in gaze density (Δ) between right- and left-cued items (right minus left) for external (top panel) and internal (bottom panel) first-cue trials. Plots depict the difference in gaze density between-400 and 1400 ms relative to the first cue. **(F)** Gaze bias expressed as towardness for external and internal first-cue trials. Horizontal lines denote significant clusters (grey line denotes external vs. internal). **(G)** Cross-correlation coefficients of the external and internal towardness time courses. **(H)** Average lag at which the correlation-coefficient reached its maximum. Error bars and shaded areas indicate ±1 *SEM*. Dots represent individual participants.

4.3.2.2 Gaze: Staying vs. Shifting

To obtain further verification that participants utilised the first cue, I next focussed the analysis on fixational gaze behaviour relative to the first-cue onset (**Figure 4.2C**). **Figure 4.2D** depicts the

average horizontal gaze position following external and internal cues (top and bottom panels, respectively) for trials with first-cued items to the left (orange) and right (green) of fixation. I observed noticeable gaze shifts towards the direction of the cued location; that is, when the first-cued item occupied a left (right) position on the screen, gaze became biased to the left (right). Density maps of gaze positions contrasting right- and left-cued trials show that this bias was primarily composed of small displacements towards the location of the cued content (Figure 4.2E; top panel: external cue, bottom panel: internal cue), consistent with previously demonstrated biases in fixational gaze behaviour (Draschkow et al., 2022; Engbert & Kliegl, 2003; Hafed et al., 2011; van Ede, Chekroud, & Nobre, 2019; van Ede et al., 2021).

To quantify these lateral gaze biases, I combined the right- and left-cued gaze time courses into a single towardness metric for each domain (Figure 4.2F). Cluster-based permutation analyses demonstrated a large significant cluster for external (blue horizontal line; cluster $p < 0.001$) and two significant clusters for internal first-cue trials (red horizontal lines; 1st cluster $p = 0.006$; 2nd cluster $p = 0.020$), indicating that both domains biased gaze towards the cued location. This result complements the behavioural findings and can be taken as additional evidence that participants utilised the first cue when performing the task. The towardness metric additionally allowed to compare the time course of attentional shifts between the two domains directly. Interestingly, when comparing gaze biases between the external and internal domains, I observed a significant cluster only after the initial peak that was common across domains (Figure 4.2F; grey horizontal line; cluster $p = 0.014$), suggesting that external attention evoked a more sustained bias than internal attention. The towardness time courses of all conditions are illustrated in Supplementary Figure 8.14.

The onset timing for external and internal gaze biases appeared similar by visual inspection. Since statements regarding the exact onset of an effect are not supported by permutation tests (Sassenhagen & Draschkow, 2019), I calculated cross-correlations to characterise the temporal relationship between external- and internal-orienting mechanisms

(Figure 4.2G). A high cross-correlation coefficient indicates a strong similarity between two signals at a particular time lag, while a low coefficient indicates dissimilarity. In other words, it determines how much one signal lags or leads another signal at various points in time. Figure 4.2H shows the time lags of the maximal cross-correlations when comparing first-cue trials between domains. Notably, there was no temporal shift in gaze compared to lag zero ($t_{(24)}$ = 0.140, p = 0.890, d = 0.028), suggesting that there was likely no latency difference in gaze when orienting attention externally versus internally. To quantify the observed cross-correlations more formally, I also applied a jack-knife approach which yielded the same result (M = 0; *SEM*, *t*-, *p*-value not determinable due to lack of variance).

These findings indicate that the latencies of gaze biases were comparable for external and internal first-cued items. However, I observed that external compared to internal orienting elicited a more sustained gaze bias. Most important for the current purpose is that, collectively, these results suggest that participants relied on the first cue to successfully complete the task.

4.3.2.3 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts

Having established that participants utilised the first cue, I next turned to the behavioural dynamics of shifting attention within versus between domains (Figure 4.3A and B). Thus, the following analyses focus on shift trials, in which participants were required to reorient attention within or between domains. If shifting attention between was more costly than shifting within domains, performance should be slower and worse for internal-to-external than external-to-external and for external-to-internal than internal-to-internal trials. I tested this using a repeated-measures ANOVA with shift type (within vs. between) and target domain (external vs. internal) as factors. RTs were slower for between- compared to within-domain shifts (Figure 4.3A; $F_{(1,24)}$ = 16.388, p < 0.001, η^2_G = 0.005) and for reporting external compared to internal targets ($F_{(1,24)}$ = 4.947, p = 0.036, η^2_G = 0.003). The interaction between target domain and shift type was not

significant ($F_{(1,24)}$ = 0.003, p = 0.957, η^2_G < 0.001). The same statistical pattern was observed when analysing response-completion times instead of RTs (**Supplementary Figure 8.13D**).

As depicted in **Figure 4.3B**, reproduction errors were higher when participants had to shift between domains ($F_{(1,24)}$ = 29.859, p < 0.001, η^2_G = 0.023). Again, I found a main effect of target domain ($F_{(1,24)}$ = 9.660, p = 0.005, η^2_G = 0.093), with errors being higher for internal than external targets. Unlike for RTs, for errors there was also a significant interaction between target domain and shift type ($F_{(1,24)}$ = 25.040, p < 0.001, η^2_G = 0.014), suggesting that the shift type had a different effect on errors depending on the target domain. The between-domain shift cost was higher for reporting internal targets compared to external targets, that is, external-to-internal versus internal-to-internal shifts were less accurate than internal-to-external versus external-to-external shifts ($t_{(24)}$ = 5.004, $p_{Bonferroni}$ < 0.001, d = 1.001). Additional pairwise post-hoc comparisons revealed that there was a significant difference between internal-to-internal and external-to-internal shifts ($t_{(24)}$ = 7.740, $p_{Bonferroni}$ < 0.001, d = 1.548). However, the difference between external-to-external and internal-to-external shifts was not significant ($t_{(24)}$ = 0.936, $p_{Bonferroni}$ = 1.000, d = 0.187).

In sum, I found evidence for higher performance costs in RTs when shifting attention between as opposed to within domains. For accuracy, there was only a cost of shifting between domains when the final target was selected from the internal attentional domain. A consistent result pattern emerged when utilising a percentage-change measure to normalise the between-domain shift-cost effect across each target domain (i.e., removing any influence of between-domain baseline performance levels; **Supplementary Figure 8.15**).

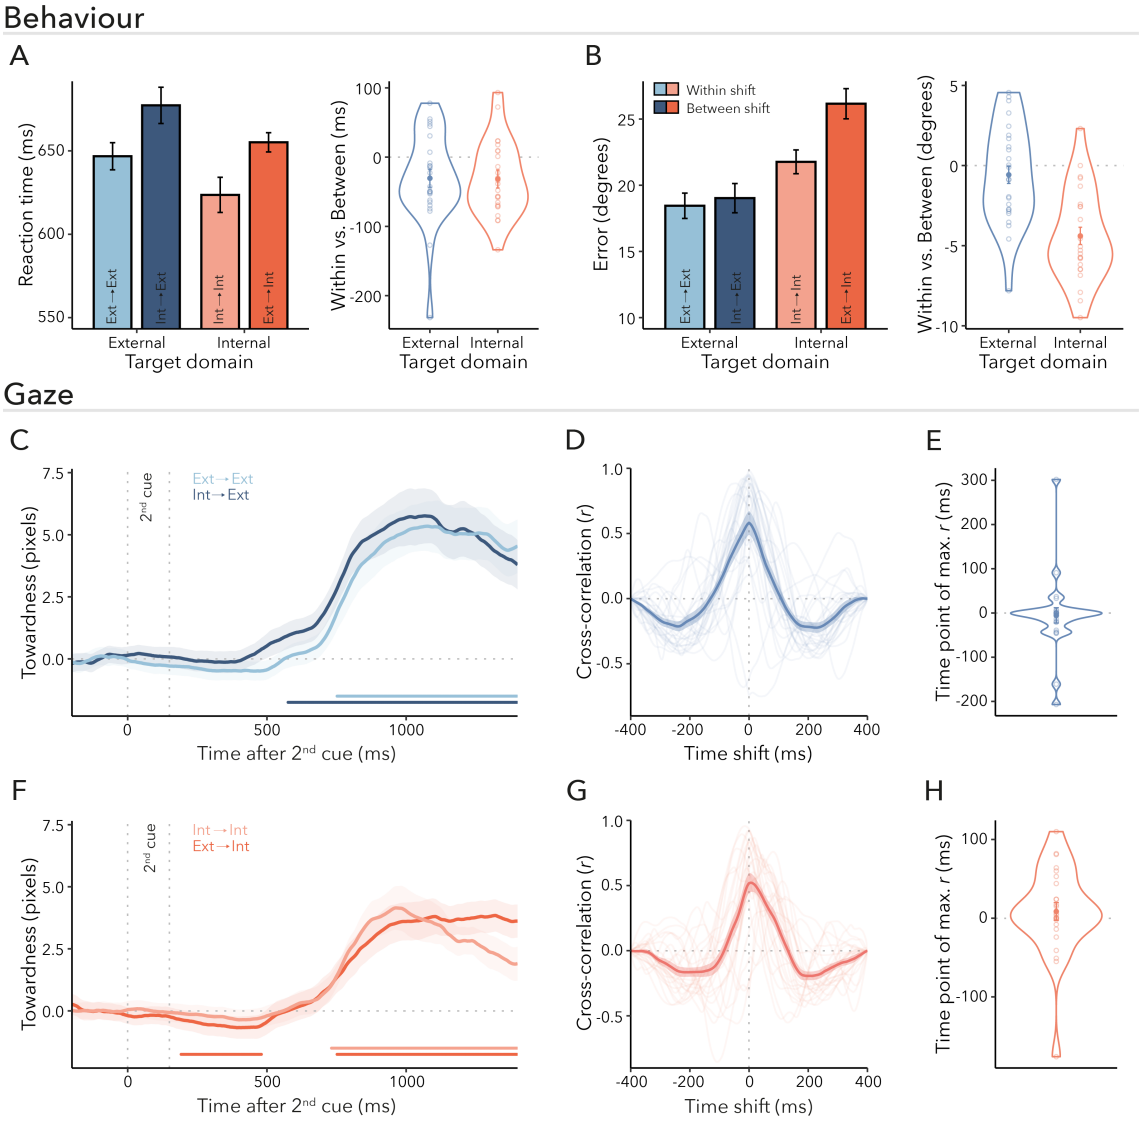

Figure 4.3. Between-domain shift costs for behaviour but not for gaze biases. (A) RTs were slower when shifting between as compared to within domains and for external- as compared to internal-target reports. **(B)** The effect of shift type on errors depended on the level of target domain. There was only a between-domain shift cost when reporting internal but not external targets. **(C)** Gaze bias expressed as towardness for external within- and between-domain shift trials. Horizontal lines denote significant clusters. **(D)** Cross-correlation coefficients for the within- and between-shift time courses, evaluated for external-target trials. **(E)** Average time point (i.e., lag) at which the correlation-coefficient reached its maximum for external-target trials. **(F)** Same as (C) but for internal-target trials. **(G)** Same as (D) but for internal-target trials. **(H)** Same as (E) but for internal-target trials. Error bars and shaded areas indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

4.3.2.4 Gaze: Within-Domain Shifts vs. Between-Domain Shifts

After having observed a higher performance cost for between-domain shifts compared to within-domain shifts, I asked whether this cost was also reflected in the gaze bias. Specifically, I aimed to investigate whether there was evidence supporting the notion that shifting attention between domains would result in delayed gaze biases as compared to shifting within domains. This effect

could potentially indicate the presence of an additional obstacle or boundary to overcome during the attentional shifting process, aligning with the observed impact on RTs. I again focussed on shift trials.

First, I computed the aggregated measure of towardness and analysed the time of within- and between-domain shifts following the second-cue onset for external-target (**Figure 4.3C**) and internal-target trials (**Figure 4.3F**). After attention was prompted to shift between contents, late clusters emerged for all the four shift types (external-to-external: cluster $p < 0.001$; internal-to-internal: cluster $p < 0.001$; internal-to-external: cluster $p < 0.001$; external-to-internal: cluster $p < 0.001$; light blue, light red, dark blue, dark red horizontal lines, respectively), indicating the second cue biased gaze towards the direction of the target item. In addition, I observed a gaze bias in the direction of the first-cued item for external-to-internal trials (dark red horizontal line; cluster $p = 0.047$). No clusters were found when comparing within- versus between-domain shift trials within either domain.

Besides comparing the magnitude of gaze biases, I was particularly interested in the latencies of these biases following external and internal within- and between-domain shifts. **Figure 4.3D** and **G** illustrate the cross-correlations when examining within-domain and between-domain shifts for each respective domain. There was no evidence for a temporal shift in gaze compared to lag zero for either external- (**Figure 4.3E**; $t_{(24)} = -0.308$, $p = 0.761$; $d = 0.062$) or internal-target trials (**Figure 4.3H** *;* $t_{(24)} = 0.756$, $p = 0.457$; $d = 0.151$). These results were replicated using a jack-knife approach (external-target trials: $M = 0$, *SE*, *t*-, *p*-value not determinable due to lack of variance; internal-target trials: $M = 6.120$, $SEM = 8.465$, $t_{(24)} = 0.723$, $p = 0.477$). Therefore, these findings did not reveal a significant delay in gaze biases when shifting attention between domains relative to within domains, despite the observed slowing of behavioural reporting following between-domain shifts.

4.3.3 Interim Summary

Experiment 1 led to several important findings. First, the task design was effective at eliciting attention shifts to external and internal contents as evidenced by behavioural and gaze-bias findings. Second, shifting attention between domains was associated with a higher behavioural cost than shifting attention within domains. Third, the between-domain shift cost was asymmetric for reproduction errors; within- versus between-domain shifts differed for internal- but not for external-target reports. And lastly, in contrast to performance measures, gaze-position measures were insensitive to whether shifts occurred within or between domains. The divergence between behavioural results and the gaze bias suggests that the cost associated with shifting attention between domains is not primarily driven by the spatial orienting functions that the gaze bias likely reflects.

4.4 Experiment 2: Temporal Dynamics of Shifting Between Attentional Domains

The second experiment was designed to uncover putative differences in the temporal dynamics of shifting attention within versus between domains by systematically varying the duration between the second cue and the external display. This allowed me to disambiguate potential explanations of the observed between-domain shift-cost effects. For example, specific patterns of performance costs could be accounted for by an extended period being necessary to complete attentional orienting, selection, prioritisation, and/or response gating, when shifting between domains, or from one domain to another. If this were the case, providing additional time to shift attention after the second cue should mitigate any increased or asymmetric costs. Furthermore, the presence of the asymmetric shift cost in reproduction errors might have been contingent on participants having minimal time to recall the internal target after the second cue. A similar memory-retrieval account of between-domain shifts has been proposed by Dark (1990).

In addition to manipulating the intervals after the second cue, Experiment 2 also encouraged the reliance on cues in a different and more economical manner. Rather than including a large proportion of stay trials, the design included one third of trials in which the external display appeared after the first cue, prompting a response to the initially cued item. The different approach was also aimed at testing the robustness of the pattern of results from Experiment 1.

4.4.1 Methods

4.4.1.1 Participants

Again, all experimental procedures were approved by the Central University Research Ethics Committee of the University of Oxford. Participant numbers were derived from applying the same power analysis I used previously, however this time to the full dataset of Experiment 1 (n = 25). The power analysis indicated that a sample size of n = 25 would yield power exceeding the 80% threshold even if the RT effect size was 30% smaller (Supplementary Figure 8.12B).

I collected data from 26 participants. One participant was excluded following an a-priori behavioural trial-removal procedure (see 4.4.1.4 Analysis). All remaining 25 participants (age range: 19 to 39; mean age: 26.92; gender: 16 female, 7 male, 2 non-binary; handedness: 23 right-handed, 2 left-handed) had normal or corrected-to-normal vision. Individuals provided written informed consent before participating in the study and were paid £15 per hour.

4.4.1.2 Apparatus and Data Acquisition

Participants completed the experiment in a group testing room with a capacity of 20 people, yet the testing sessions were conducted with no more than six participants simultaneously. Participants sat in front of a monitor (23.6-inch DELL E2417H; resolution 1920×1080 pixels; refresh rate 60 Hz; screen width 53 cm) at a viewing distance of 80 cm. The experimental script was generated in Python using the PsychoPy package (version 2021.2.3; Peirce et al., 2019).

4.4.1.3 Task and Procedure

Experiment 2 closely adhered the procedures from Experiment 1, with three notable deviations (Figure 4.4A and B). Firstly, to reveal the temporal dynamics of the between-domain shift cost, the inter-stimulus-interval (ISI) between the second cue and the external display was manipulated in shift trials. The ISI could either be 0 ms (same as Experiment 1), 250 ms, or 750 ms. Secondly, the stay condition contained only a single cue (i.e., single-cue condition). The cue in stay trials appeared at the time of the first cue in shift trials and was then followed by an ISI of 750 ms before the onset of the external display. Participants were instructed to always report the last-cued item. That is, in shift trials, participants reported the second-cued item (same as Experiment 1). However, in stay trials, participants reported the item indicated by the single cue. Lastly, the proportion of stay and shift trials differed in comparison to Experiment 1 (Figure 4.4C).

The experiment consisted of 720 trials divided across 20 blocks, each including 36 trials. Of the total trials, ⅓ (240) were stay trials. The remaining ⅔ (480) were shift trials. Participants were equally likely to report an external or internal item in both stay trials (120 each) and shift trials (240 each). Within shift trials, within- and between-domain shifts were equiprobable (240 each). Thus, the four shift conditions each included 120 trials, which were further divided into 40 trials of each equally possible ISIs. Each block consisted of ⅓ (12) stay and ⅔ (24) shift trials. All other task variables remained unchanged between Experiment 1 and Experiment 2.

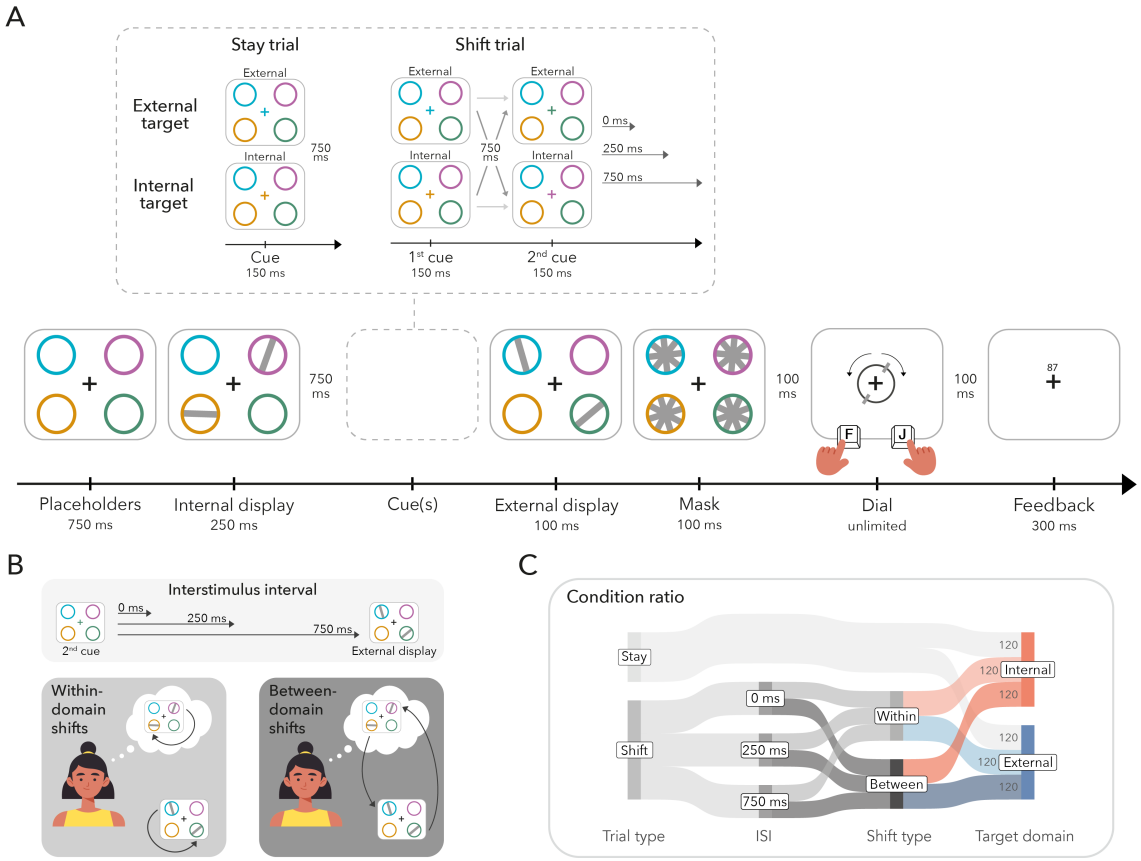

Figure 4.4. Task schematic and condition details of Experiment 2. (A) Illustration of stay and shift trials. In stay trials, only a single cue was presented. In shift trials, the presentation of the second cue was followed by an inter-stimulus interval of either 0 ms, 250 ms, or 750 ms. Light-grey arrows depict within-domain shifts and dark-grey arrows represent between-domain shifts. **(B)** In within-domain shift trials, participants either had to shift attention external-to-external or internal-to-internal. In between-domain shift trials, participants either had to shift attention external-to-internal or internal-to-external. Additionally, I manipulated the inter-stimulus interval between the second cue and the external display in shift trials. **(C)** Ratio of conditions and number of trials across the experiment.

4.4.1.4 Analysis

Data were first pre-processed according to the pipeline matching that of Experiment 1. During pre-processing, trials were removed when RTs exceeded 5000 ms. Next, I removed trials for which the remaining RTs were 2.5 *SD* above the individual mean across all conditions. None of the participants was removed for loss of more than 10% of trials. One dataset was removed due to average reproduction errors exceeding 45° in a subset of conditions. After these exclusion steps, datasets from 25 participants remained in the main behavioural analysis, with an average of 97.01% (*SD* = 0.55) of trials retained. The conditions of interest were target domain (external vs.

internal), and shift type (within vs. between), and ISI (0 ms vs. 250 ms vs. 750 ms). As in Experiment 1, the dependent variables of interest were RTs and reproduction errors.

While trial numbers were too low in Experiment 1, Experiment 2 allowed me to assess the proportion of trials in which participants incorrectly reported the angle of the first-cued item (i.e., swaps, non-target reports) as well as the proportion of trials in which random guessing occurred. For each participant and condition, I fitted a probabilistic mixture model and estimated the relative proportions of the non-target (i.e., first-cued item) distribution as well as the guessing distribution (Bays et al., 2009). The mixture modelling analysis was performed using the mixtur package (version 1.2.1; Grange & Moore, 2022).

When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples *t*-test and report Cohen's *d* as a measure of effect size. For post-hoc *t*-tests, I report Bonferroni-corrected *p*-values that I denote as $p_{Bonferroni}$.

4.4.2 Results

4.4.2.1 RTs and Errors: 0 ms vs. 250 ms vs. 750 ms

The following analyses focussed on the shift trials in which participants reoriented attention within or between domains. I used a repeated-measures ANOVA with shift type (within vs. between), target domain (external vs. internal), and ISI (0 ms vs. 250 ms vs. 750 ms) as factors. As in Experiment 1, RTs were slower for between- compared to within-domain shifts (**Figure 4.5A**; $F_{(1,24)}$ = 29.124, $p < 0.001$, $\eta^2_G = 0.004$) and for reporting external compared to internal targets ($F_{(1,24)}$ = 7.635, $p = 0.011$, $\eta^2_G = 0.012$). Moreover, there was a significant effect of ISI (**Figure 4.5C**; $F_{(2,48)}$ = 59.565, $p < 0.001$, $\eta^2_G = 0.150$). Post-hoc comparisons revealed that RTs became faster with increasing ISI; participants responded faster in trials with a 250-ms ISI as compared to a 0-ms ISI ($t_{(24)} = 9.891$, $p_{Bonferroni} < 0.001$, $d = 1.978$), and faster still in trials with a 750-ms ISI as compared to a 250-ms ISI $t_{(24)} = 3.244$, $p_{Bonferroni} = 0.010$, $d = 0.649$). Whether shift type interacted with ISI

was a central question in Experiment 2. No interactions involving ISI were observed and no other interaction effect reached significance (see **Supplementary Table 8.9A** for the full set of interferential statistics).

Reproduction errors also showed a similar pattern as in Experiment 1. Participants made more errors when they shifted attention between compared to within domains (**Figure 4.5B**; $F_{(1,24)}$ = 11.019, p = 0.003, η^2_G = 0.015). Again, I found a main effect of target domain ($F_{(1,24)}$ = 34.158, p < 0.001, η^2_G = 0.186), with errors being higher for internal than external targets. Additionally, I observed a main effect of ISI (**Figure 4.5D**; $F_{(2,48)}$ = 82.596, p < 0.001, η^2_G = 0.306). Errors became smaller with increasing ISI. Errors were smaller at an ISI of 750 ms than at an ISI of 250 ms ($t_{(24)}$ = 3.201, $p_{Bonferroni}$ = 0.011, d = 0.640), which, in turn, were smaller than at an ISI of 0 ms ($t_{(24)}$ = 9.162, $p_{Bonferroni}$ < 0.001, d = 1.832).

As in Experiment 1, there was also a significant interaction between target domain and shift type ($F_{(1,24)}$ = 10.282, p = 0.004, η^2_G = 0.012). The between-domain shift cost was higher for reporting internal targets compared to external targets, that is, external-to-internal versus internal-to-internal shifts were less accurate compared to internal-to-external versus external-to-external shifts ($t_{(24)}$ = 3.162, $p_{Bonferroni}$ = 0.013, d = 0.632). Additional pairwise post-hoc comparisons revealed a significant difference between internal-to-internal and external-to-internal shifts ($t_{(24)}$ = 4.010, $p_{Bonferroni}$ = 0.002, d = 0.802). However, the difference between external-to-external and internal-to-external shifts was not significant ($t_{(24)}$ = 0.303, $p_{Bonferroni}$ = 1.000, d = 0.061). The pattern of effects, therefore, matched that in Experiment 1.

The critical interaction between shift type and ISI did not reach significance and there was also no three-way interaction (**Supplementary Table 8.9B** for the full set of interferential statistics). I did however observe an interaction between ISI and target domain ($F_{(2,48)}$ = 12.726, p < 0.001, η^2_G = 0.077). Errors were significantly higher for internal as compared to external targets at the 250-ms ($t_{(24)}$ = 7.943, $p_{Bonferroni}$ < 0.001, d = 1.589) and 750-ms ISI ($t_{(24)}$ = 11.566, $p_{Bonferroni}$ < 0.001,

d = 2.313), but there was no difference between target domains at the 0-ms ISI ($t_{(24)}$ = 0.385, $p_{Bonferroni}$ = 1.000, d = 0.077).

Taken together, I replicated the pattern of findings from Experiment 1. Additionally, I found no evidence to suggest that the between-domain shift costs were modulated by the ISI.

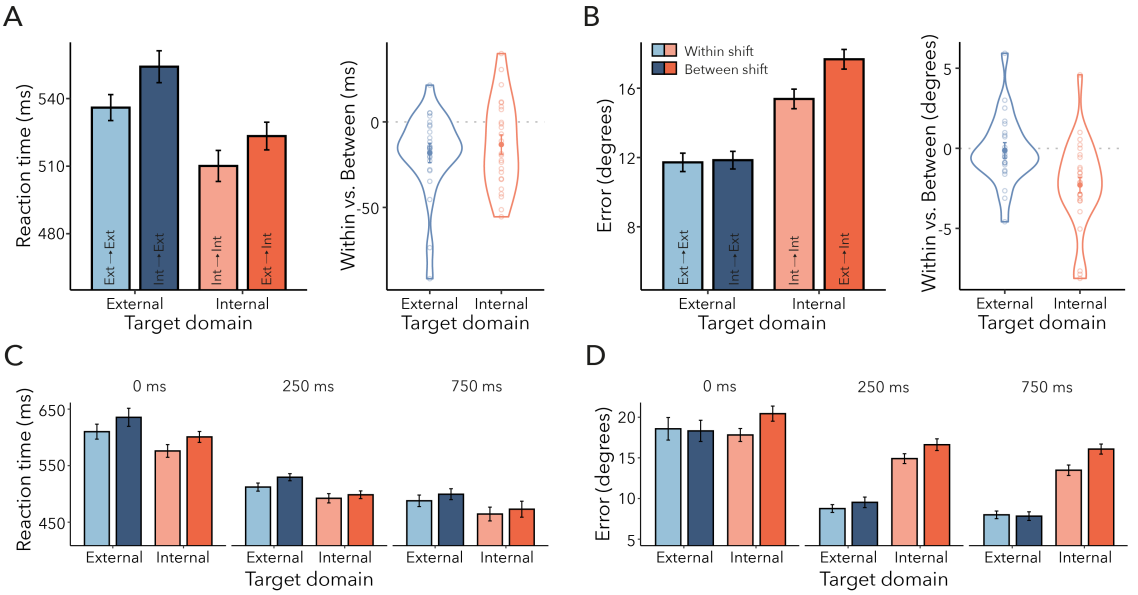

Figure 4.5. No modulation of between-domain shift costs across time. (A) RTs were slower when shifting between as compared to within domains and for external as compared to internal reports. **(B)** The effect of shift type on reproduction errors depended on the level of target domain. There was only a between-domain shift cost when reporting internal but not external targets. **(C)** The shift-cost effect in RTs was not modulated by the ISI condition. **(D)** The shift-cost effect in reproduction errors was not modulated by the ISI condition. Error bars indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

4.4.2.2 Swaps and Guesses: Within-Domain Shifts vs. Between-Domain Shifts

The mixture-modelling approach provided an additional perspective on the factors contributing to performance effects. According to mixture modelling, errors in a continuous-report task can arise due to two potential sources (Bays et al., 2009). Participants could accidentally swap the target with one of the other items or randomly guess. Hence, one plausible explanation for the asymmetric between-domain cost in reproduction errors could be that participants committed more swap errors when shifting from the external to the internal domain. Since the shift cost in reproduction errors was not modulated by the ISI and to improve the fit of the model, I trained a

probabilistic mixture model for each combination of target domain and shift type (i.e., without further splitting the data by ISI).

Figure 4.6A depicts the rate of swap errors. The proportion of swaps was very low across all conditions (M = 0.019; SEM = 0.005) and could not account for the larger reproduction-error costs when shifting between as compared to within domains in internal-target trials. Neither the effect of shift type ($F_{(1,24)}$ = 2.182, p = 0.153, η^2_G = 0.016) nor the interaction between shift type and target domain ($F_{(1,24)}$ = 0.699, p = 0.411, η^2_G = 0.010) reached significance. I did, however, observe more swaps for internal- as compared to external-target trials ($F_{(1,24)}$ = 15.408, p < 0.001, η^2_G = 0.072).

The proportion of guesses was approximately five-times larger than the proportion of swaps (M = 0.104, SEM = 0.013) and followed the asymmetric performance costs (Figure 4.6B). A significant interaction between shift type and target domain was observed ($F_{(1,24)}$ = 4.795, p = 0.038, η^2_G = 0.020). Post-hoc pairwise comparisons showed significantly more guesses for external-to-internal than internal-to-internal shifts ($t_{(24)}$ = 2.947, $p_{Bonferroni}$ = 0.014, d = 0.589). However, no difference in guesses occurred between internal-to-external and external-to-external shifts ($t_{(24)}$ = 1.017, $p_{Bonferroni}$ = 0.638, d = 0.203). A main effect of shift type showed a larger probability of guesses in between- than within- domain shift trials ($F_{(1,24)}$ = 10.425, p = 0.004, η^2_G = 0.037). The target domain did not influence the guessing rate significantly ($F_{(1,24)}$ = 1.612, p = 0.216, η^2_G = 0.021).

In sum, the mixture-model analysis suggested that misreporting the first-cued item did not explain the asymmetric between-domain shift cost observed in reproduction errors. Instead, a greater proportion of guessing occurred when shifting from the external to internal domain as compared to shifting within the internal domain.

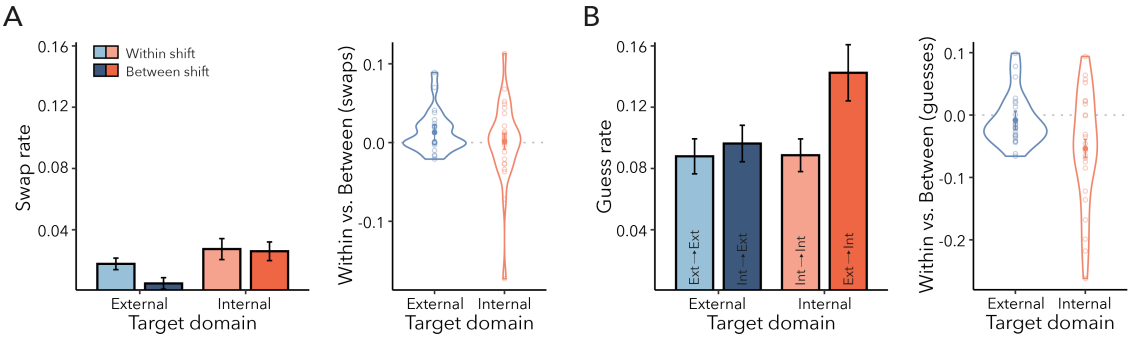

Figure 4.6. Swap and guess rates. (A) Proportion of swaps as a function of target domain and shift type. Swaps with the first-cued item were more likely in internal- as compared to external-target trials. **(B)** Proportion of guesses as a function of target domain and shift type. Between- as opposed to within-domain shifts significantly increased the probability of guesses. Moreover, the effect of shift type was modulated by the level of target domain. Error bars indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

4.4.3 Interim Summary

In Experiment 2, I successfully replicated the results obtained in Experiment 1. Participants consistently exhibited slower RTs when shifting their attention between domains compared to shifting within the same domain. Furthermore, I again observed an asymmetric shift cost in reproduction errors, indicating that external-to-internal versus internal-to-internal shifts incurred higher costs than internal-to-external versus external-to-external shifts. Given the variations in task design and condition ratio between Experiment 1 and Experiment 2, the current findings demonstrate good generalisability of the within- versus between-shift domain effects. In Experiment 2 it also became evident that the observed effects remained consistent regardless of the allocated time for attentional shifts. This implies that the attentional system did not fully recover from the consequences of shifting between domains within the tested ISIs. Moreover, the asymmetric between-domain shift cost in reproduction errors could not be attributed to swaps with the initially cued item. Taken together, the findings from Experiment 2 suggest the possible existence of a distinct, additional control function engaged during shifts of attention between domains.

4.5 Experiment 3: Spatial Dynamics of Shifting Between Attentional Domains

In the previous two experiments, every item was associated with a uniquely coloured placeholder in a spatially distinct location. Consequently, in shift trials, the second cue always required a covert shift of attention to a different location. Experiment 3 was designed to investigate whether between-domain shifts are modulated according to the spatial arrangement of the external and internal items, particularly when these items occur at distinct versus identical locations. Unlike the prior two experiments in which an external and an internal item were always positioned separately within each hemifield, this experiment placed both items at the same location. This task design enabled comparing the costs of between-domain shifts when these required or did not require any spatial reorienting. Hence, this third experiment was designed to investigate the potential role of spatial factors when shifting attention between domains.

4.5.1 Methods

4.5.1.1 Participants

As in Experiment 1 and 2, the desired sample size was $n = 25$. I collected data from 26 participants. One participant was excluded following an a-priori behavioural trial-removal procedure (see 4.5.1.4 Analysis). All remaining 25 participants (age range: 18 to 37; mean age: 22.92; gender: 13 female, 10 male, 2 non-binary; handedness: 21 right-handed, 2 left-handed) had normal or corrected-to-normal vision. Individuals provided written informed consent before participating in the study and were paid £15 per hour. All experimental procedures were approved by the Central University Research Ethics Committee of the University of Oxford.

4.5.1.2 Apparatus and Data Acquisition

Again, participants completed the experiment in a group testing room designed to accommodate up to 20 individuals. However, to maintain optimal testing conditions, no more than six participants were tested simultaneously. Participants sat in front of a monitor (23.6-inch DELL E2417H; resolution 1920×1080 pixels; refresh rate 60 Hz; screen width 53 cm) at a viewing distance of 80 cm. The experimental script was generated in Python using the PsychoPy package (version 2021.2.3; Peirce et al., 2019).

4.5.1.3 Task and Procedure

In Experiment 3, similar to Experiment 1 and 2, participants performed a combined perception and working-memory task (**Figure 4.7A**). This third experiment featured three key task manipulations. First, either one or two cues were presented in the interval separating the internal and external display (i.e., stay vs. shift), analogous to Experiment 2. Second, either cue could indicate an internal representation or upcoming sensory information or (i.e., external vs. internal; same as in Experiment 1 and 2). Again, this resulted in four possible shift conditions: external-to-external, internal-to-internal, external-to-internal, and internal-to-external; the former two being within-domain shifts and the latter two between-domain shifts. Lastly, between-domain shifts trials were divided into those that either did or did not require a spatial shift of attention (i.e., different vs. same location). This final manipulation distinguishes Experiment 3 from the first two experiments. Manipulating the overlap versus differentiation of the cued location in between-domain shift trials required modifications in the task design relative to Experiments 1 and 2.

Trials started with the presentation of two distinctly coloured circular placeholders, appearing on opposite sides of the fixation cross (e.g., a blue placeholder on the left and a green one on the right). Following a 750 ms interval, two randomly tilted and distinctly coloured bars appeared within the placeholders for a duration of 250 ms (e.g., an orange bar in the blue

placeholder, and a pink bar in the green placeholder). The colours of the bars appearing in this internal display were always distinct from those of the placeholders.

The central fixation cross then changed colour for 150 ms, cueing the item most likely to be reported. The colour of this cue could direct attention internally by indicating the colour of a bar presented in the previous internal display (i.e., a retro-cue). Alternatively, the cue could indicate the colour of a placeholder, directing attention externally to the correspondingly coloured bar appearing in the upcoming external display (i.e., a pre-cue). Hence, a cue matching the colour of a placeholder prompted external focus, whereas a cue matching the colour of a previously presented item directed attention internally.

In stay trials, only one cue was presented, similar to Experiment 2. In shift trials, a second cue was presented for 150 ms, appearing 750 ms after the disappearance of the first cue. As in Experiment 1 and 2, the shift condition was further divided into two sub-conditions: within-domain and between-domain shifts. In within-domain shift trials (i.e., external-to-external or internal-to-internal), the second cue always indicated the item appearing within the opposing hemifield of the same domain as indicated by the first cue. For instance, in an external-to-external trial of the example sequence provided in **Figure 4.7A**, if the first cue indicated the blue placeholder on the left, the second cue would indicate the green placeholder on the right. Conversely, in an internal-to-internal trial, if the first cue indicated the orange memory item that was previously presented on the left, the subsequent cue would indicate the pink memory item previously appearing on the right. As items from the same domain were never presented on the same side, within-domain shift trials always demanded attention to shift from one hemifield to the other.

In contrast, between-domain shift trials (i.e., internal-to-external or external-to-internal) could be subdivided into trials that did or did not necessitate a spatial shift of attention (**Figure 4.7B**). For instance, in an external-to-internal trial that required a shift of attention, the first cue could have signalled the left, blue placeholder, which was then superseded by a cue indicating the

pink memory item on the opposite hemifield (Figure 4.7A). Conversely, an external-to-internal trial that did not require a spatial shift of attention could have cued the left, blue placeholder initially, followed by a cue indicating the orange memory item, which was also presented to the left previously within the same blue placeholder.

The first cue in stay trials and the second cue in shift trials were 100% instructive, therefore always indicating the target item to be reported. Immediately following the single cue in the stay condition and the second cue in the shift condition, the second set of randomly oriented bars appeared within each placeholder for 100 ms (i.e., external display). These bars always matched the colour of the placeholder in which they were presented. Following the external display, a dynamic mask, the response dial, and feedback were presented in sequence, mirroring the procedure in Experiments 1 and 2. The colours, sizes, and horizontal distance between stimuli remained consistent with those used in the previous two experiments. In contrast to Experiments 1 and 2, the colour-location mappings remained constant throughout the experiment; however, these mappings were randomised across participants.

The experiment consisted of 672 trials divided across 14 blocks, each including 48 trials (Figure 4.7C). Exactly ⅓ (224) of the total trials were stay trials. The remaining ⅔ (448) were shift trials. Participants were equally likely to report an external or internal item in both stay trials (112 each) and shift trials (224 each). Within shift trials, within- and between-domain shifts were equiprobable (224 each). Thus, the four shift conditions included 112 trials each. Between-domain shift trials were equally likely to include sequential cues for items appearing in the same or in different lateral location(s) (112 each). In contrast, all within-domain shift trials required a spatial shift of attention (224). Each block consisted of ⅓ (16) stay and ⅔ (32) shift trials.

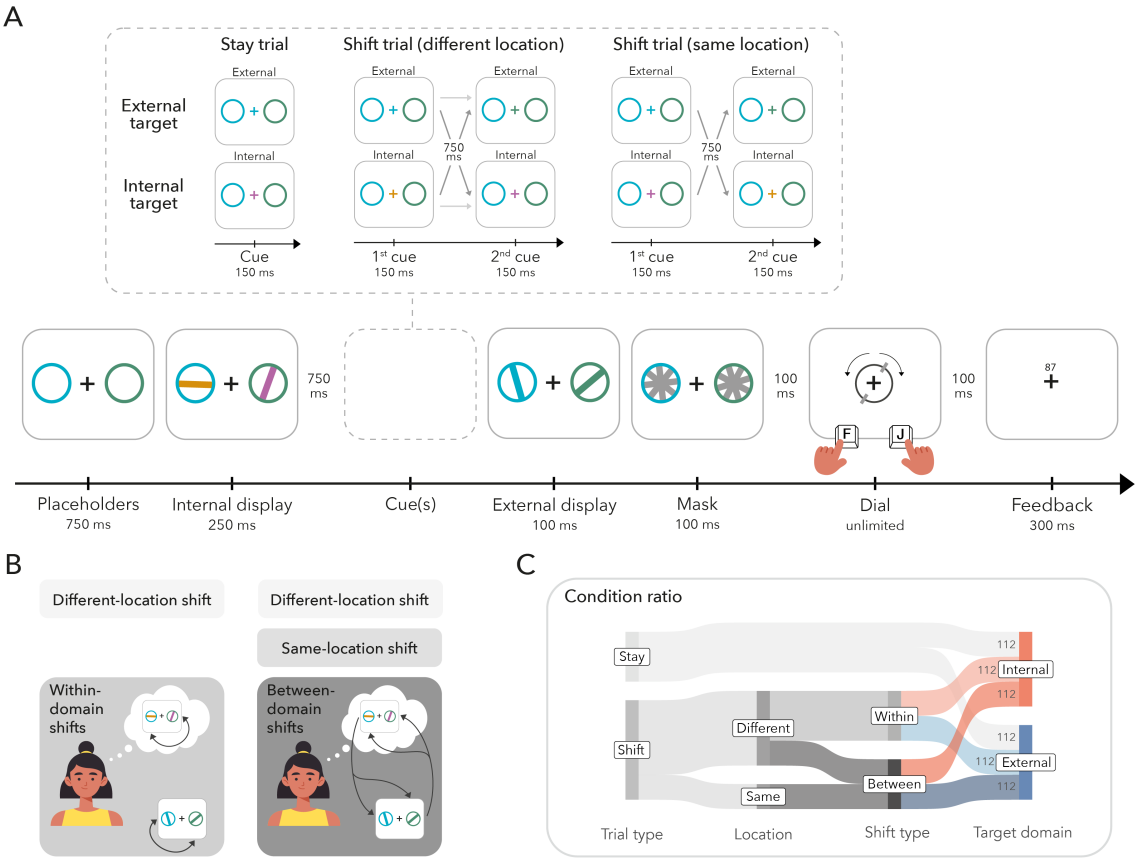

Figure 4.7. Task schematic and condition details of Experiment 3. (A) Illustration of stay and shift trials. In stay trials, only a single cue was presented. In shift trials, two consecutive cues were presented. Light-grey arrows depict within-domain shifts and dark-grey arrows represent between-domain shifts. Within-domain shifts always necessitated a spatial shift between hemifields. Between-domain shifts were categorised into those necessitating a spatial shift of attention (different location) and those not requiring such a spatial shift (same location). **(B)** In within-domain shift trials, participants either had to shift attention external-to-external or internal-to-internal. In between-domain shift trials, participants either had to shift attention external-to-internal or internal-to-external. Additionally, the necessity to execute a spatial shift was manipulated in the between-domain shift condition. **(C)** Ratio of conditions and number of trials across the experiment.

4.5.1.4 Analysis

Data were first pre-processed according to the pipeline matching that of Experiment 1 and 2. During pre-processing, trials were removed when RTs exceeded 5000 ms. Next, I removed trials for which the remaining RTs were 2.5 *SD* above the individual mean across all conditions. None of the participants was removed for loss of more than 10% of trials. One dataset was removed due to average reproduction errors exceeding 45° in a subset of conditions. After these exclusion steps, datasets from 25 participants remained in the main behavioural analysis, with an average of 96.85% (*SD* = 0.42) of trials retained. The conditions of interest were target domain (external

vs. internal), and shift type (within vs. between) and location shift (same vs. different). As in the previous two experiments, the dependent variables of interest were RTs and reproduction errors.

Moreover, for each participant and condition, I fitted a probabilistic mixture model and estimated the relative proportions of the non-target (i.e., first-cued item) distribution as well as the guessing distribution (Bays et al., 2009). The mixture-modelling analysis was performed using the mixtur package (version 1.2.1; Grange & Moore, 2022)

When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples *t*-test and report Cohen's *d* as a measure of effect size. For post-hoc *t*-tests, I report Bonferroni-corrected *p*-values that I denote as $p_{Bonferroni}$.

4.5.2 Results

4.5.2.1 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts

The subsequent analysis concentrated on the shift trials, where participants redirected their attention either within or between domains. My initial aim was to ascertain if the overarching pattern of outcomes paralleled those observed in Experiments 1 and 2, especially considering the modifications made in the task design. To this end, I employed a repeated-measures ANOVA with shift type (within vs. between) and target domain (external vs. internal) as factors.

RTs were slower for between- compared to within-domain shifts (**Figure 4.8A**; $F_{(1,24)}$ = 20.530, $p < 0.001$, $\eta^2_G = 0.006$). In contrast to Experiment 1 and 2 responses to internal compared to external targets were slower ($F_{(1,24)}$ = 4.857, $p = 0.037$, $\eta^2_G = 0.016$). Moreover, there was a significant interaction between factors ($F_{(1,24)}$ = 14.306, $p < 0.001$, $\eta^2_G = 0.004$). Post-hoc comparisons revealed that the between-domain shift cost was higher for reporting internal targets compared to external targets, that is, external-to-internal versus internal-to-internal shifts were more delayed than internal-to-external versus external-to-external shifts ($t_{(24)}$ = 3.782, $p_{Bonferroni} = 0.003$, $d = 0.756$). Additional pairwise post-hoc comparisons revealed that there was a

significant difference between internal-to-internal and external-to-internal shifts ($t_{(24)}$ = 4.698, $p_{Bonferroni}$ < 0.001, d = 0.940). However, the difference between external-to-external and internal-to-external shifts was not significant ($t_{(24)}$ = 0.812, $p_{Bonferroni}$ = 1.000, d = 0.162).

Reproduction errors showed a similar pattern of results as in Experiment 1 and 2. Participants made more errors when they shifted attention between compared to within domains (**Figure 4.8C**; $F_{(1,24)}$ = 15.866, p < 0.001, η^2_G = 0.009). Again, I found a main effect of target domain ($F_{(1,24)}$ = 13.664, p = 0.001, η^2_G = 0.178), with errors being higher for internal than external targets. There was also a significant interaction between target domain and shift type ($F_{(1,24)}$ = 24.440, p < 0.001, η^2_G = 0.018). Shifting between versus within domains incurred a higher cost when reporting internal targets compared to external targets, that is, accuracy declined more in external-to-internal versus internal-to-internal shifts compared to internal-to-external versus external-to-external shifts ($t_{(24)}$ = 4.944, $p_{Bonferroni}$ < 0.001, d = 0.989). Moreover, the difference between internal-to-internal and external-to-internal shifts reached significance ($t_{(24)}$ = 5.565, $p_{Bonferroni}$ < 0.001, d = 1.113), while the difference between external-to-external and internal-to-external shifts was not significant ($t_{(24)}$ =-1.402, $p_{Bonferroni}$ = 0.522, d = 0.280). Taken together, both RTs and reproduction errors revealed an asymmetric between-domain shift cost in Experiment 3.

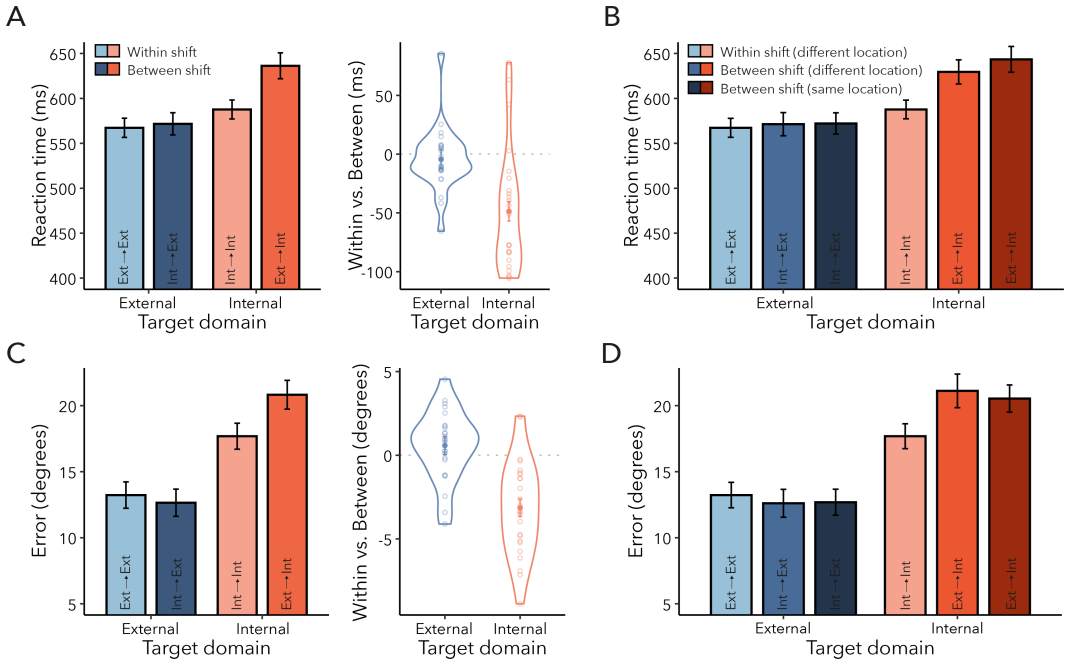

Figure 4.8. Asymmetric between-domain shift cost is not modulated by the necessity to perform a spatial shift of attention. (A) Significant between-domain shift cost in RTs when reporting internal but not external targets. **(B)** RTs were comparable for between-domain shifts, regardless of whether they were to the same or to a different location. **(C)** Significant between-domain shift cost in reproduction errors when reporting internal but not external targets. **(D)** Reproduction errors were comparable for between-domain shifts, regardless of whether they were to the same or to a different location. Error bars indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

4.5.2.2 RTs and Errors: Spatial Shift vs. Non-Spatial Shift

Next, I addressed whether the higher cost of shifting between domains might be modulated by the necessity to perform a spatial shift of attention. Considering the unbalanced nature of the task design, notably that shifts within domains were always accompanied by a spatial shift in attention (unlike between-domain shifts), I merged the factors shift type and location shift into a single factor with three levels (shift type + location: within-different vs. between-different vs. between-same). I then employed a repeated-measures ANOVA with shift type + location and target domain as factors.

Figure 4.8B shows that RTs were faster for external than internal targets ($F_{(1,24)}$ = 6.102, p = 0.021, η^2_G = 0.211). Moreover, there was a significant main effect of shift type + location ($F_{(1,24)}$ = 15.901, p < 0.001, η^2_G = 0.006) and a significant interaction between factors ($F_{(1,24)}$ = 8.361, p < 0.001, η^2_G = 0.004). Post-hoc comparisons indicated a significant difference solely for internal-

report trials. In internal-report trials, participants exhibited slower RTs for between-domain shifts that required a spatial shift of attention ($t_{(24)}$ = 5.157, $p_{Bonferroni}$ < 0.001, d = 1.031), as well as for those that did not necessitate such a spatial shift ($t_{(24)}$ = 3.436, $p_{Bonferroni}$ = 0.013, d = 0.687), as compared to within-domain shifts. All other comparisons, including those comparing between-domain shifts with and without spatial attention shifts, yielded no significant differences (see **Supplementary Table 8.10A** for all relevant post-hoc comparisons).

As depicted in **Figure 4.8D**, reproduction errors revealed a similar pattern of results with both main effects (target domain: $F_{(1,24)}$ = 15.990, p < 0.001, η^2_G = 0.193; shift type + location: $F_{(1,24)}$ = 4.615, p = 0.015, η^2_G = 0.007) and the interaction effect ($F_{(1,24)}$ = 10.029, p < 0.001, η^2_G = 0.015) being significant. As for RTs, post-hoc comparisons only revealed significant differences for internal-target trials. Reproduction errors were higher in between-domain shift trials that entailed a spatial shift of attention ($t_{(24)}$ = 4.068, $p_{Bonferroni}$ = 0.003, d = 0.814), and in between-domain shift trials not requiring spatial attention shifts ($t_{(24)}$ = 4.677, $p_{Bonferroni}$ = < 0.001, d = 0.935), in contrast to within-domain shift trials. All additional comparisons, including those contrasting between-domain shifts with and without spatial attention shifts, did not produce any significant differences (see **Supplementary Table 8.10B** for all relevant post-hoc comparisons).

Taken together, these findings suggest that whether or not a spatial shift occurs during transitions between domains does not influence the behavioural costs associated with between-domain shifts.

4.5.2.3 Swaps and Guesses: Within-Domain Shifts vs. Between-Domain Shifts

As in Experiment 2, I applied a mixture-model approach to the data. Since the shift cost in reproduction errors was not modulated by the shift location and to improve the fit of the model, I trained a probabilistic mixture model for each combination of target domain and shift type (i.e., without further splitting the data by shift location).

Figure 4.9A depicts the rate of swap errors. Again, the proportion of swaps was very low across all conditions (M = 0.042; SEM = 0.014) and could not account for the larger reproduction-error costs when shifting between as compared to within domains in internal-target trials. Neither the effect of shift type ($F_{(1,24)}$ = 0.070, p = 0.793, η^2_G < 0.001) nor the interaction between shift type and target domain ($F_{(1,24)}$ = 3.925, p = 0.059, η^2_G = 0.021) reached significance. I did, however, observe more swaps for internal- compared to external-target trials ($F_{(1,24)}$ = 8.285, p = 0.008, η^2_G = 0.090).

The proportion of guesses was approximately three-times larger than the proportion of swaps (M = 0.134, SEM = 0.026; Figure 4.9B). A main effect of shift type showed a larger probability of guesses in between- than within- domain shift trials ($F_{(1,24)}$ = 8.019, p = 0.009, η^2_G = 0.016) and for internal- compared to external-target trials ($F_{(1,24)}$ = 4.455, p = 0.045, η^2_G = 0.083). The interaction between shift type and target domain was not significant ($F_{(1,24)}$ = 3.525, p = 0.073, η^2_G = 0.008).

In sum, the mixture-model analysis suggested that neither swaps nor guesses alone can entirely replicate the asymmetric between-domain shift cost. Therefore, it seems plausible that in Experiment 3 the combined influence of random guessing and misreporting the target item was responsible for the shift-cost imbalance between the external and internal domains (see Supplementary Figure 8.16 for combined effect of swaps and guesses). Moreover, I would like to note the potential significance of contrasting the ratio of swaps and guesses in spatial versus non-spatial between-domains shift trials. Yet, the limited number of between-domain shift trials, amounting to only 56 trials for each target domain, restricted the feasibility of accurately modelling the distributions of guesses and swaps.

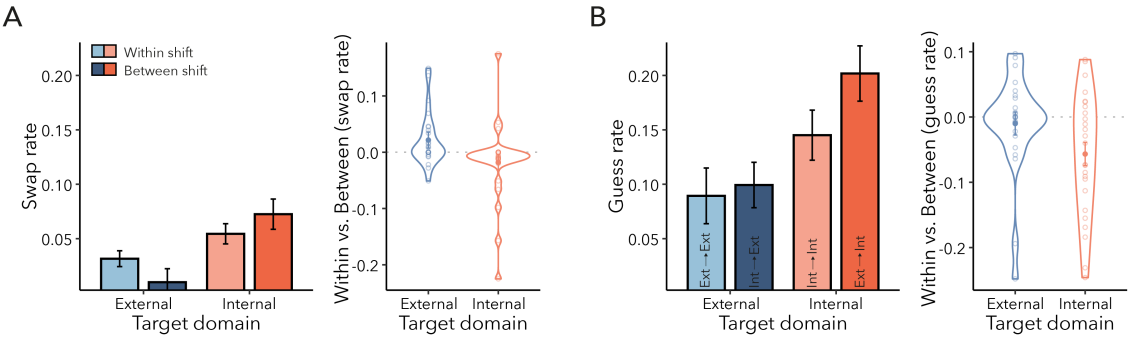

Figure 4.9. Swap and guess rates. (A) Proportion of swaps as a function of target domain and shift type. Swaps with the first-cued item were more likely in internal- as compared to external-target trials. **(B)** Proportion of guesses as a function of target domain and shift type. Between- as opposed to within-domain shifts significantly increased the probability of guesses. Error bars indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

4.5.3 Interim Summary

Experiment 3 replicated the asymmetric shift cost in reproductions errors observed in the previous two experiments. The current study additionally revealed a similar shift-cost imbalance in RTs. Specifically, responses were significantly more delayed following external-to-internal versus internal-to-internal shifts as compared to internal-to-external versus external-to-external shifts, mirroring the asymmetric shift cost in reproduction errors. Similar to Experiment 2, the mixture-modelling results suggested that swaps with the initially cued item cannot alone account for the asymmetric shift cost.

The main aim of Experiment 3 was to investigate whether spatial factors impacted between-domain shifts. The modified task design allowed the costs of between-domain shifts which required a spatial shift of attention to be contrasted against those that did not require any spatial reorienting. The current findings unveiled a consistent pattern in RTs and reproduction errors during between-domain shift trials, irrespective of whether external and internal items were positioned in the same or different location(s). Taken together, these findings indicate that performance in between-domain shift trials was unaffected by the necessity to shift attention spatially.

4.6 Discussion

The aim of the present studies was to launch a new line of research investigating a fundamental yet little explored aspect of everyday cognition; that is, shifting attention between external sensory signals and internal contents in working memory. In three experiments, participants were asked to reorient their attention to equivalent stimulus attributes in perception and working memory to achieve a common task goal, thereby echoing numerous real-life situations. Take, for example, the task of building a Lego set, in which our attention constantly shifts between the various sensory contents (i.e., the instructions, the pile of Lego pieces, and the partially assembled model) and mental representations (i.e., visualising the required pieces, how they fit together, and their placement within the model). Thus, the current experiments both complement and differ from previous studies, which have focussed on the costs associated with shifting between trials relying on external versus internal information (Calzolari et al., 2022; Carlson et al., 1993; Dark, 1990; Gilbert et al., 2005; Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019; Weber et al., 1986) and utilised different stimulus material for the external and internal domain (Verschooren et al., 2021; Verschooren, Schindler, et al., 2019).

The newly developed tasks proved effective in assessing the performative consequences of shifting attention within versus between perceptual and mnemonic contents. In Experiment 1, I observed significantly higher costs in both RTs and errors for shift as compared to stay trials. Additionally, this experiment revealed that the costs associated with shifting attention were notably higher when reorienting between different domains as opposed to within the same domain. The continuous proxy of spatial orienting in the form of small fixational gaze biases confirmed participants utilised both cues during the task and further uncovered interesting dissociations between the spatial orienting functions tracked by the oculomotor biases and the ultimate consequences to task performance as reflected in the behavioural performance measures. Experiment 2 replicated the behavioural findings of Experiment 1 and additionally demonstrated that participants were not able to recover from the between-domain shift costs

even when more time was provided to shift attention. Lastly, in Experiment 3, I uncovered that spatial between-domain shifts did not impose higher behavioural costs than non-spatial between-domain shifts.

4.6.1 Validation of a Novel Task Design

The present results yielded compelling evidence that participants utilised the cues to guide their attention. In Experiment 1, RTs were slower and reproduction errors greater when required to reorient attention to a different item. These findings align with previous research showing behavioural costs when being invalidly cued (Griffin & Nobre, 2003; Posner, 1980). Hence, the highly predictable first cue incentivised participants to focus their attention on the cued location, enabling the detection of attentional shift costs when the dial appeared. Eye-tracking measures provided convergent evidence that participants utilised the first cue. Fixational eye movements offer a continuous proxy for the dynamics of spatial shifts of attention in the external and internal domains (Draschkow et al., 2022; Engbert & Kliegl, 2003; Hafed et al., 2011; van Ede, Chekroud, & Nobre, 2019; van Ede et al., 2021), even though they may not be obligatory for causing attention shifts (B. Liu et al., 2022; Willett & Mayo, 2023). The results of Experiment 1 showed that gaze was robustly biased towards the first-cued location regardless of whether upcoming sensory information or working-memory representations were cued.

Moreover, Experiment 1 represents the first attempt to compare directly gaze biases towards external and internal stimuli within a single task. The latency of gaze biases did not differ between external and internal first cues, but the bias was more sustained when the first cue oriented attention to the external domain. The relatively smaller magnitude of the gaze bias observed after internal first cues could signify a transient access event, during which the sensory properties of the cued item are reactivated to enhance its prioritisation in memory. In contrast, when attention is cued to the location of an upcoming stimulus, it may be beneficial to sustain covert attention towards the location of the anticipated item. This pattern aligns with earlier

neurophysiological research reporting alpha-power lateralisation to be sustained following pre-cues but transient following retro-cues (Mok et al., 2016; Wallis et al., 2015). However, it is important to note that the observed pattern may not generalise to all instances of external and internal attention. In this study, external attention required spatial anticipation of an upcoming perceptual event, while in everyday situations, our attention is often focussed on currently available sensory content. Similar to recent advances on a neurophysiological level (Chota et al., 2023), it would be intriguing to investigate external gaze biases when the to-be-reported sensory contents are perceptually available at the time of the cue. This would allow for a compelling comparison against the external gaze biases observed in the present study, wherein attention was proactively guided to an upcoming sensory event.

4.6.2 Costs of Between-Domain Shifts

A central finding of the present study was that behaviour was more strongly impaired when shifting between as compared to within attentional domains. This is in line with results from previous literature wherein participants were required to shift within and between trials relying on perceptual information and trials relying on internal representations involving imagination (Gilbert et al., 2005), self-oriented cognition (Calzolari et al., 2022), verbal working-memory (Carlson et al., 1993; Dark, 1990; Weber et al., 1986), or visual memory (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019). Importantly, the current results reveal that the between-domain shift cost is present in continuous and dynamic behaviour drawing from short-term visual representations and sensory stimulation towards a single behavioural response.

While RTs showed a similar between-domain shift cost for external- and internal-target reports in Experiment 1 and 2, errors (and RTs in Experiment 3) suggested that the costs of shifting between domains might depend on the direction of the shift, with larger costs when shifting from perception to working memory than vice versa. The asymmetrical pattern I observed closely

resembles those reported in previous research. A number of studies have demonstrated greater RT costs when transitioning from an external to an internal trial as compared to the inverse (Calzolari et al., 2022; Dark, 1990; Hautekiet et al., 2023; Verschooren et al., 2020), with some evidence for comparable asymmetries in reproduction errors (Verschooren, Liefooghe, et al., 2019). However, not all studies have shown the same directional bias (Gilbert et al., 2005). It is worth noting that these previous studies exclusively utilised binary responses instead of continuous reports, which may have limited sensitivity in detecting such asymmetries in error rates. Furthermore, direct comparisons between the current experiments and previous research may be challenging due to differences in task demands. In particular, most previous research compared performance when shifting between trials requiring responses based on external and internal stimuli. In contrast, the experiments used here compared shifting between external and internal contents within a single trial to facilitate a unified behavioural response.

One potential explanation for the asymmetric shift cost could be attributed to memory-retrieval costs when transitioning from perception to memory. In a list-completion task, the between-domain shift cost disappeared when participants were cued about the to-be-relevant memory item 750 ms in advance, and thus could select the appropriate domain before being required to respond (Dark, 1990). Based on these findings, it was argued that the asymmetric shift cost could not be purely explained by the involvement of control mechanisms coordinating transitions between external and internal tasks, but instead must additionally reflect memory retrieval processes triggered by the internal task. Similarly, it could be argued that the asymmetrical cost seen in the current experiments may be due to the duration required to shift attention back and retrieve information from the internal domain. However, manipulating the time between the second cue and the external display in Experiment 2 demonstrated that the asymmetric shift cost remained consistent across all tested ISIs (i.e., 0 ms, 250 ms, and 750 ms). That is, allowing for more retrieval time did not diminish or extinguish the asymmetric shift cost. Consequently, the memory retrieval account seems unlikely to explain the current findings.

Alternatively, incomplete shifts and the resulting interference from the previously attended domain could have acted as a potential source of the shift-cost imbalances between domains. None of the three experiments provided reminders regarding which target to report at the response phase and instead forced participants to report whichever item was cued last before the external array. Hence, participants may have been susceptible to target confusions and biases. The current findings revealed that transitioning between as compared to within domains was more costly for internal- than external-target reports. Thus, initially engaging external attention might have been more impairing on the final report than first being cued about an internal representation. Indirect evidence comes from the gaze-bias findings showing that the initially cued item continued to attract attention even after the second cue was presented and this was especially pronounced if the first cue indicated an external item. However, swap rates were generally very low and did not (Experiment 2) or not exclusively (Experiment 3) account for the asymmetric shift cost observed in reproduction errors. Instead, these swap errors consistently occurred more often for internal- than external-target reports, without being modulated by the type of shift. With this is mind, the proportion of guesses may provide a more compelling explanation for the asymmetric shift cost seen in reproduction errors, suggesting a loss of information regarding the internal display after attention was shifted to the external domain.

An alternative explanation of the shift-cost asymmetries is offered by the associative interference account (Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019), which is closely related to the IDEA hypothesis (Verschooren & Egner, 2023). Both frameworks propose a dominant default bias towards internal as compared to external attention, resulting in more efficient shielding of internal compared to external contents in most situations. An exception occurs when there is a need to shift attention from externally available information to internal information. Future studies that systematically investigate the nature of errors caused by both external and internal interference in either domain would go a long way in improving our understanding of interference effects and their contribution to shift-cost asymmetries.

4.6.3 Behaviour-Gaze Dissociations When Shifting Within vs. Between Domains

Although I observed significantly slower RTs when shifting attention between as compared to within domains, directional biases in fixational gaze behaviour showed no corresponding discrepancy in the latency or magnitude following the second cue. The gaze-bias results converge with recent research showing that microsaccade direction may not necessarily fully mimic the ultimate behavioural and neural effects of visuospatial attention (Willett & Mayo, 2023; but see van Ede, 2023). Of course, the null findings observed when contrasting gaze biases across within- and between-domain shifts should be interpreted with caution; although there are multiple factors that could explain the lack of an effect. One could for instance infer that gaze biases are a noisier measure compared to RTs and reproduction errors. However, other explanations are also possible. Gaze biases are most closely related to spatial orienting within attention, but the final behavioural performance is a cumulative measure of various attention-related functions, beyond mere orienting. Both external and internal attention involve multiple functions, such anticipating, orienting, selecting, prioritising and reprioritising, and configuring contents in preparation for upcoming response demands. These are not yet fully agreed by researchers and are likely to differ between external and internal attention by virtue of the affordances of the different domains (Myers et al., 2017; van Ede & Nobre, 2023). The present results suggest that the behavioural effects measured at the end of the trials were not purely driven by the orienting functions proxied by gaze measures but might involve additional functions that ultimately serve as bottlenecks to performance.

For example, additional challenges can occur during the prioritisation of relevant contents after a spatial shift of attention has been performed. While the gaze bias continuously tracked spatial orienting of attention, RTs could have been delayed because it took longer to prioritise the second-cued information in between-domain shift trials. Alternatively, the difference in RTs between shift types could be due to challenges in response preparation. However, if this was the case, I would have expected the levels of shift type to be modulated by the duration of the ISI in

Experiment 2. Namely, when given enough time to recover, the shift cost should have diminished. Instead, differences between shift types did not differ across ISIs, suggesting that the cognitive system was not able to recover from the detrimental effects of between-domain shifts within the tested ISIs.

The results of Experiment 2 may suggest the presence of a distinct aspect associated with transitioning between domains. The processes of deactivating the currently irrelevant attentional domain and activating the relevant one may require additional control mechanisms compared to adjustments within the same domain, almost as if an additional switch or checkpoint occurs during shifts between these two attentional dimensions – a dynamic that may not be encountered when shifting within the same domain. We could imagine this extra checkpoint as a portal being activated whenever we shift between the external and internal dimensions (**Figure 4.10**).



Figure 4.10. Shifts between domains might be affected by an additional control process. The current findings suggest the presence of an additional control point when shifting attention between domains, a feature not encountered during within-domain shifts. We can imagine this additional checkpoint as a portal mediating between external and internal attention.

4.6.4 Spatial Reorienting During Between-Domain Shifts

Experiment 3 did not reveal any additional costs of spatial reorienting during between-domain shifts. However, prior research has indicated costs associated with covert shifts of spatial attention in the external (Carrasco, 2011; Posner, 1980) and internal domain (Griffin & Nobre, 2003; Landman et al., 2003; Souza & Oberauer, 2016). These costs are not exclusive to tasks where spatial shifts of attention occur reactively following an invalid spatial cue. For instance, valid cues prompting participants to proactively shift external attention to a new location have been found

to delay RTs as compared to cues instructing them to maintain focus on a single location (Ibos et al., 2009). Furthermore, these RT delays were more pronounced for spatial shifts to distant locations compared to proximate ones (Geweke et al., 2021), and for shifts across opposite hemifields compared to those within the same hemifield (Ibos et al., 2009). As such, the mere act moving attention through space, whether in a proactive or reactive manner, seems to be inherently costly. Consequently, it was plausible to expect decreased performance in trials involving between-domain shifts that required spatial reorienting, compared to those without such spatial shifts.

On the other hand, maintaining attention at a single location may have come with its own costs compared to situations where attention shifted between locations. In particular, interference effects could have been induced by the external and internal items occupying the same placeholder. This is supported by the working-memory literature, wherein perceptual distractors appearing at the location of a task-relevant memory item (or close to it) induce higher performance costs compared to those occurring further away (van der Stigchel et al., 2007; S. Wang & van Ede, 2023). However, such a spatial overlap of external and internal items may also hold differential influences depending upon which attentional domain is behaviourally relevant. For example, Awh et al. (1998) has demonstrated benefits for external attention to locations currently maintained in working memory as opposed to other locations.

Additionally, or alternatively, the temporal segregation of stimulus input, achieved by presenting sequential displays, may have further influenced behavioural performance and the typical effects associated with shifting spatial attention. Consistent with this idea, the working-memory literature suggests that dividing the presentation of to-be-remembered items into two sequential displays tends to significantly improve performance compared to presenting the same items simultaneously (Ihssen et al., 2010). In Experiment 3, the detrimental effects of spatial attention shifts might have been diminished because shifts occurred across two separate displays when transitioning between domains. At this moment, I can only provide conjectures about the

factors underlying the results in Experiment 3. Nevertheless, it seems clear that a combination of multiple factors must have contributed to produce the observed pattern of results.

4.6.5 Behavioural Differences in Reporting External and Internal Targets

Even though I tried to match the difficulty of external- and internal-report trials as closely as possible, the present findings revealed that responses to internal targets were generally faster but less precise than responses to external targets (except for the reproduction errors in the 0-ms ISI condition in Experiment 2, and RTs in Experiment 3). This discrepancy can be attributed to an inherent asymmetry between the two attentional domains in terms of their temporal dynamics and their respective competition throughout the duration of a trial. For instance, RTs in internal-report trials may have been generally faster because participants had an extended window for response preparation compared to external-report trials. After being cued to the internal domain, participants could select the item they would report and ignore the external display. Furthermore, the prolonged time gap between the internal display and the reporting phase may have caused some memory decay (J. Brown, 1958; Conrad, 1967; Cowan, 1988) and thereby contributed to greater reproduction errors when internal items had to be reported.

Alternatively, or additionally, interference of the stimuli within and between domains may be crucial in explaining the performance differences observed between external and internal attention. Previous literature has shown that internal representations can be biased and disrupted by external interference (Lorenc et al., 2021; Rademaker et al., 2015; Teng & Kravitz, 2019; van Ede, Board, et al., 2020) as well as by other memory representations (Bae & Luck, 2017; Brady & Alvarez, 2011; Chunharas et al., 2022; Dubé et al., 2014; J. Huang & Sekuler, 2010; Teng & Kravitz, 2019). While external attention can also be modulated by working-memory contents (Olivers et al., 2006; Soto et al., 2008; Teng & Kravitz, 2019) or other perceptual events (Duncan & Humphreys, 1989; Treisman & Gelade, 1980; Wolfe, 1994), it may be possible that in the current tasks, internal attention was generally more susceptible to interference compared to external

attention. In line with this idea, recent research has demonstrated that the influence of internal representations on perception decreases with sufficient time between an encoding display and a perceptual display (i.e., from 550 ms onwards), while internal representations become increasingly biased by perception (Teng et al., 2023). Therefore, in the current study, as the memory stimuli recede into the past, they might have become less distinguishable from the external stimulation, increasing the overall swap rate and reproduction errors in internal- as compared to the external-target trials. This interference effect might have been especially pronounced in Experiment 3, where pairs of external and internal items were presented at the same location. Consequently, this spatial overlap between external and internal information might have additionally led to slower RTs when reporting internal targets as compared to external ones (a pattern that was reversed in Experiment 1 and 2 where each item occurred at a distinct location).

Furthermore, the act of reorienting attention from one item to another may be more detrimental to the accuracy of target reports when shifting (whether spatially or non-spatially) to an internal as opposed to an external target. Studies employing two consecutive retro-cues have demonstrated that shifting attention between mnemonic contents can lead to a retro-cue benefit compared to a no-cue condition (Guo et al., 2024; Landman et al., 2003; Q. Li & Saiki, 2014; R. Liu et al., 2024; Rerko & Oberauer, 2013; van Moorselaar, Olivers, et al., 2015). However, some of these studies revealed that performance in double retro-cue trials does not reach the same level as in single retro-cue trials (Guo et al., 2024; R. Liu et al., 2024; van Moorselaar, Olivers, et al., 2015). In contrast reorienting between two pre-cues, hence within the external attentional domain, does not appear to exert a significant impact on external-task performance (Dowd & Golomb, 2019). Nevertheless, generalising these findings to a multi-domain task setting is challenging, as previous studies focussed solely on reorienting within a single domain.

While a naturalistic study of shifting attention between domains necessitates a dynamic task setting, I recognise that various factors may have contributed to the performance differences

between reporting items based on perception and working memory. In the current set of studies, I can only highlight and offer speculative insights into the potential roles of these factors.

4.6.6 Conclusion

The three experiments presented in this chapter open new avenues for investigating the fundamental property of shifting attention between perception and working memory. My findings demonstrate that shifting between as compared to within domains incurs behavioural costs. Notably, these costs persisted even when additional time was provided to complete the shift and when the necessity to spatially shift attention was removed. Gaze biases, acting as a proxy of spatial orienting functions, did not mirror the behavioural performance patterns, suggesting bottlenecks at additional attention functions. To shed light on these mechanisms, in Chapter 5, I will employ MEG to compare the neural dynamics underlying within- and between-domain shifts. This enables a more comprehensive understanding of the cognitive and neural mechanisms involved in this ubiquitous act of shifting attention between the external environment and working memory to fulfil goals in everyday cognition.

5 Neural Dynamics of Shifting Between Attentional Domains

5.1 Abstract

The findings of Chapter 4 demonstrated additional costs associated with shifting between contents in the sensory environment and internal memory representations. Here, I used MEG to compare the neural dynamics of shifting attention to visual contents within versus between the external and internal domains. Participants performed a combined perception and working-memory task, similar to the tasks described in Chapter 4. In the experiment, two sequential cues guided attention to upcoming (external) or memorised (internal) sensory information. Critically, the second cue could redirect attention to visual content within the same or alternative domain as the first cue. Multivariate decoding unveiled distinct patterns of human brain activity when shifting attention within versus between domains. Brain activity associated with between-domain shifts was broadly distributed, highly dynamic, and predictive of behavioural performance. Intriguingly, within- and between-domain shifts did not differ during the early attentional orienting stage. This was evidenced by alpha lateralisation, which showed similar time courses when cues redirected attention to the same versus the alternative domain. Moreover, the current findings suggest that neural states associated with a given domain lingered and influenced subsequent shifts of attention within versus between domains. In sum, Chapter 5 provides first insights into the neural dynamics that govern attentional shifts between perception and working memory.

The findings presented in Chapter 5 can be found on the preprint server bioRxiv as Gresch et al. (2024a).

5.2 Introduction

While the previous chapter explored behavioural correlates of transitions between contents in perception and working memory, Chapter 5 investigates the neural dynamics underlying these shifts. To date, only one EEG study has attempted to compare the neural processes associated with within- and between-domain shifts (Verschooren et al., 2021). This previous study was limited to investigating between-domain shifts in only one direction (i.e., external-to-external vs. internal-to-external) and relied on different stimulus materials and task demands across the domains, leaving many questions unanswered.

The only research providing a systematic comparison of the neural correlates associated with within- versus between-domain shifts is an fMRI study by Gilbert et al. (2005). Utilising a task design where participants transitioned between phases relying more heavily on sensory or internalised stimuli, this study underscored the significance of the lateral rostral PFC in shifting attention between the external and internal domains. Other brain imaging and intracranial studies have provided more anecdotal evidence, suggesting that the hippocampus (Poskanzer & Aly, 2023; Treder et al., 2021) or the fronto-insular cortex (Sridharan et al., 2008; Tang et al., 2012) might serve as the interface between perception and memory. Yet, beyond these initial clues, the dynamic neural mechanisms of cross-domain attention shifts are all to be discovered.

The aim of this chapter was to extend these initial steps by investigating the brain dynamics associated with between-domain shifts, and directly comparing them to the dynamics linked to attention shifts within domains. Building on Chapter 4, I used a combined perception and working-memory task in which informative cues indicated the location of sensory or mnemonic contents for subsequent report. Successive cues necessitated shifts of attention either within or between perception and working memory. I recorded MEG to track brain activity with high temporal resolution while participants performed the task and used multivariate pattern analysis to compare activity related to attentional shifts of different types. Based on my findings and those of others (Dark, 1990; Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren,

Liefooghe, et al., 2019; Weber et al., 1986) demonstrating higher performance costs when shifting attention between domains as compared to within domains, I hypothesised that distinct processes may be engaged when attention shifts between domains.

To preview the results, I demonstrate that dissociable patterns of brain activity accompany shifts between perception and working memory as opposed to shifts within domains. These unique patterns of brain activity invoked by between-domain attention shifts are broadly distributed, temporally dynamic, predict behavioural performance, and occur without interfering with the timing of spatial attention shifts as indexed by the lateralisation of posterior 8-12 Hz alpha activity. In addition, I show that distinct patterns of brain activity are associated with external and internal attentional states themselves. These findings provide important first clues about the dynamic neural bases that govern shifts of attention between perception and working-memory contents.

5.3 Methods

5.3.1 Participants

The study was approved by the Central University Research Ethics Committee of the University of Oxford. The sample size was set to n = 25 as in the previous experiments presented in Chapter 4. To yield the targeted number of participants, I collected data from 26 participants in total. One participant was excluded following an a-priori behavioural trial-removal procedure (see 5.3.4 Behavioural Analysis). Of the remaining 25 participants (age range: 19 to 36; mean age: 26.24; 14 female, 11 male), 22 reported being right-handed (three left-handed). Individuals provided written informed consent before participating in the study and were paid £15 per hour.

5.3.2 Task and Procedure

The current task design builds on Experiments 1 and 2 presented in Chapter 4. Two separate bilateral arrays each containing two randomly oriented bars were presented, with one array occurring before and the other after one or two spatially informative colour cues (Figure 5.1). That is, two bars appeared in an earlier internal working-memory array (encoded before the cues), while another two bars appeared in a later external perceptual array (appearing after the cues). At the end of each trial, participants reproduced the orientation of the bar that was last cued. Cues indicating the location of a previously encoded bar guided attention internally (i.e., retro-cues, internal cues), instructing which item in working memory was relevant for reporting. In contrast, cues indicating a previously unoccupied location directed attention externally (i.e., pre-cues, external cues), instructing which item in the upcoming perceptual array would be relevant. The orthogonal manipulation of the to-be-reported target domain (external vs. internal) and the shift type (within-domain shift vs. between-domain shift) enabled the independent investigation of these two factors in a systematic manner.

At the start of each trial, a white central fixation cross (RGB value: [255, 255, 255]) surrounded by four coloured circular placeholders appeared against a black background (RGB value: [0, 0, 0]) for 750 ms. Each of the placeholders appeared in one of four highly distinguishable colours (RGB values: light blue [0, 159, 183], orange [217, 120, 0], green [47, 167, 0], pink [247, 35, 255]). The location mapping of the four placeholder colours was randomly determined on each trial. Placeholders subtended 3.5° visual angle and were centred in the top-left, top-right, bottom-left, and bottom-right at a distance of 4.5° visual angle from fixation. Placeholders stayed on the screen until the response was required at the end of the trial.

Next, two tilted grey bars (RGB value: [50, 50, 50]) subtending approximately 0.5×3.4° appeared for 500 ms inside two of these placeholders (referred to as “internal display”). The orientation of the bars was randomly determined between 0 and 180°. Participants were instructed to memorise the orientation of these two bars. There were four possible configurations

of the internal display, which occurred equally often throughout the experiment: the bars could appear in the placeholders located in the (1) top-left and top-right, (2) bottom-left and bottom-right, (3) top-left and bottom-right, and (4) top-right and bottom-left. As such, at the time of the internal display, there was always one bar presented in each visual hemifield.

The internal display was followed by a delay of 750 ms, during which only the fixation cross and the placeholders remained on the screen. The central fixation cross then changed colour to match that of one of the placeholders for 250 ms. The colour change acted as an informative cue, indicating the potential item to be reported. This cue could indicate a placeholder that previously contained an item during the internal display (i.e., retro-cue) or a placeholder in which no item had been presented yet (i.e., an item of the upcoming external display; pre-cue). The cued location and attentional domain were pseudo-randomised such that each of the four placeholder locations and each domain was equally often indicated by the cue.

In 50% of the trials, the central fixation cross changed colour again after 1250 ms to match the colour of another placeholder for 250 ms (i.e., double-cue or shift trials). In the remaining 50% of the trials, only a single cue (i.e., single-cue or stay trials) was presented. In single-cue trials, the cue appeared at the same time as the first cue in double-cue trials and was 100% instructive, always indicating the target for report. These randomly intermixed single-cue trials were included only to ensure that participants relied on the initial cue when performing the task. However, single-cue trials did not serve a purpose in terms of the behavioural or neural comparisons relevant to the research question.

The double-cue trials were of central interest since these trials required a shift of attention. In particular, the double-cue condition was further divided into two sub-conditions: within- and between-domain shifts. In within-domain shift trials (i.e., external-to-external or internal-to-internal), the second cue indicated the item appearing within the opposing hemifield of the same domain as indicated by the first cue. In between-domain shift trials (i.e., external-to-internal or internal-to-external), the second cue indicated the item in the alternative domain

occurring in the other hemifield. Thus, in within- and between-domain shift trials, attention always needed to be reoriented from one hemifield to the other. The second-cued location and domain were pseudo-randomised such that each of the four placeholder locations and each of the two domains were equally often indicated. In double-cue trials, the second cue instructed the target item for report with a certainty of 100%, thus overruling the first cue. Accordingly, the target item was always identical to the second-cued item. The colour of the first and second cue always differed in double-cue trials.

The single cue in single-cue trials and the second cue in double-cue trials were followed by an ISI of 1250 ms before the onset of a perceptual display (referred to as “external display”), containing two randomly tilted bars. The new items appeared in the yet-unoccupied placeholders for 50 ms. For example, if the bars in the internal display appeared in the top-left and bottom-right placeholders, bars in the external display occurred the top-right and bottom-left placeholders. To make the perceptual discrimination challenging, items in the external array were masked. The masking array consisted of four overlayed titled bars presented in each of the placeholders for 100 ms. At each mask location the four tilted bars differed 45° from the neighbouring bars. The overall orientation of each mask was randomly drawn at 0 ms and 50 ms, thus creating the impression of a dynamic display. The masking parameters were set to equate the difficulty of internal and external reports, based on piloting efforts.

The offset of the dynamic mask was followed by a 100-ms delay, after which a visual response dial was displayed, always starting in a random position. The response dial surrounding the fixation cross was the same size as the placeholders and included markers along a circle that corresponded to the ends of a bar. The dial rotated leftwards when pressing the left button and rightwards when pressing the right button of the response device held in the dominant hand. The dial rotated with a speed 1/10° per millisecond. Participants had unlimited time to complete the orientation report. When the dial aligned with the orientation of the target item, participants

pressed either of the two buttons of the response device held in the non-dominant hand to confirm their response and to continue with the task.

After a 100-ms delay in which only the fixation cross stayed on the screen, participants received feedback in the form of a number ranging from 0 to 100, with 100 indicating a perfect report and 0 indicating that the adjusted orientation was perpendicular to the angle of the target item. Feedback was presented 0.7° above the central fixation cross for 300 ms. Trials were separated by an ITI randomly drawn between 750 and 1000 ms. Between blocks, participants were presented with their average reproduction accuracy in the previous block.

The experiment consisted of 512 trials divided across 8 blocks, each including 64 trials. Of the total trials, 50% (256) were single-cue trials and 50% (256) were double-cue trials. Participants were equally likely to report an internal or external item in single-cue trials (128 each). Double-cue trials were categorised based on the type of shift performed. That is, 50% (128) required a within-domain shift and the other 50% (128) a between-domain shift. Within-domain and between-domain shift trials were further split by the target domain. Ergo, the within-domain shift trials included 50% (64) internal-to-internal and 50% (64) external-to-external shifts, while the between-domain shift trials included 50% (64) external-to-internal and 50% (64) internal-to-external shifts. In each of the four shift conditions, participants were equally likely to shift their spatial attention between any of the eight possible cross-hemifield combinations of placeholders. Each block consisted of 50% (32) single- and 50% (32) double-cue trials.

To become familiarised with the procedure of the experiment, participants performed a practice block before the data collection started. The whole experiment lasted approximately 65 minutes.

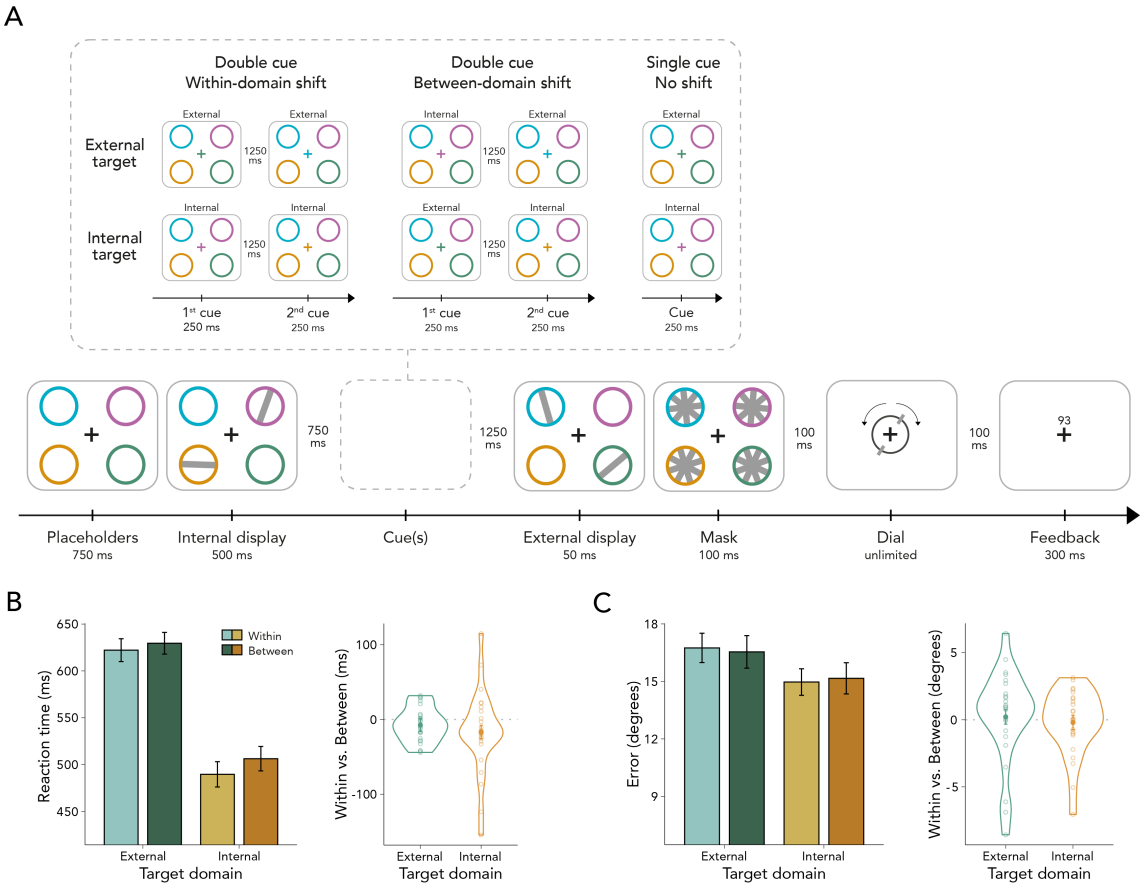

Figure 5.1. Task schematic and behavioural results. (A) Tilted bars were presented across a working-memory array, encoded before the cues (i.e., internal display), and a perceptual array, appearing after the cues (i.e., external display). In double-cue (i.e., shift) trials, two cues appeared between the presentation of the internal and external display. At the time of the second cue, attention needed to be reoriented either to an item of the same domain (i.e., within-domain shift) or to an item in the other domain (i.e., between-domain shift). To ensure participants relied on the first cue, 50% single-cue trials were included in which only one cue was presented, and thus no attentional reallocation was required. **(B)** Left panel: Average RTs as a function of target domain and shift type. RTs were faster for internal than external targets and for within- than between-domain shifts. The interaction between target domain × shift type was not significant. Right panel: Shift cost in RTs defined as the difference between shift types (within minus between) for each target domain. **(C)** Left panel: same as left panel in (B) but for reproduction errors. There were neither a significant main effect nor a significant interaction effect. Right panel: same as right panel in (B) but for reproduction errors. Error bars represent ±1 *SEM*. Dots represent individual participants.

5.3.3 Apparatus and Data Acquisition

I obtained whole-head MEG recordings in a magnetically shielded room using a 306-sensor MEG system (204 first-order planar gradiometers, 102 magnetometers; TRIUX neo, MEGIN OY, Espoo, Finland) at the Oxford Centre for Human Brain Activity. The MEG signal was sampled at 1000 Hz, with a high-pass filter at 0.1 Hz and a low-pass anti-aliasing filter at 330 Hz. Each experimental block was recorded separately, resulting in eight runs for most participants. However, due to

technical issues, two individuals had fewer recordings, with six and seven recordings obtained, respectively.

Prior to data acquisition, I digitised the individual head shape with an electromagnetic position and orientation monitoring system (FASTRAK, Polhemus, Colchester Vermont, USA). Head-shape digitalisation included fiducial landmarks (i.e., nasion, right and left preauricular points) and about 250 additional points evenly distributed across the scalp. Five Head Position Indicator (HPI) coils were placed on the mastoid bones and forehead to keep track of the participant's head position inside the dewar through electromagnetic induction before and after each recording block.

Participants performed the task in a seated position. Participants were instructed to avoid head, body, and limb movements and to keep strict eye fixation during the experimental blocks. In addition to electrocardiography (ECG) and horizontal and vertical EOGs, eye positions were continuously monitored by an MEG-compatible eye-tracking device with a sampling frequency of 1000 Hz (Eyelink 1000, SR-Research Ltd., Ottawa, Ontario, Canada). At the beginning of each experimental session, participants performed an eye-tracking calibration task to verify their gaze position on the screen. Calibration was repeated if drift was noticed during the experimental session.

I used the PsychoPy package (version 2021.2.3; Peirce et al., 2019) in Python for stimulus generation and stimulus delivery. The stimuli were projected on a translucent whiteboard using a DLP LED projector (ProPixx, VPixx Technologies Inc., Saint-Bruno, Quebec, Canada) at a 120-Hz refresh rate. The whiteboard was located at 120-cm distance from the participant, and it provided a projection area of 55×31-cm and a 1920×1080-pixel resolution. A bimanual fibre-optic response device was used to collect manual responses.

5.3.4 Behavioural Analysis

Behavioural data were analysed in R Studio (Posit team, 2022). During pre-processing, which followed the pipeline detailed in Chapter 4, trials were removed when RTs exceeded 5000 ms (calculated from dial to response onset). Next, I removed trials for which the remaining RTs were 2.5 *SD* above the individual mean across all conditions. Moreover, I removed datasets with average reproduction errors (calculated by averaging the absolute difference between the angle of the target item and the reported angle) equal to or higher than 45° in any of the conditions (n = 1). After these exclusion steps, datasets from 25 participants remained in the behavioural analysis, with an average of 16.64 ± 3.91 (M ± SD) trials removed. Additionally, I removed the bad MEG trials from the behavioural data analysis.

The conditions of interest were target domain (having to report an external vs. internal item) and shift type (within vs. between). When comparing more than two means, I applied a repeated-measures ANOVA and reported η^2_G as a measure of effect size. When evaluating only two means, I applied a paired samples *t*-test and report Cohen's *d* as a measure of effect size.

5.3.5 MEG Pre-Processing

MEG data were analysed in Python using MNE-Python (version 1.4.2; Gramfort et al., 2013) and OSL (version 0.2.0; Quinn et al., 2022) combined with custom code. For each individual run, I applied the Elekta MaxFilter (version 2.2) implementation of Spatiotemporal Signal Source Separation to remove external noise (TSSS; Taulu & Simola, 2006), to continuously compensate for head movements occurring during the run, and to realign all runs per participant to a participant-specific head position. The reference head position for each participant was determined by minimising the head-position distance between runs. Next, data were downsampled to 250 Hz and low-pass filtered at 50 Hz. Importantly, no high-pass filter was applied since high-pass filtering can temporally displace multivariate information (van Driel et al., 2021). Bad sensors were identified and interpolated using an ESD test (Rosner, 1983) at a 0.05

significance threshold. On average, 1.62 ± 3.11 bad gradiometers and 1.42 ± 2.02 bad magnetometers were identified and interpolated per run. I applied an ICA to regress out eye-movement- and cardiac-related activity. Note that the maximum number of ICA components to be extracted was determined by the residual degrees of freedom after TSSS rank reduction during Maxfiltering (Woolrich et al., 2011). Data were epoched from -500 to 3000 ms relative to the first cue onset without applying baseline correction (van Driel et al., 2021). Next, I combined all individual runs for each participant by concatenating the data. Additionally, I remapped magnetometers onto gradiometers, enabling all subsequent analyses to be conducted using a unified sensor type. Lastly, I employed an ESD trial rejection approach to detect and exclude trials with high variance. Per participant, 5.24 ± 6.96 (*M* ± *SD*) trials were removed. In addition, I removed the trials that were excluded in the behavioural data analysis.

5.3.6 Time-Resolved Multivariate Pattern Analysis

I performed a time-resolved multivariate pattern analysis on the broadband MEG data using the Scikit-learn library (version 1.3.0; Pedregosa et al., 2018) and built-in decoding functions provided by MNE-Python (version 1.4.2; Gramfort et al., 2013). The decoding pipeline closely followed the approach outlined in previous literature (Hajonides et al., 2021).

I extracted patterns of brain activity separately for each participant on a time-point-by-time-point basis (mne.decoding.SlidingEstimator). First, features (i.e., data at each time point for each sensor) were centred and scaled based on data from all trials (sklearn.preprocessing.StandardScaler). The true spatial dimensionality of the MEG data after the Maxfilter pre-processing algorithm is around 70, rather than the 204 gradiometer sensors (Woolrich et al., 2011). Therefore, to reduce the number of redundant features for decoding, I performed Principal Component Analysis (PCA) on the standardised data while maintaining 99% of the variance (sklearn.decomposition.PCA). Training and testing were done on the same data using a 10-fold stratified cross-validation procedure (sklearn.model_selection.StratifiedKFold).

First, trials were randomised in order, and divided into ten equal-sized folds. Next, a leave-one-out procedure was used on the 10 folds, such that the classifier was trained on 9 folds and tested on the remaining fold. This procedure was repeated 10 times until each fold was used once for testing. A Linear Discriminant Analysis (LDA; sklearn.discriminant_analysis.LinearDisciminantAnalysis) classifier was utilised to train based on the provided training data and labels. LDA is a supervised dimensionality reduction technique that aims to find a linear combination of features that maximises class separability and has been specifically recommended for neuroscientific decoding studies (Grootswagers et al., 2017). As a scoring metric, I used the Area Under the Curve of the Receiver Operating Characteristic Curve (ROC-AUC). This scoring metric takes into account the trade-off between true and false positive rates and is considered a sensitive, nonparametric, and criterion-free measure of classification (Hand & Till, 2001; Thölke et al., 2023). A ROC-AUC value of 0.5 means chance-level classification performance. The obtained classifier scores were averaged over all 10 folds, yielding a single decoding estimate per participant, time point, and condition. Lastly, decoding time courses were smoothed with a one-dimensional Gaussian filter with a standard deviation of 10.

In addition, I used temporal-generalisation methods (King & Dehaene, 2014) to examine whether there are recurrent neural activation patterns over time, that is, whether an activation pattern observed at one time point is observed again at a later time point. This approach results in time-point-by-time-point decoding matrices in which each cell corresponds to a classification accuracy at a unique training-test-time combination. Generalisation off the diagonal in the temporal-generalisation matrix is driven by an overlap in activation patterns at two different time points.

To test which sensors contributed most to the classifier likelihoods, I ran a sensor-space searchlight decoding analysis (Grootswagers et al., 2017; Kaiser et al., 2016; Teichmann et al., 2022). For this analysis, I conducted the same analysis as outlined above across small clusters of 15 neighbouring MEG sensors, resulting in decoding scores for each sensor and time point. These

decoding scores were averaged into time clusters of 300 ms each, resulting in 10 topographical maps between 0 and 3000 ms. These topographical plots show whether decoding results were primarily driven by specific sensor clusters. I did not test any specific hypotheses regarding the spatial distributions of the effects and therefore did not run any statistical tests on the results.

Even though stratified cross-validation is highly sensitive to class imbalances (Thölke et al., 2023), I additionally applied between-class balancing using under-sampling, such that the number of trials were equal across classes. This further ensured that during training the classifier would not only achieve a high decoding performance merely by systematically and blindly voting for the majority class (Y. Sun et al., 2009). Note that as the design was balanced in terms of trial counts, the between-class balancing was employed to account for small imbalances due to trial-rejection procedures used for data cleaning.

5.3.7 Relationship Between Decoding and Behaviour

To investigate the relationship between the observed decoding of within-domain versus between-domain shifts of attention and the behavioural shift costs, I employed the following approach. Firstly, I identified trials with low and high classifier performance for each participant by computing time-point-by-time-point decoding scores and averaging them across the interval of interest. Subsequently, I performed a median split based on these average decoding scores, creating two groups of trials: one with low decoding scores and the other one with high decoding scores. I then included the obtained median-split factor in the behavioural analysis when comparing within- to between-domain shifts.

5.3.8 Control Analysis: Eye Movements

Involuntary eye movements are a common problematic covariate in studies using multivariate approaches (Mostert et al., 2018; Quax et al., 2019). To test for an influence of gaze on the MEG findings, I ran the same MEG decoding analyses on the eye-tracking data. I acquired a total of 23

eye-tracking datasets, consisting of 15 binocular recordings and 8 monocular recordings. For the binocular recordings, I merged the x- and y-coordinates from both eyes to create a single metric. The data were then epoched from -500 ms to 3000 ms relative to the first-cue onset and were baseline-corrected prior to the cue onset (-500 to 0 ms). Next, I trained the classifier to distinguish between different conditions at each time point using the x- and y-coordinates of the eye-tracking data. Instead of an LDA classifier, I used a Support Vector Machine classifier (SVC; sklearn.svm.SVC) with Radial Basis Function (RBF) kernel. The reason behind opting for SVC with RBF kernel is the likelihood that the distribution of x- and y-coordinates of gaze is not linearly separable. In all conditions, participants were instructed to maintain fixation, making it challenging to distinguish between different classes or conditions using a simple linear boundary. The RBF kernel is a powerful tool for transforming non-linearly separable data into a higher-dimensional space, making it possible to separate the data using a hyperplane (Hsu et al., 2003). This method is particularly effective when dealing with datasets containing a smaller number of features than of samples.

By computing mutual information scores over correctly classified MEG trials on a time-point-by-time-point basis, I assessed how much information was shared between the MEG decoding results and eye-movement decoding results (Quax et al., 2019). Mutual information quantifies the extent to which knowledge of one variable reduces uncertainty about another variable. Higher mutual information values indicate a stronger dependence between two variables, suggesting that changes in one variable are informative about the other.

5.3.9 Time-Frequency Analysis

I generated time-frequency representations of power by convolving the data with Morlet wavelets across the frequency range of 3 to 40 Hz. For each frequency, I used a fixed 300-ms time window such that the number of cycles changed with the frequency. The power time series in the planar gradiometer pairs were then combined (i.e., root mean square), resulting in a 102-sensor planar

gradiometer map. For the comparison between contralateral and ipsilateral power, I computed activity following left- and right-directing cues, separately for left sensor clusters (MEG1632+1633, MEG1642+1643, MEG1912+1913, MEG1922+1923, MEG1942+1943, MEG2042+2043) and right sensor clusters (MEG2032+2033, MEG2312+2313, MEG2322+2323, MEG2342+2343, MEG2432+2433, MEG2442+2443), and subsequently pooled the contralateral versus ipsilateral contrasts between them. I expressed the contra- versus ipsilateral contrast as a normalised difference (i.e., [(contra-ipsi)/(contra + ipsi)] × 100). Sensors were chosen based on prior MEG studies that have consistently implicated the same set of posterior sensors for capturing neural activity in response to lateralised visual stimuli (Heideman, Rohenkohl, et al., 2018; Heideman, van Ede, et al., 2018). While I selected these sensor clusters based on independent data, I could confirm their validity and appropriateness in the current data (Figure 5.6J, K, and L show an overlay of these sensor clusters on the relevant topographies).

To extract time courses of alpha lateralisation, I averaged the contra- versus ipsilateral activity across the predefined alpha band (i.e., 8 to 12 Hz) in the visual sensors. Topographies of the lateralised visual activity were obtained by contrasting trials in which cues indicated visual content in the right versus left hemifield. To depict the topographies associated with relevant condition comparisons, I focussed on the alpha frequency bands between 0 and 3000 ms in predefined steps of 500 ms.

5.3.10 Statistical Analysis

Statistical evaluation of MEG and eye-tracking data was conducted by employing a cluster-based non-parametric permutation approach with 1024 permutations and a cluster alpha of 0.05. This approach circumvents the multiple-comparisons problem by evaluating temporal clusters observed in the original (non-permuted) data against a single permutation distribution of the largest temporal cluster found after each permutation of the condition labels (Maris & Oostenveld, 2007).

In a supplementary analysis, I additionally utilised a jack-knife approach to identify differences in decoding onset times. For the jack-knife approach, I computed decoding time-course averages by excluding each of the *n* datasets once, which allowed me to obtain *n* leave-one-out grand averages. Next, I identified the time point at which the decoding time course reached 50% of its peak. Jack-knifed standard errors, *t*-values, and *p*-values were adjusted according to J. Miller et al. (1998).

5.4 Results

5.4.1 RTs and Errors: Within-Domain Shifts vs. Between-Domain Shifts

In the following analyses, I focussed on double-cue trials, as these trials uniquely enabled the comparison of within- and between-domain shifts. **Figure 5.1B** depicts the average RTs relative to dial onset as a function of the target domain and shift type. Similar to the results in Chapter 4, I found RTs to be increased when reorienting attention between domains compared to reorienting within the same domain ($F_{(1,24)} = 4.581$, $p = 0.043$, $\eta^2_G = 0.002$). This between-domain shift cost replicates previous behavioural findings (Dark, 1990; Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019; Weber et al., 1986). Moreover, RTs to internal targets were faster than to external targets ($F_{(1,24)} = 40.411$, $p < 0.001$, $\eta^2_G = 0.192$), likely owing to the fact that behavioural responses to the former could be prepared ahead of the response stage (Boettcher et al., 2021; van Ede, Chekroud, Stokes, et al., 2019). The interaction between target domain and shift type was not significant ($F_{(1,24)} = 0.534$, $p = 0.472$, $\eta^2_G < 0.001$), suggesting that the effect of shift type did not depend on target domain (i.e., whether participants reported an internal or an external item).

As shown in **Figure 5.1C**, reproduction accuracy was consistently high, yielding no significant main effect (target domain: $F_{(1,24)} = 1.697$, $p = 0.205$, $\eta^2_G = 0.019$; shift type: $F_{(1,24)} < 0.001$, $p = 0.990$, $\eta^2_G < 0.001$) or interaction ($F_{(1,24)} = 0.280$, $p = 0.601$, $\eta^2_G < 0.001$).

5.4.2 MEG Decoding: Within-Domain Shifts vs. Between-Domain Shifts

My main aim was to investigate whether and when within- and between-domain shifts evoked reliably different neural activity patterns. To this end, I applied a time-resolved cross-validation decoding approach to MEG data collected while participants performed the task. To target the primary question, I first trained a LDA classifier to distinguish external-to-external and internal-to-internal trials (i.e., within-domain shifts) from external-to-internal and internal-to-external trials (i.e., between-domain shifts) on a time-point-by-time-point basis. I used the ROC-AUC as the evaluation metric (Thölke et al., 2023). The orthogonal manipulation of the target domain, shift type, and cued location ensured that the decoding of within- versus between-domain shifts could be isolated, without being influenced by merely detecting the currently cued attentional domain or location.

I found reliable decoding of within- versus between-domain shifts. As expected, no decoding of reorienting attention within versus between domains was observed until after the second cue (Figure 5.2A). Within- versus between-shift decoding reached significance approximately 360 ms after the second-cue onset and persisted until external-display presentation (~1860-3000 ms, cluster $p < 0.001$).

Since involuntary eye movements may contaminate MEG recordings and lead to decodable non-brain signals (Mostert et al., 2018; Quax et al., 2019), I trained another classifier to distinguish within- versus between-domain shifts but this time using only the x- and y-coordinates of the eye-tracking data. I could not decode the shift type based on the eye-tracking signal (Supplementary Figure 8.17A). Moreover, I assessed whether there was mutual information between the eye-tracking and the MEG data. The mutual information is shown in Supplementary Figure 8.17B, with mutual information values of zero indicating that the MEG and eye-position signals are independent and larger values indicating a greater relationship between them. There was almost no mutual information between the eye-tracking and the MEG data when decoding within- versus between-domain shifts. However, I noted that two participants were outliers,

exhibiting stronger mutual information than the rest (Supplementary Figure 8.17C). I therefore repeated the MEG decoding analysis while excluding these participants. This revealed that the MEG-based decoding was not driven by participants with strong mutual information (Supplementary Figure 8.17D).

To assess if the pattern of neural activity underlying the classification performance was stable or evolved dynamically over time, I performed a generalisation-across-time analysis (Grootswagers et al., 2017; King & Dehaene, 2014). Significant points off the diagonal would indicate that the classifier, when trained on data from time point A, can generalise to data from time point B. In the present data, the evolving pattern of neural activity that distinguished within- from between-domain attention shifts was highly dynamic. The classifier only significantly generalised over brief amounts of time (Figure 5.2B; 1st cluster $p < 0.001$, 2nd cluster $p = 0.027$). This indicates that, although I can consistently distinguish between trials in which participants shifted attention within versus between domains, the neural code driving this distinction varied moment-to-moment. My primary decoding analysis included data from all MEG sensor locations as features in the classifier. A searchlight analysis (i.e., iterative decoding for subsets of MEG sensors) was used to characterise the scalp distribution of sensors sensitive to within- versus between-domain shifts. Figure 5.2C shows that the neural patterns distinguishing within- versus between-domain shifts were broadly distributed.

Together, these data thus reveal distinct patterns of brain activity associated with between-domain shifts as compared to within-domain shifts. Moreover, the results demonstrate how these patterns occur early, are highly dynamic, and are broadly distributed.

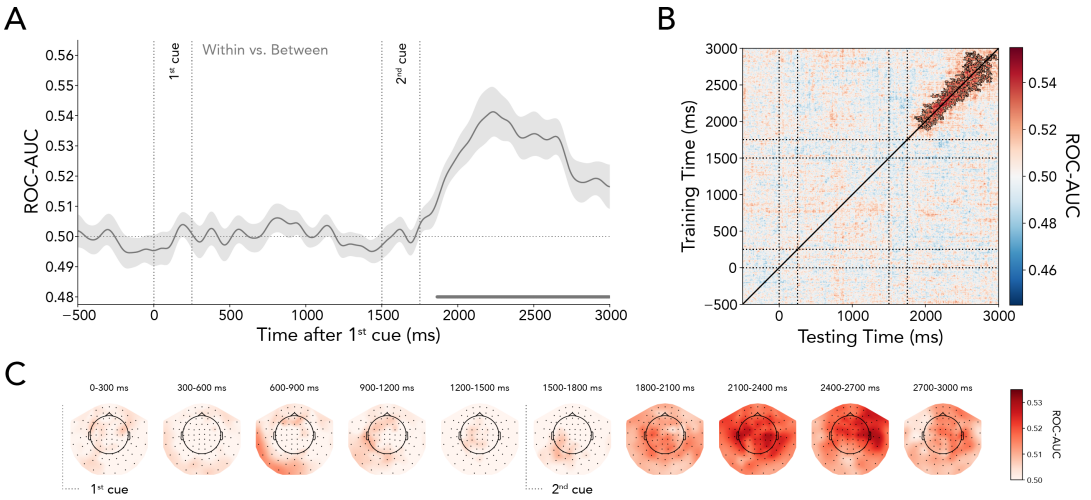

Figure 5.2. Time-resolved MEG decoding of within- versus between-domain shifts of attention. (A) Average classifier performance for discriminating within- from between-domain shifts plotted over time. Time course shows *M* ± *SEM* across participants. Horizontal lines indicate significant clusters. **(B)** Temporal-generalisation matrix reflecting classifier performance in discriminating the two shift types. The diagonal corresponds to the time course illustrated in (A). Black outlines indicate significant decoding-clusters. **(C)** The MEG topography averages in 300-ms steps showing which sensors most strongly contribute to the overall decoding performance.

5.4.3 MEG Decoding and RTs: Within-Domain Shifts vs. Between-Domain Shifts

To ascertain whether patterns of brain activity discriminating within- from between-domain shifts had functional consequences for behavioural performance, I obtained single-trial decoding scores and compared RTs as a function of low versus high decoding performance. I expected a higher behavioural between-domain shift cost when classifier performance was high as compared to low. To address this hypothesis, I categorised trials based on their average classification score within the interval between 360 to 1500 ms after the second cue (i.e., significant cluster interval in **Figure 5.2A**). **Figure 5.3** depicts RTs as a function of shift type (within vs. between) and median split (low shift-type decoding vs. high shift-type decoding). Planned comparisons revealed a significant between-domain shift cost only when comparing trials with high classifier performance ($t_{(24)}$ = 2.348, p = 0.027, d = 0.470). No difference occurred between shift types when considering trials with low classifier performance ($t_{(24)}$ =-0.049, p = 0. 961, d = 0.010). These results tentatively suggest that the ability to decode within- versus between-domain shifts of attention had consequences for later behavioural performance.

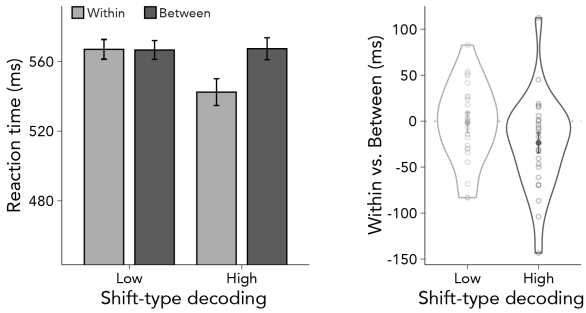

Figure 5.3. Decoding performance is behaviourally relevant. Left panel: RTs as a function of shift type (within vs. between) and median split (low vs. high). Right panel: Shift cost in RTs defined as the difference of within- and between-shift trials for each half of decoding strength (within minus between). Error bars represent ±1 *SEM*. Dots represent individual participants.

5.4.4 MEG Decoding: External Attention vs. Internal Attention

Having uncovered evidence for distinct patterns of brain activity during shifts of attention within versus between domains, I investigated whether and when the brain differentiated the attended domain before and after the shifts occurred. As I noted previously, owing to the experimental design, decoding of the currently attended domain (external vs. internal) was orthogonal to the above decoding of within- versus between-domain shifts.

I started the investigation by decoding external versus internal attentional states based on the labels of the first cue. To this end, I trained a classifier to distinguish external first-cue trials (i.e., external-to-external, external-to-internal) from internal first-cue trials (i.e., internal-to-internal, internal-to-external). Figure 5.4A shows successful decoding of the attentional domain from approximately 310 ms after first-cue onset (cluster $p < 0.001$). Interestingly, the significant decoding period extended into the second-cue phase, revealing that information about the initially cued attentional state persisted throughout the delay following the second cue.

Temporal-generalisation analysis unveiled an off-diagonal shape of generalisation following the first cue, indicating temporal stability of the neural pattern evoked by the initially cued domain (Figure 5.4B; cluster $p < 0.001$). In contrast, after second cue onset, decoding performance occurred mainly along the diagonal, with a narrower spread to neighbouring time points. In addition, I found off-diagonal below-chance decoding (i.e., the classifier systematically

guessed incorrectly; 1st cluster $p = 0.024$, 2nd cluster $p = 0.025$). The searchlight decoding analysis demonstrated that the contribution of sensors to the external versus internal decoding performance was broadly distributed, with peaks occurring between 600 and 1200 ms in the posterior and left-central sensors (Figure 5.4C).

To determine whether decoding of the first-cued attentional domain was predominantly influenced by a particular trial type, I employed two additional classifiers: one to distinguish between external-to-external and internal-to-external shifts, and another to distinguish between internal-to-internal and external-to-internal shifts. Supplementary Figure 8.18A shows that classifier performance did not differ between these two decoding analyses, thus resembling the general external- versus internal-domain decoding time course shown in Figure 5.4A. In addition, I again took precautions to ensure that eye movements did not drive the external versus internal first-cue MEG decoding results. The strength of the decoding based on gaze position was relatively weak and far from sufficient to explain the effects when decoding brain activity (Supplementary Figure 8.17E). Furthermore, exclusion of participants who exhibited the highest mutual information between eye-position and MEG decoding did not alter the MEG decoding results (Supplementary Figure 8.17F-H).

Next, I aimed to decode the domain indicated by the second cue. To this end, I trained a classifier to distinguish between external second-cue trials (i.e., external-to-external, internal-to-external) versus internal second-cue trials (i.e., internal-to-internal, external-to-internal). The results revealed significant decoding of the second-cued attentional domain from approximately 280 ms after second-cue onset (i.e., 1780 ms after first-cue onset; cluster $p < 0.001$; Figure 5.4D). As was intended in the task design via orthogonally manipulating the first- and second-cued attentional domain, the second-cued domain could not be decoded prior to the onset of the second cue. The temporal-generalisation matrix in Figure 5.4E shows a temporal spread, implying some stability in the activation pattern (cluster $p < 0.001$). Similar to the first-cue analysis,

sensor-searchlight analysis indicated a distributed pattern of sensor contribution to the decoding external versus internal attentional domains (Figure 5.4F).

To ensure that the decoding performance of the second-cued domain was not driven by the contrast between specific trial types, I trained one classifier to discriminate between external-to-external and external-to-internal shifts and another one to discriminate between internal-to-internal and internal-to-external shifts. While both decoding time courses reached significance following the second cue, the internal-to-internal versus internal-to-external contrast reached significance earlier than the external-to-external and external-to-internal contrast (Supplementary Figure 8.18B). These findings revealed earlier domain differentiation when initially cued about the internal domain compared to when previously cued about an external item. In line with the first-cue analysis, I ensured that the second-cue MEG decoding results were not dependent on eye movements (Supplementary Figure 8.17I-L).

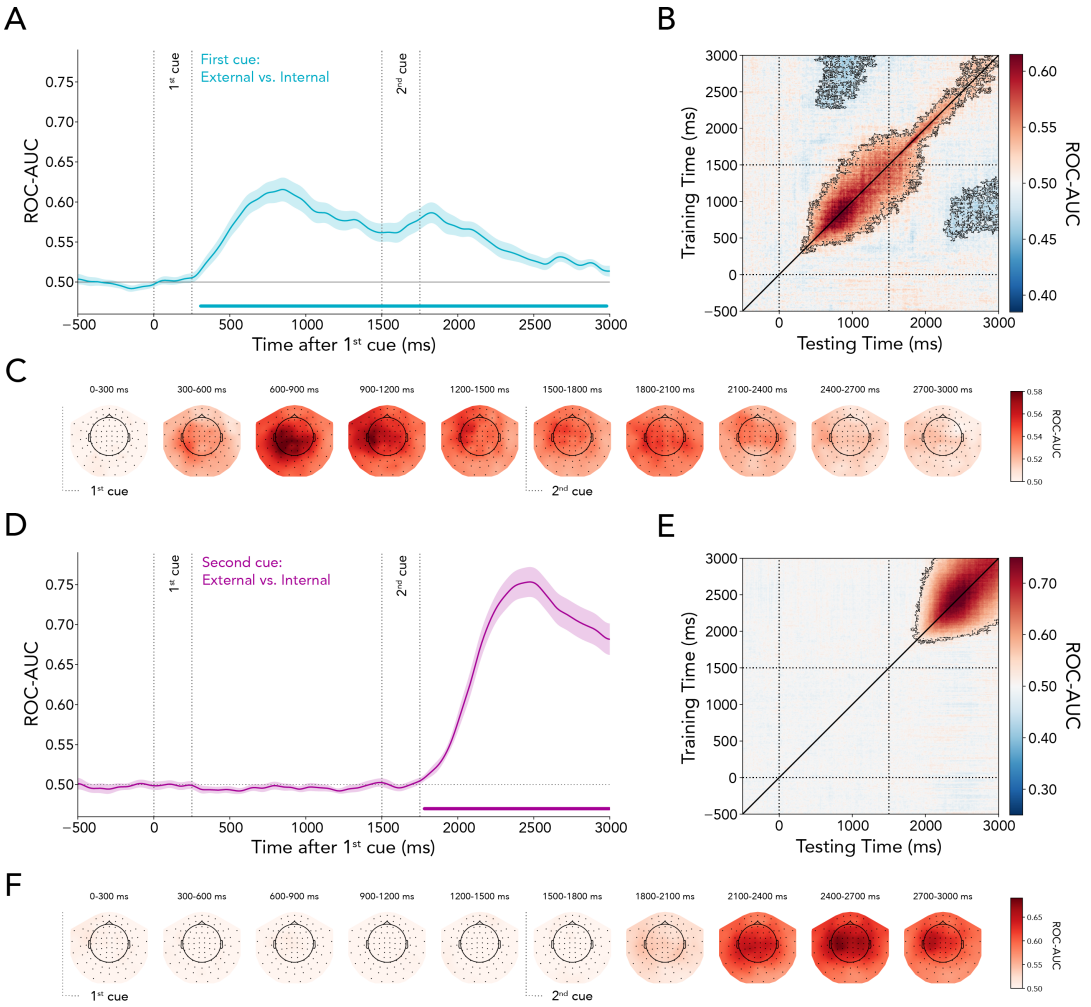

Figure 5.4. Time-resolved MEG decoding of external versus internal attention. (A) Average classifier performance for distinguishing external from internal first cues plotted over time. Time course shows *M ± SEM* across participants. Horizontal lines indicate significant clusters. **(B)** Temporal-generalisation matrix reflecting classifier performance in discriminating the first-cued attentional domain. The diagonal corresponds to (A). Black outline indicates significant decoding clusters. **(C)** The MEG topography averages in 300-ms steps showing which sensors most strongly contribute to the overall decoding performance in discriminating external- and internal first cues. **(D)** Same as (A) but based on second-cue labels. **(E)** Same as (B) but based on second-cue labels. **(F)** Same as (C) but based on second-cue labels.

5.4.5 MEG Decoding: Shift Type and Attentional Domain

The current results demonstrate that attentional domains indicated by the first and second cue could be successfully decoded. Interestingly, the significant decoding of the first-cued attentional state persisted even after the second domain was cued. Supplementary analyses revealed that during the period following the second cue, the decodability of the continued information concerning the first cue was not contingent on the quality of processing of the second cue (**Supplementary Figure 8.19**). Hence, this finding suggests that lingering processing related to the

attentional domain of the first cue and instantiating the attentional domain associated with the second cue proceeded in tandem. In other words, information regarding the past and current attentional state must have co-existed after the second cue.

Previous task-switching research has argued that switch costs can arise from lingering activity associated with the previous, but now task-irrelevant information (Wylie & Allport, 2000; Yeung et al., 2006). I therefore tested whether a similar interaction might arise between lingering activity related to the first-cued domain and the ability to decode within- versus between-domain shifts. To this end, I median-split the data into trials with bad versus good first-cue domain decoding following the onset of the second cue (i.e., 1500-3000 ms). Next, I trained a classifier to distinguish within- from between-domain shifts based on these two data subsets. The within- versus between-domain shift-type decoding only yielded a significant cluster in trials with high decodability of the lingering first-cued domain (**Figure 5.5A**; ~1910-2890 ms, cluster $p < 0.001$). This was corroborated by a direct comparison, showing that trials with stronger lingering first-cue decodability had significantly more successful within- versus between-domain decoding (~2330-2870 ms, cluster $p = 0.005$). **Figure 5.5B** displays the within- versus between-shift decoding topographies for low and high lingering first-cue decoding. Thus, the present data suggest that the lingering first-cued attentional state is linked to the decodability of the performed shift type.

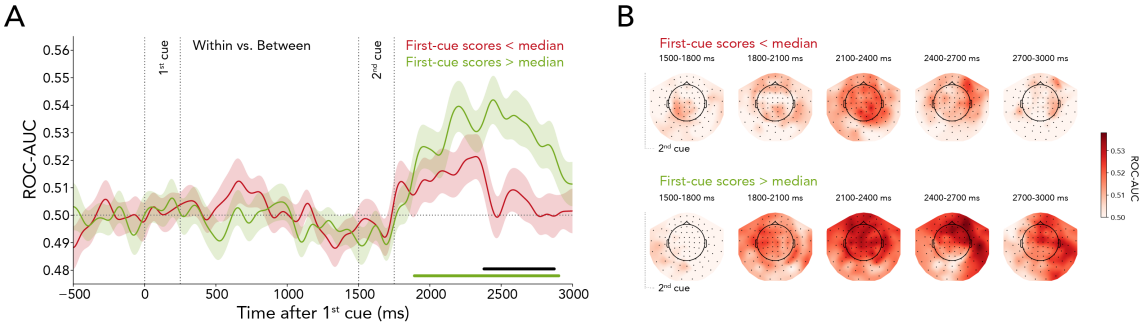

Figure 5.5. Relationship between lingering first-cued attentional domain and performed shift type. (A) Time-resolved average classifier performance for discriminating within- from between-domain shifts as a function of bad (red) versus good (green) decoding of the lingering first-cued attention state. Time courses show *M* ± *SEM* across participants. Coloured horizontal lines indicate significant clusters, black line denotes difference between time courses. **(B)** MEG topography averages in 300-ms steps showing which sensors most strongly contribute to the overall decoding performance in discriminating within- and between-domain shifts. Top panel: trials with low decodability of the lingering first-cue domain. Bottom panel: trials with high decodability of the lingering first-cue domain.

5.4.6 MEG Time-Frequency: External Attention vs. Internal Attention

Thus far, I have considered the concept of shifting from two essential perspectives. Firstly, I contrasted brain activity when transitioning between domains in comparison to shifting attention within a single domain. Secondly, I examined whether brain activity distinguished between the attentional domain following the first and the second cue. The task design additionally enabled the investigation of attentional shifts from a third perspective, involving the spatial orienting of attention between locations. I intentionally designed the task to guarantee that shifts occurred between items positioned in distinct hemifields, enabling the study of activity lateralisation according to the location to which attention was directed.

A canonical marker for spatial attention shifts is the lateralisation of posterior 8-12 Hz alpha activity. Contra- versus ipsilateral attenuation of alpha oscillations in sensory areas has been observed during shifts of spatial attention in anticipation of relevant sensory stimuli as well as during selection of relevant working-memory content (Gould et al., 2011; Jensen & Mazaheri, 2010; Thut et al., 2006; van Ede, 2018; van Ede, Chekroud, Stokes, et al., 2019; Wallis et al., 2015; Worden et al., 2000). Here, I therefore focussed on alpha lateralisation to delineate spatial shifts of attention after external and internal first and second cues. These additional analyses not only complement the primary findings derived from the decoding analyses, but also relate the present work to the existing neuroscience literature on external and internal spatial shifts of attention.

Figure 5.6A and D show the time- and frequency-resolved lateralisation of spectral power in contra- versus ipsilateral visual sensors for external and internal first cues, respectively. At the time of the second cue, participants were always required to shift attention from one hemifield to the other. Consequently, relative suppression of contralateral activity in response to the first cue was followed by relative enhancement of contralateral activity (i.e., a reversal of lateralisation) after second-cue onset (external first cued: 1st cluster $p < 0.001$, 2nd cluster $p < 0.001$; internal first cued: 1st cluster $p < 0.001$, 2nd cluster $p = 0.004$).

To chart the time course of alpha lateralisation, I averaged the time-frequency data along the classical alpha band (8-12 Hz; Figure 5.6G). I confirmed the typical contra- versus ipsilateral alpha attenuation in visual sensors following both external and internal first cues. Clusters of lateralised alpha attenuation alternated between hemispheres after the first (external first cued: ~330-1110 ms, cluster $p < 0.001$; internal first cued: ~430-1730 ms, cluster $p < 0.001$) and second cue (external first cued: ~1910-3000 ms, cluster $p < 0.001$; internal first cued: ~1850-2980 ms, cluster $p = 0.004$). The time course of the initial contra- versus ipsilateral alpha attenuation was similar for both external- and internal first-cue trials. However, in line with prior research (Mok et al., 2016; Wallis et al., 2015), the duration of the alpha modulation differed between first-cued domains. Relative suppression of contralateral activity was more transient following internally directing first cues and more sustained following externally directing first cues, as indicated by a significant cluster from approximately 1030 to 1740 ms after first-cue onset (cluster $p < 0.001$). A topographic inspection of external and internal first-cued trials confirmed a clear posterior topography for these effects in alpha lateralisation (Figure 5.6J upper and lower panel, respectively).

The same time-frequency analysis was conducted based on the attentional domain associated with the second cue. Figure 5.6B and Figure 5.6E illustrate that the relative suppression of contralateral activity in response to the second cue was preceded by a relative enhancement of contralateral activity (i.e., a reversed lateralisation) after first-cue onset (external second cued: 1st cluster $p < 0.001$, 2nd cluster $p < 0.001$; internal second cued: 1st cluster $p = 0.002$, 2nd cluster $p = 0.002$). This pattern of results was mirrored in the alpha-band time courses in Figure 5.6H (external second cued: ~360-1400 ms and ~1920-3000 ms, 1st cluster $p < 0.001$, 2nd cluster $p = 0.003$; internal second cued: ~360-1300 ms and ~1920-2940 ms, 1st cluster $p = 0.002$, 2nd cluster $p < 0.001$). Interestingly, there was no difference between external and internal second-cue trials. Topographies showing the difference between trials with right- versus left-cued items confirmed

that these visual signatures were prominent in the corresponding right and left visual sensors (Figure 5.6K; top panel: external cue, bottom panel: internal cue).

Together, these complementary analyses of alpha lateralisation reveal that spatial attention lingered longer after the first cue when directed externally versus internally, while reorienting after the second cue proceeded similarly regardless of whether spatially reorienting to an external or internal item.

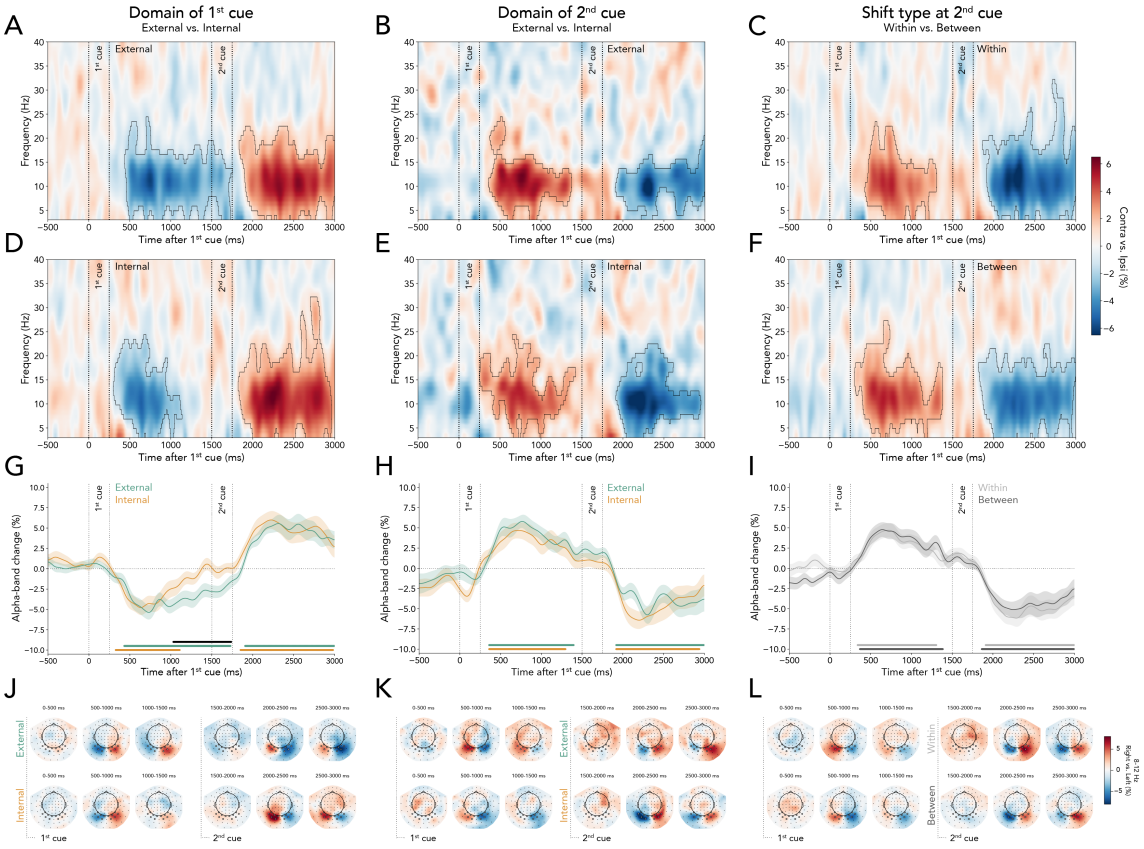

Figure 5.6. Lateralised alpha activity relative to the location of the first-cued and second-cued item. (A) Time-frequency plots of neural lateralisation following external first cues. Plots show the difference in spectral power contra- versus ipsilateral to the cued side, extracted over the predefined left and right sensor clusters. Black outlines indicate significant clusters. **(B)** Same as (A) but for external second cues. **(C)** Same as (A) but for second cues indicating within-domain shifts **(D)** Same as (A) but for internal first cues. **(e)** Same as (a) but for internal second cues. **(F)** Same as (a) but for second cues indicating between-domain shifts. **(G)** Time courses of neural lateralisation in the 8-12 Hz alpha band relative to the first cue. Time courses show *M ± SEM* across participants. Horizontal lines indicate significant clusters, black line denotes difference between time courses. **(H)** Same as (G) but relative to second cues. There was no significant difference between time courses. **(I)** Same as (G) but relative to second cues indicating within- and between-domain shifts. There was no significant difference between time courses. **(J)** Topographies of power differences following right- versus left-directing external first cues in the 8-12 Hz alpha band. Top panel shows topographies for external first cues, bottom panel for internal first cues. **(K)** Same as (J) but for second cues. **(L)** Same as (J) but for second cues indicating within- and between-domain shifts. Top panel shows topographies for within-domain shifts, bottom panel for between-domain shifts.

5.4.7 MEG Time-Frequency: Within-Domain Shifts vs. Between-Domain Shifts

Given the effectiveness of alpha attenuation in tracking the spatial orienting and reorienting of attention, I finally turn to the important question of whether spatial orienting following the second cue would be delayed when shifting attention between as compared to within domains. If shifting between domains engages additional processes prior to implementing a spatial shift, I would expect to see a delay in the onset of spatial orienting when shifting attention between domains.

Figure 5.6C and F depict alpha lateralisation relative to the second cue, split by whether the second cue indicated a within- or a between-domain shift (within-domain shift: 1st cluster p = 0.003, 2nd cluster $p < 0.001$; between-domain shift: 1st cluster $p < 0.001$, 2nd cluster $p < 0.001$). As evident from the alpha-band time courses (Figure 5.6I; within-domain shift: 1st cluster ~340-1300 ms, cluster p = 0.002 2nd cluster ~1910-3000 ms, cluster $p < 0.001$; between-domain shift: 1st cluster ~370-1380 ms, cluster $p < 0.001$, 2nd cluster ~1860-3000 ms, cluster $p < 0.001$), I observed no difference between shift types in alpha lateralisation. Again, topographical maps confirmed the distribution of the effects over corresponding right and left visual sensors (Figure 5.6L, top panel: within-domain shift, bottom panel: between-domain shift). These findings suggest that the initiation of the spatial shift remained unaffected, despite the decoding analyses revealing distinguishable neural processing engaged during within- versus between-domain shifts. The results offer no evidence for an additional stage invoked by between-domain shifts that must be resolved before attention can be shifted in space.

5.5 Discussion

This study uncovers direct neurophysiological evidence for distinct brain activity when shifting between perception and working-memory contents as compared to when shifting between contents within either domain. The utilisation of a carefully balanced task design – which orthogonally manipulated shift type and target domain while also necessitating spatial shifts of

attention – provided valuable insights into the dynamic neural bases that govern shifts between perception and working memory.

In principle, multiple scenarios could explain the behavioural costs and neural differences between shift types, and these need not be mutually exclusive. One obvious scenario is the presence of an additional supervisory control process that acts upon the neural systems that control between-domain attentional shifts (Burgess et al., 2006, 2007; Gilbert et al., 2005; Verschooren et al., 2021). This supraordinate control mechanism could be invoked when initiating a shift between perception and working memory but not when shifting between contents within either domain. Although it is most intuitive to consider a supraordinate control mechanism occurring when attentional shifts between domains are initiated, it is also possible to imagine supraordinate control mechanisms operating at other stages of attention deployment. Attention comprises several functions to guide adaptive behaviour, such as disengaging, orienting, selecting, prioritising, and output gating (Myers et al., 2017; van Ede & Nobre, 2023). Accordingly, additional supraordinate mechanisms could operate at any or multiple points along the way. An alternative type of scenario is one that invokes no additional control process(es). Instead, the systems for controlling external versus internal attention may coexist and compete for access to the relevant sensory and/or motor representations. In this scenario, the additional cost of shifting between perception and memory (and vice versa) would arise from resolving a stronger degree of competition between as compared to within domains. As I outline below, the current findings offer important initial clues regarding the factors that distinguish within- and between-domain shifts.

5.5.1 Distinct Neural Patterns of Within- and Between-Domain Shifts

The present results unequivocally demonstrated distinct neural processes during within- and between-domain shifts. By tracking neural dynamics, I showed that differences between shift types occur early following the second cue (**Figure 5.2A**), suggesting that they unlikely emerge

exclusively at the response stage. In line with that, I showed that this differential cue-related activity is behaviourally relevant, predicting later between-domain shift costs in RTs (Figure 5.3). I cautiously argue, therefore, that previously observed between-domain shift costs in performance (Calzolari et al., 2022; Dark, 1990; Gilbert et al., 2005; Hautekiet et al., 2023; Verschooren et al., 2020; Verschooren, Liefooghe, et al., 2019) are unlikely explained solely by interference at a late response-level stage, but rather emerge already during the attentional shifting process itself.

Interestingly, by analysing alpha-band lateralisation, I found no evidence indicating differences in visual activity when orienting spatial attention within versus between domains (Figure 5.6C, F, I, L). This observation challenges the intuitive scenario proposing that a supraordinate control mechanism is invoked when initiating an attentional shift that crosses domains. Such a scenario would predict a delayed shift of spatial attention. With regard to the spatial orienting functions of attention, as reflected in alpha lateralisation, I found no evidence for this scenario. These alpha-lateralisation findings complement a similar observation made in Chapter 4.3, where between-domain shifts did not modulate gaze biases. Both of these results reinforce the suggestion that no additional, time-consuming gating step occurs during spatial shifts between domains.

Neural patterns differentiating within- versus between-domain shifts evolved relatively dynamically and were broadly distributed (Figure 5.2B). Differences between shift types are thus not solely attributable to a single tonic state that differs between within- and between-domains shifts, but likely result from the dynamic interplay of multiple states. Moreover, the distributed nature of the within- versus between-shift decoding suggests that large-scale networks may be involved (Figure 5.2C). Future brain-imaging investigations using fMRI may substantiate this broad distribution and/or pinpoint the precise spatial underpinnings of the findings reported here.

5.5.2 Distinct Neural Patterns of External and Internal Attention

It is useful to consider that each within- and between-domain shift is a consequence of the attended domains before and after the shift. Therefore, to understand the nature of the two shift types, it is also essential to investigate the neural dynamics evoked by the first- and second-cued attentional domains. While some studies have emphasised the shared neural systems and dynamics between selection mechanisms in external and internal attention (Panichello & Buschman, 2021; Y. Zhou et al., 2022), others have also highlighted important differences (Chota et al., 2023; Esterman et al., 2009; Myers et al., 2017; Nobre & Stokes, 2019; Tamber-Rosenau et al., 2011; van Ede & Nobre, 2023; Vortmann et al., 2019). Underscoring the differences, the current data demonstrated successful decoding of external versus internal attention (Figure 5.4). Moreover, I show that decoding of the first-cued domain persisted even after the presentation of the second cue.

The carry-over decoding effect of the first-cued domain occurred regardless of the domain indicated by the second cue. This finding highlights the fact that, depending on the nature of the preceding attentional domain, the brain exists in a different attentional state after the presentation of the second cue (i.e., the attentional state invoked by the second cue did not overrule or wipe out the lingering state invoked by the first cue). Multiple processes could have contributed to the lingering decoding of the first-cued domain, and I am unable to adjudicate among these without further investigation. For example, anticipating the spatial location or sensory features of a target in the external domain may have elicited sustained neural activity according to spatial or feature-based receptive field properties (Kastner et al., 1999; Silver et al., 2007; Stokes et al., 2009). In contrast, cueing internal representations in working memory often triggers more transient sensory modulations (Mok et al., 2016; Wallis et al., 2015). The alpha-lateralisation findings hint at these effects, with prolonged modulation following external as compared to internal first cues (Figure 5.6A, D, G, K). Alternatively, the presentation of the second cue may have activated latent functional states in working memory through a pinging mechanism

(Stokes, 2015; Wolff et al., 2015, 2017). This neural reactivation can also occur in the presence of active working memory representations, causing an increase in their decodability (Barbosa et al., 2021). Although the second cue indicated a different item than the first cue, its associated sensory input may have interacted with latent or activated working-memory states, potentially acting as a ping in trials in where attention was initially directed to an internal item. Another possibility is that the first cue may have placed the system in different states of sensorimotor preparation. If an internal item was cued, its representation could have been readily transformed into a motor code (Boettcher et al., 2021; D. Schneider, Barth, & Wascher, 2017; van Ede, Chekroud, Stokes, et al., 2019), whereas, this was not possible when sensory items were anticipated by external attention. Therefore, the sustained decoding of the first-cued attentional domain following the second cue could reflect a lasting effect of output gating following internal but not external first cues. Any, or all, of these factors may have contributed to the successful decoding of the first-cued domain after the second cue.

5.5.3 Relationship Between Shift-Type and Attentional-Domain Decoding

The decodability of the lingering first-cued attentional domain did not trade-off against the decoding of the second-cued attentional domain (Supplementary Figure 8.19). Instead, after the second cue, information processing unfolded for both the first- and second-cued attentional domains in tandem, implying the co-existence of multiple attentional states in the brain. Though the present study did not directly tap into the nature and functional relevance of these co-existing states, the results revealed that the extent to which the first state lingered (i.e., could still be decoded after the second cue) was predictive of the decodability of within- versus between-domain shifts (Figure 5.5). This suggests that preceding attentional states shape upcoming attentional shifts in the nearby future. What role such lingering states play, and whether/how they mediate the parallel processing of contents from preceding and current attentional states remains an exciting endeavour for future studies.

5.5.4 Limitations and Future Directions

The current findings pave the way for investigating the dynamic mechanisms orchestrating the seamless shifts of attention between external and internal contents to guide flexible behaviour. My investigation serves as an initial exploration, opening the door to relevant unanswered questions and encouraging further task manipulations. For instance, in the present study, when participants were internally cued, they were able to select the relevant content from working memory. Conversely, when externally cued, attention could be directed to a specific spatial location, but the actual attentional selection occurred only after the external-display presentation. In future studies, it would also be interesting to investigate the dynamics of transitioning between external and internal content in tasks where external information is presented concurrently with external cues while remaining perceptually available until the reporting stage. Another important avenue for future studies would be to situate and study attention shifts in more naturalistic settings. Such additional manipulations will likely yield complementary insights, contributing to a more comprehensive understanding of the principles and mechanisms that govern attention shifts between perception and working memory.

5.5.5 Conclusion

In summary, this experiment revealed distinct neural dynamics for shifting within versus between perception and working-memory contents. These distinct neural processes evolved dynamically, were broadly distributed, predicted behavioural performance, and occurred without delaying the timing of spatial orienting of attention. Besides successfully decoding the shift type, I was also able to decode the attentional domain both before and after the shift, thereby capturing the origin and destination of attention. Lastly, the present results uncovered a direct link between the previously attended domain and the performed shift type, suggesting that the lingering activity related to the first-cued domain influences processing following the second cue in a way that distinguishes within- from between-domain shifts.

6 General Discussion

Our behaviours are shaped by the dynamic interplay between external attention directed to signals in the environment and internal attention directed to past sensations stored within our memories. Exclusively relying on either external or internal information to guide behaviour would likely result in inefficiencies if not potentially dangerous errors. During my doctoral studies, I have acknowledged this circumstance, investigating scenarios where individuals encounter both external and internal information in quick succession, thereby necessitating rapid shifts in attention between the two attentional domains. Collectively, this thesis advocates a holistic approach for studying external and internal attention. By considering external and internal information in concert, we can gain deeper insights into the intricacies of everyday human cognition and behaviour. In this last section of the thesis, I will synthesise the findings from the preceding empirical chapters and discuss their broader implications. Building upon the insights gained from this research, I will explore future directions for studying the dynamic interaction between external and internal attention.

In situations demanding shifts between perception and working memory, modulatory influences emerge from external attention to internal attention and vice versa (Teng et al., 2023; Teng & Kravitz, 2019). In Chapter 2, I presented two online studies investigating these cross-domain consequences of attentional shifts. The work diverged from previous studies by introducing an additional dimension to the dynamic interplay between domains. Specifically, I employed predictable temporal structures as a tool to modulate the attentional processing of both external sensory input (Chapter 2.4) and internal representations (Chapter 2.3), while monitoring for potential consequences within the opposing domain, respectively. The results of Chapter 2.3 suggested that anticipating when internal contents will be required for an upcoming working-memory task not only improved working-memory-guided behaviour, but also affected the performance of an external task occurring during the period of anticipation. The findings of

Chapter 2.4 revealed that being able to anticipate when an external event will occur can effectively ameliorate the detrimental effects thereof on the quality of internal representations. Furthermore, the results of Chapter 2.4 suggested that temporal expectations enable a protective mechanism which shields internal content from external interference, rather than suppressing such interference. Taken together, the empirical work presented in Chapter 2 illustrated that temporal expectations not only impact the domain to which they are applied but also extend their influence to the other domain.

While Chapter 2 focussed on cross-domain modulations resulting from attentional shifts, Chapter 3 marked the beginning of my investigation into the shifting process itself. I utilised a task design similar to that in Chapter 2.4, but this time I investigated the timing of when internal representations were reactivated after engaging in an external task, thus focusing on external-to-internal shifts. Although the study of internal attention has predominantly focussed on past sensory representations, it has become increasingly recognised that the purpose of retaining sensory information in working memory is to anticipate and guide future behaviour (Chatham & Badre, 2015; Myers et al., 2017; Nobre & Stokes, 2019; Olivers & Roelfsema, 2020; van Ede, 2020; van Ede & Nobre, 2023). In Chapter 3, I acknowledged the multifaceted nature of internal attention. By employing EEG time-frequency analyses, I was able to track the reselection of both visual and motor information related to working-memory content. This revealed two key findings: that visual and motor signatures reoccurred promptly following the external-task response and that changes in the latency of visual reselection after the external task were associated with variability in the onset of memory-guided behaviour. Therefore, Chapter 3 demonstrated that attentional shifts are manifested across multiple representational dimensions within a common time frame.

To get at the nature of between-domain shifts, Chapter 4 introduced a combined perception and working-memory task that allowed within- and between-domain shifts to be contrasted in each respective domain. The results revealed additional costs associated with

between- as compared to within-domain shifts (Chapter 4.3). Strikingly, these costs persisted even when more time was given to complete the attentional shift (Chapter 4.4) and when the need to shift attention in space was removed (Chapter 4.5). In summary, the findings of Chapter 4 unveiled unique characteristics of between-domain shifts, which may arise from the engagement of a control function exclusively related to this shift type.

In Chapter 5, I built on the behavioural results of Chapter 4 and used MEG to compare the neural dynamics of shifting attention to visual contents within versus between the external and internal domains. Multivariate pattern analysis revealed distinct neural dynamics for shifting within versus between perception and working-memory contents. These neural processes evolved dynamically, were broadly distributed, predicted behavioural performance, and occurred without delaying the spatial orienting of attention. Besides successfully decoding the shift type, I was also able to decode the attentional domain before and after a shift. Further, the first-cued attentional state was found to have lingered following the presentation of the second cue.

As a whole, this doctoral thesis has demonstrated that the dynamic interplay between attentional domains is (1) behaviourally costly, (2) prone to (some) modulation, (3) triggered by various events, and (4) initiated promptly, but can be leaky.

6.1 The Dynamic Interplay Is Behaviourally Costly

Throughout this thesis, I have provided multiple lines of evidence to demonstrate that the dynamic interplay between external and internal attention induces behavioural costs. In line with accounts positing that perception and working memory likely draw upon some shared neural substrates (D'Esposito & Postle, 2015; Harrison & Tong, 2009; Pasternak & Greenlee, 2005; Postle, 2006; Serences, 2016), my findings in both Chapters 2.4 and 3 indicated that internal representations are susceptible to being disrupted by external events which occur during the retention interval.

When seeking a comprehensive understanding of the dynamic interplay of attention and the potential costs it entails, it is crucial to explore all potential combinations of shifts, both within and between domains. In Chapters 4 and 5, I found higher costs for shifting between as compared to within contents of the same domain; that is, higher costs when shifting attention from internal-to-external as compared to external-to-external, or from external-to-internal as compared to internal-to-internal. That being considered, further pinpointing the origins of these additional between-domain shift costs presents a critical avenue for future research, one to which I can offer some initial answers.

Prior literature has proposed the existence of a distinct control system responsible for orchestrating the coordination of external and internal information in contexts where shifting between these domains becomes necessary (Burgess et al., 2006, 2007; Gilbert et al., 2005). If such a control system not only exists but also accounts for the variation in shift costs within and between domains, the findings of this thesis suggest that it is unlikely to function during the initial stage of attentional orienting. The biases I observed in fixational gaze behaviour, which tracked attentional orienting in both domains, did not reveal any latency or magnitude differences for within- versus between-domain shifts (Chapter 4.3). Similarly, in Chapter 5, I did not observe any differences between shift types in terms of alpha lateralisation relative to the location of the cued item. Thus, it appears that such a control mechanism involved in shifts between domains must target attentional stages other than the early orienting stage. Alternatively, no such additional control system may exist. Rather, the systems for controlling external versus internal attention may simply concurrently compete for access to the same relevant sensory representations. In this scenario, additional costs would arise from the need to resolve a stronger degree of competition between as compared to within domains. Future brain imaging studies utilising fMRI will be instrumental in contrasting various alternative explanations for the costs associated with shifts between domains.

The costs of the dynamic interplay of external and internal attention raise yet another interesting question. Could the observed costs be due to the fact that external and internal attentional contents never cooperated to serve a common behavioural purpose? In the tasks outlined in this thesis, participants were required to shift attention between independent external and internal tasks, each requiring distinct responses (Chapter 2 and 3). In the remaining experiments, participants had to shift between different external and internal contents, while entirely disregarding the previously attended content as they transitioned to new content (Chapter 4 and 5). This suggests that, at least in the experiments employed here, the dynamic interplay between domains might have been inherently competitive in nature. Nevertheless, this does not always hold true. In our daily experiences, shifts between external and internal attention are often likely to cooperate in a synergistic manner. For example, when building a piece of flat-pack furniture, we encode the necessary pieces, maintain them, and consecutively search for them – with each piece initiating recall of the assembly step to which it is associated as it is found. Attentional contents therefore transcend their individual domain boundaries, allowing for cooperation between perception and working memory. Experimental tasks studying these naturally unfolding behaviours have become to increasingly popular (Draschkow et al., 2021; Kumle et al., 2024). Mapping out the potential behavioural costs of between-domain shifts in a scenario where external and internal attention operate synergistically will be a natural progression for this line of research.

6.2 The Dynamic Interplay Is Subject to (Some) Modulations

The current work further extends prior literature by showing that the behavioural costs arising from situations involving both external and internal attentional demands are not set in stone. Rather, the interplay between attentional domains is dynamic and can be subjected to modulatory influences from multiple sources. Within this thesis, I have provided evidence that such modulations can arise based on the task demands encountered during shifts (i.e., what?), the

temporal regularities underlying shifts (i.e., when?), and the direction of shifts (i.e., where from/to?).

I provided evidence that the costs of handling external and internal information in a common timeframe are modulated by the *task demands* being encountered (Chapter 2.4). While both perceptual distractions and external-task interruptions led to decreased memory-guided performance, internal representations were more severely impacted by additional external-task demands. These modulatory differences are likely due to the fact that interrupters require the reactivation of action plans and visual information (Chapter 3), which is by contrast unnecessary during perceptual distractions (Clapp et al., 2010; Sakai et al., 2002; Zickerick, Kobald, et al., 2021). In addition, external-task and working-memory demands like draw upon similar attentional control processes and mental workspaces (Bae & Luck, 2019). This finding also resonates with the notion that the level of competition between external and internal attention increases with rising demands for attentional control in either domain (Dixon et al., 2014). In regard to the experiments detailed in Chapters 2 and 3, deciphering whether the adverse impacts of external tasks primarily stem from increased visual processing or from motor processes proves challenging, as sensory input and motor demands were directly coupled (i.e., left-tilted items always corresponded to left-hand responses, while right-tilted items corresponded to right-hand responses). Looking ahead, it would be compelling to systematically manipulate the degree of motor and sensory overlap between external and internal tasks. This approach would facilitate a thorough exploration of the individual and combined effects of sensory and motor demands during attentional shifts.

Furthermore, I demonstrated that the ability to anticipate the appearance of external or internal contents, and therefore predict *when to shift* attention, can further influence the dynamic interplay between domains. In particular, I showed that temporal expectations applied to one domain can hold significant consequences on the opposing domain. This expands upon previous research which only considered the effects of temporal regularities within one domain (Denison, 2024; Nobre, 2001; Nobre & van Ede, 2018, 2023). In Chapter 2.3, I showed that temporal

expectations regarding when to use internal contents improved memory-guided behaviour, even though an external task occurred during the retention interval. Moreover, I demonstrated that having a predictable temporal structure in a working-memory task can improve the performance of an external task appearing during the retention interval. Further, Chapter 2.4 demonstrated that temporal expectations regarding external events not only enhance the performance of external tasks, but also improve memory-guided behaviour through enhanced shielding of internal contents. Taken together, the findings of Chapter 2 revealed that knowing when to shift attention to one domain additionally modulates attentional processing in the other domain. Moreover, these results imply a degree of flexibility in attentional shifts, rather than strict adherence to a rigid cycling scheme between the external and internal domains, as suggested by theoretical and empirical work (Honey et al., 2017; Weilnhammer et al., 2023). Even if such fixed cycling schemes exist, the findings of this thesis reveal that they must allow for some level of adaptability (for further discussion on the automatic cycling between domains, refer to the next section; Chapter 6.3).

In Chapters 4 and 5, I observed that shifts between separate domains impose greater behavioural costs compared to shifts within the same domain, with these effects appearing to be influenced by the *direction of the shift*. In particular, shifting from external to internal attention was more costly than internal-to-external shifts. While these findings resonate with previous studies employing task-switching-inspired paradigms (Calzolari et al., 2022; Dark, 1990; Hautekiet et al., 2023; Verschooren et al., 2020), the explanations provided by those studies are unlikely to fully account for the effects observed in the present study. I found no evidence to indicate that internal attention holds dominance over external attention (Verschooren & Egner, 2023), nor that the asymmetrical shift cost can be attributed to the prolonged time required for memory retrieval during transitions to the internal domain (Dark, 1990). The MEG data in Chapter 5 suggested that the previously attended domain may play an important role during between-domain shifts, yet the mixture-modelling results from Chapters 4.4 and 4.5 indicated that swaps with the initially

cued item could not fully explain the asymmetric shift cost. Given the findings of this thesis, I am unable to offer a precise answer regarding the factors that led to these directional differences. However, I intend to address this question further in future studies.

While the dynamic interplay between external and internal attention may be subject to various modulatory influences, the research in this thesis also offers evidence that certain aspects of this interplay appear to be resistant to modification. Specifically, the between-domain shift cost and its included directional asymmetries appeared notably resistant to modulation, making it of unique interest within attention research. These shift costs persisted even when participants were granted additional time to complete the attentional shift (Chapter 4.4), or when the requirement to shift attention in space was eliminated (Chapter 4.5). In the future, investigating whether specific manipulations can modulate between-domain shift costs will prove intriguing. Moreover, it will be of interest to examine how different types of modulations interact. Some combinations of modulatory influences might alleviate shift costs, while others could exacerbate them, and yet others still may not have any discernible impact at all.

6.3 The Dynamic Interplay Is Triggered by Various Events

This thesis also identified various triggers for initiating the dynamic interplay between domains. In prior research, pre-cues and retro-cues have proven invaluable for studying external and internal attention, respectively (Carrasco, 2011; Griffin & Nobre, 2003; Landman et al., 2003; Posner, 1980; Souza & Oberauer, 2016). In Chapters 4 and 5, I employed a novel task design that incorporated both types of cues to explicitly guide attention between perception and working memory. While explicit cues are the dominant sources of attentional orienting in the laboratory, we rarely encounter such obvious and direct prompts during natural behaviour. In most everyday scenarios, we instead shift our attention within dynamically unfolding contexts based on our internally generated goals and learned expectations. In line with this, I have provided evidence that implicitly learned temporal expectations support attentional shifts from perception to

working memory and vice versa (Chapter 2). Besides such explicitly instructed and implicitly learned triggers of attentional shifts, I have also investigated the trajectory of other natural transitions of attention – those that are self-initiated following the completion of an external task (Chapter 3). In other words, Chapter 3 has shown that actions in of themselves can trigger the reselection of internal contents.

When considering the different factors that prompt shifts of attention across domains, a plethora of intriguing questions emerge. The first aspect I would like to highlight is the degree to which shifts are voluntarily triggered. Within this thesis, I have considered voluntary shifts of attention and those driven by learned experiences. However, our external attention is also often involuntarily captured by stimuli in the external world. Similarly, recent studies have revealed a correspondingly involuntary and stimulus-driven capture of internal attention within working memory (Ester & Nouri, 2023; van Ede, Board, et al., 2020). Exploring and comparing voluntary and involuntary between-domain shifts presents an intriguing avenue for future research. Likewise, it would be compelling to investigate whether between-domain shift costs might be mitigated when external and internal contents share bottom-up features. At the present moment, I would expect this to be the case, given that working-memory contents automatically draw attention to matching sensory stimuli in the world (Y. Huang et al., 2019; Kiyonaga et al., 2012; Olivers, 2009; Olivers et al., 2006; Sasin & Fougnie, 2020; Soto et al., 2005, 2008), and that sensory stimuli automatically draw attention to matching contents in working memory (Ester & Nouri, 2023; van Ede, Board, et al., 2020). Hence, transitions between domains in either direction should proceed more smoothly when the contents within one domain can help to guide attention towards contents in the other domain.

The factors that trigger between-domain shifts and determine how long we remain in one domain or the other are highly dependent on the context and goals of a given situation. For example, the dynamic interplay between domains might significantly vary during learning versus high-performance situations (Nobre, 2018). According to Nobre (2018), learning situations favour

a high degree of uncertainty; placing greater emphasis on sampling information from the external world while relying less on stored representations. In contrast, high-performance situations value certainty, where internal contents and predictions about upcoming stimuli are more important for facilitating performance than the sampling of new external information. When gathering information from the environment becomes more effortful, for example by increasing the time investment or locomotive effort needed to sample external information, we might implicitly rely more strongly on guidance by internal contents (Draschkow et al., 2021; Kumle et al., 2024), thus decreasing the overall between-domain shift rate. Taken together, what might serve as a trigger for a between-domain shift in one situation may not necessarily operate as such in another. If the dynamic interplay between domains is not adaptively optimised for the given specific context, it is far less likely to efficiently guide natural behaviour.

Another point of interest is whether shifts between external and internal attention could also happen automatically, independent of external events or internal states serving as triggers. As such, cycling between external and internal attention might be an intrinsic property of how attention is organised, representing its default mode of operation. Theoretical and empirical work provide various channels of support for the notion that we may automatically alternate between sampling of external and internal information (Honey et al., 2017; Kragel et al., 2020; Weilnhammer et al., 2023). For instance, recent findings from computational modelling demonstrated that, across a series of multiple trials, humans and mice automatically and periodically fluctuate between externally and internally directed cognition in antiphase (Weilnhammer et al., 2023). In fact, most of the empirical evidence supporting the concept of automatic cycling between domains comes from hippocampal research in rodents, which has demonstrated that the hippocampal area CA1 is differently activated by to-be-encoded and retrieved information. In particular, CA1 receives sensory input directly through a monosynaptic pathway from the entorhinal cortex, while also receiving mnemonic signals from area CA3 through a trisynaptic pathway during memory retrieval (via entorhinal cortex to dentate gyrus to CA3; see

Honey et al., 2017; Tarder-Stoll et al., 2020; van Strien et al., 2009). When a rodent navigates through space, the hippocampal area CA1 receives alternating input from these two pathways across functionally different phases of the theta rhythm: CA1 firing induced by the monosynaptic pathway represents information immediately in front of animal, as opposed to retrospective information which is transmitted via the trisynaptic pathway (Colgin, 2016; Colgin et al., 2009; Hasselmo et al., 2002; Honey et al., 2017; Itskov et al., 2008; O'Keefe & Recce, 1993). Intracranial recordings in neurosurgical patients preforming a spatial memory task have also suggested that hippocampal theta could coordinate the rapid sampling between novel content in the environment and internally stored memory contents (Kragel et al., 2020).

While research on hippocampal encoding and retrieval modes is currently under thorough investigation, the exploration of rhythmic sampling between external and internal attention is still in its infancy. As of now, there is evidence suggesting that external attention samples the environment in theta-rhythmic cycles, leading to alternating periods of either enhanced or diminished perceptual sensitivity (Fiebelkorn et al., 2013; Fiebelkorn & Kastner, 2019; Helfrich et al., 2018; Landau et al., 2015; Landau & Fries, 2012). In addition, recent studies have demonstrated that internal attention directed towards items in working memory is modulated in a similar theta-rhythmic manner (Chota et al., 2022; Peters et al., 2021; Pomper & Ansorge, 2021). The next step for future research is to integrate these two lines of inquiry to ascertain how the sampling of external and internal attention relate to one another. If there is indeed evidence supporting automatic sampling between external and internal attention, it will be intriguing to investigate how these attentional oscillations are nested within voluntary attentional shifts that gradually determine the dominance of one attentional domain over the other. In particular, building on the discussion in the previous section (Chapter 6.2), examining how automatic fluctuations between domains might integrate with flexible and adaptable priority shifts of attention presents an interesting avenue for further exploration.

6.4 The Dynamic Interplay Is Prompt but Leaky

The work presented in this thesis revealed the promptness of the dynamic interplay between attentional domains. In Chapter 3, I investigated shifts from external to internal attention using EEG. In particular, I studied the reselection of visuo-spatial and motor signatures following external-task engagement, acknowledging the multifaceted and future-oriented nature of internal attention (Chatham & Badre, 2015; Myers et al., 2017; Nobre & Stokes, 2019; Olivers & Roelfsema, 2020; van Ede, 2020; van Ede & Nobre, 2023). More specifically, I observed attentional shifts to working-memory contents, indicated by both visuo-spatial and motor signatures, which promptly followed the completion of an external task (as opposed to occurring just-in-time at the moment when memory-guided behaviour was required). Moreover, the results of Chapter 4 and 5 suggested that visuo-spatial reorienting during within- and between-domain shifts happened early after the shift was prompted, as indicated by biases in fixational gaze behaviour (Chapter 4) and alpha lateralisation (Chapter 5). While I did not find any differences in spatial reorienting between within- and between-domain shifts, I observed reliable decoding of the performed shift type, which again occurred only briefly after the cued reorienting of attention (Chapter 5).

In addition to relatively rapid reorienting responses, I have also found evidence suggesting that processing related to the previously attended domain persisted even after a shift had been prompted. For example, gaze biases and alpha lateralisation were sustained longer following externally directing cues as compared to internally directing cues (Chapters 4 and 5). Most notably, multivariate pattern analyses revealed that the previously attended domain could still be successfully decoded even after attention had already shifted to the new content (Chapter 5). This finding is particularly fascinating. Previous literature has suggested that the dynamic interplay between domains is prohibited by our inability to simultaneously monitor sensory information from the environment and internal representations stored in memory (Awh et al., 1998; Brooks, 1968; Dixon et al., 2014; M. D. Fox et al., 2005; Handy & Kam, 2015; Kiyonaga & Egner, 2013;

Macedo-Pascual et al., 2023; Servais et al., 2023; Unsworth & Robison, 2016). That is, whenever internal attention is engaged, the processing of external attention must be suppressed, and vice versa. However, if we are unable to simultaneously process external and internal information, why can we then decode both the currently and previously attended domains at the same time? According to Dixon et al. (2014), there is a chance that external and internal contents are not intrinsically antagonistic, but rather, their relationship depends on the current processing demands in either domain. Specifically, it is suggested that external and internal attention might co-occur with minimal interference when one or both involve spontaneous processing. In the event of an attentional shift, the previously attended domain might necessitate minimal processing, thus allowing it to be traced on a neural level even though attention has shifted to the next event. Exploring whether this lingering activity is merely a by-product of the attentional shifting process or if it plays a more active role during transitions between domains will be a compelling avenue for future research. At present, the findings of this thesis indicate a link between the attentional state of the previous domain and the performed shift type, with increased lingering being associated with enhanced discrimination of within- versus between-domain shifts.

Though not the focus of this doctoral thesis, exploring how (coexisting) contents can be attributed to their respective domains and upheld as distinct entities represents a vital area for future research. Prior research has already documented instances of coexisting mnemonic and sensory representations in posterior areas of the brain (Rademaker et al., 2019). However, our understanding of how exactly the brain distinguishes external reality from internal states is still gradually unfolding (Dijkstra & Fleming, 2023). To fully understand shifts between domains, it will be important to grasp how the brain can differentially target external and internal contents and how these contents are attributed to their relevant domains.

6.5 Final Conclusions

During almost every natural behaviour, our attention is in constant flux, seamlessly transitioning between contents available in the external environment and internal representations maintained in working memory. Yet, despite the importance of this fundamental aspect of everyday cognition, previous research has predominantly examined external and internal selective attention in isolation, leaving a significant gap in our understanding of their dynamic interplay. Apart from expanding upon previous literature on cross-domain modulations, this thesis offers vital insights into the cognitive and neural mechanisms that underlie shifts of attention between our perception of the external world and the representations thereof in working memory. Throughout four empirical chapters and seven experiments, I have presented evidence suggesting that the dynamic interplay between attentional domains can impose costs on behavioural performance, which can, in turn, be subjected to modulatory influences. Furthermore, I have demonstrated that between-domain shifts can be triggered by various events and occur promptly, but are also leaky, resulting in lingering traces of previously attended domains. Taken together, this thesis advocates for a holistic approach to examine external and internal attention. Studying external and internal attention in such an integrative way presents additional challenges not encountered when investigating attentional domains separately. However, considering external and internal information in concert has and will continue to provide a far more nuanced elucidation of the intricacies of everyday human cognition and behaviour.

7 References

Albers, A. M., Kok, P., Toni, I., Dijkerman, H. C., & de Lange, F. P. (2013). Shared Representations for Working Memory and Mental Imagery in Early Visual Cortex. *Current Biology*, *23*(15), 1427–1431. https://doi.org/10.1016/j.cub.2013.05.065

Allport, D. A., Styles, E. A., & Hsieh, S. (1994). Shifting intentional set: Exploring the dynamic control of tasks. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing* (pp. 421–452). The MIT Press.

Aly, M., & Turk-Browne, N. B. (2016a). Attention promotes episodic encoding by stabilizing hippocampal representations. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(4), E420-429. https://doi.org/10.1073/pnas.1518931113

Aly, M., & Turk-Browne, N. B. (2016b). Attention Stabilizes Representations in the Human Hippocampus. *Cerebral Cortex (New York, N.Y.: 1991)*, *26*(2), 783–796. https://doi.org/10.1093/cercor/bhv041

Anderson, B. (2011). There is no Such Thing as Attention. *Frontiers in Psychology*, *2*. https://www.frontiersin.org/articles/10.3389/fpsyg.2011.00246

Anderson, B. A., Kim, H., Kim, A. J., Liao, M.-R., Mrkonja, L., Clement, A., & Grégoire, L. (2021). The past, present, and future of selection history. *Neuroscience & Biobehavioral Reviews*, *130*, 326–350. https://doi.org/10.1016/j.neubiorev.2021.09.004

Andrews-Hanna, J. R., Saxe, R., & Yarkoni, T. (2014). Contributions of episodic retrieval and mentalizing to autobiographical thought: Evidence from functional neuroimaging, resting-state connectivity, and fMRI meta-analyses. *NeuroImage*, *91*, 324–335. https://doi.org/10.1016/j.neuroimage.2014.01.032

Astle, D. E., Summerfield, J., Griffin, I., & Nobre, A. C. (2012). Orienting attention to locations in mental representations. *Attention, Perception, & Psychophysics*, *74*(1), 146–162. https://doi.org/10.3758/s13414-011-0218-3

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. https://doi.org/10.1016/j.tics.2012.06.010

Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Sciences*, *5*(3), 119–126. https://doi.org/10.1016/S1364-6613(00)01593-X

Awh, E., Jonides, J., & Reuter-Lorenz, P. A. (1998). Rehearsal in spatial working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(3), 780–790. https://doi.org/10.1037/0096-1523.24.3.780

Awh, E., Vogel, E. K., & Oh, S.-H. (2006). Interactions between attention and working memory. *Neuroscience*, *139*(1), 201–208. https://doi.org/10.1016/j.neuroscience.2005.08.023

Backer, K. C., Buchsbaum, B. R., & Alain, C. (2020). Orienting Attention to Short-Term Memory Representations via Sensory Modality and Semantic Category Retro-Cues. *eNeuro*, *7*(6). https://doi.org/10.1523/ENEURO.0018-20.2020

Baddeley, A. D. (1992). Working Memory. *Science*, *255*(5044), 556–559. https://doi.org/10.1126/science.1736359

Baddeley, A. D. (2012). Working Memory: Theories, Models, and Controversies. *Annual Review of Psychology*, *63*(1), 1–29. https://doi.org/10.1146/annurev-psych-120710-100422

Baddeley, A. D., & Hitch, G. (1974). Working Memory. In G. H. Bower (Ed.), *Psychology of Learning and Motivation* (Vol. 8, pp. 47–89). Academic Press. https://doi.org/10.1016/S0079-7421(08)60452-1

Bae, G.-Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception, & Psychophysics*, *79*(8), 2376–2395. https://doi.org/10.3758/s13414-017-1404-8

Bae, G.-Y., & Luck, S. J. (2019). What happens to an individual visual working memory representation when it is interrupted? *British Journal of Psychology*, *110*(2), 268–287. https://doi.org/10.1111/bjop.12339

Bahmani, Z., Daliri, M. R., Merrikhi, Y., Clark, K., & Noudoost, B. (2018). Working Memory Enhances Cortical Representations via Spatially Specific Coordination of Spike Times. *Neuron*, *97*(4), 967-979.e6. https://doi.org/10.1016/j.neuron.2018.01.012

Balestrieri, E., Ronconi, L., & Melcher, D. (2022). Shared resources between visual attention and visual working memory are allocated through rhythmic sampling. *European Journal of Neuroscience*, *55*(11–12), 3040–3053. https://doi.org/10.1111/ejn.15264

Barak, O., Tsodyks, M., & Romo, R. (2010). Neuronal population coding of parametric working memory. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(28), 9424–9430. https://doi.org/10.1523/JNEUROSCI.1875-10.2010

Barbosa, J., Lozano-Soldevilla, D., & Compte, A. (2021). Pinging the brain with visual impulses reveals electrically active, not activity-silent, working memories. *PLOS Biology*, *19*(10), e3001436. https://doi.org/10.1371/journal.pbio.3001436

Barrouillet, P., Bernardin, S., Portrat, S., Vergauwe, E., & Camos, V. (2007). Time and cognitive load in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(3), 570–585. https://doi.org/10.1037/0278-7393.33.3.570

Barth, A., & Schneider, D. (2018). Manipulating the focus of attention in working memory: Evidence for a protection of multiple items against perceptual interference. *Psychophysiology*, *55*(7), e13062. https://doi.org/10.1111/psyp.13062

Battistoni, E., Stein, T., & Peelen, M. V. (2017). Preparatory attention in visual cortex. *Annals of the New York Academy of Sciences*, *1396*(1), 92–107. https://doi.org/10.1111/nyas.13320

Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, *9*(10), 7–7. https://doi.org/10.1167/9.10.7

Beck, D. M., & Kastner, S. (2009). Top-down and bottom-up mechanisms in biasing competition in the human brain. *Vision Research*, *49*(10), 1154–1165. https://doi.org/10.1016/j.visres.2008.07.012

Berry, A. S., Zanto, T. P., Rutman, A. M., Clapp, W. C., & Gazzaley, A. (2009). Practice-Related Improvement in Working Memory is Modulated by Changes in Processing External Interference. *Journal of Neurophysiology*, *102*(3), 1779–1789. https://doi.org/10.1152/jn.00179.2009

Bettencourt, K. C., & Xu, Y. (2016). Decoding the content of visual short-term memory under distraction in occipital and parietal areas. *Nature Neuroscience*, *19*(1), 150–157. https://doi.org/10.1038/nn.4174

Block, R. A., Hancock, P. A., & Zakay, D. (2010). How cognitive load affects duration judgments: A meta-analytic review. *Acta Psychologica*, *134*(3), 330–343. https://doi.org/10.1016/j.actpsy.2010.03.006

Boettcher, S. E. P., Gresch, D., Nobre, A. C., & van Ede, F. (2021). Output planning at the input stage in visual working memory. *Science Advances*, *7*(13), eabe8212. https://doi.org/10.1126/sciadv.abe8212

Boettcher, S. E. P., Shalev, N., Wolfe, J. M., & Nobre, A. C. (2022). Right place, right time: Spatiotemporal predictions guide attention in dynamic visual search. *Journal of Experimental Psychology: General*, *151*(2), 348–362. https://doi.org/10.1037/xge0000901

Bonnefond, M., & Jensen, O. (2012). Alpha Oscillations Serve to Protect Working Memory Maintenance against Anticipated Distracters. *Current Biology*, *22*(20), 1969–1974. https://doi.org/10.1016/j.cub.2012.08.029

Brady, T. F., & Alvarez, G. A. (2011). Hierarchical Encoding in Visual Working Memory: Ensemble Statistics Bias Memory for Individual Items. *Psychological Science*, *22*(3), 384–392. https://doi.org/10.1177/0956797610397956

Broadbent, D. E. (1958). *Perception and communication*. Pergamon Press.

Brody, C. D., Hernández, A., Zainos, A., & Romo, R. (2003). Timing and Neural Encoding of Somatosensory Parametric Working Memory in Macaque Prefrontal Cortex. *Cerebral Cortex*, *13*(11), 1196–1207. https://doi.org/10.1093/cercor/bhg100

Brooks, L. R. (1968). Spatial and verbal components of the act of recall. *Canadian Journal of Psychology / Revue Canadienne de Psychologie*, *22*(5), 349–368. https://doi.org/10.1037/h0082775

Brown, J. (1958). Some Tests of the Decay Theory of Immediate Memory. *Quarterly Journal of Experimental Psychology*, *10*(1), 12–21. https://doi.org/10.1080/17470215808416249

Brown, S. W. (1997). Attentional resources in timing: Interference effects in concurrent temporal and nontemporal working memory tasks. *Perception & Psychophysics*, *59*(7), 1118–1140. https://doi.org/10.3758/BF03205526

Brown, S. W. (2006). Timing and executive function: Bidirectional interference between concurrent temporal production and randomization tasks. *Memory & Cognition*, *34*(7), 1464–1471. https://doi.org/10.3758/BF03195911

Brown, S. W., Collier, S. A., & Night, J. C. (2013). Timing and executive resources: Dual-task interference patterns between temporal production and shifting, updating, and inhibition tasks. *Journal of Experimental Psychology: Human Perception and Performance*, *39*(4), 947–963. https://doi.org/10.1037/a0030484

Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A Neural Theory of Visual Attention: Bridging Cognition and Neurophysiology. *Psychological Review*, *112*(2), 291–328. https://doi.org/10.1037/0033-295X.112.2.291

Burgess, P. W., Dumontheil, I., & Gilbert, S. J. (2007). The gateway hypothesis of rostral prefrontal cortex (area 10) function. *Trends in Cognitive Sciences*, *11*(7), 290–298. https://doi.org/10.1016/j.tics.2007.05.004

Burgess, P. W., Gilbert, S. J., Okuda, J., & Simons, J. S. (2006). Rostral Prefrontal Brain Regions (Area 10): A Gateway between Inner Thought and the External World? In N. Sebanz & W. Prinz (Eds.), *Disorders of Volition* (pp. 373–396). The MIT Press. https://doi.org/10.7551/mitpress/2457.003.0021

Calzolari, S., Boneva, S., & Fernández-Espejo, D. (2022). Investigating the shift between externally and internally oriented cognition: A novel task-switching paradigm. *Neuroscience of Consciousness*, *2022*(1), niac016. https://doi.org/10.1093/nc/niac016

Capizzi, M., Correa, Á., & Sanabria, D. (2013). Temporal orienting of attention is interfered by concurrent working memory updating. *Neuropsychologia*, *51*(2), 326–339. https://doi.org/10.1016/j.neuropsychologia.2012.10.005

Capizzi, M., Sanabria, D., & Correa, Á. (2012). Dissociating controlled from automatic processing in temporal preparation. *Cognition*, *123*(2), 293–302. https://doi.org/10.1016/j.cognition.2012.02.005

Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional Templates in Visual Working Memory. *Journal of Neuroscience*, *31*(25), 9315–9322. https://doi.org/10.1523/JNEUROSCI.1097-11.2011

Carlson, R. A., Wenger, J. L., & Sullivan, M. A. (1993). Coordinating information from perception and working memory. *Journal of Experimental Psychology. Human Perception and Performance*, *19*(3), 531–548. https://doi.org/10.1037//0096-1523.19.3.531

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research*, *51*(13), 1484–1525. https://doi.org/10.1016/j.visres.2011.04.012

Chao, L. L., & Knight, R. T. (1998). Contribution of Human Prefrontal Cortex to Delay Performance. *Journal of Cognitive Neuroscience*, *10*(2), 167–177. https://doi.org/10.1162/089892998562636

Chatham, C. H., & Badre, D. (2015). Multiple gates on working memory. *Current Opinion in Behavioral Sciences*, *1*, 23–31. https://doi.org/10.1016/j.cobeha.2014.08.001

Chauvin, J. J., Gillebert, C. R., Rohenkohl, G., Humphreys, G. W., & Nobre, A. C. (2016). Temporal orienting of attention can be preserved in normal aging. *Psychology and Aging*, *31*(5), 442–455. https://doi.org/10.1037/pag0000105

Chawla, D., Rees, G., & Friston, K. J. (1999). The physiological basis of attentional modulation in extrastriate visual areas. *Nature Neuroscience*, *2*(7), 671–676. https://doi.org/10.1038/10230

Chelazzi, L., Miller, E. K., Duncan, J., & Desimone, R. (1993). A neural basis for visual search in inferior temporal cortex. *Nature*, *363*(6427), Article 6427. https://doi.org/10.1038/363345a0

Chota, S., Gayet, S., Kenemans, J. L., Olivers, C. N. L., & Van der Stigchel, S. (2023). A matter of availability: Sharper tuning for memorized than for perceived stimulus features. *Cerebral Cortex (New York, N.Y.: 1991)*, bhad064. https://doi.org/10.1093/cercor/bhad064

Chota, S., Leto, C., Van Zantwijk, L., & Van Der Stigchel, S. (2022). Attention rhythmically samples multi-feature objects in working memory. *Scientific Reports*, *12*(1), 14703. https://doi.org/10.1038/s41598-022-18819-z

Christoff, K., Irving, Z. C., Fox, K. C. R., Spreng, R. N., & Andrews-Hanna, J. R. (2016). Mind-wandering as spontaneous thought: A dynamic framework. *Nature Reviews Neuroscience*, *17*(11), Article 11. https://doi.org/10.1038/nrn.2016.113

Christophel, T. B., Klink, P. C., Spitzer, B., Roelfsema, P. R., & Haynes, J.-D. (2017). The Distributed Nature of Working Memory. *Trends in Cognitive Sciences*, *21*(2), 111–124. https://doi.org/10.1016/j.tics.2016.12.007

Chun, M. M. (2011). Visual working memory as visual attention sustained internally over time. *Neuropsychologia*, *49*(6), 1407–1409. https://doi.org/10.1016/j.neuropsychologia.2011.01.029

Chun, M. M., Golomb, J. D., & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*, 73–101. https://doi.org/10.1146/annurev.psych.093008.100427

Chunharas, C., Rademaker, R. L., Brady, T. F., & Serences, J. T. (2022). An adaptive perspective on visual working memory distortions. *Journal of Experimental Psychology: General*, *151*(10), 2300–2323. https://doi.org/10.1037/xge0001191

Clapp, W. C., Rubens, M. T., & Gazzaley, A. (2010). Mechanisms of Working Memory Disruption by External Interference. *Cerebral Cortex*, *20*(4), 859–872. https://doi.org/10.1093/cercor/bhp150

Cohen, M. R., & Maunsell, J. H. R. (2014). Neuronal Mechanisms of Spatial Attention in Visual Cerebral Cortex. In A. C. (Kia) Nobre & S. Kastner (Eds.), *The Oxford Handbook of Attention* (p. 0). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199675111.013.007

Colgin, L. L. (2016). Rhythms of the hippocampal network. *Nature Reviews Neuroscience*, *17*(4), 239–249. https://doi.org/10.1038/nrn.2016.21

Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., & Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, *462*(7271), 353–357. https://doi.org/10.1038/nature08573

Compte, A., Brunel, N., Goldman-Rakic, P. S., & Wang, X.-J. (2000). Synaptic Mechanisms and Network Dynamics Underlying Spatial Working Memory in a Cortical Network Model. *Cerebral Cortex*, *10*(9), 910–923. https://doi.org/10.1093/cercor/10.9.910

Conrad, R. (1967). Interference or decay over short retention intervals? *Journal of Verbal Learning and Verbal Behavior*, *6*(1), 49–54. https://doi.org/10.1016/S0022-5371(67)80048-3

Constantinidis, C., & Steinmetz, M. A. (1996). Neuronal activity in posterior parietal area 7a during the delay periods of a spatial memory task. *Journal of Neurophysiology*, *76*(2), 1352–1355. https://doi.org/10.1152/jn.1996.76.2.1352

Corbetta, M., & Shulman, G. L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, *3*(3), Article 3. https://doi.org/10.1038/nrn755

Correa, Á., Lupiáñez, J., Milliken, B., & Tudela, P. (2004). Endogenous temporal orienting of attention in detection and discrimination tasks. *Perception & Psychophysics*, *66*(2), 264–278. https://doi.org/10.3758/BF03194878

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

Coull, J. T. (2009). Neural Substrates of Mounting Temporal Expectation. *PLoS Biology*, *7*(8), e1000166. https://doi.org/10.1371/journal.pbio.1000166

Coull, J. T., Frith, C. D., Büchel, C., & Nobre, A. C. (2000). Orienting attention in time: Behavioural and neuroanatomical distinction between exogenous and endogenous shifts. *Neuropsychologia*, *38*(6), 808–819. https://doi.org/10.1016/S0028-3932(99)00132-3

Coull, J. T., & Nobre, A. C. (1998). Where and when to pay attention: The neural systems for directing attention to spatial locations and to time intervals as revealed by both PET and fMRI. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *18*(18), 7426–7435. https://doi.org/10.1523/JNEUROSCI.18-18-07426.1998

Courtney, S. M., Petit, L., Haxby, J. V., & Ungerleider, L. G. (1998). The role of prefrontal cortex in working memory: Examining the contents of consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *353*(1377), 1819–1828.

Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information-processing system. *Psychological Bulletin*, *104*(2), 163–191. https://doi.org/10.1037/0033-2909.104.2.163

Cravo, A. M., Rohenkohl, G., Santos, K. M., & Nobre, A. C. (2017). Temporal Anticipation Based on Memory. *Journal of Cognitive Neuroscience*, *29*(12), 2081–2089. https://doi.org/10.1162/jocn_a_01172

Cristescu, T. C., & Nobre, A. C. (2008). Differential Modulation of Word Recognition by Semantic and Spatial Orienting of Attention. *Journal of Cognitive Neuroscience*, *20*(5), 787–801. https://doi.org/10.1162/jocn.2008.20503

Curtis, C. E., & D'Esposito, M. (2003). Persistent activity in the prefrontal cortex during working memory. *Trends in Cognitive Sciences*, *7*(9), 415–423. https://doi.org/10.1016/s1364-6613(03)00197-9

Dark, V. J. (1990). Switching between memory and perception: Moving attention or memory retrieval? *Memory & Cognition*, *18*(2), 119–127. https://doi.org/10.3758/BF03197088

Davranche, K., Nazarian, B., Vidal, F., & Coull, J. (2011). Orienting Attention in Time Activates Left Intraparietal Sulcus for Both Perceptual and Motor Task Goals. *Journal of Cognitive Neuroscience*, *23*(11), 3318–3330. https://doi.org/10.1162/jocn_a_00030

de Vries, I. E. J., Savran, E., van Driel, J., & Olivers, C. N. L. (2019). Oscillatory Mechanisms of Preparing for Visual Distraction. *Journal of Cognitive Neuroscience*, *31*(12), 1873–1894. https://doi.org/10.1162/jocn_a_01460

Denison, R. N. (2024). Visual temporal attention from perception to computation. *Nature Reviews Psychology*, 1–14. https://doi.org/10.1038/s44159-024-00294-0

Denison, R. N., Carrasco, M., & Heeger, D. J. (2021). A dynamic normalization model of temporal attention. *Nature Human Behaviour*, *5*(12), 1674–1685. https://doi.org/10.1038/s41562-021-01129-1

Denison, R. N., Heeger, D. J., & Carrasco, M. (2017). Attention flexibly trades off across points in time. *Psychonomic Bulletin & Review*, *24*(4), 1142–1151. https://doi.org/10.3758/s13423-016-1216-1

Denny, B. T., Kober, H., Wager, T. D., & Ochsner, K. N. (2012). A meta-analysis of functional neuroimaging studies of self- and other judgments reveals a spatial gradient for mentalizing in medial prefrontal cortex. *Journal of Cognitive Neuroscience*, *24*(8), 1742–1752. https://doi.org/10.1162/jocn_a_00233

Desimone, R. (1998). Visual attention mediated by biased competition in extrastriate visual cortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *353*(1373), 1245–1255. https://doi.org/10.1098/rstb.1998.0280

Desimone, R., & Duncan, J. (1995). Neural Mechanisms of Selective Visual Attention. *Annual Review of Neuroscience*, *18*(1), 193–222. https://doi.org/10.1146/annurev.ne.18.030195.001205

D'Esposito, M., Detre, J. A., Alsop, D. C., Shin, R. K., Atlas, S., & Grossman, M. (1995). The neural basis of the central executive system of working memory. *Nature*, *378*(6554), 279–281. https://doi.org/10.1038/378279a0

D'Esposito, M., & Postle, B. R. (2015). The Cognitive Neuroscience of Working Memory. *Annual Review of Psychology*, *66*(1), 115–142. https://doi.org/10.1146/annurev-psych-010814-015031

Dijkstra, N., & Fleming, S. M. (2023). Subjective signal strength distinguishes reality from imagination. *Nature Communications*, *14*(1), Article 1. https://doi.org/10.1038/s41467-023-37322-1

Dixon, M. L., Andrews-Hanna, J. R., Spreng, R. N., Irving, Z. C., Mills, C., Girn, M., & Christoff, K. (2017). Interactions between the default network and dorsal attention network vary across default subsystems, time, and cognitive states. *NeuroImage*, *147*, 632–649. https://doi.org/10.1016/j.neuroimage.2016.12.073

Dixon, M. L., De La Vega, A., Mills, C., Andrews-Hanna, J., Spreng, R. N., Cole, M. W., & Christoff, K. (2018). Heterogeneity within the frontoparietal control network and its relationship to the default and dorsal attention networks. *Proceedings of the National Academy of Sciences of the United States of America*, *115*(7), E1598–E1607. https://doi.org/10.1073/pnas.1715766115

Dixon, M. L., Fox, K. C. R., & Christoff, K. (2014). A framework for understanding the relationship between externally and internally directed cognition. *Neuropsychologia*, *62*, 321–330. https://doi.org/10.1016/j.neuropsychologia.2014.05.024

Dodwell, G., Nako, R., & Eimer, M. (2024). A new method for tracking the preparatory activation of target templates for visual search with high temporal precision. *Psychophysiology*, e14582. https://doi.org/10.1111/psyp.14582

Dowd, E. W., & Golomb, J. D. (2019). Object-Feature Binding Survives Dynamic Shifts of Spatial Attention. *Psychological Science*, *30*(3), 343–361. https://doi.org/10.1177/0956797618818481

Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, *11*(6), 467–473. https://doi.org/10.1111/1467-9280.00290

Draschkow, D., Kallmayer, M., & Nobre, A. C. (2021). When Natural Behavior Engages Working Memory. *Current Biology*, *31*(4), 869-874.e5. https://doi.org/10.1016/j.cub.2020.11.013

Draschkow, D., Nobre, A. C., & van Ede, F. (2022). Multiple spatial frames for immersive working memory. *Nature Human Behaviour*, *6*(4), Article 4. https://doi.org/10.1038/s41562-021-01245-y

Drazin, D. H. (1961). Effects of foreperiod, foreperiod variability, and probability of stimulus occurrence on simple reaction time. *Journal of Experimental Psychology*, *62*(1), 43–50. https://doi.org/10.1037/h0046860

Dreisbach, G., & Haider, H. (2008). That's what task sets are for: Shielding against irrelevant information. *Psychological Research*, *72*(4), 355–361. https://doi.org/10.1007/s00426-007-0131-5

Dreisbach, G., & Wenke, D. (2011). The shielding function of task sets and its relaxation during task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(6), 1540–1546. https://doi.org/10.1037/a0024077

Dube, B., Emrich, S. M., & Al-Aidroos, N. (2017). More than a filter: Feature-based attention regulates the distribution of visual working memory resources. *Journal of Experimental Psychology. Human Perception and Performance*, *43*(10), 1843–1854. https://doi.org/10.1037/xhp0000428

Dubé, C., Zhou, F., Kahana, M. J., & Sekuler, R. (2014). Similarity-based distortion of visual short-term memory is due to perceptual averaging. *Vision Research*, *96*, 8–16. https://doi.org/10.1016/j.visres.2013.12.016

Dudukovic, N. M., & Wagner, A. D. (2006). Attending to Remember and Remembering to Attend. *Neuron*, *49*(6), 784–787. https://doi.org/10.1016/j.neuron.2006.03.008

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General*, *113*(4), 501–517. https://doi.org/10.1037/0096-3445.113.4.501

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, *96*(3), 433–458. https://doi.org/10.1037/0033-295X.96.3.433

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General*, *123*, 161–177. https://doi.org/10.1037/0096-3445.123.2.161

Emrich, S. M., Lockhart, H. A., & Al-Aidroos, N. (2017). Attention mediates the flexible allocation of visual working memory resources. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(7), 1454–1465. https://doi.org/10.1037/xhp0000398

Engbert, R., & Kliegl, R. (2003). Microsaccades uncover the orientation of covert attention. *Vision Research*, *43*(9), 1035–1045. https://doi.org/10.1016/S0042-6989(03)00084-1

Ester, E. F., & Nouri, A. (2023). Internal selective attention is delayed by competition between endogenous and exogenous factors. *iScience*, *26*(7), 107259. https://doi.org/10.1016/j.isci.2023.107259

Ester, E. F., & Weese, R. (2023). Temporally Dissociable Mechanisms of Spatial, Feature, and Motor Selection during Working Memory–guided Behavior. *Journal of Cognitive Neuroscience*, *35*(12), 2014–2027. https://doi.org/10.1162/jocn_a_02061

Esterman, M., Chiu, Y.-C., Tamber-Rosenau, B. J., & Yantis, S. (2009). Decoding cognitive control in human parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *106*(42), 17974–17979. https://doi.org/10.1073/pnas.0903593106

Esterman, M., & Yantis, S. (2010). Perceptual Expectation Evokes Category-Selective Cortical Activity. *Cerebral Cortex*, *20*(5), 1245–1253. https://doi.org/10.1093/cercor/bhp188

Failing, M., Feldmann-Wüstefeld, T., Wang, B., Olivers, C., & Theeuwes, J. (2019). Statistical regularities induce spatial as well as feature-specific suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(10), 1291–1303. https://doi.org/10.1037/xhp0000660

Failing, M., & Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review*, *25*(2), 514–538. https://doi.org/10.3758/s13423-017-1380-y

Feldmann-Wüstefeld, T., & Schubö, A. (2016). Intertrial priming due to distractor repetition is eliminated in homogeneous contexts. *Attention, Perception, & Psychophysics*, *78*(7), 1935–1947. https://doi.org/10.3758/s13414-016-1115-6

Feldmann-Wüstefeld, T., & Vogel, E. K. (2019). Neural Evidence for the Contribution of Active Suppression During Working Memory Filtering. *Cerebral Cortex (New York, N.Y.: 1991)*, *29*(2), 529–543. https://doi.org/10.1093/cercor/bhx336

Feredoes, E., Heinen, K., Weiskopf, N., Ruff, C., & Driver, J. (2011). Causal evidence for frontal involvement in memory target maintenance by posterior brain areas during distracter interference of visual working memory. *Proceedings of the National Academy of Sciences*, *108*(42), 17510–17515. https://doi.org/10.1073/pnas.1106439108

Ferrante, O., Patacca, A., Di Caro, V., Della Libera, C., Santandrea, E., & Chelazzi, L. (2018). Altering spatial priority maps via statistical learning of target selection and distractor filtering. *Cortex*, *102*, 67–95. https://doi.org/10.1016/j.cortex.2017.09.027

Fiebelkorn, I. C., & Kastner, S. (2019). A Rhythmic Theory of Attention. *Trends in Cognitive Sciences*, *23*(2), 87–101. https://doi.org/10.1016/j.tics.2018.11.009

Fiebelkorn, I. C., Saalmann, Y. B., & Kastner, S. (2013). Rhythmic Sampling within and between Objects despite Sustained Attention at a Cued Location. *Current Biology*, *23*(24), 2553–2558. https://doi.org/10.1016/j.cub.2013.10.063

Fortin, C., Champagne, J., & Poirier, M. (2007). Temporal order in memory and interval timing: An interference analysis. *Acta Psychologica*, *126*(1), 18–33. https://doi.org/10.1016/j.actpsy.2006.10.002

Fox, K. C. R., Spreng, R. N., Ellamil, M., Andrews-Hanna, J. R., & Christoff, K. (2015). The wandering brain: Meta-analysis of functional neuroimaging studies of mind-wandering and related spontaneous thought processes. *NeuroImage*, *111*, 611–621. https://doi.org/10.1016/j.neuroimage.2015.02.039

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Van Essen, D. C., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings of the National Academy of Sciences*, *102*(27), 9673–9678. https://doi.org/10.1073/pnas.0504136102

Frossard, J., & Renaud, O. (2021). Permutation Tests for Regression, ANOVA, and Comparison of Signals: The permuco Package. *Journal of Statistical Software*, *99*(15). https://doi.org/10.18637/jss.v099.i15

Fukuda, K., Pereira, A. E., Saito, J. M., Tang, T. Y., Tsubomi, H., & Bae, G.-Y. (2022). Working Memory Content Is Distorted by Its Use in Perceptual Comparisons. *Psychological Science*, *33*(5), 816–829. https://doi.org/10.1177/09567976211055375

Funahashi, S., Bruce, C. J., & Goldman-Rakic, P. S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. *Journal of Neurophysiology*, *61*(2), 331–349. https://doi.org/10.1152/jn.1989.61.2.331

Fuster, J. M., & Alexander, G. E. (1971). Neuron Activity Related to Short-Term Memory. *Science*, *173*(3997), 652–654. https://doi.org/10.1126/science.173.3997.652

Gaspar, J. M., & McDonald, J. J. (2014). Suppression of Salient Objects Prevents Distraction in Visual Search. *Journal of Neuroscience*, *34*(16), 5658–5666. https://doi.org/10.1523/JNEUROSCI.4161-13.2014

Gaspelin, N., Lamy, D., Egeth, H. E., Liesefeld, H. R., Kerzel, D., Mandal, A., Müller, M. M., Schall, J. D., Schubö, A., Slagter, H. A., Stilwell, B. T., & van Moorselaar, D. (2023). The Distractor Positivity Component and the Inhibition of Distracting Stimuli. *Journal of Cognitive Neuroscience*, *35*(11), 1693–1715. https://doi.org/10.1162/jocn_a_02051

Gaspelin, N., & Luck, S. J. (2018). Distinguishing among potential mechanisms of singleton suppression. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(4), 626–644. https://doi.org/10.1037/xhp0000484

Gayet, S., Guggenmos, M., Christophel, T. B., Haynes, J.-D., Paffen, C. L. E., Stigchel, S. V. der, & Sterzer, P. (2017). Visual Working Memory Enhances the Neural Response to Matching Visual Input. *Journal of Neuroscience*, *37*(28), 6638–6647. https://doi.org/10.1523/JNEUROSCI.3418-16.2017

Gayet, S., & Peelen, M. V. (2022). Preparatory attention incorporates contextual expectations. *Current Biology*, *32*(3), 687-692.e6. https://doi.org/10.1016/j.cub.2021.11.062

Gazzaley, A. (2011). Influence of early attentional modulation on working memory. *Neuropsychologia*, *49*(6), 1410–1424. https://doi.org/10.1016/j.neuropsychologia.2010.12.022

Gazzaley, A., Cooney, J. W., McEvoy, K., Knight, R. T., & D'Esposito, M. (2005). Top-down enhancement and suppression of the magnitude and speed of neural activity. *Journal of Cognitive Neuroscience*, *17*(3), 507–517. https://doi.org/10.1162/0898929053279522

Gazzaley, A., & Nobre, A. C. (2012). Top-down modulation: Bridging selective attention and working memory. *Trends in Cognitive Sciences*, *16*(2), 129–135. https://doi.org/10.1016/j.tics.2011.11.014

Geng, J. J. (2014). Attentional Mechanisms of Distractor Suppression. *Current Directions in Psychological Science*, *23*(2), 147–153. https://doi.org/10.1177/0963721414525780

Geng, J. J., & Witkowski, P. (2019). Template-to-distractor distinctiveness regulates visual search efficiency. *Current Opinion in Psychology*, *29*, 119–125. https://doi.org/10.1016/j.copsyc.2019.01.003

Geng, J. J., Won, B.-Y., & Carlisle, N. B. (2019). Distractor Ignoring: Strategies, Learning, and Passive Filtering. *Current Directions in Psychological Science*, *28*(6), 600–606. https://doi.org/10.1177/0963721419867099

Getzmann, S., Wascher, E., & Schneider, D. (2018). The role of inhibition for working memory processes: ERP evidence from a short-term storage task. *Psychophysiology*, *55*(5), e13026. https://doi.org/10.1111/psyp.13026

Geweke, F., Pokta, E., & Störmer, V. S. (2021). Spatial distance of target locations affects the time course of both endogenous and exogenous attentional deployment. *Journal of Experimental Psychology: Human Perception and Performance*, *47*(6), 774–783. https://doi.org/10.1037/xhp0000909

Giesbrecht, B., Woldorff, M. G., Song, A. W., & Mangun, G. R. (2003). Neural mechanisms of top-down control during spatial and feature attention. *NeuroImage*, *19*(3), 496–512. https://doi.org/10.1016/S1053-8119(03)00162-9

Gilbert, S. J., Frith, C. D., & Burgess, P. W. (2005). Involvement of rostral prefrontal cortex in selection between stimulus-oriented and stimulus-independent thought. *European Journal of Neuroscience*, *21*(5), 1423–1431. https://doi.org/10.1111/j.1460-9568.2005.03981.x

Gjorgieva, E., Geib, B. R., Cabeza, R., & Woldorff, M. G. (2022). The influence of imagery vividness and internally-directed attention on the neural mechanisms underlying the encoding of visual mental images into episodic memory. *Cerebral Cortex*, bhac270. https://doi.org/10.1093/cercor/bhac270

Goldman-Rakic, P. S. (1995). Cellular basis of working memory. *Neuron*, *14*(3), 477–485. https://doi.org/10.1016/0896-6273(95)90304-6

Goschy, H., Bakos, S., Müller, H. J., & Zehetleitner, M. (2014). Probability cueing of distractor locations: Both intertrial facilitation and statistical learning mediate interference reduction. *Frontiers in Psychology*, *5*. https://www.frontiersin.org/articles/10.3389/fpsyg.2014.01195

Gould, I. C., Rushworth, M. F., & Nobre, A. C. (2011). Indexing the graded allocation of visuospatial attention using anticipatory alpha oscillations. *Journal of Neurophysiology*, *105*(3), 1318–1326. https://doi.org/10.1152/jn.00653.2010

Goulden, N., Khusnulina, A., Davis, N. J., Bracewell, R. M., Bokde, A. L., McNulty, J. P., & Mullins, P. G. (2014). The salience network is responsible for switching between the default mode network and the central executive network: Replication from DCM. *NeuroImage*, *99*, 180–190. https://doi.org/10.1016/j.neuroimage.2014.05.052

Grabenhorst, M., Michalareas, G., Maloney, L. T., & Poeppel, D. (2019). The anticipation of events in time. *Nature Communications*, *10*(1), 5802. https://doi.org/10.1038/s41467-019-13849-0

Gramfort, A., Luessi, M., Larson, E., Engemann, D., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., & Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, *7*. https://www.frontiersin.org/articles/10.3389/fnins.2013.00267

Grange, J. A., & Moore, S. B. (2022). mixtur: An R package for designing, analysing, and modelling continuous report visual short-term memory studies. *Behavior Research Methods*, *54*(5), 2071–2100. https://doi.org/10.3758/s13428-021-01688-1

Greenwood, P. M., Blumberg, E. J., & Scheldrup, M. R. (2018). Hypothesis for cognitive effects of transcranial direct current stimulation: Externally- and internally-directed cognition. *Neuroscience and Biobehavioral Reviews*, *86*, 226–238. https://doi.org/10.1016/j.neubiorev.2017.11.006

Gresch, D., Boettcher, S. E. P., Nobre, A. C., & van Ede, F. (2022). Consequences of predictable temporal structure in multi-task situations. *Cognition*, *225*, 105156. https://doi.org/10.1016/j.cognition.2022.105156

Gresch, D., Boettcher, S. E. P., van Ede, F., & Nobre, A. C. (2021). Shielding working-memory representations from temporally predictable external interference. *Cognition*, *217*, 104915. https://doi.org/10.1016/j.cognition.2021.104915

Gresch, D., Boettcher, S. E. P., van Ede, F., & Nobre, A. C. (2024a). *Neural dynamics of shifting attention between perception and working-memory contents*. bioRxiv. https://doi.org/10.1101/2024.02.14.580280

Gresch, D., Boettcher, S. E. P., van Ede, F., & Nobre, A. C. (2024b). Shifting attention between perception and working memory. *Cognition*, *245*, 105731. https://doi.org/10.1016/j.cognition.2024.105731

Griffin, I. C. (2002). Multiple mechanisms of selective attention: Differential modulation of stimulus processing by attention to space or time. *Neuropsychologia*, *40*(13), 2325–2340. https://doi.org/10.1016/S0028-3932(02)00087-8

Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, *15*(8), 1176–1194. https://doi.org/10.1162/089892903322598139

Grootswagers, T., Wardle, S. G., & Carlson, T. A. (2017). Decoding Dynamic Brain Patterns from Evoked Responses: A Tutorial on Multivariate Pattern Analysis Applied to Time Series Neuroimaging Data. *Journal of Cognitive Neuroscience*, *29*(4), 677–697. https://doi.org/10.1162/jocn_a_01068

Guo, R., Wang, J., Fu, K., & Liu, Q. (2024). Exploring retro-cue effects on visual working memory: Insights from double-cue paradigm. *Frontiers in Neuroscience*, *17*. https://www.frontiersin.org/articles/10.3389/fnins.2023.1338075

Gupta, S., & Sridharan, D. (2023). *Distractors induce space-specific neural biases in visual working memory* (p. 2023.12.24.573161). bioRxiv. https://doi.org/10.1101/2023.12.24.573161

Hafed, Z. M., Lovejoy, L. P., & Krauzlis, R. J. (2011). Modulation of Microsaccades in Monkey during a Covert Visual Attention Task. *Journal of Neuroscience*, *31*(43), 15219–15230. https://doi.org/10.1523/JNEUROSCI.3106-11.2011

Hajonides, J. E., Nobre, A. C., van Ede, F., & Stokes, M. G. (2021). Decoding visual colour from scalp electroencephalography measurements. *NeuroImage*, *237*, 118030. https://doi.org/10.1016/j.neuroimage.2021.118030

Hajonides, J. E., van Ede, F., Stokes, M. G., & Nobre, A. C. (2020). Comparing the prioritization of items and feature-dimensions in visual working memory. *Journal of Vision*, *20*(8), 25. https://doi.org/10.1167/jov.20.8.25

Hakim, N., Feldmann-Wüstefeld, T., Awh, E., & Vogel, E. K. (2020). Perturbing Neural Representations of Working Memory with Task-irrelevant Interruption. *Journal of Cognitive Neuroscience*, *32*(3). https://doi.org/10.1162/jocn_a_01481

Hakim, N., Feldmann-Wüstefeld, T., Awh, E., & Vogel, E. K. (2021). Controlling the Flow of Distracting Information in Working Memory. *Cerebral Cortex (New York, N.Y. : 1991)*, *31*(7). https://doi.org/10.1093/cercor/bhab013

Hallenbeck, G. E., Sprague, T. C., Rahmati, M., Sreenivasan, K. K., & Curtis, C. E. (2021). Working memory representations in visual cortex mediate distraction effects. *Nature Communications*, *12*(1), 4714. https://doi.org/10.1038/s41467-021-24973-1

Hand, D. J., & Till, R. J. (2001). A Simple Generalisation of the Area Under the ROC Curve for Multiple Class Classification Problems. *Machine Learning*, *45*(2), 171–186. https://doi.org/10.1023/A:1010920819831

Handy, T. C., & Kam, J. W. Y. (2015). Mind Wandering and Selective Attention to the External World. *Canadian Journal of Experimental Psychology*, *69*(2), 183–189.

Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*(7238), 632–635. https://doi.org/10.1038/nature07832

Hasselmo, M. E., Bodelón, C., & Wyble, B. P. (2002). A proposed function for hippocampal theta rhythm: Separate phases of encoding and retrieval enhance reversal of prior learning. *Neural Computation*, *14*(4), 793–817. https://doi.org/10.1162/089976602317318965

Hautekiet, C., Niklaus, M., & Oberauer, K. (2024). *Susceptibility to Visual Interference in Working Memory: Different Results Depending on the Used Prioritization Strategy?* https://doi.org/10.31234/osf.io/hn9wp

Hautekiet, C., Verschooren, S., Langerock, N., & Vergauwe, E. (2023). Attentional switching between perception and memory: Examining asymmetrical switch costs. *Attention, Perception, & Psychophysics*. https://doi.org/10.3758/s13414-023-02665-w

Heideman, S. G., Rohenkohl, G., Chauvin, J. J., Palmer, C. E., van Ede, F., & Nobre, A. C. (2018). Anticipatory neural dynamics of spatial-temporal orienting of attention in younger and older adults. *NeuroImage*, *178*, 46–56. https://doi.org/10.1016/j.neuroimage.2018.05.002

Heideman, S. G., Van Ede, F., & Nobre, A. C. (2016). Early behavioural facilitation by temporal expectations in complex visual-motor sequences. *Journal of Physiology-Paris*, *110*(4), 487–496. https://doi.org/10.1016/j.jphysparis.2017.03.003

Heideman, S. G., van Ede, F., & Nobre, A. C. (2018). Temporal alignment of anticipatory motor cortical beta lateralisation in hidden visual-motor sequences. *European Journal of Neuroscience*, *48*(8), 2684–2695. https://doi.org/10.1111/ejn.13700

Helfrich, R. F., Fiebelkorn, I. C., Szczepanski, S. M., Lin, J. J., Parvizi, J., Knight, R. T., & Kastner, S. (2018). Neural Mechanisms of Sustained Attention Are Rhythmic. *Neuron*, *99*(4), 854-865.e5. https://doi.org/10.1016/j.neuron.2018.07.032

Hemmes, N. S., Brown, B. L., & Kladopoulos, C. N. (2004). Time perception with and without a concurrent nontemporal task. *Perception & Psychophysics*, *66*(2), 328–341. https://doi.org/10.3758/BF03194882

Hermann, P., Weiss, B., Knakker, B., Madurka, P., Manga, A., Nárai, Á., & Vidnyánszky, Z. (2021). Neural basis of distractor resistance during visual working memory maintenance. *NeuroImage*, *245*, 118650. https://doi.org/10.1016/j.neuroimage.2021.118650

Heuer, A., & Rolfs, M. (2021). Incidental encoding of visual information in temporal reference frames in working memory. *Cognition*, *207*, 104526. https://doi.org/10.1016/j.cognition.2020.104526

Heuer, A., & Rolfs, M. (2022). A direct comparison of attentional orienting to spatial and temporal positions in visual working memory. *Psychonomic Bulletin & Review*, *29*(1), 182–190. https://doi.org/10.3758/s13423-021-01972-3

Hickey, C., Di Lollo, V., & McDonald, J. J. (2009). Electrophysiological indices of target and distractor processing in visual search. *Journal of Cognitive Neuroscience*, *21*(4), 760–775. https://doi.org/10.1162/jocn.2009.21039

Honey, C. J., Newman, E. L., & Schapiro, A. C. (2017). Switching between internal and external modes: A multiscale learning principle. *Network Neuroscience*, *1*(4), 339–356. https://doi.org/10.1162/NETN_a_00024

Hopfinger, J. B., Buonocore, M. H., & Mangun, G. R. (2000). The neural mechanisms of top-down attentional control. *Nature Neuroscience*, *3*(3), 284–291. https://doi.org/10.1038/72999

Hsu, C.-W., Chang, C.-C., & Lin, C.-J. (2003). *A Practical Guide to Support Vector Classification*. https://www.csie.ntu.edu.tw/~cjlin/papers/guide/guide.pdf

Huang, C., Vilotijević, A., Theeuwes, J., & Donk, M. (2021). Proactive distractor suppression elicited by statistical regularities in visual search. *Psychonomic Bulletin & Review*, *28*(3), 918–927. https://doi.org/10.3758/s13423-021-01891-3

Huang, J., & Sekuler, R. (2010). Distortions in recall from visual memory: Two classes of attractors at work. *Journal of Vision*, *10*(2), 24. https://doi.org/10.1167/10.2.24

Huang, Y., Zuo, K., Asthana, M. K., Comfort, W. E., & Xu, Z. (2019). How robust and rapid can the memory-driven attentional capture be? *Current Psychology*, *38*(5), 1190–1203. https://doi.org/10.1007/s12144-019-00332-8

Ibos, G., Duhamel, J.-R., & Ben Hamed, S. (2009). The Spatial and Temporal Deployment of Voluntary Attention across the Visual Field. *PLoS ONE*, *4*(8), e6716. https://doi.org/10.1371/journal.pone.0006716

Ihssen, N., Linden, D. E. J., & Shapiro, K. L. (2010). Improving visual short-term memory by sequencing the stimulus array. *Psychonomic Bulletin & Review*, *17*(5), 680–686. https://doi.org/10.3758/PBR.17.5.680

Itskov, V., Pastalkova, E., Mizuseki, K., Buzsaki, G., & Harris, K. D. (2008). Theta-Mediated Dynamics of Spatial Information in Hippocampus. *The Journal of Neuroscience*, *28*(23), 5959–5964. https://doi.org/10.1523/JNEUROSCI.5262-07.2008

Jacob, S. N., & Nieder, A. (2014). Complementary Roles for Primate Frontal and Parietal Cortex in Guarding Working Memory from Distractor Stimuli. *Neuron*, *83*(1), 226–237. https://doi.org/10.1016/j.neuron.2014.05.009

Jacob, S. N., Stalter, M., & Nieder, A. (2016). Cell-type-specific modulation of targets and distractors by dopamine D1 receptors in primate prefrontal cortex. *Nature Communications*, *7*(1), 13218. https://doi.org/10.1038/ncomms13218

James, W. (1890). *The principles of psychology, Vol I.* Henry Holt and Co. https://doi.org/10.1037/10538-000

Jensen, O., & Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. *Frontiers in Human Neuroscience*, *4*. https://www.frontiersin.org/articles/10.3389/fnhum.2010.00186

Jha, A. P., Fabian, S. A., & Aguirre, G. K. (2004). The role of prefrontal cortex in resolving distractor interference. *Cognitive, Affective & Behavioral Neuroscience*, *4*(4), 517–527. https://doi.org/10.3758/cabn.4.4.517

Jin, W., Nobre, A. C., & van Ede, F. (2020). Temporal Expectations Prepare Visual Working Memory for Behavior. *Journal of Cognitive Neuroscience*, *32*(12), 2320–2332. https://doi.org/10.1162/jocn_a_01626

Jones, M. R. (2019). *Time Will Tell: A Theory of Dynamic Attending* (1st ed.). Oxford University Press. https://doi.org/10.1093/oso/9780190618216.001.0001

Jung, S., Yoon, Y., & Han, S. W. (2018). Working Memory-Driven Attention in Real-World Search. *Perception*, *47*(9), 966–975. https://doi.org/10.1177/0301006618791688

Kaiser, D., Oosterhof, N. N., & Peelen, M. V. (2016). The Neural Dynamics of Attentional Selection in Natural Scenes. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *36*(41), 10522–10528. https://doi.org/10.1523/JNEUROSCI.1385-16.2016

Kam, J. W. Y., Solbakk, A.-K., Endestad, T., Meling, T. R., & Knight, R. T. (2018). Lateral prefrontal cortex lesion impairs regulation of internally and externally directed attention. *NeuroImage*, *175*, 91–99. https://doi.org/10.1016/j.neuroimage.2018.03.063

Kamiński, J., & Rutishauser, U. (2020). Between persistently active and activity-silent frameworks: Novel vistas on the cellular basis of working memory. *Annals of the New York Academy of Sciences*, *1464*(1), 64–75. https://doi.org/10.1111/nyas.14213

Kang, M.-S., Hong, S. W., Blake, R., & Woodman, G. F. (2011). Visual working memory contaminates perception. *Psychonomic Bulletin & Review*, *18*(5), 860–869. https://doi.org/10.3758/s13423-011-0126-5

Karlin, L. (1959). Reaction time as a function of foreperiod duration and variability. *Journal of Experimental Psychology*, *58*(2), 185–191. https://doi.org/10.1037/h0049152

Kastner, S., Pinsk, M. A., De Weerd, P., Desimone, R., & Ungerleider, L. G. (1999). Increased activity in human visual cortex during directed attention in the absence of visual stimulation. *Neuron*, *22*(4), 751–761. https://doi.org/10.1016/s0896-6273(00)80734-5

Kelly, A. M. C., Uddin, L. Q., Biswal, B. B., Castellanos, F. X., & Milham, M. P. (2008). Competition between functional brain networks mediates behavioral variability. *NeuroImage*, *39*(1), 527–537. https://doi.org/10.1016/j.neuroimage.2007.08.008

Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching—A review. *Psychological Bulletin*, *136*(5), 849–874. https://doi.org/10.1037/a0019842

Kim, H., Daselaar, S. M., & Cabeza, R. (2010). Overlapping brain activity between episodic memory encoding and retrieval: Roles of the task-positive and task-negative networks. *NeuroImage*, *49*(1), 1045–1054. https://doi.org/10.1016/j.neuroimage.2009.07.058

King, J.-R., & Dehaene, S. (2014). Characterizing the dynamics of mental representations: The temporal generalization method. *Trends in Cognitive Sciences*, *18*(4), 203–210. https://doi.org/10.1016/j.tics.2014.01.002

Kiyonaga, A., & Egner, T. (2013). Working memory as internal attention: Toward an integrative account of internal and external selection processes. *Psychonomic Bulletin & Review*, *20*(2), 228–242. https://doi.org/10.3758/s13423-012-0359-y

Kiyonaga, A., Egner, T., & Soto, D. (2012). Cognitive control over working memory biases of selection. *Psychonomic Bulletin & Review*, *19*(4), 639–646. https://doi.org/10.3758/s13423-012-0253-7

Kok, P., Mostert, P., & de Lange, F. P. (2017). Prior expectations induce prestimulus sensory templates. *Proceedings of the National Academy of Sciences*, *114*(39), 10473–10478. https://doi.org/10.1073/pnas.1705652114

Konishi, M., McLaren, D. G., Engen, H., & Smallwood, J. (2015). Shaped by the Past: The Default Mode Network Supports Cognition that Is Independent of Immediate Perceptual Input. *PLOS ONE*, *10*(6), e0132209. https://doi.org/10.1371/journal.pone.0132209

Kragel, J. E., VanHaerents, S., Templer, J. W., Schuele, S., Rosenow, J. M., Nilakantan, A. S., & Bridge, D. J. (2020). Hippocampal theta coordinates memory processing during visual exploration. *eLife*, *9*, e52108. https://doi.org/10.7554/eLife.52108

Kreither, J., Papaioannou, O., & Luck, S. J. (2022). Active Working Memory and Simple Cognitive Operations. *Journal of Cognitive Neuroscience*, *34*(2), 313–331. https://doi.org/10.1162/jocn_a_01791

Kritzman, L., Eidelman-Rothman, M., Keil, A., Freche, D., Sheppes, G., & Levit-Binnun, N. (2022). Steady-state visual evoked potentials differentiate between internally and externally directed attention. *NeuroImage*, *254*, 119133. https://doi.org/10.1016/j.neuroimage.2022.119133

Kumle, L., Võ, M. L.-H., & Draschkow, D. (2021). Estimating power in (generalized) linear mixed models: An open introduction and tutorial in R. *Behavior Research Methods*, *53*(6), 2528–2543. https://doi.org/10.3758/s13428-021-01546-0

Kumle, L., Vo, M. L.-H., Nobre, K., & Draschkow, D. (2024). *Multifaceted consequences of visual distraction during natural behaviour*. OSF. https://doi.org/10.31234/osf.io/v7dwp

Kushleyeva, Y., Salvucci, D. D., & Lee, F. J. (2005). Deciding when to switch tasks in time-critical multitasking. *Cognitive Systems Research*, *6*(1), 41–49. https://doi.org/10.1016/j.cogsys.2004.09.005

Kwak, Y., & Curtis, C. E. (2022). Unveiling the abstract format of mnemonic representations. *Neuron*, *110*(11), 1822-1828.e5. https://doi.org/10.1016/j.neuron.2022.03.016

Lamy, D. (2005). Temporal expectations modulate attentional capture. *Psychonomic Bulletin & Review*, *12*(6), 1112–1119. https://doi.org/10.3758/BF03206452

Landau, A. N., & Fries, P. (2012). Attention Samples Stimuli Rhythmically. *Current Biology*, *22*(11), 1000–1004. https://doi.org/10.1016/j.cub.2012.03.054

Landau, A. N., Schreyer, H. M., van Pelt, S., & Fries, P. (2015). Distributed Attention Is Implemented through Theta-Rhythmic Gamma Modulation. *Current Biology*, *25*(17), 2332–2337. https://doi.org/10.1016/j.cub.2015.07.048

Landman, R., Spekreijse, H., & Lamme, V. A. F. (2003). Large capacity storage of integrated objects before change blindness. *Vision Research*, *43*(2), 149–164. https://doi.org/10.1016/s0042-6989(02)00402-9

Lara, A. H., & Wallis, J. D. (2014). Executive control processes underlying multi-item working memory. *Nature Neuroscience*, *17*(6), 876–883. https://doi.org/10.1038/nn.3702

LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding attended information in short-term memory: An EEG study. *Journal of Cognitive Neuroscience*, *25*(1), 127–142. https://doi.org/10.1162/jocn_a_00305

LaRocque, J. J., Riggall, A. C., Emrich, S. M., & Postle, B. R. (2016). Within-Category Decoding of Information in Different Attentional States in Short-Term Memory. *Cerebral Cortex*, cercor;bhw283v1. https://doi.org/10.1093/cercor/bhw283

Lebedev, M. A., Messinger, A., Kralik, J. D., & Wise, S. P. (2004). Representation of Attended Versus Remembered Locations in Prefrontal Cortex. *PLOS Biology*, *2*(11), e365. https://doi.org/10.1371/journal.pbio.0020365

Leber, A. B., Gwinn, R. E., Hong, Y., & O'Toole, R. J. (2016). Implicitly learned suppression of irrelevant spatial locations. *Psychonomic Bulletin & Review*, *23*(6), 1873–1881. https://doi.org/10.3758/s13423-016-1065-y

Lee, S.-H., Kravitz, D. J., & Baker, C. I. (2013). Goal-dependent dissociation of visual and prefrontal cortices during working memory. *Nature Neuroscience*, *16*(8), 997–999. https://doi.org/10.1038/nn.3452

Lee, T. G., & D'Esposito, M. (2012). The Dynamic Nature of Top-Down Signals Originating from Prefrontal Cortex: A Combined fMRI–TMS Study. *Journal of Neuroscience*, *32*(44), 15458–15466. https://doi.org/10.1523/JNEUROSCI.0627-12.2012

Lennie, P. (2003). The cost of cortical computation. *Current Biology: CB*, *13*(6), 493–497. https://doi.org/10.1016/s0960-9822(03)00135-0

Lepsien, J., Griffin, I. C., Devlin, J. T., & Nobre, A. C. (2005). Directing spatial attention in mental representations: Interactions between attentional orienting and working-memory load. *NeuroImage*, *26*(3), 733–743. https://doi.org/10.1016/j.neuroimage.2005.02.026

Lepsien, J., & Nobre, A. C. (2006). Cognitive control of attention in the human brain: Insights from orienting attention to mental representations. *Brain Research*, *1105*(1), 20–31. https://doi.org/10.1016/j.brainres.2006.03.033

Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2012). Neural Evidence for a Distinction between Short-term Memory and the Focus of Attention. *Journal of Cognitive Neuroscience*, *24*(1), 61–79. https://doi.org/10.1162/jocn_a_00140

Li, Q., Joo, S. J., Yeatman, J. D., & Reinecke, K. (2020). Controlling for Participants' Viewing Distance in Large-Scale, Psychophysical Online Experiments Using a Virtual Chinrest. *Scientific Reports*, *10*(1), 904. https://doi.org/10.1038/s41598-019-57204-1

Li, Q., & Saiki, J. (2014). The effects of sequential attention shifts within visual working memory. *Frontiers in Psychology*, *5*. https://www.frontiersin.org/articles/10.3389/fpsyg.2014.00965

Li, Y. P., Wang, Y., Turk-Browne, N. B., Kuhl, B. A., & Hutchinson, J. B. (2023). Perception and memory retrieval states are reflected in distributed patterns of background functional connectivity. *NeuroImage*, *276*, 120221. https://doi.org/10.1016/j.neuroimage.2023.120221

Liesefeld, H. R., Liesefeld, A. M., Sauseng, P., Jacob, S. N., & Müller, H. J. (2020). How visual working memory handles distraction: Cognitive mechanisms and electrophysiological correlates. *Visual Cognition*, *28*(5–8), 372–387. https://doi.org/10.1080/13506285.2020.1773594

Liljefors, J., Almeida, R., Rane, G., Lundström, J. N., Herman, P., & Lundqvist, M. (2023). *Distinct Functions for Beta and Alpha Bursts in Gating of Human Working Memory* (p. 2023.11.17.566386). bioRxiv. https://doi.org/10.1101/2023.11.17.566386

Lin, Y., Kong, G., & Fougnie, D. (2021). Object-based selection in visual working memory. *Psychonomic Bulletin & Review*, *28*(6), 1961–1971. https://doi.org/10.3758/s13423-021-01971-4

Liu, B., Nobre, A. C., & van Ede, F. (2022). Functional but not obligatory link between microsaccades and neural modulation by covert spatial attention. *Nature Communications*, *13*(1), Article 1. https://doi.org/10.1038/s41467-022-31217-3

Liu, R., Guo, L., Lin, X., Nie, D., Astikainen, P., & Ye, C. (2024). *Lack of attention shift cost for features in visual working memory*. OSF. https://doi.org/10.31234/osf.io/vsxka

Liu, T., Stevens, S. T., & Carrasco, M. (2007). Comparing the time course and efficacy of spatial and feature-based attention. *Vision Research*, *47*(1), 108–113. https://doi.org/10.1016/j.visres.2006.09.017

Lively, Z., Ng, G. J. P., Buetti, S., & Lleras, A. (2021). Irrelevant features of distractors in intervening visual search tasks cause active visual working memory interference – the more difficult the search task, the more interference it causes. *Attention, Perception, & Psychophysics*, *83*(6), 2410–2429. https://doi.org/10.3758/s13414-021-02318-w

Logie, R. H., Zucco, G. M., & Baddeley, A. D. (1990). Interference with visual short-term memory. *Acta Psychologica*, *75*(1), 55–74. https://doi.org/10.1016/0001-6918(90)90066-O

Lorenc, E. S., Mallett, R., & Lewis-Peacock, J. A. (2021). Distraction in Visual Working Memory: Resistance is Not Futile. *Trends in Cognitive Sciences*, *25*(3), 228–239. https://doi.org/10.1016/j.tics.2020.12.004

Lorenc, E. S., Sreenivasan, K. K., Nee, D. E., Vandenbroucke, A. R. E., & D'Esposito, M. (2018). Flexible Coding of Visual Working Memory Representations during Distraction. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *38*(23), 5267–5276. https://doi.org/10.1523/JNEUROSCI.3061-17.2018

Los, S. A. (1996). On the origin of mixing costs: Exploring information processing in pure and mixed blocks of trials. *Acta Psychologica*, *94*(2), 145–188. https://doi.org/10.1016/0001-6918(95)00050-X

Los, S. A. (2010). Foreperiod and sequential effects: Theory and data. In A. C. Nobre & J. T. Coull (Eds.), *Attention and Time* (1st ed., pp. 289–302). Oxford University PressOxford. https://doi.org/10.1093/acprof:oso/9780199563456.003.0021

Los, S. A., Knol, D. L., & Boers, R. M. (2001). The foreperiod effect revisited: Conditioning as a basis for nonspecific preparation11Parts of this article were presented on the symposium "Looking for stages", organized on the occasion of Andries F. Sanders becoming Professor Emeritus. *Acta Psychologica*, *106*(1–2), 121–145. https://doi.org/10.1016/S0001-6918(00)00029-9

Los, S. A., Kruijne, W., & Meeter, M. (2014). Outlines of a multiple trace theory of temporal preparation. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.01058

Los, S. A., Kruijne, W., & Meeter, M. (2017). Hazard versus history: Temporal preparation is driven by past experience. *Journal of Experimental Psychology: Human Perception and Performance*, *43*(1), 78–88. https://doi.org/10.1037/xhp0000279

Lout, E., Golomb, J., & Dube, B. (2023). *Items in visual working memory are more susceptible to visual interference while in-use.* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/28xds

Luce, R. D. (1991). *Response Times: Their Role in Inferring Elementary Mental Organization*. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780195070019.001.0001

Luck, S. J., Chelazzi, L., Hillyard, S. A., & Desimone, R. (1997). Neural mechanisms of spatial selective attention in areas V1, V2, and V4 of macaque visual cortex. *Journal of Neurophysiology*, *77*(1). https://doi.org/10.1152/jn.1997.77.1.24

Macedo-Pascual, J., Capilla, A., Campo, P., Hinojosa, J. A., & Poch, C. (2023). Selection within working memory impairs perceptual detection. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-022-02238-2

Magnussen, S., & Greenlee, M. W. (1992). Retention and disruption of motion information in visual short-term memory. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *18*(1), 151–156. https://doi.org/10.1037//0278-7393.18.1.151

Makovski, T. (2019). Preparing for distraction: Attention is enhanced prior to the presentation of distractors. *Journal of Experimental Psychology: General*, *148*(2), 221–236. https://doi.org/10.1037/xge0000509

Makovski, T., & Jiang, Y. V. (2007). Distributing versus focusing attention in visual short-term memory. *Psychonomic Bulletin & Review*, *14*(6), 1072–1078. https://doi.org/10.3758/BF03193093

Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 369–380. https://doi.org/10.1037/0278-7393.34.2.369

Mallett, R., & Lewis-Peacock, J. A. (2018). Behavioral decoding of working memory items inside and outside the focus of attention. *Annals of the New York Academy of Sciences*, *1424*(1), 256–267. https://doi.org/10.1111/nyas.13647

Mallett, R., & Lewis-Peacock, J. A. (2019). Working memory prioritization impacts neural recovery from distraction. *Cortex*, *121*, 225–238. https://doi.org/10.1016/j.cortex.2019.08.019

Mallett, R., Mummaneni, A., & Lewis-Peacock, J. A. (2020). Distraction biases working memory for faces. *Psychonomic Bulletin & Review*, *27*(2), 350–356. https://doi.org/10.3758/s13423-019-01707-5

Marini, F., Scott, J., Aron, A. R., & Ester, E. F. (2017). Task-irrelevant distractors in the delay period interfere selectively with visual short-term memory for spatial locations. *Attention, Perception, & Psychophysics*, *79*(5), 1384–1392. https://doi.org/10.3758/s13414-017-1320-y

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

Marr, D. (1971). Simple memory: A theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *262*(841), 23–81. https://doi.org/10.1098/rstb.1971.0078

Masse, N. Y., Rosen, M. C., & Freedman, D. J. (2020). Reevaluating the Role of Persistent Neural Activity in Short-Term Memory. *Trends in Cognitive Sciences*, *24*(3), 242–258. https://doi.org/10.1016/j.tics.2019.12.014

Mathewson, K. E., Fabiani, M., Gratton, G., Beck, D. M., & Lleras, A. (2010). Rescuing stimuli from invisibility: Inducing a momentary release from visual masking with pre-target entrainment. *Cognition*, *115*(1), 186–191. https://doi.org/10.1016/j.cognition.2009.11.010

Maunsell, J. H. R., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, *29*(6), 317–322. https://doi.org/10.1016/j.tins.2006.04.001

Mayr, U., Kuhns, D., & Hubbard, J. (2014). Long-term memory and the control of attentional control. *Cognitive Psychology*, *72*, 1–26. https://doi.org/10.1016/j.cogpsych.2014.02.001

McFarland, D. J. (2000). Mu and beta rhythm topographies during motor imagery and actual movements. *Brain Topography*, *12*(3), 177–186. https://doi.org/10.1023/A:1023437823106

McNab, F., & Dolan, R. J. (2014). Dissociating distractor-filtering at encoding and during maintenance. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(3), 960–967. https://doi.org/10.1037/a0036013

McNab, F., & Klingberg, T. (2008). Prefrontal cortex and basal ganglia control access to working memory. *Nature Neuroscience*, *11*(1), Article 1. https://doi.org/10.1038/nn2024

Menceloglu, M., Grabowecky, M., & Suzuki, S. (2017). Temporal expectation weights visual signals over auditory signals. *Psychonomic Bulletin & Review*, *24*(2), 416–422. https://doi.org/10.3758/s13423-016-1109-3

Merrikhi, Y., Clark, K., Albarran, E., Parsa, M., Zirnsak, M., Moore, T., & Noudoost, B. (2017). Spatial working memory alters the efficacy of input to visual cortex. *Nature Communications*, *8*(1), Article 1. https://doi.org/10.1038/ncomms15041

Meyers, E. M. (2018). Dynamic population coding and its relationship to working memory. *Journal of Neurophysiology*, *120*(5), 2260–2268. https://doi.org/10.1152/jn.00225.2018

Meyers, E. M., Freedman, D. J., Kreiman, G., Miller, E. K., & Poggio, T. (2008). Dynamic Population Coding of Category Information in Inferior Temporal and Prefrontal Cortex. *Journal of Neurophysiology*, *100*(3), 1407–1419. https://doi.org/10.1152/jn.90248.2008

Miller, B. T., Vytlacil, J., Fegen, D., Pradhan, S., & D'Esposito, M. (2011). The Prefrontal Cortex Modulates Category Selectivity in Human Extrastriate Cortex. *Journal of Cognitive Neuroscience*, *23*(1), 1–10. https://doi.org/10.1162/jocn.2010.21516

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. *Journal of Neuroscience*, *16*(16), 5154–5167. https://doi.org/10.1523/JNEUROSCI.16-16-05154.1996

Miller, E. K., Li, L., & Desimone, R. (1993). Activity of neurons in anterior inferior temporal cortex during a short- term memory task. *Journal of Neuroscience*, *13*(4), 1460–1478. https://doi.org/10.1523/JNEUROSCI.13-04-01460.1993

Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*(2), 463–475. https://doi.org/10.1016/j.neuron.2018.09.023

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior.* Henry Holt and Co. https://doi.org/10.1037/10039-000

Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*(1), 99–115.

Miniussi, C., Marzi, C. A., & Nobre, A. C. (2005). Modulation of brain activity by selective task sets observed using event-related potentials. *Neuropsychologia*, *43*(10), 1514–1528. https://doi.org/10.1016/j.neuropsychologia.2004.12.014

Miniussi, C., Wilding, E. L., Coull, J. T., & Nobre, A. C. (1999). Orienting attention in time: Modulation of brain potentials. *Brain: A Journal of Neurology*, *122*(8), 1507–1518. https://doi.org/10.1093/brain/122.8.1507

Mishra, J., Zanto, T., Nilakantan, A., & Gazzaley, A. (2013). Comparable mechanisms of working memory interference by auditory and visual motion in youth and aging. *Neuropsychologia*, *51*(10), 1896–1906. https://doi.org/10.1016/j.neuropsychologia.2013.06.011

Mok, R. M., Myers, N. E., Wallis, G., & Nobre, A. C. (2016). Behavioral and Neural Markers of Flexible Attention over Working Memory in Aging. *Cerebral Cortex*, *26*(4), 1831–1842. https://doi.org/10.1093/cercor/bhw011

Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic Theory of Working Memory. *Science*, *319*(5869), 1543–1546. https://doi.org/10.1126/science.1150769

Moon, J., & Anderson, J. R. (2013). Timing in multitasking: Memory contamination and time pressure bias. *Cognitive Psychology*, *67*(1–2), 26–54. https://doi.org/10.1016/j.cogpsych.2013.06.001

Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science (New York, N.Y.)*, *229*(4715), 782–784. https://doi.org/10.1126/science.4023713

Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005). *Tutorials in Quantitative Methods for Psychology*, *4*(2), 61–64. https://doi.org/10.20982/tqmp.04.2.p061

Morgan, A. (2014). Representations gone mental. *Synthese*, *191*(2), 213–244. https://doi.org/10.1007/s11229-013-0328-7

Mostert, P., Albers, A. M., Brinkman, L., Todorova, L., Kok, P., & Lange, F. P. de. (2018). Eye Movement-Related Confounds in Neural Decoding of Visual Working Memory Representations. *eNeuro*, *5*(4). https://doi.org/10.1523/ENEURO.0401-17.2018

Muhle-Karbe, P. S., Myers, N. E., & Stokes, M. G. (2021). A Hierarchy of Functional States in Working Memory. *Journal of Neuroscience*, *41*(20), 4461–4475. https://doi.org/10.1523/JNEUROSCI.3104-20.2021

Murray, A. M., Nobre, A. C., & Stokes, M. G. (2011). Markers of preparatory attention predict visual short-term memory performance. *Neuropsychologia*, *49*(6), 1458–1465. https://doi.org/10.1016/j.neuropsychologia.2011.02.016

Murray, J. D., Bernacchia, A., Roy, N. A., Constantinidis, C., Romo, R., & Wang, X.-J. (2017). Stable population coding for working memory coexists with heterogeneous neural dynamics in prefrontal cortex. *Proceedings of the National Academy of Sciences*, *114*(2), 394–399. https://doi.org/10.1073/pnas.1619449114

Myers, N. E., Stokes, M. G., & Nobre, A. C. (2017). Prioritizing Information during Working Memory: Beyond Sustained Internal Attention. *Trends in Cognitive Sciences*, *21*(6), 449–461. https://doi.org/10.1016/j.tics.2017.03.010

Myers, N. E., Stokes, M. G., Walther, L., & Nobre, A. C. (2014). Oscillatory Brain State Predicts Variability in Working Memory. *Journal of Neuroscience*, *34*(23), 7735–7743. https://doi.org/10.1523/JNEUROSCI.4741-13.2014

Myers, N. E., Walther, L., Wallis, G., Stokes, M. G., & Nobre, A. C. (2015). Temporal dynamics of attention during encoding versus maintenance of working memory: Complementary views from event-related potentials and alpha-band oscillations. *Journal of Cognitive Neuroscience*, *27*(3), 492–508. https://doi.org/10.1162/jocn_a_00727

Narhi-Martinez, W., Dube, B., Chen, J., Leber, A. B., & Golomb, J. D. (2024). Suppression of a salient distractor protects the processing of target features. *Psychonomic Bulletin & Review*, *31*(1), 223–233. https://doi.org/10.3758/s13423-023-02339-6

Narhi-Martinez, W., Dube, B., & Golomb, J. D. (2023). Attention as a multi-level system of weights and balances. *Wiley Interdisciplinary Reviews. Cognitive Science*, *14*(1), e1633. https://doi.org/10.1002/wcs.1633

Nasrawi, R., Boettcher, S. E. P., & van Ede, F. (2023). Prospection of Potential Actions during Visual Working Memory Starts Early, Is Flexible, and Predicts Behavior. *Journal of Neuroscience*, *43*(49), 8515–8524. https://doi.org/10.1523/JNEUROSCI.0709-23.2023

Nasrawi, R., Mautner-Rohde, M., & van Ede, F. (2024). *Memory load influences our preparedness to act on visual representations in working memory without affecting their accessibility* (p. 2024.02.23.581707). bioRxiv. https://doi.org/10.1101/2024.02.23.581707

Nasrawi, R., & van Ede, F. (2022). Planning the Potential Future during Multi-item Visual Working Memory. *Journal of Cognitive Neuroscience*, *34*(8), 1534–1546. https://doi.org/10.1162/jocn_a_01875

Nee, D. E., & Jonides, J. (2009). Common and distinct neural correlates of perceptual and memorial selection. *NeuroImage*, *45*(3), 963–975. https://doi.org/10.1016/j.neuroimage.2009.01.005

Neely, J. H. (1976). Semantic priming and retrieval from lexical memory: Evidence for facilitatory and inhibitory processes. *Memory & Cognition*, *4*(5), 648–654. https://doi.org/10.3758/BF03213230

Neisser, U. (1967). *Cognitive psychology*. Meredith.

Nelissen, N., Stokes, M., Nobre, A. C., & Rushworth, M. F. S. (2013). Frontal and parietal cortical interactions with distributed visual representations during selective attention and action selection. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *33*(42), 16443–16458. https://doi.org/10.1523/JNEUROSCI.2625-13.2013

Nemes, V. A., Parry, N. R. A., Whitaker, D., & McKeefry, D. J. (2012). The retention and disruption of color information in human short-term visual memory. *Journal of Vision*, *12*(1), 26. https://doi.org/10.1167/12.1.26

Neuper, C., Wörtz, M., & Pfurtscheller, G. (2006). ERD/ERS patterns reflecting sensorimotor activation and deactivation. In *Progress in Brain Research* (Vol. 159, pp. 211–222). Elsevier. https://doi.org/10.1016/S0079-6123(06)59014-4

Niemi, P., & Näätänen, R. (1981). Foreperiod and simple reaction time. *Psychological Bulletin*, *89*(1), 133–162. https://doi.org/10.1037/0033-2909.89.1.133

Nobre, A. C. (2001). Orienting attention to instants in time. *Neuropsychologia*, *39*(12), 1317–1328. https://doi.org/10.1016/S0028-3932(01)00120-8

Nobre, A. C. (2018). Attention. In J. T. Wixted (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (1st ed., pp. 1–75). Wiley. https://doi.org/10.1002/9781119170174.epcn206

Nobre, A. C., Coull, J. T., Maquet, P., Frith, C. D., Vandenberghe, R., & Mesulam, M. M. (2004). Orienting Attention to Locations in Perceptual Versus Mental Representations. *Journal of Cognitive Neuroscience*, *16*(3), 363–373. https://doi.org/10.1162/089892904322926700

Nobre, A. C., & Kastner, S. (2014). Attention: Time Capsule 2013. In A. C. Nobre & S. Kastner (Eds.), *The Oxford Handbook of Attention* (p. 0). Oxford University Press. https://doi.org/10.1093/oxfordhb/9780199675111.013.040

Nobre, A. C., & Stokes, M. G. (2019). Premembering Experience: A Hierarchy of Time-Scales for Proactive Attention. *Neuron*, *104*(1), 132–146. https://doi.org/10.1016/j.neuron.2019.08.030

Nobre, A. C., & van Ede, F. (2018). Anticipated moments: Temporal structure in attention. *Nature Reviews Neuroscience*, *19*(1), 34–48. https://doi.org/10.1038/nrn.2017.141

Nobre, A. C., & van Ede, F. (2023). Attention in flux. *Neuron*, *111*(7), 971–986. https://doi.org/10.1016/j.neuron.2023.02.032

Noonan, M. P., Crittenden, B. M., Jensen, O., & Stokes, M. G. (2018). Selective inhibition of distracting input. *Behavioural Brain Research*, *355*, 36–47. https://doi.org/10.1016/j.bbr.2017.10.010

Oberauer, K. (2019). *Working Memory and Attention – A Conceptual Analysis and Review* (1). *2*(1), Article 1. https://doi.org/10.5334/joc.58

O'Keefe, J., & Recce, M. L. (1993). Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus*, *3*(3), 317–330. https://doi.org/10.1002/hipo.450030307

Olivers, C. N. L. (2009). What drives memory-driven attentional capture? The effects of memory type, display type, and search type. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1275–1291. https://doi.org/10.1037/a0013896

Olivers, C. N. L., Meijer, F., & Theeuwes, J. (2006). Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *Journal of Experimental Psychology. Human Perception and Performance*, *32*(5), 1243–1265. https://doi.org/10.1037/0096-1523.32.5.1243

Olivers, C. N. L., Peters, J., Houtkamp, R., & Roelfsema, P. R. (2011). Different states in visual working memory: When it guides attention and when it does not. *Trends in Cognitive Sciences*, S1364661311000854. https://doi.org/10.1016/j.tics.2011.05.004

Olivers, C. N. L., & Roelfsema, P. R. (2020). Attention for action in visual working memory. *Cortex*, *131*, 179–194. https://doi.org/10.1016/j.cortex.2020.07.011

Olmos-Solis, K., van Loon, A. M., Los, S. A., & Olivers, C. N. L. (2017). Oculomotor measures reveal the temporal dynamics of preparing for search. In *Progress in Brain Research* (Vol. 236, pp. 1–23). Elsevier. https://doi.org/10.1016/bs.pbr.2017.07.003

Olson, I. R., & Chun, M. M. (2001). Temporal contextual cuing of visual attention. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *27*(5), 1299–1313. https://doi.org/10.1037/0278-7393.27.5.1299

O'Reilly, J. X., McCarthy, K. J., Capizzi, M., & Nobre, A. C. (2008). Acquisition of the Temporal and Ordinal Structure of Movement Sequences in Incidental Learning. *Journal of Neurophysiology*, *99*(5), 2731–2735. https://doi.org/10.1152/jn.01141.2007

Panichello, M. F., & Buschman, T. J. (2021). Shared mechanisms underlie the control of working memory and attention. *Nature*, *592*(7855), 601–605. https://doi.org/10.1038/s41586-021-03390-w

Parthasarathy, A., Tang, C., Herikstad, R., Cheong, L. F., Yen, S.-C., & Libedinsky, C. (2019). Time-invariant working memory representations in the presence of code-morphing in the lateral prefrontal cortex. *Nature Communications*, *10*(1). Scopus. https://doi.org/10.1038/s41467-019-12841-y

Pasternak, T., & Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nature Reviews Neuroscience*, *6*(2), 97–107. https://doi.org/10.1038/nrn1603

Pasternak, T., Lui, L. L., & Spinelli, P. M. (2015). Unilateral Prefrontal Lesions Impair Memory-Guided Comparisons of Contralateral Visual Motion. *Journal of Neuroscience*, *35*(18), 7095–7105. https://doi.org/10.1523/JNEUROSCI.5265-14.2015

Pasternak, T., & Zaksas, D. (2003). Stimulus Specificity and Temporal Dynamics of Working Memory for Visual Motion. *Journal of Neurophysiology*, *90*(4), 2757–2762. https://doi.org/10.1152/jn.00422.2003

Payne, L., Guillory, S., & Sekuler, R. (2013). Attention-modulated Alpha-band Oscillations Protect against Intrusion of Irrelevant Information. *Journal of Cognitive Neuroscience*, *25*(9), 1463–1476. https://doi.org/10.1162/jocn_a_00395

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Müller, A., Nothman, J., Louppe, G., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, É. (2018). *Scikit-learn: Machine Learning in Python* (arXiv:1201.0490). arXiv. https://doi.org/10.48550/arXiv.1201.0490

Peelen, M. V., & Kastner, S. (2011). A neural basis for real-world visual search in human occipitotemporal cortex. *Proceedings of the National Academy of Sciences*, *108*(29), 12125–12130. https://doi.org/10.1073/pnas.1101042108

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, *51*(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology*, *72*(2), 184–187. https://doi.org/10.1016/0013-4694(89)90180-6

Peters, B., Kaiser, J., Rahm, B., & Bledowski, C. (2015). Activity in Human Visual and Parietal Cortex Reveals Object-Based Attention in Working Memory. *Journal of Neuroscience*, *35*(8), 3360–3369. https://doi.org/10.1523/JNEUROSCI.3795-14.2015

Peters, B., Kaiser, J., Rahm, B., & Bledowski, C. (2021). Object-based attention prioritizes working memory contents at a theta rhythm. *Journal of Experimental Psychology. General*, *150*(6), 1250–1256. https://doi.org/10.1037/xge0000994

Polti, I., Martin, B., & Van Wassenhove, V. (2018). The effect of attention and working memory on the estimation of elapsed time. *Scientific Reports*, *8*(1), 6690. https://doi.org/10.1038/s41598-018-25119-y

Pomper, U., & Ansorge, U. (2021). Theta-Rhythmic Oscillation of Working Memory Performance. *Psychological Science*, *32*(11), 1801–1810. https://doi.org/10.1177/09567976211013045

Posit team. (2022). *RStudio: Integrated Development Environment for R*. Posit Software, PBC. http://www.posit.co/

Poskanzer, C., & Aly, M. (2023). Switching between External and Internal Attention in Hippocampal Networks. *Journal of Neuroscience*, *43*(38), 6538–6552. https://doi.org/10.1523/JNEUROSCI.0029-23.2023

Posner, M. I. (1980). Orienting of Attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. https://doi.org/10.1080/00335558008248231

Posner, M. I., & Snyder, C. R. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information Processing and Cognition: The Loyola Symposium* (pp. 54–85). Lawrence Erlbaum.

Postle, B. R. (2005). Delay-period activity in prefrontal cortex: One function is sensory gating. *Journal of Cognitive Neuroscience*, *17*(11), 1679–1690. https://doi.org/10.1162/089892905774589208

Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*(1), 23–38. https://doi.org/10.1016/j.neuroscience.2005.06.005

Puri, A. M., Wojciulik, E., & Ranganath, C. (2009). Category expectation modulates baseline and stimulus-evoked activity in human inferotemporal cortex. *Brain Research*, *1301*, 89–99. https://doi.org/10.1016/j.brainres.2009.08.085

Quax, S. C., Dijkstra, N., van Staveren, M. J., Bosch, S. E., & van Gerven, M. A. J. (2019). Eye movements explain decodability during perception and cued attention in MEG. *NeuroImage*, *195*, 444–453. https://doi.org/10.1016/j.neuroimage.2019.03.069

Quinn, A. J., van Es, M. W. J., Gohil, C., & Woolrich, M. W. (2022). *OHBA Software Library in Python (OSL)* (0.1.1) [Computer software]. Zenodo. https://doi.org/10.5281/ZENODO.6875060

Rademaker, R. L., Bloem, I. M., De Weerd, P., & Sack, A. T. (2015). The impact of interference on short-term memory for visual orientation. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(6), 1650–1665. https://doi.org/10.1037/xhp0000110

Rademaker, R. L., Chunharas, C., & Serences, J. T. (2019). Coexisting representations of sensory and mnemonic information in human visual cortex. *Nature Neuroscience*, *22*(8), 1336–1344. https://doi.org/10.1038/s41593-019-0428-x

Raichle, M. E., MacLeod, A. M., Snyder, A. Z., Powers, W. J., Gusnard, D. A., & Shulman, G. L. (2001). A default mode of brain function. *Proceedings of the National Academy of Sciences*, *98*(2), 676–682. https://doi.org/10.1073/pnas.98.2.676

Reder, L. M., Weber, K., Shang, J., & Vanyukov, P. M. (2003). The adaptive character of the attentional system: Statistical sensitivity in a target localization task. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(3), 631–649. https://doi.org/10.1037/0096-1523.29.3.631

Rerko, L., & Oberauer, K. (2013). Focused, unfocused, and defocused information in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*, 1075–1096. https://doi.org/10.1037/a0031172

Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measured with functional magnetic resonance imaging. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *32*(38), 12990–12998. https://doi.org/10.1523/JNEUROSCI.1892-12.2012

Rohenkohl, G., Cravo, A. M., Wyart, V., & Nobre, A. C. (2012). Temporal Expectation Improves the Quality of Sensory Information. *Journal of Neuroscience*, *32*(24), 8424–8428. https://doi.org/10.1523/JNEUROSCI.0804-12.2012

Rohenkohl, G., Gould, I. C., Pessoa, J., & Nobre, A. C. (2014). Combining spatial and temporal expectations to improve visual perception. *Journal of Vision*, *14*(4). https://doi.org/10.1167/14.4.8

Rose, N. S., LaRocque, J. J., Riggall, A. C., Gosseries, O., Starrett, M. J., Meyering, E. E., & Postle, B. R. (2016). Reactivation of latent working memories with transcranial magnetic stimulation. *Science*, *354*(6316), 1136–1139. https://doi.org/10.1126/science.aah7011

Rosner, B. (1983). Percentage Points for a Generalized ESD Many-Outlier Procedure. *Technometrics*, *25*(2), 165–172. https://doi.org/10.1080/00401706.1983.10487848

Sakai, K., Rowe, J. B., & Passingham, R. E. (2002). Parahippocampal Reactivation Signal at Retrieval after Interruption of Rehearsal. *The Journal of Neuroscience*, *22*(15), 6315–6320. https://doi.org/10.1523/JNEUROSCI.22-15-06315.2002

Salet, J. M., Kruijne, W., van Rijn, H., Los, S. A., & Meeter, M. (2022). FMTP: A unifying computational framework of temporal preparation across time scales. *Psychological Review*, *129*(5), 911–948. https://doi.org/10.1037/rev0000356

Salmelin, R., & Hari, R. (1994). Spatiotemporal characteristics of sensorimotor neuromagnetic rhythms related to thumb movement. *Neuroscience*, *60*(2), 537–550. https://doi.org/10.1016/0306-4522(94)90263-1

Sasin, E., & Fougnie, D. (2020). Memory-driven capture occurs for individual features of an object. *Scientific Reports*, *10*(1), Article 1. https://doi.org/10.1038/s41598-020-76431-5

Sassenhagen, J., & Draschkow, D. (2019). Cluster-based permutation tests of MEG/EEG data do not establish significance of effect latency or location. *Psychophysiology*, *56*(6), e13335. https://doi.org/10.1111/psyp.13335

Sauter, M., Liesefeld, H. R., Zehetleitner, M., & Müller, H. J. (2018). Region-based shielding of visual search from salient distractors: Target detection is impaired with same- but not different-dimension distractors. *Attention, Perception, & Psychophysics*, *80*(3), 622–642. https://doi.org/10.3758/s13414-017-1477-4

Sawaki, R., & Luck, S. J. (2010). Capture versus suppression of attention by salient singletons: Electrophysiological evidence for an automatic attend-to-me signal. *Attention, Perception & Psychophysics*, *72*(6), 1455–1470. https://doi.org/10.3758/APP.72.6.1455

Schneider, D., Barth, A., Getzmann, S., & Wascher, E. (2017). On the neural mechanisms underlying the protective function of retroactive cuing against perceptual interference: Evidence by event-related potentials of the EEG. *Biological Psychology*, *124*, 47–56. https://doi.org/10.1016/j.biopsycho.2017.01.006

Schneider, D., Barth, A., & Wascher, E. (2017). On the contribution of motor planning to the retroactive cuing benefit in working memory: Evidence by mu and beta oscillatory activity in the EEG. *NeuroImage*, *162*, 73–85. https://doi.org/10.1016/j.neuroimage.2017.08.057

Schneider, D., Herbst, S. K., Klatt, L.-I., & Wöstmann, M. (2022). Target enhancement or distractor suppression? Functionally distinct alpha oscillations form the basis of attention. *The European Journal of Neuroscience*, *55*(11–12), 3256–3265. https://doi.org/10.1111/ejn.15309

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*(1), 1–66. https://doi.org/10.1037/0033-295X.84.1.1

Schooler, J. W., Smallwood, J., Christoff, K., Handy, T. C., Reichle, E. D., & Sayette, M. A. (2011). Meta-awareness, perceptual decoupling and the wandering mind. *Trends in Cognitive Sciences*, *15*(7), 319–326. https://doi.org/10.1016/j.tics.2011.05.006

Schroeder, S. C. Y., Ball, F., & Busch, N. A. (2018). The role of alpha oscillations in distractor inhibition during memory retention. *European Journal of Neuroscience*, *48*(7), 2516–2526. https://doi.org/10.1111/ejn.13852

Serences, J. T. (2016). Neural mechanisms of information storage in visual short-term memory. *Vision Research*, *128*, 53–67. https://doi.org/10.1016/j.visres.2016.09.010

Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*(2), 207–214. https://doi.org/10.1111/j.1467-9280.2009.02276.x

Servais, A., Hurter, C., & Barbeau, E. J. (2023). Attentional switch to memory: An early and critical phase of the cognitive cascade allowing autobiographical memory retrieval. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02270-w

Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., & Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*, *146*(3), 1082–1108. https://doi.org/10.1016/j.neuroscience.2006.12.072

Shalev, N., Boettcher, S. E. P., & Nobre, A. C. (2019). Right time, right place: Implicit learning of target onsets in a visual search task. *Journal of Vision*, *19*(10), 255b. https://doi.org/10.1167/19.10.255b

Shalev, N., & Nobre, A. C. (2022). Eyes wide open: Regulation of arousal by temporal expectations. *Cognition*, *224*, 105062. https://doi.org/10.1016/j.cognition.2022.105062

Sherman, B. E., & Turk-Browne, N. B. (2020). Statistical prediction of the future impairs episodic encoding of the present. *Proceedings of the National Academy of Sciences*, *117*(37), 22760–22770. https://doi.org/10.1073/pnas.2013291117

Shibata, K., Yamagishi, N., Goda, N., Yoshioka, T., Yamashita, O., Sato, M., & Kawato, M. (2008). The Effects of Feature Attention on Prestimulus Cortical Activity in the Human Visual System. *Cerebral Cortex*, *18*(7), 1664–1675. https://doi.org/10.1093/cercor/bhm194

Shulman, G. L., Fiez, J. A., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. *Journal of Cognitive Neuroscience*, *9*(5), 648–663. https://doi.org/10.1162/jocn.1997.9.5.648

Silver, M. A., Ress, D., & Heeger, D. J. (2007). Neural correlates of sustained spatial attention in human early visual cortex. *Journal of Neurophysiology*, *97*(1), 229–237. https://doi.org/10.1152/jn.00677.2006

Smallwood, J. (2013). Distinguishing how from why the mind wanders: A process–occurrence framework for self-generated mental activity. *Psychological Bulletin*, *139*(3), 519–535. https://doi.org/10.1037/a0030010

Snyder, A. C., & Foxe, J. J. (2010). Anticipatory attentional suppression of visual features indexed by oscillatory alpha-band power increases: A high-density electrical mapping study. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *30*(11), 4024–4032. https://doi.org/10.1523/JNEUROSCI.5684-09.2010

Soto, D., Heinke, D., Humphreys, G. W., & Blanco, M. J. (2005). Early, involuntary top-down guidance of attention from working memory. *Journal of Experimental Psychology. Human Perception and Performance*, *31*(2), 248–261. https://doi.org/10.1037/0096-1523.31.2.248

Soto, D., Hodsoll, J., Rotshtein, P., & Humphreys, G. W. (2008). Automatic guidance of attention from working memory. *Trends in Cognitive Sciences*, *12*(9), 342–348. https://doi.org/10.1016/j.tics.2008.05.007

Souza, A. S., & Oberauer, K. (2016). In search of the focus of attention in working memory: 13 years of the retro-cue effect. *Attention, Perception, & Psychophysics*, *78*(7), 1839–1860. https://doi.org/10.3758/s13414-016-1108-5

Souza, A. S., & Oberauer, K. (2017). The contributions of visual and central attention to visual working memory. *Attention, Perception, & Psychophysics*, *79*(7), 1897–1916. https://doi.org/10.3758/s13414-017-1357-y

Spaak, E., Watanabe, K., Funahashi, S., & Stokes, M. G. (2017). Stable and Dynamic Coding for Working Memory in Primate Prefrontal Cortex. *Journal of Neuroscience*, *37*(27), 6503–6516. https://doi.org/10.1523/JNEUROSCI.3364-16.2017

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring Latent Visual Working Memory Representations in Human Cortex. *Neuron*, *91*(3), 694–707. https://doi.org/10.1016/j.neuron.2016.07.006

Spreng, R. N. (2012). The Fallacy of a "Task-Negative" Network. *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00145

Spreng, R. N., Mar, R. A., & Kim, A. S. N. (2009). The Common Neural Basis of Autobiographical Memory, Prospection, Navigation, Theory of Mind, and the Default Mode: A Quantitative Meta-analysis. *Journal of Cognitive Neuroscience*, *21*(3), 489–510. https://doi.org/10.1162/jocn.2008.21029

Spreng, R. N., Sepulcre, J., Turner, G. R., Stevens, W. D., & Schacter, D. L. (2013). Intrinsic Architecture Underlying the Relations among the Default, Dorsal Attention, and Frontoparietal Control Networks of the Human Brain. *Journal of Cognitive Neuroscience*, *25*(1), 74–86. https://doi.org/10.1162/jocn_a_00281

Spreng, R. N., Stevens, W. D., Chamberlain, J. P., Gilmore, A. W., & Schacter, D. L. (2010). Default network activity, coupled with the frontoparietal control network, supports goal-directed cognition. *NeuroImage*, *53*(1), 303–317. https://doi.org/10.1016/j.neuroimage.2010.06.016

Sreenivasan, K. K., Curtis, C. E., & D'Esposito, M. (2014). Revisiting the role of persistent neural activity during working memory. *Trends in Cognitive Sciences*, *18*(2), 82–89. https://doi.org/10.1016/j.tics.2013.12.001

Sreenivasan, K. K., & Jha, A. P. (2007). Selective Attention Supports Working Memory Maintenance by Modulating Perceptual Processing of Distractors. *Journal of Cognitive Neuroscience*, *19*(1), 32–41. https://doi.org/10.1162/jocn.2007.19.1.32

Sreenivasan, K. K., Vytlacil, J., & D'Esposito, M. (2014). Distributed and Dynamic Storage of Working Memory Stimulus Information in Extrastriate Cortex. *Journal of Cognitive Neuroscience*, *26*(5), 1141–1153. https://doi.org/10.1162/jocn_a_00556

Sridharan, D., Levitin, D. J., & Menon, V. (2008). A critical role for the right fronto-insular cortex in switching between central-executive and default-mode networks. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(34), 12569–12574. https://doi.org/10.1073/pnas.0800005105

Steel, A., Silson, E. H., Garcia, B. D., & Robertson, C. E. (2024). A retinotopic code structures the interaction between perception and memory systems. *Nature Neuroscience*, 1–9. https://doi.org/10.1038/s41593-023-01512-3

Steinborn, M. B., Rolke, B., Bratzke, D., & Ulrich, R. (2008). Sequential effects within a short foreperiod context: Evidence for the conditioning account of temporal preparation. *Acta Psychologica*, *129*(2), 297–307. https://doi.org/10.1016/j.actpsy.2008.08.005

Stilwell, B. T., Bahle, B., & Vecera, S. P. (2019). Feature-based statistical regularities of distractors modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(3), 419–433. https://doi.org/10.1037/xhp0000613

Stokes, M. G. (2015). 'Activity-silent' working memory in prefrontal cortex: A dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394–405. https://doi.org/10.1016/j.tics.2015.05.004

Stokes, M. G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., & Duncan, J. (2013). Dynamic Coding for Cognitive Control in Prefrontal Cortex. *Neuron*, *78*(2), 364–375. https://doi.org/10.1016/j.neuron.2013.01.039

Stokes, M. G., Thompson, R., Nobre, A. C., & Duncan, J. (2009). Shape-specific preparatory activity mediates attention to targets in human visual cortex. *Proceedings of the National Academy of Sciences*, *106*(46), 19569–19574. https://doi.org/10.1073/pnas.0905306106

Sun, S. Z., Fidalgo, C., Barense, M. D., Lee, A. C. H., Cant, J. S., & Ferber, S. (2017). Erasing and blurring memories: The differential impact of interference on separate aspects of forgetting. *Journal of Experimental Psychology: General*, *146*(11), 1606–1630. https://doi.org/10.1037/xge0000359

Sun, Y., Wong, A. K. C., & Kamel, M. S. (2009). Classification of imbalanced data: A review. *International Journal of Pattern Recognition and Artificial Intelligence*, *23*(4), 687–719. Scopus. https://doi.org/10.1142/S0218001409007326

Suzuki, M., & Gottlieb, J. (2013). Distinct neural mechanisms of distractor suppression in the frontal and parietal lobe. *Nature Neuroscience*, *16*(1), 98–104. https://doi.org/10.1038/nn.3282

Taatgen, N., & van Rijn, H. (2011). Traces of times past: Representations of temporal intervals in memory. *Memory & Cognition*, *39*(8), 1546–1560. https://doi.org/10.3758/s13421-011-0113-0

Tamber-Rosenau, B. J., Esterman, M., Chiu, Y.-C., & Yantis, S. (2011). Cortical Mechanisms of Cognitive Control for Shifting Attention in Vision and Working Memory. *Journal of Cognitive Neuroscience*, *23*(10), 2905–2919. https://doi.org/10.1162/jocn.2011.21608

Tang, Y.-Y., Rothbart, M. K., & Posner, M. I. (2012). Neural correlates of establishing, maintaining, and switching brain states. *Trends in Cognitive Sciences*, *16*(6), 330–337. https://doi.org/10.1016/j.tics.2012.05.001

Tarder-Stoll, H., Jayakumar, M., Dimsdale-Zucker, H. R., Günseli, E., & Aly, M. (2020). Dynamic internal states shape memory retrieval. *Neuropsychologia*, *138*, 107328. https://doi.org/10.1016/j.neuropsychologia.2019.107328

Taulu, S., & Simola, J. (2006). Spatiotemporal signal space separation method for rejecting nearby interference in MEG measurements. *Physics in Medicine & Biology*, *51*(7), 1759. https://doi.org/10.1088/0031-9155/51/7/008

Teichmann, L., Moerel, D., Rich, A. N., & Baker, C. I. (2022). The nature of neural object representations during dynamic occlusion. *Cortex*, *153*, 66–86. https://doi.org/10.1016/j.cortex.2022.04.009

Teng, C., Kaplan, S. M., Shomstein, S., & Kravitz, D. J. (2023). Assessing the interaction between working memory and perception through time. *Attention, Perception, & Psychophysics*. https://doi.org/10.3758/s13414-023-02785-3

Teng, C., & Kravitz, D. J. (2019). Visual working memory directly alters perception. *Nature Human Behaviour*, *3*(8), Article 8. https://doi.org/10.1038/s41562-019-0640-4

Theeuwes, J. (2010). Top–down and bottom–up control of visual selection. *Acta Psychologica*, *135*(2), 77–99. https://doi.org/10.1016/j.actpsy.2010.02.006

Theeuwes, J., Bogaerts, L., & van Moorselaar, D. (2022). What to expect where and when: How statistical learning drives visual selection. *Trends in Cognitive Sciences*, *26*(10), 860–872. https://doi.org/10.1016/j.tics.2022.06.001

Thölke, P., Mantilla-Ramos, Y.-J., Abdelhedi, H., Maschke, C., Dehgan, A., Harel, Y., Kemtur, A., Mekki Berrada, L., Sahraoui, M., Young, T., Bellemare Pépin, A., El Khantour, C., Landry, M., Pascarella, A., Hadid, V., Combrisson, E., O'Byrne, J., & Jerbi, K. (2023). Class imbalance should not throw you off balance: Choosing the right classifiers and performance metrics for brain decoding with imbalanced data. *NeuroImage*, *277*, 120253. https://doi.org/10.1016/j.neuroimage.2023.120253

Thomson, E., & Piccinini, G. (2018). Neural Representations Observed. *Minds and Machines*, *28*(1), 191–235. https://doi.org/10.1007/s11023-018-9459-4

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

Treder, M. S., Charest, I., Michelmann, S., Martín-Buro, M. C., Roux, F., Carceller-Benito, F., Ugalde-Canitrot, A., Rollings, D. T., Sawlani, V., Chelvarajah, R., Wimber, M., Hanslmayr, S., &

Staresina, B. P. (2021). The hippocampus as the switchboard between perception and memory. *Proceedings of the National Academy of Sciences*, *118*(50), e2114171118. https://doi.org/10.1073/pnas.2114171118

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Trübutschek, D., Marti, S., Ueberschär, H., & Dehaene, S. (2019). Probing the limits of activity-silent non-conscious working memory. *Proceedings of the National Academy of Sciences*, *116*(28), 14358–14367. https://doi.org/10.1073/pnas.1820730116

Tukey, J. W. (1977). *Exploratory data analysis*. Addison-Wesley.

Uddin, L. Q. (2015). Salience processing and insular cortical function and dysfunction. *Nature Reviews Neuroscience*, *16*(1), 55–61. https://doi.org/10.1038/nrn3857

Ülkü, S., Getzmann, S., Wascher, E., & Schneider, D. (2024). Be prepared for interruptions: EEG correlates of anticipation when dealing with task interruptions and the role of aging. *Scientific Reports*, *14*(1), 5679. https://doi.org/10.1038/s41598-024-56400-y

Unsworth, N., & Robison, M. K. (2016). The influence of lapses of attention on working memory capacity. *Memory & Cognition*, *44*(2), 188–196. https://doi.org/10.3758/s13421-015-0560-0

Vallesi, A., & Shallice, T. (2007). Developmental dissociations of preparation over time: Deconstructing the variable foreperiod phenomena. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(6), 1377–1388. https://doi.org/10.1037/0096-1523.33.6.1377

van der Mijn, R., & van Rijn, H. (2021). Attention Does Not Affect the Speed of Subjective Time, but Whether Temporal Information Guides Performance: A Large-Scale Study of Intrinsically Motivated Timers in a Real-Time Strategy Game. *Cognitive Science*, *45*(3), e12939. https://doi.org/10.1111/cogs.12939

van der Stigchel, S., Merten, H., Meeter, M., & Theeuwes, J. (2007). The effects of a task-irrelevant visual event on spatial working memory. *Psychonomic Bulletin & Review*, *14*(6), 1066–1071. https://doi.org/10.3758/BF03193092

van Driel, J., Olivers, C. N. L., & Fahrenfort, J. J. (2021). High-pass filtering artifacts in multivariate classification of neural time series data. *Journal of Neuroscience Methods*, *352*, 109080. https://doi.org/10.1016/j.jneumeth.2021.109080

van Ede, F. (2018). Mnemonic and attentional roles for states of attenuated alpha oscillations in perceptual working memory: A review. *The European Journal of Neuroscience*, *48*(7), 2509–2515. https://doi.org/10.1111/ejn.13759

van Ede, F. (2020). Visual working memory and action: Functional links and bi-directional influences. *Visual Cognition*, *28*(5–8), 401–413. https://doi.org/10.1080/13506285.2020.1759744

van Ede, F. (2023). Do microsaccades track shifting but not sustaining covert attention? *Proceedings of the National Academy of Sciences*, *120*(35), e2309431120. https://doi.org/10.1073/pnas.2309431120

van Ede, F., Board, A. G., & Nobre, A. C. (2020). Goal-directed and stimulus-driven selection of internal representations. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, *117*, 24590–24598. https://doi.org/10.1073/pnas.2013432117

van Ede, F., Chekroud, S. R., & Nobre, A. C. (2019). Human gaze tracks attentional focusing in memorized visual space. *Nature Human Behaviour*, *3*(5), Article 5. https://doi.org/10.1038/s41562-019-0549-y

van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2018). Decoding the influence of anticipatory states on visual perception in the presence of temporal distractors. *Nature Communications*, *9*(1), 1449. https://doi.org/10.1038/s41467-018-03960-z

van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2019). Concurrent visual and motor selection during visual working memory guided action. *Nature Neuroscience*, *22*(3), Article 3. https://doi.org/10.1038/s41593-018-0335-6

van Ede, F., Deden, J., & Nobre, A. C. (2021). Looking ahead in working memory to guide sequential behaviour. *Current Biology*, *31*(12), R779–R780. https://doi.org/10.1016/j.cub.2021.04.063

van Ede, F., Niklaus, M., & Nobre, A. C. (2017). Temporal Expectations Guide Dynamic Prioritization in Visual Working Memory through Attenuated α Oscillations. *The Journal of Neuroscience*, *37*(2), 437–445. https://doi.org/10.1523/JNEUROSCI.2272-16.2016

van Ede, F., & Nobre, A. C. (2023). Turning Attention Inside Out: How Working Memory Serves Behavior. *Annual Review of Psychology*, *74*(1), 137–165. https://doi.org/10.1146/annurev-psych-021422-041757

van Ede, F., Rohenkohl, G., Gould, I., & Nobre, A. C. (2020). Purpose-dependent consequences of temporal expectations serving perception and action. *The Journal of Neuroscience*, *40*(41), 7877–7886. https://doi.org/10.1523/JNEUROSCI.1134-20.2020

van Lambalgen, R. M., & Los, S. A. (2008). The Role of Attention in Nonspecific Preparation. In L. K. McRae & V. M. Sloutsky (Eds.), *Proceedings of the 30th Annual Conference of the Cognitive Science Society* (pp. 1525–1530). Cognitive Science Society.

van Moorselaar, D., Battistoni, E., Theeuwes, J., & Olivers, C. N. L. (2015). Rapid influences of cued visual memories on attentional guidance. *Annals of the New York Academy of Sciences*, *1339*, 1–10. https://doi.org/10.1111/nyas.12574

van Moorselaar, D., Gunseli, E., Theeuwes, J., & N. L. Olivers, C. (2015). The time course of protecting a visual memory representation from perceptual interference. *Frontiers in Human Neuroscience*, *8*. https://doi.org/10.3389/fnhum.2014.01053

van Moorselaar, D., Olivers, C. N. L., Theeuwes, J., Lamme, V. A. F., & Sligte, I. G. (2015). Forgotten but not gone: Retro-cue costs and benefits in a double-cueing paradigm suggest multiple states in visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *41*, 1755–1763. https://doi.org/10.1037/xlm0000124

van Moorselaar, D., & Slagter, H. A. (2020). Inhibition in selective attention. *Annals of the New York Academy of Sciences*, *1464*(1), 204–221. https://doi.org/10.1111/nyas.14304

van Moorselaar, D., Theeuwes, J., & Olivers, C. N. L. (2014). In competition for the attentional template: Can multiple items within visual working memory guide attention? *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1450–1464. https://doi.org/10.1037/a0036229

van Strien, N. M., Cappaert, N. L. M., & Witter, M. P. (2009). The anatomy of memory: An interactive overview of the parahippocampal–hippocampal network. *Nature Reviews Neuroscience*, *10*(4), 272–282. https://doi.org/10.1038/nrn2614

Vangkilde, S., Coull, J. T., & Bundesen, C. (2012). Great expectations: Temporal expectation modulates perceptual processing speed. *Journal of Experimental Psychology: Human Perception and Performance*, *38*(5), 1183–1191. https://doi.org/10.1037/a0026343

Vatterott, D. B., & Vecera, S. P. (2012). Experience-dependent attentional tuning of distractor rejection. *Psychonomic Bulletin & Review*, *19*(5), 871–878. https://doi.org/10.3758/s13423-012-0280-4

Vergauwe, E., Hautekiet, C., & Langerock, N. (2023). *Distractor susceptibility in visual working memory: No evidence for particularly vulnerable mnemonic representations in the focus of attention*. https://doi.org/10.31234/osf.io/4yqjt

Verschooren, S., & Egner, T. (2023). When the mind's eye prevails: The Internal Dominance over External Attention (IDEA) hypothesis. *Psychonomic Bulletin & Review*. https://doi.org/10.3758/s13423-023-02272-8

Verschooren, S., Liefooghe, B., Brass, M., & Pourtois, G. (2019). Attentional flexibility is imbalanced: Asymmetric cost for switches between external and internal attention.

Journal of Experimental Psychology: Human Perception and Performance, *45*(10), 1399–1414. https://doi.org/10.1037/xhp0000683

Verschooren, S., Pourtois, G., & Egner, T. (2020). More efficient shielding for internal than external attention? Evidence from asymmetrical switch costs. *Journal of Experimental Psychology: Human Perception and Performance*, *46*(9), 912–925. https://doi.org/10.1037/xhp0000758

Verschooren, S., Schindler, S., De Raedt, R., & Pourtois, G. (2019). Switching attention from internal to external information processing: A review of the literature and empirical support of the resource sharing account. *Psychonomic Bulletin & Review*, *26*(2), 468–490. https://doi.org/10.3758/s13423-019-01568-y

Verschooren, S., Schindler, S., De Raedt, R., & Pourtois, G. (2021). Early reduction of sensory processing within the visual cortex when switching from internal to external attention. *Biological Psychology*, *163*, 108119. https://doi.org/10.1016/j.biopsycho.2021.108119

Villena-González, M., Moënne-Loccoz, C., Lagos, R. A., Alliende, L. M., Billeke, P., Aboitiz, F., López, V., & Cosmelli, D. (2017). Attending to the heart is associated with posterior alpha band increase and a reduction in sensitivity to concurrent visual stimuli. *Psychophysiology*, *54*(10), 1483–1497. https://doi.org/10.1111/psyp.12894

Vincent, J. L., Kahn, I., Snyder, A. Z., Raichle, M. E., & Buckner, R. L. (2008). Evidence for a Frontoparietal Control System Revealed by Intrinsic Functional Connectivity. *Journal of Neurophysiology*, *100*(6), 3328–3342. https://doi.org/10.1152/jn.90355.2008

Vogel, E. K., McCollough, A. W., & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature*, *438*(7067), 500–503. https://doi.org/10.1038/nature04171

Vortmann, L.-M., Kroll, F., & Putze, F. (2019). EEG-Based Classification of Internally- and Externally-Directed Attention in an Augmented Reality Paradigm. *Frontiers in Human Neuroscience*, *13*, 348. https://doi.org/10.3389/fnhum.2019.00348

Wallis, G., Stokes, M., Cousijn, H., Woolrich, M., & Nobre, A. C. (2015). Frontoparietal and Cingulo-opercular Networks Play Dissociable Roles in Control of Working Memory. *Journal of Cognitive Neuroscience*, *27*(10), 2019–2034. https://doi.org/10.1162/jocn_a_00838

Wang, B., Knapen, T., & Olivers, C. N. L. (2022). Visual Working Memory Adapts to the Nature of Anticipated Interference. *Journal of Cognitive Neuroscience*, *34*(7), 1148–1163. https://doi.org/10.1162/jocn_a_01853

Wang, B., & Theeuwes, J. (2018a). How to inhibit a distractor location? Statistical learning versus active, top-down suppression. *Attention, Perception, & Psychophysics*, *80*(4), 860–870. https://doi.org/10.3758/s13414-018-1493-z

Wang, B., & Theeuwes, J. (2018b). Statistical regularities modulate attentional capture. *Journal of Experimental Psychology: Human Perception and Performance*, *44*(1), 13–17. https://doi.org/10.1037/xhp0000472

Wang, B., Theeuwes, J., & Olivers, C. N. L. (2019). Momentary, Offset-Triggered Dual-Task Interference in Visual Working Memory. *Journal of Cognition*, *2*(1), 38. https://doi.org/10.5334/joc.84

Wang, B., van Driel, J., Ort, E., & Theeuwes, J. (2019). Anticipatory Distractor Suppression Elicited by Statistical Regularities in Visual Search. *Journal of Cognitive Neuroscience*, *31*(10), 1535–1548. https://doi.org/10.1162/jocn_a_01433

Wang, S., & van Ede, F. (2023). *Does attention to internal visual content in working memory selectively affect distractor intrusions at memory-congruent locations?* [Poster]. European Conference on Visual Perception, Cyprus. https://journals.sagepub.com/pb-assets/ecvp23_abstracts_UPD-1702301447867.pdf

Weber, R. J., Burt, D. B., & Noll, N. C. (1986). Attention switching between perception and memory. *Memory & Cognition*, *14*(3), 238–245. https://doi.org/10.3758/BF03197699

Weilnhammer, V., Stuke, H., Standvoss, K., & Sterzer, P. (2023). Sensory processing in humans and mice fluctuates between external and internal modes. *PLOS Biology*, *21*(12), e3002410. https://doi.org/10.1371/journal.pbio.3002410

Weinbach, N., & Henik, A. (2012). Temporal Orienting and Alerting – The Same or Different? *Frontiers in Psychology*, *3*. https://doi.org/10.3389/fpsyg.2012.00236

Weissman, D. H., Roberts, K. C., Visscher, K. M., & Woldorff, M. G. (2006). The neural bases of momentary lapses in attention. *Nature Neuroscience*, *9*(7), 971–978. https://doi.org/10.1038/nn1727

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer New York. https://doi.org/10.1007/978-0-387-98141-3

Willett, S. M., & Mayo, J. P. (2023). Microsaccades are directed toward the midpoint between targets in a variably cued attention task. *Proceedings of the National Academy of Sciences*, *120*(20), e2220552120. https://doi.org/10.1073/pnas.2220552120

Williams, G. C., Boettcher, S. E. P., & Nobre, A. C. (2023). Using regularity-based temporal predictions to shift our attentional template across time during multiple-target dynamic visual search. *Journal of Vision*, *23*(9), 5304. https://doi.org/10.1167/jov.23.9.5304

Wilsch, A., Henry, M. J., Herrmann, B., Herrmann, C. S., & Obleser, J. (2018). Temporal expectation modulates the cortical dynamics of short-term memory. *Journal of Neuroscience*, *38*(34), 7428–7439. https://doi.org/10.1523/JNEUROSCI.2928-17.2018

Wilsch, A., Henry, M. J., Herrmann, B., Maess, B., & Obleser, J. (2015). Alpha Oscillatory Dynamics Index Temporal Expectation Benefits in Working Memory. *Cerebral Cortex*, *25*(7), 1938–1946. https://doi.org/10.1093/cercor/bhu004

Wolfe, J. M. (1994). Guided Search 2.0 A revised model of visual search. *Psychonomic Bulletin & Review*, *1*(2), 202–238. https://doi.org/10.3758/BF03200774

Wolfe, J. M. (2021). Guided Search 6.0: An updated model of visual search. *Psychonomic Bulletin & Review*, *28*(4), 1060–1092. https://doi.org/10.3758/s13423-020-01859-9

Wolff, M. J., Ding, J., Myers, N., & Stokes, M. (2015). Revealing hidden states in visual working memory using electroencephalography. *Frontiers in Systems Neuroscience*, *9*. https://www.frontiersin.org/articles/10.3389/fnsys.2015.00123

Wolff, M. J., Jochim, J., Akyürek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*(6), 864–871. https://doi.org/10.1038/nn.4546

Woodrow, H. (1914). The measurement of attention. *The Psychological Monographs*, *17*(5), i–158. https://doi.org/10.1037/h0093087

Woolrich, M., Hunt, L., Groves, A., & Barnes, G. (2011). MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization. *NeuroImage*, *57*(4), 1466–1479. https://doi.org/10.1016/j.neuroimage.2011.04.041

Worden, M. S., Foxe, J. J., Wang, N., & Simpson, G. V. (2000). Anticipatory biasing of visuospatial attention indexed by retinotopically specific alpha-band electroencephalography increases over occipital cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *20*(6), RC63. https://doi.org/10.1523/JNEUROSCI.20-06-j0002.2000

Wöstmann, M., Störmer, V. S., Obleser, J., Addleman, D. A., Andersen, S. K., Gaspelin, N., Geng, J. J., Luck, S. J., Noonan, M. P., Slagter, H. A., & Theeuwes, J. (2022). Ten simple rules to study distractor suppression. *Progress in Neurobiology*, *213*, 102269. https://doi.org/10.1016/j.pneurobio.2022.102269

Wulf, G., Höß, M., & Prinz, W. (1998). Instructions for Motor Learning: Differential Effects of Internal Versus External Focus of Attention. *Journal of Motor Behavior*, *30*(2), 169–179. https://doi.org/10.1080/00222899809601334

Wylie, G., & Allport, D. A. (2000). Task switching and the measurement of “switch costs”. *Psychological Research*, *63*(3), 212–233. https://doi.org/10.1007/s004269900003

Xu, Y. (2017). Reevaluating the Sensory Account of Visual Working Memory Storage. *Trends in Cognitive Sciences*, *21*(10), 794–815. https://doi.org/10.1016/j.tics.2017.06.013

Xu, Y. (2020). Revisit once more the sensory storage account of visual working memory. *Visual Cognition*, *28*(5–8), 433–446. https://doi.org/10.1080/13506285.2020.1818659

Yeung, N., & Monsell, S. (2003). Switching between tasks of unequal familiarity: The role of stimulus-attribute and response-set selection. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(2), 455–469. https://doi.org/10.1037/0096-1523.29.2.455

Yeung, N., Nystrom, L. E., Aronson, J. A., & Cohen, J. D. (2006). Between-Task Competition and Cognitive Control in Task Switching. *The Journal of Neuroscience*, *26*(5), 1429–1438. https://doi.org/10.1523/JNEUROSCI.3109-05.2006

Yoon, J. H., Curtis, C. E., & D’Esposito, M. (2006). Differential effects of distraction during working memory on delay-period activity in the prefrontal cortex and the visual association cortex. *NeuroImage*, *29*(4), 1117–1126. https://doi.org/10.1016/j.neuroimage.2005.08.024

Yu, X., Zhou, Z., Becker, S. I., Boettcher, S. E. P., & Geng, J. J. (2023). Good-enough attentional guidance. *Trends in Cognitive Sciences*, *27*(4), 391–403. https://doi.org/10.1016/j.tics.2023.01.007

Zabelina, D. L., & Andrews-Hanna, J. R. (2016). Dynamic network interactions supporting internally-oriented cognition. *Current Opinion in Neurobiology*, *40*, 86–93. https://doi.org/10.1016/j.conb.2016.06.014

Zanto, T. P., Rubens, M. T., Thangavel, A., & Gazzaley, A. (2011). Causal role of the prefrontal cortex in top-down modulation of visual processing and working memory. *Nature Neuroscience*, *14*(5), 656–661. https://doi.org/10.1038/nn.2773

Zhang, Z., & Lewis-Peacock, J. A. (2022). Prioritization sharpens working memories but does not protect them from distraction. *Journal of Experimental Psychology: General*, No Pagination Specified-No Pagination Specified. https://doi.org/10.1037/xge0001309

Zhang, Z., & Lewis-Peacock, J. A. (2023). Bend but don’t break: Prioritization protects working memory from displacement but leaves it vulnerable to distortion from distraction. *Cognition*, *239*, 105574. https://doi.org/10.1016/j.cognition.2023.105574

Zhou, Y., Curtis, C. E., Sreenivasan, K. K., & Fougnie, D. (2022). Common Neural Mechanisms Control Attention and Working Memory. *Journal of Neuroscience*, *42*(37), 7110–7120. https://doi.org/10.1523/JNEUROSCI.0443-22.2022

Zhou, Y. J., Ramchandran, A., & Haegens, S. (2023). Alpha oscillations protect working memory against distracters in a modality-specific way. *NeuroImage*, *278*, 120290. https://doi.org/10.1016/j.neuroimage.2023.120290

Zickerick, B., Kobald, S. O., Thönes, S., Küper, K., Wascher, E., & Schneider, D. (2021). Don't stop me now: Hampered retrieval of action plans following interruptions. *Psychophysiology*, *58*(2), e13725. https://doi.org/10.1111/psyp.13725

Zickerick, B., Rösner, M., Sabo, M., & Schneider, D. (2021). How to refocus attention on working memory representations following interruptions-Evidence from frontal theta and posterior alpha oscillations. *The European Journal of Neuroscience*, *54*(11), 7820–7838. https://doi.org/10.1111/ejn.15506

Zickerick, B., Thönes, S., Kobald, S. O., Wascher, E., Schneider, D., & Küper, K. (2020). Differential Effects of Interruptions and Distractions on Working Memory Processes in an ERP Study. *Frontiers in Human Neuroscience*, *14*, 84. https://doi.org/10.3389/fnhum.2020.00084

Ziegler, D. A., Janowich, J. R., & Gazzaley, A. (2018). Differential Impact of Interference on Internally- and Externally-Directed Attention. *Scientific Reports*, *8*(1), 2498. https://doi.org/10.1038/s41598-018-20498-8

Zokaei, N., Board, A. G., Manohar, S. G., & Nobre, A. C. (2019). Modulation of the pupillary response by the content of visual working memory. *Proceedings of the National Academy of Sciences*, *116*(45), 22802–22810. https://doi.org/10.1073/pnas.1909959116

Zokaei, N., Gillebert, C. R., Chauvin, J. J., Gresch, D., Board, A. G., Rolinski, M., Hu, M. T., & Nobre, A. C. (2021). Temporal orienting in Parkinson's disease. *European Journal of Neuroscience*, *53*(8), 2713–2725. https://doi.org/10.1111/ejn.15114

Zokaei, N., Manohar, S., Husain, M., & Feredoes, E. (2014). Causal Evidence for a Privileged Working Memory State in Early Visual Cortex. *The Journal of Neuroscience*, *34*(1), 158–162. https://doi.org/10.1523/JNEUROSCI.2899-13.2014

8 Appendix

8.1 Supplementary Information for Chapter 2

Supplementary Results 8.1

Working-Memory Performance as a Function of Sensory Similarity, Block Type, and Delay Condition

As an additional analysis, I examined whether the observed temporal-expectation effects in the working-memory task were modulated by the sensory similarity between the intervening and memory item. I calculated the absolute angular differences between the intervening item and the memory item, sorted them ordinally, and binned them into three bins. This resulted in three levels of sensory similarity, that is, low (60° to 90° angular difference), medium (30° to 60° angular difference), and high (0° to 30° angular difference). As a next step, I performed a repeated-measures ANOVA with the factors delay condition, block type, and sensory similarity.

RTs did not differ between more versus less similar items, and none of the interactions including the factor sensory similarity reached significance (**Supplementary Figure 8.3A**, see **Supplementary Table 8.2** for the full set of inferential statistics). For reproduction errors, I found a main effect of sensory similarity ($F_{(2,106)} = 5.411$, $p = 0.006$, $\eta^2_G = 0.002$; **Supplementary Figure 8.3B**). Pairwise comparisons revealed that errors were higher when the similarity between items was low ($t_{(53)} = -2.911$, $p_{Bonferroni} = 0.016$, $d = 0.396$) or medium ($t_{(53)} = -3.150$, $p_{Bonferroni} = 0.008$, $d = 0.429$) as compared to high. However, there was no difference in errors between trials with low and medium sensory similarity ($t_{(53)} = -0.515$, $p_{Bonferroni} = 1.000$, $d = 0.070$). As for RTs, none of the interactions including the factor sensory similarity reached significance (see **Supplementary Table 8.3** for the full set of inferential statistics).

Supplementary Table 8.1. Inferential statistics for Experiment 1: control analysis. Pairwise comparisons of RTs to the working-memory task between the different levels of delay condition and block type. The analysis was performed following pre-processing steps described in '2.3.1.3 Analysis', including the removal of RTs to the intervening task which were faster than 200 ms and slower than 1500 ms, as well as after removing RTs to the intervening task that were faster than 200 ms and slower than 1000 ms. This resulted on average in 4.83 ± 2.78% removed trials per participant for the former and 8.42 ± 8.80% for the latter. Pairwise comparisons revealed the same pattern of results after removing RTs to the intervening task that were faster than 200 ms and slower than 1500 ms and after removing RTs to the intervening task that were faster than 200 ms and slower than 1000 ms.

	Trials removed	Fixed-early vs. Variable-early	Fixed-late vs. Variable-late	Fixed-early vs. Fixed-late	Variable-early- vs. Variable late
200 ms < intervening-task RT < 1500 ms	4.83 ± 2.78%	$t_{(53)}$ =-7.437 $p_{Bonferroni}$ < 0.001 d = 1.012	$t_{(53)}$ = 1.568 $p_{Bonferroni}$ = 0.491 d = 0.213	$t_{(53)}$ = 2.531 $p_{Bonferroni}$ = 0.058 d = 0.344	$t_{(53)}$ = 10.633 $p_{Bonferroni}$ < 0.001 d = 1.447
200 ms < intervening-task RT < 1000 ms	8.42 ± 8.80%	$t_{(53)}$ =-6.977 $p_{Bonferroni}$ < 0.001 d = 0.950	$t_{(53)}$ = 1.558 $p_{Bonferroni}$ = 0.501 d = 0.212	$t_{(53)}$ = 2.183 $p_{Bonferroni}$ = 0.134 d = 0.297	$t_{(53)}$ = 10.611 $p_{Bonferroni}$ < 0.001 d = 1.444

Supplementary Table 8.2. Inferential statistics for Experiment 1: reaction times in the working-memory task. Main and interaction effects of RTs for the factors block type (fixed vs. variable), delay condition (early vs. late), and sensory similarity (low vs. medium vs. high), tested with a 2×2×3 repeated-measures ANOVA. This analysis yielded a significant main effect of block type and delay condition, and a significant interaction between these two factors. * indicates $p < 0.05$, ** indicates $p < 0.001$.

	df	*F*	*p*	η^2_G
Block Type **	1, 53	24.909	< 0.001	0.005
Delay Condition **	1, 53	61.691	< 0.001	0.023
Sensory Similarity	2, 106	0.237	0.789	< 0.001
Block Type × Delay Condition **	1, 53	51.199	< 0.001	0.009
Block Type × Sensory Similarity	2, 106	2.492	0.088	< 0.001
Delay Condition × Sensory Similarity	2, 106	0.555	0.576	< 0.001
Block Type × Delay Condition × Sensory Similarity	2, 106	0.544	0.582	< 0.001

Supplementary Table 8.3. Inferential statistics for Experiment 1: reproduction errors in the working-memory task. Main and interaction effects of reproduction errors for the factors block type (fixed vs. variable), delay condition (early vs. late), and sensory similarity (low vs. medium vs. high), tested with a 2×2×3 repeated-measures ANOVA. This analysis yielded a significant main effect of block type and sensory similarity. * indicates $p < 0.05$, ** indicates $p < 0.001$.

	df	*F*	*p*	η^2_G
Block Type *	1, 53	5.340	0.025	0.002
Delay Condition	1, 53	0.291	0.592	< 0.001
Sensory Similarity *	2, 106	5.411	0.006	0.002
Block Type × Delay Condition	1, 53	0.596	0.443	< 0.001
Block Type × Sensory Similarity	2, 106	0.372	0.690	< 0.001
Delay Condition × Sensory Similarity	2, 106	1.360	0.261	< 0.001
Block Type × Delay Condition × Sensory Similarity	2, 106	2.562	0.082	0.001

Supplementary Table 8.4. Inferential statistics for Experiment 2: reproduction errors in the working-memory task. **(A)** Main and interaction effects of reproduction errors for the factors temporal predictability (fixed vs. variable), interference type (distraction vs. interruption), and interference onset (250 ms vs. 1000 ms vs. 1750 ms), tested with a 2×2×3 repeated-measures ANOVA. This analysis yielded a significant main effect of temporal predictability, confirming the effect that temporal expectations improve visual working-memory performance. I also found a significant effect of interference type and interference onset. Further, the interaction between interference onset and interference type was significant. **(B)** Bonferroni-corrected pairwise comparisons of reproduction errors between the three different interference onsets were conducted to break down the main effect of the factor interference onset. However, errors did not significantly differ between the three interference onsets. **(C)** To test whether the effect of interference type differed between interference onsets, I performed Bonferroni-corrected pairwise comparisons. This analysis revealed that interrupters as compared to distractors significantly decreased visual working-memory performance at all three interference onsets. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Experiment 2: Repeated-measures ANOVA of reproduction errors.

	df	*F*	*p*	η^2_G
Temporal Predictability*	1, 53	8.774	0.005	0.004
Interference Type**	1, 53	75.289	< 0.001	0.096
Interference Onset*	2, 106	3.395	0.037	0.004
Temporal Predictability × Interference Type	1, 53	0.502	0.482	< 0.001
Temporal Predictability × Interference Onset	2, 106	0.104	0.902	< 0.001
Interference Type × Interference Onset*	2, 106	3.943	0.022	0.004
Temporal Predictability × Interference Type × Interference Onset	2, 106	0.415	0.662	< 0.001

(B) Experiment 2: Bonferroni-corrected pairwise comparisons of reproduction errors between the three different interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
250 ms vs. 1000 ms	53	1.936	0.175	0.263
250 ms vs. 1750 ms	53	-0.722	1.000	0.098
1000 ms vs. 1750 ms	53	-2.192	0.098	0.298

(C) Experiment 2: Bonferroni-corrected pairwise comparisons of reproduction errors between distractors and interrupters at each of three interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Distractor vs. Interrupter: 250 ms**	53	-6.101	< 0.001	0.830
Distractor vs. Interrupter: 1000 ms**	53	-5.572	< 0.001	0.758
Distractor vs. Interrupter: 1750 ms**	53	-7.839	< 0.001	1.067

Supplementary Table 8.5. Inferential statistics for Experiment 2: reaction times in the working-memory task. (A) Main and interaction effects of RTs for the factors temporal predictability (fixed vs. variable), interference type (distraction vs. interruption), and interference onset (250 ms vs. 1000 ms vs. 1750 ms), tested with a 2×2×3 repeated-measures ANOVA. **(B)** Bonferroni-corrected pairwise comparisons of RTs between the three different interference onsets. **(C)** Bonferroni-corrected pairwise comparisons of RTs between fixed and variable blocks at each of three interference onsets. **(D)** Bonferroni-corrected pairwise comparisons of RTs between distractor and interrupter blocks at each of three interference onsets. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Experiment 2: Repeated-measures ANOVA of RTs.

	df	*F*	*p*	η^2_G
Temporal Predictability	1, 53	1.278	0.264	< 0.001
Interference Type**	1, 53	21.080	< 0.001	0.018
Interference Onset**	2, 106	45.657	< 0.001	0.065
Temporal Predictability × Interference Type	1, 53	0.843	0.363	< 0.001
Temporal Predictability × Interference Onset*	2, 106	3.482	0.003	0.005
Interference Type × Interference Onset**	2, 106	16.650	< 0.001	0.008
Temporal Predictability × Interference Type × Interference Onset	2, 106	0.754	0.473	< 0.001

(B) Experiment 2: Bonferroni-corrected pairwise comparisons of RTs between the three different interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
250 ms vs. 1000 ms	53	-0.025	1.000	0.003
250 ms vs. 1750 ms**	53	-8.211	< 0.001	1.117
1000 ms vs. 1750 ms**	53	-8.703	< 0.001	1.099

(C) Experiment 2: Bonferroni-corrected pairwise comparisons of RTs between fixed and variable blocks at each of three interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Fixed vs. Random: 250 ms	53	1.548	0.383	0.211
Fixed vs. Random: 1000 ms	53	-2.410	0.058	0.328
Fixed vs. Random: 1750 ms	53	-1.186	0.722	0.161

(D) Experiment 2: Bonferroni-corrected pairwise comparisons of RTs between distractors and interrupters at each of three interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Distractor vs. Interrupter: 250 ms	53	-0.976	1.000	0.133
Distractor vs. Interrupter: 1000 ms**	53	-3.611	< 0.001	0.491
Distractor vs. Interrupter: 1750 ms**	53	-6.436	< 0.001	0.876

Supplementary Table 8.6. Inferential statistics for Experiment 2: error rates in the interrupter task. (A) Main and interaction effects of error rates for the factors temporal predictability (fixed vs. variable) and interference onset (250 ms vs. 1000 ms vs. 1750 ms), tested with a 2×3 repeated-measures ANOVA. **(B)** Bonferroni-corrected pairwise comparisons of error rates between the three different interference onsets. Note: in the main analysis, trials with wrong interrupter responses were removed from the data. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Experiment 2: Repeated-measures ANOVA of error rates.

	df	*F*	*p*	η^2_G
Temporal Predictability	1, 53	3.636	0.062	0.006
Interference Onset**	2, 106	9.351	< 0.001	0.049
Temporal Predictability × Interference Onset	2, 106	0.166	0.847	< 0.001

(B) Experiment 2: Bonferroni-corrected pairwise comparisons of error rates between the three different interference onsets.

	df	*t*	$p_{Bonferroni}$	*d*
250 ms vs. 1000 ms*	53	3.455	0.003	0.470
250 ms vs. 1750 ms*	53	3.395	0.004	0.462
1000 ms vs. 1750 ms	53	-0.086	1.000	0.012

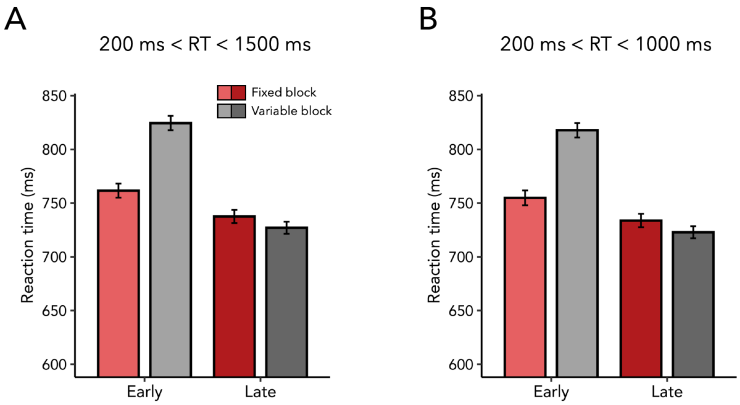

Supplementary Figure 8.1. Reaction times to the working-memory task after removing slow responses to the intervening task. Removing RTs to the intervening task that were faster than 200 ms and slower than 1500 ms (**A**; same as **Figure 2.1D**) or removing RTs that were faster than 200 ms and slower than 1000 ms **(B)** resulted in the same pattern of results. RTs to early memory probes were faster when the probe occurred in a fixed as opposed to a variable block. However, there was no difference in RTs to fixed-late versus variable-late probes. RTs to the memory probe were slower in variable-early than variable-late bocks, however, there was no difference in RTs to fixed-early versus fixed-late probes (see **Supplementary Table 8.1** for inferential statistics). Error bars indicate ±1 *SEM*.

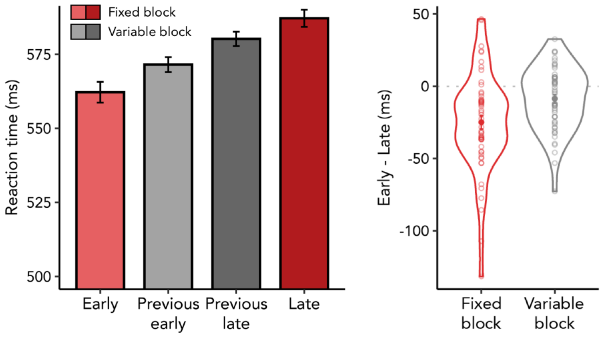

Supplementary Figure 8.2. Between-task task consequences in fixed and variable blocks. RTs to the intervening task in fixed and variable blocks. Variable-block trials are split into trials in which the previous delay condition was early versus late. Violin plots depicting the difference in RTs between early and late trials in fixed blocks (fixed-early minus fixed-late) and between previous-early and previous-late trials in variable blocks (previous-early minus previous-late). This difference was significantly stronger in fixed versus variable blocks. Dots represent individual participants. Error bars indicate ±1 *SEM*.

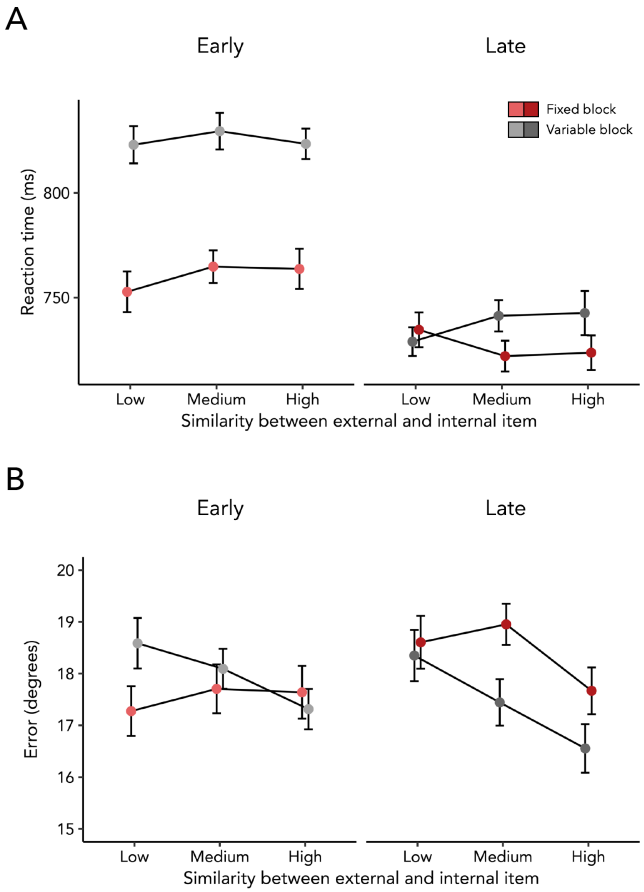

Supplementary Figure 8.3. Effect of sensory similarity on working-memory performance for early and late trials in fixed and variable blocks. (A) RTs did not differ between the different levels of sensory similarity (see **Supplementary Table 8.2** for inferential statistics). **(B)** Reproduction errors were smaller when sensory similarity was high versus low and when sensory similarity was high versus medium (see **Supplementary Table 8.3** for inferential statistics). Error bars indicate ±1 *SEM*.

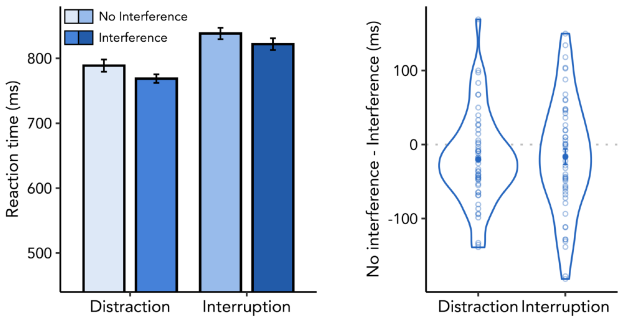

Supplementary Figure 8.4. Reaction times as a function of interference type and interference presentation. Left panel: Participants were significant slower in trials without as opposed to with interference ($F_{(1,53)}$ = 6.520, p = 0.001, η^2_G = 0.003). Moreover, participants were slower in interruption as compared to distraction blocks ($F_{(1,53)}$ = 20.339, p < 0.001, η^2_G = 0.022). There was no significant interaction between the factors ($F_{(1,53)}$ = 0.086, p = 0.770, η^2_G < 0.001). Right panel: Difference in RTs for each interference type between trials with and without interference (No interference minus interference). Dots represent individual participants. Error bars indicate ±1 *SEM*.

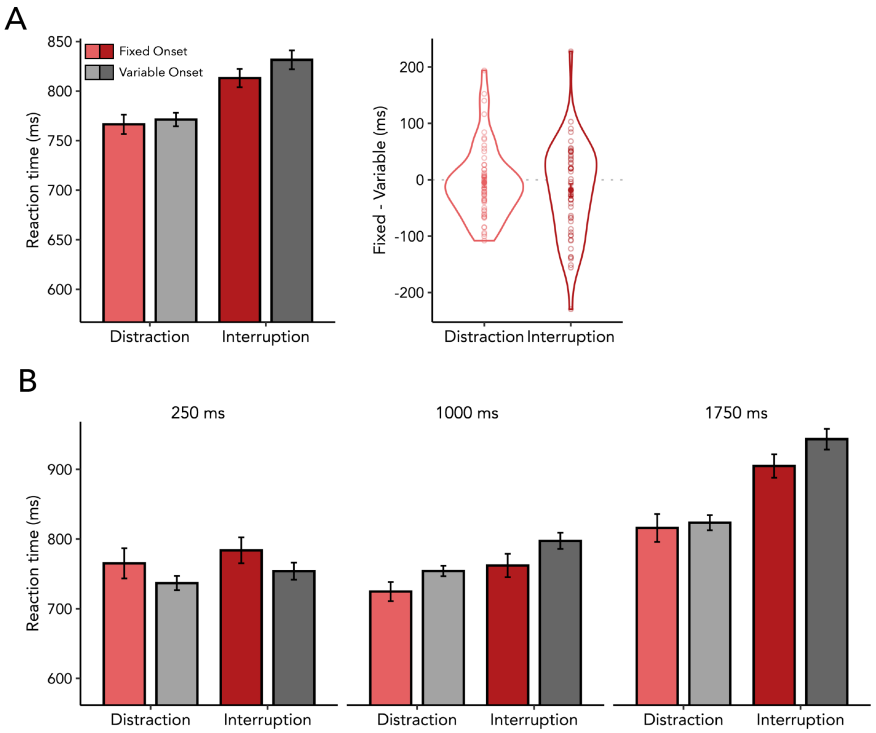

Supplementary Figure 8.5. Reaction times as a function of interference type, temporal predictability, and interference onset. **(A)** Left panel: RTs in the visual working-memory task were smaller for distractor as compared to interrupter blocks ($F_{(2,106)} = 20.546$, $p < 0.001$, $\eta^2_G = 0.022$). There was neither a main effect of temporal predictability ($F_{(1,53)} = 2.444$, $p = 0.124$, $\eta^2_G = 0.001$) nor an interaction effect between factors ($F_{(2,106)} = 1.043$, $p = 0.312$, $\eta^2_G < 0.001$). Right panel: Difference in RTs for each interference type between variable-onset blocks and fixed-onset blocks (fixed minus variable). **(B)** The temporal-expectation effect was not significant for any of the interference onsets (see **Supplementary Table 8.5** for inferential statistics). Dots represent individual participants. Error bars indicate ±1 *SEM*.

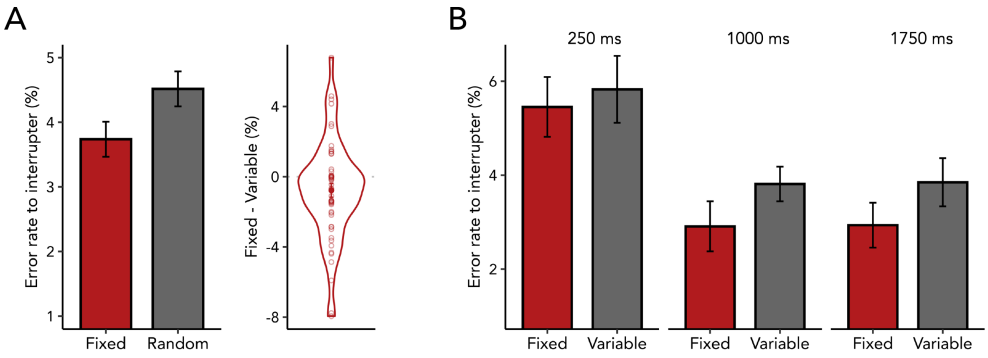

Supplementary Figure 8.6. Error rates in the interrupter task as a function of temporal predictability and interference onset. (A) Left panel: Error rates in the interrupter task were smaller for fixed as compared to variable onsets ($t_{(53)}$ =-2.026, $p_{Bonferroni}$ = 0.048, d = 0.276). Right panel shows the difference in error rates in the interrupter task between trials where the interrupter occurred at a fixed versus a variable point in time during memory retention (fixed minus variable). **(B)** The temporal-expectation effect is visually observable for all interference onsets (see **Supplementary Table 8.6** for inferential statistics). Note: for all other analyses in this experiment, trials with erroneous responses to the interrupter were removed. Dots represent individual participants. Error bars indicate ±1 *SEM*.

8.2 Supplementary Information for Chapter 3

Supplementary Table 8.7. Inferential statistics for working-memory performance: reaction times. **(A)** Main and interaction effects of RTs in the working-memory task for the factors temporal predictability (fixed vs. variable) and external-task onset (early vs. medium vs. late), tested with a 2×3 repeated-measures ANOVA. **(B)** Bonferroni-corrected pairwise comparisons of RTs in the working-memory task between the three different external-task onsets. RTs in the working-memory task were significantly faster in the early- versus medium-, early- versus late, medium- versus late-onset condition. **(C)** Bonferroni-corrected pairwise comparisons of RTs in the working-memory task to test whether the effect of temporal predictability differed between the external-task onsets. RTs were slower for variable versus fixed onsets in late trials, however faster for variable versus fixed onsets in early trials. There was no difference for medium trials. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Repeated-measures ANOVA of RTs in the working-memory task.

	df	*F*	*p*	η^2_G
Temporal Predictability**	1, 29	8.903	< 0.001	0.381
External-Task Onset*	2, 58	68.780	0.006	0.006
Temporal Predictability × External-Task Onset**	2, 58	21.382	< 0.001	0.036

(B) Bonferroni-corrected pairwise comparisons of RTs in the working-memory task between the three different external-task onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Early vs. Medium**	29	-5.200	< 0.001	0.949
Early vs. Late**	29	-8.410	< 0.001	1.536
Medium vs. Late**	29	-8.475	< 0.001	1.547

(C) Bonferroni-corrected pairwise comparisons of RTs in the working-memory task between fixed and variable trials for each of three external-task onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Fixed vs. Variable: Early*	29	2.604	0.043	0.475
Fixed vs. Variable: Medium	29	0.615	1.000	0.112
Fixed vs. Variable: Late **	29	-4.884	< 0.001	0.892

Supplementary Table 8.8. Inferential statistics for external-task performance: reaction times. **(A)** Main and interaction effects of RTs in the external task for the factors temporal predictability (fixed vs. variable) and external-task onset (early vs. medium vs. late), tested with a 2×3 repeated-measures ANOVA. **(B)** Bonferroni-corrected pairwise comparisons of RTs in the external task between the three different external-task onsets. RTs in the external task were significantly slower in the early- versus medium-, early- versus late-, and medium- versus late-onset condition. **(C)** Bonferroni-corrected pairwise comparisons of RTs in the external task to test whether the effect of temporal predictability differed between the external-task onsets. RTs were slower for variable versus fixed onsets in early trials, however faster for variable versus fixed onsets in late trials. There was no difference for medium trials. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Repeated-measures ANOVA of RTs in external task.

	df	*F*	*p*	η^2_G
Temporal Predictability*	1, 29	13.350	0.001	0.007
External-Task Onset**	2, 58	93.250	< 0.001	0.185
Temporal Predictability × External-Task Onset**	2, 58	55.931	< 0.001	0.063

(B) Bonferroni-corrected pairwise comparisons of RTs in external task between the three different external-task onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Early vs. Medium**	29	9.799	< 0.001	1.785
Early vs. Late**	29	11.092	< 0.001	2.025
Medium vs. Late**	29	4.242	< 0.001	0.774

(C) Bonferroni-corrected pairwise comparisons of RTs in external task between fixed and variable trials for each of three external-task onsets.

	df	*t*	$p_{Bonferroni}$	*d*
Fixed vs. Variable: Early**	29	-10.151	< 0.001	1.853
Fixed vs. Variable: Medium	29	0.187	1.000	0.034
Fixed vs. Variable: Late **	29	5.073	< 0.001	0.926

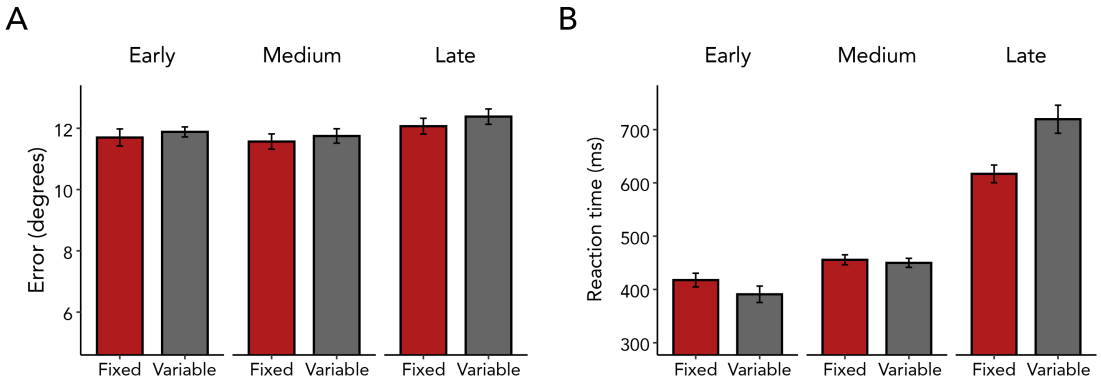

Supplementary Figure 8.7. Reproduction errors and reaction times in the working-memory task. (A) Reproduction errors in the working-memory task as a function of temporal predictability and external-task onset. There was no significant difference between conditions. **(B)** Same as (A) but for RTs (see **Supplementary Table 8.7** for inferential statistics). Error bars indicate ±1 *SEM*.

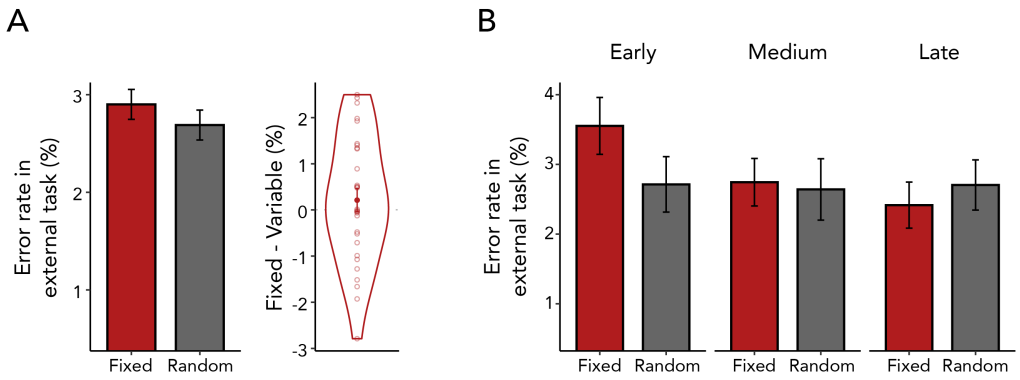

Supplementary Figure 8.8. Error rates in the external task. (A) Left panel: There was no significant difference in error rates between fixed- and random-onset blocks ($t_{(29)}$ = 0.846, $p_{Bonferroni}$ = 0.405, d = 0.154). Right panel: Difference in error rates between variable-onset blocks and fixed-onset blocks (fixed minus variable). **(B)** Error rates a function of temporal predictability and external-task onset. Neither the main effects (temporal predictability: $F_{(1,29)}$ = 10.745, p = 0.395, η^2_G = 0.002; external-task onset: $F_{(2,58)}$ = 0.876, p = 0.422, η^2_G = 0.008) nor the interaction effect ($F_{(2,58)}$ = 1.115, p = 0.335, η^2_G = 0.007) was significant. Note: for all other analyses in this experiment, trials with erroneous responses to the interrupter were removed. Error bars indicate ±1 *SEM*. Dots represent individual participants.

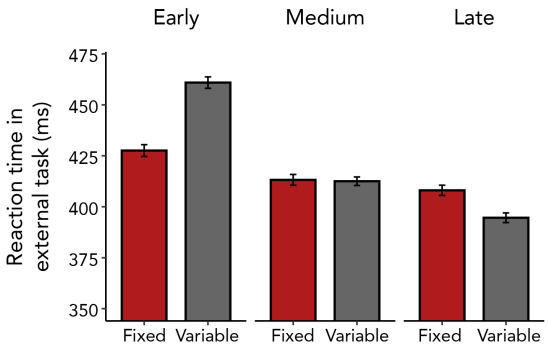

Supplementary Figure 8.9. Reaction times in the external task. RTs in the external task as a function of temporal predictability and external-task onset (see **Supplementary Table 8.8** for inferential statistics). Error bars indicate ±1 *SEM*.

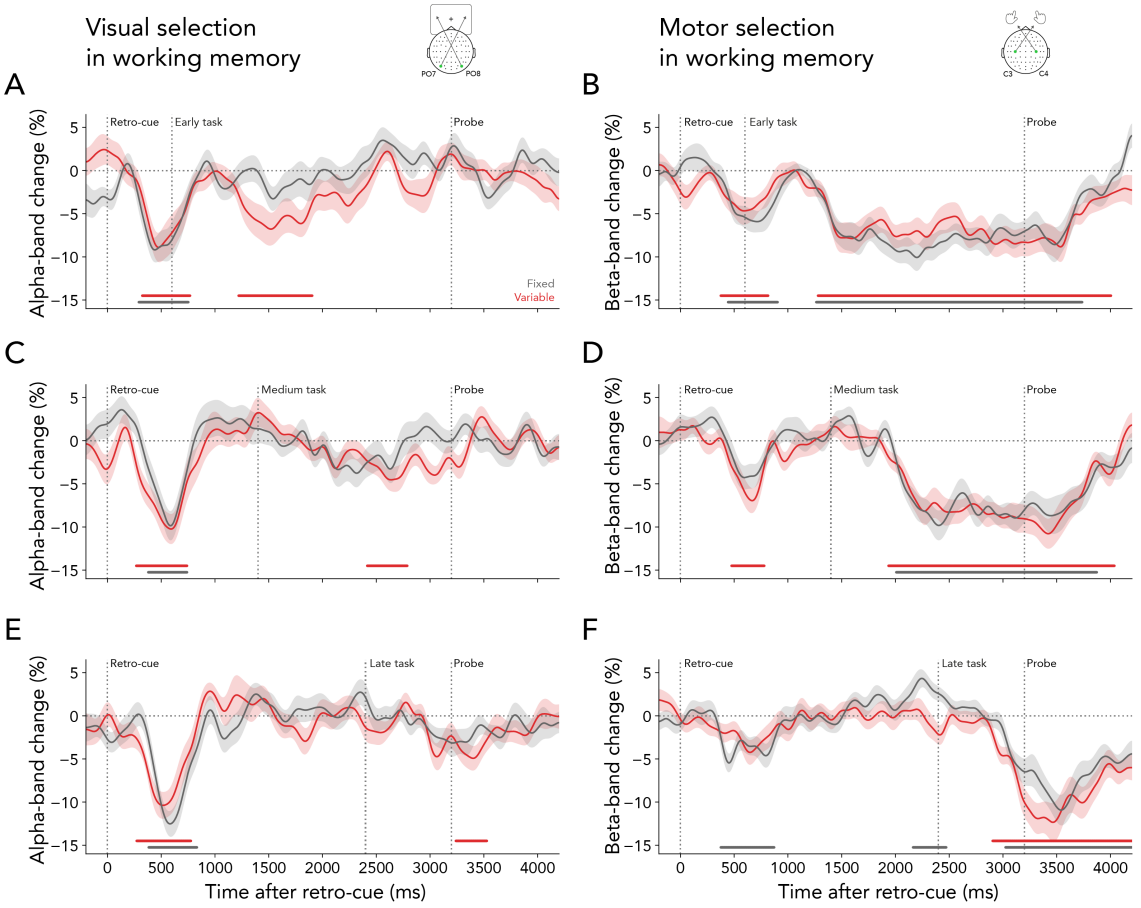

Supplementary Figure 8.10. No difference in visual and motor signatures between fixed- and variable-onset trials. (A) Time courses of lateralised visual alpha-band (8-12Hz) for fixed-early and variable-early external-task trials (fixed-early: 1st cluster $p = 0.002$, 2nd cluster $p < 0.001$; variable-early: cluster $p < 0.001$). **(B)** Time courses of lateralised motor beta-band (13-30 Hz) for fixed-early and variable-early external-task trials (fixed-early: 1st cluster $p = 0.025$, 2nd cluster $p < 0.001$; variable-early: 1st cluster $p = 0.017$, 2nd cluster $p < 0.001$). **(C)** Same as (A) but fixed-medium and variable-medium external-task trials (fixed- medium: 1st cluster $p < 0.001$, 2nd cluster $p = 0.017$; variable- medium: cluster $p = 0.002$). **(D)** Same as (B) but fixed-medium and variable-medium external-task trials (fixed- medium: 1st cluster $p = 0.034$, 2nd cluster $p < 0.001$; variable- medium: cluster $p < 0.001$). **(E)** Same as (A) but fixed-late and variable-late external-task trials (fixed-late: 1st cluster $p < 0.001$, 2nd cluster $p = 0.033$; variable-late: cluster $p < 0.001$). **(F)** Same as (B) but fixed-late and variable-late external-task trials (fixed-late: 1st cluster $p < 0.001$; variable-late: 1st cluster $p = 0.020$, 2nd cluster $p = 0.004$, 3rd cluster $p < 0.001$). Coloured horizontal lines indicate significant clusters. Black horizontal lines would indicate significant differences between conditions; however, no significant differences were observed for any of the comparisons. Time courses show *M ± SEM* across participants.

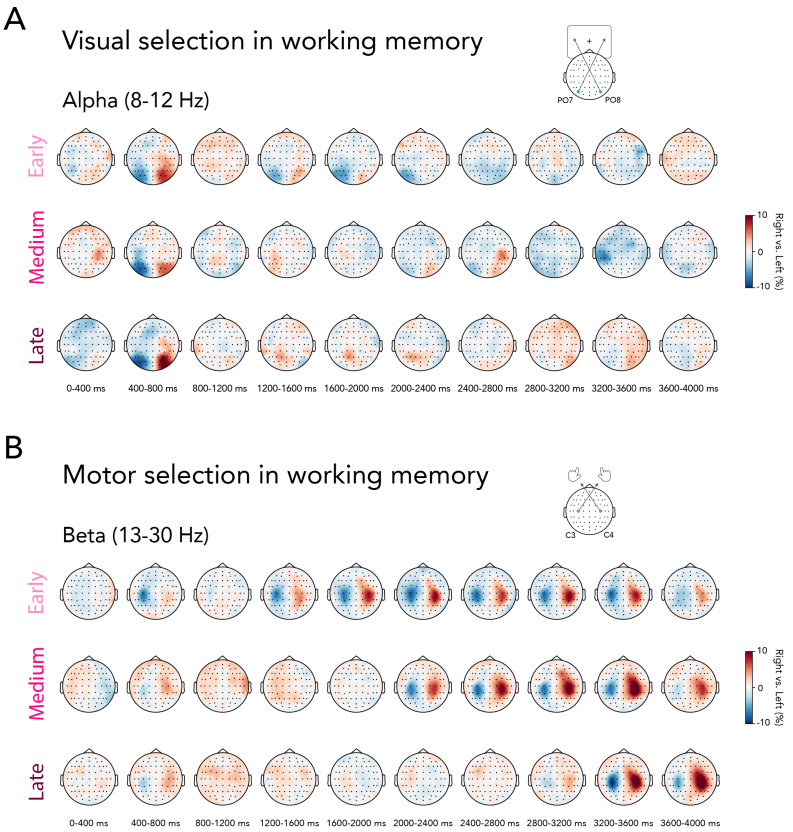

Supplementary Figure 8.11. Visual and motor topographies relative to the retro-cue onset. (A) Topographies of the difference in 8-12 Hz alpha power between trials in which the retro-cue indicated the selection of the right or left item in memory (right minus left), averaged in steps of 400 ms after retro-cue onset. **(B)** Same as (A) but showing the difference in 13-30 Hz beta power between trials in which the retro-cue indicated the right or left response hand.

8.3 Supplementary Information for Chapter 4

Supplementary Table 8.9. Inferential statistics for Experiment 2: reaction times and reproduction errors. (A) Main and interaction effects of RTs for the factors shift type (within vs. between), target domain (external vs. internal), and ISI (0 ms vs. 250 ms vs. 750 ms), tested with a 2×2×3 repeated-measures ANOVA. **(B)** same as (A) but for reproduction errors. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Experiment 2: Repeated-measures ANOVA of RTs.

	df	*F*	*p*	η^2_G
Shift Type**	1, 24	29.124	< 0.001	0.004
Target Domain*	1, 24	7.635	0.011	0.012
ISI**	2, 48	59.565	< 0.001	0.150
Shift Type × Target Domain	1, 24	0.355	0.557	< 0.001
Shift Type × ISI	2, 48	2.757	0.074	< 0.001
Target Domain × ISI	2, 48	0.296	0.745	< 0.001
Shift Type × Target Domain × ISI	2, 48	0.197	0.822	< 0.001

(B) Experiment 2: Repeated-measures ANOVA of reproduction errors.

	df	*F*	*p*	η^2_G
Shift Type*	1, 24	11.019	0.003	0.015
Target Domain**	1, 24	34.158	< 0.001	0.186
ISI**	2, 48	82.596	< 0.001	0.306
Shift Type × Target Domain*	1, 24	10.282	0.004	0.012
Shift Type × ISI	2, 48	0.002	0.998	< 0.001
Target Domain × ISI**	2, 48	12.726	< 0.001	0.077
Shift Type × Target Domain × ISI	2, 48	1.413	0.253	0.002

Supplementary Table 8.10. Inferential statistics for Experiment 3: reaction times and reproduction errors. (A) Pairwise comparisons of RTs for the factor target domain (external vs. internal) and the combined factor of shift type and location shift (within-different vs. between-different vs. between-same). **(B)** same as (A) but for reproduction errors. * indicates $p < 0.05$, ** indicates $p < 0.001$.

(A) Experiment 3: Bonferroni-corrected pairwise comparisons for RTs.

	df	*t*	$p_{Bonferroni}$	*d*
External: Within-different vs. Between-different	24	-1.274	1.000	0.255
External: Within-different vs. Between-same	24	-1.113	1.000	0.223
External: Between-different vs. Between-same	24	-0.156	1.000	0.031
Internal: Within-different vs. Between-different**	24	5.157	< 0.001	1.031
Internal: Within-different vs. Between-same*	24	3.436	0.013	0.687
Internal: Between-different vs. Between-same	24	0.586	1.000	0.117

(B) Experiment 3: Bonferroni-corrected pairwise comparisons for reproduction errors.

	df	*t*	$p_{Bonferroni}$	*d*
External: Within-different vs. Between-different	24	0.635	1.000	0.127
External: Within-different vs. Between-same	24	0.753	1.000	0.152
External: Between-different vs. Between-same	24	-0.118	1.000	0.024
Internal: Within-different vs. Between-different*	24	4.068	0.003	0.814
Internal: Within-different vs. Between-same**	24	4.677	< 0.001	0.935
Internal: Between-different vs. Between-same	24	-1.725	0.584	0.345

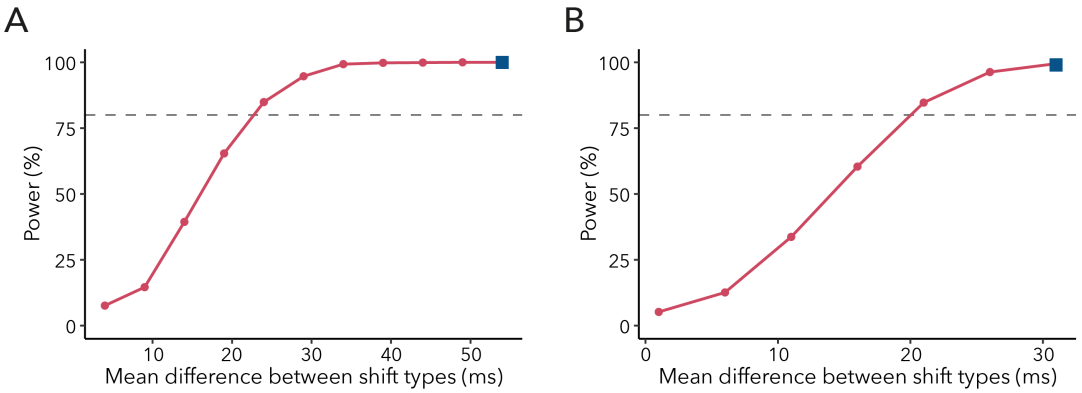

Supplementary Figure 8.12. Simulation-based power analysis for Experiment 1 and Experiment 2. (A) Experiment 1: Power simulation based on pilot data (*n* = 7). By systematically reducing the difference in RTs between the two shift-type conditions, it becomes evident that with a simulated sample size of *n* = 25, a statistical power exceeding 80% can still be attained even in the scenario where the RT effect is reduced by 50%. The RT difference observed in the pilot data was 54 ms (blue square). **(B)** Experiment 2: Power simulation based on data of Experiment 1 (*n* = 25). The RT difference between within- and between-domain shifts in Experiment 1 was 31 ms (blue square). Systematically reducing the RT effect between the two shift conditions revealed that with a simulated sample size of *n* = 25, a statistical power exceeding 80% can still be attained even in the scenario where the RT effect is reduced by 30%.

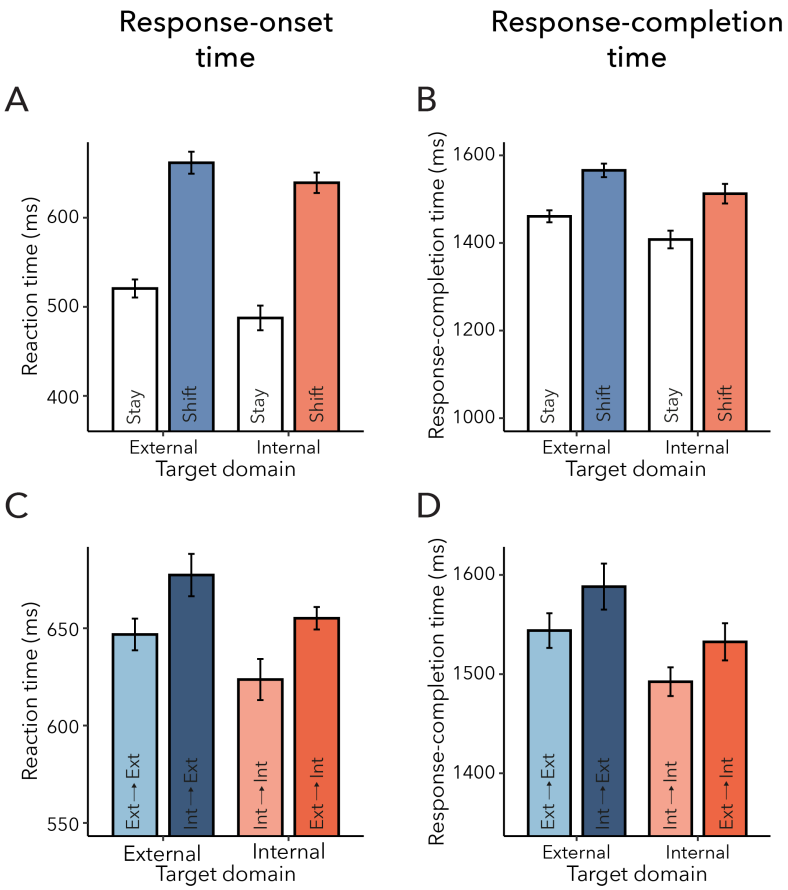

Supplementary Figure 8.13. Response-onset and response-completion times in Experiment 1. (A) RTs as a function of target domain and trial type as depicted in **Figure 4.2A**. **(B)** Response-completion time as a function of target domain and trial type. Participants showed slower response-completion times for shift than for stay trials ($F_{(1,24)}$ = 19.500, p < 0.001, η^2_G = 0.013) and for external than internal targets ($F_{(1,24)}$ = 8.737, p = 0.007, η^2_G = 0.003). The interaction between factors was not significant ($F_{(1,24)}$ < 0.001, p = 0.991, η^2_G < 0.001). **(C)** RTs as a function of target domain and shift type as depicted in **Figure 4.3A**. **(D)** Response-completion time as a function of target domain and shift type. Participants had slower dial-completion times in between- than within-shift trials ($F_{(1,24)}$ = 6.881, p = 0.015, η^2_G = 0.002) and in external- than internal-report trials ($F_{(1,24)}$ = 5.301, p = 0.030, η^2_G = 0.003). The interaction between factors was not significant ($F_{(1,24)}$ = 0.017, p = 0.898, η^2_G < 0.001). Error bars indicate ±1 *SEM*. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

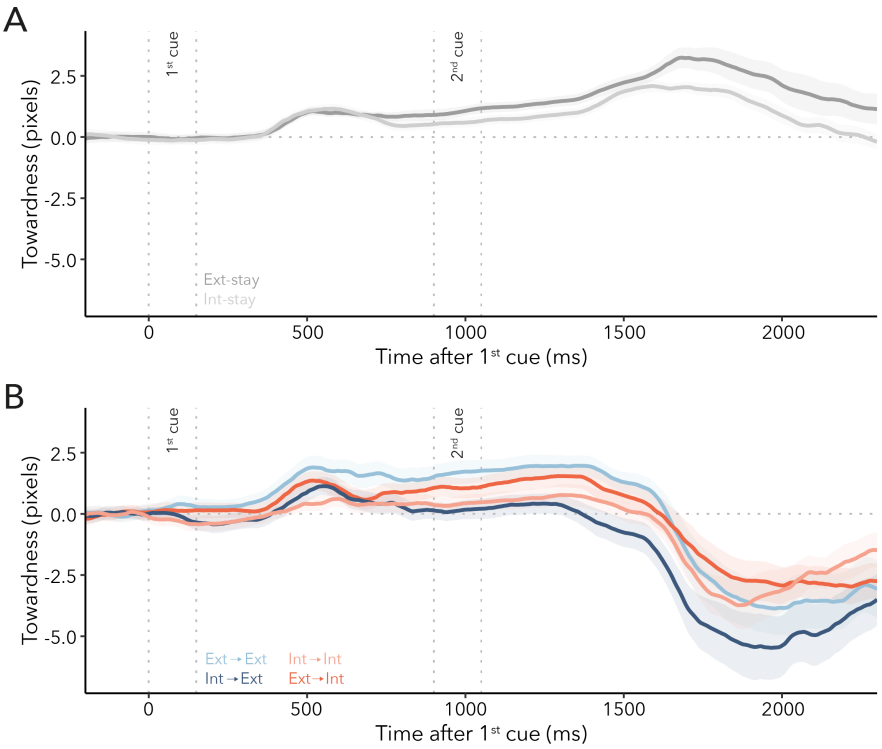

Supplementary Figure 8.14. Towardness time courses for each individual condition relative to the first cue. (A) Stay trials. **(B)** Shift trials. Shaded areas indicate ±1 *SEM*. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

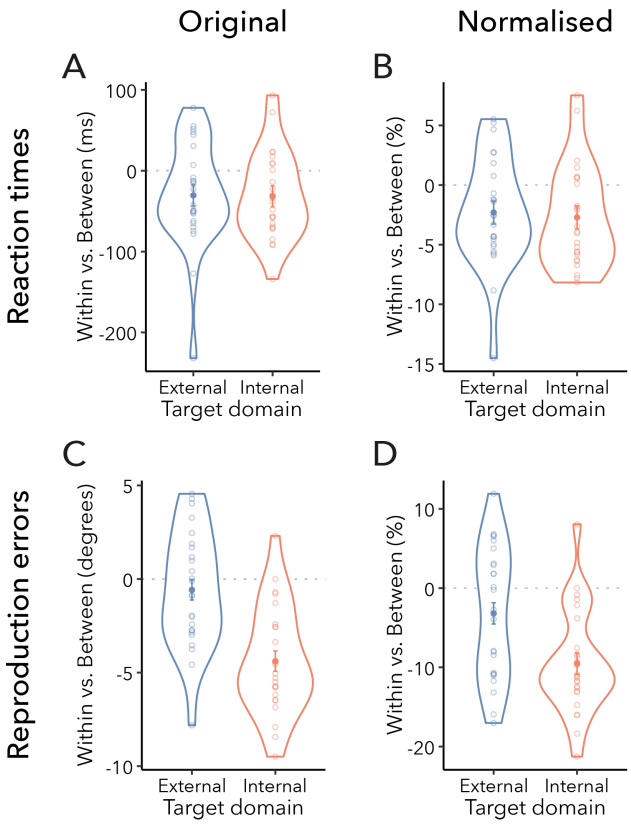

Supplementary Figure 8.15. Original and normalised between-domain shift-cost effects in Experiment 1. (A) Between-domain shift-cost effect for RTs as depicted in **Figure 4.3A**. **(B)** Normalised between-domain shift-cost effect for RTs (e.g., [(internal within – internal between) / (internal within + internal between)] × 100). Pairwise post-hoc comparisons revealed that there was a significant difference between external-to-external and internal-to-external shifts ($t_{(24)}$ = 2.501, $p_{Bonferroni}$ = 0.039, d = 0.500). Moreover, the difference between internal-to-internal and external-to-internal shifts was significant ($t_{(24)}$ = 3.206, $p_{Bonferroni}$ = 0.008, d = 0.641). **(C)** Between-domain shift-cost effect for reproduction errors as depicted in **Figure 4.3B**. **(D)** Normalised between-domain shift-cost effect for reproduction errors. Pairwise post-hoc comparisons revealed that there was a significant difference between internal-to-internal and external-to-internal shifts ($t_{(24)}$ = 7.293, $p_{Bonferroni}$ < 0.001, d = 1.549). However, the difference between external-to-external and internal-to-external shifts was not significant ($t_{(24)}$ = 1.921, $p_{Bonferroni}$ = 0.134, d = 0.384). Error bars indicate ±1 *SEM*. Dots represent individual participants.

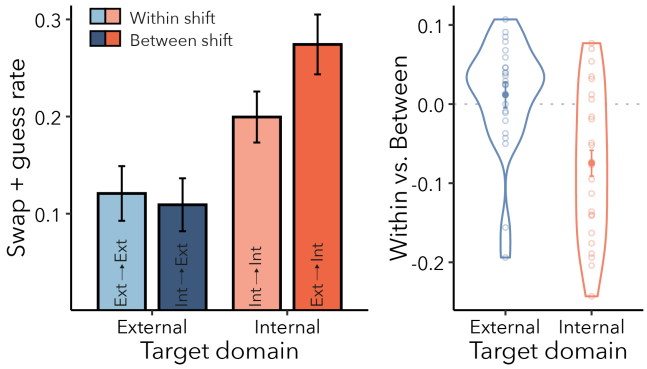

Supplementary Figure 8.16. Combined swap and guess rates in Experiment 3. Asymmetric between-domain shift cost for the combined measure of swaps and guesses. Participants made more combined swaps and guesses for internal as compared to external targets ($F_{(1,24)}$ = 7.058, p = 0.014, η^2_G = 0.126). Moreover, there was a main effect of shift type, showing more combined swaps and guesses for between- versus within-domain shifts ($F_{(1,24)}$ = 7.047, p = 0.014, η^2_G = 0.010). The interaction between target domain and shift type was significant as well ($F_{(1,24)}$ = 14.097, p < 0.001, η^2_G = 0.018). Post-hoc comparisons revealed that increased guesses and swaps occurred only for between- versus within-domain shifts when reporting internal targets ($t_{(24)}$ = 3.990, $p_{Bonferroni}$ = 0.001, d = 0.789), but not when reporting external targets ($t_{(24)}$ = 0.831, $p_{Bonferroni}$ = 0.828, d = 0.166). Error bars indicate ±1 *SEM*. Dots represent individual participants. Ext→Ext = external-to-external, Int→Int = internal-to-internal, Int→Ext = internal-to-external, Ext→Int = external-to-internal.

8.4 Supplementary Information for Chapter 5

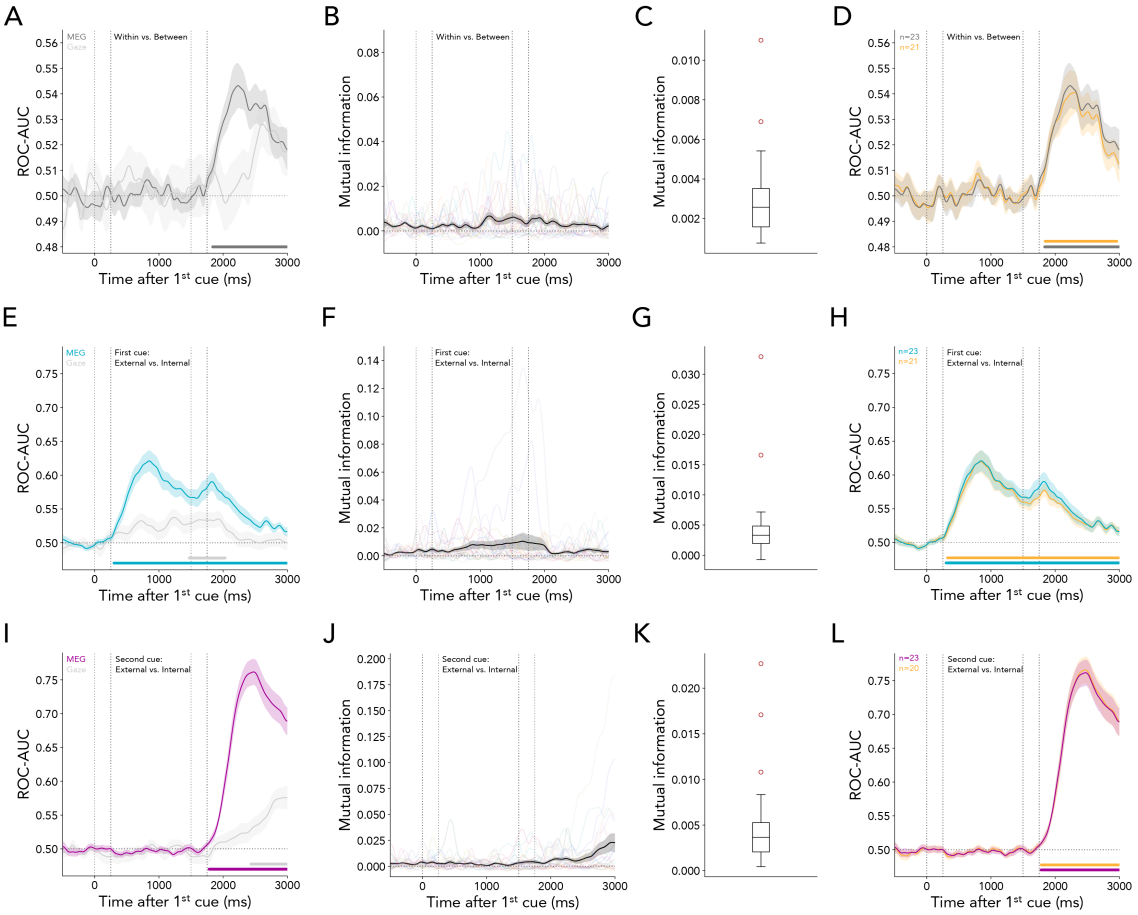

Supplementary Figure 8.17. Eye movements do not account for MEG decoding results. (A) Average time-resolved decoding of within- versus between-domain shifts based on MEG and eye-tracking data. Only MEG datasets were included that had an accompanying eye tracking dataset (n = 23). Horizontal lines indicate significant clusters (MEG decoding: ~1840-3000 ms, cluster $p < 0.001$) **(B)** Average mutual information between MEG and eye-tracking data for the within- versus between-shift decoding. Faded lines indicate individual participants **(C)** Box-and-whisker plot according to Tukey (1977) displaying distribution of average mutual information across -500 to 3000 ms relative to first-cue onset. Red dots represent outliers. **(D)** Average time-resolved decoding of within- versus between-domain shifts based on MEG after removing outliers (n=23: ~1840-3000 ms, cluster $p < 0.001$; n=21: ~1840-3000 ms, cluster $p < 0.001$). **(E)** Same as (A) but for external versus internal first cues (MEG decoding: ~300-3000 ms, cluster $p < 0.001$; eye decoding: ~1470-2030 ms, cluster $p = 0.042$). **(F)** Same as (B) but for external versus internal first cues. **(G)** Same as (C) but for external versus internal first cues. **(H)** Same as (D) but for external versus internal first cues (n=23: ~300-3000 ms, cluster $p < 0.001$; n=21: ~320-3000 ms, cluster $p < 0.001$). **(I)** Same as (A) but for external versus internal second cues (MEG decoding: ~1780-3000 ms, cluster $p < 0.001$; eye decoding: ~2430-3000 ms, cluster $p = 0.010$). **(J)** Same as (B) but for external versus internal second cues. **(K)** Same as (C) but for external versus internal second cues. **(L)** Same as (D) but for external versus internal second cues (n=23: ~1780-3000 ms, cluster $p < 0.001$; n=20: ~1780-3000 ms, cluster $p < 0.001$).

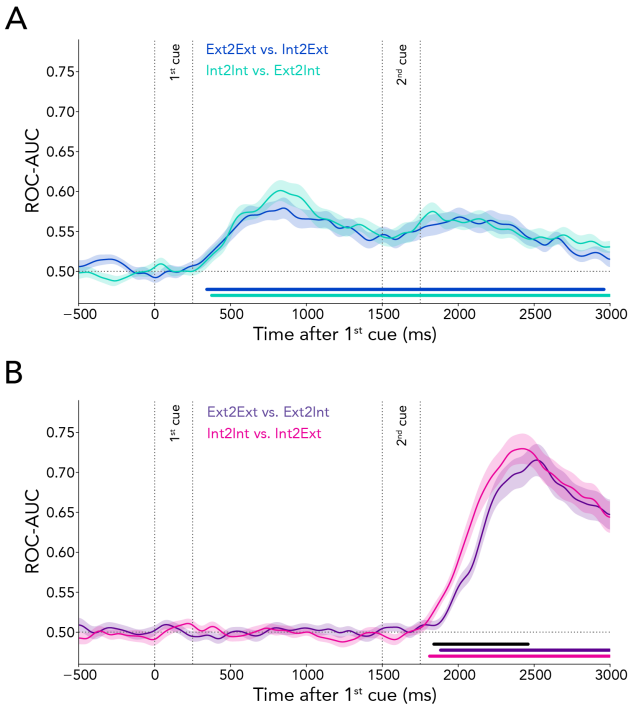

Supplementary Figure 8.18. Time-resolved MEG decoding of different trial types. (A) Average time-resolved classifier performance for distinguishing external-to-external from internal-to-external shifts, and internal-to-internal from external-to-internal shifts. Time courses show $M \pm SEM$ across participants. Coloured horizontal lines indicate significant clusters (external-to-external vs. internal-to-external cluster: ~340-2960 ms, cluster $p < 0.001$; internal-to-internal vs. internal-to-external cluster: ~380-3000 ms, cluster $p < 0.001$). **(B)** Same as (A) but decoding internal-to-internal versus internal-to-external shifts, and external-to-external versus external-to-internal shifts (external-to-external vs. external-to-internal cluster: ~1880-3000 ms, cluster $p < 0.001$; internal-to-internal vs. internal-to-external cluster: ~1810-3000 ms, cluster $p < 0.001$). Black horizontal line denotes significant difference between decoding time courses (~1840-2460 ms, cluster $p < 0.001$). Since statements regarding the exact onset of an effect are not supported by permutation tests (Sassenhagen & Draschkow, 2019), I quantified the timings of the decoding time courses more formally by applying a jack-knife approach to compare the point at which the decoding time courses reached 50% of their peak value. This revealed significantly earlier onset times for internal-to-internal versus internal-to-external decoding (M = 2046.240, SEM = 20.800) as compared to external-to-external versus external-to-internal decoding (M = 2133.92, SEM = 16.738; $t_{(24)}$ =-3.379, p = 0.002).

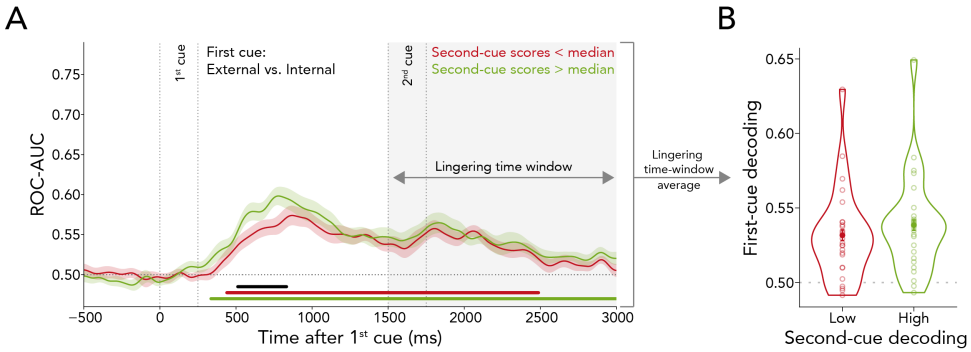

Supplementary Figure 8.19. No evidence for a trade-off in decodability between first- and second-cued attentional domains. (A) I median-split the MEG data into trials with low versus high decoding performance in distinguishing the second-cued domain following the second cue (i.e., 0-1500 ms after second-cue onset). Next, I trained a classifier separately on both datasets to predict the first-cued domain (external vs. internal). In trials with low and high second-cue decoding, I could successfully predict the first-cued attentional domain after first-cue appearance and the decodability remained significant following the second cue (low decoding: ~440-2490 ms, cluster $p < 0.001$; high decoding: ~340-3000 ms, cluster $p < 0.001$). I did not observe any trade-off in the decoding of the first-cued attentional state between low versus high second-cue decoding following the second cue (i.e., grey area). There was an earlier difference in the decodability of external versus internal first-cue trials when the decodability of the second cue was low versus high (~510-830 ms, cluster $p = 0.012$). Coloured horizontal lines indicate significant clusters, the black line denotes significant difference between time courses. **(B)** Average first-cued domain decoding in the second-cued period (i.e. grey area in (A)) as a function of low versus high second-cued domain decoding. There was no difference between low and high second-cue decoding in predicting the first-cued domain ($t_{(24)} = 1.348$, $p = 0.190$, $d = 0.270$). Error bars indicate ±1 *SEM*. Dots represent individual participants.