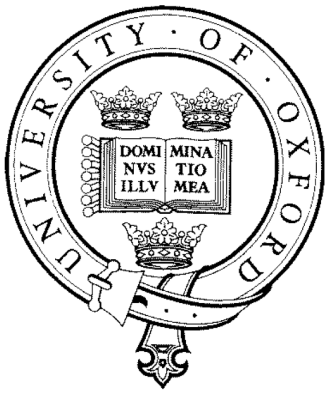
Effect of Reward on Visual Perceptual Decison-Making in Humans and Non-Human Primates

Nela Cicmil



Thesis submitted for the degree of Doctor of Philosophy

Department of Physiology, Anatomy & Genetics

University of Oxford

Trinity 2012

Abstract

When primates make decisions about sensory signals, their choices are biased by the costs and benefits associated with different possible outcomes. However, much remains unknown about the neuronal mechanisms by which reward information is integrated into the perceptual decision-making process. I used electrophysiological, psychophysical and imaging techniques to explore the effect of reward on sensory representations, whilst humans and monkeys made perceptual judgements about structure-from-motion (SFM) stimuli.

Electrical microstimulation of visual area V5/MT in the macaque monkey was used to bias perceptual judgements, under different available reward sizes for correct choices. The behavioural effect of microstimulation interacted with available reward, and, in the context of a drift diffusion model of decision-making, the results demonstrated that reward must influence sensory processing before visual signals and microstimulation signals are accumulated in sensorimotor areas. In a parallel human psychophysics experiment, viewers made decisions about SFM stimuli whilst their choices were biased towards one outcome or the other by unequal pay-offs. A full drift-diffusion model was fitted to human choice and reaction time data. There was a significant effect of reward on model drift rate, a parameter known to be dependent upon evidence represented in sensory brain regions.

A second set of experiments used magnetoencephalography (MEG) to explore activity in visual areas in human cortex and the effects of reward. Brain responses to retinotopic visual stimuli were localized with three different MEG source analysis methods, and localization accuracy was evaluated by comparison with fMRI maps obtained in the same individuals. The results demonstrated that the beamformer and minimum norm estimate (MNE) methods were most suitable for investigating early visual activity with MEG. Human brain activity was then recorded with MEG whilst viewers made perceptual judgements about SFM stimuli, under unequal pay-offs. The results revealed an effect of reward size on early MEG responses in the region of the occipital cortex and visual precuneus.

Taken together, the experiments presented in this thesis provide consistent evidence that information about reward can influence the processing of sensory information during perceptual decisions.

Contents

1. Introduction 2

1.1. Neural correlates of visual perceptual decisions 4

1.2. The drift-diffusion model of decision-making 6

1.3. Representations of reward in the brain 8

1.4. Reward and perceptual decisions 11

1.5. MEG and the human visual system 15

1.6. Aims and scope of this thesis 16

2. Reward and electrical microstimulation in V5/MT: Experimental data 18

2.1. Introduction 18

2.2. Methods 21

2.2.1. Animals 21

2.2.1.1. Locations and ethics statements 21

2.2.2. Visual stimuli 22

2.2.3. Task 23

2.2.3.1. Reward Schedule 23

2.2.4. V5/MT multiunit recording 24

2.2.5. V5/MT MU receptive field mapping 26

2.2.6. Selection and characterization of stimulation sites 26

2.2.7. Microstimulation experimental procedure 26

2.2.8. Pseudo-stimulation: psychophysical reproduction of microstimulation effect 27

2.2.9. Data Analysis 28

2.2.9.1. Quantification of microstimulation shift effect 28

2.2.9.2. Quantification of effect of expected reward size on microstimulation 30

2.2.9.3. Normalization and pooling of data across multiple experimental sessions 30

2.2.9.4. Quantification of interaction between reward and microstimulation with logistic regression . . . 30
2.3. Results . . . 31
2.3.1. Microstimulation of visual area V5/MT shifts perceptual decisions . . . 32
2.3.2. Expected reward size interacts with the effect of electrical microstimulation in area V5/MT . . . 32
2.3.3. Reward conditions do not differ in the proportion of microstimulated trials, nor in the distribution of disparity values . . . 37
2.3.4. Wearing-off of microstimulation does not explain reward effect . . . 37
2.3.5. Animal cannot adapt discrimination criterion selectivity on large reward trials . . . 38
2.4. Discussion . . . 43
2.4.1. Alternative explanations of the experimental data . . . 43
2.4.2. Limitations of the present study . . . 44
2.4.3. Future work . . . 45

3. Reward and electrical microstimulation in V5/MT: Drift-diffusion model 46
3.1. Introduction . . . 46
3.1.1. Diffusion model for perceptual decisions and microstimulation . . . 47
3.1.2. Possible interactions between reward and decision-making . . . 49
3.2. Methods . . . 52
3.2.1. Drift-diffusion model fitting . . . 52
3.2.2. Quantifying interactions between reward and drift-diffusion model parameters . . . 53
3.3. Results . . . 53
3.3.1. Reward affects sensitivity to stimulus strength (drift rate) . . . 53
3.3.2. Reward also affects the distance to the decision bounds . . . 55
3.3.3. Summary . . . 56
3.4. Discussion . . . 56
3.4.1. Reward and the sensitivity of sensory evidence to visual input (drift rate) . 57
3.4.2. Reward and distance to decision bounds . . . 59
3.4.3. Limitations of the drift-diffusion model . . . 61
3.4.4. Summary and future directions . . . 61

4. Cylinder Reward Psychophysics in Human Participants 63
4.1. Introduction . . . 63

4.2. Materials & Methods . . . 64
4.2.1. Participants . . . 64
4.2.2. Stimuli . . . 65
4.2.3. Task . . . 65
4.2.3.1. Training . . . 65
4.2.3.2. Basic task (no reward) . . . 66
4.2.3.3. Reward task . . . 67
4.2.4. Analysis procedures . . . 68
4.2.4.1. Preprocessing . . . 68
4.2.4.2. Psychometric functions & optimal strategy analysis . . . 69
4.2.4.3. Quantile-probability (QP) plots and quantile-quantile (QQ) plots 69
4.2.4.4. Model Fitting . . . 69
4.2.4.5. Model Selection . . . 70
4.3. Results . . . 72
4.3.1. Preprocessing: reaction time outliers . . . 72
4.3.2. Drift-diffusion model for cylinder discrimination in a reaction time task . 72
4.3.2.1. Psychometric functions and reaction times . . . 72
4.3.2.2. Quantile-probability and quantile-quantile plots . . . 74
4.3.2.3. Effect of cylinder disparity on drift-diffusion model parameters . 76
4.3.3. Effects of unbalanced reward conditions on perceptual decisions . . . 76
4.3.3.1. Behavioural results: psychometric functions & RT distributions . 76
4.3.3.2. Effects of unbalanced rewards on drift-diffusion model parameters 77
4.3.3.3. Strategies for reward maximization . . . 79
4.3.3.4. Effect of reward cue duration on perceptual decisions . . . 80
4.3.4. Effects of balanced reward sizes on perceptual decisions . . . 81
4.3.4.1. Psychometric functions and reaction times . . . 81
4.3.4.2. Effects on drift-diffusion model parameters . . . 81
4.3.5. Summary of results . . . 82
4.4. Discussion . . . 82
4.4.1. The cylinder discrimination task follows a drift-diffusion process . . . 82
4.4.2. Cylinder disparity affects only the drift rate of the diffusion process . . . 84
4.4.3. Reward bias affects both drift rate and distance-to-bound . . . 85
4.4.4. The ‘under-biasing’ phenomenon . . . 86
4.4.5. Reward cue duration decreases choice bias . . . 86
4.4.6. Balanced reward conditions . . . 87

4.4.7. Future work . . . 88

5. MEG Materials and Methods 104

5.1. Introduction to MEG . . . 104

5.2. MEG Data Collection . . . 106

5.2.1. Scanner & Acquisition Systems . . . 106

5.2.2. Participant Preparation & Head Position Recording . . . 107

5.2.3. Eyetracking . . . 107

5.2.4. Stimulus Presentation, Timing & Button Presses . . . 107

5.2.5. MEG Recording . . . 108

5.3. MEG Data Preprocessing . . . 108

5.3.1. Maxfiltering & HPI correction . . . 108

5.3.2. Event Assignment & Epoching . . . 109

5.3.3. Artefact Removal . . . 109

5.4. MEG Analysis: Sensor Space . . . 110

5.4.1. Time-frequency (TF) decomposition . . . 110

5.5. MEG Analysis: Source Reconstruction . . . 110

5.5.1. Structural MRI . . . 111

5.5.2. Anatomical reconstruction (Freesurfer) . . . 111

5.5.3. MEG-MRI coregistration . . . 111

5.5.3.1. Co-registration with MNE software . . . 111

5.5.3.2. Co-registration with SPM8 . . . 112

5.5.4. Minimum Norm Estimate (MNE) . . . 113

5.5.5. SPM8: multiple sparse priors (MSP) . . . 114

5.5.6. LCMV Beamformer . . . 114

5.5.6.1. General Linear Model (GLM) statistics: . . . 115

6. MEG Retinotopy 117

6.1. Introduction . . . 117

6.1.1. Retinotopy in early visual cortical areas . . . 118

6.1.2. MEG source reconstruction methods . . . 120

6.1.2.1. Minimum Norm Estimate (MNE) . . . 120

6.1.2.2. Multiple Sparse Priors (MSP) . . . 120

6.1.2.3. Beamforming . . . 121

6.1.3. Previous MEG studies of visual retinotopy . . . 122

6.2. Methods . . . 123

6.2.1. MEG Retinotopy . . . 124
6.2.1.1. Subjects and stimuli . . . 124
6.2.1.2. Data Acquisition & Preprocessing . . . 124
6.2.1.3. MNE source reconstruction . . . 125
6.2.1.4. LCMV beamformer source reconstruction: initial evoked response 125
6.2.1.5. LCMV beamformer source reconstruction: time-frequency decomposition . . . 126
6.2.1.6. SPM8 MSP source reconstruction . . . 126
6.2.1.7. Morphing 3D images to cortical surface . . . 127
6.2.2. Retinotopy: fMRI . . . 127
6.2.2.1. Stimuli . . . 127
6.2.2.2. Data Acquisition . . . 127
6.2.2.3. Analysis & ROIs . . . 128
6.2.3. MEG-fMRI Comparison . . . 130
6.2.4. Threshold-accuracy functions . . . 131
6.3. Results I: Polar Angle (Quadrants) . . . 133
6.3.1. Visually evoked responses at sensor level . . . 133
6.3.2. Evaluation of three source reconstruction methods . . . 135
6.3.3. Time-frequency beamformer: stimulus-induced oscillations . . . 137
6.3.4. SPM8's MSP: averaged responses, trial-wise responses, and statistical contrast . . . 137
6.3.5. Effect of cut-off threshold on MEG localization accuracy . . . 141
6.3.6. Summary . . . 142
6.4. Results II: Eccentricity (Rings & mini-rings) . . . 146
6.4.1. Visually evoked responses at sensor level . . . 146
6.4.2. Evaluation of source localization accuracy . . . 146
6.4.3. Mini-rings . . . 148
6.4.4. Summary . . . 155
6.5. Discussion . . . 161
6.5.1. Source reconstructions for angular retinotopy . . . 161
6.5.2. Source reconstructions for eccentricity mapping . . . 163
6.5.3. Relationship of MEG to fMRI signal . . . 164
6.5.4. Conclusions and future directions . . . 165

7. Human perceptual decision-making under reward: a MEG study 167
7.1. Introduction . . . 167

7.2. Materials & Methods . . . 168
7.2.1. Data Collection . . . 168
7.2.1.1. Participants . . . 168
7.2.1.2. Cylinder Discrimination Task . . . 169
7.2.1.3. Eye movement tracking . . . 170
7.2.2. Data Analysis . . . 171
7.2.2.1. Preprocessing: outlier reaction times . . . 171
7.2.2.2. Behavioural choices and effect of reward condition . . . 171
7.2.2.3. MEG preprocessing . . . 171
7.2.2.4. MEG source reconstruction . . . 172
7.2.2.5. Experimental contrasts . . . 172
7.2.2.6. Group statistics . . . 173
7.2.2.7. Display of MEG source analysis results . . . 174
7.2.2.8. Eye movement analysis . . . 174
7.3. Results . . . 174
7.3.1. Behavioural results: effect of reward . . . 174
7.3.2. Lateralized motor response peri-button press . . . 177
7.3.3. Visual responses to onset of supra-threshold cylinders . . . 177
7.3.4. Differential responses to biasing versus balanced reward cues . . . 180
7.3.5. Reward association affects responses to cylinders in visual precuneus and prefrontal cortex . . . 182
7.4. Discussion . . . 188
7.4.1. Lateralization of brain responses to button presses . . . 188
7.4.2. Source reconstructions of MEG responses to supra-threshold cylinders . . . 189
7.4.3. Differential activity in early visual areas, visual precuneus, and prefrontal cortex in visual perceptual decision-making that is biased by reward information . . . 190
7.4.4. Significance testing and sample-size . . . 193
7.4.5. Future directions . . . 193

8. Conclusions and future directions 194
8.1. Summary . . . 194
8.2. Future directions . . . 195
8.2.1. Further analysis of MEG data & the potential of connectivity studies . . . 195
8.2.2. Single cell electrophysiology: reward and sensory neurons . . . 197
8.2.3. Computational modeling of reward and sensory information processing . 198

Appendices 199

A. Mathematical details of a drift-diffusion process applied to microstimulation data 199

A.1. Basic drift-diffusion model 199

A.1.1. Stimulus strength C and response direction R 199

A.1.2. The diffusion process: accumulating evidence $U(t)$ 200

A.1.3. The diffusion process with two bounds, A and B 201

A.1.3.1. Special case of symmetrical bounds B and constant variance σ^2 . 202

A.1.4. Linear scaling of drift rate μ' by stimulus strenght C 202

A.2. Adaptation of the drift-diffusion model to microstimulation experiment 203

A.2.1. Representation of electrical microstimulation S 203

A.2.2. Representation of the null bias x 204

B. Cylinder task: participant instruction sheets 205

Acknowledgements

Foremost, I would like to thank my supervisors Dr Kristine Krug and Professor Andrew Parker for invaluable academic guidance.

I am also indebted to my colleagues Bruce Cumming, Stuart Gilson, Holly Bridge, Lori Minini, Betina Ip, Sven Braeutigam, Jennifer Swettenham, Anling Rao and Mark Woolrich for their help with many and various aspects of the work presented here. I thank Chris Summerfield, Roozbeh Kiani and Geoff Boynton for valuable discussions. Additionally, this work would not have been possible without the financial support of Exeter College, University of Oxford.

I would also like to thank my friends and my family, especially my mother, for their support and friendship, which made all the difference.

Declaration

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

Chapter 2: Electrophysiological and microstimulation data for one animal (Fle) was collected by Dr. Bruce G. Cumming, Professor Andrew J. Parker and Dr. Kristine Krug. Pseudo-stimulation data for one animal (Ica) was collected by Dr. Bruce G. Cumming.

Chapter 4: Psychophysics data for four participants was collected by Mr Samuel Mills.

Chapter 6: Retinotopic fMRI data was collected by Dr. Holly Bridge, Dr. Loredana Minini and Dr. Betina Ip. Angular retinotopic fMRI maps for 3 subjects were delineated by Dr. Loredana Minini.

List of Abbreviations

-

2AFC	two-alternative forced-choice
aMRI	anatomical (structural) magnetic resonance imaging
ACOPE	absolute contrast of parameter estimates
COAPE	contrast of absolute parameter estimates
COPE	contrast of parameter estimates
CW / CCW	clockwise / counter-clockwise
DLPFC	dorsolateral prefrontal cortex
DV	decision variable (of drift-diffusion model)
EEG	electroencephalography
FFX	fixed effects group level GLM statistics
fMRI	functional magnetic resonance imaging
GLM	general linear model
hMT+	human middle temporal areas
IID	independent and identically distributed (MEG source reconstruction)
LCMV	linearly constrained minimum variance (beamformer)
MEG	magnetoencephalography
MFX	mixed effects group level GLM statistics
MNE	minimum norm estimate
MNI152	Montreal Neurological Institute template brain (average 152 subjects)
MU	multi-unit group of neurons
MSP	multiple sparse priors (MEG source reconstruction)
NHP	non-human primate
OHBA	Oxford Centre for Human Brain Activity
OSL	OHBA software library
SFM	structure-from-motion
SPM	statistical parametric mapping
TF	time-frequency
V5/MT	visual area V5 / middle temporal area (macaque)

-

1. Introduction

The way we see things is affected by what we know or what we believe.

John Berger, *Ways of Seeing* (1972)

To perceive is to form a mental representation of sensory information about the external world. The biological process of sensory transduction transforms the information contained in energetic signals from the environment into neuronal activity in sensory cortical areas of the nervous system. It is well-established that perception depends upon this neuronal activity, most compellingly from studies on neurological patients who report selective or total lack of subjective visual perception after lesions to visual brain regions (Zihl *et al.*, 1983; Zeki, 1990; Weiskrantz, 1986). Nevertheless, the external world is not there to be discovered or recognized for exactly what it is, but is always accessed by the viewer in an active way. In the first instance, we quite literally choose what we want to see by directing our eyes towards the most salient sources of visual information in any given scene (Yarbus, 1967). At a higher level, the subjective percept of a particular visual stimulus, or image, can significantly differ between viewers (or within the same viewer at different times), even when the stimulus itself remains exactly the same:



Figure 1.1: An example of the old woman / young lady optical illusion.

Such visual illusions demonstrate that even when we have unimpaired access to the information depicting what is literally there in front of our eyes, we still don't necessarily see it for what it is:

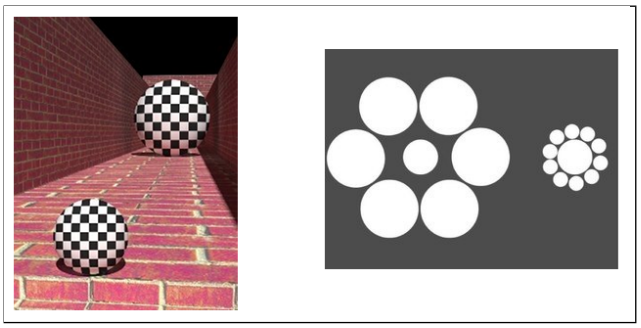


Figure 1.2.: Some examples of objects that are in fact the same size but appear to be different due to context: on the left, the two black-and-white spheres; on the right, the circles at the centre of each group.

The processes of detection, discrimination and recognition of sensory inputs allow us to respond to our environment with a particular behavioural response out of a range of possible options. Making choices on the basis of sensory information has been termed **perceptual decision-making** (Gold & Shadlen, 2007; Heekeren *et al.*, 2008).

As we have seen, the subjective approach of the individual and the context in which a particular image is presented strongly affects how it is perceived. Why might this occur? Theories of perceptual hypothesis-testing and perception as Bayesian inference argue that it is optimal for the visual system to make inferences about the world based on both image data and prior models (knowledge) about the world, with the final percept reflecting how visual information is interpreted in the context of these priors (Knill & Richards, 1996; Gregory, 1997). Many similar theories of perception exist and can apply to low-level, automatic processes or higher, cognitive-level effects (e.g. Friston, 2008).

It is a tautology that the most important perceptual decisions are those that have the most significant costs and benefits hanging in the balance of their outcome. In cognitive neuroscience, prior models about costs and benefits are investigated under the concept of **reward**, or value. Reward information has been shown to strongly affect behavioural decisions, by biasing organisms towards certain choices or causing them to implement particular strategies (Herrnstein, 1961; Sugrue *et al.*, 2005). The aim of this thesis is to investigate how the size of reward associated with perceptual outcomes influences the processing of sensory information that underlies visual perceptual decision-making.

1.1. Neural correlates of visual perceptual decisions

The end of the last century saw tremendous progress in research into the links between neural activity and the perceptual judgements of conscious observers (Newsome *et al.*, 1989a; Salzman *et al.*, 1990; Celebrini & Newsome, 1994, 1995; Britten *et al.*, 1996; reviewed by Parker & Newsome, 1998). Much of this rested on the use of two-alternative forced-choice (2AFC) perceptual tasks in experiments that simultaneously measured brain activity. The most commonly-used of these perceptual tasks has been the visual motion discrimination task, or 'motion task' (Newsome *et al.*, 1989b). In each trial, the viewer is presented with a dynamic random dot display and must decide in which of two possible directions the majority of the dots are moving (Figure 1.3). In animal studies, two reponse targets represent each of the two possible choices, and the animal indicates its perceptual choice with an eye movement towards the corresponding target. In studies with human volunteers, the viewer generally presses a button to indicate the choice. The difficulty of motion direction discrimination can be varied trial-by-trial by changing the percentage of dots moving in the correct direction, known as the motion *coherence*. At 0% coherence, the direction of the stimulus is ambiguous and viewers tend to choose each direction with equal likelihood.

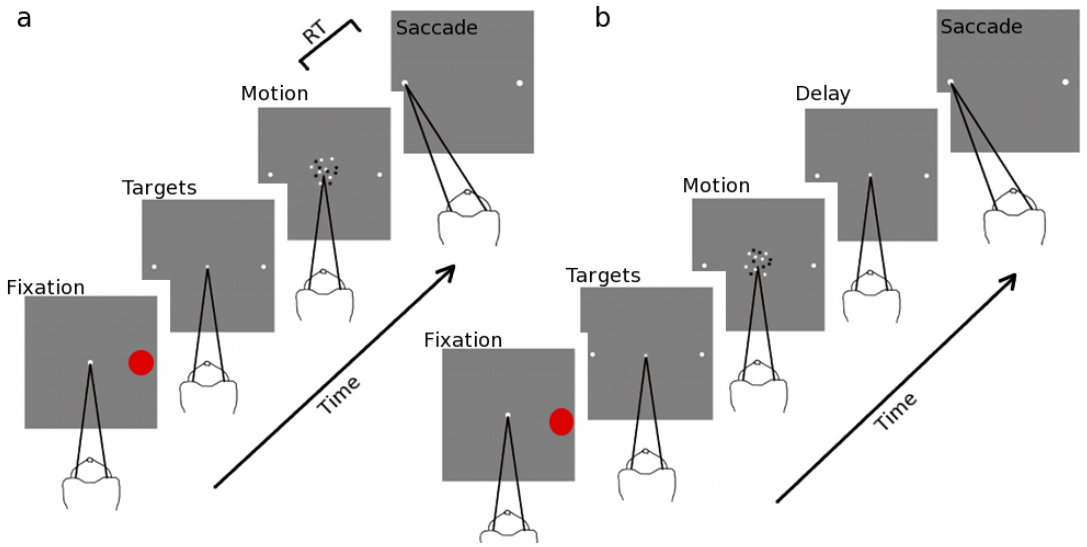


Figure 1.3.: Schema of the 'motion task', in which the viewer must discriminate the direction of motion of the majority of dots in the dynamic dot display. Monkeys indicate their perceptual decisions with an eye movements to the corresponding saccade target. **a** Reaction time (RT) motion task; **b** Fixed duration motion task. Red shading indicates the response field *RF* of a recorded LIP neuron that covers one of the response targets.

Electrophysiological studies with macaques have shown that the activity of single neurons in visual area V5/MT can predict the animals' subjective choices about ambiguous stimuli (Celebrini & Newsome, 1994; Britten *et al.*, 1996). This correlation was quantified by a measure known as

'choice probability' (CP), where CP values significantly different to 0.5 indicate a correlation between the neuronal activity and subjective perceptual choice (Parker & Newsome, 1998). Choice probability for area V5/MT in the motion task is 0.55 (Britten *et al.*, 1996), and the sensitivity of V5/MT neurons to motion coherence can account for the the precision of monkeys' judgements in the task (Britten *et al.*, 1992). The causal role of area V5/MT in perceptual judgements in the motion task was demonstrated by electrical microstimulation of direction-selective neuronal clusters, which produced a measurable shift in the animals' responses towards the preferred direction of the stimulated neurons (Salzman *et al.*, 1990, 1992). Together with results from lesion studies that show that damage to area V5/MT impairs motion discrimination (Pasternak & Merigan, 1994; Lauwers *et al.*, 2000), this comprises a strong body of evidence that signals in visual area V5/MT contribute directly to perceptual judgements in the motion task. Electrophysiology experiments have also implicated V5/MT in perceptual judgements about coarse binocular disparity (Uka & DeAngelis, 2006; DeAngelis & Newsome, 2007) and structure-from-motion discrimination (Dodd *et al.* (2001), where CP = 0.67; Krug *et al.* (2004)); area V2 in fine disparity discrimination (Nienborg & Cumming, 2006, 2009); inferotemporal cortex (IT) in face categorization (Afraz *et al.*, 2006; Kiani *et al.*, 2007); and somatosensory cortex for vibrotactile frequency discrimination (Romo *et al.*, 2000, 2002). In humans, lesion studies and non-invasive neuroimaging methods confirm the role of the analogous visual areas in perceptual decisions about visual motion and binocular disparity (Barton *et al.*, 1996; Neri *et al.*, 2004; Brouwer & van Ee, 2007; Lynn & Curran, 2010; Minini *et al.*, 2010). This body of work has greatly increased our understanding of the neuronal activity that represents sensory evidence upon which perceptual judgements are based.

Sensory neuroscience was strongly influenced by signal detection theory (SDT), a computational framework used to quantify signals amongst noise, which allowed researchers to infer underlying distributions of sensory representations from organisms' behaviour (Green & Swets, 1966; Johansson & Vallbo, 1979). Quantitative models, fitted to the electrophysiological data described above, were used to extend our understanding of the specific neuronal mechanisms involved in sensory representations that underlie perceptual decisions. Shadlen *et al.* (1996) gave a quantitative account of the number and tuning selectivity of V5/MT neurons most likely to underlie perceptual decisions in the motion task, based upon data about choice probabilities, neuronal selectivity, inter-neuronal noise correlations and organisms' motion-detection thresholds. Gold & Shadlen (2001) then demonstrated how, in 2AFC tasks, the likelihood ratio favouring one perceptual outcome (or hypothesis) over the other can be computed by comparing the average spike rates from two such groups of sensory neurons with opposite stimulus

tuning preferences.

1.2. The drift-diffusion model of decision-making

This increased understanding of the neuronal representation of sensory evidence for perceptual judgements allowed progression to the next stage of perceptual decision-making: how is the evidence represented in sensory brain areas actually integrated to form the final decision? Information encoded in neuronal spike rates is momentary and variable, and must be accumulated over time in order to account for organisms' performance accuracy in perceptual tasks. Furthermore, to discriminate sequentially-presented stimuli, a sustained representation of sensory evidence is required (Gold & Shadlen, 2001). Shadlen, Newsome and colleagues opened a highly fruitful direction of research by showing that neuronal activity in the frontal eye fields (FEF) and lateral interparietal area (LIP) contained the signature of an integration process (Gold & Shadlen, 2000; Shadlen & Newsome, 2001; Roitman & Shadlen, 2002). Neurons in these association and motor regions had previously been reported to show sustained responses during delay tasks and to contain persistent activity related to the planning of eye movements (Colby & Goldberg, 1999; Tehovnik *et al.*, 1999). In macaques performing the motion task, Shadlen and colleagues showed that LIP and FEF neuronal spikes rates increased over time, with a gradient that depended upon stimulus motion strength, and the corresponding eye movement response was triggered when the neuronal activity reached a particular 'bound'.

This pattern of neuronal activity resembled a **drift-diffusion process**, first posited to explain choice accuracy and reaction times (RTs) in a memory retrieval task (Ratcliff, 1978), and since used to explain a wide range of simple decision-making tasks, including perceptual judgements. The drift-diffusion model describes perceptual decision-making as the integration, or accumulation, of sensory evidence over time, influencing the otherwise Brownian 'drift' of a decision variable (DV) from its starting position. Thus, the DV accumulates momentary evidence towards one choice or the other (Figure 1.4). A decision is made when the decision variable has drifted far enough to reach a particular bound[1]. The full drift-diffusion model contains parameters describing the relation between the stimulus strength and drift rate, the distances between the DV's starting point and the decision bounds, and the inter-trial variability of the drift rate

[1] In a **reaction time task**, the viewer can give their response at any time after the onset of the stimulus (e.g. Fig 1.3 a). The decision model explains the choice as corresponding to the decision bound that is crossed, and reaction times as the time taken for the decision variable to cross the bound. In a **fixed duration task**, the stimulus is presented for a fixed time period and the viewer is allowed to give their response only after a given temporal deadline (e.g. Fig 1.3 b). In this case, the diffusion model can explain choices as the first bound that was crossed or the bound that the decision variable is nearest to at time of decision deadline.

and the starting point (Ratcliff & Tuerlinckx, 2002; Vandekerckhove & Tuerlinckx, 2007). Thus, the model quantitatively predicts choices, mean RTs for correct and error responses, and RT distributions, as a function of stimulus strength (Ratcliff & McKoon, 2008).

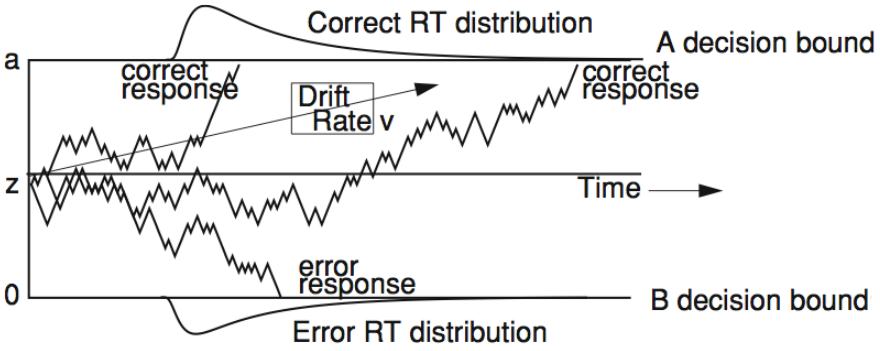


Figure 1.4.: Schema of the drift-diffusion process. Three possible paths of the decision variable, with starting point z and drift rate v, are shown; two reach the decision bound for the correct response (A). *RT* = reaction time. Taken with permission from Ratcliff & McKoon (2008).

Theoretically, the drift-diffusion model is a member of the general class of **sequential sampling models**, of which there are a number of different instantiations but which fall into two basic sub-classes for simple two-choice tasks. The first is the drift-diffusion process discussed above, where there is only one decision variable that accumulates evidence and for which positive evidence for one alternative is simultaneously negative evidence for the other (Ratcliff & McKoon (2008); Figure 1.4). The second is the sub-class of **race models**, or accumulator models, where evidence for each alternative is integrated separately rather than as a single value of relative evidence (Smith & Vickers, 1988; Reddi *et al.*, 2003). In race models, a response is initiated as soon as one of the accumulators reaches its decision bound. Recent race models allow the evidence in the accumulators to inhibit each other to a greater or lesser extent (Usher & McClelland, 2001). A race model with mutual inhibition seems to best explain the electrophysiological data from LIP neurons, each of which represent evidence accumulation towards the decision corresponding to the target in their individual response field (Mazurek *et al.*, 2003; Huk & Shadlen, 2005). Comparisons between the inhibitory race model and the single-variable drift-diffusion model show that their empirical predictions are very similar (Ratcliff & Smith, 2004; Ratcliff *et al.*, 2005). Under the simplifying assumption of perfect mutual inhibition in the race model, the models are equivalent (Palmer *et al.*, 2005). Some theoretical investigations have focused on whether there is a 'leak' in the evidence accumulation process such that over some time constant, accumulated evidence decays back to its neutral state (Usher & McClelland, 2001; Busemeyer & Townsend, 1993). However, empirical investigations have found evidence for only a minor leak

(Palmer *et al.*, 2005; Huk & Shadlen, 2005; Ratcliff & Smith, 2004). Therefore, notwithstanding the various adaptations that can be made to the sequential sampling models, the basic, non-leaky drift-diffusion model is an appropriate model to explain the main characteristics of decision-making under uncertainty and is the version employed in this thesis.

Competing hypotheses regarding the role of LIP and FEF activity, recorded from macques performing the motion task, include that it simply represents eye movement preparation (Hanes & Schall, 1996), a gain in noisy sensory evidence without integration (Cisek *et al.*, 2009), or a shift of attention towards the choice target (Gottlieb & Balan, 2010). However, based upon detailed analysis of the sources of variance of LIP neuronal spike rates, drift-diffusion accumulation still provides the best explanatory account (Churchland *et al.*, 2011). In humans, neuroimaging experiments have shown that, like in the macaque, drift-diffusion model parameters corresponding to stimulus strength and sensory evidence correlate best with brain activity in specific sensory regions sensitive to the stimulus features, whilst estimates of the decision variable itself and the distances to the decision bounds correlate with activity in parietal and prefrontal cortex (reviewed by Heekeren *et al.*, 2008). Taken together, the animal and human literature strongly argue that a drift-diffusion process underlies perceptual decision-making, where sensory evidence is represented in the activity of sensory brain areas and the integration of that evidence into a perceptual choice occurs in sensorimotor regions. To take the motion task as an example, sensory evidence about the motion direction of the random dot kinetogram is represented by firing rates in visual area V5/MT in macaque (or hMT+ in the human), whilst the integration of this evidence into a decision variable is represented by firing rates in areas LIP and FEF in macaque (or parietal and prefrontal regions in the human).

This conception of perceptual decision-making as a drift-diffusion process, with known neural correlates for its parameter values, forms a strong basis for further investigations. One outstanding question is how this model can help explain the incorporation of prior knowledge and bias into perceptual judgements. In particular, how might information about the size of reward associated with different perceptual outcomes be incorporated into the decisions?

1.3. Representations of reward in the brain

Reward is defined as a beneficial outcome, such as the reception of food or money, or the promise of this receipt in the future; a non-beneficial outcome can be conceived of as a negative reward. Information about reward affects behavioural decisions and learning, and there

are many brain areas implicated in reward circuitry (Herrnstein, 1961; Sugrue *et al.*, 2005; Rorie *et al.*, 2010; Salzman *et al.*, 2005; Rushworth *et al.*, 2009). In particular, strong signals for reward prediction and error are carried in the dopaminergic (DA) system, and representations of rewards associated with particular stimuli or actions have been investigated in the prefrontal cortex (PFC), anterior cingulate cortex (ACC), orbitofrontal cortex (ObFC), amygdala and, posterior parietal cortex (PPC).

Dopaminergic neurons in the midbrain project to to the striatum, a nucleus of the basal ganglia, and to the prefrontal cortex. In macaque monkeys, these neurons are activated by presentation of reward occurring when no reward was predicted, are depressed when a predicted reward is not presented, and do not respond to the presentation of expected rewards (Schultz, 1998). These neurons therefore appear to report rewards only when they occur differently than predicted; this 'prediction error' is likely to be an essential teaching signal for learning models, which can be used as priors for future interpretations of and responses to the environment, to maximize future reward (Schultz, 2001). Dopaminergic input to the basal ganglia appears to be necessary for efficient feed-forward cortical-striatal Hebbian learning, aiding information transfer through the basal ganglia and allowing it to play a key role in action selection (Bar-Gad & Bergman, 2001). Damage to the substantia nigra pars compacta (SNpc), which provides the dopaminergic input to the striatum, results in Parkinson's disease, characterized by poverty of movement and tremor at rest. Less is known about the role of dopaminergic projections from the ventral tegmental area (VTA) in the midbrain to to the limbic regions and prefrontal cortex, but one hypothesis suggests that DA modulates the signal-to-noise ratio in prefrontal cortex and the balance of activity between prefrontal and limbic areas (Callicott *et al.*, 2000, 2003). Human neuroimaging evidence also suggests that this DA system may be involved in the pathophysiology of schizophrenia and is a target of the dopaminergic neuroleptics used to treat the disease (Laruelle *et al.*, 1999; Abi-Dargham *et al.*, 2000; Winterer & Weinberger, 2004).

The representation of reward within the prefrontal cortex itself has also been a significant area of research, targeted by a variety of approaches that include human neuroimaging studies and animal lesion and electrophysiological studies (Rushworth *et al.*, 2009). The lateral orbitofrontal cortex represents the value of primary reinforcing and aversive stimuli, such as taste, touch, and facial expression, and via stimulus-reinforcer associations learns to associate these with more abstract stimuli, including monetary reward (Rolls & Grabenhorst, 2008; Morrison & Salzman, 2011). On the other hand, the ACC seems to be essential for learning the value of different actions (Kennerley *et al.*, 2006), whilst the ventromedial prefrontal cortex (VMPFC) plays a role

in the selection between different possible actions on the basis of their relative value (Padoa-Schioppa & Assad, 2008; Rushworth *et al.*, 2009). Within the limbic system, which supports a variety of functions including reinforcement learning, memory and emotions, single neuron activity in the amygdala is modulated by positive and negative values associated with visual stimuli (Paton *et al.*, 2006). Moreover, many amygdalar neurons respond differently to expected versus unexpected reward feedback, and this correlates with monkeys' learning behaviour (Belova *et al.*, 2007). This suggests a role for the amygdala in representing value and appropriately updating this representation according to external context and the internal state of the organism (Morrison & Salzman, 2010).

Finally, neurons in LIP of the posterior parietal cortex, which were shown above to represent the accumulation of evidence during perceptual decisions (Roitman & Shadlen, 2002), also represent information about rewards. In a seminal study using single cell electrophysiology and a forced-choice task, Platt & Glimcher (1999) showed that LIP activity is modulated in proportion to both the probability of reward and the relative amount of reward that a macaque monkey can expect to gain from an eye movement response. Moreover, when monkeys' choices in a foraging task provided a window into their subjective valuations, LIP neurons represented the relative, subjective value of competing actions (Sugrue *et al.*, 2004). These findings tie in with a prominent theory of posterior parietal cortex which suggests that it generates a relative *salience* map of different locations in external space, where salience comprises factors related to bottom-up stimulus energy as well as top-down aspects of task-relevance (Gottlieb, 2007).

Further work is necessary to show how different brain regions work together to combine various aspects of reward during learning and behaviour; nevertheless, our understanding of the neuronal representation of reward during learning, decision-making and behaviour is steadily increasing. 'Neuroeconomics' is now a flourishing field of research into the neural correlates of value-based decision-making under uncertainty (Sugrue *et al.*, 2005). The next step must be to combine these separate strands of research into the fields of reward and perceptual decision-making, in order to characterize the neuronal mechanisms by which the size reward available affects perceptual judgements and specifically to identify *where* in the decision process these effects might take place.

1.4. Reward and perceptual decisions

When we return to the drift-diffusion model of perceptual decision-making, the representation and accumulation of sensory evidence, the starting location of the decision variable (DV) and the distances to the decision bounds, all represent sites where reward might have effects (Figure 1.4). These parameters represent different stages of neuronal processing, such that the *drift rate* of the decision variable is dependent upon sensory evidence, represented in sensory brain regions, whilst the accumulation of this evidence over time, and the distance(s) between the DV starting point and the decision bound(s), are represented in later, sensorimotor brain regions. Does reward affect simply the accumulation of evidence and the motor response or does it affect sensory representations of the stimulus?

A few psychophysical studies with human participants have investigated how biasing choices towards one outcome or the other, by varying the size of rewards or costs available, affects the parameters of decision models. Diederich & Busemeyer (2006) and Diederich (2008) asked viewers to decide whether two white lines on a black background were the *same* or *different* in length, and biased choices by changing the amount of pay-off received for a correct or incorrect choice of either type. For example, to bias participants to respond *same*, the cost for erroneously responding *same* was made smaller than the cost for erroneously responding *different*. By imposing four levels of time limit on responses (ranging 450 to 900 ms after stimulus onset) and interpreting the results in the context of a modified drift-diffusion model, the authors report that the results are most consistent with a two-stage processing hypothesis, consisting of evaluation of pay-off attributes in the first stage, followed by unbiased evaluation of stimulus evidence in the second stage.

This model is similar to one in which reward affects the distance betwen the starting points of the accumulation process and decision bounds, and there was no evidence for an effect of reward on drift rate during stimulus processing (Diederich & Busemeyer, 2006; Diederich, 2008). In a similar study, Gao *et al.* (2011) asked participants to discriminate whether a light gray rectangle, presented on a dark gray background, was shifted 1-3 pixels to the left or right of a 300 pixel screen. An arrow at the start of each trial indicated which response direction would receive a larger amount of points if chosen correctly. Ten levels of time limit (ranging 0 to 2000 ms after stimulus onset) were imposed on responses. When a leaky, competing-accumulator model of decision-making was fit to the data, the results were consistent with an effect of reward to offset the higher-rewarded accumulator towards its decision bound (Gao *et al.*, 2011). This explained

the greater proportion of choices towards the higher-rewarded direction and how this changed according to response deadline without invoking effects on drift rate/sensory evidence. Liston & Stone (2008), however, found that when saccadic choices towards the brighter of two visual stimuli were biased by unequal reward probabilities, participants' responses on a separate, subsequent perceptual decision, about the brightness of the stimulus, were consistent with a sensory amplification of both the signal and noise associated with the higher-reward target. This appears to contradict the previous studies by suggesting an effect on sensory evidence itself but the interpretation of this study is limited because the reward was associated with the first, oculomotor task, and not the later, perceptual, task. Therefore the results do not speak directly to the effects of reward on perceptual decision-making when the reward bears upon the outcome of the perceptual decision itself. The only psychophysics study that reports a direct effect of reward on diffusion model drift rate is Voss *et al.* (2008), who fitted a drift-diffusion model to choices and reaction times of participants discriminating the colour of visual stimuli under different reward contingencies. This finding corroborates that of Liston & Stone (2008) and indicates that there could be an effect of reward on sensory evidence during perceptual judgements; if such an effect is small, it could be that the other psychophysics studies (Diederich & Busemeyer, 2006; Diederich, 2008; Gao *et al.*, 2011), which did not analyze shapes of reaction time distributions nor use standard drift-diffusion modeling, were not sensitive enough to detect it.

A few electrophysiological studies have investigated the neural correlates of perceptual judgements or stimulus representations under different reward conditions. Rorie *et al.* (2010) recorded directly from single neurons in area LIP whilst macaque monkeys made perceptual decisions in the motion task. At the beginning of each trial, a colour cue presented at the fixation point indicated the size of fluid reward available for a correct choice in each direction; reward conditions could either be *balanced* (equally large or small reward sizes for both directions) or *unbalanced* (one response direction had a larger reward size, if correctly chosen). Their computational analysis of LIP traces with a drift-diffusion model did not reveal any effect of reward condition on the drift rate of the decision variable for different reward sizes in either balanced or unbalanced reward conditions, only on the starting point of the accumulation process. However, in order to ensure their model fit the behavioural data, a surprisingly large value for inter-trial drift rate variability had to be set, suggesting that other features of the model could change if this parameter were adjusted (Rorie *et al.*, 2010). Other studies have shown effects of reward on rat primary visual and somatosensory cortex to stimuli associated with differing reward schedules, but did not link this to effects of reward on perceptual judgements themselves (Shuler & Bear, 2006; Pantoja *et al.*, 2007).

The effects of reward on sensory representations and perceptual decisions has been the focus of some human neuroimaging studies. Using fMRI, Summerfield & Koechlin (2010) and Mulder *et al.* (2012) investigated the effect of unequal payoff on human brain activity during perceptual decision-making, using fMRI and also analyzing behavioural choices and reaction times in the context of a drift-diffusion model. They report significant effects of reward only on the distance between the starting point and decision bounds in the model, which correlated with changes in brain activity in parietal and prefrontal regions, but no effects on the drift rate of the model nor on the sensory regions involved in the perceptual judgement. Fleming *et al.* (2010) report a role for the ACC in representing asymmetric pay-offs in a face-house discrimination task. An effect of reward on BOLD signal in the house-selective, but not face-selective, visual areas was also found; the reason for this disparate effect is unclear (Fleming *et al.*, 2010). Also using human fMRI, Serences (2008) and Serences & Saproo (2010) investigated visual cortical responses in tasks where participants select between two visual stimuli, where the reward probability associated with each stimulus was varied. They showed that selected stimuli were associated with a greater BOLD response in early visual areas than unselected stimuli, and that this difference was increased when the selected stimulus had higher value. With an elegant analysis that utilized voxel tuning functions (VTFs), Serences & Saproo (2010) suggest that high value sharpens the response profile across feature-selective neurons that encode the stimulus feature. However, these tasks did not measure perceptual decisions, as behavioural choice was related to reward history of previous choices, and not to the outcome of a visual discrimination task. These studies cannot show if and how the modulations of cortical BOLD signal due to reward might affect sensory evidence upon which perceptual judgements are based. Moreover, for these and other fMRI studies (Pleger *et al.*, 2008, 2009), the low temporal resolution of the BOLD signal means that we cannot be sure exactly at what stage of decision-making the reported brain activity change occurs and in particular, whether it occurs during the decision-making or after the choice due to feedback effects (Weil *et al.*, 2010).

Finally, the relationship between attention and reward must be considered because in attention tasks, subjects are motivated to direct attention to particular locations, features or objects through expectations about which location or feature is most salient (Maunsell, 2004). Many studies have found that directing attention towards a subpart of visual space or a specific stimulus feature enhances striate and extra-striate neuronal representations of the behaviourally-relevant spatial location or stimulus feature (for reviews see Corbetta & Shulman, 2002; Treue, 2003; Maunsell & Treue, 2006; Boynton, 2009). However, investigations into the effects of re-

ward on perceptual judgements that are considered in this thesis differ from these attentional paradigms. For example, in studies of spatial attention, attention is directed either wholly towards or wholly away from a particular region of space (Maunsell, 2004) or is cued towards a stimulus location that is *more likely* to be task-relevant (e.g. Cohen & Maunsell, 2011). In the investigations proposed in this thesis, spatial attention is always directed towards the stimulus to be discriminated, correct choices in the task are always rewarded and it is the size of available reward that is manipulated. Therefore spatial attention effects as measured in previous studies are not expected to play a role in the reward manipulations considered here. In feature attention paradigms, attention to particular features (such as stimulus colour, orientation or motion directon) is manipulated by changing which feature is 'task-relevant' i.e. needs to be responded to in order to obtain a reward, if correct (Maunsell & Treue, 2006). For example, to obtain reward the viewer might be required to report whether the colours of two stimuli match, or whether their orientations match; in both cases, the same stimuli are presented but the task relevance of their different features changes (McAdams & Maunsell, 2000). This may be closely related to a reward manipulation in a perceptual task in which one perceptual outcome is associated with higher reward if chosen correctly, since one possible effect of this manipulation could be that reward enhances the sensory representation of the higher-rewarded perceptual interpretation.

However, the perceptual judgement paradigms considered in this thesis differ significantly from feature attention studies, because the stimulus must be responded to correctly with either the higher or the lower-rewarded response to gain the reward, and in this way both perceptual outcomes are always task-relevant. Some studies have begun to differentiate the neural mechanisms of attention and reward; in a human EEG study, Baines *et al.* (2011) manipulated spatial attention by cueing the likely location of upcoming stimuli and manipulated 'motivation' (reward) by cueing whether the upcoming trial would account towards the final remuneration for task performance. They found differential effects of motivation and attention on the P1 and N1 brain potentials, which carry putative perceptual signals. Differentiating the effects of attention and reward is not tackled in this thesis, but when interpreting data the similarities of the paradigms must be borne in mind.

In summary, the vast majority of studies into the effects of reward on perceptual decisions have only reported effects on the starting point position or distance-to-bound of the evidence accumulation process, represented in sensorimotor brain regions. However, two psychophysics studies indicate that reward could affect sensory representations themselves, inferred from reward effects on the relevant parameters of fitted decision models (Voss *et al.*, 2008; Liston &

Stone, 2008) and some fMRI studies report a change in BOLD signal in visual areas according to reward associated with the stimulus (Serences, 2008; Serences & Saproo, 2010; Fleming *et al.*, 2010). Therefore, there is an indication that reward could affect sensory representations, but this has not been appropriately investigated as no study has directly measured the activity of sensory neurons with high temporal resolution nor selectively interfered with neurons' activity under different reward conditions.

1.5. MEG and the human visual system

Investigation into the neural correlates of perceptual decision-making in humans, and how they may be affected by contextual factors such as reward, has been limited by the lack of a human brain recording method that has both good spatial and good temporal resolution. As discussed above, conclusions drawn about the human brain activity from fMRI studies are limited due to the low temporal resolution of the measured BOLD signal relative to the spiking activity of neurons (Menon *et al.*, 1998; Liao *et al.*, 2002). Furthermore, the BOLD signal is an indirect indication of local neuronal activity levels and the coupling between neuronal activity and the BOLD signal in nearby blood vessels is still not fully understood (Logothetis *et al.*, 2001; Logothetis & Wandell, 2004). On the other hand, magnetoencephalography (MEG) records magnetic signals produced by electrical currents in neurons from some brain sites in the human (Chapter 5). The temporal resolution of MEG is high, on the order of 1 ms, and MEG has a superior spatial sample profile than EEG (Hansen *et al.*, 2010).

If reward affects sensory processing in a specific way during perceptual judgements, a reward signal should be found in the subparts of the sensory areas that process the stimulus upon which the judgement is based, and this must be differentiable from effects at other parts of the visual system and the decision or motor output stages. MEG has the potential to record neuronal activity with sufficient specificity in space and time to investigate this question. However, MEG analysis methods are still in development and the accurate localization of the MEG signal to sources in the brain is a non-trivial problem (Hansen *et al.*, 2010). There are many different approaches to source reconstruction: (single) dipole modeling (e.g. Chen *et al.*, 2005), minimum norm estimates (MNE) (Hamalainen & Ilmoiemi, 1994), beamformning (van Veen *et al.*, 1997) and multiple sparse priors (Friston *et al.*, 2008), to name but a few. These all work with differing assumptions about the pattern and distribution of neuronal sources in the brain. It is not yet clear which source reconstruction method is optimal for investigating and localizing the activity of the visual brain areas that process the sensory information.

1.6. Aims and scope of this thesis

This thesis aims to investigate the mechanisms that underlie the effects of reward size on visual perceptual decisions, specifically by exploring the possibility that reward affects the representation and processing of visual information in sensory brain regions. I probe this question using neurophysiological methods in humans and monkeys, combined with computational modeling approaches. The perceptual decision task used in this thesis is the **cylinder task**, in which the viewer must discriminate the direction of rotation of a cylinder stimulus defined by structure-from-motion (SFM) with varying amounts of binocular disparity (depth) information. The disparity of the stimulus determines the difficulty of the discrimination in any given trial (for details, see Figure 2.1 & Chapter 2). In the macaque, visual area V5/MT has been shown to represent perceptual information that underlies judgements in this task (Dodd *et al.*, 2001; Krug *et al.*, 2004). The cylinder task was chosen because previous studies have given us an excellent understanding of the representation of sensory evidence about these SFM stimuli in the macaque and human brains (Dodd *et al.*, 2001; Krug *et al.*, 2004; Brouwer & van Ee, 2007) and because this stimulus set contains strong Gestalt or form-generating properties since at zero disparity the SFM cylinder still gives a strong, albiet multistable, percept of a 3D object rotating either to the left or to the right. Furthermore, this task has not previously been analyzed in the context of a drift-diffusion model of decision-making such that the investigations here allow a further test of the generalizability of the model to a new stimulus set.

In Chapter 2, electrical microstimulation was applied to motion- and disparity-selective clusters of neurons in visual area V5/MT, to shift perceptual judgements of macaque monkeys performing the cylinder task. The effect of microstimulation on perceptual decisions was compared under different reward conditions. The microstimulation intervention dissects the decision-making pathway at the level of the sensory representations in V5/MT, allowing us to infer what effects of reward occurred up to this stage, affecting stimulus-dependent sensory evidence alone, or after this stage, affecting stimulus-dependent and microstimulation-dependent evidence together. In Chapter 3, a drift-diffusion model, adapted to the cylinder task with microstimulation and reward, was applied to the resultant behavioural data. The results of the microstimulation study and model-fitting provide support for the hypothesis that reward condition affects visual processing at the level of sensory regions such as area V5/MT or earlier.

The study was extended to a human psychophysics study, in which participants performed the cylinder task whilst their decisions were biased by unbalanced rewards for the possible per-

ceptual choices (Chapter 4). Full reaction time distributions, as well as behavioural choices, were collected for nine levels of stimulus difficulty. A full drift-diffusion model was fit to the behavioural and reaction time data to investigate the effects of reward on the decision process (Vandekerkhove & Tuerlinckx, 2007b). It was found that decisions in the cylinder task are as well-described by the diffusion model, and the best-fitting models contain effects of reward on the drift rate of the diffusion process, as well as the distance-to-bound. Since drift rate is dependent upon sensory evidence, represented in early visual areas, this finding agrees well with the hypothesis that reward affects sensory processing.

The second half of the thesis focuses on using magnetoencephalography (MEG) to investigate responses in the human visual system. Chapter 5 contains information on the background of MEG and the MEG methods used in this thesis. The potential of this technique to be applied to questions of human perceptual decision-making motivated a magnetoencephalographic retinotopy mapping study, the focus of Chapter 6. This study evaluated and compared three contrasting MEG source analysis methods against visual source locations recorded with fMRI in the same individuals. This allowed me to choose analysis methods that can localize visual activity correctly to subregions of early visual areas in the human brain. A further MEG study recorded brain activity from human volunteers as they performed the cylinder task under different reward conditions (Chapter 7). The results of this MEG investigation implicate a range of brain areas across sensory and association regions, including early visual areas and visual precuneus, as being affected by reward in this perceptual task.

2. Reward and electrical microstimulation in V5/MT: Experimental data

2.1. Introduction

Since the origin of electrophysiological experiments, electrical microstimulation of the brain has been an important method of investigation of sensory and motor systems in animals and humans, first used to delineate topographic maps in somatosensory and motor cortices and oculomotor structures (Fritsch & Hitzig, 1870; Grunbaum & Sherrington, 1902; Penfield & Boldrey, 1937; Robinson & Fuchs, 1969; Robinson, 1972). In general, experimental interference with the activity of a particular brain area, especially through electrical or chemical activation of the neurons therein whilst observing the resultant effect on the subject's behaviour or perceptual report, is taken to be strong evidence towards a causal role of that brain region in generating that behaviour or report (Parker & Newsome, 1998).

In the past two decades, elegant experiments have established that electrical microstimulation of small, selective groups of cortical neurons, in combination with sensory stimulation, can produce a reliable change in perceptual reports. The change in these reports can be predicted from the selectivity of the microstimulated neurons. For instance, microstimulation of direction-selective neuronal clusters in visual area V5/MT in macaque monkey produced a measurable shift in monkeys' perceptual choices, towards the direction-selectivity of the stimulated neurons, in a visual motion-discrimination task in which the perceptual report was indicated by saccadic eyemovements (Salzman *et al.*, 1990, 1992). As such, the animals' behaviour suggests that visual information arising from both sensory stimulation and electrical stimulation is treated in the same way and combined to produce the final perceptual report. More recently, using the same task, Hanks et al. (2006) showed that microstimulation of sensorimotor neurons in lateral intraparietal area (LIP), which represent the direction of upcoming eye-movements, speeds up the reaction time for eye movements towards the target in the stimulated neurons'

response field and slows the reaction time for eye movements towards targets outside the response field. Similarly, work in somatosensory cortex of the mouse has shown that activation of even a single sensory neuron can produce a significant behavioural response (Houweling & Brecht, 2008). Overall, electrical microstimulation is an important and highly effective experimental technique for probing the roles of specific neuronal populations in the perceptual perceptual decisions of awake, behaving animals.

In this study, I use electrical microstimulation of a visual area to investigate whether reward affects the incorporation of artificial sensory information into perceptual decisions. Two macaque monkeys were trained to report the direction of rotation of structure-from-motion (SFM) cylinder stimuli (Figure 2.1). The size of the fluid reward received for a correct choice in each trial was variable, depending upon the outcome of the previous trial. Although it has been been shown that reward can affect decisions at the stages of sensorimotor integration (Platt & Glimcher, 1999; Rorie *et al.*, 2010; Summerfield & Koechlin, 2010; Mulder *et al.*, 2012), effects at the earlier level of sensory representations are controversial. In this study, I show that the effect of electrical microstimulation in sensory area V5/MT on perceptual reports is influenced by the size of the reward available on a specific trial; thus, reward differentially affected the incorporation of sensory stimulation and electrical microstimulation on perceptual choices. The results agree with the hypothesis that reward can affect visual processing during perceptual decision-making.

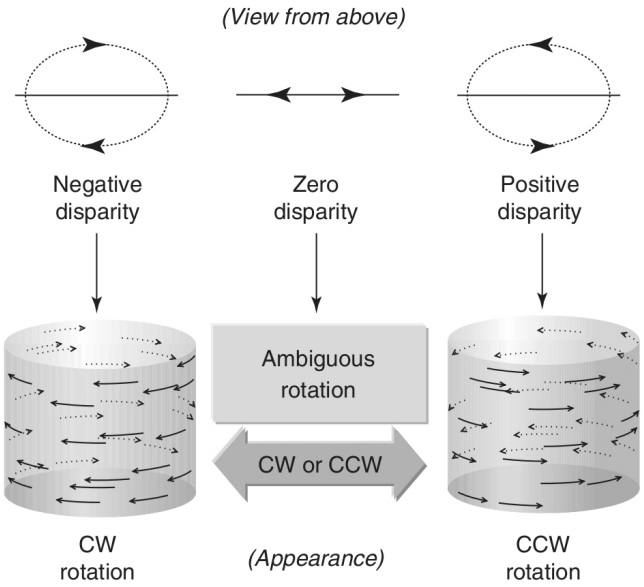


Figure 2.1.: Diagrammatic representation of the cylinder stimuli (taken with permission from Parker & Krug, 2003). At zero disparity (middle), dots moving in opposite directions are at the same visual depth and the stimulus is bistable such that the perceived direction of rotation is either clockwise (CW) or counterclockwise (CCW), as viewed from above. When dots are separated in depth by the application of different binocular disparities to the dots moving in opposite directions, cylinder rotation becomes unambiguous. For example, when the leftward moving dots are at a near disparity and the rightward moving dots are at a far disparity, the cylinder is seen to rotate in a CW direction (left example).

2.2. Methods

The experiments described in this chapter were carried out in collaboration with Dr. Bruce Cumming, Dr. Andrew Parker and Dr. Kristine Krug. Dr. Cumming, Dr. Parker and Dr. Krug performed the training and collected electrophysiological and microstimulation data for monkey Fle. Dr. Cumming performed the training for monkey Ica, and Dr. Cumming, Dr. Krug and myself collected the electrophysiological and microstimulation data for monkey Ica. Dr. Cumming collected the pseudo-stimulation data for monkey Ica. I analyzed all the electrophysiological and microstimulation data presented in this chapter.

2.2.1. Animals

The microstimulation experiments were conducted in a single cortical hemisphere in each of two awake behaving male rhesus macaque monkeys (*Macaca mulatta*). Prior to the experiments, each monkey was surgically implanted under general anaesthesia with a head-holding device and a stainless steel recording chamber placed over a craniotomy above the occipital lobe. In monkey Fle, scleral magnetic search coils were implanted in both eyes to monitor eye position and vergence. In monkey Ica, a scleral magnetic search coil was implanted in one eye, to monitor eye position. Monkeys were trained to fixate on a binocularly presented marker and to perform a visual discrimination task that required the animals to respond to the motion and binocular disparity of a rotating structure-from-motion (SFM) cylinder stimulus (Figure 2.1).

2.2.1.1. Locations and ethics statements

Experiments were conducted at two locations: monkey Fle was tested at the Laboratory of Physiology at Oxford University, UK; monkey Ica was tested at the National Eye Institute, National Institutes of Health, Bethesda, Maryland, USA. These locations are henceforth referred to as Oxford and NIH, respectively. Except where specifically described in this section, the protocols were identical in the two locations.

At Oxford, all of the procedures complied with the United Kingdom Home Office regulations on animal experimentation. At NIH, all of the procedures complied with the US Public Health Service policy on the humane care and use of laboratory animals and all protocols were approved by the National Eye Institute Animal Care and Use Committee.

2.2.2. Visual stimuli

The visual stimulus for perceptual discrimination was a rotating cylinder, rendered in structure-from-motion (SFM) by two transparent planes of random dots moving in opposite directions (random dot kinematogram, Figure 2.1) (Treue *et al.*, 1991). The dots had a sinusoidal velocity profile as would be expected in a two-dimensional projection of a three-dimensional rotating cylinder. When the two dot planes are presented at the same binocular disparity, the cylinder rotation direction is bistable and over time will undergo spontaneous fluctuations in the perceived direction of rotation (Wallach & O'Connell, 1953; Ullman, 1979). Adding horizontal disparity separates the front and back surfaces of the cylinder in depth, and thereby disambiguates the direction of rotation. The disparity of each dot was determined by its position on the depicted cylinder by being scaled sinusoidally to match the velocity profile, so that the maximum disparity difference between the two dot planes occurred at the midline of the front and back surfaces of the cylinder. This maximum disparity was called the stimulus disparity. Disparity was always added equally to the front and back surfaces of the cylinder, in effect expanding the cylinder in depth around its principal axis, which remained in the plane of fixation (Figure 2.1). In this chapter and the following chapter, the convention of the stimulus program and data analysis is that a *negative* disparity refers to a *clockwise* (CW) (front surface moving left) direction of rotation of the cylinder, and a *positive* disparity refers to *counterclockwise* (CCW) rotation (front surface moving right). Cylinder stimuli consisted of 50% black dots and 50% white dots presented at maximum contrast (99%) on a mid-gray background. Dot locations were randomly selected on the cylinder surface and dot density was 25%. Half of the dots (randomly selected) moved in one direction and the rest of the dots moved in the opposite direction. When a dot reached the edge of the cylinder, it reversed direction.

In Oxford, stimuli were displayed binocularly using a Wheatstone stereoscope configuration (Wheatstone, 1838). Two monitors (Eizo FlexScan 78) displayed images, with mirrors positioned in front of the monkey each reflecting the image from one screen into one eye so that the images could be combined into a single stereogram. Frame rate was 72 Hz. Monitor screens were positioned 84 cm from the monkey, and covered approximately 21x17° of the visual field. Mean luminance was 42 cd/m^2 and dot diameter was 0.2. The dot refresh rate was 2%, i.e. on each video frame, 2% of the dots were relocated to a randomly chosen location on the cylinder surface, so that dots had a mean lifetime of 615 ms in each location.

At NIH, stimuli were displayed binocularly using two DLP projectors (Projection Design Ev02sx+)

with polarizing filters. The image was projected onto a polarization-preserving screen (Steward Filmscreen 150). Frame rate was 60 Hz. The screen was 112 cm away from the monkey, and covered approximately 41x32° of the visual field. Mean luminance was 54 cd/m^2 and dot diameter was 0.18°. The dot refresh rate was 1%. The size, orientation and speed of the cylinder stimulus were matched to the receptive field preference of the neurons at the microstimulation site (described below).

2.2.3. Task

The psychophysical task was a two-alternative forced choice (2AFC) discrimination of the direction of rotation of the SFM cylinder. The cylinder stimulus disparities used were matched to the psychophysical threshold of the animal in order to obtain the full psychometric function: discrimination was easiest at the largest disparity values where performance was near 100% correct. At smaller disparities, discrimination was more difficult and performance was usually between 30-60% correct. At zero disparity, the rotation direction was ambiguous and animals' behavioural choice was around 50% for each of the two possible directions. On these trials, the 'correct' answer was chosen by the computer program to reward half of the trials at random.

The animal first acquired the fixation point to initiate the trial. The visual stimulus was then presented for 2 seconds (Figure 2.2). At stimulus offset, the fixation point also disappeared and two choice targets corresponding to the two possible directions of cylinder rotation appeared. In the Oxford set-up, the targets were located on either side of the fixation point. In the NIH set-up, the targets were located on either side of where the cylinder stimulus had been, normal to the axis of cylinder orientation. The animals indicated their perceptual decision by making a saccadic eye-movement to one of the two targets. If a saccade was made to the correct target, a fluid reward was received. If the choice was incorrect, the animal received no reward and there was a short time-out before the start of the next trial. If the monkey broke fixation at any time during stimulus presentation in any trial, the trial was terminated and no reward was given.

2.2.3.1. Reward Schedule

Animals worked on the task to gain fluid rewards to meet their daily requirements for intake of fluid. Therefore the incentive for task performance was a fluid reward upon a correct response. The reward size available for a correct choice on each trial depended upon the number of previous consecutive correct responses that the animal had made immediately prior to the

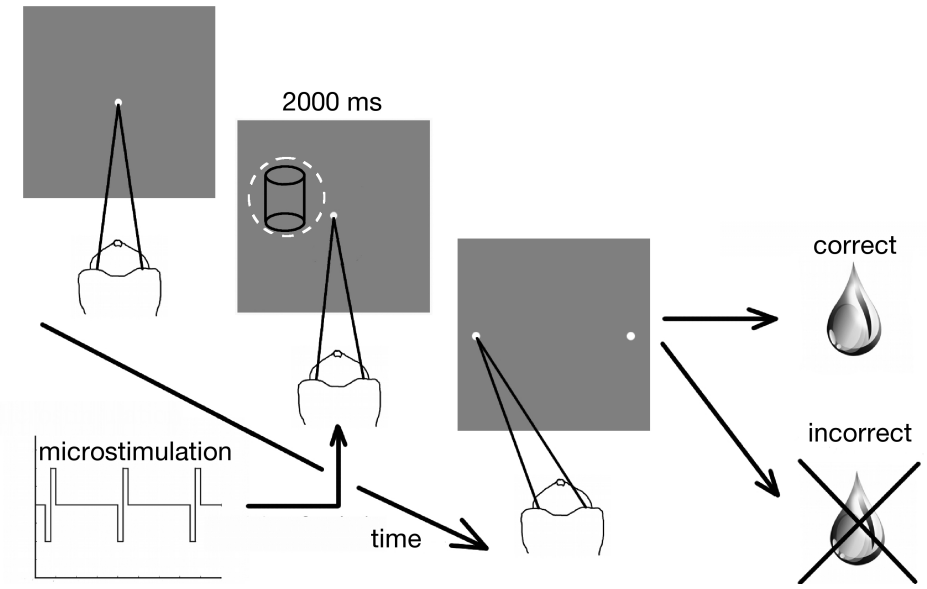


Figure 2.2.: Trial sequence; each trial began with the acquisition of the fixation point. The visual stimulus was then presented for 2 s in the receptive field (RF) of the V5/MT site (RF boundary is indicated by white dashed line). Upon the offset of the stimulus, two choice targets appeared. The monkey indicated its perceptual decision with an eye movement to the corresponding choice target. If the response was correct, the animal received a fluid reward. If incorrect, it received no reward and a time-out. An electrical microstimulating current was introduced into the V5/MT site in 50% of trials. Microstimulation trials were pseudo-randomly interleaved with the non-stimulated trials.

trial, and increased in two steps up to a maximum. For monkey Ica, reward size was 0.08 ml on the first and second consecutive correct choices after an error, 0.12 ml for the third consecutive correct choice, and 0.2 ml on the fourth and all subsequent consecutive correct trials (Figure 2.3). For monkey Fle, reward size was 1/3 of maximum for the first correct choice, 2/3 of maximum for the second, and reached maximum size for the third and all subsequent consecutive correct choices; thus reward size increased more quickly for subject Fle than subject Ica as a function of consecutive correct responses. The maximal reward size for monkey Fle varied site-by-site. I categorized the trials into two conditions: maximal (large) available reward size and sub-maximal (small and medium) available reward size, where for both animals the average size of the two sub-maximal reward sizes was half that of the maximal reward size. This reward schedule was used throughout training and all recording sessions, it was not introduced just for these microstimulation experiments.

2.2.4. V5/MT multiunit recording

We recorded from multi-unit (MU) neuronal clusters to characterize and select the cortical microstimulation sites. Tungsten micro-electrodes coated with parylene with an exposed tip

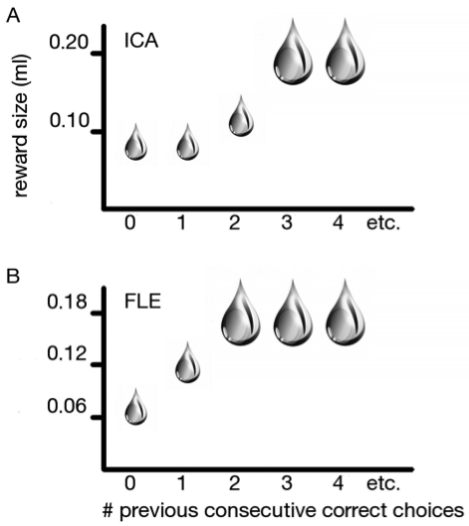


Figure 2.3.: Reward schedules for monkeys *Ica* (A) and *Fle* (B). The size of fluid reward available for a correct choice in each trial increased in two steps up to a maximum, depending upon the number of previous consecutive correct responses that the animal had made.

length of 10 mm were used (MicroProbe, Inc., USA; 0.1-0.3 M impedance at 1 kHz). A hydraulic microdrive was mounted on the recording cylinder at the beginning of each recording session. The electrode was advanced through a guide tube inserted through the dura. The guide tube was supported by a 2D grid with 1mm spacing to allow the systematic mapping of recording sites. Electrical signals from the electrode were bandpass filtered (0.2 - 15 kHz), amplified and displayed on a computer monitor and audio monitor, and stored to computer disk. For the signal recording, online classification of activity and the pre-processing steps, Datawave Discovery system (DataWave Technologies, Minneapolis, MN) or Spike 2 (Cambridge Electronic Design Ltd.) data acquisition program were used at Oxford and NIH respectively.

Area V5/MT was approached posteriorly through the recording chamber in incremental steps of 50 μm. Area V5/MT was identified by established physiological criteria: 1) the gray/white matter pattern comprising the striate cortex, lunate sulcus and the posterior bank of the superior temporal sulcus (STS) (Zeki, 1974); 2) the response profiles of the neurons encountered were identified as from V5/MT by their characteristic direction selectivity and binocular disparity selectivity and its clustering (van Essen *et al.*, 1981; Maunsell & van Essen, 1983; Albright & Desimone, 1987; Bradley *et al.*, 1995, 1998; DeAngelis *et al.*, 1998; DeAngelis & Newsome, 1999); 3) the known relationship between the size and eccentricity of their receptive fields (Desimone & Ungerleider, 1986); and 4) from penetration to penetration, the relationship between receptive field positions and known topography of V5/MT (Albright & Desimone, 1987).

2.2.5. V5/MT MU receptive field mapping

Monkeys maintained fixation on a central fixation point whilst passively viewing the visual stimuli used for MU receptive field mapping. Motion direction selectivity was established for an MU cluster at a particular electrode location by presenting a quick succession of kinetic dots moving coherently in different directions on the screen. The average MU firing rate in response to each direction of motion was calculated online. The direction evoking the greatest response on average was assigned to be the preferred motion direction for the cluster. The receptive field boundaries were then quantitatively mapped using a patch of dots moving in the preferred direction. The binocular disparity preference of the MU cluster was measured by presenting a plane of moving dots in the preferred direction at different binocular disparities around the fixation plane. Only sites with discernible motion and disparity tuning were further explored. By aligning the axis of rotation of the cylinder stimulus orthogonally to the preferred direction of the MU cluster, the directions of motion of the two transparent planes of dots that made up the cylinder were aligned with the preferred and the opposite directions of the MU cluster. The size of the cylinder was matched to the MU receptive field. Cylinder stimuli were presented at a range of disparities and the MU clusters average firing rate to each cylinder disparity was calculated online to measure the cylinder rotation direction preference for the cluster, termed the 'PREF' rotation direction. The opposite direction was termed 'NULL'.

2.2.6. Selection and characterization of stimulation sites

If and when the direction- and cylinder disparity-selectivity of MU cluster activity were found to remain constant over 300-400 μm of cortex, the electrode was slowly and carefully retracted 150-200 μm to a position in the middle of that stretch, for electrical microstimulation. At the proposed stimulation site, the MU receptive field size, direction-selectivity and cylinder preference were quantitatively re-confirmed. The cylinder stimulus position, orientation and size properties were carefully matched to the response selectivity of the MU at the stimulation site.

2.2.7. Microstimulation experimental procedure

Once a potential microstimulation site had been characterized, the animal performed blocks of trials of the psychophysical task during which microstimulation was not applied. The purpose of this block of trials was to allow the selection of a range of stimulus disparities appropriately matched to the animals threshold (see below). Once the animal's behavioural performance was

stable, the first block of electrical microstimulation experiment trials was commenced. Electrical microstimulation was applied through the same electrode as used for MU recording at the stimulation site. Half of the trials were microstimulated, and these microstimulation trials were pseudo-randomly interleaved with non-stimulation trials in each experimental block. Microstimulation was applied only during the 2 s of stimulus presentation in the stimulated trials. The stimulating pulses were 20 μA biphasic pulses of 500 μs cathodal stimulation followed by 500 μs anodal stimulation, and were delivered at 200 Hz. At least 10 microstimulation trials and 10 non-stimulation trials at each of at least 5 different stimulus disparities were necessary for the data from a given microstimulation site to be included in subsequent analyses.

2.2.8. Pseudo-stimulation: psychophysical reproduction of microstimulation effect

A control psychophysics experiment was run to evaluate whether monkeys were able to counteract the effect of an additional, microstimulation-like cylinder disparity signal, when this additional signal was overtly cued in those trials in which it occurred. This control experiment was performed on monkey Ica, at NIH, USA. Methods were identical to the real microstimulation experiment, except that no electrical microstimulation of the nervous system took place. Instead, a fake stimulation signal of +0.005° disparity was inserted additively to the cylinder disparity during stimulus presentation, in a pseudo-randomly selected half of the trials. The monkey's perceptual reports were rewarded according to the stimulus disparity as assigned by the experimental condition without the addition of the extra, pseudo-stimulation signal.

In the electrical microstimulation experiment, the cylinder stimulus consisted of dots (50% white dots, 50% black dots) and there was no change to the cylinder stimulus between microstimulated and non-stimulated trials. Importantly, in the pseudo-stimulation psychophysical experiment, the trials containing the fake stimulation signal were overtly cued by a change in the colour of the cylinder stimulus, from 100% white dots to 100% black dots. The reward schedules were kept the same as in the real microstimulation experiment. Therefore, for maximal reward accumulation, the animals optimal strategy would be to bias perceptual choices away from the direction of the fake signal on those trials that were overtly cued as being 'stimulated'.

Over 5 days, 15 experimental blocks of fake stimulation psychophysical data were collected, with all experimental parameters remaining constant. There were approximately 1000-2000 trials in each experimental block. By contrast, there were on average 500 trials in each experimen-

tal block at the electrical microstimulation sites.

2.2.9. Data Analysis

2.2.9.1. Quantification of microstimulation shift effect

To evaluate the psychophysical performance in the cylinder task and the effect of microstimulation on perceptual decision-making at each stimulation site, I calculated the proportion of choices made by the monkey towards the PREF direction for each cylinder disparity. This was done separately for microstimulated and non-microstimulated trials. To generate the psychometric functions (PMFs), these data were then fitted with a cumulative Gaussian distribution using maximum likelihood estimation (MLE), implemented with the *fminsearch* function in MATLAB 2009a (MathWorks). The cumulative distribution function for the Gaussian distribution is

$$P_{PREF} = \frac{1}{2}\left(1 + \text{erf}\left[\frac{x-\mu}{\sigma\sqrt{2}}\right]\right) \tag{2.1}$$

where μ is the mean of the distribution, σ is the standard deviation (SD), erf is the error function, and P_{PREF} corresponds to the probability of making a decision in the PREF direction of the microstimulated MU neurons. To measure the effect of microstimulation on behavioural performance at each site, the psychometric curves for microstimulated and non-microtimulated trials were fitted with a pair of cumulative Gaussian functions that had the same SD but whose means were allowed to vary. This is a reasonable procedure based upon earlier evidence that microstimulation in this brain region does not substantially affect the slope of the psychometric curve (Salzman *et al.*, 1992).

For clarity in observing the direction of the microstimulation effect, in this chapter all psychometric functions are plotted with the y-axis representing the proportion of perceptual reports in the MU PREF direction. The absolute microstimulation effect can be measured as the difference in the means of the psychometric function for microstimulated trials compared to non-stimulated trials; this is equivalent to a horizontal shift in PMF between the two conditions. The shift effect is in the units of binocular disparity (DeAngelis *et al.*, 1998; Uka & DeAngelis, 2006). Therefore, the size of the electrical microstimulation effect can be directly compared to the binocular disparity that would need to be added to the stimulus to shift perceptual behaviour to the same degree.

To evaluate whether electrical microstimulation significantly changed behavioural report at each stimulation site, the behavioural data were fit with two competing cumulative Gaussian models: 1) the two-mean model, in which microstimulated and non-stimulated PMFs were fitted with separate models that had the same SD but different means, and 2) the null model: all the data points, from both microstimulation and non-stimulation trials, were fitted together with a single cumulative Gaussian. The best-fit parameters for both models to the data were fitted by maximum likelihood estimation (MLE). The goodness-of-fit, for each model, is thus given by the negative log likelihood that the model parameters explain the observed data points:

$$log_e(L) = log_e\Big(L(p|n,y)\Big) = log_e\binom{n}{r} + y \cdot log_e(p) + (n-y) \cdot log_e(1-p) \tag{2.2}$$

As model parameters are equivalent to degrees of freedom, the two-mean model, which has more parameters (2 means, 1 SD), will always fit at least as well as the null model, which has 1 fewer parameters (1 mean, 1 SD). The model with fewer parameters can be denoted the restricted model (H_0) whilst the model with additional parameters can be denoted the full model (H_1). Whether the model with more parameters fits significantly better must therefore be determined by the significance of the difference in their likelihoods. The likelihood-ratio test statistic D takes a value that is twice the difference between the log likelihoods of the models:

$$D = 2(log_e(\text{likelihood } H_1) - log_e(\text{likelihood } H_0)) \tag{2.3}$$

The probability distribution of the test statistic D can be approximated by a χ^2 distribution, where the degrees of freedom (d.f.) are the difference between the number of parameters of the two models being tested. Thus if the p value of the difference in goodness of fits of the two models is less than 0.05, a significant microstimulation effect is indicated, because a model with two different means for microstimulated and non-stimulated psychometric curves describes the data significantly better than a model with a single mean. A site where microstimulation had a significant effect on the psychometric response curve is referred to as a significant ('sig') site.

2.2.9.2. Quantification of effect of expected reward size on microstimulation

At each microstimulation site, the trials were separated into two groups according to whether the available reward size for a correct choice was large or small. The size of the horizontal shift of the psychometric function induced by electrical microstimulation was calculated separately for each reward size group. To compare the effect of microstimulation between the stimulation sites, for which different stimulus disparity ranges were used (see section Psychophysical Task above), the horizontal shifts of the PMFs due to microstimulation were divided by the standard deviation of the fitted cumulative Gaussian at each site (Uka & DeAngelis, 2006). The resultant normalized shift is unitless. Since the underlying distribution of shift values was not normal, non-parametric Wilcoxon sign-rank tests were used to ascertain whether there was an overall significant difference in the size of the shift effect between the two reward conditions across sites.

2.2.9.3. Normalization and pooling of data across multiple experimental sessions

For model-fitting analysis, data were pooled over multiple microstimulation sites to increase the number of samples in each condition, separately for each monkey. To allow pooling, the stimulus disparity values for each site were normalized by division by the maximum disparity value in the range used at that site. As a result, the maximum normalized disparity was equal to 1 (unitless), the minimum was -1, and the other stimulus disparity values took relative values between 1 and -1. Data were pooled by adding together the total number of trials and the number of each perceptual report at each normalized stimulus disparity value across the multiple sessions, separately for each monkey.

2.2.9.4. Quantification of interaction between reward and microstimulation with logistic regression

A logistic regression model was used to evaluate whether the addition of an extra model parameter (β_7 in Equation 2.4), which allowed the size of the microstimulation shift effect to change with reward condition, significantly improved the fit of the model to the data. The equation for the full logistic regression model is:

$$P_{PREF}(x) = \frac{1}{1 + e^{-(\beta_1 + \beta_2 \cdot C + \beta_3 \cdot S + \beta_4 \cdot C \cdot S + \beta_5 \cdot R_{val} + \beta_6 \cdot C \cdot R_{val} + \beta_7 \cdot S \cdot R_{val} + \beta_8 \cdot C \cdot S \cdot R_{val})}} \tag{2.4}$$

where C is the normalized stimulus disparity value, R_{val} = 1 for large reward trials and 0 for small reward trials, S = 1 for microstimulated (or pseudo-stimulated) trials and 0 for non-stimulated trials, β_1 represents the y-intercept (refereed to as the *null bias*, Salzman *et al.*, 1992), β_2 represents the effect of cylinder stimulus disparity, β_3 represents the effect of microstimulation, β_4 represents the interaction between cylinder stimulus disparity and microstimulation, β_5 represents the effect of reward on the null bias, β_6 describes the interaction between cylinder stimulus disparity and reward, β_7 describes the interaction between microstimulation and reward, β_8 describes the interaction between cylinder stimulus disparity, microstimulation and reward. The corresponding equation for the restricted logistic regression model that does not include the reward-microstimulation interaction parameter β_7, is:

$$P_{PREF}(x) = \frac{1}{1 + e^{-(\beta_1+\beta_2 \cdot C+\beta_3 \cdot S+\beta_4 \cdot C \cdot S+\beta_5 \cdot R_{val}+\beta_6 \cdot C \cdot R_{val}+\beta_8 \cdot C \cdot S \cdot R_{val})}} \tag{2.5}$$

The significance of interaction between microstimulation and reward was evaulated using a χ^2 test (see above) between the restricted model (Equation 2.5) and full model (Equation 2.4).

2.3. Results

Electrical microstimulation and changing reward contingencies were used to investigate how reward size affects perceptual decision-making. Animals were trained to perform a two-alternative forced choice task requiring the discrimination of the direction of rotation of a transparent, rotating cylinder, defined by structure-from-motion in a random dot kinetogram (Figure 2.1). When both dot surfaces are at zero disparity, the depth order of such a stimulus is ambiguous and the percept is bistable. The cylinder stimulus was disambiguated by addition of binocular disparities to the dots to define the depth order (Figure 2.1). In each trial, animals viewed the cylinder stimulus for 2 seconds and at stimulus offset, indicated their decision about the direction of cylinder rotation with a saccadic eye movement to one of two choice targets (Figure 2.2). Multi-unit (MU) clusters of neurons in V5/MT, which were consistently selective for a particular direction of rotation over a stretch of at least 200-300 μm of cortex, were electrically microstimulated during stimulus presentation in half of the trials.

2.3.1. Microstimulation of visual area V5/MT shifts perceptual decisions

Figure 2.4A shows the cylinder disparity tuning curve for the MU in visual area V5/MT from monkey Ica (site ICA266). Neurons at this site were selective for, or 'preferred', a negative cylinder disparity (CW rotation). Microstimulation was applied to the MU cluster on half of the trials pseudo-randomly, during visual stimulus presentation. The proportion of the subject's perceptual reports in the MU's preferred (PREF) direction was plotted as a function of the cylinder disparity (Figure 2.4B). Psychometric functions were fit with a cumulative Gaussian, for which the mean (but not the standard deviation) was allowed to change according to whether electrical microstimulation was applied. If the activitation of the V5/MT MU cluster contributes to the perceptual decision, we would expect that an increased proportion of the monkey's perceptual reports would be in the PREF direction in microstimulated trials, producing a rightwards horizontal shift in the psychometric function (PMF). The shift of the PMF due to electrical microstimulation was in fact highly significant at this site ($p << 0.001$, χ^2-test). The size of the shift was equivalent to adding -0.01° to the cylinder disparity (Figure 2.4B).

We microstimulated 48 such sites in two monkeys. In a majority of sites (27/48 sites), microstimulation induced a significant shift of the psychometric function towards the stimulus tuning preference of the stimulated site ($p < 0.05$, χ^2-test). By contrast, only 3/48 sites showed significant shifts in the opposite direction. This effect of microstimulation was also present when each monkey was considered separately: 12/20 sites in monkey Fle and 15/28 sites in monkey Ica showed a significant microstimulation effect in the direction of the MU tuning preference. Overall, the average magnitude of these significant shifts was equivalent to adding a disparity of 0.03° between the front and back of the cylinder stimulus for monkey Fle, and 0.02° for monkey Ica.

2.3.2. Expected reward size interacts with the effect of electrical microstimulation in area V5/MT

The size of the fluid reward available for a correct choice at the end of each trial was variable. As the number of consecutive correct choices increased, the size of the available fluid reward also increased, in two steps up to a maximum (Figure 2.3). The maximum reward size ('large' reward) was twice the average sub-maximal reward size ('small' reward). To investigate the effect of expected reward size on microstimulation effect, the psychometric functions for microstimulated and non-stimulated trials for site ICA266 are shown separately for small reward trials

(Figure 2.5A) and large reward trials (Figure 2.5B). At this this site, the size of the horizontal shift of the psychometric function due to electrical microstimulation is smaller when the expected reward size is larger.

To compare the effect of expected reward size on microstimulation over all stimulation sites across both monkeys, the trials at each site were separated by microstimulation condition and reward condition, and the resulting psychometric functions were fitted with cumulative Gaussians, as for example site ICA266 (Figure 2.5). The size of the microstimulation shift effect for each reward condition at each site was normalized by dividing the shift by the animal's performance threshold in the reward condition, measured as the standard deviation (SD) of the cumulative Gaussians fitted to the psychometric functions. At those sites where there was a significant overall microstimulation effect in the direction predicted by the MU sensory preference, the size of the normalized microstimulation effect was significantly smaller in those trials where expected reward size was large, compared to those trials where expected reward size was small (large reward trials: mean normalized shift size = 0.78; small reward trials: mean normalized shift size = 1.03; $p < 0.001$, n=27, Wilcoxon sign-rank test). The microstimulation effect was smaller in the large reward trials than the small reward trials for 22/27 of these sites (Figure 2.6A). The effect of reward on microstimulation was also significant when considered separately for each monkey (Fle: p = 0.0093, n = 12; Ica: p = 0.0215, n = 15; Wilcoxon sign-rank tests), and also when all microstimulation sites were taken into account regardless of shift effect direction or significance ($p < 0.001$, n = 48, Wilcoxon sign-rank test).

To check whether the method of cross-site normalization of shift size was responsible for the

Figure 2.4. *(following page)*: A. Multi-unit (MU) cylinder tuning at example site ICA266. *Filled circles*: mean MU firing rate response to cylinder stimuli of different disparities (errorbars $\pm$ 1SD). *Open circle*: baseline firing rate (blank screen). This site has a preference for CW-rotating cylinders (PREF direction = CW). B. Psychometric functions (PMFs) at this microstimulation site. The proportion of PREF responses are plotted for each stimulus disparity (*black* = non-stimulated trials; *gray* = microstimulated trials). Electrical microstimulation consistently increased the proportion of choices in the PREF direction. PMFs were fitted with cumulative Gaussians constrained such that both curves had the same SD but means were allowed to vary. The microstimulation effect was quantified as the difference between the means (i.e. the horizontal shift between the PMFs). In this example, the shift was equivalent to adding 0.0103° binocular disparity to the cylinder. Division by the SD (SD 0.01) gives the normalized microstimulation effect at this site: 1.03. The significance of the microstimulation effect at this site was $p<<0.001$, (χ^2-test).

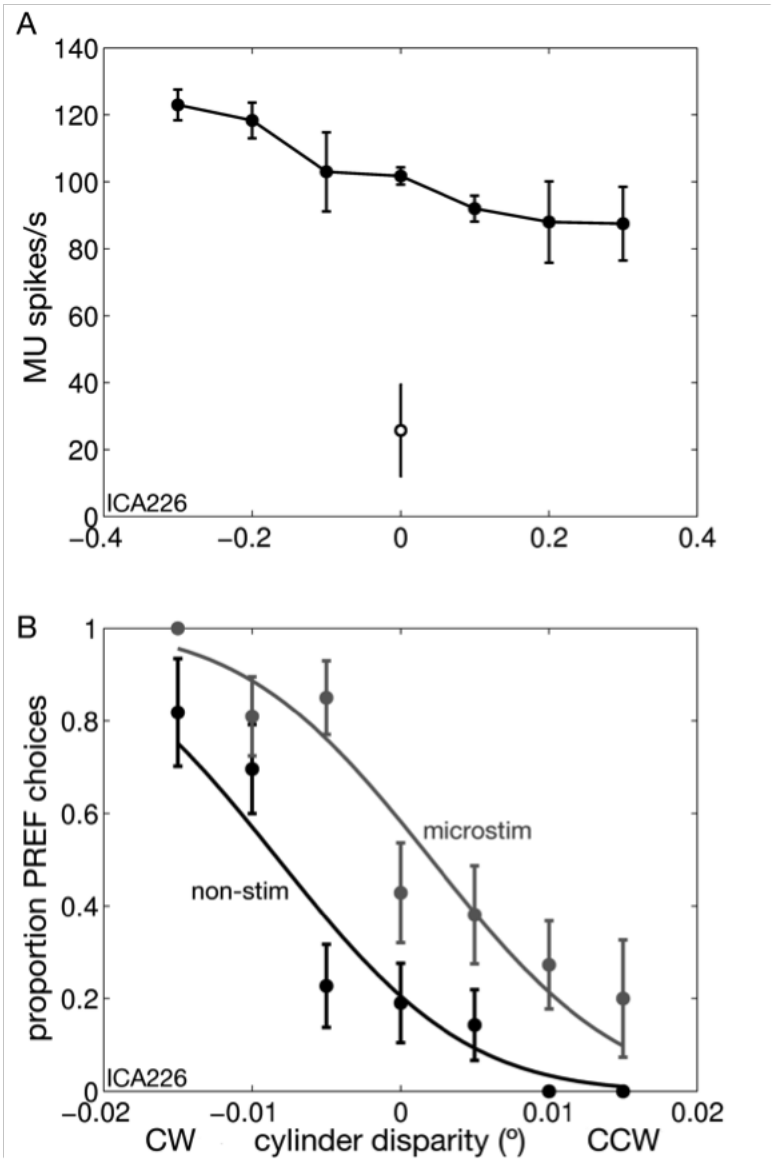

A
MU spikes/s
140
120
100
80
60
40
20
0
ICA226
−0.4
−0.2
0
0.2
0.4
B
proportion PREF choices
1
0.8
0.6
0.4
0.2
0
microstim
non-stim
ICA226
−0.02
−0.01
0
0.01
0.02
CW
cylinder disparity (°)
CCW

Figure 2.4.

result above, the analysis was repeated using the absolute, non-normalized microstimulation effect size at each site. Overall, absolute shift size was significantly smaller in the large reward trials compared to the small ($p < 0.001$, n = 27, Wilcoxon sign-rank test). In 22/27 of those significant microstimulation sites with shift in MU PREF direction, the size of the absolute microstimulation effect in the large reward condition was smaller than in the small reward condition (Figure 2.6B). The difference between large and small reward conditions was also significant for each monkey separately (Fle: p = 0.0210, n = 12; Ica: p = 0.0054, n = 15; Wilcoxon sign-rank tests) and also when all sites were taken into account regardless of shift effect direction or significance ($p < 0.001$, n = 48, Wilcoxon sign-rank test). The effect of reward was therefore not a result of the normalization procedure.

Thus, the data show that the size of the shift induced by electrical microstimulation of a motion- and disparity-selective site in cortical area V5/MT is affected by the reward size for a correct choice in the trial. Specifically, when expected reward size is large, the microstimulation shift effect is smaller. If visual signals are unaffected by reward size, the effect of microstimulation in visual area V5/MT should be equivalent in large and small reward trials. Instead, these data show that the effect of microstimulation changes according to reward, suggesting that the way the artificial visual signal is incorporated into the perceptual decision depends significantly upon the expected reward size at the end of the decision.

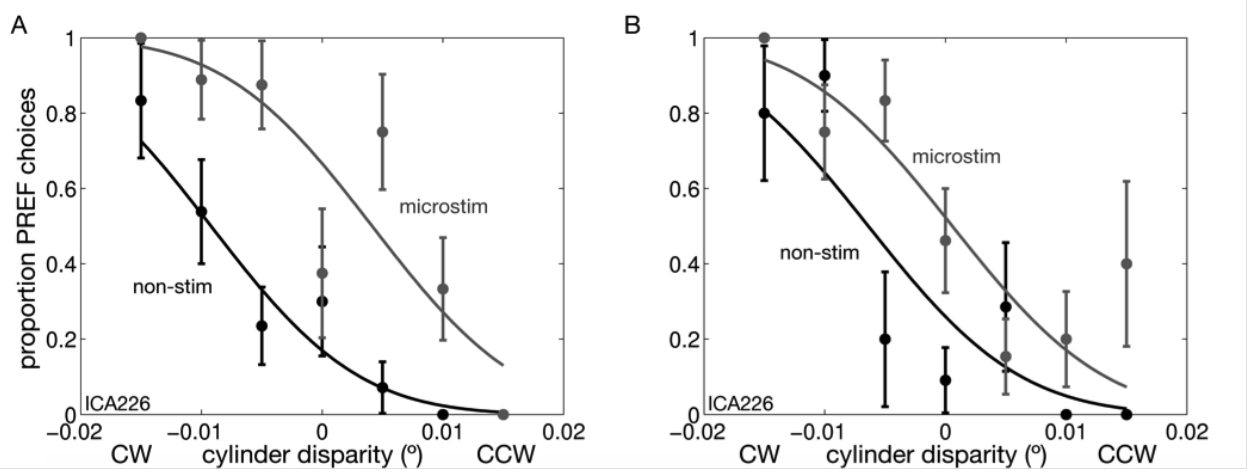


Figure 2.5.: PMFs for microstimulation session at site ICA266 (see Figure 2.4), separated according to whether reward size for correct choice was small (A) or large (B). For small reward trials, microstimulation was equivalent to adding 0.0133° binocular disparity (normalized effect = 1.38); for large reward trials, it was equivalent to adding 0.007° binocular disparity (normalized effect = 0.70).

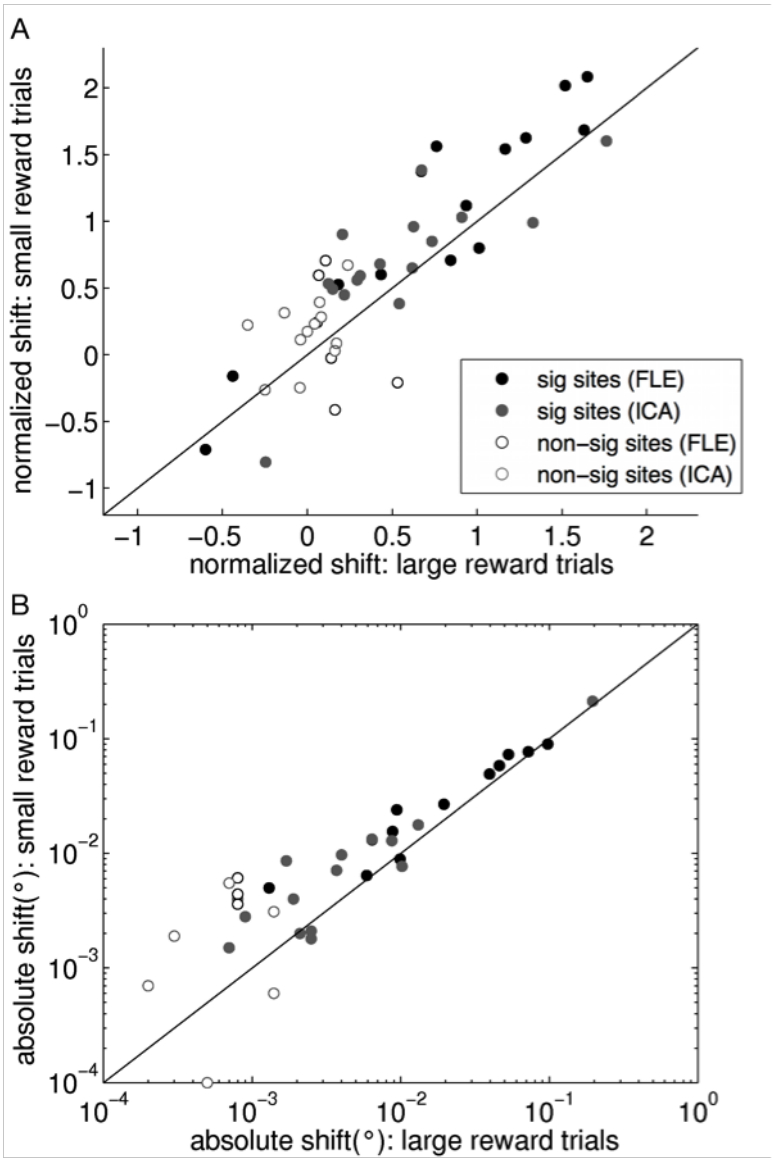


Figure 2.6.: A. Comparison of normalized microstimulation shifts for large and small reward trials for all sites across the two monkeys (*black* = Fle; *grey* = Ica). *Filled circles* indicate sites where microstimulation had a significant overall effect (χ^2-test, $p < 0.05$); *open circles* indicate sites with no significant effect of microstimulation. Overall, normalized microstimulation shift size is greater for small reward trials than for large reward trials (sites with significant shift in MU PREF direction: $p < 0.001$, n = 27; all sites: $p < 0.001$, n = 48 (Wilcoxon sign-rank test)). This is emphasized by the way that the majority of sites lie above the line of unity (*black line*). B. Absolute microstimulation effect for large versus small expected reward trials, legend as for A. Log-scale used, hence three negative shifts, i.e. in MU NULL direction, are omitted from the plot. Overall, the absolute microstimulation shift size is greater for small reward trials than for large reward trials (sites with significant shift in MU PREF direction: $p < 0.001$, n = 27; all sites: $p < 0.001$, n = 48, Wilcoxon sign-rank test).

2.3.3. Reward conditions do not differ in the proportion of microstimulated trials, nor in the distribution of disparity values

The nature of the reward contingencies in this experiment meant that the reward size for any particular trial was dictated by the monkeys accumulated performance on the previous trials, and was not predetermined pseudo-randomly with other experimental parameters. Therefore, experimental factors such as stimulus disparity values and the proportion of microstimulated trials were not necessary equally balanced between the large and small reward trial conditions. Any systematic differences between the distributions of these factors in the two reward conditions could have contributed to the reported difference in the size of the microstimulation effect on large and small reward trials. However, there was no significant difference in the proportion of microstimulated/non-stimulated trials between the two reward conditions in either monkey across all sites (Fle: n = 20, p = 0.627; Ica: n = 28, p = 0.064, Wilcoxon sign-rank test), nor when significant microstimulation sites were considered alone (Fle: n = 12, p = 0.909; Ica: n = 15, p = 0.092). Similarly there was no significant difference in the distribution of stimulus disparities between the reward conditions in either monkey across all sites (Fle: n = 20, p = 0.2627; Ica: n = 28, p = 0.1866; Wilcoxon sign-rank test) nor when significant microstimulation sites were considered alone (Fle: n = 12, p = 0.1099; Ica: n = 15, p = 0.1688) (Table 2.1).

	mean (absolute) cylinder disparity			% microstimulated trials		
FLE	**small rew**	**large rew**	**p val**	**small rew**	**large rew**	**p val**
sig. PREF sites	0.023	0.025	p>0.110	0.5	0.5	p>0.909
all sites	0.018	0.019	p>0.263	0.5	0.5	P >0.627
ICA						
sig. PREF sites	0.014	0.013	p>0.169	0.51	0.49	P>0.92
all sites	0.01	0.01	p>0.187	0.51	0.49	p>0.064

Table 2.1.: The mean absolute disparity, and the percentage of microstimulated trials, for the small reward condition versus the large reward condition. A cross-site Wilcoxon sign-rank test was used to ascertain whether any differences between the reward conditions were significant, considering significant microstimulation effect sites (where the direction of effect was in MU PREF direction) alone (Fle: n = 12; Ica: n = 15) and also over all sites (Fle: n = 20; Ica: n = 28). None of the tests showed significance at the $p < 0.05$ level.

2.3.4. Wearing-off of microstimulation does not explain reward effect

The proportion of large and small reward trials of each condition towards the end, relative to the beginning, of each experimental session could influence the size of the microstimulation effect in the two conditions because the microstimulation effect has been reported to wane over time in at least one study (Salzman *et al.*, 1992). If the effect of microstimulation were to decrease

towards the end of each experimental session or block of trials, and the relative proportion of large reward trials changed towards the end of a session/block, for example due to a training effect on the monkeys performance, then this could explain the reduced microstimulation effect in the large reward condition.

To investigate this, a running average of the proportion of large reward trials was computed for each stimulation site. The gradient of the best-fitting straight line to this running average was calculated for each site, providing a measure of the increase or decrease of the proportion of large reward trials over the experimental session. The mean gradient calculated for the 27 significant microstimulation sites across two monkeys was 0.0015, a small positive gradient suggesting a slight tendency for the proportion of large reward trials to increase towards the end of the session. However, there was no significant correlation between the proportion of large reward trial conditions and the difference in the normalized microstimulation effect between the small and large reward conditions across sites (Spearmans rank correlation: R = 0.22, p = 0.2753, n = 27; linear correlation: R = 0.21, p = 0.2859; Figure 2.7). Therefore, even though there is a small trend towards a greater proportion of large reward trials at the end of a microstimulation session, possibly due to improved monkey performance due to learning, this trend does not explain the effect of reward on microstimulation reported in the main experimental analyses.

2.3.5. Animal cannot adapt discrimination criterion selectivity on large reward trials

A purely behavioural control experiment was performed with one monkey (Ica) in order to investigate the possibility that the reported result arose because the animals were able by some means to distinguish trials with microstimulation from those without, and adopted a strategy of counteracting the effect of microstimulation and implemented this strategy more effectively on large reward trials. The experimental procedure was exactly the same as for the electrical microstimulation experiment, except that no electrical microstimulation of the nervous system took place; instead, a pseudo stimulation signal of +0.005° disparity was added to the cylinder stimulus disparity in the psuedo-stimulated trials during stimulus presentation. Importantly, on these trials the animal was rewarded for perceptual reports that would have been correct had the additional disparity signal not been added to the visual stimulus. These pseudo-stimulated trials were overtly cued by a change in the color, from white to black, of the cylinder dots (see Methods). The resulting psychometric function for these trials was horizontally shifted towards

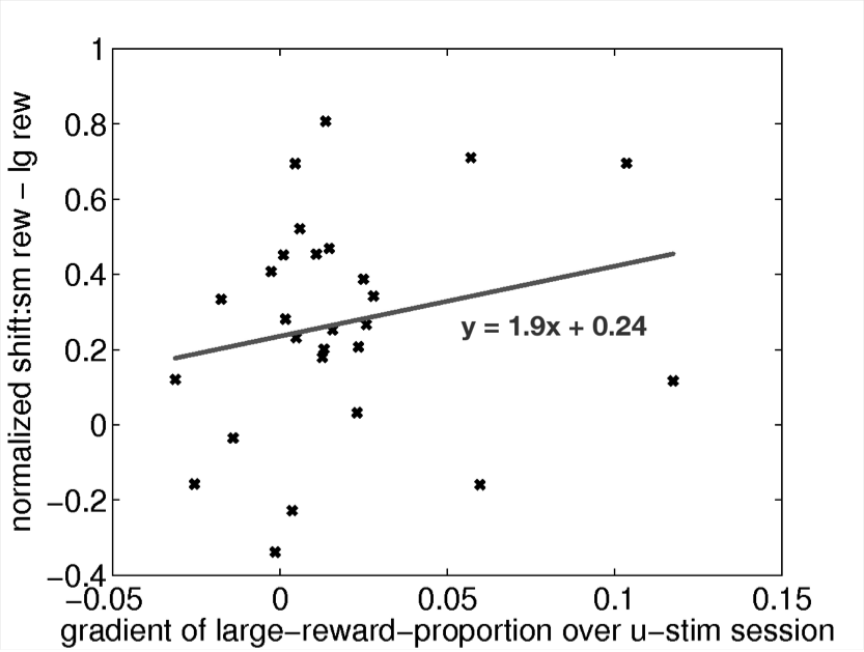


Figure 2.7.: For each significant microstimulation site where effect was in MU PREF direction (n = 27) across 2 monkeys, the difference in the normalized microstimulation shift between small and large reward trials (i.e. the size of the effect of reward on shift) is plotted against the gradient of the line fitted to the running-average proportion of large reward trials over time in the microstimulation experimental session (a positive value indicates an increased proportion of large reward trials near the end). The gray line shows the best-fit linear regression, which is positive but does not reach significance in correlation tests (Spearmans rank correlation: R = 0.22, p = 0.2753, n = 27; linear correlation: R = 0.21, p = 0.2859). Correlation does not reach significance when all microstimulation sites are taken into consideration (Spearmans rank correlation: R = 0.22, p = 0.1408, n = 48; linear correlation: R = 0.13, p = 0.3668; scatter points not shown).

Monkey / Dataset	estimated value of parameter β_7	p-val (χ^2-test)
FLE: electrical microstimulation	-0.3295	p = 0.032
ICA: electrical microstimulation	-0.5014	p = 0.0007
ICA: pseudo-stimulation	-0.0531	p = 0.6547

Table 2.2.: Estimated values for parameter β_7, which represents the interaction between reward condition and microstimulation/psuedo-stimulation shift size in the logistic model used to fit the psychometric response curves (Equation 2.4). The interaction was significant at the $p<0.05$ level for the electrical microstimulation data of Fle and Ica, but not for the pseudo-stimulation data of Ica.

the pseudo-stimulated direction, roughly by the amount predicted by the additional signal (absolute shift = $+0.007^\circ$, normalized value 0.78; compared to electrical stimulation effects in significant sites: Ica: mean shift 0.02°, normalized value 0.69; Fle: mean shift 0.03°, normalized value 0.12). Thus, the pseudo-stimulation experiment successfully emulated the horizontal shift of the psychometric function due to V5/MT electrical microstimulation (Figure 2.8C).

The effect of reward on the fake stimulation effect could not be quantified by using the cross-site Wilcoxon sign-rank test used in the main experiment because there is no equivalent of multiple microstimulation sites in the pseudo-stimulation experiment. Instead, a basic logistic regression model (Equations 2.4 & 2.5) was fitted to the psychometric functions. A χ^2-test was used to evaluate whether the inclusion of the model parameter which allowed the size of the stimulation effect to change with reward condition (parameter β_7 in Equation 2.4) significantly improved the model fit. If it did, then the null hypothesis, that there was no interaction between fake stimulation effect and reward condition, would have to be rejected. For comparison, the same test was also performed on pooled data (across significant microstimulation sites in MU PREF direction) for the electrical microstimulation experiment from both monkeys separately. Figure 2.8 shows the fit of the logistic regression model to the pooled electrical microstimulation datasets for monkeys Fle (A) and Ica (B) and the pseudo-stimulation dataset for monkey Ica (C). The inclusion of parameter β_7 did not significantly improve the fit of the model to the pseudo-stimulation data, although it did significantly improve the fit of the model to the real electrical microstimulation data for both monkeys (Table 2.2). Hence, there was no significant effect of reward condition on the size of the bias induced by the fake stimulation signal in the control data and we cannot account for the effect of reward on the microstimulation shift, reported for the electrical microstimulation data above, by the hypothesis that the animal is aware of being microstimulated, and overtly implements a counteractive strategy more effectively in large reward microstimulation trials.

Figure 2.8. *(following page)*: PMFs fitted by the logistic model (Equation 2.4) for Fle (A) and Ica (B) for electrical microstimulation data and Ica pseudo-stimulation data (C). Electrical microstimulation data was pooled over sites where there was a significant shift in MU PREF direction; pseudo-stimulation data was pooled over all sessions. *Black lines* indicate non-stimulated trials, *gray lines* indicate stimulated trials. *Smooth lines* indicate trials with small available reward, *dashed lines* indicate trials with large available reward. The numbers on the psychometric functions refer to the number of trials at each data point. The interaction between stimulation effect and reward was tested by evaluating whether the inclusion of parameter $\beta 7$ (Equations 2.4 & 2.5) in the model significantly improved the fit of the model to the data (χ^2-test). This interaction was significant for the electrical microstimulation data but not for the fake stimulation data (see Table 2.2).

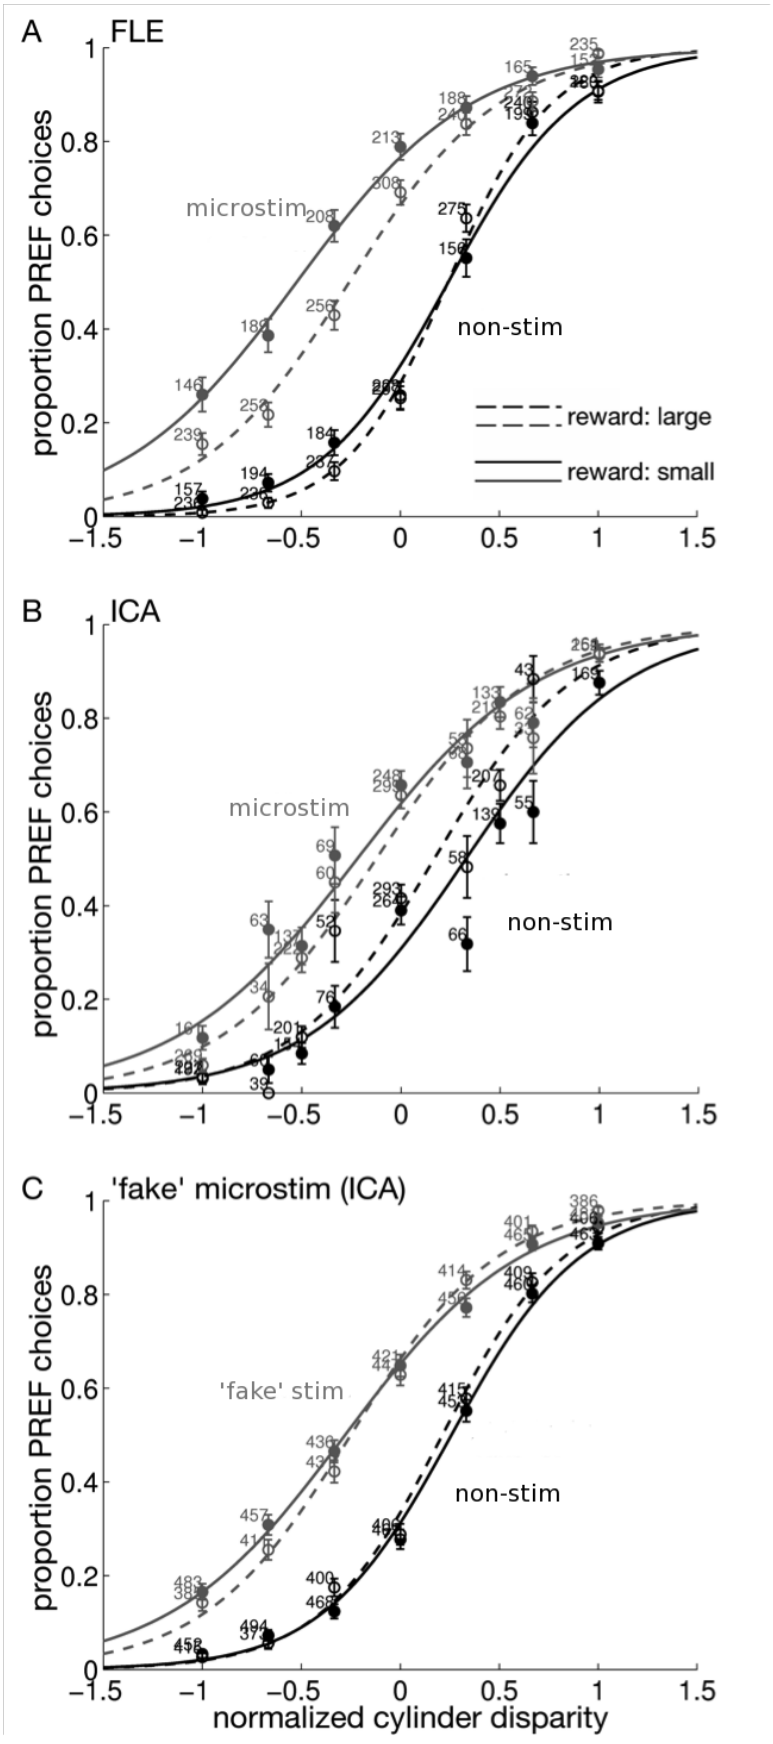

A
FLE
proportion PREF choices
microstim
non-stim
reward: large
reward: small
B
ICA
proportion PREF choices
microstim
non-stim
C
'fake' microstim (ICA)
proportion PREF choices
'fake' stim
non-stim
normalized cylinder disparity

Figure 2.8

2.4. Discussion

Neurons in extra-striate visual area V5/MT were selectively microstimulated whilst two macaque monkeys performed a 2AFC visual discrimination task on a structure-from-motion visual stimulus. As expected, electrical microstimulation in V5/MT shifted the perceptual reports of the animals towards the stimulus properties represented by the microstimulated neurons. The size of the reward available for a correct choice at the end of each trial changed the effect of microstimulation on perceptual decisions, decreasing the size of the microstimulation effect when available reward was larger. These results demonstrate that there is an effect of expected reward size on the incorporation of a specific, artificial electrical signal in sensory cortex into a perceptual decision. This is the first electrophysiological evidence that specific visual information used for a perceptual decision is treated differently according to the expected reward size available for a correct decision. This result is particularly important for understanding how sensory and reward information is combined in the perceptual decision-making process because it implies an effect of reward on visual processing and representation in perceptual decision-making.

2.4.1. Alternative explanations of the experimental data

Systematic differences in the proportions of microstimulated trials or particular stimulus disparity values in the small and large reward conditions could have contributed to the reported difference in the size of the microstimulation effect by reward; however, control analyses found no significant different in the distributions of these factors. Similarly, there was no evidence that the possible wearing off of the microstimulation effect over the course of a microstimulation session could account for the difference in microstimulation effect between reward conditions.

An alternative explanation of the data could be that animals were able by some means to distinguish trials with microstimulation from those without: for example, the monkey may be able to adopt a strategy of counteracting the effect of microstimulation, and may implement this strategy more effectively on large reward trials (when it cares more about being correct) than small reward trials. The lesser effect of electrical microstimulation in the large reward condition compared to the small reward condition might then be explained as a change in the decision criterion for the respective choices, such that the criteria for choosing CW over CCW cylinder rotation directions are shifted. However, in a control experiment, we tested whether monkey Ica could adjust his choice strategy when given an overt cue about which trials contained 'microstimulation'. An pseudo-stimulation signal was incorporated into the visual stimulus by the

addition of extra cylinder disparity, and was overtly cued by a change in the colour of the dots making up the stimulus from black to white. Even with training and feedback through reward, there was no significant interaction between reward size and the shift in perceptual choices induced by the additional stimulus disparity (Figure 2.8). Monkey Ica did not differentially respond to the false cue in the large reward trials to obtain maximal reward, despite the trials with the fake signal being overtly cued, and despite the same stimulus location/axis orientation and fake disparity signal biasing direction being used continuously for 5 days. This may be compared to the electrical microstimulation experiment, in which it is not obvious that microstimulated trials can be distinguished by the monkey, and in which the location, axis orientation and direction of the microstimulated biasing effect were changed once or twice every day according to the microstimulation site MU characteristics, making the counter-acting strategy harder to learn. These argues strongly against the hypothesis that the interation between reward and electrical microstimulation is explained by selective discounting of the microstimulated signal by the animals.

2.4.2. Limitations of the present study

The most significant short-coming of this experiment is the implicit, rather than explicit, nature of the reward regimen and that reward was correlated with the number of previous consecutive correct choices. It was assumed that the animals were aware of the reward contingencies for sequences of correct trials because they were trained upon these sequences for 1-3 years prior to the experiments. Indeed, this reward regime is commonly used in such experiments and training because it successfully motivates the animals to perform the task consistently and in order for this to be the case, animals' behaviour must be influenced by the reward regime. The results reported here hold regardless of whether the monkeys were explicitly aware of the reward size available at the end of each trial; an unconscious expectation of reward size still makes this an interesting and important result. Alternatively, it could be that monkeys shift through different levels of 'arousal' during the experimental session and that in states of lower arousal their performance is worse resulting in greater numbers of low-reward trials. Then, a differential effect of microstimulation in different states of arousal might explain the data. However, 'arousal' is a term that is poorly defined and understood, and has not been investigated with microstimulation; we can only guess, but do not know, whether microstimulation in visual areas should be more or less effective in states of high or low 'arousal'. Furthermore, microstimulation sessions were short (on average 300 trials per experiment), so it is unlikely that animals underwent significant shifts in arousal level in such a short space of time. Regardless, it would be important to

replicate this experiment with pseudo-randomly assigned reward conditions, indicated by an overt cue at the start of each block or trial, in a method similar to Rorie *et al.* (2010), in order to conclusively de-correlate reward size from arousal or previous choice effects. Due to practical limitations, I was not able run such an experiment for this thesis. In addition, it would be very interesting to investigate the effect on microstimulation of a reward cue that biases choices towards a particular rotation direction. Such 'unbalanced' reward conditions might affect microstimulation in a similar way to the balanced conditions reported here and this would further our understanding of different effects of reward on perceptual decision-making. Future work could also combine microstimulation with balanced and/or unbalanced reward regimens in a reaction time task to provide more behavioural information about animals' choices. For example, reaction times could be used to infer states of animals' arousal or to fit detailed models of decision-making (see Chapter 4).

2.4.3. Future work

These results present a previously unreported effect of reward on the incorporation of visual sensory information into a perceptual choice. Further work is needed to understand exactly what affect reward size has on visual processing and the decision-making process. The drift-diffusion model of perceptual decision-making explains the decision-making process very successfully in terms of a small number of constrained parameters that have biological underpinnings and has been extremely useful in interpreting experimental results from many datasets (Gold & Shadlen, 2007). The data presented in this Chapter should be further analyzed in the context of a drift-diffusion in order to illuminate how reward might affect the incorporation of the artificial visual signal into the perceptual decision. This will be the focus of the next Chapter.

3. Reward and electrical microstimulation in V5/MT: Drift-diffusion model

3.1. Introduction

In the previous Chapter, I presented data showing that the size of the reward available for a correct choice in a perceptual decision task affects the incorporation of an artificial, microstimulated signal in visual cortex into the decision process. Here, I apply a simplified drift-diffusion model of decision-making to that data in order to illuminate where in the decision process these interactions between visual information and reward might be taking place.

As described in Chapter 1, the drift-diffusion model of decision-making is a powerful tool for analyzing and interpreting results from psychophysics studies because its parameters have been closely linked with real, discrete brain processes (see Gold & Shadlen, 2007, for a review). Electrophysiological, psychophysical and neuroimaging studies provide strong evidence that the *drift rate* of the decision variable in the model relates to the strength of sensory evidence, which is represented in sensory brain regions (Mazurek *et al.*, 2003; Voss *et al.*, 2004; Heekeren *et al.*, 2004; Hanks *et al.*, 2006; Summerfield & Koechlin, 2010). On the other hand, the increment of the decision variable itself, from its starting point towards the decision bound(s), is represented in the activity of neurons in the lateral interparietal region (LIP) or frontal eye fields (FEF) in monkey electrophysiology studies (Gold & Shadlen, 2000; Roitman & Shadlen, 2002; Mazurek *et al.*, 2003; Huk & Shadlen, 2005; Hanks *et al.*, 2006) and in parietal and prefrontal regions in human neuroimaging studies (Heekeren *et al.*, 2004, 2008; Summerfield & Koechlin, 2010; Mulder *et al.*, 2012). In particular, activity in these regions has been correlated with changes in the distance between the starting point of the DV and the decision bound(s) (Figure 3.1 A). Therefore, fitting the model to experimental data and observing how experimental conditions affect drift rate and distance-to-bound parameters illuminates the brain regions potentially involved in the observed effects. In the following passage, I describe how the drift-diffusion model was

simplified and adapted to fit the microstimulation experiment datasets described in Chapter 2. I refer the reader to Chapter 1 for a further description and discussion of drift-diffusion models generally.

3.1.1. Diffusion model for perceptual decisions and microstimulation

In the drift-diffusion model, sensory information, also called evidence, about the visual stimulus is represented as the difference in activity between two pools of sensory neurons, where the visual stimulus tuning preference of each pool represents one the two possible interpretations of the stimulus (Shadlen *et al.*, 1996; Gold & Shadlen, 2001) (Figure 3.1 B). This information is integrated over time to influence the drift rate of a decision variable (DV), which moves stochastically from its starting point towards one or the other of the decision bounds (Figure 3.1). In our experiment, the perceptual decision to be made was to discriminate whether the cylinder was rotating in a clockwise (CW) or counter-clockwise (CCW) direction (Chapter 2, Figure 2.1). Upon stimulus offset, the animal indicated its perceptual choice by making a saccade towards the target that corresponds to the the decision bound that was reached (Figure 3.1 B). In half of the trials, electrical microstimulation was applied to a neuronal multi-unit (MU) in visual area V5/MT, which contains the visual motion and disparity information used to make choices about the cylinders (Dodd *et al.*, 2001). Microstimulation is therefore modeled as additional sensory evidence towards the cylinder rotation direction preferred by the MU (Figure 3.1 B). Under this design, the probability of a perceptual choice in preferred (PREF) direction of the MU can be modeled as a logistic function:

$$P_{PREF}(C) = \frac{1}{1 + e^{-(B(kC+S)+x)}} \tag{3.1}$$

where C represents the cylinder disparity, or stimulus strength, relative to MU PREF direction, S represents microstimulation, and x represents the *null bias*. Parameter k controls the sensitivity of the DV *drift rate* to the visual stimulus strength C whilst parameter B represents the distance from the starting point of the drift-diffusion process to the decision bounds, along with the overall gain of the accumulation rate (Figure 3.1 B). The *null bias* refers to the tendency of animals to overall bias their responses away from the microstimulated (PREF) direction (Salzman *et al.*, 1992) (see Appendix A for mathematical model derivation). Without the microstimulation intervention S, effects on parameters k and B would be indistinguishable with behavioural choice data alone, because the outcome of an increase or decrease in either parameter would be

to increase or decrease the accuracy of the perceptual discrimination. Normally, these effects on parameters controlling stimulus sensitivity and distance-to-bound can only be differentiated with a full distribution of chronometric (reaction time) data, which is not available from this fixed-duration task dataset (Palmer *et al.*, 2005; Ratcliff & McKoon, 2008; Feng *et al.*, 2009). Here, it is the microstimulation intervention that allows the important differentiation between effects of reward on sensory evidence, represented in early visual areas, and effects on the decision variable, represented in later (e.g. sensorimotor) stages of the decision process.

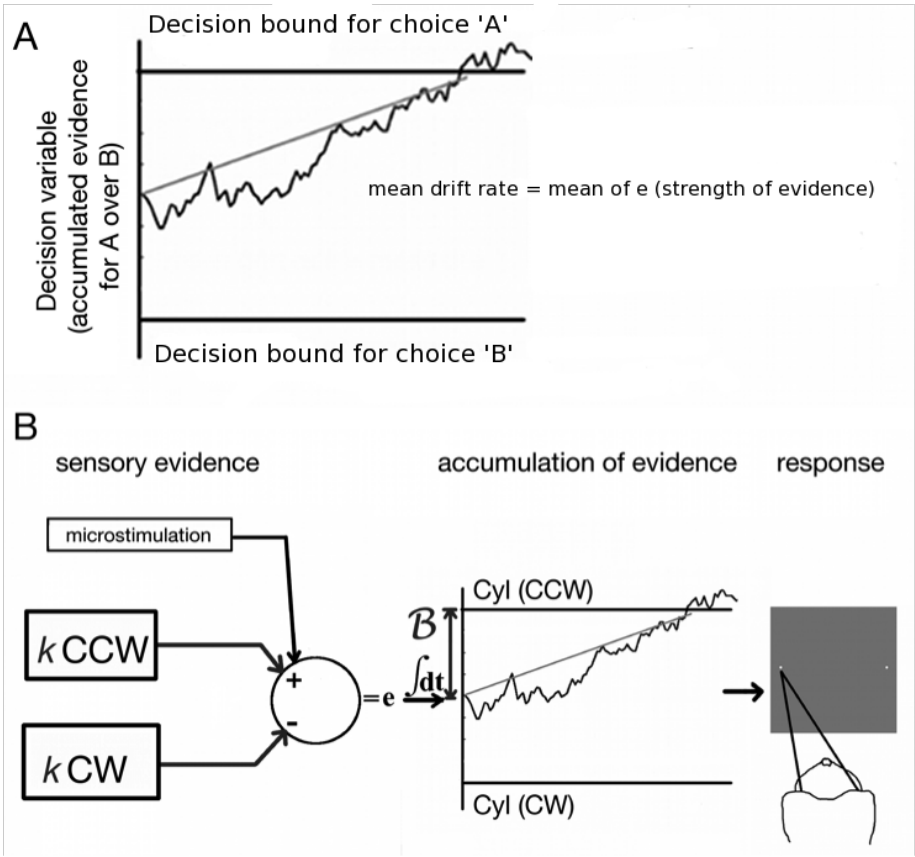


Figure 3.1.: **A**. Diffusion-to-bound model of the decision process. Momentary evidence (e) in favor of choice A over B is accumulated over time. The evidence is normally distributed, with mean μ proportional to stimulus strength C via the relationship $\mu = kC$, where k represents the stimulus scaling sensitivity (see Appendix A). The accumulation of evidence is termed the *decision variable* (DV), and on a single trial follows a stochastic 'diffusion' path (drift) with the *drift rate* depending upon kC. **B**. Drift-diffusion model of the decision process in the cylinder rotation discrimination task with microstimulation. Momentary evidence (e) is represented as the difference in activity between two pools of sensory neurons, where each pool is selective for one of the cylinder rotation directions (CW or CCW). Parameter k thus controls the sensitivity of the neurons to the cylinder disparity. Equal distances to the decision bounds are assumed. Parameter B represents this distance and the overall gain of evidence accumulation. Microstimulation is represented as additional evidence towards one of the cylinder rotation directions, according to the sensory preference of the stimulated V5/MT multiunit.

3.1.2. Possible interactions between reward and decision-making

Effects of reward on decision-making could arise through a variety of interactions with the parameters of this drift-diffusion model (Equation 3.1). Reward may affect the distance between the starting point of the drift-diffusion process and the decision bound and/or the overall gain of the diffusion process (represented by B), or it may affect the sensitivity of the decision variable to the visual stimulus (represented by k), or have combinatorial effects. Intuitions about the effects of changes to these parameters can be gained by observing their effects on the psychometric function (PMF) (Figure 3.2). If the parameter for electrical microstimulation is excluded from the model (by letting $S = 0$), then changes to k and B are indistinguishable because they both affect the accuracy of the perceptual discrimination. For example, improved accuracy results in an increase in the slope of the PMF (Figure 3.2 A). In contrast, the effect of electrical microstimulation in area V5/MT, modeled by letting S equal some number greater than zero, does not affect overall accuracy but instead shifts the position of the psychometric function along the x-axis (Figure 3.2 B). When $S > 0$, the effect of changing B is to change the slope of the psychometric function, without changing the microstimulation effect (Figure 3.2 C). When $S > 0$, the effect of changing k is to modulate both the slope of the PMF *and* the size of the microstimulation effect relative to stimulus strength C, because k selectively affects how stimulus strength, but not microstimulation, is incorporated into the decision process (Equation 3.1, Figure 3.2 D). Alternatively, the reward condition may affect both k and B in tandem (Figure 3.2 E). Finally, the effect of the null bias x is to shift both the microstimulated and non-stimulated PMFs along the x-axis (Figure 3.2).

Two additional potential outcomes must be considered in this model in order to ensure a best fit to the data. Firstly, microstimulation alone might affect the performance accuracy, regardless of reward condition. This has been occasionally observed in previous V5/MT microstimulation studies (Uka & DeAngelis, 2006; Salzman *et al.*, 1992). For example, this could occur if the microstimulation current spreads further than the selected stimulated column to neighbouring cells of different stimulus preference properties, adding general noise to the sensory evidence and thus impairing performance accuracy. This possible effect is modeled as a modulation of stimulus sensitivity k. Also included for completion is the possibility for an interaction between reward condition and null bias x.

The parameter values for the simple drift-diffusion model described above were estimated by fitting the model to the behavioural data for the microstimulation effect in large reward and

small reward conditions. A likelihood ratio test was used to ascertain whether the addition of the parameters for the effect of reward condition significantly improve the fit of the model to the psychometric data. Therefore, the analysis identifies the best-fitting model that explains effect of trial reward size on microstimulation shift effect.

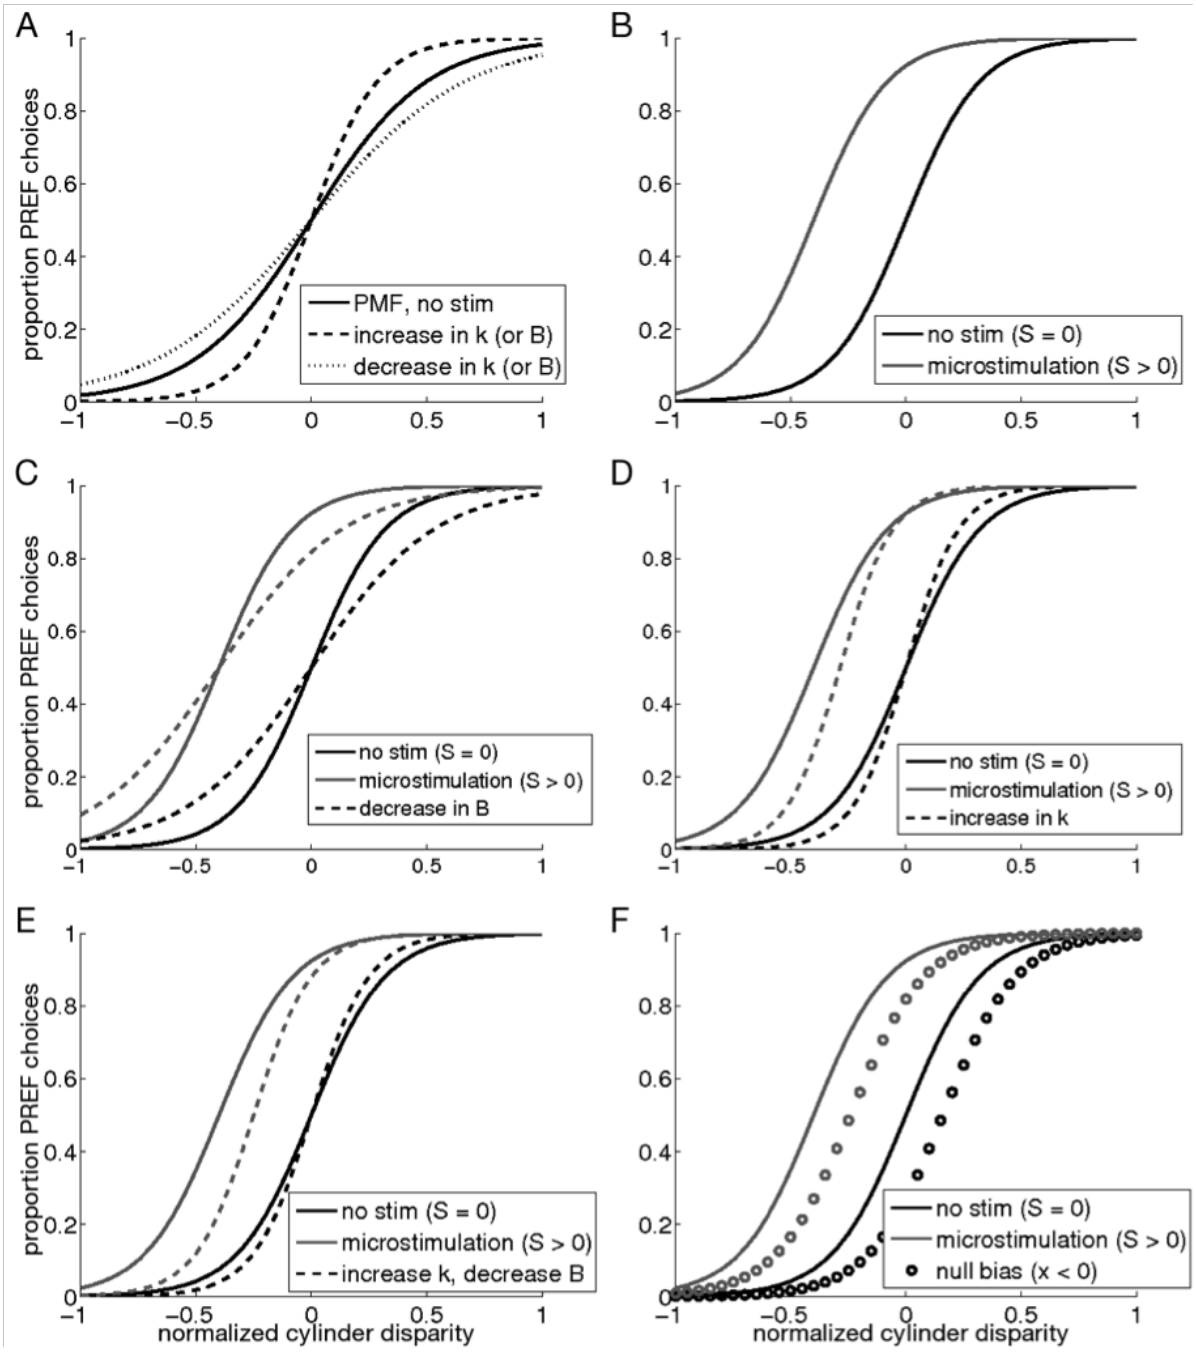


Figure 3.2.: Model simulations presented to provide intuition about the effects on the psychometric functions (PMFs) of changing the the decision model parameters. *Black*: non-stimulated trials; *gray*: microstimulated trials. **A**. Effect on the non-microstimulated PMF of increasing or decreasing parameter k or B. **B**. Effect of microstimulation parameter S. **C**. Effect of decreasing parameter B, on the microstimulated and non-stimulated PMFs (*dashed lines*). **D**. Effect of increase in parameter k (*dashed lines*). **E**. Effects of opposing changes in parameters B and k (*dashed lines*). **F**. Effects of changing the value of null bias parameter x (*dotted lines*).

3.2. Methods

In this chapter, a computational model of perceptual decision-making is fitted to the experimental data analyzed and presented in the previous chapter. Please refer to that chapter's Methods for more information about the task and data analysis.

3.2.1. Drift-diffusion model fitting

A logistic regression model (Equation 3.1) representing the effect of cylinder stimulus disparity, microstimulation and trial reward condition on the psychometric functions was derived from the simple, one-dimensional drift-diffusion model of accumulative decision-making (see Appendix A, after Mazurek *et al.*, 2003; Palmer *et al.*, 2005; Feng *et al.*, 2009). The model was fitted to the psychometric data using maximum likelihood estimate (MLE), implemented using MATLAB's *fminsearch* function, as in the previous chapter.

Four possible interactions between the drift-diffusion model parameters and the experimental conditions were included, as described in the Introduction: 1) interaction between stimulus scaling sensitivity k and reward condition; 2) interaction between bound parameter B and reward condition; 3) interaction between stimulus scaling sensitivity k and microstimulation condition S; and 4) interaction between null bias parameter x and reward condition. The values of the model parameters were estimated by fitting the following equation to the psychometric functions obtained for each of the microstimulation/reward condition combinations:

$$P_{PREF}(C) = \frac{1}{1 + e^{-((B_0+B_1R_{val})((k_0+k_1R_{val}+k_2S)C+S)+(x_0+x_1R_{val}))}} \tag{3.2}$$

where R_{val} = 1 for large reward trials, 0 for small reward trials; S = 1 for microstimulated trials, 0 for non-stimulated trials; B_0 represents the bound parameter in the small reward condition; B_1 represents the change in the bound parameter in the large reward condition; k_0 represents the stimulus scaling sensitivity in the small reward condition; k_1 represents the change in stimulus scaling sensitivity in the large reward condition; k_2 represents change in stimulus scaling sensitivity in microstimulation condition; x_0 represents the null bias in the small reward condition; and x_1 represents the change in null bias in the large reward condition. As such, parameters k, B and x from Equation 3.1 are represented by $(k_0 + k_1 + k_2)$, $(B_0 + B_1)$ and $(x_0 + x_1)$ from Equation 3.2, respectively.

3.2.2. Quantifying interactions between reward and drift-diffusion model parameters

Using Equation 3.2, the significance of a particular parameter, and hence particular condition interaction, was ascertained by testing whether the full model, in which all parameters are allowed to vary freely, fit the data significantly better than a nested model in which the value of the parameter in question is restricted to 0. For example, to test whether allowing stimulus scaling sensitivity k to change with reward condition resulted in a better fit to the data, the full model was compared to a nested model in which the value of k_1 was restricted to 0. Similar steps were taken to evaluate the significance of decision-to-bound parameter B. Each model was fitted to the data with the MLE method. A χ^2-likelihood ratio test was used to evaluate whether the full model fitted the data better than the restricted, nested model, as described in detail in Methods of Chapter 2.

3.3. Results

Experimental data analyzed and presented in Chapter 2 demonstrated that the incorporation of an artificial, microstimulated signal in visual cortex into a perceptual decision was influenced by the reward size available for a correct choice. The simplified drift-diffusion model (Equation 3.1) was fit to the data, pooled across significant stimulation sites separately for each monkey (Ica: n = 15; Fle: n = 12). I investigated whether it was sufficient to model the effects of reward on perceptual decision-making as acting only on integrating sensorimotor structures, represented by parameter B, or, alternatively, whether a reward effect on stimulus-dependent sensory evidence, represented by stimulus scaling sensitivity parameter k, was additionally necessary to explain the data. A χ^2-likelihood ratio test was used to test whether allowing the estimated value for parameter k to vary with reward condition significantly improved the fit of this model to the data, when the estimated value of parameter B was already allowed to vary by reward condition. The value of the null bias parameter x was allowed to vary with reward condition in all cases, in order to not unnecessarily constrain model fit.

3.3.1. Reward affects sensitivity to stimulus strength (drift rate)

The drift-diffusion model was fit to the pooled data separately for monkeys Fle and Ica (Figure 3.3A & B). When stimulus scaling sensitivity k was allowed vary with reward condition, its estimated value increased by 33.2% for monkey Fle and by 80.4% for monkey Ica in the large

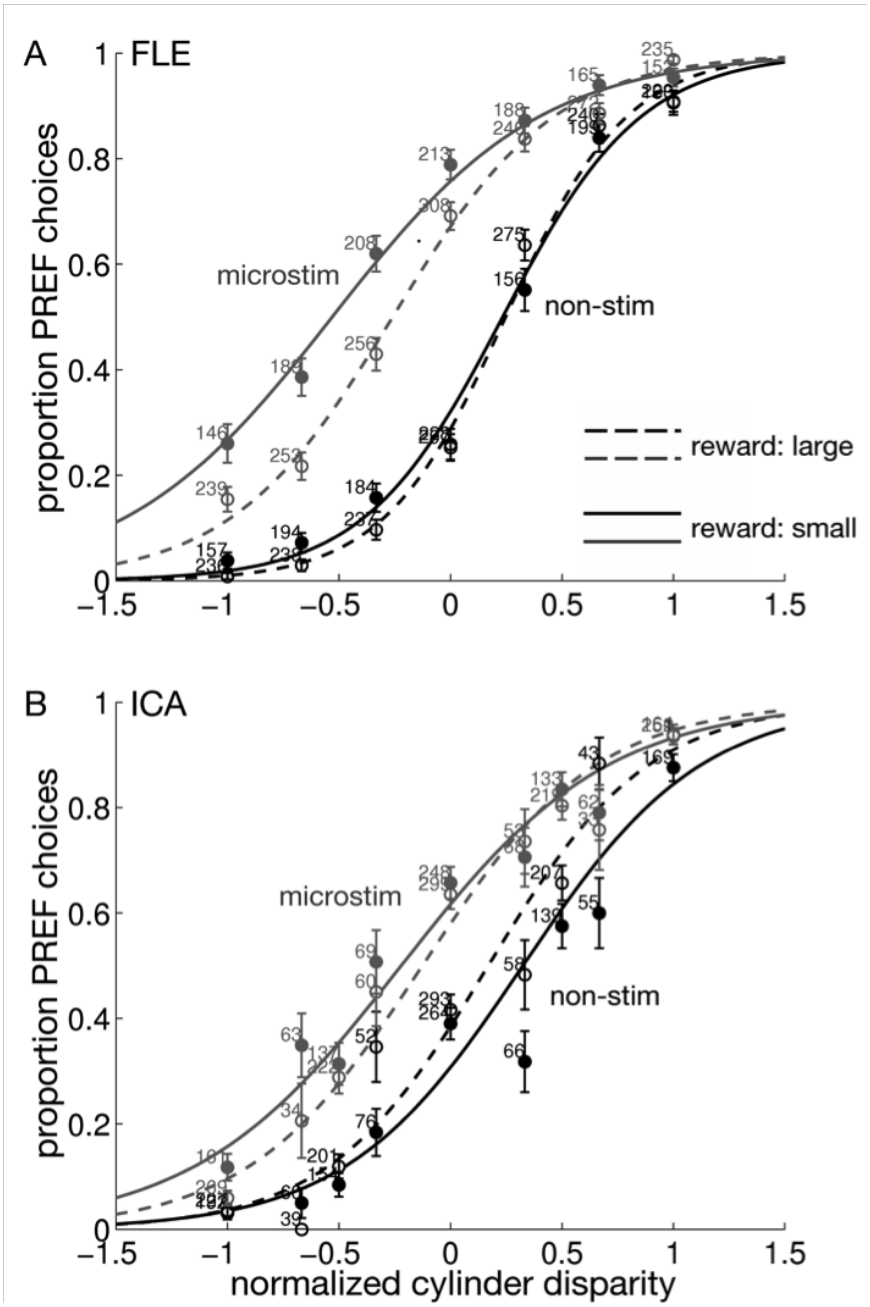


Figure 3.3.: Drift-diffusion model fitted to the pooled reward-microstimulation, for monkeys **(A)** Fle and **(B)** Ica. Data is pooled across the microstimulation sites that had a significant effect in MU PREF direction (χ^2-test, $p<0.05$) and normalized across sites by division by each site's maximum disparity. Numbers refer to the number of trials corresponding to each data point. *Black*: non-stimulated trials; *gray*: microstimulated trials; *dashed lines*: large reward size for correct choice; *smooth lines*: small reward size for correct choice.

compared to small reward condition (Table 3.1). Allowing k to vary by reward condition significantly improved the fit of the model to the data for both monkeys (Fle: p $<<$ 0.001; Ica: p $<<$ 0.001; χ^2-test). This indicates that in the context of this drift-diffusion model, the stimulus scaling sensitivity k, which addresses the relationship between visual stimulus strength and the internal representation of sensory evidence pertaining to the perceptual decision, is significantly increased in those trials where the expected reward size is larger. According to the model, such an increase in stimulus sensitivity would decrease the relative effect of microstimulation on the perceptual choice, explaining the relatively smaller microstimulation effect in large reward trials reported in Chapter 2.

The drift-diffusion model was also fit to to the data pooled across all microstimulation sites, not just significant sites, and the interaction between reward and parameter k was tested again, separately for each monkey (Ica: n = 28; Fle: n = 20). The results were the same: for both monkeys, the estimated value of k increased in the large reward condition and allowing this parameter to change with reward condition significantly improved the fit of the model to the data (Fle: 40.4% increase, p $<$ 0.01; Ica: 71.5% increase, p $<$ 0.01)

		Estimated parameter value:			
Monkey	Parameter	small rew trials	large rew trials	% change	p value
FLE	k	1.70	2.26	+33.2	$p << 0.001$
ICA	k	1.94	3.50	+80.4	$p << 0.001$
FLE	B	1.88	1.62	−14.3	$p = 0.074$
ICA	B	1.29	0.79	−38.4	$p < 0.001$

Table 3.1.: Parameter values for the drift-diffusion model fit (data pooled over significant microstimulation sites). Estimated values of parameters were compared between large and small reward conditions; p-values were obtained from a χ^2-likelihood ratio test comparing the fit of the full drift-diffusion model (Figure 3.3A&B) to a restricted version of the drift-diffusion model in which the parameter in question was not allowed to change by reward condition. Allowing the estimated value of stimulus scaling sensitivity parameter k to change by reward condition significantly improved the fit of the model to the data for both monkeys.

3.3.2. Reward also affects the distance to the decision bounds

I also investigated whether allowing the estimated value for parameter B to vary with reward condition significantly improved the fit of the model to the data, if the stimulus scaling sensitivity k was already allowed to vary by reward condition. The estimated value of the decision bound parameter B decreased by 14.3% for monkey Fle and decreased by 38.5% in monkey Ica in the large expected reward condition compared to the small (Table 3.1). Allowing B to vary by reward condition significantly improved the fit of the model for Ica, but missed significance

for Fle (Fle: p = 0.0736; Ica: p<0.001; χ^2-test). When the model was fit to data pooled over all microstimulation sites, not just significant ones, the estimated value of B again decreased in the large reward condition and this time, this change was significant for both monkeys (Fle: 23.7% decrease, $p < 0.01$; Ica: 30.7% decrease, $p < 0.05$) (Table 3.2).

		Estimated parameter value:			
Monkey	Parameter	small rew trials	large rew trials	% change	p value
FLE	k	2.37	3.33	+40.4	$p << 0.001$
ICA	k	3.18	5.45	+71.5	$p = 0.001$
FLE	B	1.28	0.98	−23.7	$p = 0.005$
ICA	B	0.73	0.51	−30.7	$p = 0.032$

Table 3.2.: Parameter values for the drift-diffusion model fit to the data normalized and pooled over all microstimulation sites. Further details are as for Table 3.1. Allowing the estimated values of stimulus scaling sensitivity parameter k and decision-to-bound parameter B to change by reward condition significantly improved the fits of the models to the data for both monkeys.

3.3.3. Summary

This analysis shows that a simple drift-diffusion model, which allows effects of reward size on the relative gain of sensory evidence and on the position of the decision bounds, can account for the interaction of electrical microstimulation and reward on the psychophysical responses of both monkeys tested here. The fit of the model to the data showed a significant increase in the value of the stimulus scaling sensitivity parameter k (stimulus-dependent drift rate) in large reward trials compared to small reward trials for both monkeys. Additionally, in large reward trials there was a decrease in the value of the bound parameter B. The change in stimulus scaling k provides an explanation for the smaller influence of the evidence introduced by cortical microstimulation relative to the visually evoked evidence, when the expected reward is large. This strongly suggests that the effects of reward act on visual representations prior to their integration towards a decision in sensorimotor structures.

3.4. Discussion

An adapted drift-diffusion model of perceptual decision-making was fitted to the behavioural data presented in Chapter 2, which showed that the effect of microstimulation on a visual perceptual decision was influenced by the size of reward available for a correct choice. The sensitivity of sensory evidence to the visual stimulus was represented by parameter k in the drift-diffusion model, and the analysis showed that an effect of reward on k was necessary in order for the model to explain the change in efficacy of microstimulation intervention by reward. For

both monkeys, an increase in reward size was associated with an increase in the value of k (sensitivity). These reward effects must act prior to the sensorimotor integration of visual stimulus-evoked and microstimulation-evoked sensory evidence into the decision variable, as k pertains only to visual stimulus-evoked sensory representations (Figure 3.1 B). The results therefore indicate that increased reward size increases the sensitivity of visual sensory evidence at the early stages of the perceptual decision process, namely in sensory cortex at or before area V5/MT (Bradley *et al.*, 1998; Dodd *et al.*, 2001; Krug *et al.*, 2004). That this scaling is not applied to the electrically-evoked cortical activity demonstrates for the first time that animals' descision-making differentiates these two types of input and provides the first clear evidence that reward affects perceptual decision-making at the sensory processing stage. In addition, reward affected the decision variable integration itself, which has been reported to be represented in sensorimotor regions (Roitman & Shadlen, 2002; Mazurek *et al.*, 2003; Rorie *et al.*, 2010).

3.4.1. Reward and the sensitivity of sensory evidence to visual input (drift rate)

For both animals, a change in the stimulus scaling sensitivity k with reward condition significantly improved the fit of the drift-diffusion decision model to the data. The value of k is tightly constrained by the data, because it is the only parameter that influences visually-evoked sensory evidence differentially from the microstimulation signal in V5/MT, prior to the combination and integration in of all sensory signals in sensorimotor cortex (Figure 3.1 B). This points to an effect of reward on sensory processing at the level of, or prior to, area V5/MT.

These results significantly extend previous electrophysiological and neuroimaging studies that have investigated brain processing of visual stimuli associated with rewards. Using human fMRI, Serences (2008) and Serences & Saproo (2010) have investigated visual cortical responses in tasks where participants select between two visual stimuli, where the reward probability associated with each stimulus was varied. They showed that selected stimuli were associated with a greater BOLD response in early visual areas than unselected stimuli, and that this difference was increased when the selected stimulus had higher value. With an elegant analysis that utilized voxel tuning functions (VTFs), Serences & Saproo (2010) suggest that high value sharpens the response profile across feature-selective neurons that encode the stimulus feature. This is in line with our findings of increased sensory sensitivity (k) with larger reward. However, for these and other fMRI studies (Pleger *et al.*, 2008, 2009), the low temporal resolution of the BOLD signal limits our interpretation of when in the decision-making process the reported brain activity change occurs. Furthermore, these tasks did not measure perceptual decisions because

the behavioural choice related to reward history and previous choices and not to a visual discrimination. Hence, those studies cannot show if and how the modulations of cortical BOLD signal due to reward might affect perceptual decisions themselves. Similarly, other studies that have investigated responses of visual brain to stimuli associated with differing reward schedules did not use tasks in which the effects of reward on perceptual decisions themselves could be tested (Shuler & Bear, 2006; Pantoja *et al.*, 2007). Rorie *et al.* (2010) directly recorded LIP neuronal activity whilst monkeys performed a perceptual decision task under changing balanced and unbalanced available reward sizes. However, their computational analysis of LIP traces did not reveal any effect of reward condition on the drift rate of the decision variable. The current study therefore extends previous findings by using a perceptual task associated with different available reward sizes, with a specific experimental intervention in an early visual area (V5/MT) with results analyzed in the context of a drift-diffusion decision model. This allows us to show that reward affects the way in which signals in sensory cortex are incorporated into the perceptual decision with which reward is associated, and that this suggests that reward modulates sensory activity and visual perception in the awake, behaving primate.

A discussed in Chapter 1, the relationship between attention and expected value must be considered in all studies that investigate effects of reward associations (Maunsell, 2004). In the current study, model-fitting suggests that in larger-reward trials, the sensitivity (k) of visually-evoked sensory evidence to the stimulus was increased. This resembles the effect on neuronal responses when spatial attention is directed into versus out of their visual receptive field (Treue & Maunsell, 1999) or when feature attention is directed towards their preferred stimulus feature (Maunsell & Treue, 2006). However, such effects of attention cannot trivially explain our findings. Our task always directed spatial attention to the location of the visual stimulus, as the monkey responded to the visual stimulus in every trial, so results from spatial attention studies, which measure differences in neuronal activity when attention is directed into and away from their receptive field, cannot be applied. We are not aware of any studies that attempt to categorically vary the amount the spatial attention applied to the receptive field in the way that the size of reward available was categorically varied in our task. Furthermore, the trial-by-trial reward schedule was balanced, such that both possible perceptual interpretations of the cylinder were equally valuable if correct. Finally, although the neuronal mechanisms of attention are not fully understood, the gain modulation and suppressive drive that are thought to underlie attentional processes (Reynolds & Heeger, 2009) likely act via changes to the overall conductance or synaptic drive to the cells, which affect their responsivity to excitatory inputs (Abbott & Chance, 2005; Chance *et al.*, 2002). If electrical microstimulation activates cells by providing excitatory

input at the dendritic pathways, cell body or axon hillock, then such modulations/suppressions at the level of area V5/MT might be expected to affect the microstimulation signal as well. This is not in line with our results, but further research is required into the cellular mechanisms that mediate the effects of attention and of electrical microstimulation.

3.4.2. Reward and distance to decision bounds

When significant microstimulation sites were considered alone, a change in parameter B, which represents the distance between the starting point of the drift-diffusion and the absorbing bound(s), significantly improved the fit of the model to the behavioural data for monkey Ica and showed a strong trend for Fle. When all sites were considered, the effect of reward on B was significant for both monkeys. Importantly, the instantiation of the diffusion model considered here is a special case of the set of race, or accumulator, models, in which evidence for each alternative is integrated separately in two accumulators, or decision variables (DVs), under the simplifying assumption that the inputs to the accumulators are perfectly negatively correlated (Palmer *et al.*, 2005; Ratcliff & McKoon, 2008). Biologically, the evidence points towards a race accumulator model, with correlated inputs, underlying the decision variable in 2AFC perceptual decision tasks like the one used here, at least at the level of sensorimotor area LIP (Mazurek *et al.*, 2003; Huk & Shadlen, 2005). Because of the simplifying assumption used, it is not possible to ascertain from our model whether the decrease in B means that both the decision bounds have moved closer to the starting point of the DVs or whether it represents the move of the starting point of each decision-variable towards its decision bound in the *race accumulator* model. Physiological firing rates of neurons in the brain structures responsible for accumulating the sensory evidence would need to be measured to resolve this issue. Indeed, the lack of behavioural chronometric (reaction time) measurements in the present data also leaves open the possibility that changes to the parameter B may arise from an effect on the overall gain of the accumulation rate. However, from recent electrophysiological studies (Rorie *et al.*, 2010), human fMRI studies (Summerfield & Koechlin, 2010; Mulder *et al.*, 2012), and psychophysical studies (Diederich & Busemeyer, 2006; Diederich, 2008), a consensus is emerging that reward effects on distance-to-bound occur in part through an offset to the starting position of DV towards the decision bounds in a race model. In particular, Rorie *et al.* (2010) recorded LIP neuronal activity whilst monkeys performed a perceptual decision task under changing balanced and unbalanced available reward sizes and reported that the representation of relative value was most pronounced at the start of the stimulus presentation, consistent with a bias in the starting point of the stimulus integration process. Therefore, the most likely interpretation of the effect on parameter B

reported here would be an offset in the starting position of the decision variable towards the decision bounds, thereby shortening the distance between the decision variable starting points and the decision bounds.

What could be the purpose of decreasing the distance to the decision bound in this way? For unbalanced reward conditions, where one correct choice is more greatly rewarded than the other, an offset of the decision variable starting position towards the choice with greater reward is strategic because the perceptual choices will err towards the choice of greater reward when stimulus discrimination is difficult in conditions of ambiguous or weak stimulus strengths, resulting in greater accumulation of reward (Rorie *et al.*, 2010; Feng *et al.*, 2009; Gao *et al.*, 2011). However, it is less clear why an increased offset towards the decision bound(s) in a balanced large reward condition compared to a balanced small reward condition, used here, is a useful strategy. In fact, this would be expected to make performance accuracy worse in the large reward condition, when the animal might care most about being correct, because the closer proximity of the initial value of the accumulator to the bound(s) would result in a greater effect of noise in the accumulation process on the final decisions. Indeed, in order to account for their data, Rorie *et al.* (2010) had to neutralize the effects of reward-related distance-to-bound changes on accuracy, by including a large 'inter-trial' variability parameter in their model. However, large inter-trial variability is not required to account for the data presented in this study: in the model, the concomitant increase in stimulus scaling sensitivity k increases performance accuracy, balancing the decrease in accuracy that results from the shortening of the distance between the starting point and the decision bound(s). In other words, increasing the proximity of the decision bound to the starting point of the accumulator decreases the mean time taken to reach a decision and whilst also decreasing the accuracy of the response based upon the available information. This latter effect on accuracy is balanced by a simultaneous increase in stimulus scaling sensitivity k that simulateously improves performance accuracy. Therefore, if performance accuracy is maintained overall (or even increased overall by large relative increase in k), in a reaction time task the faster mean response time due to smaller B would increase the rate of reward collection. Maximizing the overall rate of reward collection in a given experiment is presumably the most important motivator for the subject in these tasks, as discussed in (Gold & Shadlen, 2002). Although the task in the present study was a fixed duration task and not a reaction time task, recent findings suggest that even in fixed stimulus viewing duration tasks, the decision is made when the evidence accumulator is absorbed by a decision bound (Kiani *et al.*, 2008), and therefore would be expected to be subject to the same effects of reward.

3.4.3. Limitations of the drift-diffusion model

The simplified logistical derivation of the drift-diffusion model used here is only a basic approximation to the neuronal dynamics that underlie perceptual decision-making in the primate brain. As discussed above, the present task design, which incorporated a fixed stimulus-viewing duration rather than a reaction time task, limited the parameters of the model that could be fitted to the present data. Also, due to the need to make inductions about reward effects from changes in model parameters, it remains potentially unclear whether these effects may be due to strength of coupling of visual evidence to response, or reflect a genuine change in perceptual processing. Furthermore, the drift-diffusion model does not include a representation of changes in the size of the pool of the sensory neurons used in the perceptual choice. A further evaluation of the effects of changing pool size versus changing multiplicative gain is out of the scope of the present analysis. Regardless, the important distinction between a reward effect on visually evoked cortical activity and electrically microstimulated cortical activity could still be made.

More complex models of perceptual decision making, as reviewed extensively elsewhere (Bogacz *et al.*, 2006; Eckhoff *et al.*, 2008), may provide better explanations of the canon of perceptual decision-making data than the simple, adapted drift-diffusion model used in this study. These models include leaky diffusion models such a the leaky competitive accumulator (Usher & McClelland, 2001), competing race models without perfectly correlated inputs (Smith & Vickers, 1988; Reddi *et al.*, 2003), independent sampling models (Maloney & Wandell, 1984), or even a recurrent competitive network between V5/MT and MST which forms the decision (Grossberg & Pilly, 2008). However, except for the latter two, under simplifying assumptions when reaction time data is unavailable, all of these models generalize to the simple drift-diffusion model used in this analysis (Palmer *et al.*, 2005; Feng *et al.*, 2009). More data obtained from the electrophysiological recordings of neuronal activity and more complex task designs, focusing on collecting reaction time data as well as perceptual choices, is required to differentiate between these model groups in any conclusive way.

3.4.4. Summary and future directions

In conclusion, electrical microstimulation was used to investigate the effect of expected reward size on perceptual choices, in the context of a simplified drift-diffusion model. The experiment showed that the microstimulation shift effect was smaller when the available reward size for a correct response is larger and this result was explained in context of a drift-diffusion model by an

increase in the stimulus scaling sensitivity for visually-evoked cortical activity when the available reward is larger. Future work must evaluate the effect of reward on perceptual decisions in a reaction time task that records the time taken for the viewer to make a perceptual choice. This would add additional information about the way that reward affects the decision-making process, in particular in terms of fitting the drift-diffusion model and teasing out the effects of reward on different model parameters and hence the corresponding underlying brain processes. The use of a reaction time task and full drift-diffusion model in a human psychophysics task is the focus of the next Chapter. Most importantly, further electrophysiological experiments in animals, and neuroimaging experiments in humans, are needed to directly record the activity of brain regions that represent visual sensory information, in order to delineate exactly how this information representation is affected by reward context. Such investigations into the neural correlates of reward in early cortical representations of sensory evidence would provide the perfect compliment to the results revealed here by the electrical microstimulation approach; a human MEG experiment on this theme is described in Chapter 7.

4. Cylinder Reward Psychophysics in Human Participants

4.1. Introduction

In Chapters 2 & 3, we have seen that the size of reward available for a correct choice affects the integration of an electrical microstimulation signal in visual cortical area V5/MT of macaque and that, in the context of a simplified version of the drift-diffusion model of decision-making, this indicates an effect of reward prior to the integration of sensory evidence into sensorimotor structures. In the present Chapter, I further investigate the effect of reward on sensory decisions in with human psychophysics experiment using the same cylinder discrimination task.

In psychophysics experiments, full reaction time distributions, in addition to behavioural choice data, are required to separate effects of reward on diffusion model drift rate and distance-to-bound parameters (Figure 1.4). This is because changes in the drift rate and distance-to-bound parameters can give the same response bias in behavioural choice data considered alone, regardless of the fact that one indicates an effect on sensory representations and the other on sensorimotor and output stages (Chapter 1). These parameters do however make different predictions about the reaction time (RT) distribution associated with the behavioural bias (Ratcliff & McKoon, 2008). For example, changes to drift rates predict changes to the leading edge and the tail of the RT distributions (0.1 and 0.9 percentiles) to be in the ratio of about 4:1, whilst changes to the distance-to-bound predict a ratio of about 2:1 (Ratcliff & McKoon, 2008). The use of a fully-parameterized diffusion model, which includes parameters for inter-trial variability in the drift rate, starting point position, and non-decision time, helps ensure that effects in these variables do not contaminate estimates of the main parameters of interest. The present study contained enough trial repetitions to produce full RT distributions at the level of the individual participant, and a fully-parameterized drift-diffusion model was fit according to verified methods to maximize the sensitivity and accuracy of effects on model parameters (Vandeker-

ckhove & Tuerlinckx, 2007, 2008).

As reviewed in Chapter 1, most previous studies into effects of reward on perceptual judgements have found significant effects of reward only on the distance-to-bound parameter of the drift-diffusion process, which has been correlated with activity in sensorimotor brain regions of macaques and humans (Roitman & Shadlen, 2002; Summerfield & Koechlin, 2010). Based upon the results of Chapters 2 & 3, here I test the specific hypothesis that reward bias on perceptual choices acts through an effect on the drift rate parameter of the diffusion model, in addition to any effects on the starting point of the decision-variable integration to the decision bounds. Before applying the drift-diffusion model to investigate the effects of reward, it was necessary to show that the drift-diffusion model adequately described perceptual decisions in the structure-from-motion (SFM) cylinder task and that the drift rate parameter of the model corresponds to the processing of sensory information about the stimulus (Voss *et al.*, 2004).

4.2. Materials & Methods

4.2.1. Participants

Ten human volunteers (5 female, 5 male, 1 left-handed, 9 right-handed, age mean 29.3 years $\pm$ 12.4, range 20 - 59 years) participated. All had normal or corrected-to-normal visual acuity and static binocular disparity threshold above 60° of arc (TNO stereo test, Lameris, Utrecht). All participants had no reported history of neurological problems and no current psychoactive medication. The research was approved by the University of Oxford's Central University Research Ethics Committee (CUREC) and written informed consent was obtained from all participants.

Seven participants (*crw1-7*) took part in experimental sessions of the cylinder discrimination task, in which reward conditions were manipulated; participant *mrw3* was excluded because they did not complete enough trials. All participants apart from *crw1* and *crw2* were fully naive to the experiment research goals. Participant *crw1* had previous experience conducting visual psychophysics studies, but was not involved in research into reward biases and was naive to the goals of this study. Participant *crw2* was non-naive (the author). Four further participants *srt1-4* took part in experimental sessions of the cylinder discrimination task in which there were no reward manipulations. These datasets were collected by an undergraduate project student (Samuel Mills). Participants *srt1* and *srt2* had previous experience doing the cylinder task but were naive to the goals of this study; participant *srt3* was fully naive; *srt4* was non-naive.

4.2.2. Stimuli

The task involved trial-by-trial perceptual discrimination of the direction of rotation of a structure-from-motion (SFM) cylinder, similar to the task in Chapter 2. Stereoscopic cylinders were generated by a custom C++ program running on Qt software platform (*http://qt.nokia.com/*).The program was originally written by Stuart Gilson, University of Oxford; modifications for investigating effects of reward were designed by the author. Stimuli were presented with a Samsung R710 laptop (Centrino 2 P7450 processor, nVIDIA GeForce 9300M graphics card, Linux), which provided image signals to two identical CRT monitors (Eizo FlexScan T961) via a three-channel splitter cable. The visual stimulus for the left eye was presented on the left monitor and for the right eye on the right monitor. Screens were separately directed to the left and right eyes of viewers by a Wheatstone stereoscope configuration (Wheatstone, 1838). Binocular disparities were produced by presenting an image displaced by the appropriate distance to each eye. Total viewing distance, 2480 mm; monitor size, 385x290 mm; resolution, 1024x768 pixels. Cylinders were presented in the centre of each screen, 2.35° in height and diameter, formed of 50 black and 50 white dots of size 0.1°. 2% of dots were re-plotted in a new location in each frame (screen refresh rate 60Hz). Rotation speed was 83% of full rotation (298.8°) per second. 'Cylinder disparity' refers to the binocular distance between the centre of the cylinder and its front wall at the point of largest disparity. In this chapter, the convention is that a positive disparity refers to a clockwise (CW) rotating cylinder as viewed from above. Mean background luminance was 80 cd/m^2; mean stimulus luminance was 95 cd/m^2; mean Michelson contrast between black and white dots was 56.5%.

4.2.3. Task

4.2.3.1. Training

Participants performed 3-4 training blocks (144 trials/block) of the basic cylinder reaction time task. A central fixation point was presented at the beginning of each trial for 450 ms, followed by onset of the cylinder stimulus (Figure 4.1 a). Participants pressed the left mouse button to indicate leftward cylinder rotation (CW) or the right mouse button to indicate rightward cylinder rotation (CCW) (see Appendix B for example instruction sheet). The stimulus and fixation point disappeared once the response was given. The next trial commenced after an inter-trial interval (ITI) of 1000 ms. No points or feedback were given. Nine cylinder disparity values were presented for participants in group *crw*. The initial cylinder disparity range using in training was (-0.04, -0.02, -0.01, 0.005, 0, 0.01, 0.005, 0.01, 0.02, 0.04)°. This disparity range was either kept,

or increased by doubling each disparity value, or decreased by halving each disparity value, according to individual sensitivity, to as to give a full psychometric function (PMF) for each individual. For *srt* participants, seven cylinder stimulus disparity strengths were presented, centered on zero disparity, with the range selected according to the individual participants' stereothreshold in a similar manner.

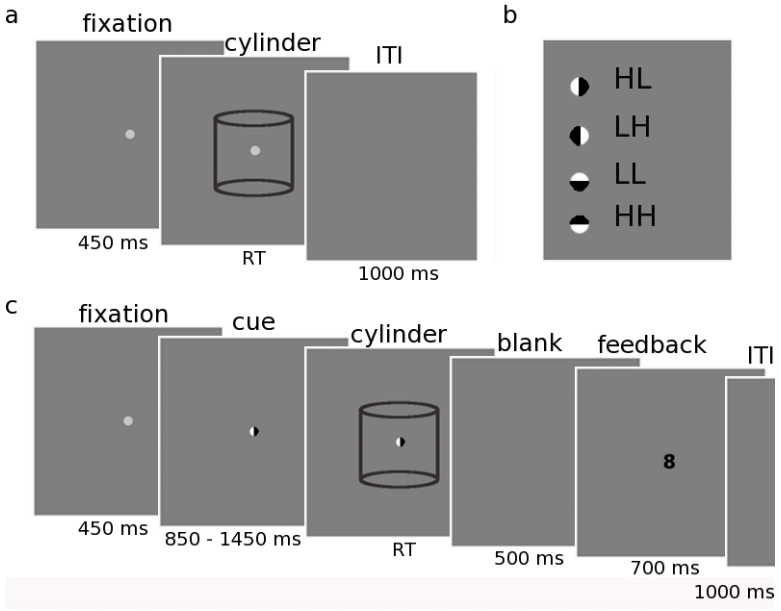


Figure 4.1.: **a** Trial sequence for the basic cylinder reaction time task, using in all training sessions and in the *srt* participants' datasets. **b** Points cues for the four reward conditions, see Table 7.1. **c** Example trial sequence for the reward task, depicting a 'HL' points cue. ITI = inter-trial interval; RT = reaction time.

Table 4.1.: Points assignments in the four reward conditions:

condition:	Points for correct CW choice (left button)	Points for correct CCW choice (right button)
HH	**8**	**8**
LL	1	1
HL	**8**	1
LH	1	**8**

4.2.3.2. Basic task (no reward)

After training, participants in group *srt* completed 15 blocks (140 trials/block) of the same reaction time task (Figure 4.1) in 3 experimental sessions on separate days, giving a total of 300 trials per cylinder disparity, except for participant *srt4* who completed 10 blocks, giving a total of 200 trials per cylinder disparity. Participants were remunerated for their time as no points/feedback were involved in the basic task.

4.2.3.3. Reward task

After training, participants in group *crw* completed 15 blocks of trials (150-200 trials/block) across 3-4 experimental sessions on separate days, in which reward conditions were manipulated on a trial-by-trial basis. Four possible reward conditions were high-high (HH), low-low (LL), high-low (HL) and low-high (LH) (Table 4.1). Reward conditions were crossed with cylinder disparity to produce 4 x 9 = 36 possible experimental conditions. For three participants (*crw2*, *crw4* and *crw5*) all four reward conditions were pseudo-randomly interleaved across trials. The number of trials collected in total for each reward condition and each stimulus disparity was 120, the minimum needed to obtain the full reaction time (RT) distribution necessary for fitting the full drift-diffusion model for an individual participant (Ratcliff & McKoon, 2008). Due to the large number of trials required and the consequent long experimental duration, only three reward conditions (LL, HL and LH) were presented for the remaining three participants (*crw1*, *crw6* and *crw7*). The instruction sheet informed participants that points earned were converted to money at the end of the experiment. For participants who were presented with all 4 reward conditions, points-to-money remuneration rate was 150 points:1 GBP. For participants who were presented with only 3 reward conditions, remuneration rate was 250 points:1 GBP, to compensate for the absence of the highest-points condition (HH). Participant *crw2* (myself) was not remunerated.

Trial Sequence: The fixation point (size 0.2°) appeared at the beginning of each trial, replaced after 450 ms by the *points cue* (size 0.2°) which indicated trial reward condition (Figure 4.1 b & c). The points cue always consisted of a circle, same size as the fixation point, bisected into a white half and a black half. A vertical bisection indicated different numbers of points available for each response (conditions HL & LH). A horizontal bisection indicated that equal points were available for each response direction, but the overall points available could be either high (HH) or low (LL) (Table 4.1). Points cue luminance and contrast was therefore the same across trials and differed only in orientation; the association of high reward with the white or black cue side was counterbalanced across participants. The cue was presented for a variable duration (850 - 1450 ms; uniform distribution) to prevent anticipation of the stimulus onset, which could encourage fast guess responses or induce rhythmicity in viewers' response times, contaminating RT data. The cue remained on screen when the cylinder stimulus appeared, to minimize visual changes at stimulus onset. Participants were instructed to respond as 'speedily and accurately' as possible and to use the points cue to get as many points as possible overall (see Appendix B for example instruction sheet). Stimuli offset when the response was given. 500 ms

after offset, visual feedback (duration 700 ms) indicated points gained, in the form of a number presented in place of the fixation point: 0 = incorrect response; 1 or 8 = correct response. A blank screen inter-trial interval (ITI) of duration 1000 ms separated trials. Every 20 trials, the sum of points earned so far in the block was presented on the screen for 4000 ms. The reward manipulation experimental design was specifically chosen to be comparable to that of Rorie *et al.*, 2010. Eye movements were not tracked but participants were instructed to fixate on the fixation point/points cue at all times.

4.2.4. Analysis procedures

4.2.4.1. Preprocessing

Removal of outlier reaction times (RTs) is essential for reliable fitting of the drift-diffusion model; contaminant reaction time values significantly affect parameter estimates (Ratcliff, 1993; Ratcliff & Tuerlinckx, 2002). Minimum RT cut-offs to remove fast guesses were calculated for individual participants using the exponentially-weighted moving average (EWMA) method (Vandekerckhove & Tuerlinckx, 2007): RTs for all trials across all conditions (except zero-disparity) were sorted in ascending order and the proportion of correct choices was calculated for increasing RT values, using an exponentially weighted window. The RT value at which response accuracy first became significantly different from chance (50% correct) was taken as the minimum cut-off and applied to all cylinder disparity conditions, including zero disparity. Performance was judged to be above chance if it surpassed the upper control limit (UCL) of a random generator process based upon a Bernoulli(0.5) distribution (Chandra, 2001; Roberts, 1959). Any subsequent dips below the chance were ignored.

To remove long-duration outliers, the RT distribution was transformed to resemble a normal distribution by applying the commonly-used logarithm transformation. A normal distribution was fit to the resultant log(RT) histogram and the maximum RT cut-off was assigned as mean(log(RT)) + 2.5*SD(log(RT)) where SD = standard deviation of the fitted normal distribution. This value was chosen to remove approximately 2-3% of the extreme outlier reaction times, which is the commonly-reported, reasonable proportion of rejected trials (Ratcliff & Tuerlinckx, 2002; Ratcliff & McKoon, 2008).

4.2.4.2. Psychometric functions & optimal strategy analysis

Psychometric functions (PMFs) in Figures 4.4 & 4.9, which describe behavioural choices to different cylinder disparities, were fitted with a cumulative Gaussian distribution (for details see Methods, Chapter 2). Each PMF is defined by two parameters: the *mean* (μ) and the *standard deviation* (SD; σ). To combine results across participants for average PMFs, cylinder disparity values were normalized by division by the largest disparity value used for each participant. To predict optimum strategy for maximal reward collection, a custom Matlab script calculated the total number of points collected under each reward condition for any given μ and σ. Matlab's *fminsearch* function (Nelder-Mead method) was used to find the optimum bias shifts for the PMFs in the unbalanced reward conditions, HL and LH.

4.2.4.3. Quantile-probability (QP) plots and quantile-quantile (QQ) plots

QP plots and QQ plots (Figure 4.7) were generated by assigning trials into groups according to the response direction (CW or CCW) and cylinder disparity value. For each response-disparity group, the 0.1, 0.3, 0.5 (median), 0.7, and 0.9 proportion percentiles of the cumulative RT distribution were calculated using the Matlab *percentile* function. The resultant RT percentile values for each group were plotted against the associated proportion of choices in the respective direction (CW or CCW) (Ratcliff & McKoon, 2008). QQ plots were generated by plotting the RT percentile values for each disparity group against the values of the group with the second-largest disparity value. The linearity of the relationship between RT percentile values between groups was then tested with Matlab's *corrcoef* linear correlation regression analysis. To generate averaged QP and QQ plots, RT percentiles were first established for each individual, then averaged across subjects (Ratcliff, 1979).

4.2.4.4. Model Fitting

Methods for fitting the full drift-diffusion model to psychophysical data efficiently and accurately are a separate research topic in their own right (Ratcliff & Tuerlinckx, 2002; Tuerlinckx, 2004; Vandekerckhove & Tuerlinckx, 2007). To ensure valid and optimal model fitting in this study, the model was fit to individual participant data using the Diffusion Model Analysis Toolbox (DMAT) version 0.4 (11 Apr 2007) (Vandekerckhove & Tuerlinckx, 2008). The DMAT package has been extensively theoretically validated, and applied to psychophysics and neuroimaging data in a number of recent publications (Tuerlinckx, 2004; Milosavljevic *et al.*, 2010; Wenzlaff *et al.*, 2011). Beyond trivial bug fixes, no changes to the DMAT software were necessary for ap-

plication to this study.

The parameter estimation method for model-fitting minimized the negative log of the 'multinomial likelihood' (MLF) value, a measure of model goodness-of-fit (Vandekerckhove & Tuerlinckx, 2007). The negative log of the MLF is symbolized by 'Δ'. Sixteen RT percentile values obtained from the data (0.01, 0.02, 0.05, 0.1, 0.3, 0.5, 0.7 and 0.9 RT percentiles for both correct and error RT distributions) were used to fit model parameter estimates to each dataset.

4.2.4.5. Model Selection

(1) To investigate how *cylinder disparity* (stimulus strength) affected model parameters, the following four models were successively evaluated for each participant:
Model 1: 'No effects' (no effect of cylinder disparity on any parameter);
Model 2: 'Linear drift rate' (drift rate allowed to vary linearly with disparity);
Model 3: 'Free drift rate' (drift rate allowed to vary freely with disparity);
Model 4: 'Unconstrained parameters' (all parameters allowed to vary freely with disparity).

(2) To investigate how reward bias (*unbalanced* conditions: HL vs. LH) affected model parameters, the following four models were successively evaluated for each participant:
Model 1: 'No effects' (drift rate allowed to vary linearly with cylinder disparity, and no parameters allowed to vary with reward condition);
Model 2: 'Starting point effects' (in addition to the effect allowed in Model 1, the starting point of evidence accumulation (decision variable), and its inter-trial variability, was also allowed to vary freely with reward condition);
Model 3: 'Drift rate effects' (in addition to the effects allowed in Models 1 & 2, the linear relationship between the drift rate and cylinder disparity was also allowed to vary with reward condition);
Model 4: 'Unconstrained parameters' (all parameters allowed to vary freely with cylinder disparity and reward condition).

(3) To investigate how overall reward size (*balanced* conditions: HH vs. LL) affected model parameters, the following four models were successively evaluated for each participant:
Model 1: 'No effects' (drift rate allowed to vary linearly with cylinder disparity, and no parameters allowed to vary with reward condition);
Model 2: 'Distance-to-bound effects' (in addition to effect allowed in Model 1, the starting point

of evidence accumulation (decision variable) and its inter-trial variability, and the distance between the decision bounds, were also allowed to vary freely with reward condition);
Model 3: 'Drift rate effects' (in addition to the effects allowed in Models 1 & 2 above, the linear relationship between the drift rate and cylinder disparity was also allowed to vary with reward condition);
Model 4: 'Unconstrained parameters' (all parameters allowed to vary freely with cylinder disparity and reward condition).

The χ^2 assumption, used in Chapter 3 to evaluate significance of difference of fit of nested models, does not hold when fitting the drift-diffusion model with data-driven RT percentiles (Speckman & Rouder, 2004; Vandekerckhove & Tuerlinckx, 2007). Instead, the Bayes Information Criterion (BIC; Schwarz (1978)) was used, which has been applied in previous studies (Milosavljevic *et al.*, 2010; Wenzlaff *et al.*, 2011):

$$BIC = \Delta + d(\log N) \tag{4.1}$$

where Δ is the minimal value of the multinomial log likelihood function after model-fitting (MLF, see above), N is the total number of observations, d is the total number of free parameters (Vandekerckhove & Tuerlinckx, 2007). The Akaike Information Criterion (AIC; (Hurvich & Tsai, 1989)) was also inspected for agreement with the BIC:

$$AICc = \Delta + 2dN/(N - d - 1) \tag{4.2}$$

The model with the lowest BIC (or AIC) value is the best-fitting. To find the best-fitting model across all participants, the difference in BIC scores was computed for each model pair for each participant. Populations statistics were computed with a two-sided paired t-test over the individuals' score differences. If the difference was significant at the 0.05 level (Bonferroni correction for multiple comparisons), then the model with the lower BIC score was the better fit (Milosavljevic *et al.*, 2010).

4.3. Results

4.3.1. Preprocessing: reaction time outliers

Minimum RT cut-offs were calculated for individual participants as the shortest reaction time at which the proportion of correct responses differs from 0.5 (guess level), with the EWMA method (see *Methods*). The resultant EWMA plots indicate the RT at which the proportion of correct responses first escapes chance level with a change in function colour from red to blue (Figure 4.2). Mean minimum RT cut-off across participants was 391 ms ($\pm$ 41 ms). The low variability suggests that this value could be used to inform minimum RT cut-offs in future psychophyics studies with the cylinder task. These results also indicate that human observers can reliably discriminate these cylinder stimuli in under half a second. Across participants, the mean maximum RT cut-off was 2685 ms ($\pm$ 1267 ms) (Figure 4.3). The large variability in RT distributions and maximum cut-off values across participants suggests that these should be individually determined for participants in future studies.

4.3.2. Drift-diffusion model for cylinder discrimination in a reaction time task

4.3.2.1. Psychometric functions and reaction times

Figure 4.4 shows the psychometric functions (PMFs) for individual participants for the cylinder reaction time task, fit with cumulative normal distributions. The sigmoid shape of the psychometric functions, depicting the relationship between cylinder disparity (stimulus strength) and viewers' choices, is in line with previous studies and with the predictions of the diffusion model (Dodd *et al.*, 2001; Mazurek *et al.*, 2003; Palmer *et al.*, 2005; Ratcliff & McKoon, 2008). Mean RTs were generally under 1000-1500 ms and RT distributions showed a strong positive skew (gamma-like distributions) for both correct and error trials for all participants (Figure 4.3), again in line with the predictions of the drift-diffusion model.

Figure 4.2. *(following page)*: Exponentially-weighted moving average (EWMA) charts, for statistically identifying each participant's fast-guess reaction time cut-off. The pink area shows the confidence bound, or upper control limit (UCL), for fast guesses such that P(*correct response*) = 0.5. The point at which the function changes colour from red to blue indicates the reaction time at which the EWMA statistic first exceeds the UCL.

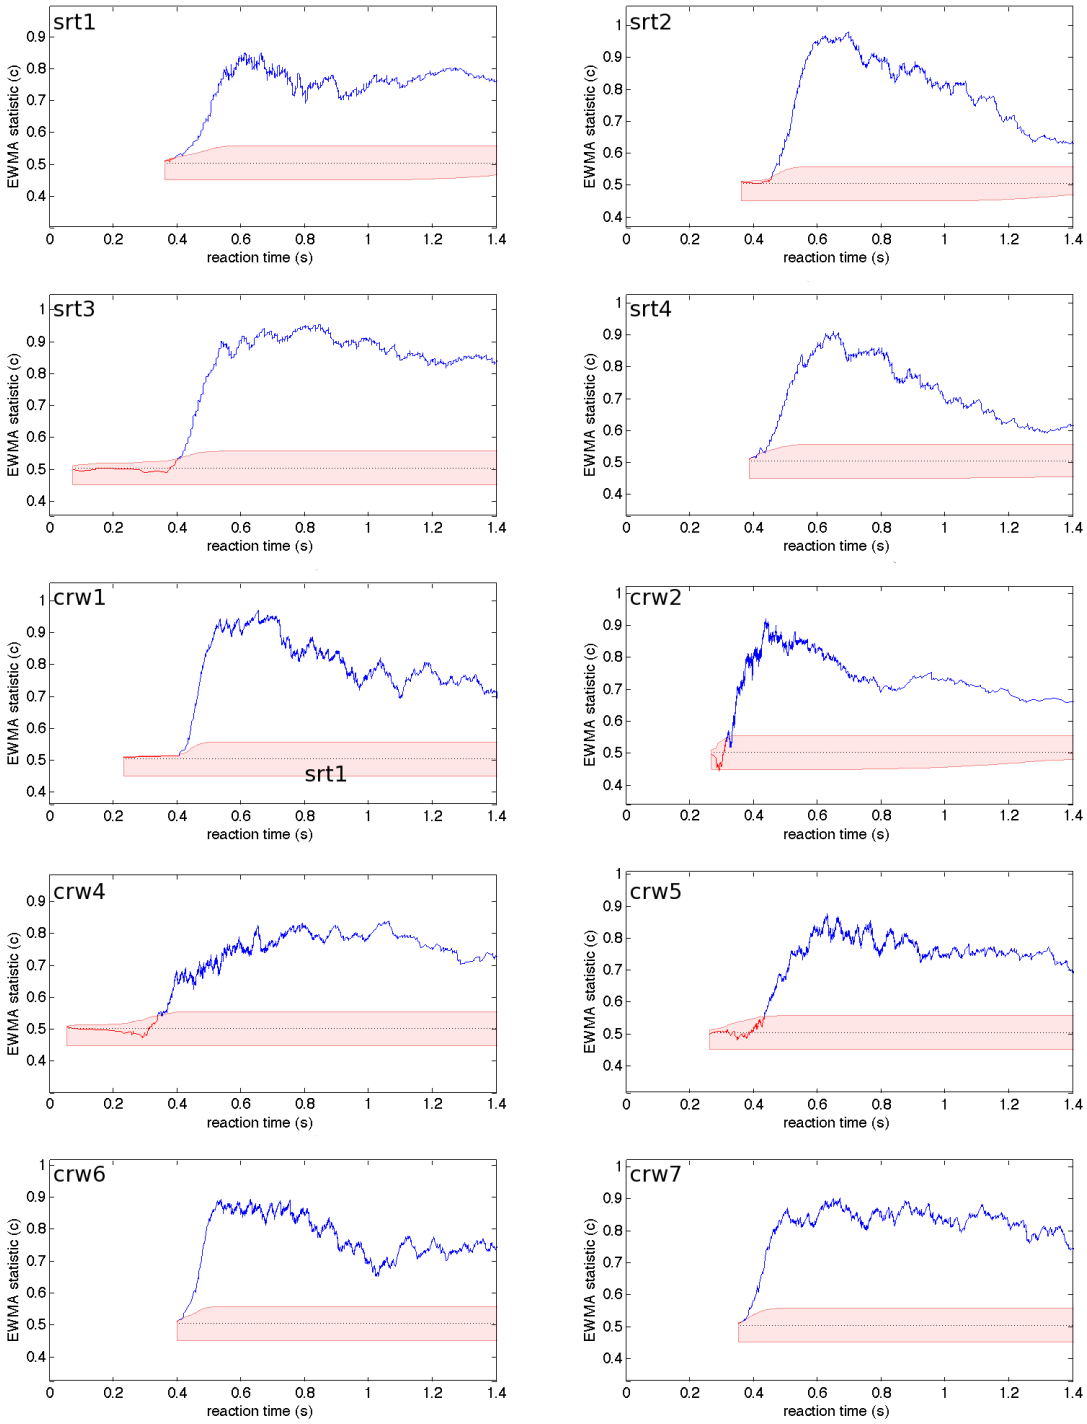

srt1
srt2
srt3
srt4
crw1
srt1
crw2
crw4
crw5
crw6
crw7
EWMA statistic (c)
reaction time (s)

Fig. 4.2

4.3.2.2. Quantile-probability and quantile-quantile plots

Quantile-probability (QP) functions display RT distribution percentiles plotted against the proportion of correct responses for each cylinder disparity, giving a pictorial summary of RT distribution size and shape for CW and CCW responses to different cylinder disparities (Ratcliff & McKoon, 2008). Here, QP plots showed that RT percentiles tended to decrease as cylinder disparity increased, especially for correct trials (Figures 4.5 a&b, 4.6 a&b & 4.7). This effect was small compared to the spread of the RT distributions. Error trials (the percentiles plotted left-of-centre) tended to have longer RTs than correct trials (plotted right-of-centre); this was generally true of all participants. Since the drift-diffusion model predicts longer reaction times for error trials compared to correct trials if the inter-trial variability of drift rate is large relative to the inter-trial variability of the starting point (Ratcliff & McKoon, 2008), this suggests that noise sources for perceptual judgements in the cylinder task are dominated by inter-trial variability of drift rate rather than accumulation start point. This is similar to the findings for human observers performing the motion task (Ratcliff & McKoon, 2008).

Quantile-quantile (QQ) plots show RT percentiles for each cylinder disparity condition plotted against RT percentiles from the condition with the second-largest disparity. The drift-diffusion model predicts linear QQ functions because changes in model parameter values do not predict significant changes to the overall shape of the RT distributions (Ratcliff & McKoon, 2008). Here, QQ relationships between different cylinder disparity conditions were generally linear (Figures 4.5 and 4.6). This was verified with a linear correlation analysis at the 0.05 significance level (Bonferroni-correction for n=20 comparisons; $p<0.0025$). All twenty pairwise QQ comparisons showed a significant linear relationship, except for the largest-disparity CW error condition and the two largest-disparity CCW error conditions in the *crw* subject group which did not reach significance at the corrected level ($p=0.01$, $p=0.0121$, $p=0.0044$ respectively). However, RT distribution percentile estimates for these conditions are unreliable because of the low numbers of error trials at these large cylinder disparities, and therefore may not be expected to follow model predictions (Ratcliff & McKoon, 2008). Overall, QP and QQ plots show that viewers' decisions in the cylinder task fitted the predictions of the drift-diffusion process.

Figure 4.3. *(following page)*: Histograms of reaction times (RTs) (*in red, upper panels*) and log(RTs) (*in black, lower panels*) with normal distribution fit *red lines*), for each participant. The vertical lines indicate maximum RT cut-offs, calculated separately for each individual participant, 2.5 standard deviations above the mean of the log(RT) distribution.

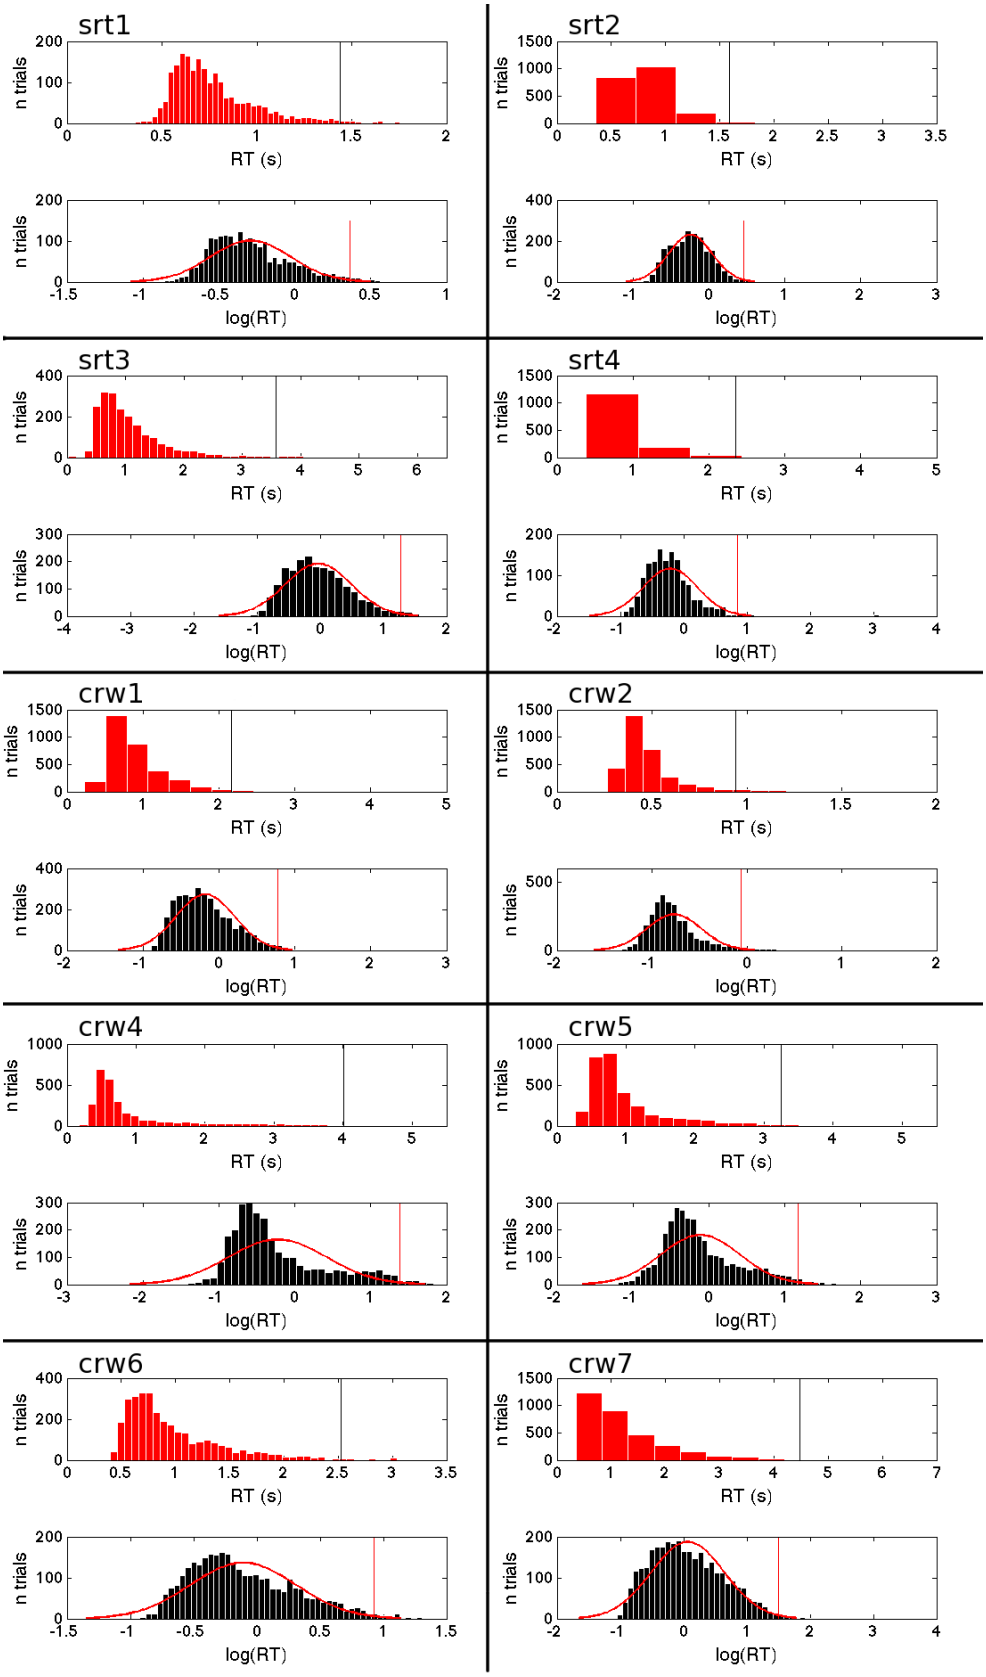

srt1
srt2
srt3
srt4
crw1
crw2
crw4
crw5
crw6
crw7
n trials
RT (s)
log(RT)

Fig. 4.3

4.3.2.3. Effect of cylinder disparity on drift-diffusion model parameters

I investigated whether the effect of cylinder disparity (stimulus strength) on choices and RTs could be explained purely by changes in the drift rate parameter of the model and whether the relationship between stimulus strength and drift rate was linear. Four models were sequentially applied to the data and the Bayesian Information Criterion (BIC) was used to infer the best model for each individual participant and rank the models across participants (see *Methods* for details). For 10/10 participants, the best-fitting model was one in which only the drift rate parameter was varied varied by cylinder disparity: Model 2 ('linear drift rate') or Model 3 ('free drift rate') (Table 4.2). Model 4 ('unconstrained parameters') performed significantly better than Model 1 ('no effects') ($p < 0.0056$). Model 2 ('linear drift rate') and Model 3 ('free drift rate') both performed significantly better than Model 4 ($p \ll 0.001$ in both cases). Model 2 had a better fit than Model 3, but this difference was not significant ($p < 0.59$). This shows that changes to the drift rate parameter alone are sufficient to explain the effect of stimulus strength on decision-making in the cylinder discrimination task, as expected from the standard interpretation of the drift-diffusion process when applied to perceptual decisions (Gold & Shadlen, 2007; Voss *et al.*, 2004; Hanks *et al.*, 2006).

Since Model 2 ('linear regression') was the optimal model for 6/10 participants (Table 4.2), a linear relationship between cylinder disparity and drift rate was the most parsimonious explanation of the data for the majority of participants (Figure 4.8). Moreover, for participants for whom Model 3 ('free drift') was the best model, a linear regression analysis showed a significant linear relationship between drift rate estimates and cylinder disparity in all cases ($p \leq 0.0002$). Therefore, a linear approximation for the relationship between cylinder disparity and drift rate parameter was used in all subsequent model-fitting tests, to minimize time and complexity of model fitting since linear approximations use fewer parameters.

4.3.3. Effects of unbalanced reward conditions on perceptual decisions

4.3.3.1. Behavioural results: psychometric functions & RT distributions

In the unbalanced reward conditions (HL & LH), participants' PMFs were horizontally shifted towards the higher-rewarded direction, showing that the experimental manipulation of unbalanced reward was effective in biasing responses toward the direction with larger number of available points (Figure 4.9). Mean RTs for correct choices for a particular cylinder disparity were affected by the reward condition of the trial in a predictable manner: responses were faster

participant:	Model 1	Model 2	Model 3	Model 4
srt1	10,791	**10,029**	10,066	10,268
srt2	11,130	9,622	**9,609**	9,749
srt3	10,732	9,356	**9,343**	9,549
srt4	7,343	**6,721**	6,754	6,976
crw1	5,805	4,855	**4,783**	5,012
crw2	5,559	**4,906**	4,931	5,196
crw4	4,868	**4,724**	4,754	5,059
crw5	5,444	**4,888**	4,914	5,190
crw6	5,665	**5,145**	5,161	5,433
crw7	5,576	4,875	**4,868**	5,131

Table 4.2.: Bayesian Information Criteria (BIC) values of model fits for effect of cylinder disparity on the drift-diffusion process. For each participant, the best model (lower criterion value) is indicated in bold. Model 1: no effect; Model 2: linear regression; Model 3: free drift; Model 4:unconstrained.

in the direction of greater available points, and slower in the direction of fewer available points (Figure 4.10). Mean RTs in the unbiased condition (LL condition) tended to lie between those of the biased reward conditions. QP plots confirmed that the effects of reward manipulation were to increase the proportion of choices towards the more highly-reward stimuli: when the location on the x-axis of each given colour column (cylinder disparity condition) is compared between the upper and lower panels, the locations are shifted towards having a greater response proportion in the more highly-rewarded direction (Figure 4.11). This QP plot modulation is similar to that presented by Ratcliff & McKoon (2008, their Figure 10), where human observers' choices in the motion task were biased by manipulation of likelihoods of different stimuli being presented.

These effects of reward on choices and reaction times must be be underpinned by changes in parameter values of the underlying decision process. Biasing the drift rate parameter towards the higher-rewarded direction, or decreasing the distance between the starting point and higher-rewarded decision-bound (Figure 1.4), or both, could result in the observed reward bias. The effects of these parameters have different signatures on the RT distributions and therefore can be distinguished by RT percentile-based drift-diffusion model-fitting. Based upon the results of the previous two Chapters, I predicted that reward bias acts in part through an effect on the drift rate parameter of the diffusion model.

4.3.3.2. Effects of unbalanced rewards on drift-diffusion model parameters

Four models were fit to the datasets obtained in the unbalanced reward conditions HL and LH (see *Methods* for details). The Bayesian Information Criterion (BIC) was used to test the specific prediction that reward condition affects the drift rate parameter of the drift-diffusion process, in addition to effects on the starting point. The BIC values for Model 2 ('starting point effects')

participant:	model with lowest BIC score:	model with lowest AICc score:
crw1	drift effect (model 3)	unconstrained (model 4)
crw2	drift effect (model 3)	drift effect (model 3)
crw4	drift effect (model 3)	drift effect (model 3)
crw5	starting point effect only (model 2)	drift effect (model 3)
crw6	drift effect (model 3)	drift effect (model 3)
crw7	drift effect (model 3)	drift effect (model 3)

Table 4.3.: Best model fit for each participant for the effect of reward bias direction in the unbalanced reward conditions (HL vs. LH). The best model fit by both BIC and AICc is reported.

and Model 3 ('drift rate effects') was compared across participants (two-sided paired t-test). BIC comparisons across the other models fits are also reported. The Akaike Information Criterion was also inspected to verify the results.

Model 3 ('drift rate effects') was the best-fitting model for 5/6 participants and was better-fitting than Model 2 across participants ($p < 0.047$) (Table 4.3). This finding was corroborated by the AIC measure, for which Model 3 was best-fitting in 5/6 participants (Table 4.3). These results demonstrate that effects of reward on starting point offset are not sufficient to explain choices biases in this task; a change to the drift rates by reward is also necessary to best-fit the drift-diffusion model to the data.

In additional comparisons, Model 1 ('no effects') had a slightly lower BIC score (better fit) than Model 4 ('unconstrained model'), but this was not significant. Model 2 ('starting point effects') and Model 3 both performed significantly better than Models 1 & 4 ($p \ll 0.05$ in all cases). This suggests that, of all models tested, Model 3 ('drift rate effects') best explained the data.

For all participants, the effect of reward on the starting point was to shift it closer to the decision bound corresponding to the higher-rewarded choice, as expected (Figure 4.12). In 5/6 participants, drift rate estimates were biased towards the higher-reward response direction (Figure 4.13, *orange and blue lines*). This effect was seen both when the cylinder disparity-drift rate relationship was linearly constrained (Model 3) (Figure 4.13, *smooth lines*) and when drift rates, along with all other parameters, were allowed to vary entirely (Model 4) (*dashed lines*). This shows that when responses in the cylinder task are biased towards one perceptual outcome by unequal pay-offs, in the majority of participants the drift rate of the underlying diffusion process is biased towards the higher-reward cylinder rotation direction.

4.3.3.3. Strategies for reward maximization

All participants modulated their responses in the 'unbalanced' reward conditions by biasing their choices towards CW or CCW responses according to the higher-rewarded direction. How does their strategy compare to other possible strategies that could have been implemented in this task? To what extent was their strategy optimal?

A participant's goal in this task is to maximize the total number of points that they earned, because points were converted to money at the end of the experiment. Figure 4.14 shows some examples of strategies that could be implemented, and the total number of points expected to be earned (*E(rew)*) from each strategy. Since points are only given for correct choices, a perfect observer would respond correctly to all non-zero disparity stimuli (regardless of reward condition), whilst their responses to zero-disparity stimuli would always be in the more highly-rewarded direction (Figure 4.14a, *blue, orange & grey lines*). A perfect observer would thus collect 98.3% of the available points. At the opposite extreme, an observer whose strategy was to simply respond randomly in all trials would collect 50.0% of available points (Figure 4.14a, *black dashed line*). Alternatively, an observer whose strategy was to only utilize reward cue information, and ignore the visual stimulus, would respond CW in all HL trials, CCW in all LH trials, and arbitrarily in LL trials (Figure 4.14 b). This observer would collect 85% of available points, a large proportion of the total rewards because the difference between the higher- and lower-rewarded directions in the biased reward conditions is very large (8:1).

In reality, the proportion of correct choices that a participant can make is limited by their sensitivity, or threshold, with respect to the presented cylinder disparities. The measure of each participant's sensitivity was taken to be the average of the standard deviations of the fitted psychometric functions (PMFs) in the three reward conditions (LL, HL, LH, Figure 4.9). Viewers may additionally have an intrinsic 'bias' towards one perceptual outcome. The measure of each participant's intrinsic bias, if any, was taken to be the mean of the PMF in the LL condition (i.e. the size of its horizontal shift away from the (0,0.5) coordinate, in units of cylinder stimulus disparity). The percentage of total available points that would be collected simply according to the individual's disparity threshold and intrinsic bias towards a particular perceptual choice (without reward-related response bias) was calculated for each individual. On average, this value was 75.6% (Figure 4.14 c, *black scatter point*). The amount of points each participant could have collected had they implemented the optimal reward response bias strategy was on average 85.6% (Figure 4.14 c, *red scatter point*). The actual percentage of total available points collected

by each participant in this experiment was calculated; on average, this value was 77.3% (Figure 4.14 c, *blue scatter point*). Therefore, although participants did bias their responses towards the more highly-rewarded direction according to the reward condition (see shifts in Figure 4.9), these shifts were small compared to the optimum strategy, and only earned them an additional $\approx$ 2% of total available points.

I calculated, for each participant, the total amount of points gained, as a percentage of the total that would have been gained under the optimal strategy. Then, I calculated, for each participant, the size of their actual reward bias (the normalized horizontal distance between the psychometric functions in the LH and HL conditions, see Figure 4.9), as a percentage of the reward bias in the optimal strategy. For all participants, the actual reward bias was less than 50% of the optimal strategy (Figure 4.14). Unsurprisingly, the closer their reward bias was to that of the optimum strategy, the closer the amount of points they collected was to the total amount available under the optimal strategy. All participants received fewer points than they would have had they implemented the optimum strategy; indeed they received fewer points than they would have had they completely ignored the visual stimulus and responded purely according to the points cue direction (85.0%).

4.3.3.4. Effect of reward cue duration on perceptual decisions

In each trial, the temporal duration of the presented points cue was randomly selected from a uniform distribution with range 850 - 1450 ms. The effect of the duration d of the points cue on the responses to the subsequently-presented cylinder stimulus was investigated. Cue durations were grouped into three conditions: 1) *short duration*: $850 \leq d < 1050$ ms (mean 950 ms); 2) *medium duration*: $1050 \leq d < 1250$ ms (mean 1150 ms); 3) *long duration*: $1250 \leq d < 1450$ ms (mean 1350 ms). For each group, PMFs were generated for each of the LL, HL and LH reward conditions. Since only a small number of trials remained after separation of the data into cue duration groups, datasets was combined across participants to generate the PMFs. The size of the reward bias for each duration group was measured as the absolute difference between the HL and LH psychometric function means. It was expected that the size of the reward bias would increase with longer exposure to the points cue, as this might give the points cue more influence over the upcoming decision. Contrary to expectation, the size of the reward bias decreased as the points cue duration increases (Figure 4.15 d, *black line*). Additionally, the slope of the average PMF decreased with increasing cue duration, suggesting that responses became more accurate with greater cue duration times (Figure 4.15 d, *red line*). Reward bias still de-

creased with increasing cue duration even when the reward bias was normalized by division by the average slope (Figure 4.15 d, *green line*), showing that changes in accuracy cannot account for the decrease in bias.

4.3.4. Effects of balanced reward sizes on perceptual decisions

4.3.4.1. Psychometric functions and reaction times

For 3 participants (*crw2, crw4 & crw5*), HH reward conditions, in which 8 points were given for a correct choice in either direction, were interleaved with the LL, HL and LH reward conditions (see Table 7.1). Comparing the 'balanced' reward conditions, LL and HH, could indicate how the total points available for a correct choice affects decision-making in this cylinder task. In this case, neither one nor the other direction was more highly-rewarded so no bias was expected towards a particular response direction, but overall accuracy or RT distributions could be affected, relating to changes in model parameters. This comparison relates to the macaque microstimulation study in Chapters 2 & 3, in which the overall amount of reward available for a correct choice was greater in some trials than others.

For one participant (*crw2*), there was a steeper slope in the PMF, and a longer mean RT, for the HH condition compared to the LL condition, but this was not found in the other two participants (Figure 4.16). Overall, there was no consistent pattern of effects between the LL and HH conditions across participants, and therefore no evidence for differences in task accuracy with different reward levels.

4.3.4.2. Effects on drift-diffusion model parameters

Four models were compared to investigate the effects of balanced reward conditions. The Bayesian Information Criterion (BIC) was used to infer the best model fit for each individual participant and to rank the models across participants (see *Methods* for details).

Model 1 ('no effects') was the best-fitting model in 2/3 participants (Table 4.4). When all possible pairwise comparisons were made, Model 1 had the lowest overall BIC score, but was not significantly different from Model 2 ('bound effects') or Model 3 ('drift effects') ($p > 0.4$ in all cases). Models 1, 2 and 3 were all significantly better-fitting than Model 4 ('unconstrained model') ($p \ll 0.001$ in all cases). These results do not support any effect of balanced reward size on model parameters. When parameter estimates were inspected, there was no consistent effect on drift

participant:	Model 1	Model 2	Model 3	Model 4
crw2	9,235	**9,205**	9,219	9,894
crw4	**9,239**	9,253	9,267	9,957
crw5	**9,729**	9,749	9,764	10,437

Table 4.4.: Bayesian Information Criteria (BIC) values of model fits for effect of balanced reward condition (HH versus LL) on the drift-diffusion process. For each participant, the best model (lower criterion value) is indicated in bold. Model 1: no effects; Model 2: bound effects only; Model 3: additionally drift effects; Model 4: unconstrained.

rate by balanced reward condition across participants (Figure 4.17 a-c). However, for all participants, the distance between the two response bounds increased in the HH condition compared to the LL condition (Figure 4.17 d), although this did not reach statistical significance in model comparison above.

4.3.5. Summary of results

In summary, the effect of cylinder disparity (stimulus strength) on decisions was fully explained by changes solely to the drift rate of the drift-diffusion process. The relationship between cylinder disparity and drift rate was well-approximated by a linear function. When viewers' responses were biased towards one cylinder direction or another by unequal pay-offs, an effect of reward condition on drift rate was necessary to explain the data, in addition to a decrease in the distance between the accumulation starting point and the higher-rewarded bound. The drift rate was biased towards the direction of rotation of the more highly-rewarded cylinder in the majority of subjects. These results demonstrate that perceptual decisions in the cylinder task are well-explained by a drift-diffusion process and that reward bias can affect both the drift-rate of the decision process, thought to be based upon stimulus-dependent cortical activity, as well as the starting point offset, thought to be instantiated in later, sensorimotor areas.

4.4. Discussion

4.4.1. The cylinder discrimination task follows a drift-diffusion process

I evaluated whether the drift diffusion model was an appropriate description of perceptual decision-making in the cylinder task. The general class of sequential sampling models, of which the drift-diffusion model is a sub-class, have been previously applied to a wide range of decision tasks (Ratcliff, 1978; Palmer *et al.*, 2005; Ratcliff & McKoon, 2008; Milosavljevic *et al.*, 2010; Summerfield & Koechlin, 2010; Wenzlaff *et al.*, 2011), particularly the random-dot motion discrim-

ination task, where there is striking evidence for specific neural correlates of different model parameters (Roitman & Shadlen, 2002; Mazurek *et al.*, 2003; Hanks *et al.*, 2006; Kiani *et al.*, 2008; Rorie *et al.*, 2010; Churchland *et al.*, 2011). However, it was not certain *a priori* that the drift-diffusion model would also be applicable to the cylinder task used in this study. Firstly, in most previous experiments, noise was explicitly added to the visual stimulus to impair detection and discrimination. For example, in the motion task, randomly-moving dots mask the coherent motion direction (Britten *et al.*, 1992); in face-house discrimination and oriented gabor detection tasks, filtered white noise is added to the stimuli to increase task difficulty (Summerfield & Koechlin, 2010; Wenzlaff *et al.*, 2011). In the cylinder task no such is added, suggesting that the majority of noise arises internally, presumably from motion- and disparity-detection mechanisms in the visual system. This could affect the nature of the decision-making mechanism. Additionally, this is the first study to directly measure reaction times (RTs) for correct and error responses to the cylinder stimuli, so it was not known whether these would follow a distribution predicted by the diffusion model, nor whether mean RTs would fall within the range to which the drift-diffusion model is applicable (Ratcliff & McKoon, 2008). Finally, in the cylinder task, the zero-disparity condition gives rise to a coherent percept of a 3-dimenstional cylinder rotating in a particular direction. The direction of rotation is perceptually bistable, but not ambiguous; indeed the percept cannot be distinguished readily from the unambiguous case (Nawrot & Blake, 1993). This contrasts with most other perceptual tasks, in which the condition where there is no evidence towards one direction or the other is perceptually ambiguous.

The results showed that choices and RTs in the cylinder task contain the hallmarks of a drift-diffusion process. Mean RTs were under 1000-1500 ms, a necessary condition for the diffusion model as it should be applied only to relatively fast two-choice decisions that are likely to have a single-stage decision process (Ratcliff & McKoon, 2008). Secondly, both correct and error trial RT distributions were positively skewed, like gamma distributions, as predicted by the drift-diffusion process. For the random dot motion task, human participant RT distributions are also positively skewed (Ratcliff & McKoon, 2008). Quantile-quantile (QQ) plots for the cylinder task showed a significantly linear relationship between percentiles of RT distributions from different experimental conditions, indicating that the RT distribution shapes were similar between experimental conditions, even though model parameter estimates may change. This is again consistent with the predictions of the diffusion model (Ratcliff & McKoon, 2008). Finally, the percentile probability (QP) plots showed that, generally, error responses were longer than correct responses, indicating a greater inter-trial drift rate variability relative to starting-point offset variability. This was also found when human psychophysics was evaluated in the random-dot

motion task (Ratcliff & McKoon, 2008).

This is the first application of the drift diffusion model to the cylinder task and indeed to any perceptual decision task in which the zero-evidence case (here the zero disparity condition) gives rise to a coherent percept. There were many similarities between the results from human psychophysics in this cylinder task and in the commonly-used motion task. The results strongly argue that the drift diffusion model of perceptual decision-making is at least as applicable to perceptual decision-making in the cylinder task as for the motion task and the results confirm that the application of a simplified drift-diffusion model in the previous Chapter (microstimulation and reward in the cylinder task for monkeys) was appropriate.

4.4.2. Cylinder disparity affects only the drift rate of the diffusion process

The the drift rate parameter of the drift-diffusion process represents the strength of evidence towards one choice or the other (Voss *et al.*, 2004; Palmer *et al.*, 2005). Electrophysiological recording and microstimulation studies in monkeys performing the motion task have demonstrated that the drift rate parameter corresponds closely to the integral of the sensory evidence in V5/MT neurons (Mazurek *et al.*, 2003; Huk & Shadlen, 2005; Hanks *et al.*, 2006). Similarly, in human neuroimaging studies, when perceptual decisions are modeled by the diffusion process, drift rate parameters correlate with activity in visual cortical areas, for example the fusiform face area (FFA) for face stimuli (Heekeren *et al.*, 2004; Wenzlaff *et al.*, 2011) and lateral occipital area (LO) for gabor contrasts (Summerfield & Koechlin, 2010). Here, I found that the main effect of cylinder disparity was on the drift rate parameter of the diffusion process. When drift rates were plotted against disparity, the relationship was linear (verified by linear regression analysis). Therefore, future studies can confidently approximate a linear relationship between cylinder disparity and drift-diffusion model drift rate. Constraining the effect of disparity on drift to a linear function simplifies model specification and means that only two values (the linear function's intercept and gradient) need to be estimated for model fitting, decreasing computation time and parameter space complexity. These results also confirm that the linear assumption in the simplified drift-diffusion model applied in the previous Chapter was appropriate and valid.

4.4.3. Reward bias affects both drift rate and distance-to-bound

I investigated how biasing viewers' responses with unequal pay-offs affected parameters of the drift-diffusion model, revealing the possible implications for underlying brain processes. Reward manipulations affected both the choices and reaction times of perceptual judgements. A model which contained changes to the drift rate parameter, as well as the starting point of the diffusion process, was significantly best-fitting across participants. This result was supported by two measures of model fit (Bayesian Information Criterion (BIC) and Aikake Information Criterion (AIC)). In the majority of participants (5/6), drift rate was biased towards the higher-rewarded cylinder direction. Since diffusion model drift rate is closely related to stimulus strength and activity in visual areas, this strongly suggests that reward condition affects the sensory representations of the visual stimulus. Specifically, this effect corresponds to an increase in the gain, or sensitivity, of the representation of visual evidence for the more highly rewarded cylinder rotation direction. This is in line with the effect of overall large reward size to increase the sensitivity of visual representations of the stimulus generally, presented in the previous Chapters 2 & 3 with a microstimulation manipulation with non-human primates.

The only participant in whom drift rates were biased towards the lower-reward cylinder direction was *crw2*, who was the author and the only participant who was not naive to the goals of the study. For this participant, no monetary remuneration was given for task participation. It could be the case that participation in the task in good faith and for real remuneration is necessary for true effects of reward bias on perceptual decision-making to be present.

There are a number of possible reasons why a significant effect of reward on drift rate was found in this study, whilst the majority of previous studies found only effects on the starting point and decision bounds (with the exception of Liston & Stone (2008) and Voss *et al.* (2008); see Chapter 1). The present study fitted a fully-parameterized drift-diffusion model to dataests of individual participants, according to methods verified in previous literature (Vandekerckhove & Tuerlinckx, 2007, 2008; Milosavljevic *et al.*, 2010; Wenzlaff *et al.*, 2011). This contrasts with most previous studies in which variations on the sequential sampling process were fit with a custom approach (Summerfield & Koechlin, 2010; Rorie *et al.*, 2010) or fit to averaged datasets (Diederich & Busemeyer, 2006; Diederich, 2008). In the present study, both correct and error 0.01, 0.02 and 0.05 RT distribution percentiles were used for model-fitting, in addition to the usual 0.1, 0.3, 0.5 (median), 0.7 and 0.9 percentiles. As explained above, changes in drift rate and distance-to-bound parameters predict different relative effects on the leading edge versus

the tail of the distribution, even when the overall bias effect on choices is the same (Ratcliff & McKoon, 2008). Densely sampling the leading edge of the RT distribution may have given this study the sensitivity to drift rate effects that other studies lacked (e.g. Mulder *et al.* (2012)). Furthermore, I verified *a priori* the appropriateness of modeling drift rate as a linear function of stimulus strength, and nine different stimulus strength conditions were presented (a larger range than all previous studies). Finally, outlier RT trials were removed from the analysis in a principled way, eliminating short RT 'fast guess' trials in which stimulus evidence was not integrated at all, which might have contaminated model fitting.

4.4.4. The 'under-biasing' phenomenon

In psychophysics studies that manipulate reward, human participants tend to significantly under-bias their responses (Gao *et al.*, 2011; Diederich, 2008). It was for this reason that a large difference between the 'high' and 'low' rewards was chosen in this study (8:1 points ratio). I found that the size of the reward bias, measured as the distance between the means of the psychometric functions for HL and LH conditions, was less than 50% of the optimal strategy for all participants. Indeed, participants would have collected more points just by responding according to the more highly-rewarded direction and ignoring the cylinder stimulus completely. One hypothesis to explain the phenomenon of under-biasing is that participants place an intrinsic value C on being correct, independent of the rewards available in the experiment, thus changing the participantive reward ratio to $(R_H + C/R_L + C)$ which is always less than R_H/R_L as long as C is positive (Gao *et al.*, 2011; Maddox & Bohil, 2003). In this study, instructions required participants to be as 'speedy and accurate' as they could in their responses (Appendix B). It is an interesting question to what extent viewers would place intrinsic value C on 'correctly' observing the world and still under-bias their choices in more 'natural' decision-making contexts i.e. when away from the instructions and motivations present during a visual psychophysics experiment.

4.4.5. Reward cue duration decreases choice bias

I investigated whether the duration of the points cue affected the response bias to subsequent cylinder. My hypothesis was that increased cue duration would result in an increased reward bias in the subsequent response, perhaps due to increased integration of the reward-biasing influence over time. Contrary to expectations, the size of response bias decreased as cue duration increased. It could have been the case that responses generally increased in accuracy with in-

creased cue duration (for example due to stimulus onset expectation/preparatory effects) and that reward bias changed relative to this. However, when accuracy thresholds were measured, it was found that although accuracy of responses did increase as cue duration increased, the reward bias normalized by threshold still decreased as a function of points cue duration. The small differences in points cue durations limit conclusions drawn from this observation, but this warrants further exploration. For example, the increased duration of the reward cue may give increased preparation time, affect temporal orienting, or allow bias-compensatory mechanisms to be invoked more easily.

4.4.6. Balanced reward conditions

In three participants, decison-making was compared in two balanced reward conditions: LL in which correct choices earned 1 point regardless of response direction, and HH in which correct choices earned 8 points regardless of response direction. This comparison relates most closely to that of the macaque microstimulation study presented in Chapters 2 & 3, the results of which indicated an effect of reward size on sensory evidence. In the present study, no consistent effect of overall reward size was seen in the psychometric functions (PMFs) nor in the mean RTs. Drift-diffusion model comparisons did not show any significant effect of reward size on any parameter. The model of no effects by reward condition was the best-fitting across participants, but could not be statistically distinguished from those with changes to bound distance and drift rate. There are a number of possible reasons for this lack of effect. Firstly, there may genuinely be no effect of this reward manipulation on human decision-making. Secondly, the small number of participants means that statistical model comparisons are under-powered. Alternatively, since the balanced reward conditions were pseudo-randomly interleaved with unbalanced conditions, it could be that it was too demanding for participants to change strategies on a trial-by-trial basis as well as performing the demanding cylinder task. In the end, strategies for the unbalanced reward conditions may have been favoured. An additional study with more participants and fewer total reward conditions is needed to fully investigate this effect.

Although there were no significant effects on model parameters by reward size, there was a consistent trend across participants for the distance between the decision bounds to increase in the high reward condition (HH) compared to the low (LL) (Figure 1.4). Previous studies on reward and perceptual decisions invariably focused on unequal payoff manipulations; the investigation into the effect of available reward size might correspond better to previous studies that manipulated the speed/accuracy trade-off and also reported that in the increased accuracy

condition, the distance between the decision bounds is increased (Palmer *et al.*, 2005; Ratcliff & McKoon, 2008; Bogacz *et al.*, 2010; Wenzlaff *et al.*, 2011; Balci *et al.*, 2011). If we imagine that the participants are motivated to be more accurate in the HH condition compared to LL since more points are at stake, then this trend is in line with the previous literature. On the other hand, this finding is in the direction opposite to the effect found for macaques in Chapter 3, where the best-fitting diffusion model showed a decrease in the distance between decision bounds in the higher-rewarded trials. The findings can be reconciled when we consider that there might be unknown differences in temporal utility functions between macaques and humans in these experiments. In Chapter 3, animals worked for a fluid droplet at the end of each trial to satiate an immediate thirst which may have motivated them to collect rewards as quickly as possible. The speed of response should presumably be maintained or increased when large reward was available, even as the animals maintained or even improved their response accuracy. This could be achieved by decreasing the distance between the starting points of the diffusion process and the decision bounds, whilst concurrently increasing sensitivity to the stimulus by increasing drift rate gain, as seen in the results of Chapters 2 & 3. On the other hand, human volunteers worked to earn points that were indirectly related to a monetary remuneration given long after the end of the experiment. They could not bring about a faster reception of the reward by changing response speed, and for this reason may not have implemented the same strategies as macaques even though reward conditions were ostensibly the same. Human psychophysics experiments designed to exactly replicate the time-dependent value of reward as in the monkey studies are needed to explore this issue further.

4.4.7. Future work

This study confirmed the drift-diffusion model is an appropriate model of perceptual decision-making for the cylinder task. The next step would to measure underlying neural activity whilst participants perform the cylinder task under different reward conditions and relate this to the findings from model fitting reported here. The use of an experimental method with high temporal resolution would be particularly appropriate, to closely observe accumulation of evidence over time and to pinpoint the timing of brain activity changes according to reward condition. In humans, magnetoencephalography (MEG) is a non-invasive method with the appropriate temporal resolution. In animals, such as macaque monkeys, direct electrophysiological recordings from candidate brain regions such as area V5/MT (Dodd *et al.*, 2001) and LIP (Roitman & Shadlen, 2002) would be ideal to investigate these questions.

Figure 4.4. *(following page)*: Psychometric functions for cylinder discrimination choices, all participants. For participants *srt1-4*, no points or performance feedback was involved. For *crw* participants, data are taken from the balanced low reward condition (LL) in which 1 point is received for a correct choice regardless of choice direction. Choice proportion varies monotonically as a function of cylinder disparity; positive cylinder disparities signify CW rotation.

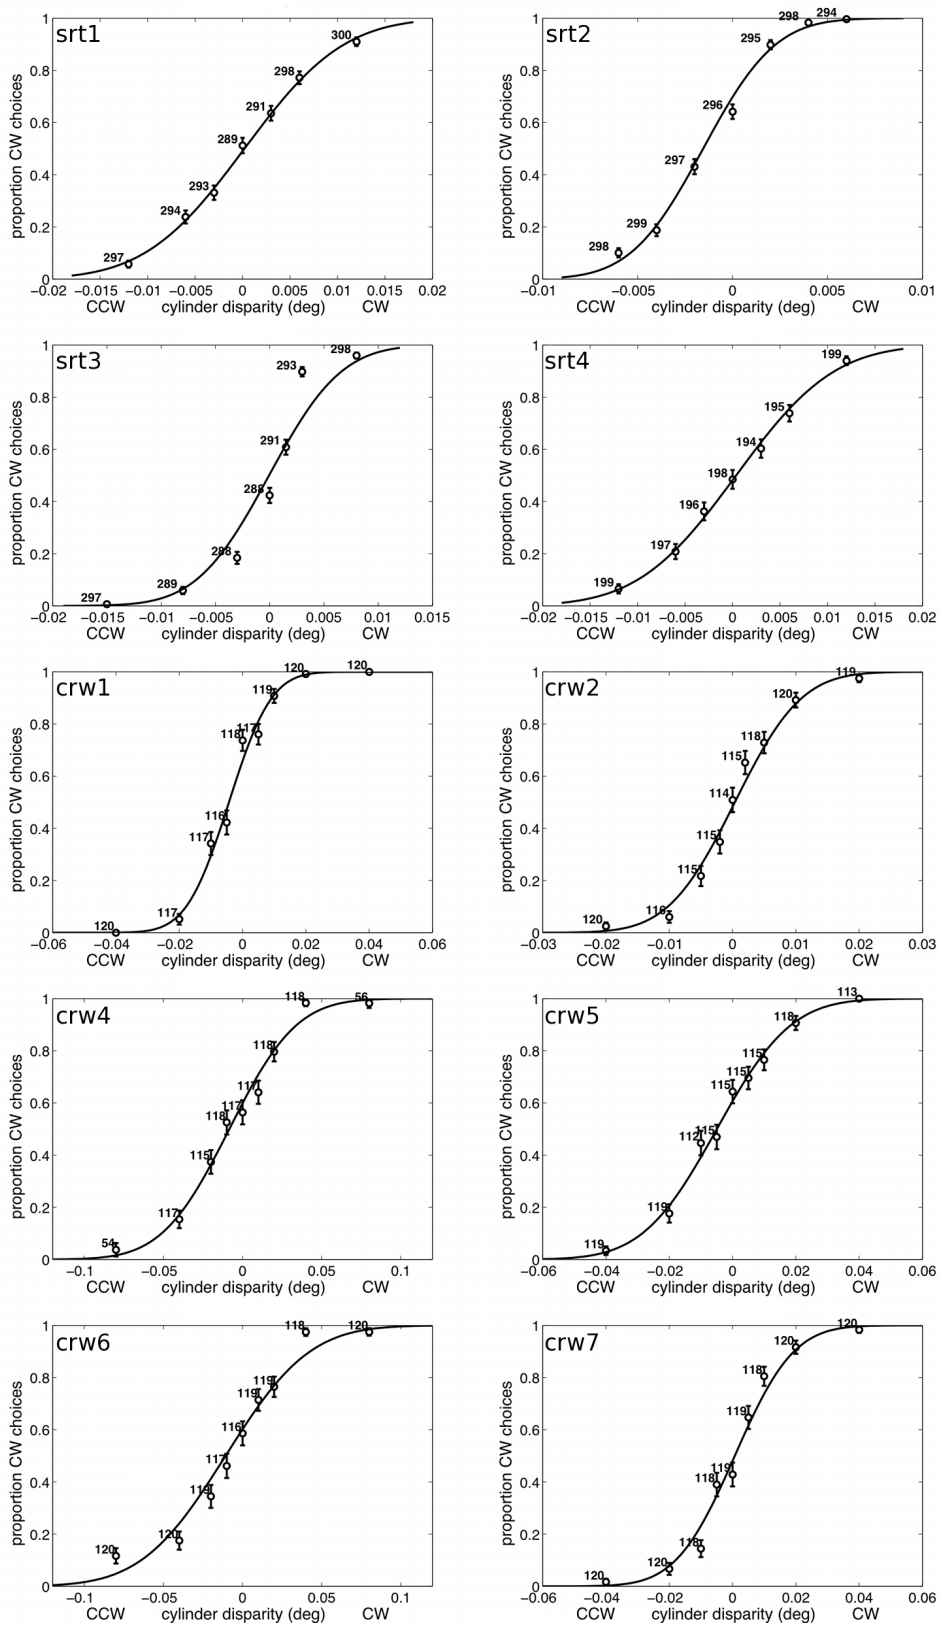

srt1
srt2
srt3
srt4
crw1
crw2
crw4
crw5
crw6
crw7
proportion CW choices
CCW
cylinder disparity (deg)
CW

Fig. 4.4

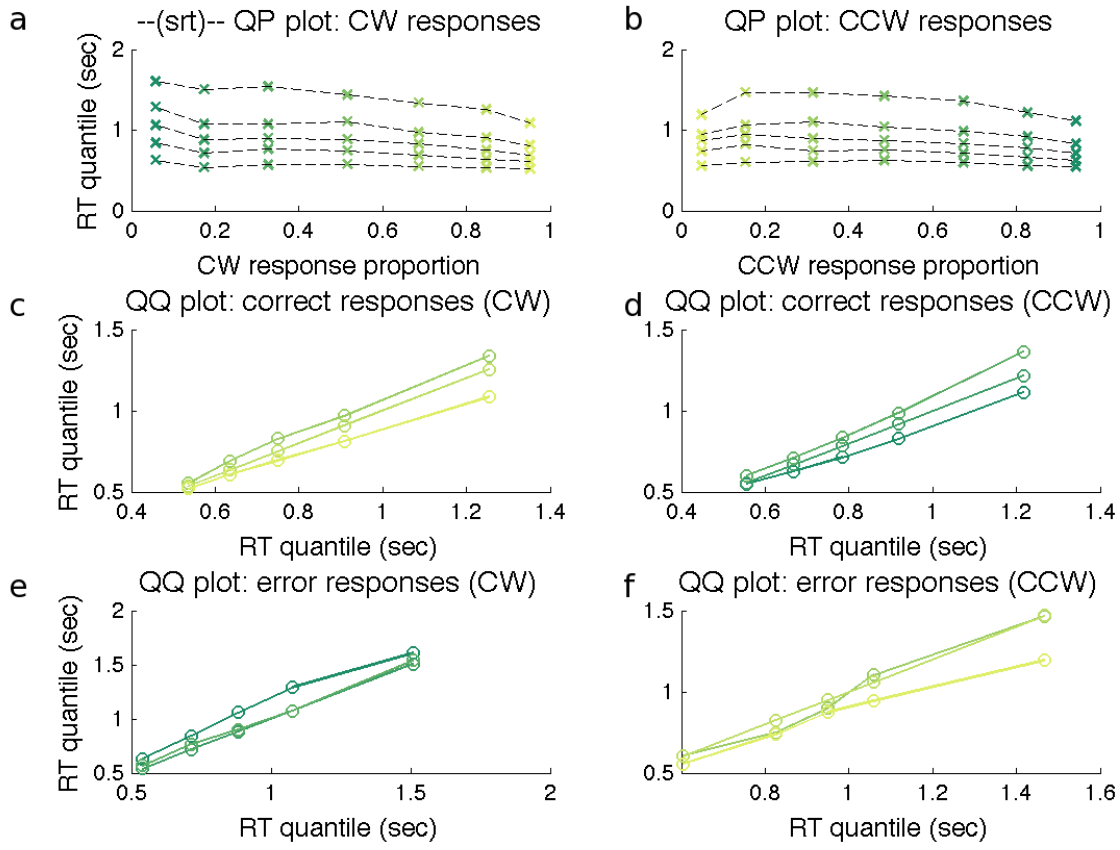


Figure 4.5.: **a, b** percentile probability (QP) twin plots, averaged over *srt* participants. *Dashed lines* in QP plots connect the RTs that correspond to the same percentile (0.1, 0.3, 0.5 (median), 0.7 and 0.9) across disparity conditions. Scatter point colours indicate the stimulus disparity conditions and range in 7 steps from *turquoise* = maximum CCW (negative) disparity, to *bright yellow* = maximum CW (positive) disparity. **c - f** quantile-quantile (QQ) plots, averaged over *srt* participants, displayed separately for correct and error trials. QP and QQ plots for the two response directions (CW and CCW) have been drawn separately side-by-side to lay clear any differences in reaction time distribution by response direction

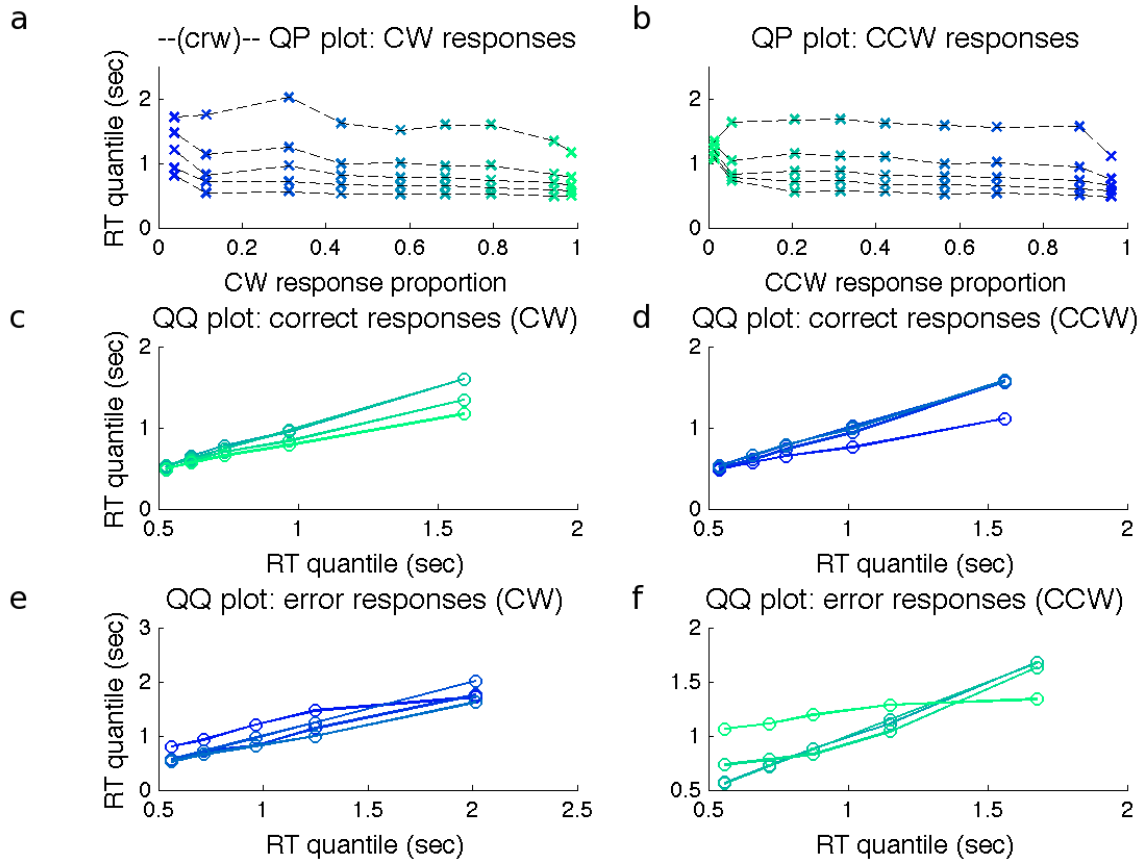


Figure 4.6.: **a, b** percentile probability (QP) twin plots, averaged over *crw* participants. *Dashed lines* in QP plots connect the RTs that correspond to the same percentile (0.1, 0.3, 0.5 (median), 0.7 and 0.9) across disparity conditions. Scatter point colours indicate the stimulus disparity conditions and range in 9 steps from *deep blue* = maximum CCW (negative) disparity, to *bright green* = maximum CW (positive) disparity. **c - f** quantile-quantile (QQ) plots, averaged over *crw* participants, displayed separately for correct and error trials. QP and QQ plots for the two response directions (CW and CCW) have been drawn separately side-by-side to lay clear any differences in reaction time distribution by response direction

Figure 4.7. *(following page)*: percentile propability (QP) twin plots (*upper panels*) and quantile-quantile (QQ) plots (*lower panels*) for two example individual participants, *srt3* & *crw4*. *Dashed lines* in QP plots connect the RTs that correspond to the same percentile (0.1, 0.3, 0.5 (median), 0.7 and 0.9) across disparity conditions. For the QP plot of participant *srt3*, turquoise indicates the zero-disparity condition; purple, ochre and blank indicate CW (positive) disparity conditions; red, green and blue indicate CCW (negative) disparities. For the QP plot of participant *crw4*, purple indicates the zero-diparity condition; ochre, black, brown and yellow indicate the CW disparities; turquoise, red, green and blue indicate CCW disparities.

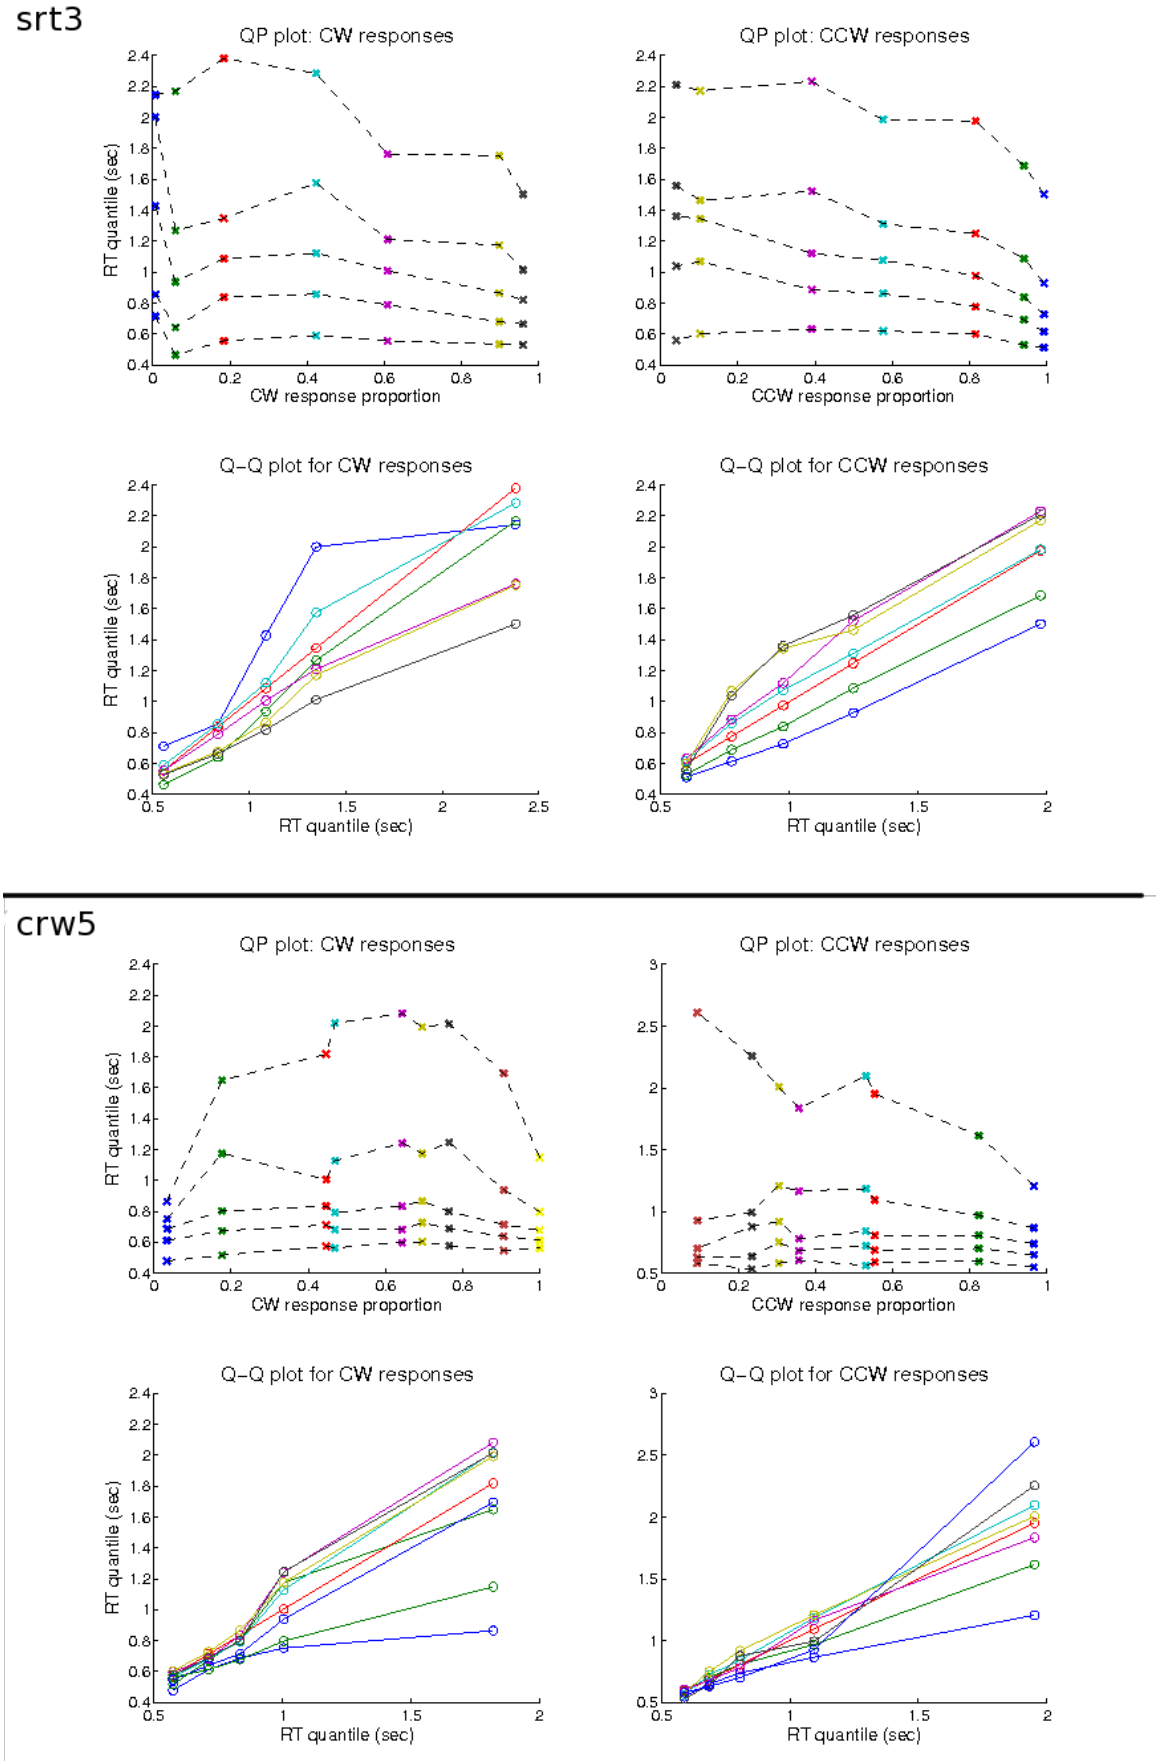

srt3
QP plot: CW responses
RT quantile (sec)
CW response proportion
QP plot: CCW responses
CCW response proportion
Q–Q plot for CW responses
RT quantile (sec)
RT quantile (sec)
Q–Q plot for CCW responses
RT quantile (sec)
crw5
QP plot: CW responses
RT quantile (sec)
CW response proportion
QP plot: CCW responses
CCW response proportion
Q–Q plot for CW responses
RT quantile (sec)
RT quantile (sec)
Q–Q plot for CCW responses
RT quantile (sec)

Fig. 4.7

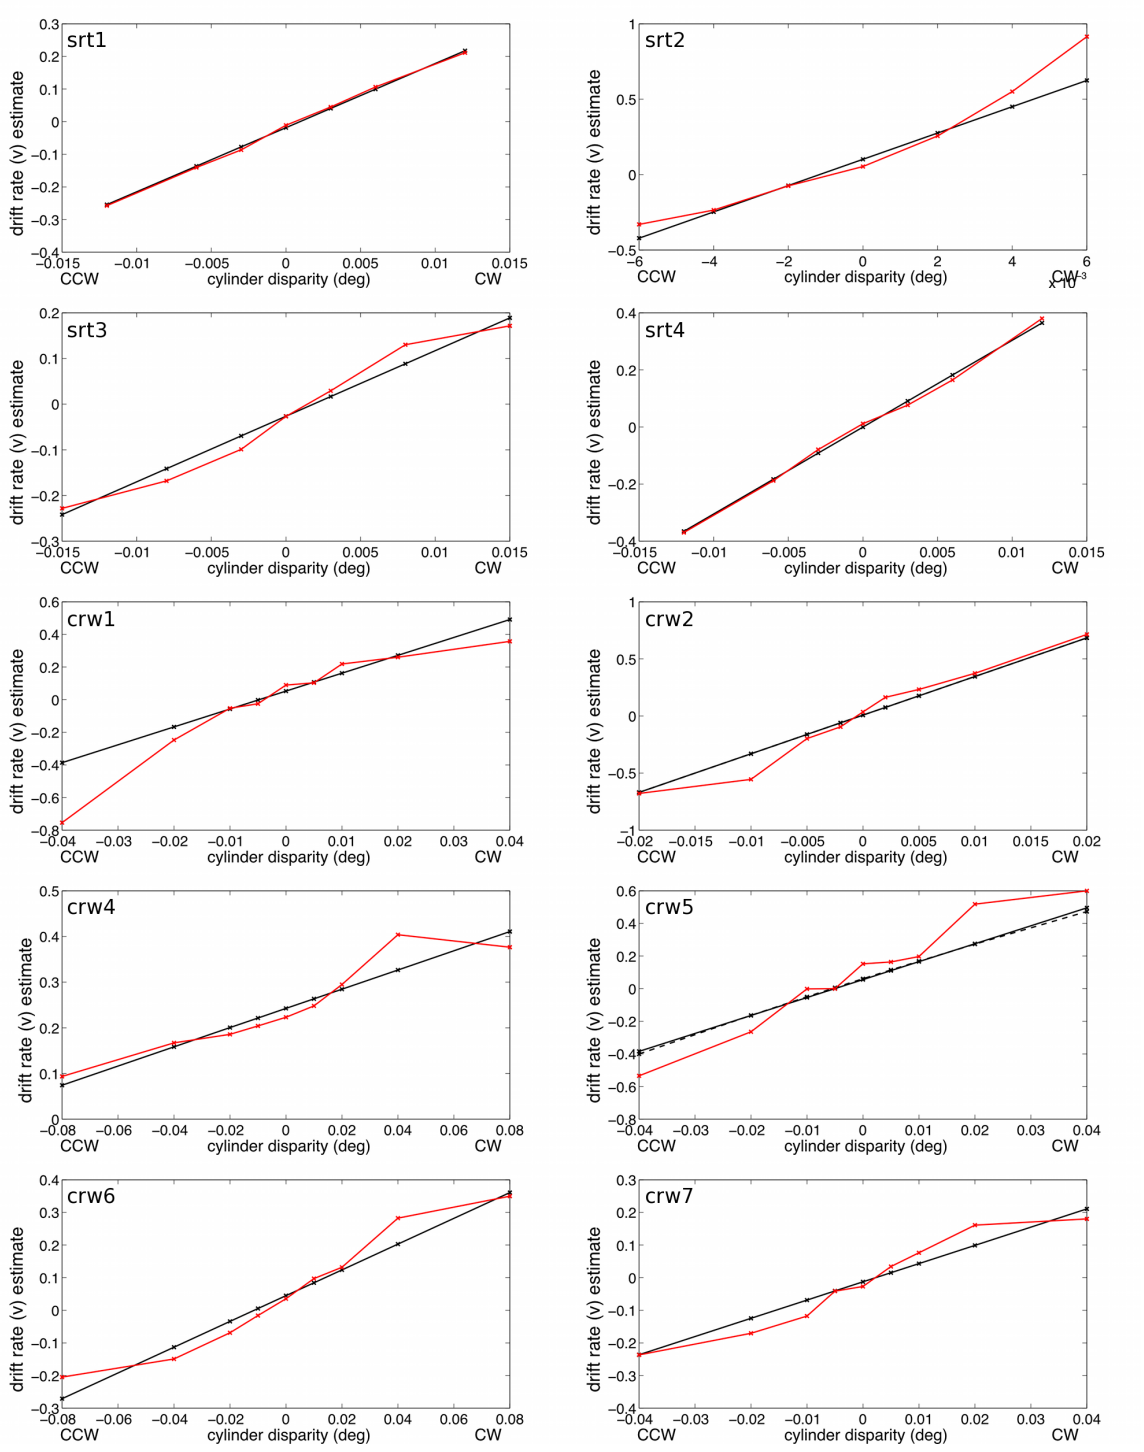


Figure 4.8.: Drift rate (v) parameter estimates plotted against cylinder disparity (stimulus strength), for all 10 participants for the basic cylinder reaction time task (no reward bias manipulations). *Black lines*: drift rate estimates from the diffusion model in which the relationship between the drift rate and stimulus disparity was constrained to be linear; *red lines*: drift rate estimates from the model in which drift rate was unconstrained.

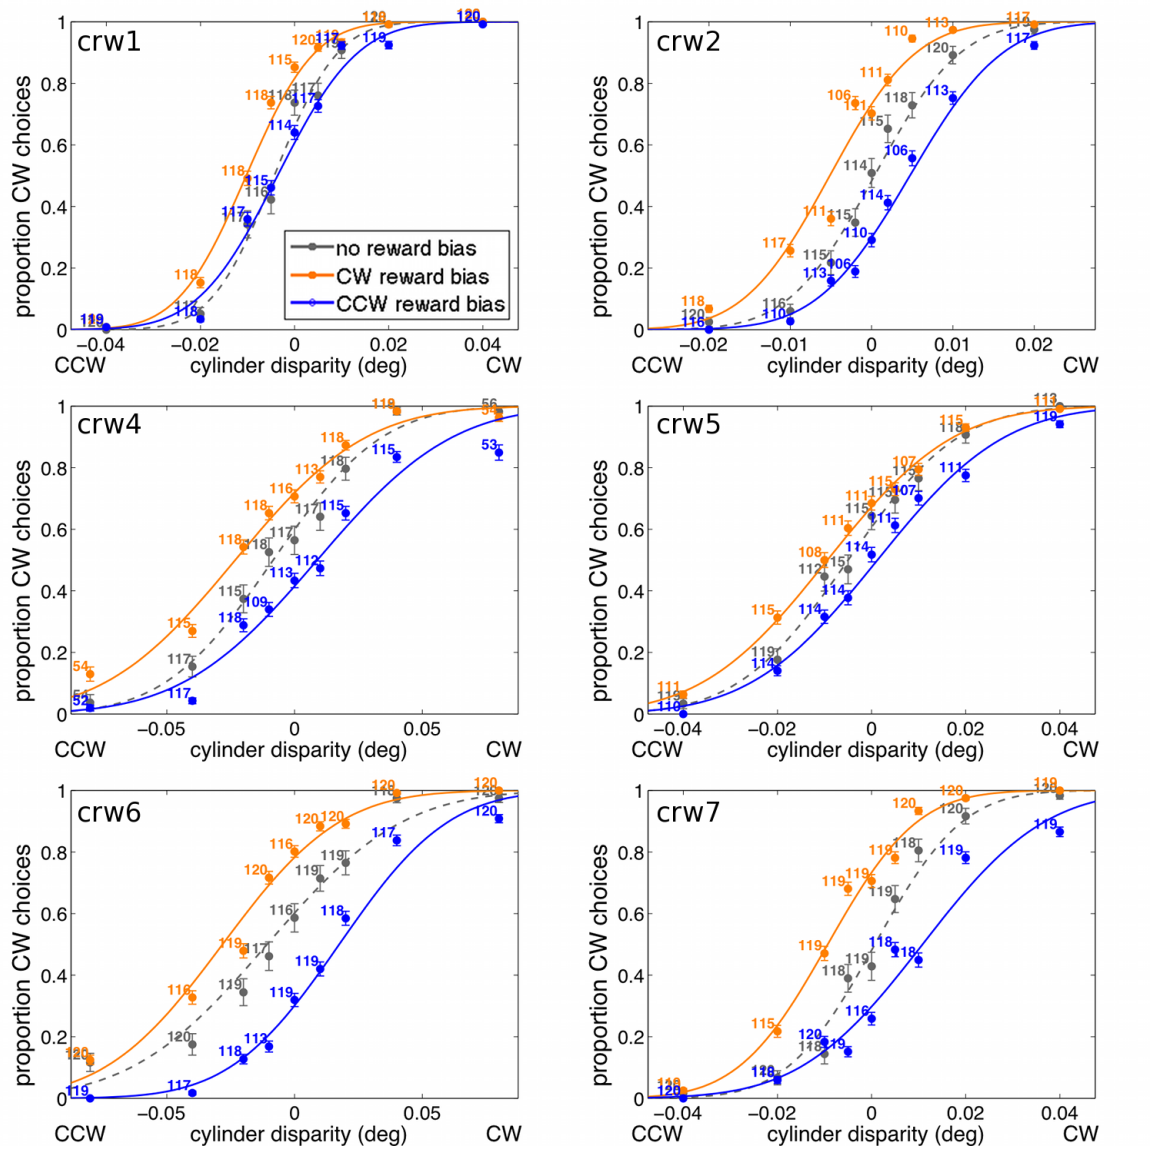


Figure 4.9.: Psychometric functions in each reward condition fit with cumulative normal distributions, shown for each participant: *orange lines*: HL condition (CW direction is highly rewarded); *blue lines* LH condition (CCW direction is highly rewarded); *grey lines*: neutral LL condition (both directions equally rewarded).

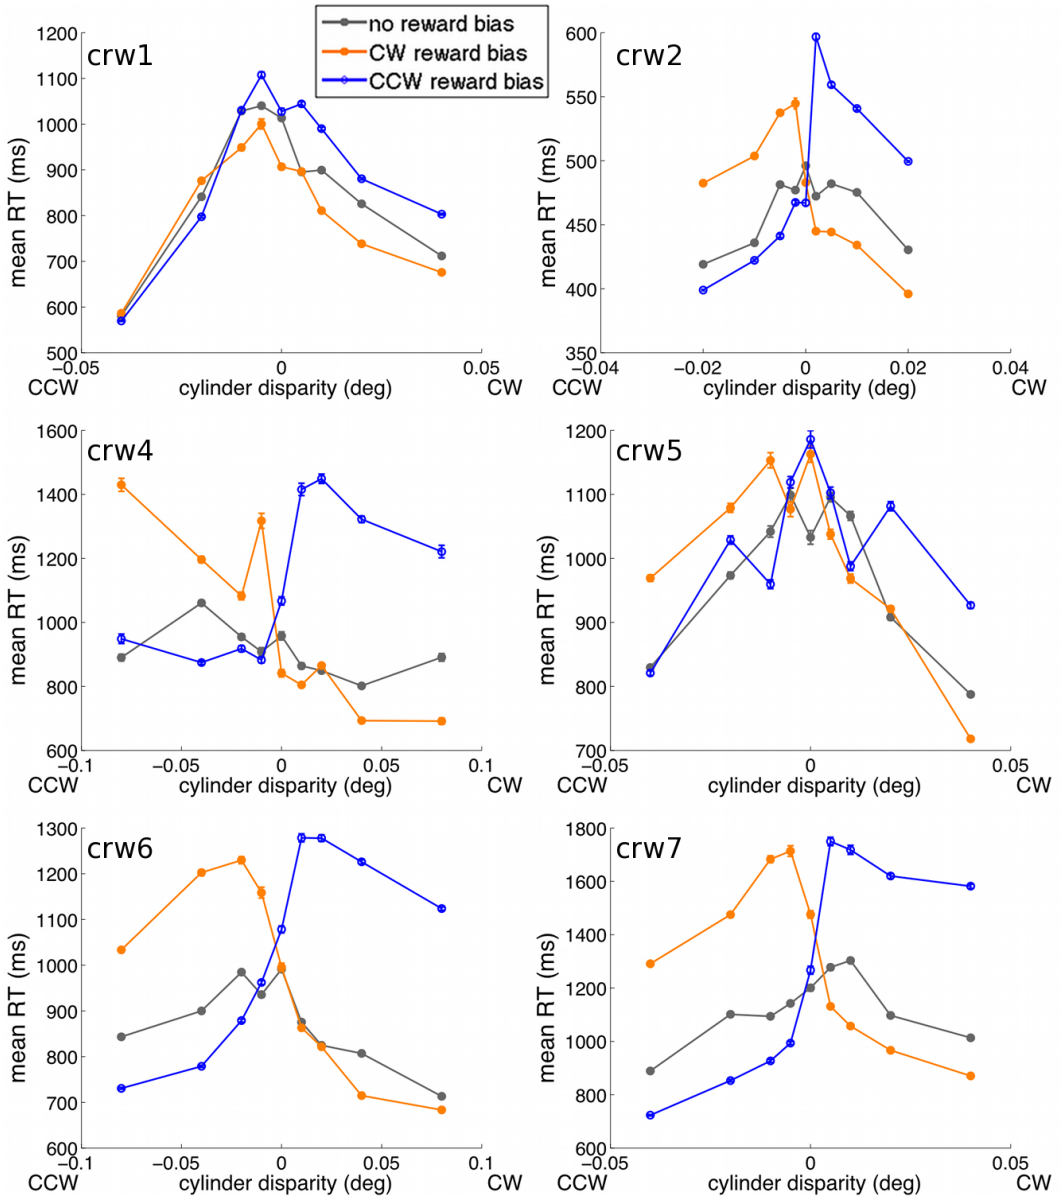


Figure 4.10.: Mean reaction times for correct responses to each cylinder stimulus disparity, by reward condition, for each participant. Errrorbars = ± SEM.

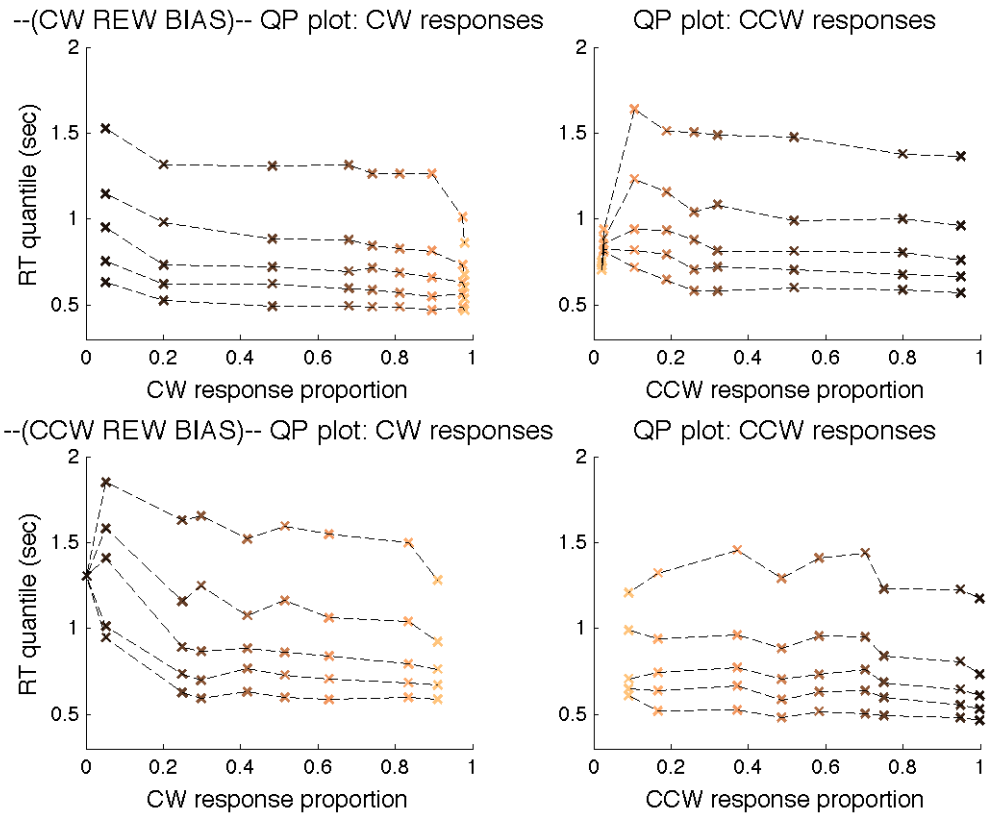


Figure 4.11.: Quantile-probability (QP) twin plots, averaged over *crw* participants. *Upper panels:* HL condition, where CW is the more highly-rewarded direction; *lower panels:* LH condition, where CCW is the more highly-rewarded direction. Scatter points correspond to the RT percentiles: 0.1, 0.3, 0.5 (median), 0.7 and 0.9. Colours indicate the stimulus disparity conditions and range in 9 steps from *black* = maximum CCW disparity, to *brightest copper* = maximum CW disparity. *Dashed lines* connect the same RT percentiles across disparity conditions.

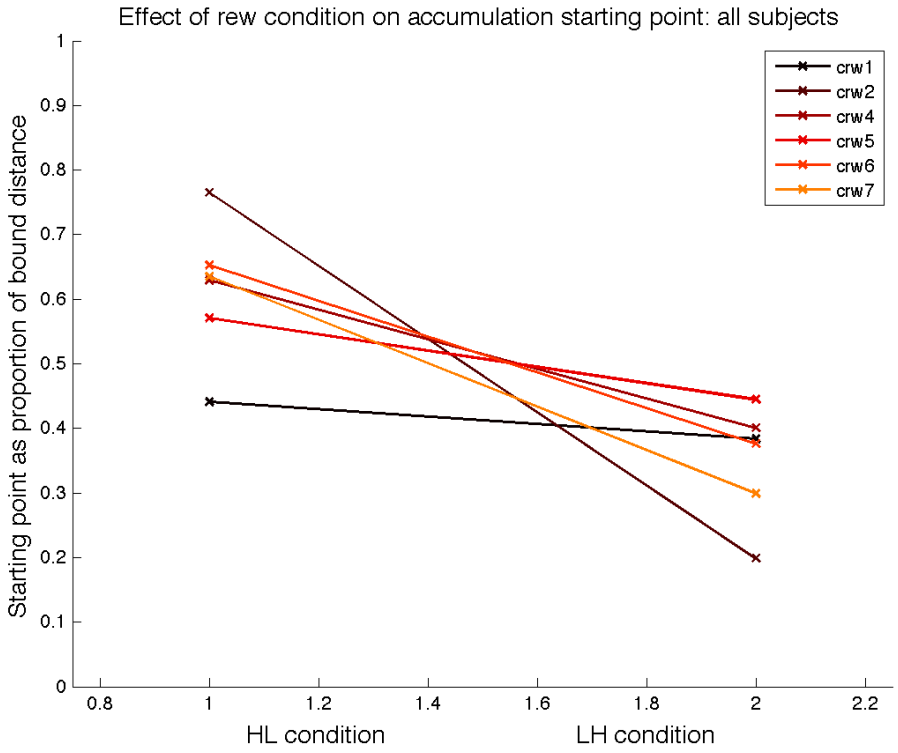


Figure 4.12.: The starting point of the drift-diffusion process was closer to the more-rewarded boundary in the best-fitting drift-diffusion model for all six participants. HL condition: the leftwards (higher bound) response was more highly-rewarded, corresponding to a CW rotating cylinder; LH condition: the rightwards (lower bound) response was more highly-rewarded, corresponding to a CCW rotating cylinder.

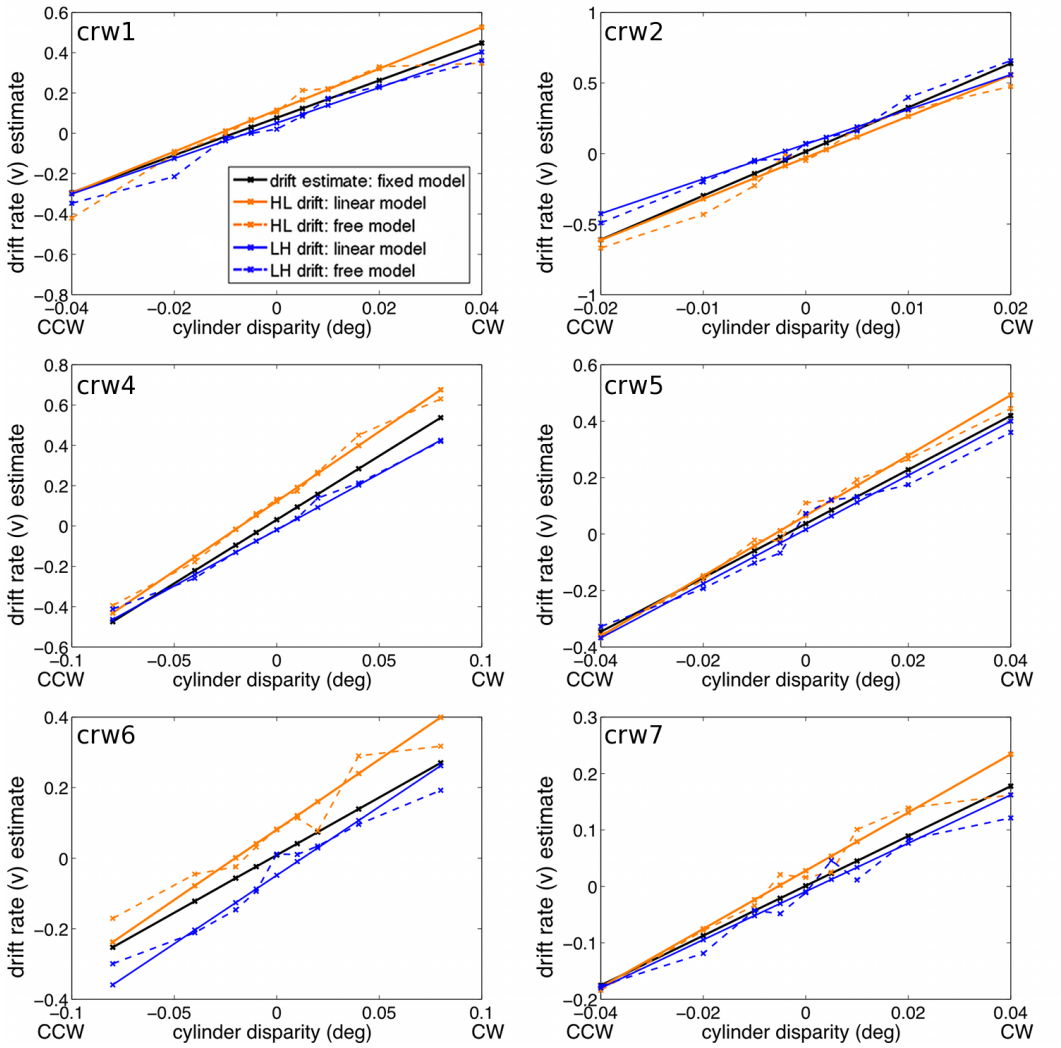


Figure 4.13.: The relationship between the drift rate parameter (v) and cylinder disparity shows a difference by reward condition in the best-fitting model for all six participants (see Table 2). In 5/6 participants, the drift rate was biased towards the higher-rewarded response direction: *orange*: CW reward bias condition (HL); *blue*: CCW reward bias condition (LH); *smooth lines*: drift rate estimates from linear drift rate model; *dotted lines*: drift rate estimates from unconstrained model; *black lines*: drift rate estimates from the model in which drift rate was forced to be constant across reward conditions.

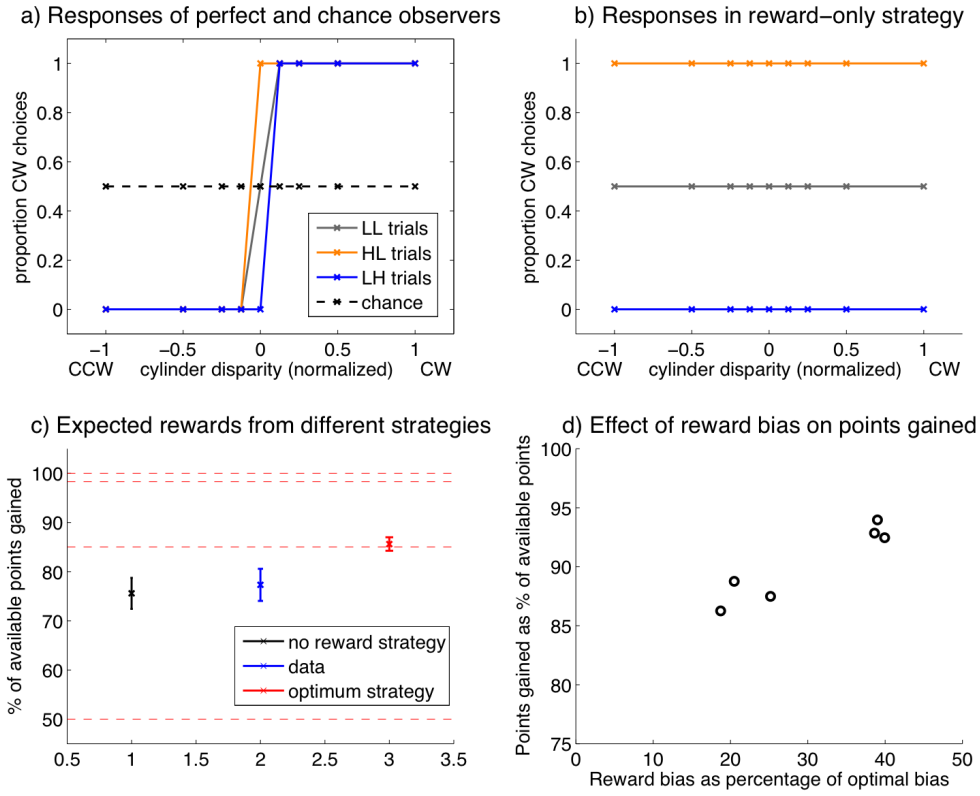


Figure 4.14.: **a)** Responses of a perfect observer (*grey lines* = LL condition; *orange lines* = HL condition; *blue lines* = LH condition) and a chance observer (*black dashed line*). **b)** Responses of a reward-only strategist who ignores information about the visual stimulus (colours as for **(a)**). **c)** The percentage of total points gained by different strategies: *black*: expected gain if no reward-related bias is implemented; *blue* points actually collected in the experiment across participants; *red* expected gain in the optimal reward biasing strategy; *red dashed lines* show the 50% (chance), 85% (reward-only strategy), 98.3% (perfect observer) and 100% points collection levels. **d)** Points collected in experiment, as a % of the total available under the optimal strategy, plotted against the reward bias, for each participant. The reason the points do not lie on a straight line is because only the shift between the HL and LH PMFs, and not the absolute locations of these PMFs, is represented by these calculations. The overall perceptual bias of each participant, which determines the absolute locations of the PMFs, also affects the reward collection rate and is variable between participants, leading to the variability seen here.

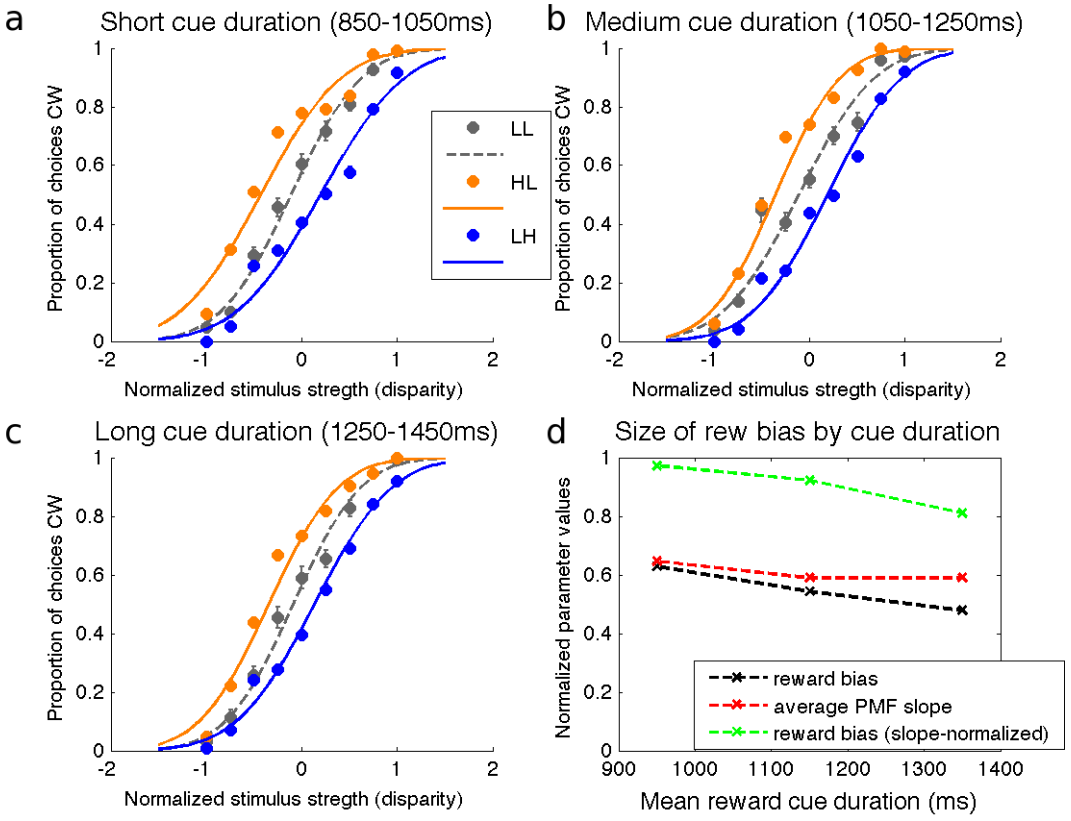


Figure 4.15.: Psychometric functions (PMFs), combined across participants, separated by reward condition and fit with cumulative normal distributions, are shown for short (**a**), medium (**b**) and long (**c**) points cue duration groups (*grey*: LL condition; *orange*: HL condition; *blue*: LH condition). **d**) Reward bias (*black*), measured as the absolute difference between the means of the HL and LH PMFs; average slope (SD) of the PMFs (*red*); slope-normalized reward bias (*green*); all plotted against mean cue duration of each cue duration group. To combine choice results across participants, cylinder disparity values were normalized by division by the largest disparity value used for each participant.

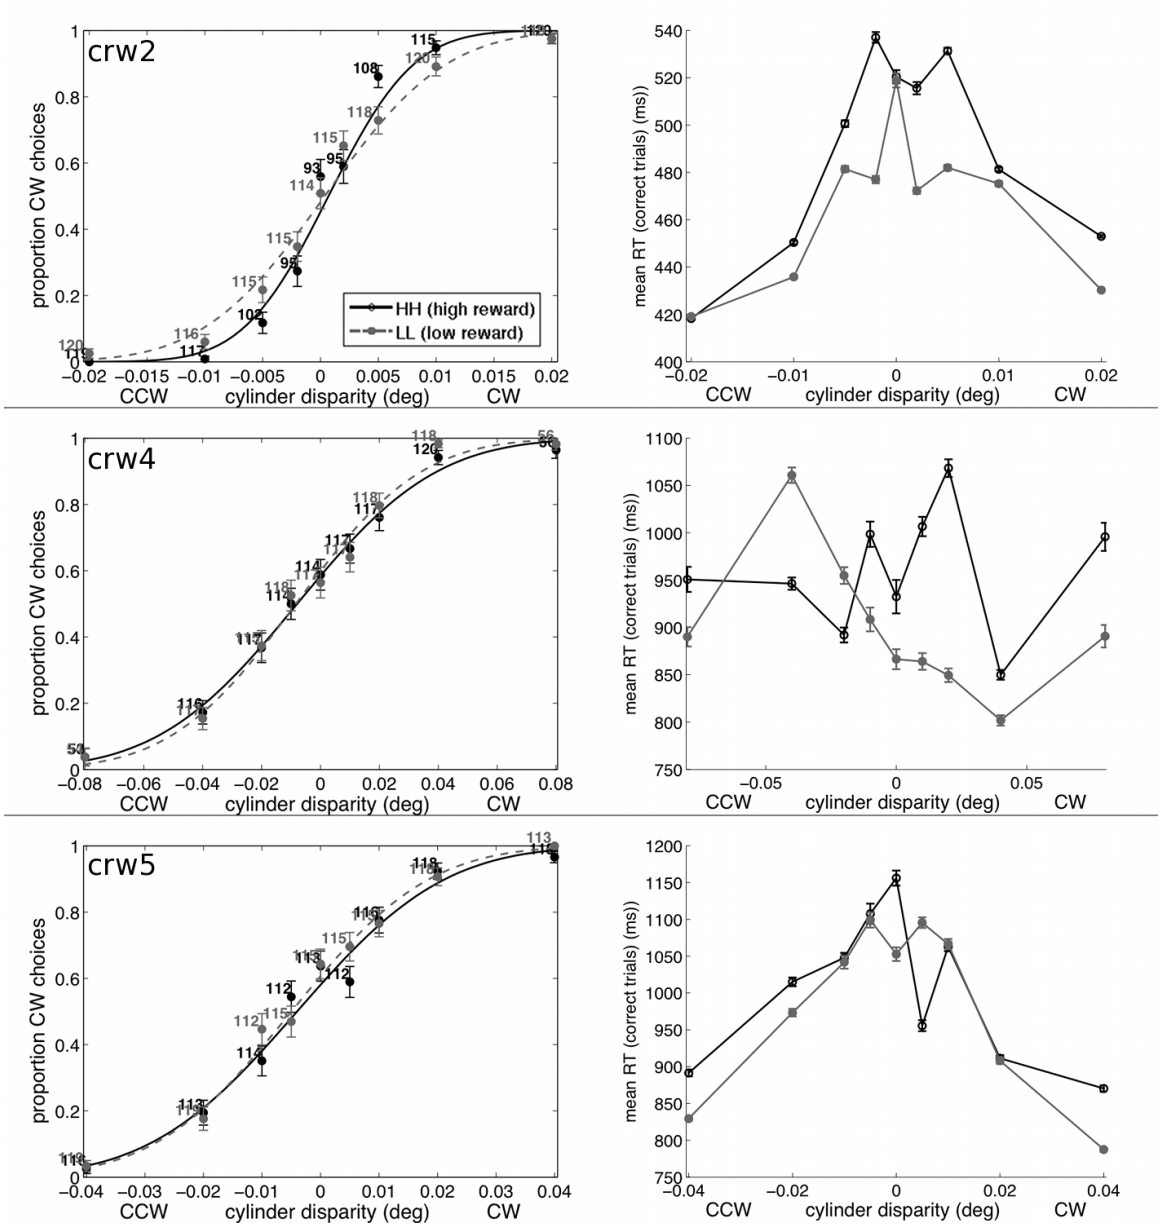


Figure 4.16.: Psychometric functions, fit with cumulative normal distributions (*left panels*), and mean RTs (*right panels*) are shown for three participants (*crw2*, *crw4*, *crw5*). *Black lines*: HH condition, 8 points for correct choice in either direction; *grey lines*: LL condition, 1 point for correct choice in either direction.

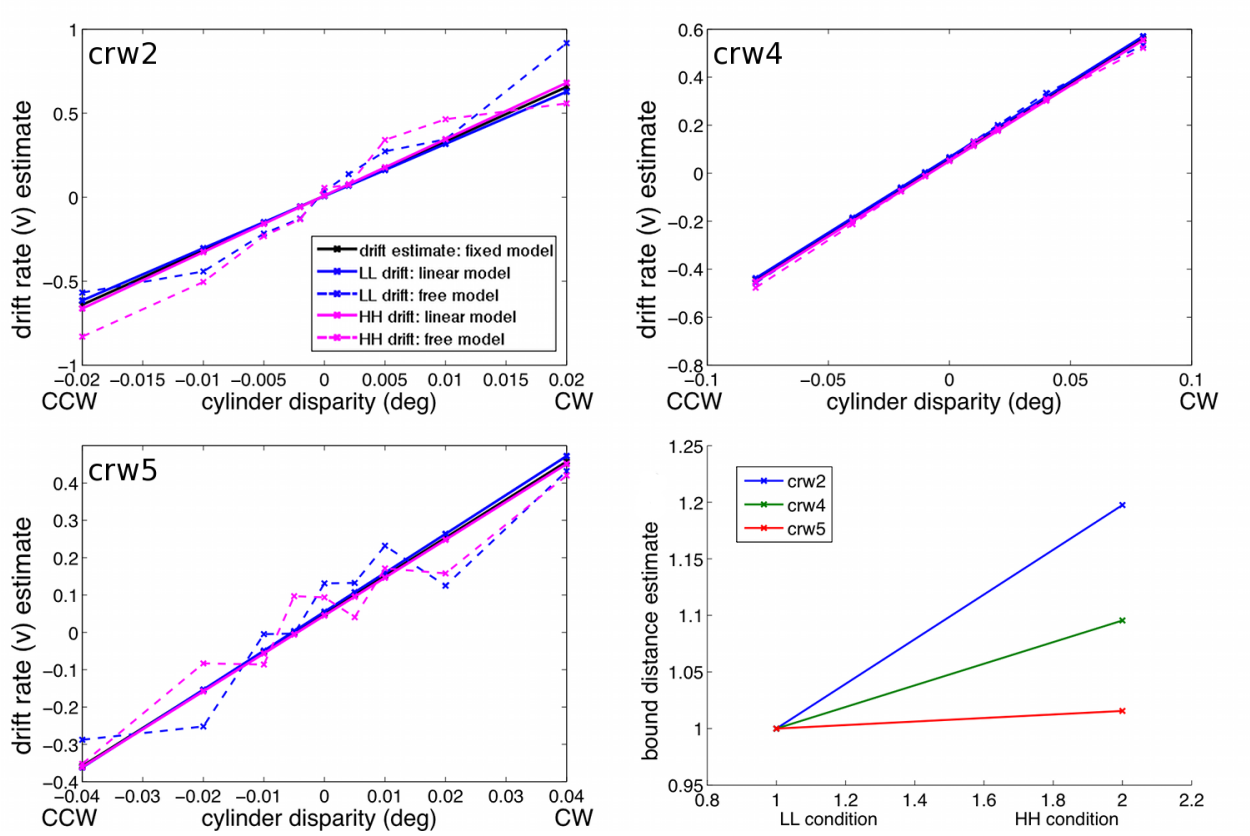


Figure 4.17.: *Upper left, upper right, lower left:* Estimates of drift rate plotted against cylinder disparity for the different models tested in HH and LL conditions, for the three participants (*crw2*, *crw4*, *crw5*). *Lower right:* Estimated distance between decision bounds in HH and LL conditions, normalized by the estimate in the LL condition.

5. MEG Materials and Methods

The following sections describe the common experimental and analytical methods applied in Chapters 6 & 7.

All MEG studies were performed on healthy human volunteers with normal or corrected-to-normal visual acuity. No participants had a reported history of neurological problems or received current psychoactive medication. MEG research projects were approved by the University of Oxford Central University Research Ethics Committee (CUREC). Written informed consent was obtained from all volunteers, except those directly collaborating on the research. Participants were recruited by poster advertisments throughout the University of Oxford, by online advertisments, or through personal communication.

5.1. Introduction to MEG

Magnetoencephalography (MEG) is a human neuroimaging technique that has received increasing recent interest. With highly sensitive magnetic sensor coils, known as super-conducting quantum interference devices (SQUIDs), MEG detects and records magnetic fields produced by electrical brain activity at the level of the femto-Tesla (10^{-15} T) (Hansen *et al.*, 2010). The main neural source of the MEG signal is likely to be post-synaptic currents in the apical dendrites of tens of thousands of simultaneously active neocortical pyramidal cells, although other sources such as basal dendrites, bushy cells and subcortical structures may make a small contribution (Hamalainen *et al.*, 1993; Okada *et al.*, 1997; Murakiami & Okada, 2006). Due to the plane of orientation of the scanner coils, MEG can only detect magnetic fields that have a component perpendicular to the skull. Since pyramidal cells are oriented perpendicular to the cortical surface, activity in neurons located in the walls of sucli, rather than within the crowns of gyri, will be most 'visible' to MEG (Figure 5.1).

The temporal resolution of MEG measurements is high, limited only by the sampling rate of

the electronic acquisition system, which is usually on the order of 1000 Hz. This compares very favourably with the temporal resolution of functional magnetic resonance imaging (fMRI), a more commonly-used neuroimaging technique which generally has a temporal resolution of 1-2 seconds (Menon *et al.*, 1998; Liao *et al.*, 2002). Therefore, MEG signals may be correlated with psychophysical events shorter than 1000 ms in duration. On the other hand, accurate localization of brain sources of MEG signals is a non-trivial problem and techniques for source analysis an active research topic. An evaluation of some approaches for localizing early visual cortical activity in humans is the subject of Chapter 6. Still, MEG has a better spatial sample profile across the head than electroencephalography (EEG), another commonly-used recording technique, because MEG records from more sensors than are found on EEG head caps and, unlike the electrical currents recorded by EEG, magnetic fields are minimally smeared by passage through the inhomogenously conducting pial, skull and skin surfaces (Hansen *et al.*, 2010). All in all, MEG is a promising non-invasive technique for investigating human brain activity with excellent temporal resolution and reasonable spatial resolution.

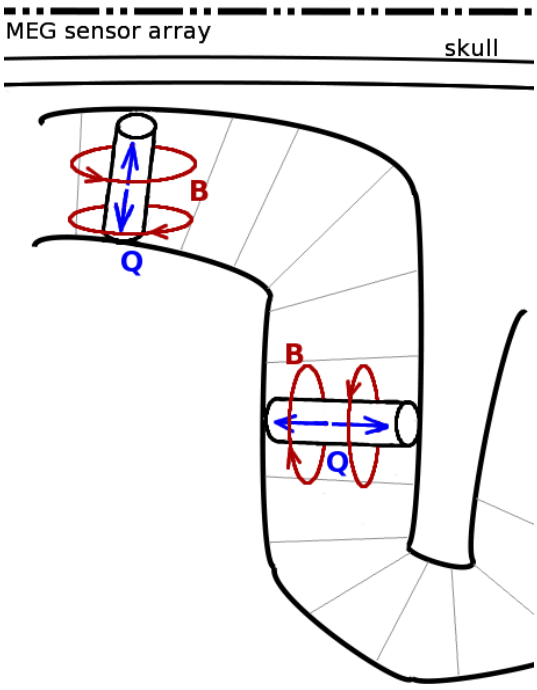


Figure 5.1.: Diagrammatic representation of magnetic fields caused by neuronal activity sources in the crown of a gyrus and the wall of a sulcus in the cerebral cortex (after Hansen, Kringelbach & Salmelin, 2010). Two spindles indicate the relationship between the direction of intracellular currents (**Q**) and the resultant magentic fields (**B**). The sulcal source produces a magnetic flux that is directed out of the head and detected by the MEG sensor array; the magnetic flux from the gyral source remains mainly within the head and parallel to the sensor array, and is not well-detected.

5.2. MEG Data Collection

The main focus of MEG data collection is to record human brain activity whilst removing/minimizing sources of noise, and monitoring sources of noise where they cannot be removed.

5.2.1. Scanner & Acquisition Systems

MEG data were collected with an Elekta Neuromag VectorView® MEG scanner at The Oxford Centre for Human Brain Activity (OHBA), Department of Psychiatry, University of Oxford. The Elekta scanner has 306 MEG-channel sensors, comprised of 102 magnetometers and 204 planar gradiometers. Magnetometers measure the magnetic field component perpendicular to the surface of their coils, and are sensitive to both near and far sources (Figure 5.2 a). Gradiometers measure a spatial gradient of the magnetic field component, rather than the field component itself, resulting in greatest sensitivity to sources immediately beneath the sensors (Figure 5.2 b). Combining information from both sensor types gives a more detailed picture of the magnetic field distributions and accounts better for noise. Sensors were tuned daily prior to the first MEG recording session, to take noise levels down to approximately 2.5 fT/cm. If necessary, sensors that became very noisy during a recording session would be individually re-tuned between blocks by a quick automatized heating process.

MEG signals, electrooculogram (EOG), electrocardiogram (ECG), response pad and photo-diode signals (see below for details), were recorded by the Neuromag Acquisition system. Graphical interfaces on a GNU/Linux workstation allowed the experimenter to monitor the data acquisition process on-line. Continuous data were recorded with a sampling rate of 1000 Hz, with a bandpass filter of 0.3 - 330 Hz.

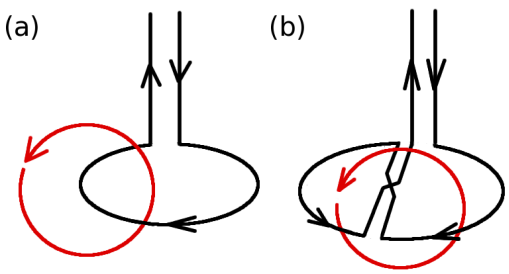


Figure 5.2.: Magnetic flux pick-up geometries for (a) magnetometer & (b) gradiometer. Red lines indicate the magnetic field position and orientation detected by the sensor.

5.2.2. Participant Preparation & Head Position Recording

All metal was removed from the participant prior to entering the scanner room, to minimize magnetic interference. EOG and ECG electrodes were connected to the participant at the appropriate locations on the left and right wrist and above and below the outer-most corner of the right eye. A ground electrode was placed in the centre of the forehead. Horizontal EOG (an electrode placed at the outer corner of each eye) was additionally used for recording eye-movements in some cases.

Accurately recording participant's head position in the scanner is necessary to later infer the spatial profile of measured magnetic signals with respect to brain structures. Four head position indicator (HPI) coils were attached to the participant's head, placed behind each ear and wide apart on each side of the forehead and firmly stuck on with surgical tape. A Polhemus stylus and digitizer device were used to record the locations of fiducial points (right and left preauricular points (RPA, LPA) and nasion), HPI coils, and 50 to 150 extra points on the head surface. Prior to each recording, head location in the scanner was measured with an automatic process that detected the coils. From early 2010, the acquisition system also supported continuous HPI measurement, which was used.

5.2.3. Eyetracking

An EyeLink 1000 infra-red eyetracker (SR Research Ltd.) measured eye position and movement during MEG recording sessions. The eyetracker was positioned on a desk tray in front of the participant and a lens was used to focus the eyetracker onto the pupil and cornea. Eye recording calibration was performed at the start of each session with EyeLink's *elcl.exe* program. Data were recorded at a sample rate of 1000 Hz.

5.2.4. Stimulus Presentation, Timing & Button Presses

Visual stimuli were projected onto a back-projection screen in front of the participant, with a Panasonic® DLP (Digital Light Processing) based projector (PT-D7700E). Refresh rate was 60 Hz. The distance between viewers's eyes and the screen was 1500 mm. The stimulus onset times were accurately recorded with a photodiode (sampling rate 1000 Hz), placed over a small (8 x 8 mm) black box located in the bottom-left corner of the stimulus screen; this box flashed white on the first frame of each stimulus onset.

Participant's responses were recorded with optical response devices for finger lifts and presses. These 'button press' events were relayed near-instantly to the MEG acquisition system via fibre-optic cables.

5.2.5. MEG Recording

The participant was seated in the scanner under the MEG sensor helmet and the chair raised until the top of their head touched the helmet. The door of the magnetically shielded room was kept closed during experiments and we continuously monitored the participant via video camera connection and audio link and the MEG/EOG/ECG recordings via the Neuromag acquisition system GUI.

5.3. MEG Data Preprocessing

Preprocessing stages are required to improve the low signal-to-noise ratio of the data prior to further analysis. This includes noise reduction and artefact removal (Figure 5.4).

5.3.1. Maxfiltering & HPI correction

Data were preprocessed with Electa Neuromag® *MaxFilter* software (version 2.1, May 2009). *MaxFilter* (Maxwell filtering) software reduces noise in the data by suppressing magnetic interference coming from outside and inside the sensory array, using signal-space separation (SSS). SSS is a theory-driven method that separates raw data into subspaces comprising sources from outside the sensor helmet, sources from inside the helmet, and sources of noise generated at or very close to the sensors. Interference suppression is achieved by neglecting the the outside-space component and reducing the noise component as far as possible. *Maxfilter* inherently transforms measured MEG signals into virtual channels that are independent of the recording device. All data were maxfiltered using default command line parameters.

The *MaxMove* sub-command was used to transform data between different head positions. *MaxMove* registers multiple MEG datasets to the same spatial coordinates and thus allows combination/averaging of the data across blocks and/or participants. *MaxMove* utilizes the virtual channels to estimate the signals corresponding to the desired new head position and resaves

the data to these values. Where continuous HPI was measured, continuous movement compensation by *MaxMove* was additionally applied.

5.3.2. Event Assignment & Epoching

Each stimulus onset event was discerned as the time of the corresponding onset in the photodiode (see above). The photodiode sequence was aligned with order of stimuli/conditions presented during the experiment (Matlab custom scripts). This assigned an identity to each stimulus event. Button press times were discerned as the time of the onset in the response trigger channel of the MEG acquisition system. All event times were listed in an 'events' file (.txt format) for each dataset and imported into softwares used for subsequent MEG analysis. Data were epoched by cutting continuous MEG data around the timepoint(s) of events, to create many individual epochs (trials).

5.3.3. Artefact Removal

Post-maxfiltering, the main remaining artefacts arise from eye blinks and movements, participant noise source such as muscle movements, and 'bad' (noisy) channels. Those channels with constant high noise levels as identified by visual inspection were labelled as 'bad' and removed from further analysis. Eye-related artifacts such as blinks were identified as deviations in the EOG recording trace. Epochs (trials) containing artifacts arising from the eyes or intermittent sensor noise were removed from further analysis:

(i) Automatic removal of artefactual epochs:
In Minimum Norm Estimate analysis software (MNE, see below), the EOG trace was visually inspected, and an appropriate peak-to-peak threshold level was specified to remove eye blinks and overt eye movements (thresholds were normally within the range 100 to 200 x 10^{-6} V). A similar process was used to set the maximum noise level threshold for magnetometer and gradiometer activity (normally within range 2 to 3 x 10^{-12} T and 1.5 to 2.0 x 10^{-10} T/m, respectively). Epochs in which one or more of these thresholds were surpassed were deemed artifactual.

(ii) Visual artifact rejection:
For SPM8 software-based analyses, SPM8's visual artifact rejection (VAR) scripts were used to remove artifactual epochs. The variance of the EOG, gradiometers, and magnetometers traces in each epoch was presented on a GUI plot and those epochs that contained high noise lev-

els were visually identified, selected and removed from further processing. This resulted in the same events being removed as method *(i)* above in almost 100% of cases.

5.4. MEG Analysis: Sensor Space

MEG signals were analyzed in sensor space to inspect the topographic pattern of activity over the scalp in different experimental conditions (Figure 5.4).

5.4.1. Time-frequency (TF) decomposition

Identifying the frequency bands with the greatest power in sensor space can usefully inform the subsequent specification of frequency bins for beamformer source reconstruction of the MEG signals. In this thesis, such TF analyses were performed using FieldTrip scripts (Oostenvald *et al.*, 2011) incorporated into the OHBA Software Library toolbox (OSL, see below). Preprocessed, epoched data was passed through a Hilbert transform frequency decomposition for a specified number of frequency bins spanning the relevant frequency range. Both magnetomers and gradiometers were included.

5.5. MEG Analysis: Source Reconstruction

Reconstructing the sources of the MEG signal in the brain is a non-trivial, ill-posed problem known as the *inverse problem* (see Chapter 6 for a detailed discussion). A number of different MEG source reconstruction algorithms are in development, each with different assumptions about underlying brain activity patterns. Source analysis methods can be split into two broad groups, *dipole localization* and *distributed localization*. Dipole approaches assume that brain activity is adequately described by one or a few dipole sources in specific locations. Distributed, or *imaging*, approaches aim to recover parameters that describe the amplitude of mass neural currents that spatially sample the whole brain. Three of the main distributed source analysis methods are used in this thesis: the MNE software (Minimum Norm Estimate, MGH/MIT Martinos Centre for BioMedical Brain Imaging), the SPM8 software for M/EEG (FIL Methods Group, UCL) and a linearly-constrained minimum variance (LCMV) beamformer implemented in the OSL toolbox (OHBA Analysis Group, Oxford University). Analysis pipelines for each are described below (Figure 5.4). In all cases, source analyses were run on a `tcsh` shell on a Unix-based (Apple MacBook Pro) or on a GNU/Linux operating system.

5.5.1. Structural MRI

Information about individual subjects' head and brain anatomy is necessary to most accurately localize MEG signals. Anatomical magnetic resonance imaging (aMRI) data were collected with a 3.0 Tesla TIM Trio scanner, located at the University of Oxford Centre for Clinical Magnetic Resonance Research (OCMR), by the author or by collaborators Dr Holly Bridge, Dr Lori Minini or Dr Betina Ip. One such T1 scan was taken for each participant, lasting around 5 minutes. A standard structural magnetization-prepared rapid gradient echo (MPRAGE) scan procedure was used (130 Hx/pixel, flip angle = 8 degrees, TR/TE/TI = 2040 ms / 4.7 ms / 900ms). The orientation of scan acquisition was transverse (192 slices); sagittal slices numbered 174. Slice thickness was 1.00 mm and voxel size was 1.0 x 1.0 x 1.0 mm.

5.5.2. Anatomical reconstruction (Freesurfer)

Individuals' anatomical surfaces, onto which MEG data were coregistered and reconstructed, were created from using Freesurfer software (`http://www.surfer.nmr.mgh.harverd.edu`) (Dale *et al.*, 1999; Fischl *et al.*, 1999). FreeSurfer's *recon-all* process was run with default parameters to reconstruct cortical surfaces from aMRI images. White/grey matter segmentations were checked by eye and manual edits were made where necessary. FreeSurfer's *watershed* algorithm was used to reconstruct the brain, inner skull, outer skull and outer skin surfaces from the aMRI data, to estimate the boundary element model (BEM) which describes the compartments that the magnetic field passes through before reaching MEG sensors.

5.5.3. MEG-MRI coregistration

The fiducial points and other HPI information from the 3D digitizer (Polhemus) data were used to align the coordinate frames of the MEG data and the structural brain data in one of two ways:

5.5.3.1. Co-registration with MNE software

The HPI data and the structural surface data were loaded into the *mne_analyze* graphic interface of the MNE software package (*version 2.7*). Locations of fiducial points were identified by eye on the structural surfaces to produce an initial alignment between digitizer points and head anatomy, which was fine-tuned by eye with linear shifts and rotations along the three coordinate axes. HPI points that remained >5mm away from the head surface were discarded as

outliers. An automatic alignment procedure used the Iterative Closest Point (ICA) algorithm to non-linearly converge to the optimal alignment (Figure 5.3 a).

5.5.3.2. Co-registration with SPM8

In SPM8-based software, surface meshes used for source reconstruction were created by *inverse normalization* of the canonical mesh (Mattout *et al.*, 2007). Firstly, the individual's aMRI image was normalized to the canonical MNI152 template brain. Then, the inverse of this deformation (normalization) field was applied to a high resolution cortical mesh, derived from the MNI152 template. This produced a high resolution cortical mesh and headshape morphed to the individual's anatomy. Fiducial points used for MEG-MRI co-registration were then aligned to pre-labelled fiducial points on the inverse-normalized headshape. HPI points were also used in the alignment procedure which was based upon an ICP algorithm as above (Figure 5.3 b).

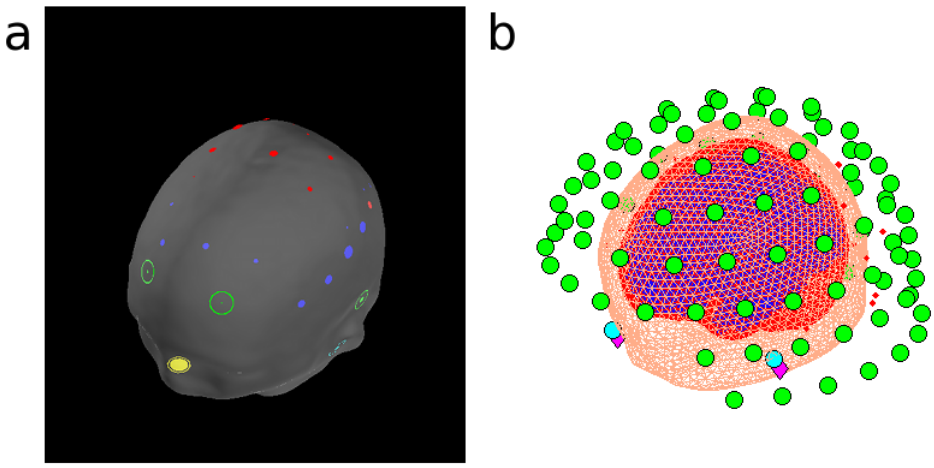


Figure 5.3.: Graphical outputs of co-registration of MEG and anatomical MRI dataa) Co-registration output from the MNE software, showing the headshape points (*red, blue*) and the Polhemus fiducials and HPI (*yellow, green*). (b) Co-registration output from the SPM8 software, showing the headshape points (*red diamonds*), Polhemus fiducials and HPI (*blue circles, pink diamonds*) and MEG sensor locations (*green circles*).

5.5.4. Minimum Norm Estimate (MNE)

Anatomically-constrained minimum norm estimate source reconstruction was implemented with the MNE software, version 2.7 (`http://www.nmr.mgh.harvard.edu/martinos/userInfo/data/sofMNE.php`; (Dale & Sereno, 1993; Hamalainen & Ilmoiemi, 1994). The minimum norm approach to source reconstruction of the MEG signal uses a distributed source model, in which the locations of a large number of dipoles are kept fixed and their amplitudes are estimated based upon measured MEG data. A noise-covariance matrix (NCM) is estimated from the baseline data, such that the output of the n-channel MEG sensor array during baseline is represented as a vector in the n-dimensional signal space (each channel spans one dimension of that virtual space). This information is used to inform the source estimation process by placing a smaller weight on model-data mismatches that occur in 'noisier' signal-space directions. To overcome the ill-posedness of the inverse problem, the current distribution with the minimum mean amplitude is selected as the unique solution.

Due to close integration between MNE software and Freesurfer software, individuals' cortical meshes, produced by segmentation with Freesurfer, could be used to create the source space for MNE analysis. The cortical surface was decimated at either 3.1 mm or 4.9 mm spacing (depending upon quality of individual subject's MRI), resulting in either 10242 or 4098 vertices (sources) per cortical hemisphere, respectively.

MEG data were bandpass filtered 1-40 Hz and averaged over epochs by condition; magnetometer and gradiometer signals were combined; the NCM was estimated from averaged, prestimulus baseline data; a single compartment model (brain only) was used for BEM estimation in the forward model; a *loose* orientation constraint was employed such that the variance of source components tangential to the cortical surface were reduced by a factor of 0.2; standard MNE *depth* weighting was applied; source space points closer than 0.5 mm to the inner skull surface, and all points outside it, were omitted. The resulting current density map was normalized by baseline (noise) levels with the dSPM method (Dale *et al.*, 2000), resulting in an F statistic representing activity each cortical location. Source localizations were summarized as smoothed images on the cortical surface (*w* format, Freesurfer) either at a particular time point or arithmetically averaged across time.

5.5.5. SPM8: multiple sparse priors (MSP)

A *multiple sparse priors* (MSP) (Friston *et al.*, 2008) approach to source reconstruction was implemented with the SPM8 graphical user interface (GUI) software (updated release r4290, April 2011, running on MATLAB R2010b (v7.11.0)); Litvak *et al.* (2011)). Unlike the MNE method described above, the MSP method assumes that brain activity can be sparse and bilateral and uses a hierarchical Bayes method to select the best-fitting priors depending upon the data (Friston *et al.*, 2008). This results in more focal source patterns than MNE (Henson *et al.*, 2009). MSP constrains sources to the cortical mantle and constrains their orientation to be normal. The released SPM8 GUI-based software does not have the option to directly utilize individuals' cortical surfaces for source space. Instead, an inverse-normalized template mesh is used (see above). Cortical mesh size used was 'normal' (total 8196 vertices i.e. 4098 per cortical hemisphere).

MEG data were bandpass filtered at 1-48 Hz; magnetometer and gradiometer signals were combined in inversion; a single compartment (brain only) model was used for BEM forward model estimation; inversion of MEG signals to brain sources was done with the MSP Greedy Search (GS) algorithm, which builds up successive combinations of source configurations until the model can no longer be improved (Friston *et al.*, 2006); the total number of patches in the MSP and the smoothness level were according to the default of spm8 release 4290 (Apr 2011); the results of inverse reconstruction were summarized as an image either at a particular time point or arithmetically averaged across time points (with a Guassian weighting, 8ms full width half maximum).

5.5.6. LCMV Beamformer

A linearly constrained minimum variance (LCMV) beamformer approach was used for MEG source reconstruction (OSL toolbox (version 1.1), developed by the OHBA Analysis Group, OHBA, University of Oxford; Woolrich *et al.* (2011)). This software was developed as an extension of the SPM8 software; therefore the source space was constructed from individual subject inverse-normalized template as for MSP analysis described above. The beamformer works by systematically testing the prediction of each source point and forward model on the observed data, whilst attenuating as far as possible the contributions of all other source points, like a spatial filter or *virtual beam*. This results in a model score map over the sampled brain source space. The beamformer tests the contribution of source throughout the brain volume, not just the cortical surface. A 4mm grid resolution was set for dipoles in the source space.

Parameters for evoked analysis: Data were bandpass filtered 3 - 40 Hz; magnetometer and gradiometer signals were combined for inversion; data was reconstructed from 0 to 200 ms post-stimulus onset to capture the evoked response.

Parameters for time-frequency analysis: Baseline was set 100 to 1 ms prior to stimulus onset; magnetometer and gradiometer signals were both used in inversion; a Hilbert transform was used for time-frequency decomposition. Data was reconstructed from 200 to 1000 ms post-stimulus onset to capture the ongoing induced oscillations.

5.5.6.1. General Linear Model (GLM) statistics:

The OSL beamformer toolbox uses trial-wise general linear model (GLM) regression to convert beamformer results into a t-statistic map. The t-statistics can then be used for significance testing of experimental effects at the individual subject level ('first level') or at the group level ('second level'). In the fitting of the GLM, the dependent variable (the activity score at each brain source in time) is regressed against a *design matrix* that indicates the experimental condition to which each trial belongs. This allows estimation of the mean and the standard deviation (std) of the *regression parameters* across trials, giving a least squares solution of how much of the source activity is explained be each regression parameter, or in other words, the effect of each experimental condition on MEG brain activity.

Contrasts of parameter estimates (COPEs) are linear combinations, or contrasts, of regression parameter estimates. For example, to evaluate the effect of experimental condition represented by parameter B1, let COPE = B1 and test the null hypothesis that COPE = 0. To evaluate any differential effect of two experimental conditions B1 and B2, let COPE = (B1 - B2), and test the null hypothesis that COPE = 0. At the individual subject level, the t-statistic is obtained from the estimated value and standard deviation of each COPE[1]:

$$t = \frac{\text{COPE}}{\text{std(COPE)}} \tag{5.1}$$

[1] In MEG beamforming, there is ambiguity regarding the directionality and sign of the activity at each source point; to over come this, parameter estimates must be rectified (converted to absolute value by the modulus function) at some stage during the analysis. Parameter estimates can either be rectified prior to contrast calculation, to form *coapes* (contrast of *absolute* parameter estimates) or the contrast itself can be rectified, to form *acopes* (*absolute* contrast of parameter estimates). Whether coape or acope is used depends upon the experimental test and prior hypothesis.

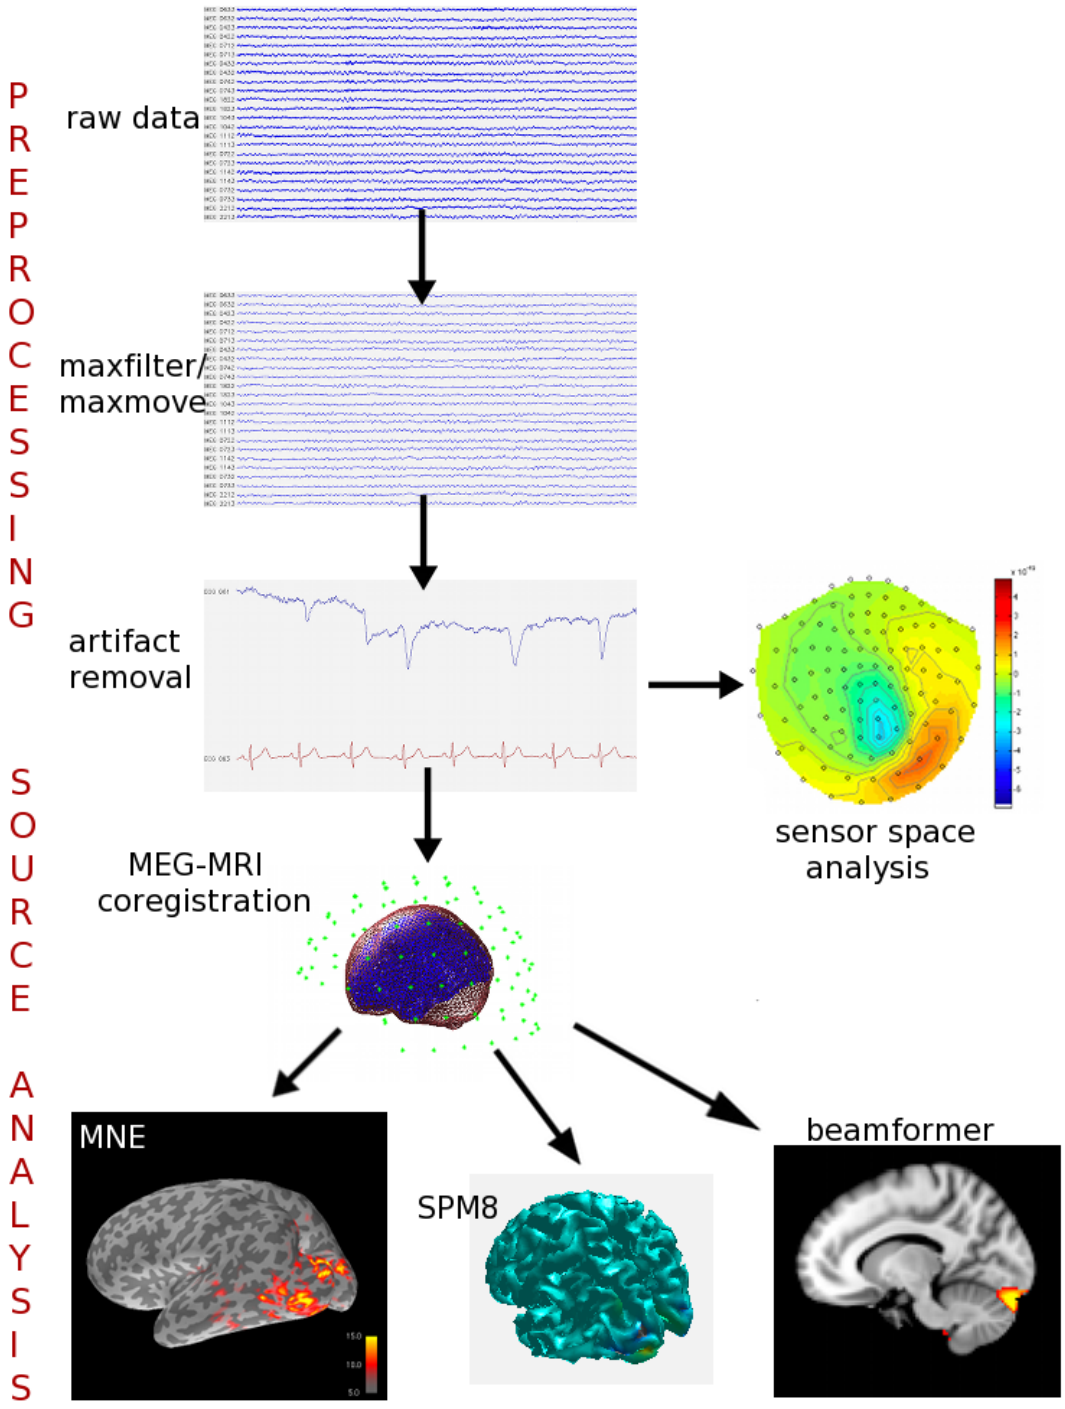


Figure 5.4.: Schematic summary of the MEG analysis pipeline from raw data preprocessing, through to source reconstruction.

6. MEG Retinotopy

6.1. Introduction

Retinotopy refers to the ordered, map-like representation of visual space in early visual cortical areas, the function of which has been extensively investigated with functional magnetic resonance imaging (fMRI). Retinotopic mapping is a fundamental part of fMRI research because it allows the delineation of boundaries between early visual areas, based upon map reversals and functional differences (Engel *et al.*, 1994; DeYoe *et al.*, 1994; Wandell, 1999). However, functional conclusions drawn from fMRI studies are limited because the measured blood-oxygenation level-dependent (BOLD) signal corresponds to changes in blood oxygenation, which follow changes in local neuronal activity. This is an indirect measure of neuronal function and has low temporal resolution relative to neuronal spiking (Logothetis *et al.*, 2001).

Magnetoencephalography (MEG) measures magnetic fields emitted by neuronal electrical activity (Chapter 5). It has high temporal resolution, limited by the sampling rate of the recording system (on the order of 1000 Hz). MEG therefore has the potential for investigating mechanistic questions about human sensory and decision-making systems that cannot be easily investigated with methods such as fMRI. However, MEG sensor activity must be projected to sources in the brain. Firstly, a **forward model**, or lead field matrix, is calculated which describes the effect of each brain source dipole on each sensor, using anatomical information about brain and headshape (skull and skin) geometries relative to the sensory array (Figure 6.1). Putative source activity patterns are passed through the forward model and optimized with respect to observed data. This reconstruction of MEG signals is known as the **inverse problem**, described by:

$$\mathbf{b}(t) = \mathbf{G}\mathbf{q}(t) + \mathbf{n}(t) \tag{6.1}$$

where **b**(t) represents the vector of measurements from the MEG sensors at each time point, **G** represents the forward model, **n**(t) represents noise, and **q**(t) represents the current ampli-

tude at each brain source. Thus **q**(t) is the unknown solution to the inverse problem (Hansen *et al.*, 2010). The inverse problem is *ill-posed* because many possible brain source configurations could give rise to the same pattern as 'seen' by the sensors. To reach a unique solution, additional constraints must be placed on the properties of brain sources, based upon prior assumptions about underlying brain activity. Different assumptions can be applied and accordingly there are many different approaches to MEG source reconstruction. In this study, I test the accuracy of different analysis approach for reconstructing the sources of MEG signals to retinotopic visual patterns.

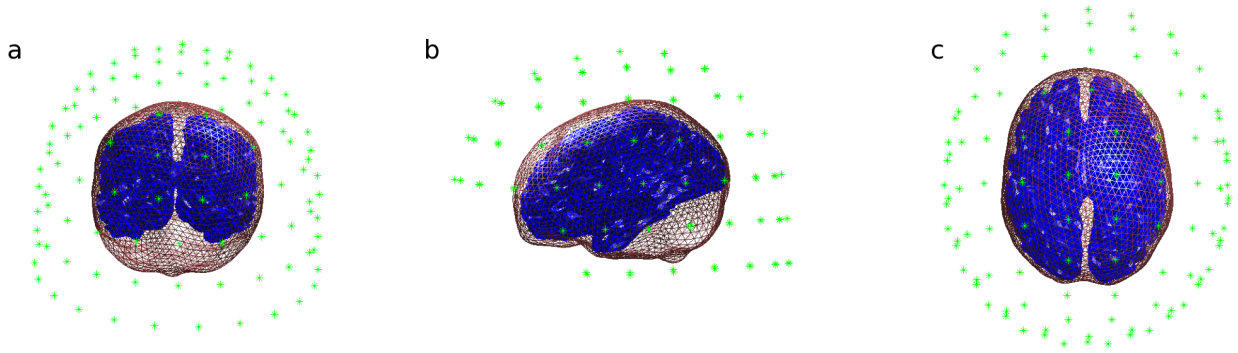


Figure 6.1.: The brain and headshape geomteries with respect to the MEG sensor array, necessary for computing the forward model for source reconstruction. Shown from the posterior (a), lateral (b) and superior (c) views.

6.1.1. Retinotopy in early visual cortical areas

Retinotopic maps in early visual cortical areas of the human brain follow a well-established pattern. For **polar angle** retinotopy, upper visual field locations are represented in ventral subregions, and lower field locations in dorsal subregions (Figure 6.2a). Left- and right-field locations are represented in the contralateral cortical hemisphere (Wandell, 1999). For **eccentricity**, foveal (central) locations are represented at the occipital pole and increasingly peripheral locations are represented more anteriorly through the calcarine sulcus and lateral side of the occipital cortex (Figure 6.2b). Source reconstructions of MEG responses to retinotopic stimuli are expected to follow this pattern; exact locations can be determined in each individual with fMRI.

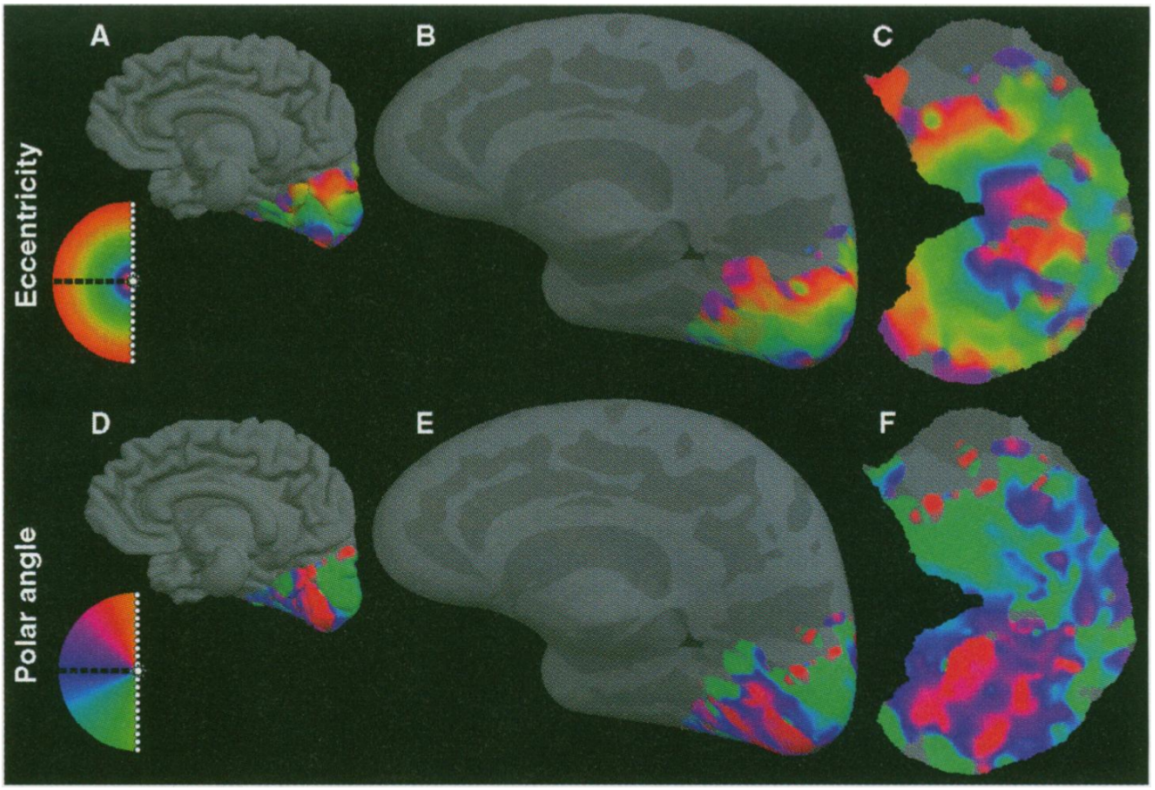


Figure 6.2.: Standard fMRI retinotopic representations in human visual cortex for a typical subject (from Sereno *et al.* (1995), with permission). The top row shows eccentricity coded by colour [red (fovea) → blue → green (parafoveal) → yellow → red (periphery)] displayed on the original cortical surface (**A**), the inflated cortical surface (**B**), and the flattened cortical surface (**C**). The bottom row show polar angle [red (lower vertical meridian) → blue (horizontal meridian) → green (upper vertical meridian)] plotted on the same three surfaces (**B, D, F**).

6.1.2. MEG source reconstruction methods

Three contrasting methods are compared in this study:

- Minimum Norm Estimate (MNE)
- Multiple Sparse Priors (MSP) in SPM8 analysis software
- Linearly constrained minimum variance (LCMV) beamformer

All methods discretize the the brain volume into a 3D grid and estimate the contribution of each source to the measured MEG signal. MNE and MSP are examples of *distributed* (imaging) methods because they simultaneously estimate activity at all possible source locations, by fitting source activity to the data with some set of prior assumptions. The beamformer will be referred to as a distributed source localization method, although technically it does not *localize* sources but rather estimates the signal at each brain location independently.

6.1.2.1. Minimum Norm Estimate (MNE)

MNE projects MEG brain activity onto a 3D source space grid derived from the individual's cortical mesh (or, if unavailable, derived from a template mesh). Source dipoles are constrained to lie within the cortex, oriented roughly perpendicular to the cortical surface, reflecting the assumption that the main origin of the MEG signal is cortical pyramidal cells (Dale & Sereno, 1993). Cortical surface and head shapes are used to estimate the forward model. A unique, linear solution to the inverse problem (Equation 6.1) is calculated with the main constraint being that amplitudes of source activity must be as small as possible on average (Hamalainen & Ilmoiemi, 1994). MNE minimizes the least squares estimation error between sensor signals predicted by the source solution and observed data. An implicit consequence of these assumptions is that results are biased towards small-amplitude, superficial, spatially spread currents, rather than large-amplitude, deep, focal currents. A depth-weighting parameter is used to correct for this superficiality bias (Lin *et al.*, 2006). Raw MNE results can be normalized with the *dSPM* statistical technique. This uses a modeled distribution of sources of noise, usually calculated from a baseline time window, to generate an F statistic for the activity level at each vertex (i.e. each point on the cortical surface) (Dale *et al.*, 2000).

6.1.2.2. Multiple Sparse Priors (MSP)

The multiple sparse priors (MSP) approach was developed as an alternative to the MNE assumptions above. MSP assumes that activity occurs sparsely and discretely, in patches, and

possibly with a bilateral correlation component (Friston *et al.*, 2008). A Bayesian procedure is used to iteratively decide which source components to include, depending upon the observed data, leading to a unique solution to the inverse problem. For example, if including an assumption of correlated bilateral sources does not significantly improve the fit to observed data, this component is switched off. Previous literature suggests that MSP is superior to MNE-like methods for localizing EEG and MEG visually-evoked responses to face stimuli (Friston *et al.*, 2008; Henson *et al.*, 2009). However, these studies did not quantitatively evaluate localization accuracy.

Like MNE, MSP uses anatomical information for the forward model and a cortical constraint for source dipoles. This is based upon an *inverse-normalized* template head and brain mesh, rather than the individual's cortex and headshape surfaces themselves (Mattout *et al.*, 2007). Firstly, the individual's anatomical MRI is normalized to the canonical MNI152 template brain. Then, the inverse of this normalization field is applied to a high resolution cortical mesh derived from the same template brain. This produces a high resolution mesh that is morphed to the individual's brain and headshape. Previous studies using simulated (Mattout *et al.*, 2007) and real data (Henson *et al.*, 2009) for responses to face stimuli found no evidence that the inverse-normalized mesh results in worse source reconstructions than the individual's mesh.

6.1.2.3. Beamforming

In contrast, the beamformer applies an adaptive spatial filter to the source space, such that the activity timecourse at each source location is systematically estimated independently of all other sources. The spatial filter is made up of a set of weights, $\mathbf{w}_i$, which define a weighted sum of the signals at all sensors, based upon the forward model, and which selectively enhance activity arising from the target source whilst minimizing interference from all other locations. This multichannel signal processing technique was originally developed for radar and antenna arrays (van Veen & Buckley, 1988) Different versions have been adapted for MEG and EEG source analyses (van Veen *et al.*, 1997; Robinson & Vrba, 1999; Hillebrand & Barnes, 2011). In this study, I apply a linearly constrained minimum variance (LCMV) beamformer, which determines the spatial filter weights $\mathbf{w}_i$ for a source $\mathbf{q}(t)$ at a given location $\mathbf{G}_i$ according to:

$$\mathbf{w}_i = \mathbf{C}^{-1}\mathbf{G}_i^T[\mathbf{G}_i^T\mathbf{C}^{-1}\mathbf{G}_i]^{-1} \tag{6.2}$$

where $\mathbf{C}^{-1}$ is the inverse of the data covariance matrix (Hansen *et al.*, 2010).

The LCMV beamformer was implemented with the OSL toolbox, based upon SPM8 software (Woolrich *et al.*, 2011).The forward model is estimated using the same inverse-normalized head-shape mesh as MSP. Sources are sampled throughout the brain volume. In the OSL toolbox, a generalized linear model (GLM) approach is used to calculate t statistics, from linear combinations of brain activity estimates in single experimental conditions or between conditions (see Chapter 5). The beamformer can also utilize time-frequency information to investigate whether power in a particular frequency band varies according to experimental variables. In this study, I investigate whether using frequency information about the stimulus-evoked response aids beamformer localization.

6.1.3. Previous MEG studies of visual retinotopy

Previous studies have investigated the efficacy of MEG source reconstructions of retinotopic responses in early visual areas, but no study has yet compared different analysis approaches directly. A number of studies used a single equivalent current dipole to model early MEG responses as a single brain source (Chen *et al.*, 2005; Nakamura *et al.*, 1997; Nishiyama *et al.*, 2004; Stephen *et al.*, 2002). They showed that the resultant single dipole is located in or close to the subregion of primary visual cortex that retinotopically represents stimulus location. However, as we know from fMRI (e.g. Wandell *et al.* (2005)), a single dipole source is not a true representation of the activations for the stimuli used, and the performance of *distributed* MEG localization methods needs to be evaluated.

A few studies have used custom-developed MEG analysis methods to localize retinotopic visual activity. Ioannides and colleagues compared MEG sources, reconstructed with a magnetic field tomography (MFT) approach, to fMRI in individual subjects, and report that MEG and fMRI activation centres were generally within 2-5 mm of each other in visual area V1 (Moradi *et al.*, 2003; Poghosyan & Ionnides, 2007). Cottereau *et al.* (2011) localized MEG signals with a frequency-based analysis and compared this to fMRI in one subject; 74.5% of activated cortical vertices were located in the correct subregions of early visual areas. However, these methods are not freely available for application by other researchers. Other studies used fMRI retinotopy to constrain MEG sources and with this successfully separated signals between adjacent visual areas (Yoshioka *et al.*, 2008; Hagler *et al.*, 2009). Methods that incorporate *a priori* fMRI constraints were not evaluated here because it is not always be possible to have fMRI priors for a given experimental question in visual research.

Sharon *et al.* (2007) compared MNE localization of retinotopic stimuli from MEG recordings alone versus MEG with simultaneous EEG. Accuracy was evaluated against fMRI in the same individuals by computing the distance between the peaks of fMRI and M/EEG sources in the 3D brain volume. MEG with simultaneous EEG were more accurate by 2.5 mm compared to MEG alone, most likely due to their independent measurements of noise. For MEG alone, peak MEG activation was 10.8 ± 0.8 mm away from peak fMRI source, suggesting that visual retinotopic localizations with MNE for MEG alone are reasonably accurate.

Brookes *et al.* (2010) used a beamformer to localize responses to a checkerboard wedge stimulus to occipital cortex in a pattern that seems consistent with known retinotopy, but accuracy was not quantified or compared to fMRI maps. Perry *et al.* (2011) used a beamformer to localize gamma-band (30 - 60 Hz) responses to a checkerboard wedge which rotated clockwise around the visual field. Viewed posteriorly, resultant brain sources moved clockwise in relation to stimulus movement, consistent with the retinotopy of primary visual cortex (V1). However, many sources were not located near the anatomical location of V1 (i.e. in the calcarine sulcus) when compared to individuals' cortical anatomies (their Figure 4). This suggests that localization accuracy with the beamformer in this frequency band is poor. There are no publications for the MSP method with retinotopic visual stimuli.

In this Chapter, I evaluate MNE, MSP and beamformer localization accuracy for retintopic visual activity by comparing MEG sources with fMRI maps in the same individuals. The results reveal which reconstruction assumptions may be best employed for MEG studies of human visual processing.

6.2. Methods

MEG and structural MRI analyses are described in *Chapter 5: MEG Materials & Methods*, which should be used as a reference. Study-specific details are described here.

6.2.1. MEG Retinotopy

6.2.1.1. Subjects and stimuli

Black-and-white chequerboard quadrant stimuli were presented with a Cambridge Research Systems VSG 2/5 graphics generator run with a Dell laptop (subjects *mcp1*, *mvp1*, *mvp2*, *mvp3*), or with Presentation® (Neurobehavioural Systems, Inc.) running on a Samsung R710 laptop (Centrino 2 P7450 processor, nVIDIA GeForce 9300M graphics card) (subjects *mrt2*, *mrt3*). Stimulus parameters were identical for both presentations. Each quadrant extended 0 to 5.4° eccentricity, presented either in the upper left, upper right, lower left or lower right visual field . Quadrants contained 6 checks along the radius and the arc, decreasing in size by a factor of $1/d$, where d is distance to apex. A black fixation point was present at the apex. Michelson contrast of the stimulus was 99%. Quadrants were presented for 1000 ms with no inter-stimulus interval.

Black-and-white checkerboard concentric ring and mini-ring stimuli were presented with Presentation® software, as above, for all subjects. Rings consisted of 24 checks along their circumference and 6 checks along their radius, presented at three eccentricities: FOV (0° - 1.5°), PARAF (2.0° - 4.0°) and PERI (6.0° - 11.0°). These eccentricity bands were selected to activate approximately equal areas of cortex due to foveal magnification (Wandell *et al.*, 2005). Mini-rings were formed from rings by masking out all but either the top-right or lower-right visual field quadrant. This gave 6 mini-rings: upper field FOV, PARAF and PERI and lower field FOV, PARAF and PERI. Ring and mini-ring eccentricities were presented in a pseudo-randomized order with a variable inter-stimulus interval (ISI; 600, 800 or 1000 ms, selected pseudo-randomly).

All stimuli flickered with 4Hz contrast reversal and were presented on a mid-gray background (mean luminance, 25 cd/m^2).

6.2.1.2. Data Acquisition & Preprocessing

Participants passively viewed stimuli in the MEG scanner whilst maintaining fixation. Datasets for quadrants was recorded for 6 volunteers (*mcp1*, *mrt2*, *mrt3*, *mvp1*, *mvp2*, *mvp3*). Quadrants were presented in blocks of 25 full-cycle rotations (upper right, upper left, lower left, lower right fields); 6 blocks were collected per participant. Datasets for rings were recorded for 5 participants (*mcp1*, *mrt2*, *mrt3*, *mvp1*, *mvp2*); 4 blocks of 150 stimuli were collected per participant. Datasets for mini-rings were recorded for 3 participants (*mcp1*, *mrt2*, *mvp1*); 5 blocks of 180 stimuli were collected per participant. After preprocessing (for details, see Chapter 5), >100

trials per stimulus remained per participant.

6.2.1.3. MNE source reconstruction

Data were analyzed with MNE software (Minimum Norm Estimate, MGH/MIT Martinos Centre for BioMedical Imaging; Hamalainen & Ilmoiemi, 1994; Dale *et al.*, 2000), time-locked to stimulus onset and averaged. A noise covariance matrix (NCM) was calculated from -500 to 0 ms prior to each stimulus onset. For quadrants, this window necessarily comprised the final 500 ms of the preceding stimulus. This implies that some genuine brain responses were included in the NCM calculation; however, all stimulus types would have been included in the NCM and this would not be expected to bias sources to any particular location.

Source reconstructions were performed on data bandpass filtered 1 - 40 Hz for 0 - 1000 ms post-stimulus, combining magnetometer and gradiometer measurements. Anatomically constrained dSPM inverse solutions (F-statistics) were generated at each cortical vertex(Dale *et al.*, 2000). The time window of interest for the initial visual-evoked response was centered on the ascension of the first response peak at 83 ms (see Figure 6.6). dSPM source estimates were averaged across a 20 ms time window (73 - 93 ms) and saved to *w file* cortical surface format (Freesurfer, Fischl (2012)).

6.2.1.4. LCMV beamformer source reconstruction: initial evoked response

MEG data were analyzed using the LCMV (linearly constrained minimum variance) beamformer from the OSL toolbox (Oxford Centre for Human Brain Activity, Oxford; van Veen *et al.*, 1997; Woolrich *et al.*, 2011). MEG data were not averaged; instead, the beamformer was used to reconstruct sources separately for each trial. A t statistic was generated at each source location. For quadrant stimuli, the t value described the difference between responses to a particular quadrant compared to all other quadrants, as no inter-stimulus interval baseline was available. For rings and mini-rings, the t value described the difference between responses to the stimuli versus baseline activity -200 to 0 ms prior to stimulus onset. Source were reconstructed for 0 - 1000 ms post-stimulus, bandpass filter 1 - 40 Hz, combining magnetometers and gradiometers. Resultant t-statistic 3D images in MNI152 standard space were averaged across a time window of 73 - 93 ms post-stimulus.

6.2.1.5. LCMV beamformer source reconstruction: time-frequency decomposition

Time-frequency decomposition was performed within 7 - 9 Hz, centered on 8Hz which is the 2nd harmonic of the stimulus contrast-reversal frequency. The 2nd harmonic is used because each contrast reversal of the stimulus actually involves two contrast *changes*, for example from black to white and then from white back to black, and visual brain areas respond to each contrast change (Figure 6.6, see also Cottereau *et al.* (2011)). Time window for source reconstruction was 200 - 1000 ms. Resultant t-statistic images were averaged across the whole time window to give the final localization; otherwise, all parameters were identical as for the initial evoked response beamforming (above).

6.2.1.6. SPM8 MSP source reconstruction

MEG data were analyzed using the multiple sparse priors (MSP) approach from the SPM8 neuroimaging analysis package (FIL Methods Group, UCL; Friston *et al.*, 2008). Source were reconstructed in three different ways:

(i) Averaged data, no statistics: The averaged MEG data utilized for MNE source reconstruction were directly converted to SPM8 (Matlab) format for use in MSP source reconstruction. The time-window of source reconstruction was 63 - 103 ms, following the SPM8 manual advice to focus source reconstruction on the main activity of interest (in this case the early visually-evoked response). Information from magnetometers and gradiometers was combined. Source results, in MNI152 standard brain space, were averaged over 63 - 103 ms, weighted by a Gaussian centered on 83 ms, as the SPM8 software did not offer any alternative way to generate time-averaged source localization result outputs.

(ii) Trial-wise data, one-sample t-test: As the localization results from the above approach were very poor (see Results), alternative approaches, based upon the beamformer procedure, were attempted. Firstly, instead of averaged data, the non-averaged data used for beamformer source reconstruction was used for trail-wise MSP source reconstruction. Localization results were averaged over 63 - 103 ms post-stimulus for each trial. A t-test was applied across trials to indicate how different the activity estimates at each cortical location were from zero (one-sample t-test, across-trial variance). All other parameters were identical to *(i)*.

(iii) Trial-wise data, t-test contrast across stimuli: Same as *(ii)* except for each stimulus, a t-test was applied across the trials to indicate how the activity estimates, at each cortical location,

differed between the stimulus-of-interest and all other stimuli in the set (quadrants, rings or mini-rings; two-sample t-test, unequal variance assumption). All other parameters were identical to *(ii)*.

6.2.1.7. Morphing 3D images to cortical surface

The beamformer and MSP results in MNI152 volumetric standard space were converted to cortical surface format (*w file*, Freesurfer) to enable comparison with fMRI retinotopic maps. The *flirt* command from FSL (FMRIB, Oxford) generated a transformation matrix from MNI152 to Freesurfer template volumetric space and transformed the MEG source reconstructions to Freesurfer volume space. Freesurfer volume images were then converted to cortical surface format *w files* with *mri_vol2surf* (Freesurfer). Finally, the surface files were morphed to the cortical anatomy of the individual subject with *mri_surf2surf* (Freesurfer).

6.2.2. Retinotopy: fMRI

6.2.2.1. Stimuli

Retinotopic quadrant and ring stimuli used for fMRI data collection were presented with the Cambridge Research Systems VSG 2/5 graphics generator with a Dell laptop. Visual stimulus parameters were identical to those used for MEG. The quadrant stimulus rotated through 30° every TR (4000 ms), producing 'traveling wave' responses in the brain necessary for analysis with standard fMRi retintopy software (Wandell, 1999). Concentric rings expanded every TR (4000 ms), taking 8 steps to cover the visual field 0° to 23 °. Therefore, although the timing of the movement of the retinotopic visual stimuli differed between the fMRI and MEG data acquisitions, the spatial positions/angles of the stimuli in the visual field were identical enabling direct comparisons between the brain source activation measurements.

6.2.2.2. Data Acquisition

Retinotopic fMRI data were acquired according to standard methods with a 3T Tesla whole-body TIM Siemens Trio scanner and a 12-channel receive-only head coil, located at the University of Oxford Centre for Clinical Magnetic Resonance Research (OCMR), by collaborators Dr Holly Bridge, Dr Lori Minini or Dr Betina Ip. EPI sequence parameters were: TE = 30 ms; TR = 4000 ms; 40 2-mm slices; 2 x 2 mm in-plane resolution; matrix = 64 x 64. Each run consisted of 72 volumes acquired continuously (288 s). A reduced (40 2-mm slices) T1-weighted image (3D

FLASH) was also included in each functional session, acquired coronally at an in-plane resolution of 1 x 1 mm^2. These slices were in the same planes as the retinotopic functional images, and were used to register functional retinotopy data to the whole brain structural MRI.

6.2.2.3. Analysis & ROIs

Retinotopic maps were analyzed for individual subjects according to standard procedures using either mrTools software (HeegerLab; `http://www.cns.nyu.edu/heegerlab/wiki/`) or mrVista software (Standford; `http://white.stanford.edu/software`). Retinotopic BOLD activity maps were displayed on flat renderings of the occipito-temporo-parietal region, allowing borders between visual areas to be identified and traced. For angular retinotopy, dorsal (lower field) and ventral (upper field) subregions were defined on the left and right hemisphere for areas V1, V2 and V3. Area V3A was also defined on each hemisphere (Figure 6.3). For eccentricity, the mapping colour guide was used to delineate regions representing the FOV, PARAF and PERI eccentricity bands (Figure 6.4). Maps were converted to Freesurfer *w file* surface format for comparison to MEG. For subjects *mvp1*, *mvp2* and *mvp3*, the polar angle ROIs and their subregions were delineated by Dr Lori Minini. All other ROIs and subregions were delineated by myself.

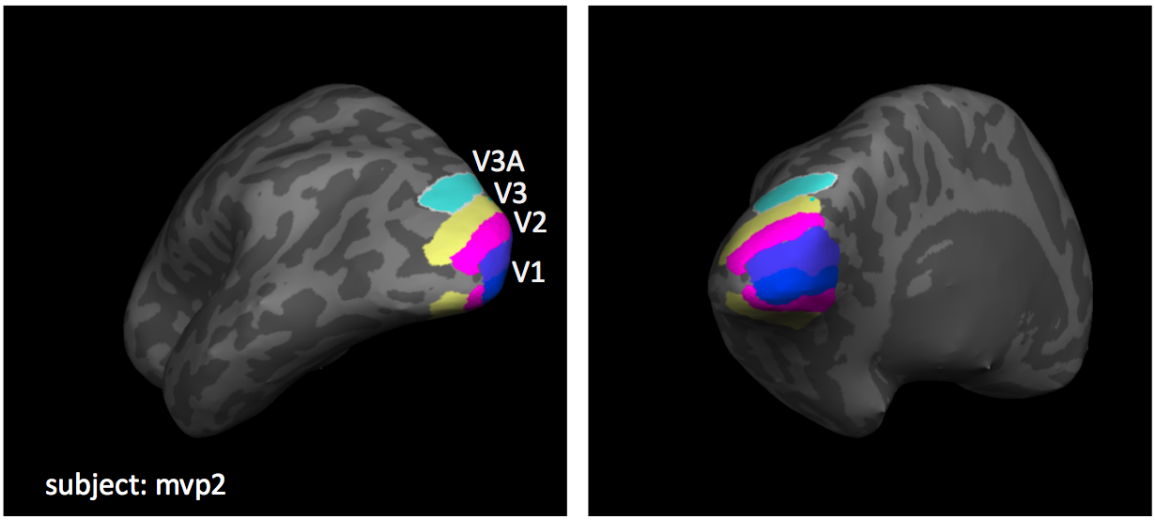


Figure 6.3.: Regions of interest (ROIs) defined from fMRI retinotopy: dorsal and ventral V1 (*indigo/blue*), V2 (*pink*) and V3 (*yellow*), and area V3A (*light blue*). These ROIs were used for polar angle MEG retinotopy source accuracy evaluation. Shown on inflated cortex (*left hemisphere*) of an example subject.

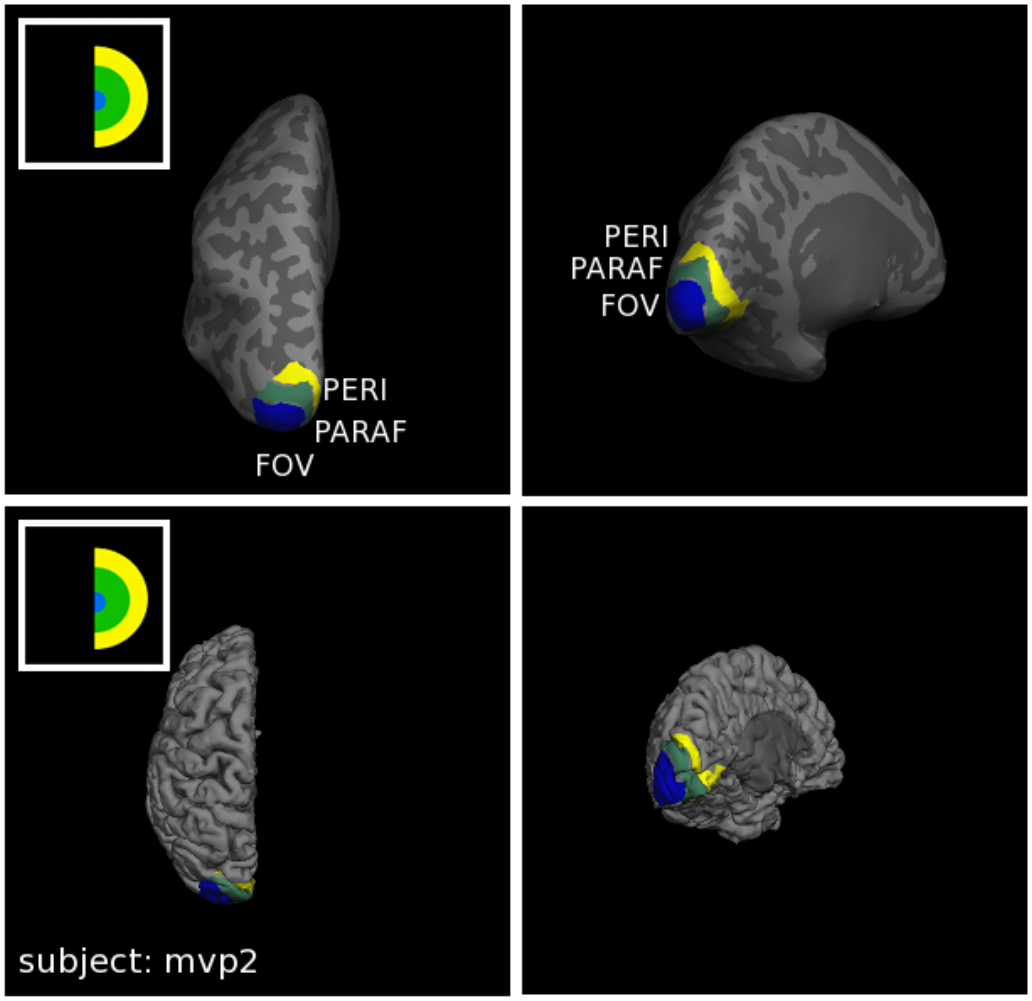


Figure 6.4.: Right visual field eccentricity bands defined from fMRI retinotopy, used for eccentricity (rings) MEG retinotopy analysis. FOV locations are shown in *blue*, PARAF in green, PERI in *yellow* (see *insets*). Superior (*left panels*) and posterior-medial (*right panels*) views are shown for both the inflated (*upper panels*) and folded (*lower panels*) left cortical hemisphere.

6.2.3. MEG-fMRI Comparison

Comparison between MEG source locations and fMRI retinotopy was done in Matlab with custom scripts. To evaluate localization accuracy, I calculated the percentage of total 'active' cortical vertices, present inside a pre-defined region-of-interest (ROI), that were located in the retinotopically expected subregion of that ROI, separately for each subject and stimulus. For example, to evaluate localization accuracy for an upper right quadrant in area V1, I calculated what percentage of all active vertices within V1 were located inside the left ventral subregion, which is the retinotopically expected subregion for that stimulus. Localization accuracies for quadrants were averaged across stimuli and subjects for each early visual area ROI. For concentric rings, eccentricity band ROIs corresponding to stimulus eccentricities were defined across areas V1, V2 and V3 combined. Of the active vertices located across all bands, I calculated the percentage located in the retinotopically-expected band, separately for each cortical hemisphere for each subject. Results were averaged within each eccentricity band. The same process was used for mini-rings, where retintopic subregions were defined both by upper/lower field location and stimulus eccentricity.

One-sample t-tests were used to calculate whether MEG sources localized into expected subregions better than chance ($p < 0.05$). For quadrants, chance level was 25% for visual areas V1, V2 and V3, which have four angular subregions each (dorsal and ventral subregions in the left and right hemisphere), and 50% for V3A which has two subregions only (left and right hemisphere). For rings, chance level was 33% as there were three eccentricity bands, each of approximately equal area. For mini-rings, chance level was 17%, as there were six possible stimulus subregions of roughly equal area. Bonferroni multiple-comparison correction was applied for testing between MNE, SPM8's MSP and beamformer methods for the initial, visually-evoked response. A one-way ANOVA tested whether localization accuracy was different between the MEG source reconstruction methods.

If a particular stimulus resulted in zero active vertices in early visual areas for a particular subject, then that case was excluded from the accuracy analysis in order to focus analysis on those brain responses that were localized to early visual areas rather than on failed localizations (Figure 6.5).

6.2.4. Threshold-accuracy functions

In contrast to single dipole MEG localization approaches, the distributed analysis methods used in this study output an activity value for each vertex (point) on the cortical surface. This is because each cortical location has some probabilistic contribution to the MEG sensor signal due to the inherent uncertainty of the inverse problem. An activation, or cut-off, threshold must therefore be designated which controls whether or not a vertex is considered 'active'. There is some arbitrariness in designating such a threshold, and a non-systematic designation may affect the retinotopic MEG-fMRI comparisons in unexpected ways and render unfair the comparison between MEG methods. I defined the *threshold* as the percentage of highest-responding vertices across the cortex which to designate 'active'. For example, a threshold of 1% means only vertices with activity values in the top 1% are designated 'active'. I systematically calculated localization accuracy as a function of threshold, for each MEG analysis method, averaged across all subjects' visual areas (area V3A was not included because it was not activated as reliably as the other regions; Figure 6.5). For a given threshold, if a stimulus resulted in zero 'active' vertices in early visual areas for a particular subject, the accuracy result for that case was set to zero. The threshold that produced the most accurate source reconstructions in the threshold-accuracy function was designated the 'optimal' threshold and was used for the main accuracy evaluations. The same threshold was applied to all subjects for a given reconstruction method.

Figure 6.5. *(following page)*: Total number of rejected cases due to zero active vertices localized to early visual areas, for quadrant (a), ring (b) and mini-ring (c) stimulus sets. MNE = minimum norm estimate; MSP = multiple sparse priors; BF = beamformer. The number of exclusions is generally low across all methods and stimulus sets, except for quadrants where the MNE method failed to localize any active vertices in 16/24 cases for area V3A.

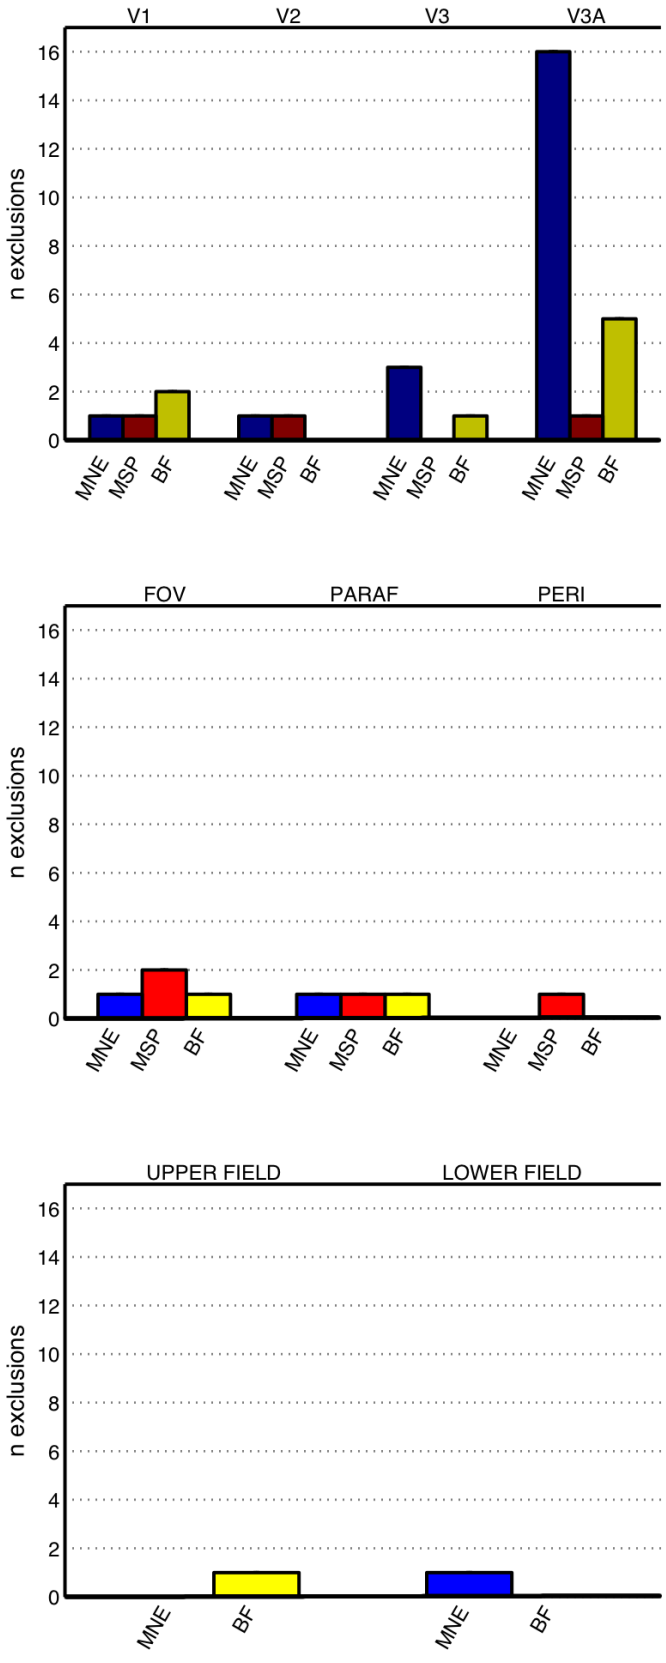

V1
V2
V3
V3A
n exclusions
0
2
4
6
8
10
12
14
16
MNE
MSP
BF
FOV
PARAF
PERI
UPPER FIELD
LOWER FIELD

Figure 6.5

6.3. Results I: Polar Angle (Quadrants)

Human brain responses to contrast-reversing retinotopic quadrant visual stimuli were measured with magnetoencephalography (MEG) in six participants. Sources of the MEG signals were localized with three different source reconstruction methods, and compared to expected locations as discerned from retinotopic fMRI in the same individuals.

6.3.1. Visually evoked responses at sensor level

In response to the onset of quadrant stimuli, brain activity showed a large deflection at 60 - 100 ms post-stimulus: the 'initial evoked response' (Figure 6.6). Further evoked responses to the 4Hz contrast reversals could be seen throughout the stimulus duration. Scalp topography of sensor activity showed the expected cruriform retinotopic pattern: responses to upper field stimuli were located below the responses to lower field stimuli and responses to right-field stimuli were located further left than responses to left-field stimuli (Figure 6.7).

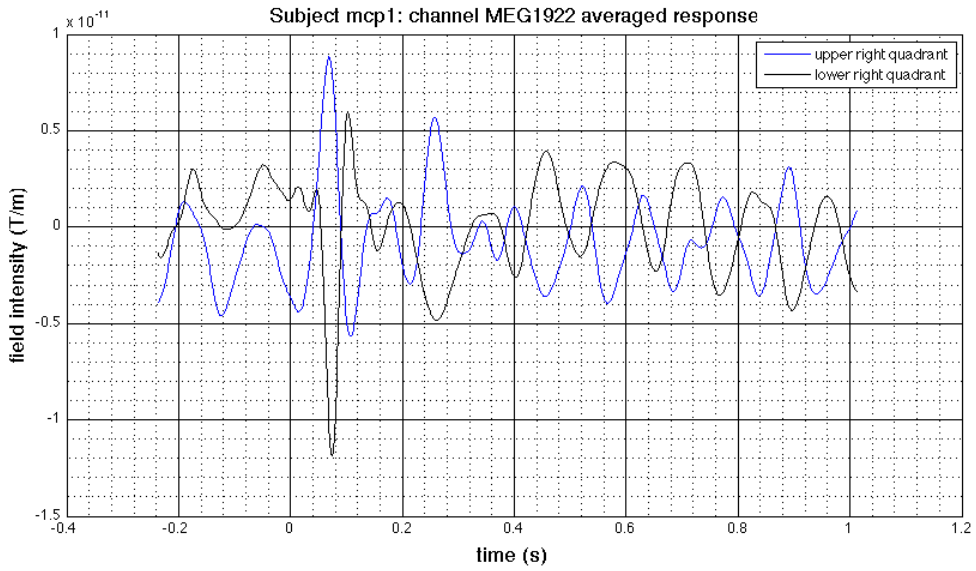


Figure 6.6.: Responses to upper and lower right quadrants, recorded by occipital channel MEG1922 of example subject (*mcp1*). Traces are time-locked to the onset of the visual stimulus (time = 0) and averaged. Evoked deflections to the upper and lower quadrants show opposite polarities, as expected from oppositely-oriented current sources in the lower and upper calcarine banks (Wandell et al., 2005). MEG source reconstructions were performed either upon the initial evoked response peaks (73 - 93 ms) or upon the ongoing stimulus-induced oscillations (200 - 1000 ms).

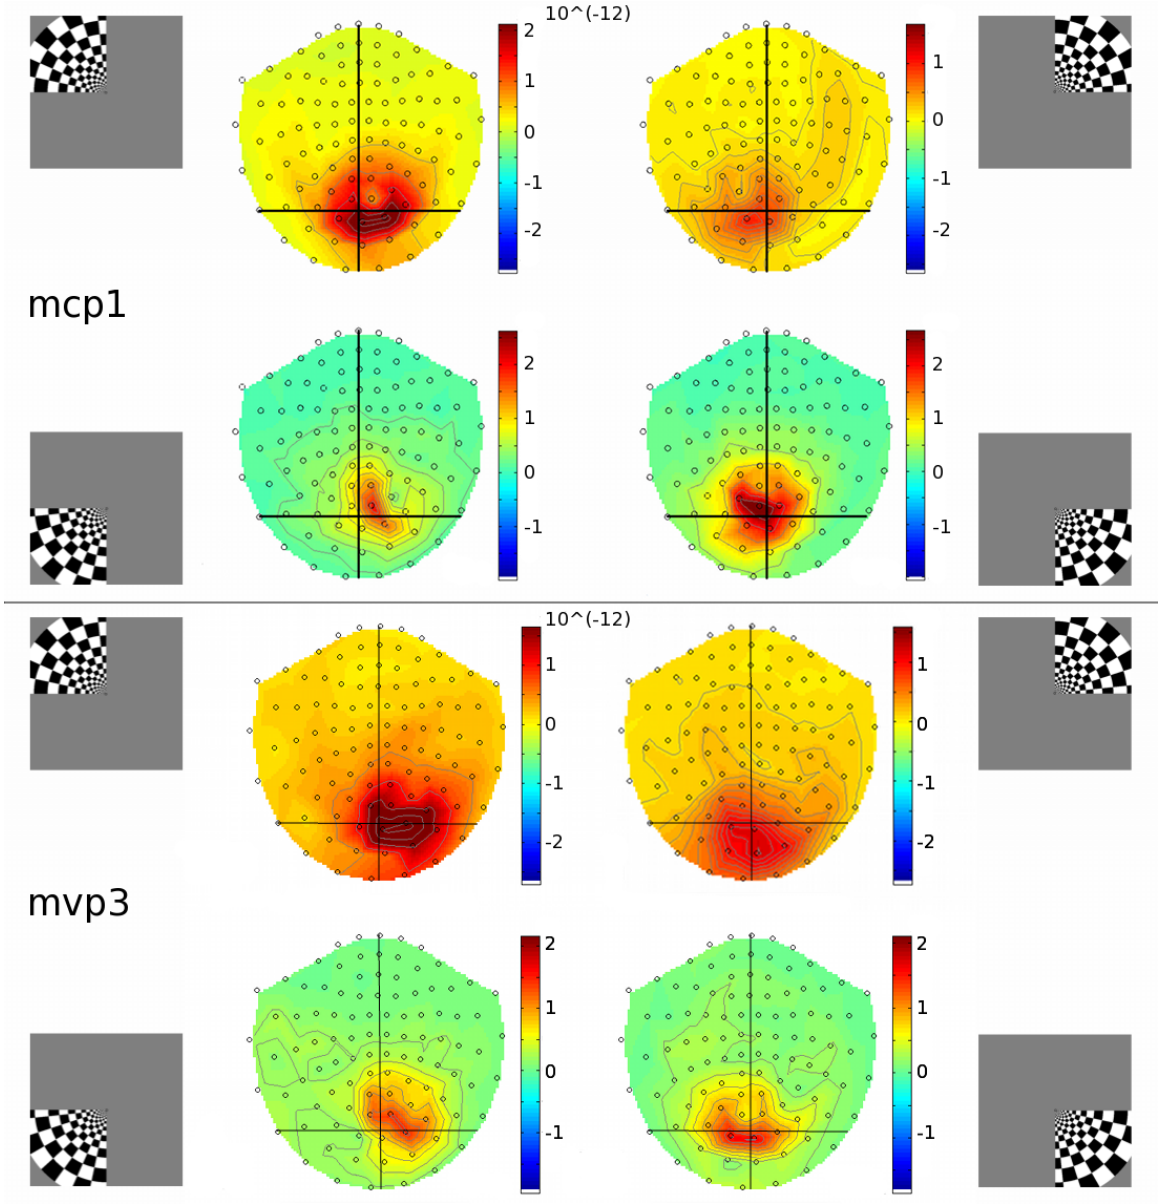


Figure 6.7.: Raw gradiometer sensor topography (T/m) of the visual evoked response at 83 ms, for subjects *mcp1* and *mvp3*, for each quadrant location. Colour scales are normalized to the peak response to each stimulus. Insets show the quadrant locations. This pattern was consistent across all subjects.

6.3.2. Evaluation of three source reconstruction methods

Three MEG source reconstruction methods (MNE, MSP, beamformer) were used to localize cortical sources of visually evoked responses to quadrant stimuli. MNE sources were consistently well-lateralized to the hemisphere contralateral to the quadrant and there was a consistent ventral-dorsal distinction for lower and upper field stimuli respectively. MSP and beamformer localizations generally followed this pattern but contained some deviations (Figures 6.8 & 6.9). For example, the MSP result for the upper left quadrant for *mvp3* localized sources bilaterally to the upper part of the occipital lobe, instead of unilaterally on the lower right (Figure 6.9). The percentage of active vertices located in the expected sub-region of each early visual area region of interest (ROI) was significantly better than chance for all ROIs with the MNE method, but only significantly better than chance for 3 of 4 ROIs for MSP and beamformer methods (Figure 6.10, Table 6.1). The MNE method gave a consistently higher percentage of active vertices localized to the correct subregion, compared to the other methods. A one-way ANOVA confirmed that the accuracy of the different methods was significantly different for early visual areas V1, V2 and V3 (F = 7.3, $p \ll 0.001$) and for area V3A (F = 7.3, p = 0.0183). As some of the underlying data were not normally distributed (Lillifors test, $p < 0.05$), these results were confirmed with a non-parametric test (Kruskal-Wallis: V1-V3: χ^2 = 17.1, $p \ll 0.001$; V3A: χ^2 = 17.1, p = 0.0032).

	% of active vertices in the expected subregion of ROI		
ROI (chance level):	MNE	MSP	beamformer
V1 (25%)	66.1 (±18.7)	49.9 (±29.0)	54.3 (±30.1)
V2 (25%)	75.2 (±29.4)	53.4 (±32.8)	47.2 (±37.1)
V3 (25%)	89.3 (±30.1)	52.2 (±41.2)	68.3 (±34.0)
mean (V1, V2, V3)	*76.5 (±27.8)*	*51.9 (±34.3)*	*56.5 (±34.8)*
V3A (50%)	100 (±0)	71.7 (±35.9)	91.8 (±20.6)

Table 6.1.: Localization accuracy for the MNE, SPM8 MSP and LCMV beamformer source reconstruction methods, averaged across subjects and quadrants. The table shows the percentage of active vertices in each early visual area that are localized in the expected subregion according to fMRI retinotopy in the same individual (± S.D.)

Figure 6.8. *(following page)*: Source reconstructions of MEG-measured early evoked responses to quadrants for example subject *mcp1*. MNE results are displayed on the individual subjects' inflated cortical surface (Freesurfer software); MSP and beamformer results are displayed volumetrically on the MNI152 standard brain.

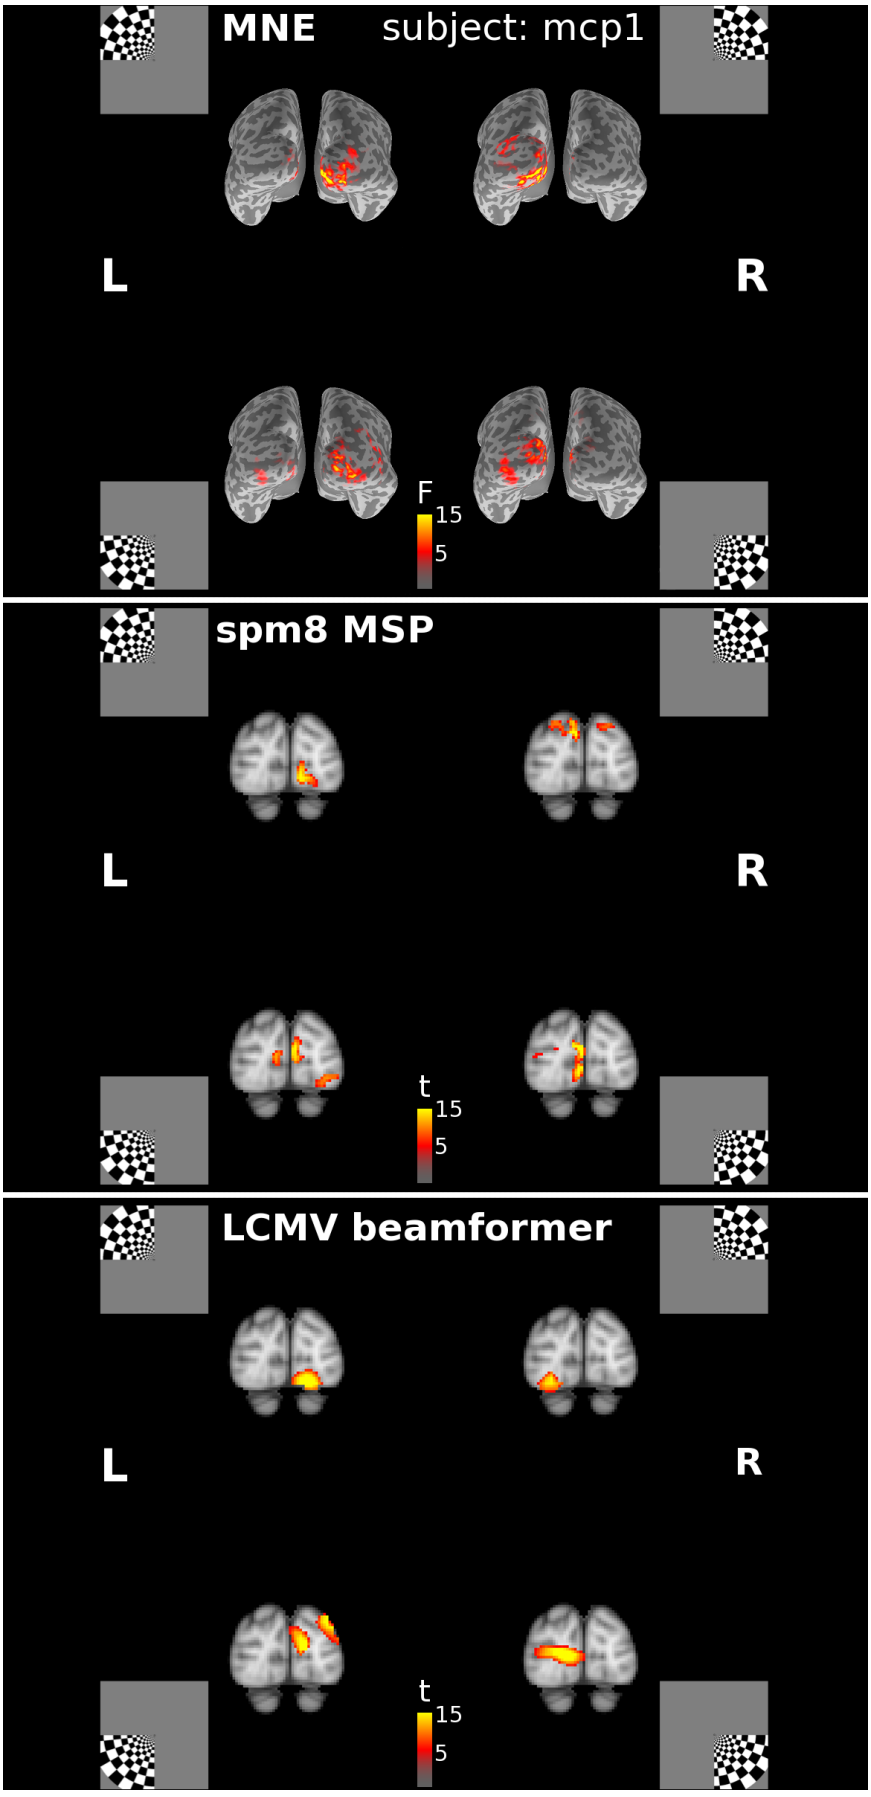

MNE
subject: mcp1
L
R
F
15
5
spm8 MSP
L
R
t
15
5
LCMV beamformer
L
R
t
15
5

Figure 6.8

6.3.3. Time-frequency beamformer: stimulus-induced oscillations

The stimuli underwent contrast reversal at a rate of 4Hz, evoking an ongoing, stimulus-induced oscillation in brain responses (Figure 6.6). The period of the evoked oscillations was 8Hz, the second harmonic of the contrast reversal frequency. In a recent study, Cottereau *et al.* (2011) suggested that induced oscillation of this kind could be utilized to accurately reconstruct MEG sources: in the one subject that they tested, on average 74.5% of MEG sources were located within the expected visual area subregions for V1, V2 and V3 (their Supplementary Data). My results for localizations of the initial evoked response peak with MNE (76.5%) compares well to their value, but results for MSP and beamformer methods were lower (Table 6.1). I investigated whether utilizing the frequency tag in the stimulus improved retinotopic source reconstruction with the LCMV beamformer. A one-way ANOVA was used to test whether accuracy was different between evoked beamformer (BF) and time-frequency beamformer (BF TF), averaged across subjects and quadrants.

Both the evoked LCMV beamformer (BF) and the time-frequency beamformer (BF TF) localized sources into the correct subregion significantly better than chance for 3 of 4 ROIs (Figure 6.11 & Table 6.2), but overall there was no significant difference between the two approaches (one-way ANOVA: V1-V3: F = 0.12, p = 0.729; V3A: F = 1.28, p = 0.265; Kruskal-Wallis: V1-V3: χ^2 = 0.02, p = 0.884; V3A: χ^2 = 0.44; p = 0.506). In this case, the use of the stimulus frequency-tag resulted in source reconstructions that were as good as, but not significantly better than, the use of the initial peak evoked response.

6.3.4. SPM8's MSP: averaged responses, trial-wise responses, and statistical contrast

Figure 6.10 and Table 6.1 show that the SPM8 MSP method localized sources to the correct subregion better than chance (25%) for 3 of 4 early visual areas, with overall accuracy of 51.9%. This result was based upon an analysis approach that applied MSP source reconstruction trial-

Figure 6.9. *(following page)*: Source reconstructions of MEG-measured early evoked responses to quadrants for example subject *mvp3*. MNE results are displayed on the individual subjects' inflated cortical surface (Freesurfer software); MSP and beamformer results are displayed volumetrically on the MNI152 standard brain.

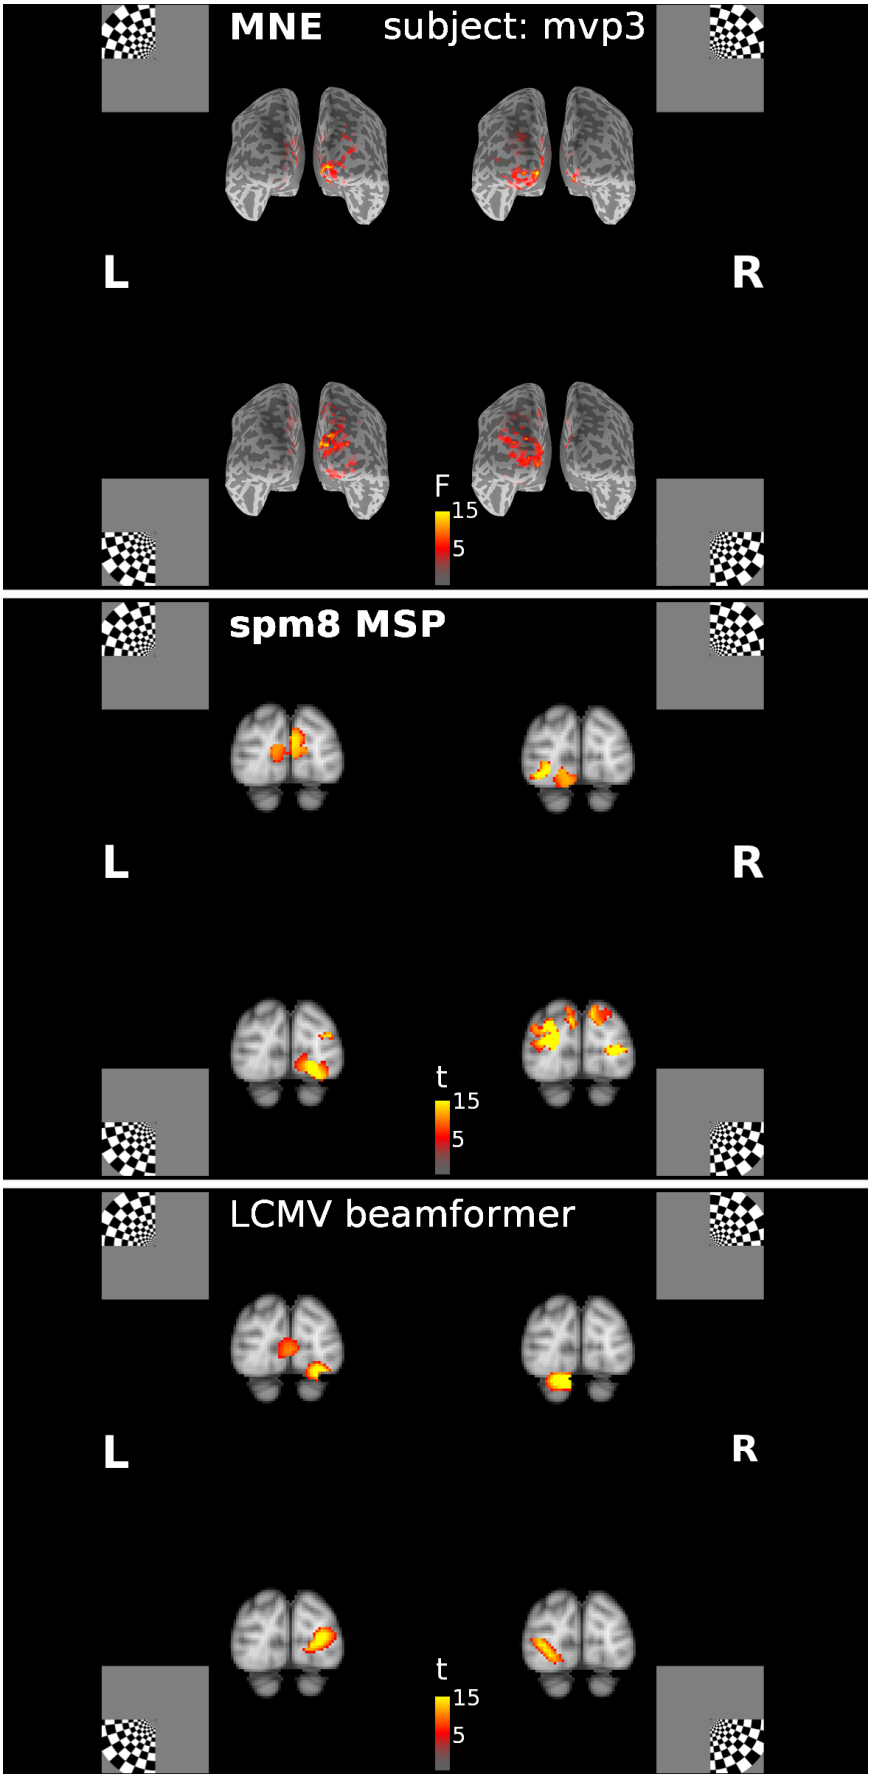

MNE subject: mvp3
L
R
F
15
5
spm8 MSP
L
R
t
15
5
LCMV beamformer
L
R
t
15
5

Figure 6.9

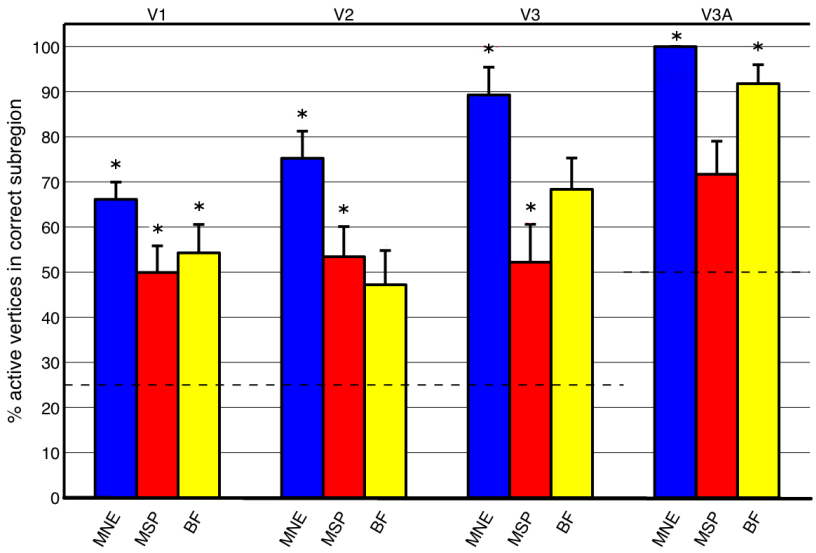


Figure 6.10.: Accuracy of retinotopic source localizations of early evoked responses to quadrants: Minimum Norm Estimate (MNE, *blue*), SPM8's Multiple Sparse Priors (sMSP, *red*), and an LCMV beamformer (BF, *yellow*). Bars show the percentage of active vertices that were in the expected subregion of the early visual area (ROI), averaged across subjects and quadrants. Error bars show S.E.M. Black dashed lines show the chance levels for each early visual area ROI. One-sample t-test of localization accuracy compared to chance level: * indicates $p < 0.004$ (equivalent to significance level $\alpha = 0.05$ under Bonferroni correction for n=12 multiple comparisons).

by-trial to the MEG data and then statistically contrasted the evoked responses to different quadrants (Methods). However, I also investigated other approaches to MSP source reconstruction, which were not as successful. Initially, MSP reconstruction was applied to averaged, time-locked data, as used for MNE source reconstruction, and no statistics were applied. This initial approach performed poorly (Figure 6.12 (*pink bars*) & Table 6.3). Sources were not localized to the retinotopically expected subregion better than chance for any ROI, even though the same MEG signals had produced excellent results with the MNE method (Table 6.1). An alternative approach was tried where MEG data was localized trial-by-trial, and a one-sample t-test (difference from 0) was applied across trials, for each cortical location. Resulting sources were still not localized to the expected subregion better than chance for any ROI (Figure 6.12 (*purple bars*) & Table 6.3). However, an MSP approach that statistically contrasted responses to different stimuli against each other did localize activity to the expected subregion significantly better than chance, for 3 of 4 ROIs (Figure 6.12 (*red bars*), Table 6.3). As the 'best case', this approach was used in the general comparison with the other source methods (above). This exploration of MSP approaches revealed that trial-wise variance improved localization accuracy and that a statistical contrast between responses to different stimuli was necessary to bring localization accuracy above chance levels.

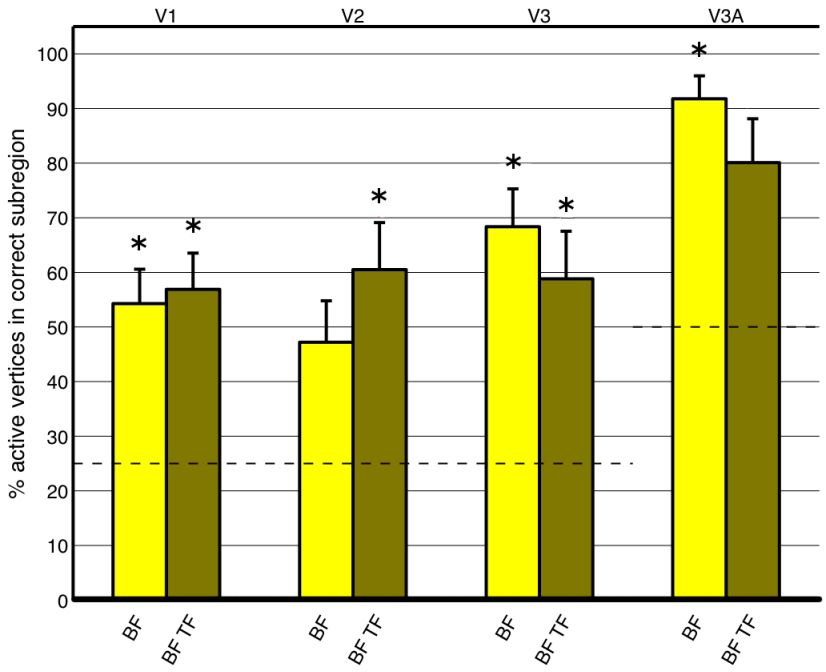


Figure 6.11.: Accuracy of retinotopic source localizations of early evoked responses to quadrants: beamformer for the early evoked response (BF, *yellow*) and time-frequency beamformer (BF TF, *brown*). Bars show the percentage of active vertices that were in the expected subregion of the early visual area (ROI), averaged across subjects and quadrants. Error bars show S.E.M. Black dashed lines show the chance levels for each early visual area ROI. One-sample t-test of localization accuracy compared to chance level: * indicates $p < 0.004$.

SPM8 software contains the option to apply an 'independently and identically distributed' (IID) source algorithm, which is equivalent to classical minimum norm estimates as implemented by MNE (Friston, 2009). This IID algorithm was applied to the averaged MEG data, and all other analysis parameters were identical to those used in the MSP method applied to the averaged data (described above and in the Methods). Over the six subjects, the MNI coordinates of peak responses to the quadrant stimulis were better lateralized to the contralateral hemisphere with the IID method than with MSP and peak responses to lower field stimuli were generally located higher in the brain than were responses to upper field stimuli for IID, but not MSP (Figure 6.13). This suggests a qualitative improvement in MEG localization of angular retinotopic stimuli with an algorithm that assumes a minimum norm estimate (IID), rather than multiple sparse priors (MSP), when all other factors, such as cortical source space and statistics, are identical.

	% of active vertices in the expected subregion of ROI	
ROI (chance level):	BF	BF TF
V1 (25%)	54.3 (±30.1))	56.9 (±32.5)
V2 (25%)	47.2 (±37.1)	60.5 (±44.2)
V3 (25%)	68.3 (±34.0)	58.8 (±42.7)
mean (V1, V2, V3)	*56.5 (±34.8)*	*58.7 (±38.7)*
V3A (50%)	91.8 (±20.6)	80.1 (±39.3)

Table 6.2.: Localization accuracy for the early evoked beamformer (BF, 73 - 93 ms) and time-frequency beamformer (TFBF, 7-9 Hz, 200 - 1000 ms), averaged across subjects and quadrants. The table shows the percentage of active vertices in each early visual area that are localized in the expected subregion according to fMRI retinotopy in the same individual (± S.D.)

	MSP method: % of active vertices in the expected subregion of ROI		
ROI (chance level):	averaged: no contrast	trial-wise: no contrast	trial-wise: t-test contrast
V1 (25%)	31.1 (±27.9)	39.0 (±24.1)	49.9 (±29.0)
V2 (25%)	28.6 (±38.0)	40.5 (±35.1)	53.4 (±32.8)
V3 (25%)	52.0 (±49.5)	47.4 (±49.3)	52.2 (±41.2)
mean (V1, V2, V3)	*35.1 (±38.3)*	*41.7 (±35.4*	*51.9 (±34.3)*
V3A (50%)	60.7 (±44.7)	75.0 (±45.2)	71.7 (±35.9)

Table 6.3.: Localization accuracy for the early evoked responses with the 3 different approaches to MSP source reconstruction, averaged across subjects and quadrants. The table shows the percentage of active vertices in each early visual area that are localized in the expected subregion according to fMRI retinotopy in the same individual (± S.D.)

6.3.5. Effect of cut-off threshold on MEG localization accuracy

For the distributed MEG source reconstruction methods used in this study, it was necessary to designate an activation, or cut-off, threshold which controls whether or not a vertex is considered 'active' in response to the stimulus. (see Methods). Localization accuracy was calculated over a range of activation thresholds, to produce threshold-accuracy functions for each source reconstruction method. These functions would be expected to follow a specific shape. As threshold tends towards 100% active vertices, all vertices in the cortex (and therefore in the ROIs) are 'active'. In this case, the percentage of active vertices located in the retinotopically correct subregion of the ROI would be at chance level (25% for areas V1, V2 and V3). As the threshold tends towards 0% active vertices, there would be fewer and fewer active vertices in any of the subregions of the ROI and at the lower limit, the percentage in the correct subregion of the ROI tends towards 0%. If the source reconstruction method localized MEG sources accurately then the theshold-accuracy function will peak at a level above chance for some threshold between the two extremes.

For all methods, threshold-accuracy functions followed this expected shape (Figure 6.14). The threshold that produced the greatest localization accuracy in the overall grand mean trace was

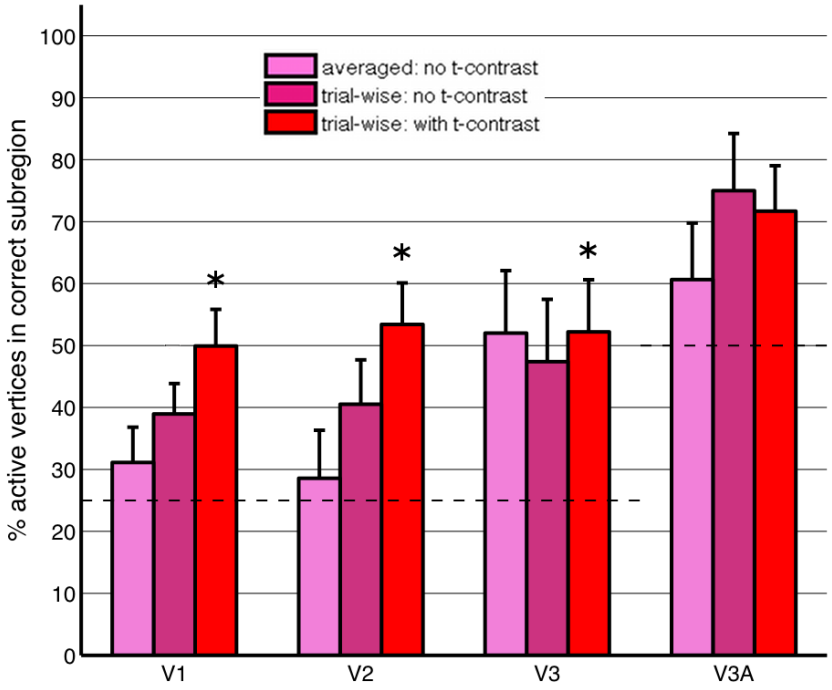


Figure 6.12.: Accuracy of retinotopic source localizations of early evoked responses to quadrants in early visual areas for the 3 different approaches to MSP source reconstruction. Bars show the percentage of active vertices that were in the expected subregion of the ROI, averaged across subjects and quadrants. Error bars show S.E.M. Black dashed lines show the chance levels for each early visual area ROI. One-sample t-test of localization accuracy compared to chance level: * indicates $p < 0.004$.

designed the 'optimal threshold'. On average, the best source reconstruction accuracies were achieved when approximately 2% of cortical vertices were active (Table 6.4).

6.3.6. Summary

For angular stimuli, the minimum norm estimate (MNE) localized the most retinotopically accurate sources when compared to fMRI maps. The other two methods, SPM8's MSP and an LCMV beamformer, localized brain activity to the appropriate early visual area subregion at a level significantly better than chance for 3 of the 4 early visual areas. MNE localized over 75% of active vertices into the correct retinotopic subregion within each ROI, whilst the other methods localized 50-60% of sources to the correct cortical subregions. Using an induced time-frequency power signal for the beamformer did not improve localization accuracy compared to localizing just the broad-band initial evoked response. Trial-wise variance, and a between-stimulus statis-

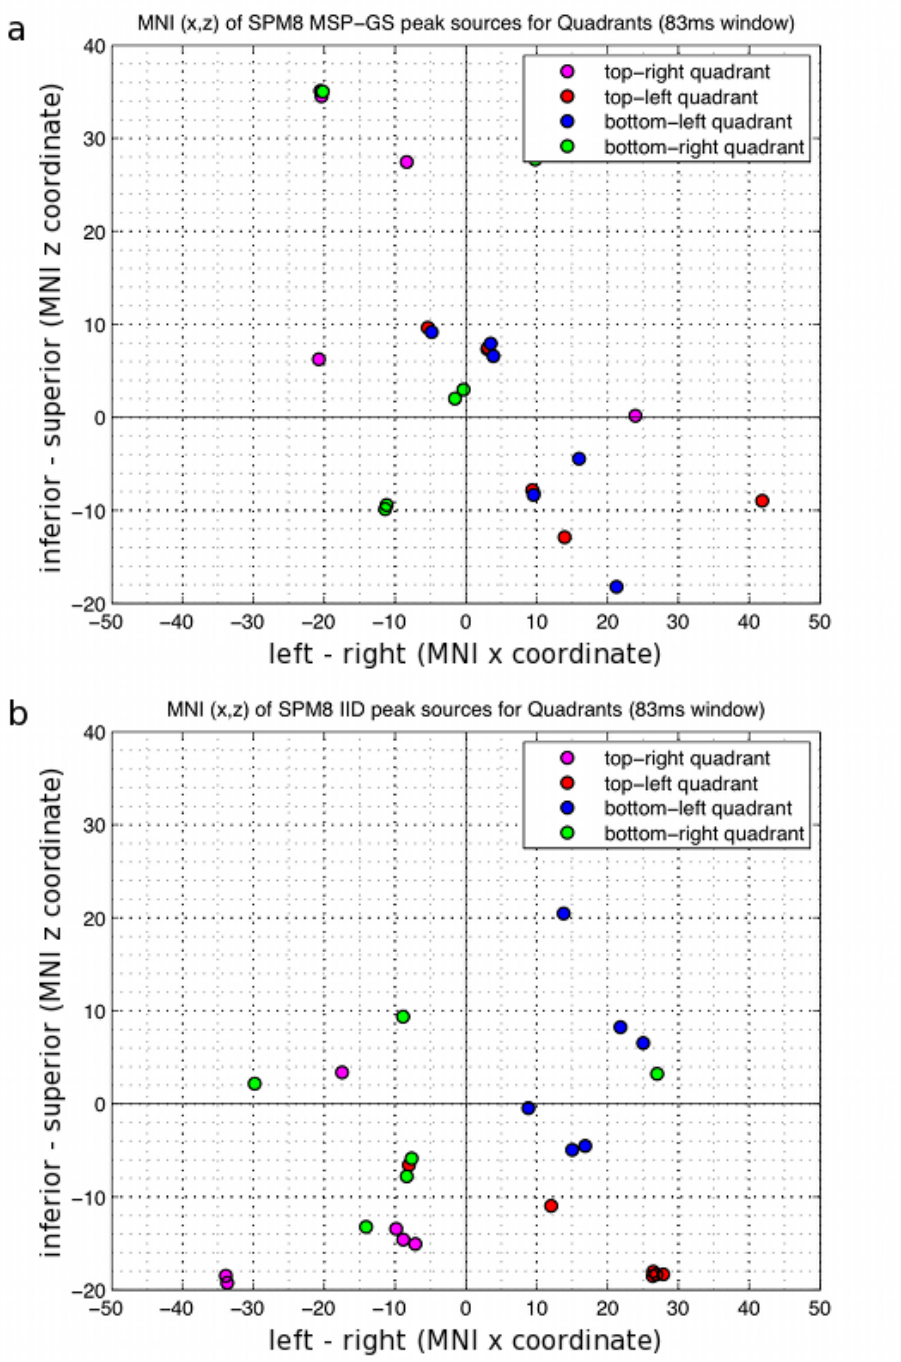


Figure 6.13.: MNI coordinates of the peak response locations to quadrant stimuli for 6 subjects, on the MNI152 standard brain. Right field quadrants (*pink*, green) are expected to localize to the left hemisphere; left field quadrants (*red, blue*) are expected to localize to the right hemisphere. Lower field stimuli (*blue, green* and expected to be located more superiorly (dorsally) in the brain than upper field stimuli (*pink, red*). IID source reconstructions follow this expected retintopic pattern (b), but MSP ones do not (a).

Reconstruction method:	optimal threshold	corresponding statistical value
MNE	2.5%	5.2 (±1.1) *(sqrt(F-val))*
SPM8's MSP	7.9% †	1.5 (±0.91) *(t-val)*
BF	2.0%	2.2 (±0.49) *(t-val)*
BF TF	1.3%	1.7 (±0.47) *(t-val)*
mean	1.9%	

Table 6.4.: The optimal threshold and corresponding statistical value for each of the four MEG source reconstruction methods. † is to indicate that for the MSP approach, the resultant statistical map did not cover all cortical vertices but was smoothed across the best-responding locations, covering an unknown percentage of the total cortex. Hence this threshold value was excluded from the group mean as it was not comparable to the others.

tical contrast, was necessary to raise the MSP method's accuracy above chance. Across methods, on average the best source reconstruction accuracies were achieved when approximately 2% of cortical vertices were active.

Figure 6.14. *(following page)*: Threshold-accuracy functions for three MEG source reconstruction methods for quadrant stimuli: a) MNE, b) sMSP, c) LCMV beamformer. Individual ROI functions are shown for early visual areas V1, V2 and V3 (*black*), and their average (*red*). In each case the accuracy values were averaged across stimulus conditions and subjects. Bars show S.E.M. Threshold-accuracy functions were also evaluated for the time-frequency beamformer (*not shown*).

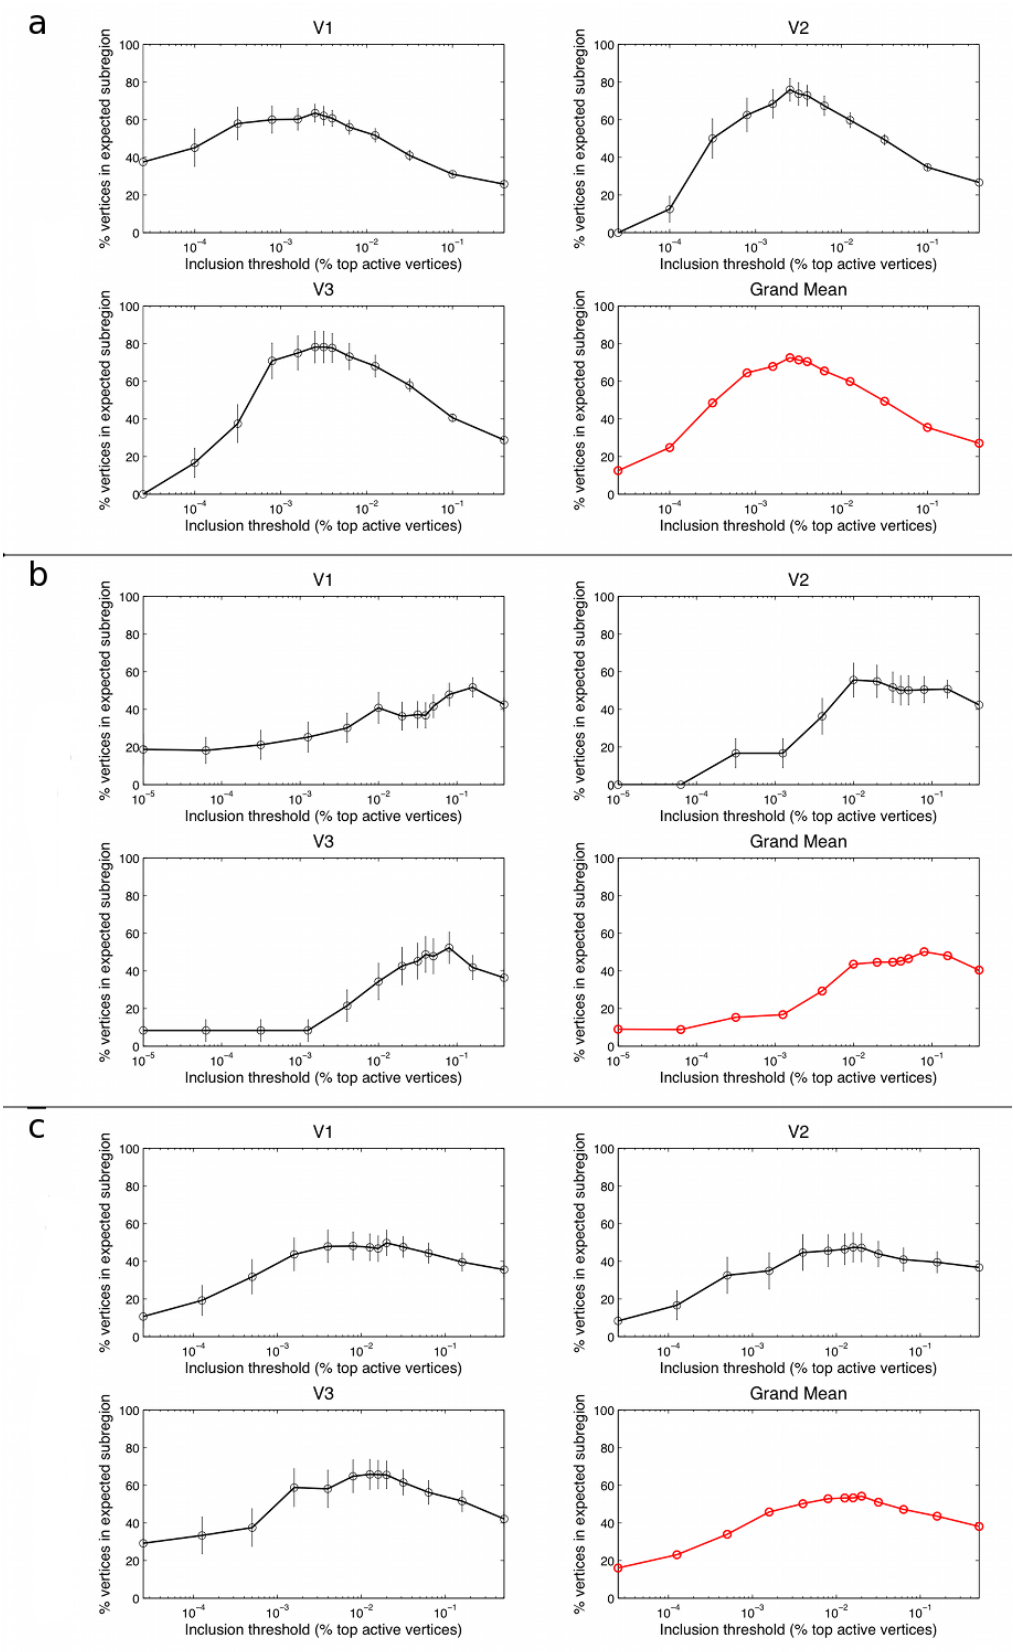

a
V1
V2
V3
Grand Mean
b
V1
V2
V3
Grand Mean
c
V1
V2
V3
Grand Mean
% vertices in expected subregion
Inclusion threshold (% top active vertices)

Fig. 6.14

6.4. Results II: Eccentricity (Rings & mini-rings)

Human brain responses to contrast-reversing concentric ring stimuli were measured with MEG in five of the volunteers who participated in the angular retinotopy study above. Concentric rings are used for mapping retinotopic representations of visual field eccentricity with respect to the fovea (Wandell *et al.*, 2005). Three ring eccentricities were used: 'FOV' ('foveal': 0° to 1.5°); 'PARAF' ('para-foveal': 2° to 4°); and PERI ('peripheral': 6° to 11°). Unlike quadrants, rings are bilateral stimuli, extending over both halves of the visual field, and are therefore expected to activate cortical locations simultaneously in both hemispheres. I compared sources of the MEG signals to the expected locations in early visual areas, based upon individuals' fMRI retinotopy.

6.4.1. Visually evoked responses at sensor level

Sensor space topographies of the early evoked responses to rings did not show a clear spatial pattern according to stimulus eccentricity (Figure 6.15). Moreover, some responses were unilaterally biased (Figure 6.15: *mrt2*, PARAF; *mvp1*, PERI). Therefore, the sensor topography appeared less clear for rings than for quadrants (above).

6.4.2. Evaluation of source localization accuracy

Three MEG source reconstruction methods (MNE, MSP, beamformer) were used to localize the cortical sources of visually evoked responses, this time to the concentric ring stimuli. Subject *mrt2* had the best localization result for the rings: sources localized with the MNE method tended towards the expected posterior-anterior progression in the calcarine region as eccentricity increased (Figure 6.16). However, this result was the exception. Subject *mvp1* is indicative of the pattern of localizations in the remaining subjects: there was no consistent posterior-anterior progression and sources were sometimes not localized in the vicinity of early visual areas at all (Figure 6.17, lower left panel). The beamformer did not result in a recognizable pattern of progression of sources as ring eccentricity increased (Figure 6.18) and similarly poor results were obtained with MSP (*data not shown*). The percentage of active vertices localized to the expected eccentricity band was not significantly better than chance for the FOV and PARAF rings for any source reconstruction method, including MNE which had performed so well in the angular retinotopy (quadrants) analysis (Figure 6.19 & Table 6.5). MNE and MSP methods performed significantly worse than chance level for PARAF rings. However, the beamformer lo-

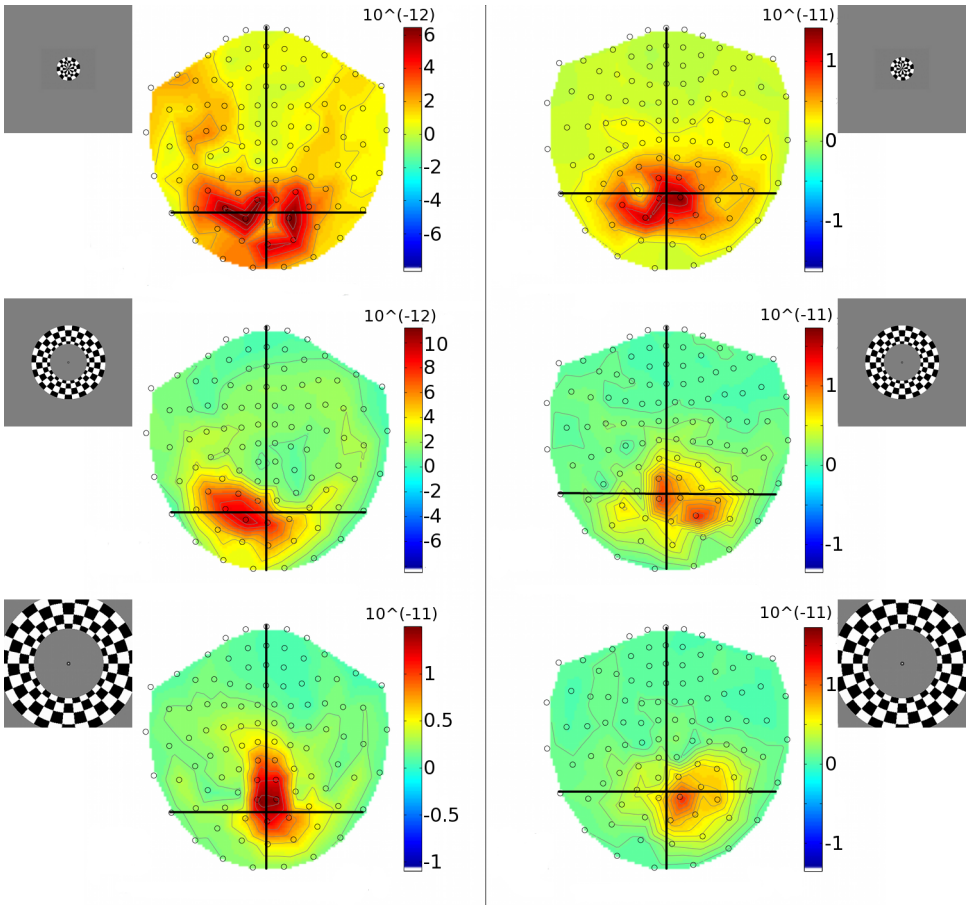


Figure 6.15.: Raw gradiometer sensor topography (T/m) of the visual evoked response at 83 ms for subjects *mcp1* and *mvp3*, for each ring eccentricity. Colour scales are normalized to the peak response to each stimulus. Insets show the ring stimuli.

calized sources preferentially to the expected eccentricity band ROI for PERI ring eccentricities.

I investigated whether using the stimulus frequency tag with the beamformer would improve localization accuracy. Accuracy was very similar between the beamformer for the initial evoked response (BF) and for the time-frequency decomposition (BF TF): both methods failed to accurately localize sources for the FOV and PARAF ring stimuli, whilst localizations for the PERI ring stimuli were better than chance (Figure 6.20 & Table 6.6).

Localization accuracy for ring stimuli was evaluated as a function of cut-off threshold, as it was for quadrant stimuli (described above). The optimal threshold was again defined as the threshold that produced the best localization accuracy overall, for each source reconstruction method. Results confirmed that the poor localization accuracy for FOV and PARAF rings was not due to the specific threshold used, because accuracy remained close to chance level across a wide range of thresholds systematically tested (Figure 6.21 & Table 6.7). The threshold-accuracy functions of PERI rings have an above-chance accuracy peak at the optimal threshold. On aver-

age, the best source reconstruction accuracies were achieved when approximately 4% of cortical vertices were active (Table 6.7). But overall, all source reconstruction methods performed very poorly for eccentricity retinotopy tested with concentric rings.

	% of active vertices in the expected eccentricity band		
band (chance level):	MNE	MSP	beamformer
FOV (33%)	34.2 (±27.5)	28.8 (±38.2)	30.5 (±40.3)
PARAF (33%)	26.9 (±4.5)	10.7 (±14.9)	17.7 (±19.1)
PERI (33%)	64.4 (±29.8)	43.5 (±44.1)	88.0 (±12.2)
mean	*42.6 (±28.4)*	*27.7 (±36.0)*	*46.9 (±40.5)*

Table 6.5.: Localization accuracy for each ring eccentricity, for the source reconstruction methods: MNE, SPM8 MSP and LCMV beamformer. Numbers refer to the percentage of active vertices across visual areas V1, V2 and V3 that are localized in the expected eccentricity band according to fMRI retinotopy (± S.D.)

	% of active vertices in the expected eccentricity band	
ROI (chance level):	BF	BF TF
FOV (33%)	30.5 (±40.3)	25.7 (±39.3)
PARAF (33%)	17.7 (±19.1)	28.1 (±33.9)
PERI (33%)	88.0 (±12.2)	84.5 (±28.8)
mean	*46.9 (±40.5)*	*47.5 (±43.2)*

Table 6.6.: Localization accuracy for each ring eccentricity, for the evoked response beamformer (BF) and time-frequency beamformer (BF TF) approaches. Numbers refer to the percentage of active vertices across visual areas V1, V2 and V3 that are localized in the expected eccentricity band according to fMRI retinotopy (± S.D.)

Reconstruction method:	optimal threshold	corresponding stat value
MNE	6.3%	4.3 (±0.65) *(sqrt(F-val))*
SPM8's MSP	4.0% †	2.2 (±0.57) *(t-val)*
BF	3.2%	1.7 (±0.55) *(t-val)*
BF TF	3.2%	2.7 (±0.05) *(t-val)*
mean	4.2%	

Table 6.7.: The optimal threshold and corresponding statistical value for each of the four MEG source reconstruction methods. † is to indicate that for the MSP approach, the resultant statistical map did not cover all cortical vertices but was smoothed across the best-responding locations, covering an unknown percentage of the total cortex. Hence this threshold value was excluded from the group mean as it was not comparable to the others.

6.4.3. Mini-rings

I re-scanned three subjects with an amended stimulus set ('mini-rings') to investigate the discrepancy between the success of source localization accuracy for MEG visual responses to quadrants and rings. Mini-rings were checkerboard rings confined to the either the upper or lower quadrant of the right visual field, with identical eccentricity positions as the rings above. Since

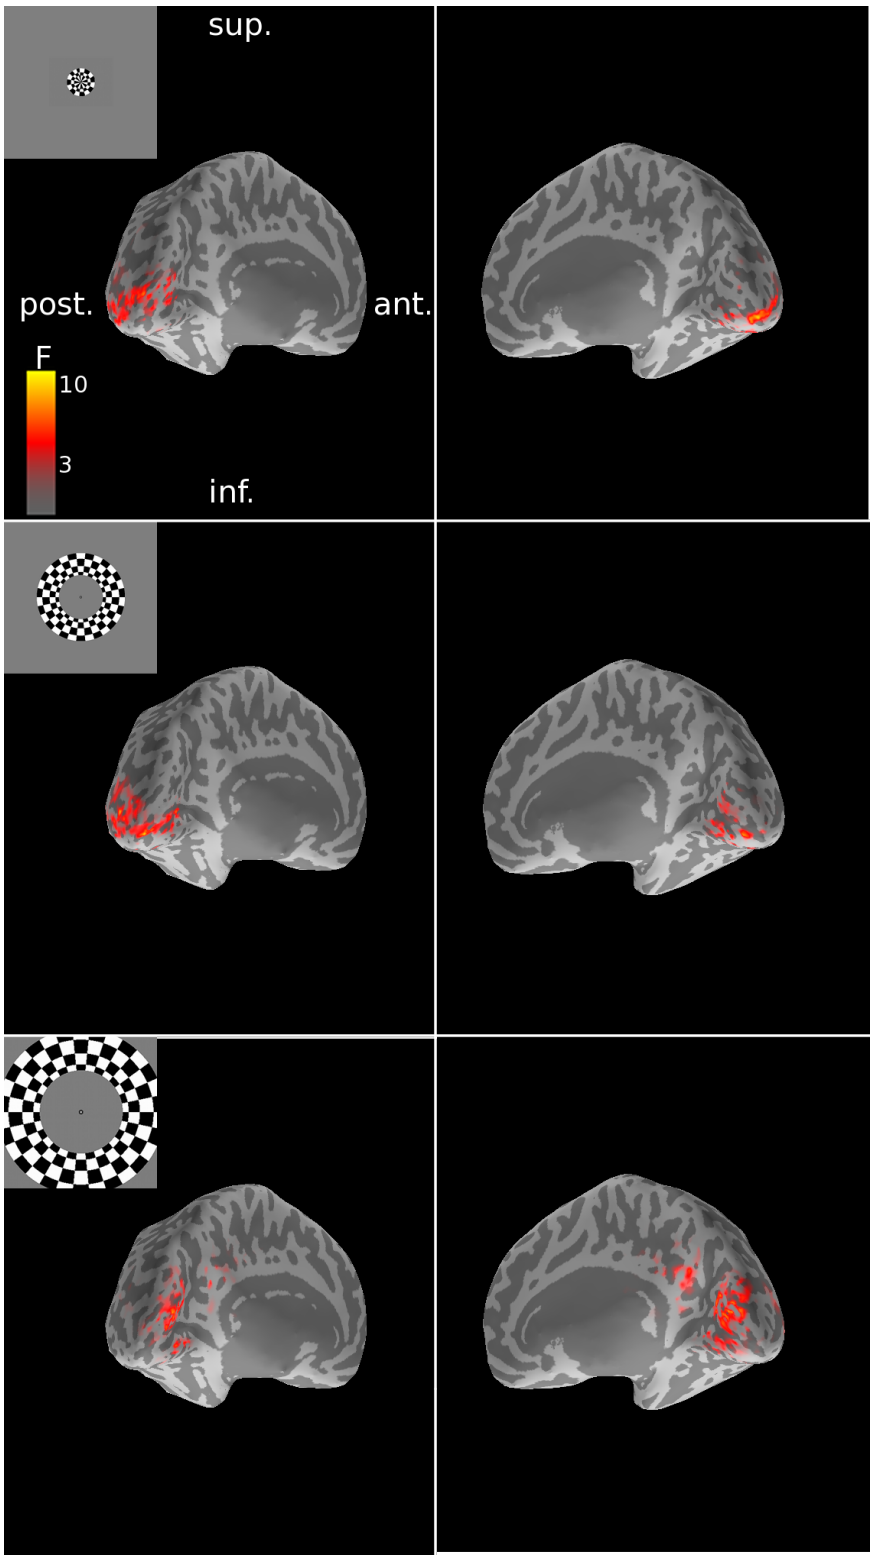


Figure 6.16.: Localizations of initial evoked responses to concentric ring stimuli (73 - 93 ms) with the MNE method, shown for an example subject *mrt2*. Activity is presented on the inflated cortical surface (Freesurfer); left and right hemispheres are on the left and right, respectively. Insets show the ring stimuli.

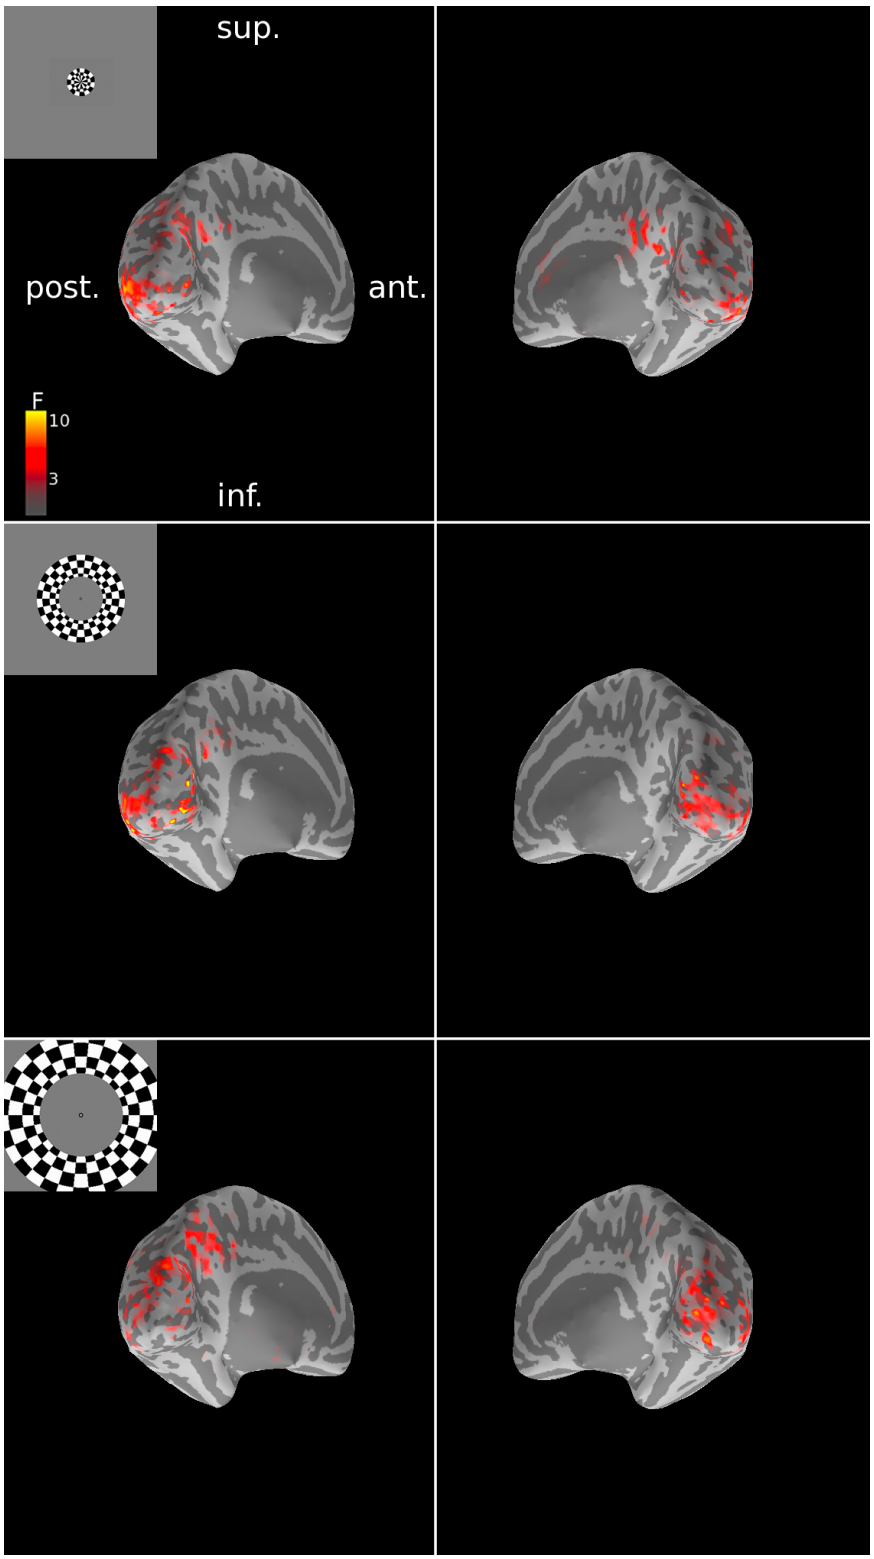


Figure 6.17.: Localizations of initial evoked responses to concentric ring stimuli (73 - 93 ms) with the MNE method, shown for an example subject *mvp1*. Activity is presented on the inflated cortical surface (Freesurfer); left and right hemispheres are on the left and right, respectively. Insets show the ring locations.

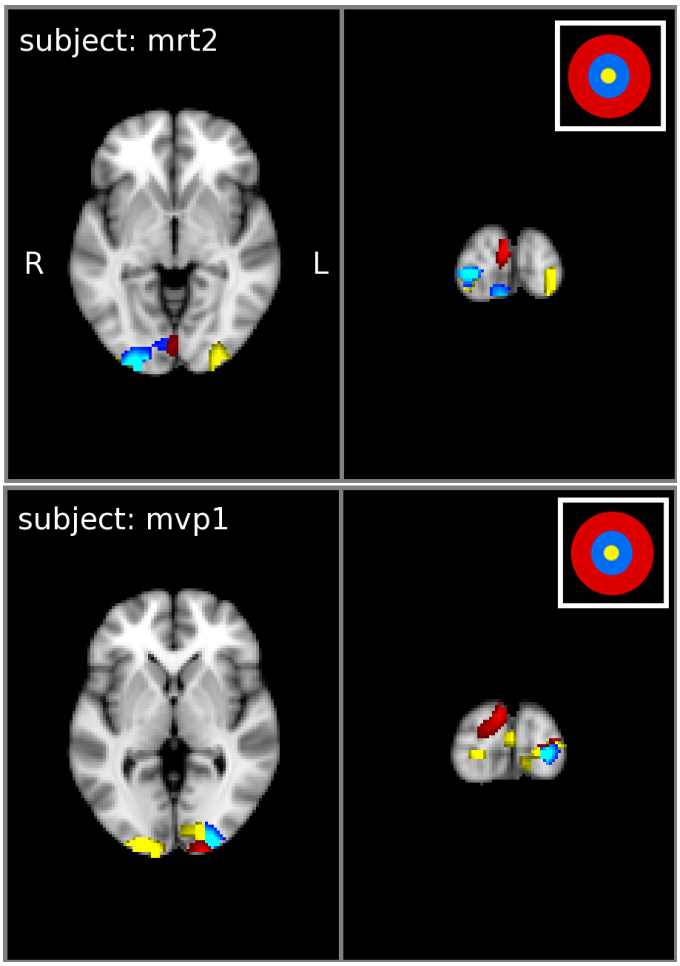


Figure 6.18.: Localizations of initial evoked responses to concentric rings stimuli (73 - 93 ms) with the LCMV beamformer method, shown for two example subjects *mrt2* and *mvp1*. Brain activity response to FOV rings is shown in *yellow*; PARAF in *blue*; PERI in *red* (see *insets* for colour code). Activity range $5 \leq t \leq 15$, displayed on the volumetric MNI152 standard brain.

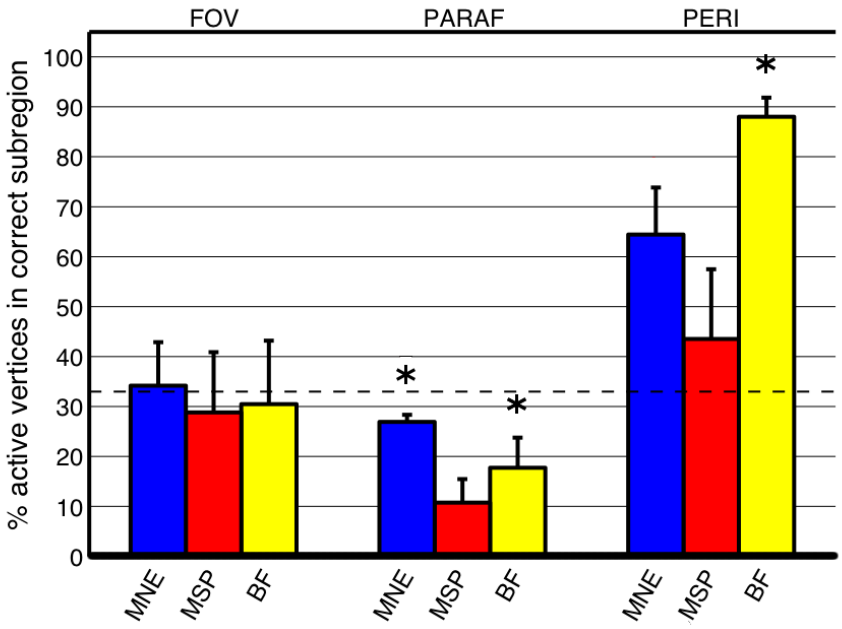


Figure 6.19.: Accuracy of source localizations in early visual area ROIs of visually evoked responses to the rings: minimum norm estimate (MNE, *blue*), SPM8's multiple sparse priors (MSP, *red*), and LCMV beamformer (BF, *yellow*). Bars show the percentage of active vertices located in the correct eccentricity band, across areas V1, V2 and V3, for each ring eccentricity. Error bars show S.E.M. Black dashed lines show chance level. One-sample t-test of localization accuracy compared to chance level: * indicates $p < 0.006$ (equivalent to significance level $\alpha = 0.05$ under Bonferroni correction for n=9 multiple comparisons).

responses to quadrants had been well-localized by the source reconstruction methods, I used the mini-rings stimuli to investigate whether eccentricity retinotopy was localized more accurately when its presentation was confined within a quadrant.

At the sensor level, most evoked responses to mini-rings were located above the contralateral hemisphere, as expected from retinotopy, and lower field stimuli were generally located more dorsally than upper field stimuli (Figure 6.22). However, again there did not appear to be a consistent pattern of source locations as a function of mini-ring eccentricity.

Localization accuracy of the MEG source reconstruction methods was evaluated in a process similar to the quadrant and ring stimulus sets above. Due to the small number of subjects (n=3), values were averaged across the eccentricities (FOV, PARAF and PERI) of lower field stimuli, and separately averaged across the eccentricities of upper field stimuli, to form two statistical groups. A one-sample t-test was used to discern whether the localization accuracy values within each of the two groups (upper field and lower field) were different from chance. Only results for

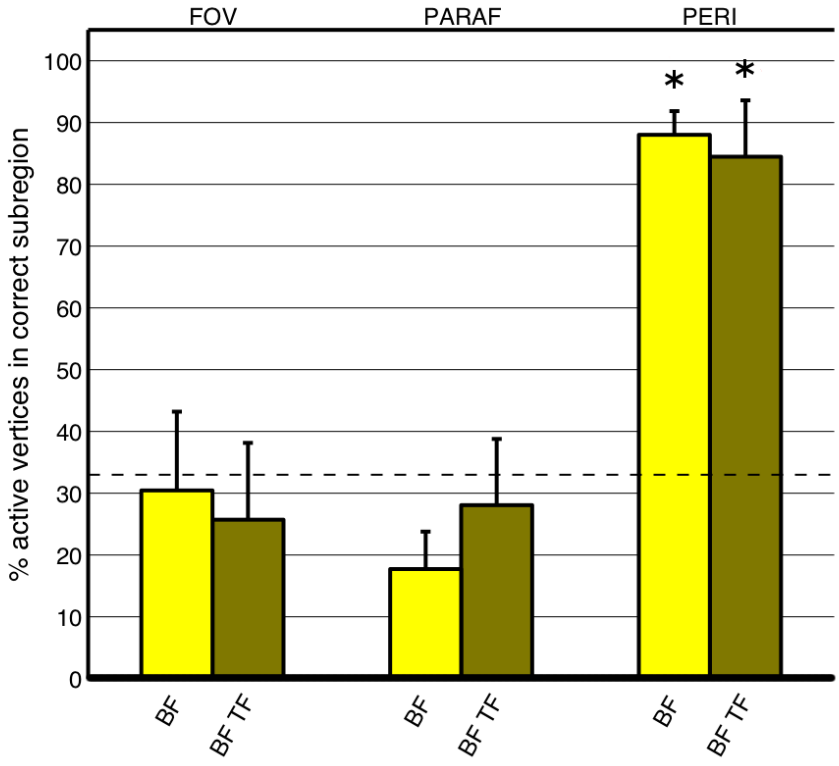


Figure 6.20.: Localization accuracy for rings with the beamformer for the evoked response (73 - 93 ms, BF, *yellow*) and time-frequency beamformer for the stimulus-induced oscillation (200 - 1000 ms, BF TF, *brown*). Bars show percentage of active vertices located in the correct eccentricity band, across areas V1, V2 and V3 for each ring eccentricity. Error bars show S.E.M. Black dashed lines show chance level. One-sample t-test of localization accuracy compared to chance level: * indicates $p < 0.006$.

the MNE and beamformer methods are reported because when the MSP method was evaluated it did not consistently localize responses to the occipital cortex of any subject (*not shown*), and was therefore excluded from further analysis.

Localization accuracy for upper field stimuli was very similar to chance level for both MNE and beamformer analysis methods (MNE: % of active vertices in the expected subregion = 25.0%, p = 0.373; BF: 7.09%, p = 0.0530) whilst the localization accuracy for lower field stimuli showed a trend for improvement from chance for both methods (MNE: % of active vertices in the expected subregion = 50.8%, p = 0.0249; beamformer: 47.1 %, p = 0.053) (Figure 6.23 & Table 6.8).

Figure 6.21. *(following page)*: Threshold-accuracy functions for three MEG source reconstruction methods for ring stimuli: a) MNE, b) sMSP, c) LCMV beamformer. Functions are shown for individual stimuli (FOV, PARAF and PERI; *black*) and the grand average (*red*). In each case, accuracy values were averaged across subjects. Bars show S.E.M.

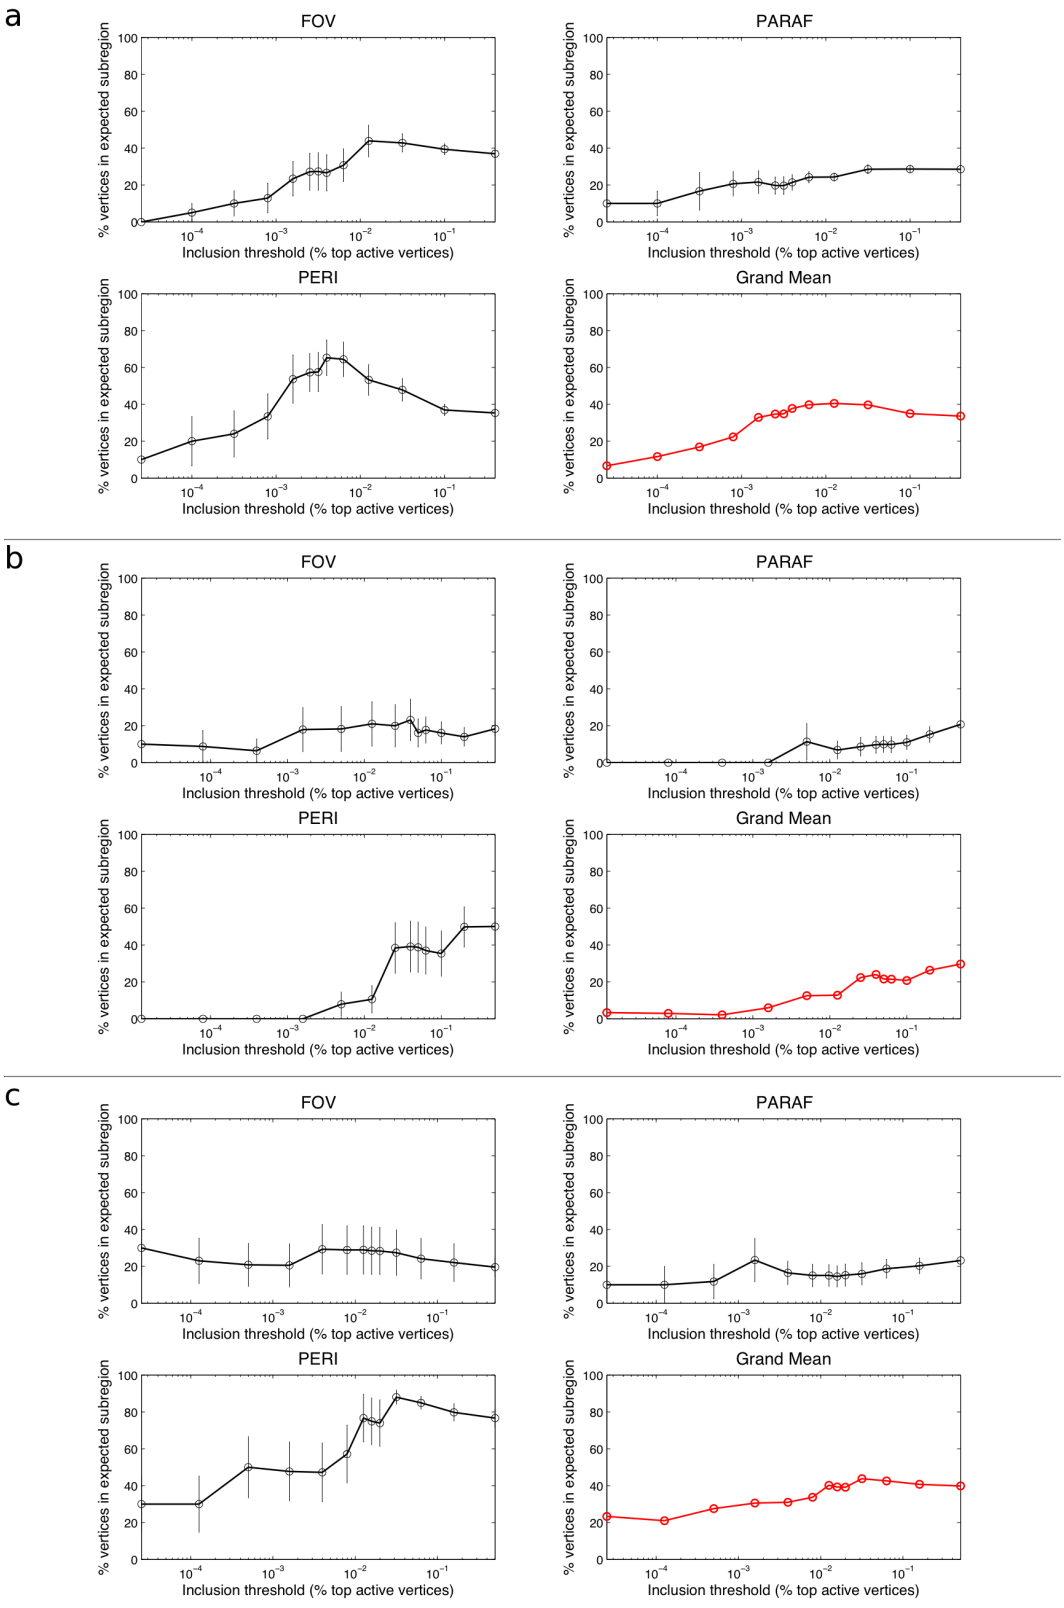

a
FOV
PARAF
PERI
Grand Mean
b
FOV
PARAF
PERI
Grand Mean
c
FOV
PARAF
PERI
Grand Mean
% vertices in expected subregion
Inclusion threshold (% top active vertices)

Figure 6.21

Although the results for lower field stimuli did not reach significance (Bonferroni correction for n=4 t-tests: p = 0.05/4 = 0.0125), the trend does suggest that eccentricity retinotopy may be better discerned by MEG source reconstructions when unilateral lower visual field stimuli are presented.

Source reconstructions are displayed for the MNE method (Figure 6.24) and beamformer (Figure 6.25) for example subject *mvp1* for whom eccentricity localizations with full rings stimulus did not work well (Figures 6.17 & 6.18). With both methods, responses to lower field mini-rings show the expected pattern of progression across the occipital lobe as stimulus eccentricity increases (compare to Figures 6.2 & 6.4). However, the upper field stimuli do not show a consistent pattern with the beamformer. With MNE, a pattern is present but displaced relative to fMRI mapping, as revealed by the quantitative comparison (Figure 6.23). Finally, threshold-accuracy functions for the mini-rings confirm that localization accuracy for upper field stimuli do not rise far above chance (17%) for any activation cut-off thresholds tested (upper panels), whilst localization accuracy for lower field stimuli had a large, above-chance peak at an 'optimal' threshold value (middle panels) (Figure 6.26, Table 6.9).

	% of active vertices in the expected subregion (mini-rings)	
ROI (chance level):	MNE	beamformer
upper field (17%)	25.0 (±25.6)	7.1 (±18.9)
lower field (17%)	50.8 (±33.6)	47.1 (±39.8)
mean	*37.2 (±31.6)*	*28.3 (±37.0)*

Table 6.8.: Source localization accuracy of initial evoked responses to upper field and lower field mini-ring stimuli, for MNE and beamformer methods. Numbers refer to the percentage of active vertices that are localized to the expected subregion according to fMRI retinotopy (± S.D.)

Reconstruction method:	optimal threshold	corresponding stat value
MNE	1.3%	2.9 (±0.42) *(sqrt(F-val))*
BF	0.8%	2.7 (±0.33) *(t-val)*

Table 6.9.: The optimal threshold and corresponding statistical value for each MEG source reconstruction method used to localize responses to mini-rings stimuli.

6.4.4. Summary

MEG localization accuracy for the eccentricity of full-field visual stimuli was poor across all source reconstruction methods, including those that worked well for the quadrants stimulus set. It was again found that using the stimulus frequency tag for the LCMV beamformer resulted

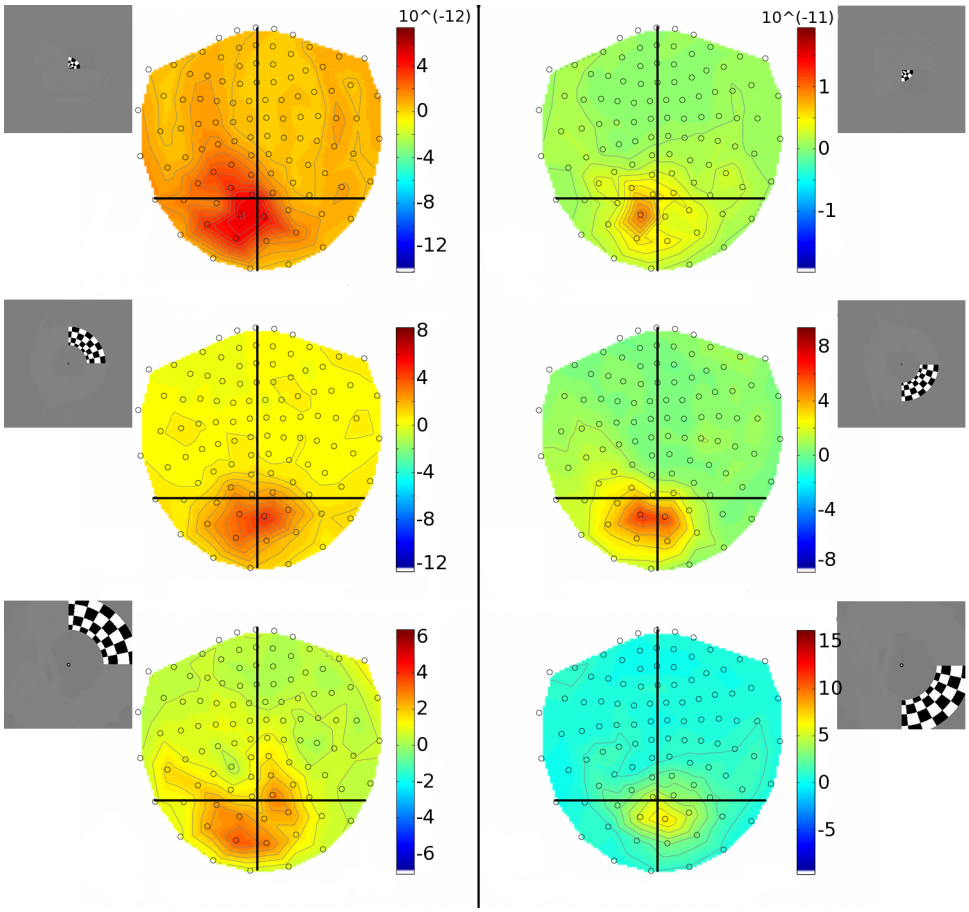


Figure 6.22.: Raw gradiometer sensor topography (T/m) of the visual evoked response to mini-rings at 83 ms for example subject *mvp1*. FOV, PARAF, PERI eccentricities are shown for the upper (*left*) and lower (*right*) visual fields; insets show stimulus locations. Colour scales are normalized to the peak response to each stimulus. This sensor topography pattern was representative of all three subjects. Compare to responses to full-field rings in Figure 6.15.

in similar localization accuracy as achieved with reconstructing the initial evoked response. When the same eccentricity retinotopy locations were evaluated with visual stimuli confined to a quadrant of the visual field, a trend toward improvement in localization accuracy with two source analysis methods (MNE and beamformer) was found for lower field but not upper field stimulus locations. However, more data is needed to confirm this result.

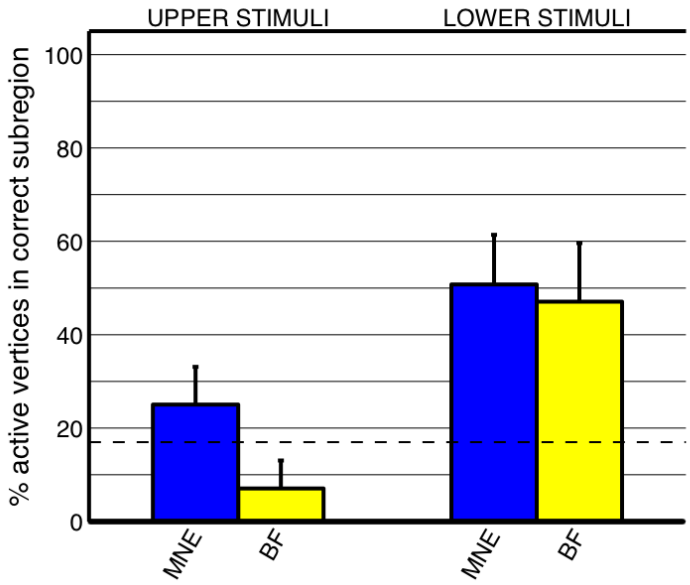


Figure 6.23.: Accuracy of source localizations in early visual area ROIs of visually evoked responses (73 - 93 ms post-stimulus) to the mini-ring stimuli: minimum norm estimate (MNE, *blue*) and LCMV beamformer (BF, *yellow*). Bars show the percentage of active vertices located in the correct subregion, across areas V1, V2 and V3. Error bars show S.E.M. Black dashed line shows chance level. No method reached statistical significance when compared to chance level with one-sample t-test (significance level $\alpha = 0.05$ under Bonferroni correction for n=4 multiple comparisons)

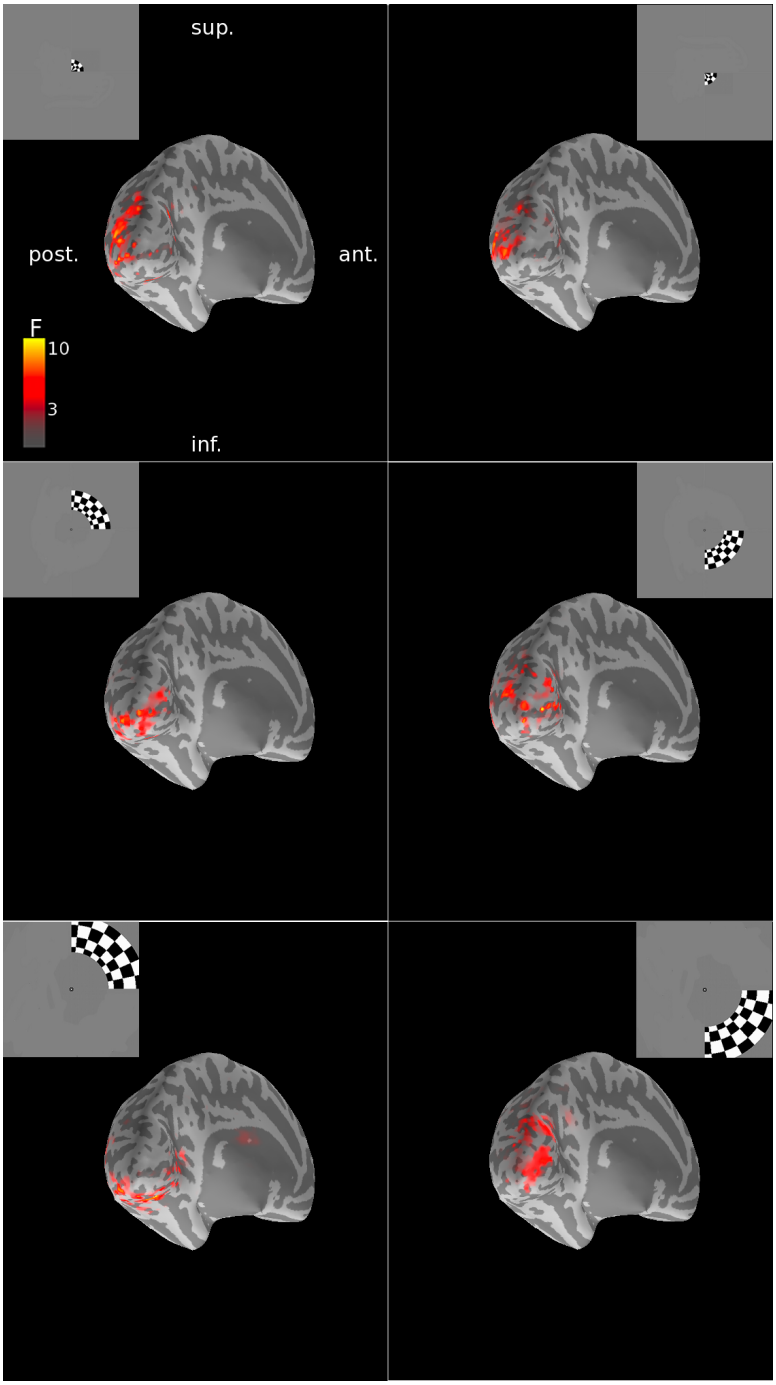


Figure 6.24.: Source reconstructions of initial evoked responses (73 - 93 ms) to mini-rings with the MNE method, shown for example subject *mvp1*. Activity is presented on the inflated surface of the left cortical hemisphere; insets show stimulus locations.

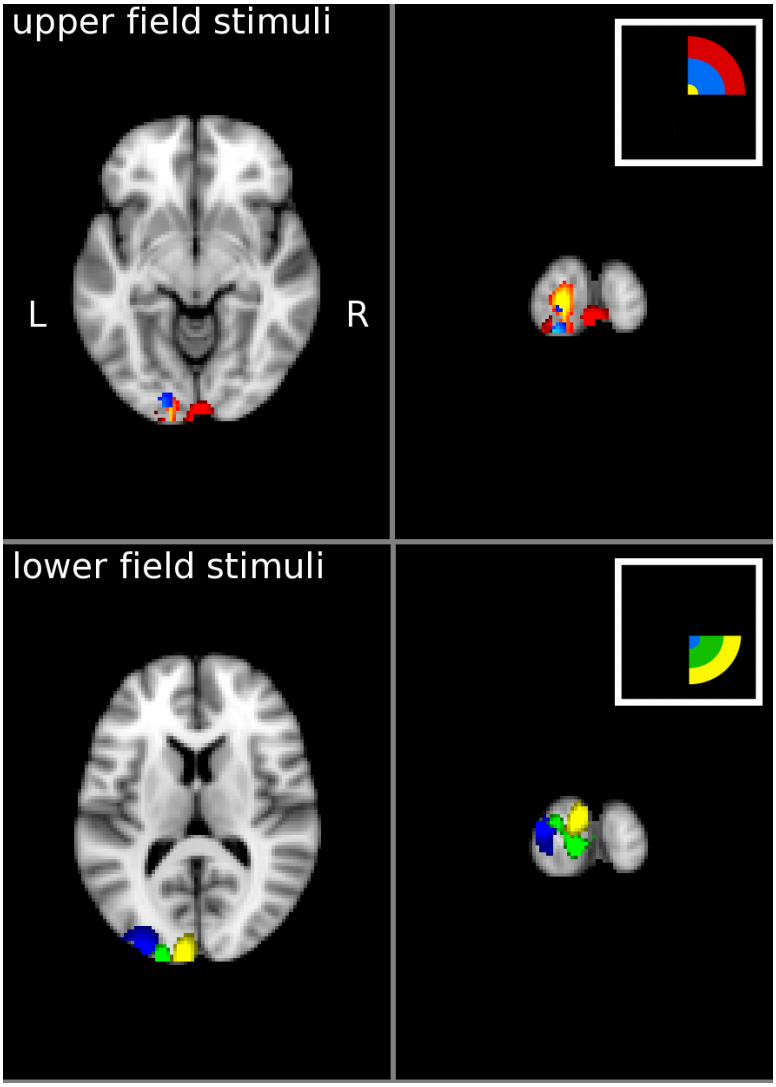


Figure 6.25.: Source reconstructions of initial evoked responses to concentric rings stimuli (73 - 93 ms) with the LCMV beamformer method, for example subject *mvp1*. Brain activity response to upper field stimuli (*top panel*): FOV is shown in *yellow*; PARAF in *blue*; PERI in *red*. Brain activity response to lower field stimuli (*lower panel*): FOV is shown in *dark blue*; PARAF in *green*; PERI in *yellow* (see *insets* for colour code). Activity is displayed on the volumetric MNI152 standard brain. Compare source localizations to lower field stimuli to the fMRI mapping pattern in Figure 6.4.

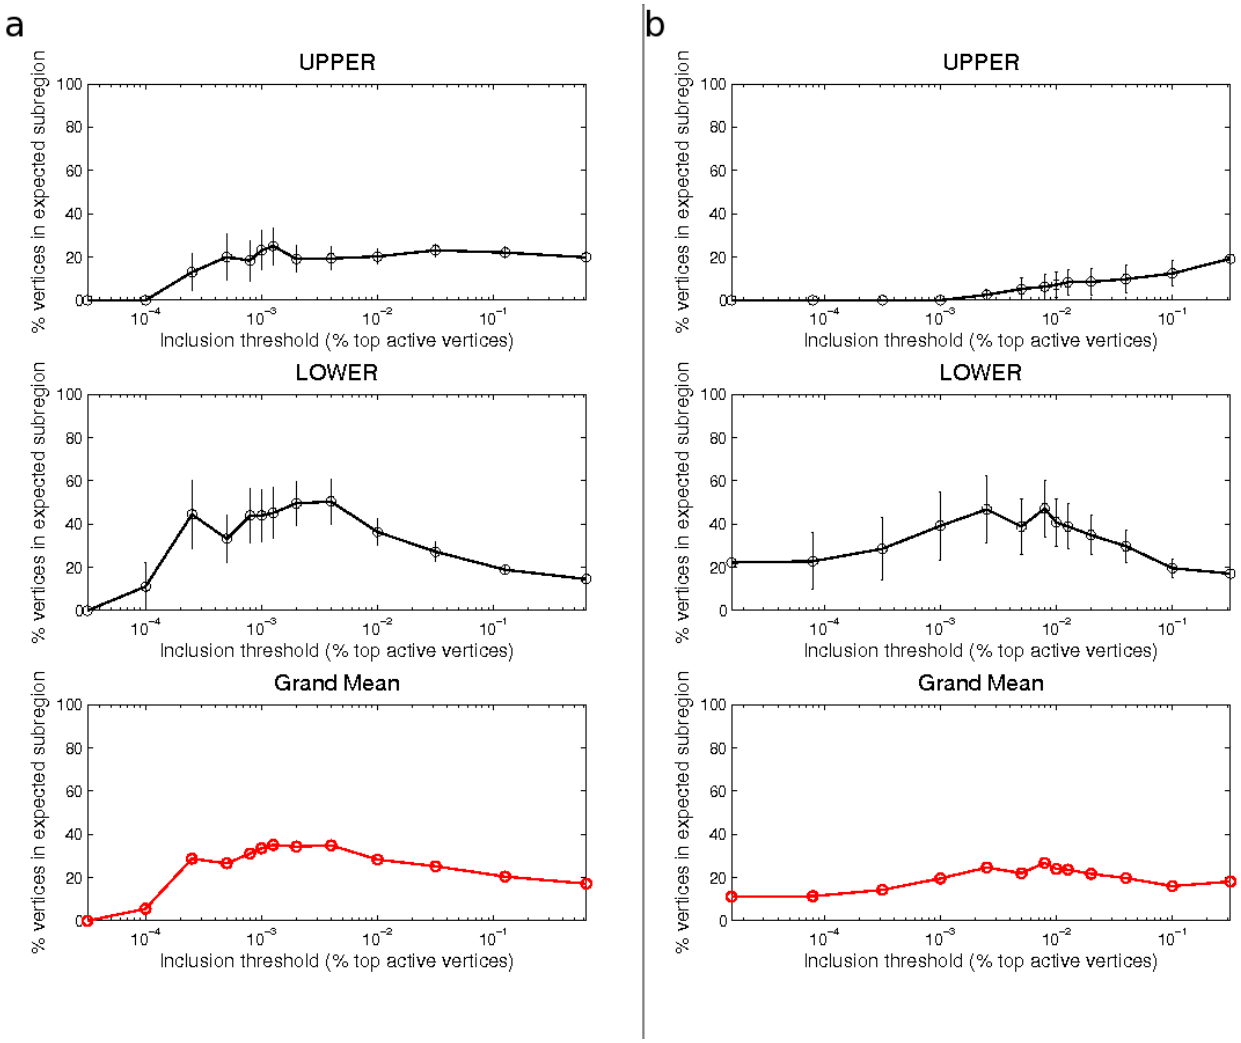


Figure 6.26.: Threshold-accuracy functions for MEG source reconstruction methods used with mini-rings stimuli: a) MNE and b) LCMV beamformer. Individual functions for upper and lower stimulus groups are shown in black, and the grand average of these is shown in red. In each case the accuracy values were averaged across mini-ring eccentricities (FOV, PARAF, PERI).

6.5. Discussion

6.5.1. Source reconstructions for angular retinotopy

Of the three MEG source analysis methods used, the MNE approach clearly resulted in the most accurate localizations of the brain sources. For early visual areas V1, V2 and V3, on average 76.5% of the active vertices localized by the MNE method were located in the retinotopically expected subregion of the ROIs; accuracies of MSP and beamformer methods were much lower. The optimum cut-off threshold for the MNE method, which is a proxy for signal-to-noise ratio, confirmed Sharon *et al.* (2007), who reported that the MNE source reconstruction accuracy is best when dSPM values ≥ 5, i.e. when SNR is high. The MNE approach differs from the other methods mainly in 1) the assumption in MNE that source amplitudes are minimal and cortical sources are many and distributed (instead of fewer sparse patches), and 2) the use of the individual's cortical surface as the source space for reconstructions (instead of an inverse-normalized template brain mesh). Since previous studies have claimed that reconstruction of visually evoked responses is not impaired by the use of the inverse-normalized template mesh rather than the individual's (Mattout *et al.*, 2007; Henson *et al.*, 2009), the main source of improved localization accuracy of MNE should be the underlying algorithmic assumptions of distributed sources of minimal amplitude. When a MNE-like algorithm, IID, was implemented on averaged data in the SPM8 package, there was a qualitatively clear improvement in localization accuracy compared to the MSP algorithm, even though in this case both used the inverse-normalize template mesh. This requires further quantitative testing, but supports the hypothesis that it is the MNE algorithm assumptions which contribute to the superior source localization. An additional factor might have been that MNE reconstruction was performed on averaged, time-locked data; however, when the MSP method was applied to the same averaged data the MSP localization results were even poorer. Interestingly, MNE localized active vertices to the area V3A in response to only 1/3 of stimuli; however, when it did, the localization accuracy was 100% such that all active vertices were within the V3A of the contralateral (retinotopically correct) cortical hemisphere.

The SPM8 MSP and LCMV beamformer methods performed very similarly in terms of angular retinotopic localization accuracy, with accuracy of 50-60%. For MSP, this was surprising as a previous study had shown good source localizations of responses to face stimuli (Henson *et al.*, 2009). The MSP method may have worked well for responses to faces because the expected responding regions (fusiform face areas) are large, bilateral clusters, matching the MSP prior

assumption of sparseness and potentially bilateral activity. However, this signal distribution pattern may be too different from that expected for the angular retinotopy stimuli, which are unilateral and spread irregularly over occipital areas. Furthermore, Henson *et al.* (2009) did not apply a quantitative comparison to fMRI locations to evaluate localization accuracy but only used 'model evidence', a value that optimized reconstruction with respect to the MEG signals alone. Finally, the morphing of the MEG source reconstruction results from the 3D volume image to the individual cortical surface, applied to MSP and beamformer methods but not MNE, may have introduced some loss of spatial resolution and further research is necessary to investigate this.

One previous study suggested that the use of a stimulus frequency tag could produce very high source reconstruction accuracy (Cottereau *et al.*, 2011). In the present study, I did not find an improvement in retinotopic source reconstruction accuracy when I analyzed source power at the second harmonic of the stimulus contrast reversal frequency with the beamformer. Cottereau *et al.* (2011) used a faster stimulus contrast-reversal rate (7.5 Hz) and a wider time window for source reconstruction (5600 ms in length), such that they focused on localizing a 'steady state' visual response. In the current study, stimulus contrast-reversal rate was 4Hz and the time window was 800 ms long, giving fewer oscillations per trial and perhaps resulting in a noisier power estimate which impaired localization accuracy. However, many experimental questions in visual research are not compatible with unnaturally long passive viewing of flickering stimuli ($>$5 s), so the utility of source reconstruction of steady-state stimulus-driven brain oscillations in MEG research generally is not clear.

I applied MSP reconstruction to the averaged MEG data that was used successfully for MNE reconstruction. The accuracy of retinotopic source reconstructions of the averaged data with the MSP method very poor. It was necessary to apply the MSP reconstruction trial-wise, and to apply a between-condition two-sample t-test across trials, to obtain the most accurate localizations. Overall, this suggests that MSP is not an ideal approach for localizing early visual responses and that trial-wise reconstruction and between-stimulus contrasts might have the potential to maximize localization accuracy of MEG signals generally. The LCMV beamformer approach implicitly included these factors to some extent, but an interesting question for future work is whether MNE reconstruction can be further improved with the inclusion of these factors.

6.5.2. Source reconstructions for eccentricity mapping

The MEG retinotopy source analysis was extended to a stimulus set comprising concentric rings, which have been used to map eccentricity in early visual areas of human subjects with fMRI (DeYoe *et al.*, 1994; Engel *et al.*, 1994; Wandell, 1999). None of the reconstruction methods localized responses to rings into the appropriate eccentricity bands in early visual areas at a level better than chance. This was a surprising result, especially for MNE which had localized angular visual stimuli with high accuracy.

There are a number of reasons why full-field eccentricity-varying stimuli may present a unique set of challenges to MEG source reconstruction. Firstly, representation of eccentricity in early visual areas varies along the lateral-medial and anterior-posterior axes of the brain, such that foveal stimuli are represented at the occipital pole and the representation of more peripheral locations progresses deeper into the brain, medially and anteriorly along the calcarine sulcus (Figures 6.2 & 6.4). This may result mainly in changes in amplitudes of activity in the same sensors, rather than different sensors becoming active to different eccentricities (see Figure 6.1). When inspecting sensor topography, indeed there was no consistant pattern across sensor space with stimulus eccentricity (Figure 6.15). In MEG source reconstruction there is inherent ambiguity in discerning low-amplitude superficial activity from higher-amplitude deep activity, which might explain the poor eccentricity results. In particular, MNE is biased towards minimum current amplitudes and hence always favours low amplitude superficial activity results over deeper ones. The MNE software incorporates a depth-weighting calculation to compensate for this (Fuchs et al., 1999; Lin et al., 2006b), but this correction may not be sensitive enough to discriminate retinotopic eccentricity progression. Secondly, concentric rings as full-field visual stimuli are expected to synchronously activate upper and lower calcarine regions in both the left and the right cortical hemispheres, which may result in some cancelation of equal and opposite magnetic fields arising from opposing cortical surfaces. Moreover, spatially extended, correlated source activity cannot easily be spatially filtered for by the beamformer algorithm and would result in localization problems (Hansen *et al.*, 2010). MSP should be best for localizing ring stimuli as it incorporates priors of bilaterality and does not impose a minimum current estimate. However, this method was not better than the others, suggesting that its spatial priors assumptions were still not appropriate for these retinotopic source patterns.

Eccentricity retinotopy was repeated with three participants, using a 'mini-rings' stimulus set that had the same eccentricity locations as the concentric rings but were confined to either the

upper or lower right visual field. Qualitatively, source activity patterns showed the expected retinotopic progression through the early visual areas for lower field stimuli with both MNE and beamformer, but not for upper field stimuli which remained poorly localized or displaced with respect to fMRI. More participants would be needed to confirm this statistically, but an improvement with mini-rings generally could be due to reduced cancellation of signals from opposing cortical surfaces since only one bank of the calcarine sulcus is expected to be active at a time for quarter-field stimulation. Moreover, lower field stimuli are represented in dorsal visual cortical subregions in the occipital brain, which are closer to and sampled by more sensors in the MEG helmet than are their ventral counterparts (Figure 6.1) and therefore have better signal-to-noise ratio (Perry *et al.*, 2011), which as previously mentioned is linked to better localization accuracy (Sharon *et al.*, 2007). The MSP method did not reliably localize the mini-rings stimulus responses to the occipital cortex at all (*data not shown*). The may be simply because MSP is not sensitive enough to localize smaller signals.

6.5.3. Relationship of MEG to fMRI signal

An important assumption of this study was that fMRI retinotopy correctly localizes the true sources of brain responses in individual subjects. Although there is a wealth of histological and lesion evidence to suggest that retinotopic mapping measured by fMRI corresponds to the true patterns (Holmes, 1918, 1944; Horton & Hoyt, 1991; Bridge *et al.*, 2005; Bridge & Clare, 2006), there might be unknown differences between the exact locations of the sources of brain activity measured by MEG and fMRI. These methods detect different underlying processes (electrophysiological versus metabolic) and the time-scale on which these processes change is different (milliseconds versus seconds). MEG signals most likely arise from synchronous synaptic current in cortical pyramidal cells (Chapter 5). There is ongoing research into the electrophysiological correlations of the fMRI BOLD signal but it seems tobe linked to the local field potential (LFP), which is also a measure of total of synaptic activity in cortical cells (Logothetis *et al.*, 2001). Early visual cortex, especially striate cortex, is well-perfused, such that blood vessels are likely to be spatially close to their neuronal sources (Engel *et al.*, 1997). These arguments indicate that an fMRI-MEG comparison is appropriate.

However, it would be useful to confirm of MEG signals using a more direct method. In human neurological patients or animal models, one could invasively measure the brain sources of visual responses and record their locations on a 3D anatomical MRI using stereotactic computational methods. Source reconstructions of separate MEG recordings to the same stimuli could

be coregistered with the anatomical MRI and compared to these locations.

6.5.4. Conclusions and future directions

The MNE method, which assumes many distributed sources of small amplitude, seems best-matched to localize the irregularly patterned, unilateral responses of retinotopic subregions of early visual areas. On the other hand, the sparse priors of the MSP method may be better matched for large, cluster-like source distributions that are bilateral, such as responses to face stimuli in the fusiform gyrus (Henson *et al.*, 2009). The spatial filter approach of the beamformer may lie somewhere in between, having greater flexibility albeit less specialization for any particular response pattern. Experimenters should match the analysis method to the distribution of the expected signal and if necessary use more than one method to analyze the same dataset, reporting the results of both. These findings also suggest that rather than using full-field visual stimuli, stimuli confined to a quadrant will improve signal to noise ratio and source localization with the MNE and beamformer methods.

We cannot exclude the possibility that the different anatomical source spaces incorporated in the different MEG analysis approaches contributed to the different localization accuracies. Future studies should directly compare MSP and beamformer localization using the individual cortical surface, in native source space, to the same using the inverse-normalized template cortical surface in MNI152 template brain space. It would be necessary for the SPM8 and OSL analysis software to be extended to enable this comparison. In spite of previous studies (Mattout *et al.*, 2007; Henson *et al.*, 2009), it could be that information about individuals' cortical folding geometry would aid localization accuracy for responses in the tightly folded early visual areas (Figure 6.4).

Localization accuracy of the different MEG analysis methods was evaluated by calculating what percentage of the active vertices in the ROI were located in the retinotopically expected subregion of the ROI. This method purposely ignores the incidence of active vertices outside the ROIs in question. This is because many other visual regions (LO, V4, V3B, hMT+, etc) might respond to these stimuli in the time windows used, but testing all these regions was outside the scope of this study. An alternative way of testing accuracy could be to evaluate how many of the cortical vertices in a particular retinotopic subregion are 'active' in response to the corresponding stimulus. However, this value would be some currently unknown low number even with perfect MEG source reconstruction, because MEG sensors are blind to sources located at certain parts of the

cortex, such as the crests of gyri, due to the geometry of magnetic fields of the brain relative to the orientation of the sensor array (Figure 5.1, Chapter 5). Future attempts at an anatomically-corrected analysis of this type would be highly interesting. Although the optimal threshold for source reconstruction accuracy was systematically evaluated and found to be 2% across the methods available here, due to the sensitivity of MEG to brain geometry it cannot be concluded that this is the actual percentage of the cortical surface responding to the stimuli. The true value is likely to be greater as it includes gyral activations invisible to MEG (Figure 5.1). Finally, in order to extend the conclusions between source pattern distribution and different MEG source analysis algorithms more generally across the brain, future work should investigate localization accuracy in other brain regions that are also known to be organized topographically, such as the body map in primary somatosensory cortex (Penfield & Boldrey, 1937) or tonotopic maps in primary auditory cortex (Langers & van Dijk, 2011).

7. Human perceptual decision-making under reward: a MEG study

7.1. Introduction

Chapters 2, 3 & 4 of this thesis provided behavioural evidence that reward can influence the processing of visual information during perceptual decisions, prior to the evidence-accumulation stage in sensorimotor regions. Here, I build upon these findings to investigate whether effects of reward on sensory processing during perceptual decisions occur in the human brain, measured with magnetoencephalography (MEG). I have already shown that MEG responses to retinotopic stimuli can be accurately localized to the corresponding subregions of early visual areas (Chapter 6). In Chapters 2 & 3, the effect of reward was related to the overall size of reward available to the animals for a correct choice in either response direction. However, in the human psychophysics study of Chapter 4, reward significantly biased behavioural choices and reaction times when unequal payoffs were available for the two possible choices, but not when overall reward sizes were equal for both directions. Therefore, in the present human MEG study, I focus on localizing differences in MEG brain activity when perceptual judgements are biased by unequal payoffs.

As reviewed in Chapter 1, a few fMRI studies have investigated effects of reward on visual processing. However, these did not test how modulations of cortical BOLD signal due to reward might affect perceptual decisions themselves, and did not have the temporal resolution to confirm at which stage of the decision process these effects occur (Serences, 2008; Serences & Saproo, 2010). Previous MEG studies have investigated perceptual decision-making and reward processing separately. Donner *et al.* (2009) measured MEG sensors above motor cortical regions whilst humans performed the motion task, and reported activity that correlated with the integration of gamma power arising near visual area hMT+ (V5). Wenzlaff *et al.* (2011) investigated the brain sources of the speed-accuracy trade-off for visual perceptual decisions with MEG and

found that activity in supplementary motor area (SMA) and left dorsolateral prefrontal cortex (DLPFC) correlated with changes in drift diffusion model parameter estimates under different strategies. This suggests that SMA and other motor regions, and the DLPFC, are locations involved in perceptual decision-making which may potentially also contain reward signals. Some studies have used MEG to investigate brain activity in value-based decision-making tasks, but not in perceptual tasks (Hunt *et al.*, 2012; Steffen *et al.*, 2011).

A recent EEG study investigated the effects of motivational (reward) and attentional cues on brain activity whilst humans performed a Posner task, and found that the peak latency of the N1 response, a putative perceptual signal occurring 150 - 200 ms post-stimulus, was modulated by whether or not the outcome of the trial contribute to overall remuneration for task performance (Baines *et al.*, 2011). I aim to extend this finding by investigating whether *unequal* reward sizes for a correct choice affect visual processing, using a MEG source localization approach.

Finally, Jiang *et al.* (2008) investigated MEG-recorded brain responses to structure-from-motion (SFM) stimuli, similar to those I present in the cylinder task. They report a progression of activity through primary visual cortex, visual area hMT+, parietal cortex, lateral-occipital (LO) regions and temporal cortex within 500 ms of stimulus onset (Jiang et al., 2008). I expected to find a similar pattern of activity in MEG datasets recorded here.

7.2. Materials & Methods

MEG and structural MRI analysis processes are described in *Chapter 5: MEG Materials & Methods*; only additional study-specific details are described in this section.

7.2.1. Data Collection

7.2.1.1. Participants

Seven volunteers (5 female, 2 male; 6 right-handed, 1 left-handed; mean age 25 $\pm$ 5 years, range 20 - 33) took part. Six volunteers (*mrw1-6*) had already participated in the psychophysics study (Chapter 4). The seventh volunteer (*mrw7*) completed 3 training blocks (no feedback) and 8 experimental blocks (with feedback and reward) of the psychophysics version of the task prior to MEG scanning. All were naive to the goals of the experiment, except *mrw2* who was the author.

7.2.1.2. Cylinder Discrimination Task

Participants discriminated the direction of rotation of a SFM cylinder (Figure 2.1). Points were available for correct choices and were converted to money at the end of the experiment. The reaction time task stimulus program was adapted from that of Chapter 4 (see Figure 4.1). In the MEG scanner participants viewed the binocular disparities of the cylinder stimuli through red-green glasses, as stimuli had to be projected onto one screen. Screen size was 285 x 355 mm (1280 x 1024 pixels), viewing distance 1500 mm. Cylinders were presented centrally (foveally), 2.35° in height and diameter, and made of 50 light and 50 dark dots. Screen luminance was 9.6 cd/m^2 and Michelson contrast between light and dark dots was 89%. All other parameters were identical to Chapter 4.

Each trial began with a central fixation point presented for 500 ms, followed by a central points cue which indicated reward condition. The cue always consisted of a circle the same size as the fixation point, bisected into a white and a black half (Figure 4.1). Vertical bisection indicated different point values (1 point or 8 points) available for each response direction (*unbalanced* conditions: high/low (HL) and low/high (LH)). Horizontal bisection indicated equal points (5 points) available for each response direction (*balanced* reward condition, medium/medium (MM)). The HH (8 points in each direction) and LL (1 point in each direction) reward conditions, used in Chapter 4, could not both be included in this MEG due to constraints on scanning time. The MM condition (5 points in each direction) was chosen instead of either HH or LL alone in order to make the combined total of points available in the trial similar to that in the HL or LH conditions. Whether white or black indicated the higher-rewarded direction was counterbalanced across subjects.

The points cue had a variable duration (uniform distribution range 1000 - 1400 ms) and was followed by cylinder stimulus onset. The stimulus remained until the participant gave a response by briefly lifting their thumb from an optic response button, one held in each hand. Left button indicated a clockwise (CW) rotation choice and right button indicated a counter-clockwise (CCW) rotation choice. The points cue remained on the screen throughout stimulus presentation, to minimize unnecessary visual changes at stimulus onset. 500 ms after stimulus offset, visual feedback, in the form of a number presented at fixation, indicated points gained: 0 = incorrect response; 1, 5 or 8 = correct response, depending upon reward condition (Table 7.1). Feedback duration was 700ms, followed by an inter-trial interval (ITI) of 1000 ms. Every 20 trials, total points earned so far in the block were presented for 4000 ms.

Table 7.1.: Points assignments in reward conditions:

reward condition	Points for correct CW choice (left button)	Points for correct CCW choice (right button)
high/low (HL)	8	1
low/high (LH)	1	8
medium/medium (MM)	5	5

On the basis of each individual's binocular disparity threshold, ascertained in prior psychophsyics testing (Chapter 4), five cylinder disparity (stimulus strength) values were presented:

- 2 x supra-threshold cylinders (CW & CCW), easily discriminated by the participant
- 2 x near-threshold cylinder (CW & CCW), close to the participant's disparity threshold
- 1 x zero-disparity cylinder (perceptually bistable)

Supra-threshold disparity was always 4x greater than near-threshold disparity; CW rotation corresponded to positive disparity. To minimize the time spent in the MEG scanner, the proportion of trials at each cylinder disparity/reward combination differed, chosen to allow the main experimental hypotheses of interest to be tested efficiently (Table 7.2). Eight blocks were presented, of 100 trials each, except the first block which had 110 trials (total 810 trials per subject). There was a 5 minute break between each block. Participants were instructed to make judgements as speedily and accurately as possible and to use the points cue to gain as many points as possible (example instruction sheet in Appendix B).

Table 7.2.: Number of trials at each cylinder disparity/reward condition combination.

cylinder disparity	MM condition	HL condition	LH condition
supra-threshold CW	60	0	0
near-threshold CW	60	75	75
zero (bistable)	120	75	75
near-threshold CCW	60	75	75
supra-threhsold CCW	60	0	0

7.2.1.3. Eye movement tracking

In addition to vertical EOG measurements to record eye-blinks, the right eye movements were recorded with eye-tracker (details in Chapter 5). Due to technical limitations, eye-tracking was only possible for 4 of 7 subjects (*mrw1*, *mrw2*, *mrw5*, *mrw6*). For 3 subjects (*mrw5*, *mrw6*, *mrw7*), a horizontal EOG measurement was additionally used as a back-up monitor of eye movements.

7.2.2. Data Analysis

7.2.2.1. Preprocessing: outlier reaction times

Methods used to calculate maximum and minimum reaction time (RT) cut-offs, to exclude outlier RT trials from analysis, were as described in Chapter 4 (Methods). Across subjects, mean minimum RT cut-off was 375 ms ($\pm$ 8.7) and mean maximum RT cut-off was 2526 ms ($\pm$ 1071.2). These values, obtained in the MEG scanner, were similar to those obtained in the equivalent psychophysics set-up for which the corresponding values were 391 ms ($\pm$ 41) and 2685 ms ($\pm$ 1267) (Chapter 4).

Table 7.3.: Outlier RT cut-offs and the number of trials rejected for each participant.

	minimum RT		maximum RT		total removed	
subjects:	**cut-off (ms)**	**n**	**cut-off (ms)**	**n**	**total n**	**%**
mrw1	529	2	4409	2	4	0.49
mrw2	301	50	1131	31	81	10.00
mrw3	330	5	2584	13	18	2.22
mrw4	308	48	2264	19	67	8.27
mrw5	329	15	1836	15	30	3.71
mrw6	462	46	2111	17	63	7.78
mrw7	364	5	3346	5	10	1.23
mean	*375*	*24.4*	*2526*	*13.1*	*39*	*4.8*
SD	*8.7*	-	*1071.2*	-	-	-

7.2.2.2. Behavioural choices and effect of reward condition

The proportion of CW choices and median RT of correct responses to each cylinder disparity were calculated. For the balanced reward condition (MM), a psychometric function (PMF) was fitted to the choice response data using a cumulative Guassian distribution (details in Chapter 4). It was not possible to fit a PMF to the unbalanced conditions (HL & LH) because only three cylinder disparity values were used (Table 7.2). Choice and reaction time data were averaged across subjects. As the data were not normally distribution ($p < 0.05$, Lilliefors test), a non-parametric Kruskal-Wallis test was used to evaluate whether choices were significantly affected by reward condition.

7.2.2.3. MEG preprocessing

Data preprocessing was as described in Chapter 5. Continuous MEG data was epoched according to reward cue onset, visual stimulus onset or button press.

7.2.2.4. MEG source reconstruction

Individual structural information from magnetic resonance imaging (MRI) was available for three subjects (*mrw1*, *mrw2* & *mrw6*). The anatomical MRI for subject *mrw1* was used as a proxy for remaining subjects. Over 100 headshape digitizer points were collected for each subject prior to MEG scanning, to optimize the co-registration between the subject's headshape in the MEG scanner and the MRI source space. Source space gridstep resolution was 5mm.

The LCMV GLM beamformer (OSL toolbox, Woolrich et al., 2011) was chosen for source reconstruction of MEG data because it automatically normalizes MEG signal source locations to an average brain source space (MNI152 template), which aids group level statistical comparisons. Group level comparisons were necessary because some of the measured effects in this study were expected to be small and not visible at the individual subject level. The beamformer was also one of the two most consistent methods for localizing MEG signals accurately to early visual areas in the study of Chapter 6, the other being MNE. However, unlike MNE, the beamformer does not impose anatomical priors on the MEG source reconstructions, and this was deemed more appropriate for this study as individual anatomical information was not available for 4 of 7 subjects. Data were bandpass filtered 3 to 40 Hz prior and the beamformer was used to estimate brain sources of MEG signals in each trial.

7.2.2.5. Experimental contrasts

Trial-wise general linear model (GLM) statistics were used to identify brain responses that corresponded with experimental factors. When fitting the GLM, the dependent variable (activity at each source) is regressed against a *design matrix* which indicates the experimental condition of each trial. This allows estimation of the mean and the standard deviation of *regression parameters* across trials, for each point in time. These parameters estimate the effect of each experimental condition on MEG brain activity. Contrasts of parameter estimates (COPEs) are linear combinations of parameter estimates (see Chapter 5 for details). In beamforming, there is ambiguity regarding the directionality and sign of the activity of a point in the source space. To over come this, parameter estimates must be rectified (converted to their absolute value by the modulus function) during the statistical process. Depending upon the experimental test and prior hypotheses, parameter estimates can either be rectified prior to contrast calculation, forming *coapes* (contrast of *absolute* parameter estimates) or the contrast itself can be rectified, forming *acopes* (*absolute* contrast of parameter estimates).

In this study, I tested the MEG data for the following experimental effects:

(i) Lateralized motor effects: Left button press (B1) contrasted against right button press (B2), such that COPE = 1∗B1 + (-1)∗B2. Time window was -100 to 100 ms peri button press; no baseline was used. Parameter estimates were rectified prior to contrasting (*coape*) to show directional effects of left button presses ($t > 0$) and right button presses ($t < 0$) in different brain areas.

(ii) Responses to cylinders: The onsets of CW (B1) and CCW (B2) supra-threshold cylinder stimuli (balanced reward condition MM) were combined, such that COPE = (1/2)∗B1 + (1/2)∗B2. Parameter estimates were rectified after contrasting (*acope*). Time window was 0 ms to 500 ms to capture the visually evoked response to stimulus onset; baseline was -250 to 0 ms.

(iii) Biased versus balanced reward cue: Onsets of HL (B1) and LH (B2) reward cues were contrasted with the onset of the neutral MM (B3) cue, such that COPE = (1/2)∗B1 + (1/2)∗B2 + (-1)∗B3. Parameter estimates were rectified after contrasting (*acope*). Time window was 0 to 600 ms, to capture evoked and ongoing responses to the reward cue; baseline was -400 to 0 ms.

(iv) Differential responses to high versus low reward cylinders: Responses to CW cylinders in the HL condition (B1) and CCW cylinders in the LH condition (B2)) were contrasted with responses to CW cylinders in the LH condition (B3) and CCW cylinders in the HL condition (B4), such that COPE = (1/2)∗B1 + (1/2)∗B2 + (-1/2)∗B3 + (-1/2)∗B4. Parameter estimates were rectified after contrasting (*acope*). Time window was 0 ms to 500 ms, to capture the visually evoked response as for section *(ii)* above. A longer baseline (-500 to 0 ms) was used in order to estimate noise measures as reliably as possible, as experimental effects were expected to be small.

7.2.2.6. Group statistics

Group level t-statistics were calculated in one of two ways: mixed effects (MFX) or fixed effects (FFX). In both cases, the observations upon which the test was performed are the COPEs (not t-stats) from individual subjects, but the two methods differ in the source of their variance estimate. In MFX, also known as 'random' effects (RFX), error variance estimation is based upon the variability of the COPE effect between subjects, and the degrees of freedom are consequently the number of subjects. This inference can be generalized to the population from which subjects were selected (Friston, 2003; Penny & Holmes, 2003). In FFX, error variance estimation is based upon the trial-by-trial variability of the COPE effect within subjects, by assuming that each trial

represents an independent observation; the degrees of freedom are the number of trials. FFX is therefore less conservative than MFX and should not be used to make population inferences; it is more appropriately used to report results of case studies. In this study, MFX is used for all group level analysis unless otherwise specified. Statistical p-values were not calculated.

7.2.2.7. Display of MEG source analysis results

Unless otherwise specified, MEG source results are displayed as t-statistics from the group level inference, smoothed and projected onto the 2x2x2 mm standard template MNI152 brain in *fslview* (FSL, FMRIB, Oxford University). Figures show either the momentary brain source activity at a given time point, or the t-statistic averaged arithmetically across a time window. Radiological orientation is used, such that the left cortical hemisphere is on the right side of the image and *v.v.*. MNI coordinates of each sagittal (x-axis), coronal (y-axis) or transverse (z-axis) slice are reported and threshold values for displaying t-statistics are labelled. Timecourses of the COPEs or t-statistics at the individual subject level or the group level are displayed, taken either from a particular voxel or averaged over a region of interest (ROI) mask. ROI masks were taken from the Juelich Histological Atlas available from FSL Tools (Eickhoff *et al.*, 2005, 2007).

7.2.2.8. Eye movement analysis

Mean and SD were analyzed for each eye movement channel (eyetracker, vertical EOG, horizontal EOG) to see whether this could account for any differences in MEG brain activity between experimental conditions. If data were not normally distributed (Lilliefors test), a Wilcoxon sign-rank test was used to test for differences between conditions of interest; otherwise, a t-test was used. No correction for multiple comparisons was applied in order to liberally detect any effects at $p<0.05$.

7.3. Results

7.3.1. Behavioural results: effect of reward

In the neutral reward condition (MM) there were 5 points available for a correct choice in either cylinder rotation direction. Participants' choices showed the expected monotonic increase with cylinder disparity when fitted with a cumulative Guassian psychometric function (PMF) (Figure 7.1 a). Reaction times (RTs) showed the expected inverted-U shape, such that RTs were

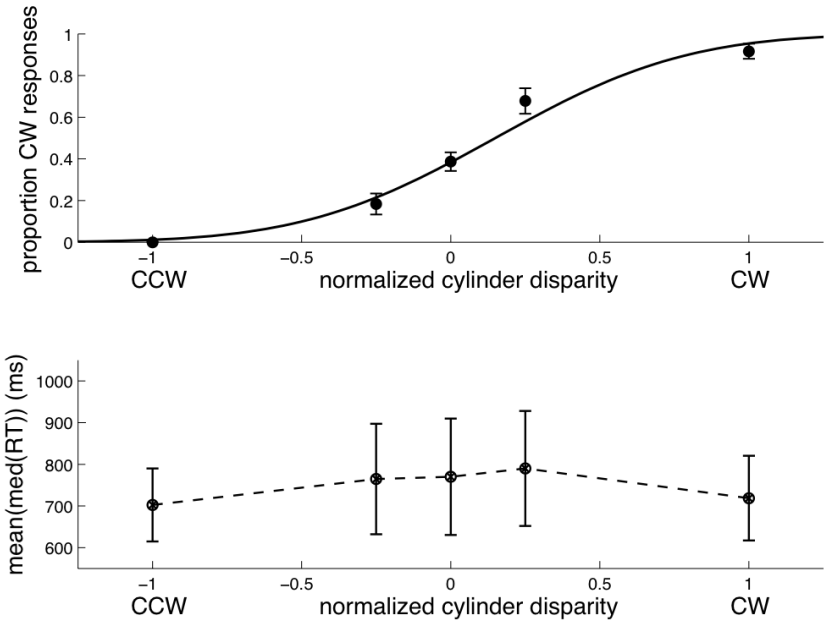


Figure 7.1.: *Upper panel*: Psychometric function (PMF) of responses in the cylinder discrimination task, averaged across subjects. *Lower panel*: median reaction times (RTs) for correct trials, averaged over subjects. Error bars: ± SEM. Different participants had different cylinder disparity ranges presented to them according to their unique performance threshold; therefore, cylinder disparities were normalized by division by the maximum CW disparity before combining data across subjects.

shorter (decisions faster) for larger cylidner disparity values (i.e. for easier trials) (Figure 7.1 b). Generally, individuals' median RTs were slightly faster than those recorded outside the scanner (Chapter 4) and disparity discrimination sensitivity was poorer (thresholds higher) (Table 7.4). This suggests that participants placed greater value on speed rather than accuracy when performing the task inside the MEG scanner.

In the unbalanced reward conditions (HL & LH), more points were available for a correct choice in one direction than the other (Table 7.1). Participants were more likely to report a CW cylinder when more points were available for the CW direction (HL condition) and more likely to report a CCW cylinder when more points were available for the CCW direction (LH condition) (Figure 7.2). At the group level, the effect of reward bias on choices was highly significant (n = 48, χ^2 = 16.4, p $<$ 0.0001, Kruskal-Wallis), confirming that the unbalanced reward conditions had the desired effect of biasing choices in this experiment, as in the psychophysics study of Chapter 4. CW responses to the onset of a CW cylinder tended to be faster when the bias was towards CW choices than CCW choices, and vice versa for CCW responses to CCW cylinders (Figure 7.2); however, this interaction was not significant at the group level.

	MEG scanner (current study)				psychophysics (Chapter4)			
	medium RT				medium RT			
subject	supra-thres	near-thres	bistable	*SD*	supra-thres	near-thres	bistable	*SD*
mrw1	1209	1523	1528	0.025	668	898	858	0.011
mrw2	383	421	437	0.012	415	448	458	0.008
mrw3	677	715	686	0.043	682	817	835	0018
mrw4	708	547	525	0.069	805	950	845	0.042
mrw5	589	622	576	0.033	989	1196	1213	0.015
mrw6	716	751	706	0.084	396	591	580	0.033
mrw7	712	854	934	0.020	-	-	-	-
mean	*713*	*776*	*770*	*0.041*	*659*	*816*	*798*	*0.021*

Table 7.4.: Median RTs of participants' responses to supra-threshold, near-threshold and bistable (zero-disparity) cylinders, and their discrimination thresholds (calculated from the standard deviation (*SD*) of the cumulative Gaussian fitted to individuals' PMFs). Values obtained in this study, where participants performed the task in the MEG scanner, are presented along with the equivalent values for the psychophysics study (Chapter 4), where available. Values for the present study are taken from the balanced reward (MM) condition only.

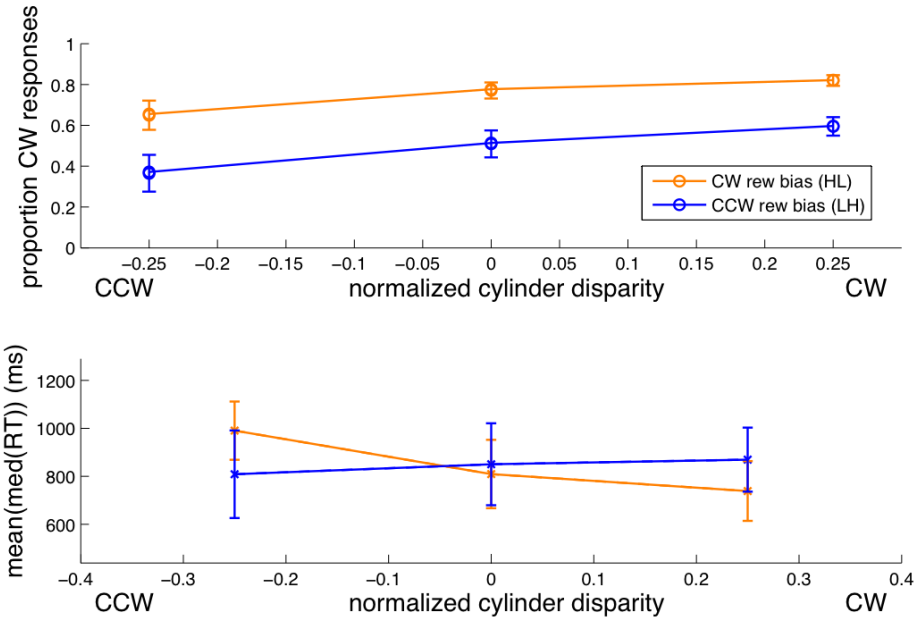


Figure 7.2.: *Upper panel*: Choices in the cylinder discrimination task under the two different unbalanced reward conditions, averaged across subjects. *Lower panel*: median reaction times (RTs) for correct trials in the two unbalanced reward conditions, averaged over subjects. Error bars: ± SEM. It was not possible to fit a psychometric function to this data because only three cylinder disparity values were presented in the unbalanced reward conditions; hence, response data are presented in a line plot. Again, cylinder disparities were normalized by division by the maximum CW disparity before combining data across subjects.

7.3.2. Lateralized motor response peri-button press

Participants indicated their perceptual choice by releasing a button with their thumb, referred to as a 'button press'. A contrast analysis between left and right button presses, -100 ms to 100 ms peri-button press, was run to test source localizations, specifically lateralization of motor brain activity, by the LCMV beamformer and group level GLM statistical approach. This approach would reveal differences in brain activity surrounding left button presses compared to right button presses. In accordance with electrical stimulation and fMRI studies (Penfield & Boldrey, 1937; Dassonville *et al.*, 1997), I expected the resulting activity to generally be located around the central gyrus (primary motor and primary somatosensory cortices), contralateral to the side of the button press.

For all subjects, responses to left button press were localized near the central gyrus of the right cortical hemisphere and responses to right button press were localized near the central gyrus of the left cortical hemisphere, with the same pattern seen at the group level (Figure 7.3). The single highest-responding vertices to each of the left and right button press conditions were selected from the group level result. The time-courses of the left-right activity contrasts at these vertices followed a similar pattern for all subjects: activity peaked at, or just prior to, the time of button press (Figure7.4). Activity contrasts also showed high values at the boundaries of the time window. This may be because the response involved the release of the button by the thumb followed by replacement over the button, resulting in on-going sensory and motor activity. Overall, sources lateralized well and closely matched predictions for brain activity surrounding unilateral button presses. The results give confidence that the MEG source reconstruction and statistical approach applied are valid for investigating further experimental questions in this study.

7.3.3. Visual responses to onset of supra-threshold cylinders

Group level (FFX) source reconstructions of visually evoked responses to the onset of supra-threshold cylinders showed strong activity localized to the occipital pole at 50 -100 ms, corresponding to the retinotopic representation of the central visual field (where the cylinder was presented) by early visual areas V1, V2 and V3 (Figure 7.5). At 100-150 ms, activity at the occipital pole increased in strength and extent. At 150 - 200 ms, activity appeared bilaterally in area hMT+, the human analogue to macaque visual area V5/MT (Huk *et al.*, 2002; Mantini *et al.*, 2012). Subsequently, activity was restricted to the RH lateral occipital areas at 200 - 250 ms. From 250 - 300 ms, activity also appeared in RH posterior parietal region (PPC). Between 300

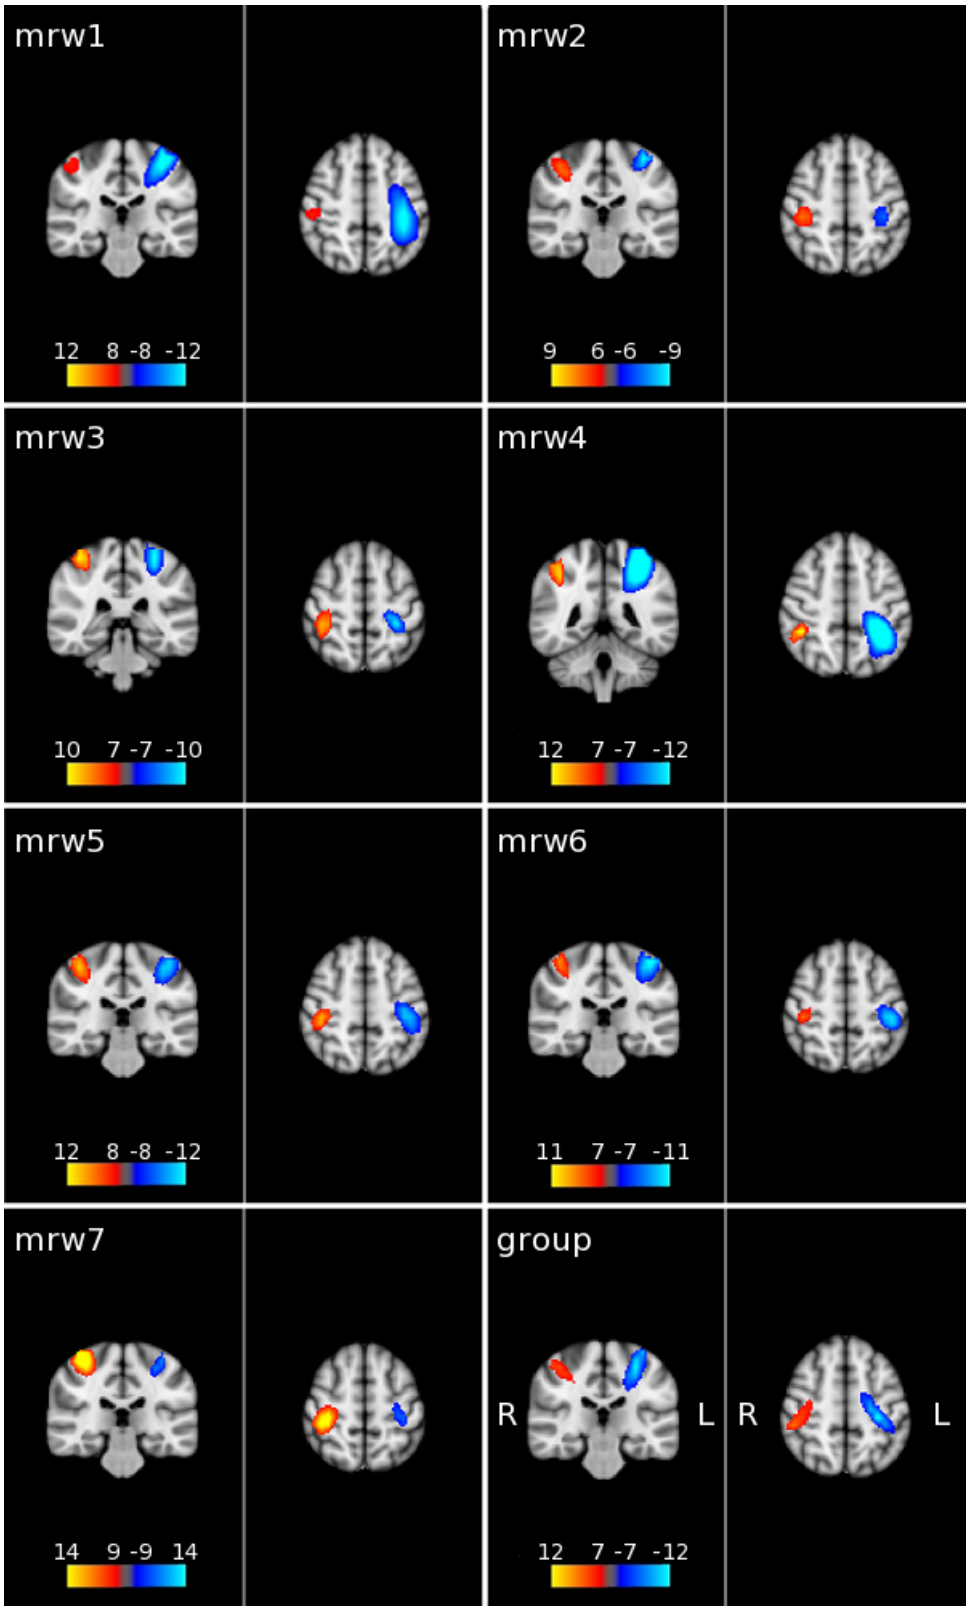


Figure 7.3.: MEG source comparison results (t-statistics) for left button press activity (positive t, *red*) and right button press activity (negative t, *blue*), shown for individual participants and the group level. Time: 10 ms before button press. Colour bar indicates t-statistic range for each plot; t thresholds were selected to show the location of the peak activity cluster for each individual and for the group level. Radiological convention places the left cortical hemisphere on the right-hand side of the image. The results show that responses to left button press (red) are localized on the right cortical hemisphere, and responses to right button on the left hemisphere.

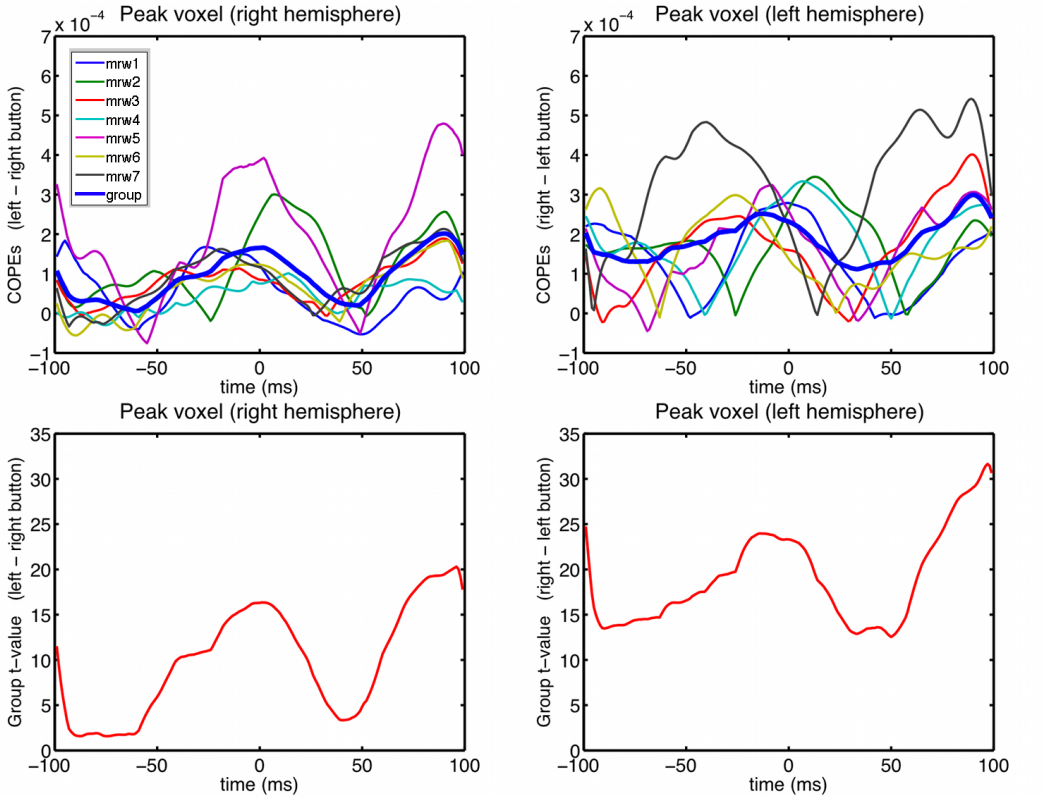


Figure 7.4.: Timecourses of contrasts of parameter estimates (COPEs) and t-statistics were taken at the voxel of greatest group level activity on each of the right and left hemispheres (Figure 7.3). MNI coordinates for these peak voxels were [48, -30, 56] and [-30, -8, 54] for the right and left hemispheres, respectively. *Upper panels*: COPE time-series for individual subjects and group level; *lower panels* group level t-statistic time-series. Time 0 refers to time of 'button press' (lift of the thumb from the optic response button). The time-series show a peak in activity at or just before the time of contralateral button press.

- 450 ms activity across the occipital brain was significantly decreased, but from 450 - 500 ms, there was a strong rebound in the region of early visual areas, dorsal occipital cortex and right PPC. Time-courses of activity t-statistics showed that activity in the region of primary visual cortex (V1) peaked between 100 - 150 ms after stimulus onset, whilst activity in the region of hMT+ (V5) peaked later, between 150 - 200 ms (Figure 7.6). V1 timecourses began to increase again at 450 ms.

The rebound of V1 activity at 450 ms could be due to a change in position or variability in eye movements in the second half of the analysis epoch, resulting in changes to visual input and thus visual brain responses. Eye movements were compared between the first half (1 - 250 ms) and second half (251 - 500 ms) of the epoch, for each of the five subjects for whom this information was available. Two subjects, *mrw2* and *mrw5*, showed significant differences in at least one measurement (Table 7.5). Source reconstructions were repeated excluding these two subjects. The activity at 450 ms was still present, at an equal t-statistic value, despite the fewer subjects

(Figure 7.7). It is therefore unlikely that eye movements can account for this rebound in activity in early visual areas.

subject:	mrw1		mrw2		mrw5		mrw6		mrw7	
eye tracking:	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**
horizontal	0.10	0.05	0.76	*0.032*	*<<0.001*	*<<0.001**	0.18*	0.50*	-	-
vertical	0.78	0.65	0.35	0.11	0.60	0.15	0.67*	0.53*	-	-
pupil size	0.97	0.86	0.40	0.43	0.92	*0.005**	0.43*	0.87*	-	-
h EOG:	-	-	-	-	0.004*	0.33	0.19*	0.29	0.85	0.72*

Table 7.5.: Table of p-values obtained from t-tests comparing the mean and standard deviation (SD) of eye movement parameters between 0 - 250 ms and 251 - 500 ms after stimulus onset, across maximal disparity cylinder trials separately for 5 subjects for whom data were available. All tests are Wilcoxon sign-rank, except where data were normally distributed ($p>0.05$, Lillifors test) for which a paired t-test was used instead (indicated by *). Significant differences in eye movements are indicated by *italics*, uncorrected $p<0.05$. Two subjects, *mrw2* and *mrw5*, showed significant differences in eye-movement measures.

The results reported so far were obtained with the fixed effects (FFX) statistical group level approach, for which variance estimates are based on trial-by-trial variability in signals within individual subjects. When the more conservative mixed effects (MFX) approach was used, for which variance estimates are based on effect variability between subjects, t-statistics were lower as expected, but the main characteristics of the data were maintained, in particular the activation of bilateral hMT+ around 175 ms and rebound activation of early visual areas and parietal cortex around 480 ms (Figure 7.8). At 175 ms, bilateral inferior temporal cortical regions also showed responses, which was not seen with the FFX approach. On the other hand, the large and extensive activation of early visual areas 100 - 150 ms in the FFX analysis was not found in the MFX analysis. This could be due to large variability in the timing of the activation of early visual areas between subjects, such that at any given point in time, the group level effect is small relative to between-subjecct variability.

7.3.4. Differential responses to biasing versus balanced reward cues

How does the brain differentially respond to the onset of a reward cue that indicates a biased reward condition (HL and LH) compared to a balanced reward condition (MM)? I contrasted

Figure 7.5. *(following page)*: Group level FFX t-statistics showing brain responses to onset of supra-threshold cylinder stimuli. Activity is averaged systematically within 50 ms time windows, in time period 50 - 350 ms and 450 - 500 ms post-stimulus. MEG activity was localized with an LCMV meanformer; t value thresholds are selected to show the main clusters of activity in response to the cylinders. Different brain regions were identified with the Juelich Histological Atlas (fslview, FSL).

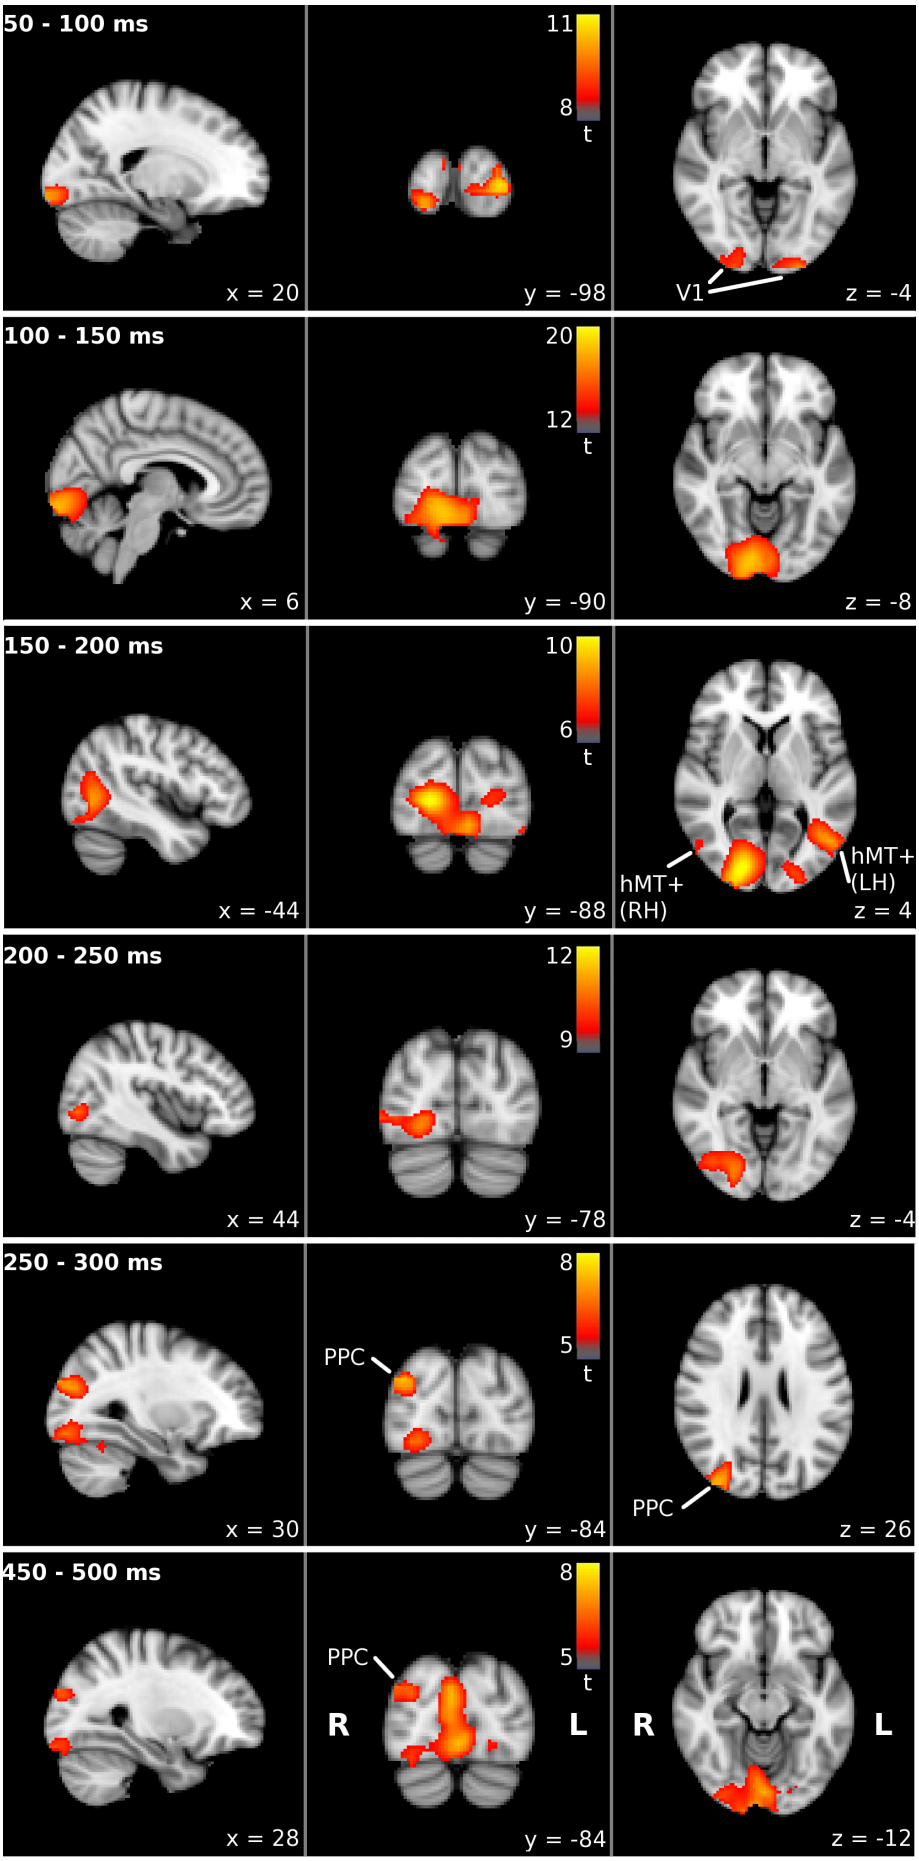

50 - 100 ms
11
8
t
x = 20
y = -98
V1
z = -4
100 - 150 ms
20
12
t
x = 6
y = -90
z = -8
150 - 200 ms
10
6
t
x = -44
y = -88
hMT+
(RH)
hMT+
(LH)
z = 4
200 - 250 ms
12
9
x = 44
y = -78
z = -4
250 - 300 ms
8
5
t
PPC
x = 30
y = -84
PPC
z = 26
450 - 500 ms
8
5
t
PPC
R
L
x = 28
y = -84
R
L
z = -12

Figure 7.5

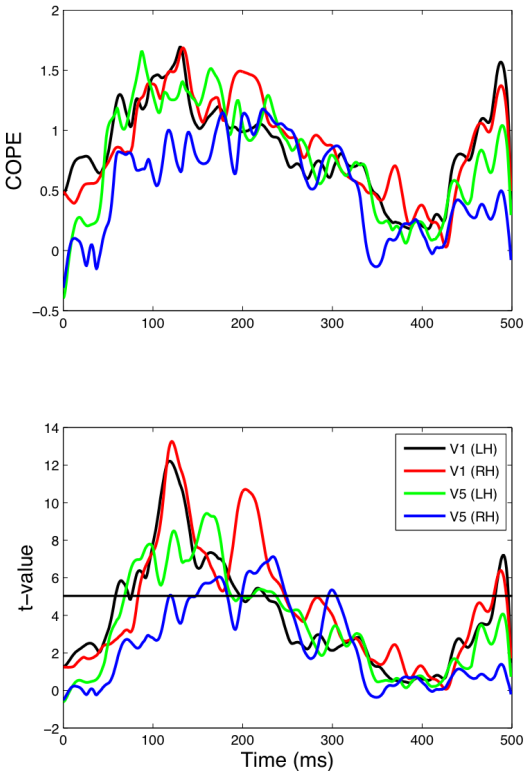


Figure 7.6: Time courses of COPEs and t-statistics for group level response to onset of supra-threshold cylinder stimuli. Activity is averaged across all voxels in the following regions of interest (ROIs): V1, left hemisphere; V1, right hemisphere; hMT+, left hemisphere; hMT+, right hemisphere. *Black line* indicate the minimum t threshold of Figure 7.5 (t≥5). Peak responses in area hMT+ occur later than in primary visual cortex (V1); this is easier to see in the t time-courses because the raw COPEs are not normalized by noise levels.

MEG brain responses to a biasing versus neutral points cue to inspect effects of bias condition *per se*, 0 - 600 ms post-cue onset, during which the viewer processes the reward cue and anticipates the onset of the cylinder stimulus.

At 175 ms, 225 ms and 275 ms post-cue, differences in signals for the contrasting conditions were present in the right and left early visual cortex (V1), as well as in the left posterior precuneus (visual part; Margulies *et al.* (2009)) (Figure 7.9). At 475 ms, a cluster was located in the left hMT+ region and at 590 ms a strong effect was present in right dorsolateral prefrontal cortex (DLPFC), concurrently with small clusters in the occipital cortex (Figure 7.9). Control analyses found no significant differences in eye movements between the two conditions that could account for differential effects in visual areas (Table 7.6), and no changes in the visual stimulus display occurred in the analysis time window that could otherwise explain the activity in visual areas.

7.3.5. Reward association affects responses to cylinders in visual precuneus and prefrontal cortex

What differences are there in brain responses to a cylinder stimulus when it corresponds to the more highly-rewarded choice direction, compared to when it corresponds to the lower-

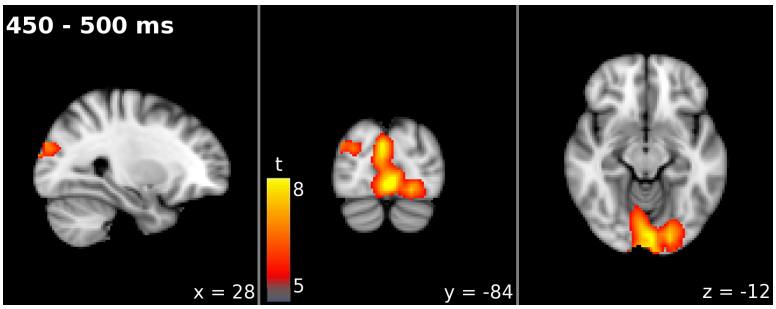


Figure 7.7.: Group level t-statistic for response to supra-threshold cylinders, 450 - 500 ms post-stimulus for five subjects (excluding subjects *mrw2* and *mrw5* who had a significance difference in eyemovements in the second half of the stimulus viewing period). Pattern does not look qualitatively different compared to Figure 7.5 for the same time window. T threshold is the same as Figure 7.5, lowest panel.

subject:	mrw1		mrw2		mrw5		mrw6		mrw7	
eye tracking:	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**
horizontal	0.855	0.609	0.650	0.0834	0.636	0.101	0.421	0.903	-	-
vertical	0.639	0.325	0.377	0.547	0.794	0.607	0.425	0.994	-	-
pupil size	0.781	0.498	0.774	0.716	0.923	0.949	0.519	0.397	-	-
h EOG:	-	-	-	-	0.177	0.395	0.587*	0.678	0.399*	0.897

Table 7.6.: Table of p-values obtained from t-tests comparing the mean and standard deviation (SD) of eye movement parameters between trials in which a neutral cue was presented and trials in which a biasing cue was presented (separately for 5 subjects). All tests are Wilcoxon sign-rank, except where data were normally distributed ($p>0.05$, Lillifors test) for which a paired t-test was used instead (indicated by *). No significant differences in eye movements were found in any measure (uncorrected $p<0.05$).

rewarded one? Brain responses to high-rewarded cylinders (CW-rotating cylinder in the HL condition & CCW-rotation cylinder in the LH condition) were contrasted with responses to low-rewarded cylinders (CW-rotating cylinder in the LH condition & CCW-rotating cylinder in the HL condition).

As early as 65 - 75 ms and 85 - 95 ms post-stimulus, clusters of differential activity were found in the left posterior precuneus region (visual part; Margulies *et al.* (2009)) (Figure 7.10). At 85 - 95 ms, there was also a small cluster located close to primary visual cortex. This was followed by clusters in the bilateral supplementary motor regions (SMA)/dorsomedial prefrontal cortex (DMPFCx) and left DLPFC at 125 - 135 ms. At 140 - 150 ms, a cluster was located in the right inferior frontal gyrus. After an interval of no differential effect, a cluster was located in the right medial anterior occipital cortex at 480 - 490 ms, between primary visual cortex and the precuneal region. Control analyses evaluated whether differential effects could be assigned to differences in eye movements between the two conditions; the only significant difference was in the standard deviation (SD) of the horizontal eye position in subject *mrw2* (Table 7.7). This suggests that for the majority of subjects, eye movements cannot account for the differences found.

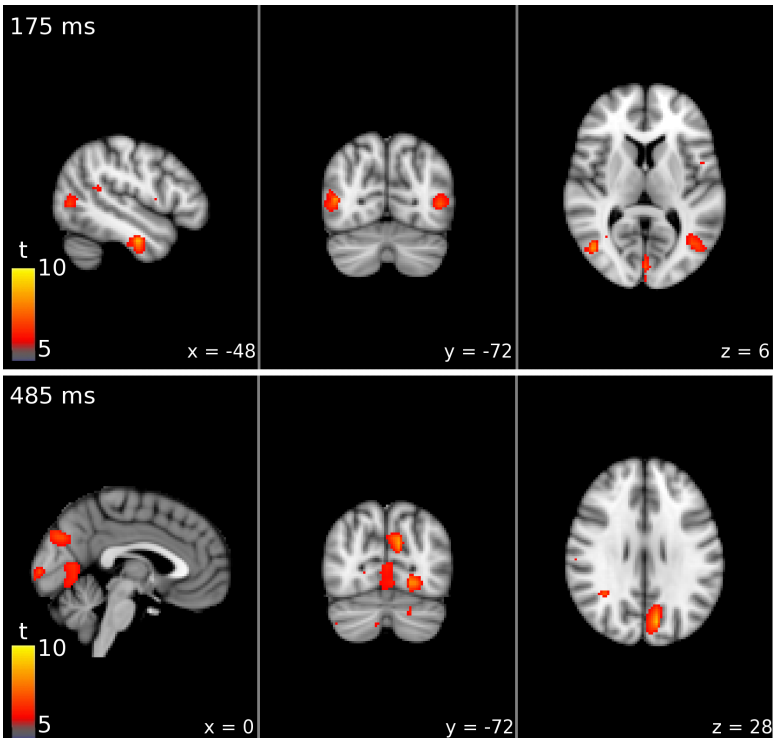


Figure 7.8.: Group level t-statistics (all subjects) for brain responses to supra-threshold cylinders, using MFX approach instead of FFX. Minimum t threshold for display is the same as in Figures 7.5 & 7.7. Although t values were generally much lower with the MFX approach than FFX, the main characteristics remained: there was bilateral activation of hMT+ at 175 ms and rebound of activity in early visual areas at 485 ms.

No differential activity was found in visual areas V1 or hMT+, which responded highly to the onset of supra-threshold stimuli. To investigate this further, brain sources of MEG responses to high-rewarded and low-rewarded cylinders were separately reconstructed, and timecourses of group level t-statistics in V1 and hMT+ ROIs were plotted separately for high and low reward stimuli (Figure 7.11). No consistent differences in the timecourses were seen.

subject:	mrw1		mrw2		mrw5		mrw6		mrw7	
eye position:	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**	**mean**	**SD**
horizontal	0.909	0.799	0.978	*0.012*	0.686*	0.792	0.997*	0.224	-	-
vertical	0.371*	0.533	0.486*	0.447	0.684	0.588	0.585*	0.212	-	-
pupil size	0.991	0.691	0.536	0.152	0.941	0.446	0.546*	0.456	-	-
h EOG:	-	-	-	-	0.831*	0.080	0.706*	0.533	0.880	0.884

Table 7.7.: Table of p-values obtained from t-tests comparing the mean and standard deviation (SD) of eye movement parameters between trials with highly-rewarded and low-rewarded cylinder stimuli (separately for 5 subjects). All tests are Wilcoxon sign-rank, except where data were normally distributed ($p>0.05$, Lillifors test) for which a paired t-test was used instead (indicated by *). Significant differences in eyemovements are indicated by *italics*, uncorrected $p<0.05$. Only one subject, *mrw2*, showed a difference in an eye movement measure (horizontal eye track position).

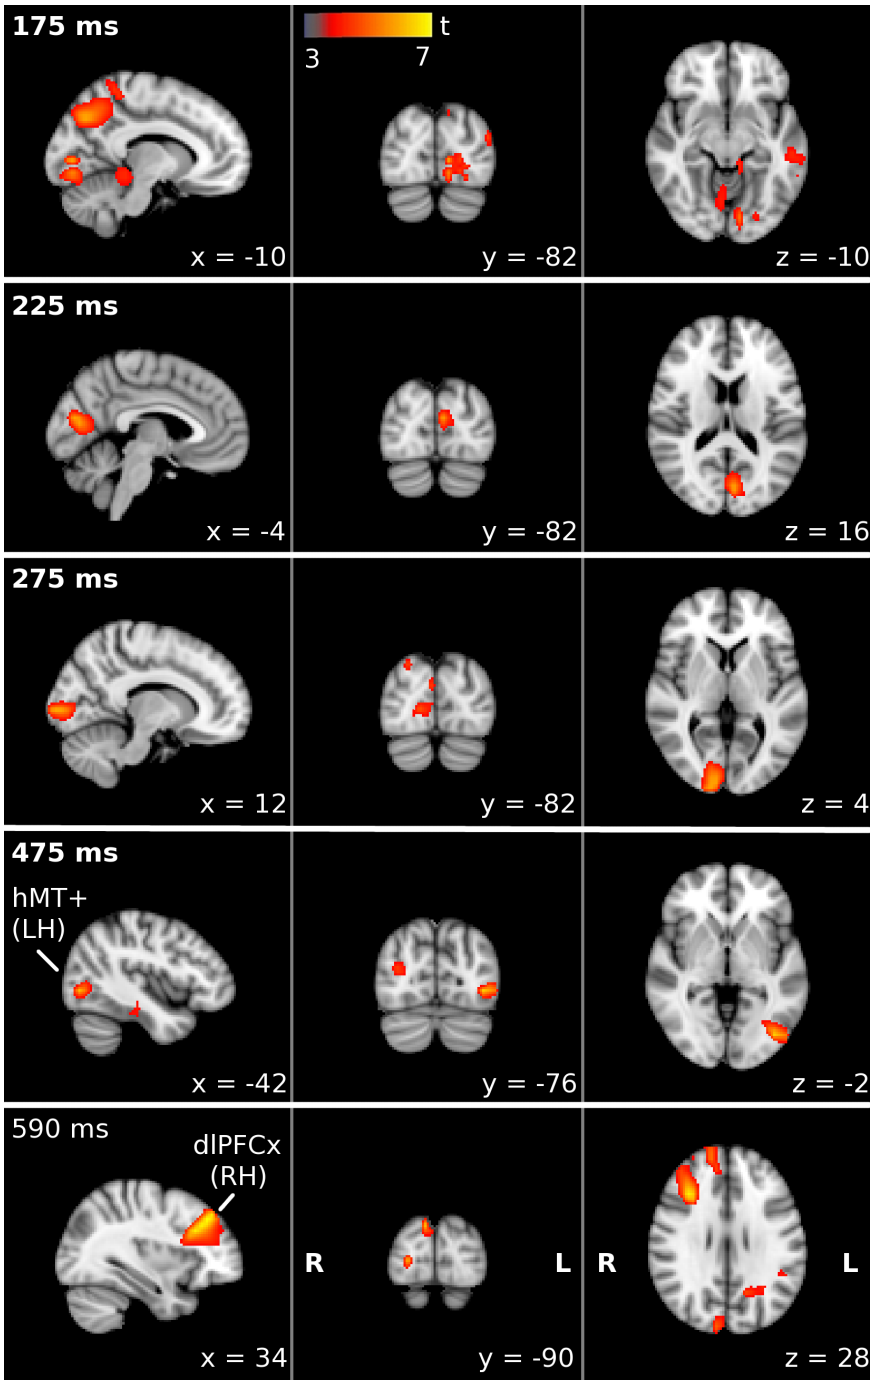


Figure 7.9.: Differential brain responses to the onset of a biasing versus balanced reward cue. Group level MFX t-statistics are displayed at selected time points after cue onset. T values > 3 are plotted. The left posterior precuneus (visual part) shows differential activity at 175 ms post-cue, and left and right early visual areas areas show differential responses at 225 and 275 ms. Differential signals are also present in left hMT+ at 475 ms post-cue and dorsolateral prefrontal cortex (dlPFCx) at 590 ms.

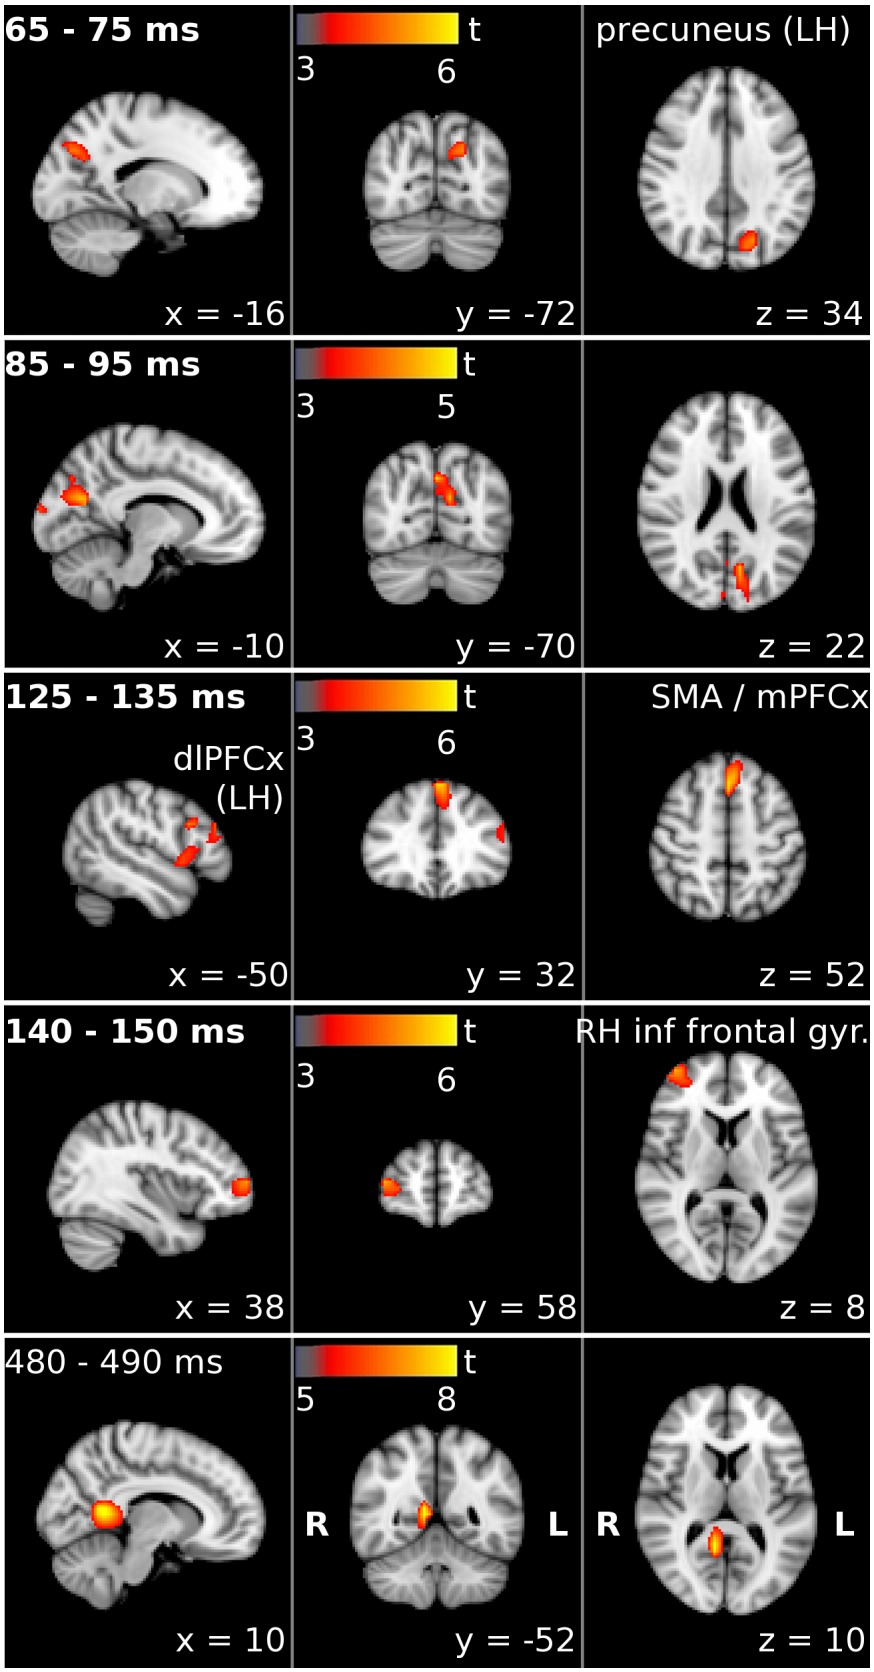


Figure 7.10.: Differential responses to the onset of high-rewarded and low-rewarded cylinder stimuli: group level MFX t-statistics are averaged over the time windows indicated. T thresholds ≥ 3 are selected to show main peaks of differential activity.

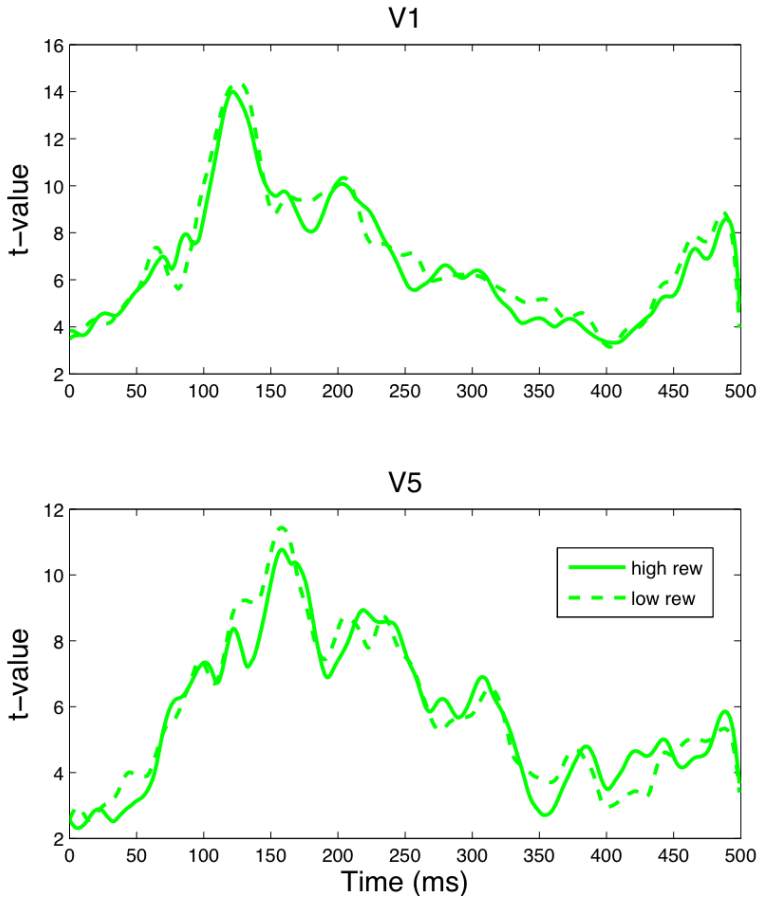


Figure 7.11.: Time-courses of group level MFX t-statistics, averaged over voxels in the V1 and hMT+ ROIs (combined over ROIs in right and left cortical hemispheres). Source reconstructions were made of MEG signals to onsets of high-rewarded and low-rewarded cylinder stimuli. The time-courses are generally very similar; further research is required to ascertain whether the small deviations seen constitute a significant difference in visual processing between conditions.

7.4. Discussion

Brain activity was recorded with magnetoencephalography (MEG) whilst seven human participants made perceptual decisions about structure-from-motion rotating cylinder stimuli. Behavioural choices were biased on a trial-by-trial basis by a *reward cue* which indicated the number of points available for a correct choice in each cylinder rotation direction. Group level source analysis of MEG signals with a GLM beamformer was verified by the localization of brain activity around the time of button presses to expected regions in the contralateral peri-central gyri. In response to the onset of a supra-threshold cylinder stimulus, MEG source analyses showed a sequence of activation in early visual areas V1, V2 & V3, hMT+ (human analogue of macaque area V5/MT), and posterior parietal cortex (PPC), amongst other regions. After reward cue onset (but prior to stimulus onset), responses in early visual areas, left visual precuneal region (parieto-occipital sulcus), and right dorsolateral prefrontal cortex (DLPFC) differed according to whether the cue indicated a biased or balanced reward condition. These effects could not be accounted for by differences in eye movements, nor in the visual symbols used for the reward cues, which differed not in contrast or luminance but only in orientation. Comparing brain responses to stimuli associated with the high-reward perceptual choice against brain responses to stimuli associated with the low-reward choice revealed effects in the occipital cortex and visual precuneus within 100 ms of stimulus onset, and subsequent effects in SMA/dorsomedial PFCx, left DLPFC, and right inferior frontal gyrus.

7.4.1. Lateralization of brain responses to button presses

When using a new method or analysis technique which is still in development, it is essential to ensure that the approach gives expected results in a test for which the outcome is known in advance. Although the LCMV beamformer has been used for MEG source reconstruction and evaluated previously (see Chapter 6's MEG retinotopy comparisons to fMRI), this application of the beamformer approach with the general linear model (GLM) to calculate t-statistics at the group level has not (see only Hunt *et al.* (2012)). I localized brain activity surrounding the time of the button press to locations in the contralateral peri-central gyrus (where the primary motor and somatosensory regions involved in button press are located), at both the individual subject and group level. In all cases, activity lateralized to the region of the contralateral peri-central gyrus as expected, giving confidence in the localization accuracy of the MEG analysis method.

7.4.2. Source reconstructions of MEG responses to supra-threshold cylinders

Results confirmed that early visual areas, the human MT+ complex, and the PPC responded to supra-threshold structure-from-motion (SFM) cylinders. ROI-based timecourses confirmed that V1 activity peaks at 100 - 150 ms, whilst activity in hMT+ (human analogue to macaque area V5/MT) peaks later, at 150 - 200 ms. These timings are similar to stimulus-selective responses of neurons in areas V1 (Super *et al.*, 2001), V5/MT (Pack *et al.*, 2001) and V4 (Bichot *et al.*, 2005) recorded in the macaque. The late parietal activity (250 - 300 and 450 - 500 ms) may reflect accumulation of sensory information into a decision variable, a human analogue of the evidence accumulation observed neurons of parietal area LIP of the macaque (Roitman & Shadlen, 2002).

Early visual area activity rebounded at 450 - 500 ms; this has not previously been seen in electrophysiological studies with the same stimuli in macaques, albeit in fixed-duration tasks (*personal observation & personal communication, Dr K Krug*). However, the findings are similar to those reported by Jiang *et al.* (2008), who recorded MEG responses to SFM spheres and cylinders whilst viewers identified the stimulus shape in a reaction time task. They found activity in V1 at 100 ms and in hMT+ at 200 - 300 ms, inferior temporal cortex at 400 - 500 ms and the same rebound of activity in visual areas and PPC around 500 ms. In the present study, rebound activity could not be accounted for by eye movements. However, even though median RTs of participants to supra-threshold cylinders were generally longer than 500 ms, a possibility remains that very short RT trials, which result in early stimulus offset and therefore a change in visual input, may underlie this visual area rebound. Alternatively, the explanation may be due to the differences in the source of the electrophysiological measurements: action potentials in the case of single cell recording in primates, and synaptic activity in the case of human MEG (see Chapter 5). A goal of future analysis should be to explain the source of this rebound activity.

Jiang *et al.* (2008) reported activity in area LO 350 ms post-stimulus, but this was not found in the present study. This may be due to task differences; in the present study, viewers discriminated the direction of cylinder rotation by using visual motion and binocular disparity information. In the task of Jiang *et al.*, viewers discriminated the *shape* of the stimulus. Regions LO and inferior temporal cortex, part of the ventral visual stream important for processing shape information and recognizing visual stimuli (Goodale & Milner, 1992), might therefore be activate in their study but not in ours. Future studies could use MEG to directly compare the location and timing of the brain responses to the onset of the same SFM stimuli in different tasks, such as rotation versus shape discrimination.

In order to increase statistical power, fixed effects (FFX) inference, which estimates variance from the variability across all trials in all subjects, was used at the group level to arrive at the findings above. When mixed effects (MFX), which uses the variability of the effect(s) between subjects to estimate variance, was applied instead, the values of the group level t-statistics were far lower. This suggests that the inter-subject variability in the responses to the cylinder stimulus onset was large and more subjects need to be tested. Regardless, the MFX analysis replicated the main findings of the FFX analysis.

7.4.3. Differential activity in early visual areas, visual precuneus, and prefrontal cortex in visual perceptual decision-making that is biased by reward information

Before the onset of the cylinder stimulus in each trial, a cue indicated to the viewer how many points were available for a correct choice in each response direction. The HL and LH reward conditions had unequal points available for the two possible choices, and resulted in behavioural biases. A biasing cue thus indicated that the upcoming perceptual decision had unequally rewarded outcomes. On the other hand, a balanced cue indicated that the reward size for a correct choice would be the same regardless of the choice direction. The different reward cues as visual stimuli were not expected to differentially affect brain activity at the level of whole-brain MEG measurement because they did not differ from each other in contrast or luminance.

Following the onset of a biased versus balanced reward cue, there were differential effects in right and left early visual areas and in the posterior precuneual area at a number of time points within 300 ms of cue onset, as well as in hMT+ (human V5) of the left hemisphere at 475 ms. At 590 ms, there was differential activity in the right dorsolateral prefrontal cortex (DLPFC). One explanation for effects on early visual areas is that the effect of the biasing cue is to modulate activity in these regions in response to information about upcoming payoffs associated with different perceptual choices, in preparation for stimulus onset. Regarding the precuneal cortical area, very little is known as it has not received much previous research attention due to its location in the medial wall of the posterior parietal lobe and its very low incidence of focal lesions. A few studies have implicated the precuneus in a variety of higher cognitive functions such as attention, visuo-spatial imagery, episodic memory retrieval, the default network and even consciousness and self-awareness (Cavanna & Trimble, 2006).

A recent resting-state MRI study in humans and macaques emphasized the specific functional

connections of the posterior part of the precuneus, the region showing activation in this study, to V1 and other adjacent visual cortical regions (Margulies *et al.*, 2009). Therefore, the visual precuneus is connected to areas involved in early sensory processing as well as being implicated in higher-level cognitive functions. The results reported here suggest that this region might play a role as a conduit between contextual factors represented in higher cortical areas and sensory processing implemented in visual cortex. This could be linked to the close proximity of the precuneus to the PPC, which has already been strongly implicated value-based decision-making functions (as reviewed in Chapter 1). Further research is required to elucidate the functions of this region, in particular what role it might play in the processing of visual information and the nature of the connection is between the precuneus, early visual areas, and higher cortical regions. The functions of the right DLPFC are slightly better-understood, and in particular it plays a role in conveying motivational incentives on the basis of rewards in decision-making (Kouneiher *et al.*, 2009). This role for right DLPFC is supported by the results presented here, as it responded differentially to a biasing cue, which carries strong reward motivations that bias responses, compared to a neutral cue, which does not bias responses.

I also investigated how reward context affected visually evoked responses to cylinder stimuli. I found that the locations in the occipital lobe, left visual precuneus, bilateral SMA/dorsomedial PFCx, left DLPFC, and right inferior frontal gyrus showed differential responses to the onset of a high-rewarded compared to a low-rewarded cylinder. Effects in visual areas occurred early, within 100 ms of stimulus onset, whilst effects in prefrontal regions occurred later, after 100 ms. In line with my earlier results from electrical microstimulation and psychophsyics (Chapters 2, 3 & 4)), this suggest that reward can affect the activity of visual regions in human occipital cortex during perceptual decision-making. This compliments the finding of Baines *et al.* (2011), where reward motivation affected putative perceptual signals at the sensor level, measured with EEG. The MEG source localization results presented here again implicate the posterior precuneal region in reward-biased perceptual judgements, and confirm previous studies that have shown the involvement of prefrontal cortical regions in decision-making tasks that involve reward (Kouneiher *et al.*, 2009; Rushworth *et al.*, 2009). Previous MEG studies have shown that SMA and DLPFC may represent the accumulation of sensory evidence into a decision variable (Donner *et al.*, 2009; Wenzlaff *et al.*, 2011) and that this activity can be affected by speed-accuracy strategy implemented by the viewer (Wenzlaff *et al.*, 2011). Here, I show that these same regions are also affected by reward biases. Furthermore, the role of response inhibition has been assigned to the right inferior frontal cortex by both fMRI and human lesion-mapping studies (Aron *et al.*, 2004). This relates well to the activation of this region in the contrast analysis

reported here, where a preparation to bias the motor response towards the higher-rewarded direction must be inhibited upon the onset of a cylinder stimulus rotating in the opposite, lower-rewarded, direction, when compared to the onset of the higher-reward cylinder.

The timings of the effects, when inspected in the context of the general brain responses to the onset of cylinder stimuli, suggest that effects in the occipital lobe, visual precuneus and cuneus occur before or during the peak response of the early visual areas, especially V1, to the cylidner onset (100 - 150 ms), whilst effects in the dlPFCx and other prefrontal areas occurred at or before the peak response of area hMT+ (V5) to the cylinder onset (150 - 200 ms). Future studies should investigate how the functional connectivity between sensory and prefrontal brain regions might change under different reward conditions. Additionally, it could be that reward signals are relayed to the visual precuneus or visual areas from non-cortical brain regions that directly represent reward and value, such as the dopaminergic systems or the amygdala, the activity of which cannot easily be measured by MEG due to their deep locations in the brain. Further work is necessary to investigate whether effects of reward on sensory evidence arise from higher cortical areas or from sub-cortical reward structure, or both.

I found no specific effect of the reward association of the cylinder stimulus in area V1 (foveal region) nor in early hMT+ (V5) either in the whole-brain analysis nor in ROI time-course comparisons. Since any reward effects would be expected to be small, the lack of finding could be because the analysis was not sensitive enough due to the small number of subjects (see below). Alternatively, the ROIs obtained from the Juelich Histological Atlas, used to inspect timecourses, may not be specific enough for individuals' brain regions in this study, decreasing signal to noise ratio. Future studies should use fMRI to functionally localize hMT+ and other visual regions in the same participants. It might also be that reward effects on selective neuronal subpopulations cannot be easily discerned at the gross level measured by MEG and that more specific electrophysiologial techniques, which record from feature-selective neuronal populations, would be more appropriate for delineating the effect of reward context on the processing of specific visual information. For example, if reward bias towards a CW choice increased the activity of visual neurons tuned to the CW cylinder direction, but decreased activity in neurons tuned to CCW direction, this might result in a very similar overall activity measurement from the visual area as a whole as measured by MEG.

7.4.4. Significance testing and sample-size

The general linear model (GLM) framework is well-established for group level significance testing in fMRI research (Friston, 2003; Penny & Holmes, 2003) but not well-established for MEG. The main challenge for significance testing of neuroimaging data is to appropriately correct for multiple comparisons. Bonferroni correction for multiple comparisons should not be applied across voxels as this would result in over-correction because brain activity is not expected to be independent across voxels. A more appropriate, well-established cluster-based method for significance testing for fMRI is Random Fields Theory (RTF), which corrects for multiple comparisons at the level of clusters that spread across voxels (Brett *et al.*, 2003). However, MEG source reconstructions have the added complication of being 4-dimensional, because they are sampled across time as well as 3D space, and RTF is not applicable in this case. The OHBA Analysis Group have recently pioneered a non-parametric 4D cluster permutation method to estimate the null distribution of cluster sizes and t-statistics (see Hunt *et al.*, 2012); however, it was beyond the scope of this study to incorporate this development and as a result only the group level t-statistics are reported here, rather than p-values. Moreover, the number of subjects in this study is quite low (n=7), which limits power of significance testing. Future work should increase the number of subjects and apply 4D permutation statistics. Further development to the GLM beamformer OSL toolbox is necessary to make the significance testing applicable across a wide range of MEG studies.

7.4.5. Future directions

A number of different analysis approaches could be applied to this dataset in future work. In particular, regression of individual behavioural parameters, such as extent of reward bias, against MEG activity could delineate different brain areas involved in the task and strategy. Further research into the relationship between attentional and reward effects on visual processing and perceptual decision-making, similar to the approach of Baines *et al.* (2011), should also be tested in this context. However, reward manipulations always have the potential to affect attention and indeed this could be one mechanism by which some reward effects on behaviour are mediated. Further research into into the effect of reward on functional connectivity between brain regions would also further our understanding of the spatio-temporal profile of brain activity underlying perceptual decisions (Siegel *et al.*, 2011). Future studies could also inspect changes in power in different frequency bands of the MEG signals from early visual areas, which might represent changes in processing activity but may be hidden in the broad-band signal analysis (3 - 40 Hz) inspected in this study (Wang, 2010).

8. Conclusions and future directions

8.1. Summary

Visual area V5/MT, which contains neuronal correlates of subjective perceptual decisions in the cylinder task (Dodd *et al.*, 2001), also has a causal role in the percept, because electrical microstimulation of stimulus-selective neuronal clusters shifted animals' perceptual categorizations towards the stimulus represented by the microstimulated neurons. This artificial, electrical addition to the sensory evidence was used as a tool to explore the mechanisms by which the size of reward available for a correct choice might influence perceptual judgements (Chapters 2 & 3). The results, interpreted in the context of an adapted drift-diffusion model of decision-making (Palmer *et al.*, 2005; Gold & Shadlen, 2007), demonstrated that effects of reward must occur before visual stimulus-dependent sensory evidence and microstimulated sensory evidence are accumulated together in sensorimotor areas. Since microstimulation occurs in V5/MT, this supports the hypothesis of an effect of reward on sensory evidence in perceptual judgements at the level of V5/MT or earlier. In a human psychophysics study in which viewers' choices in the cylinder task are biased by unequal pay-off for correct choices in each direction, the analysis of choice data and reaction time distributions with a full drift-diffusion model of perceptual decision-making revealed effects of reward on the drift rate model parameter (Chapter 4). Since drift rate depends upon sensory evidence represented in visual regions (Mazurek *et al.*, 2003), this finding also supports the hypothesis that reward can affect sensory evidence. Both these studies additionally revealed an effect of reward on the distance between the starting points of the evidence accumulation process and the decision bound(s), known to occur in sensorimotor brain regions, in line with previous studies on the effect of reward on such decisions.

A newly-established human neuroimaging technique, magnetoencephalography (MEG), has promise for investigating responses in human visual cortex with high temporal resolution. Since MEG analysis is still in development, the focus of Chapter 6 was to evaluate accuracy and contrast the ability of three MEG source analysis methods to localize signals to correct sub-regions

of early visual areas. Retinotopic visual stimuli were presented to viewers and source localizations of MEG signals were compared to retinotopic maps of visual brain areas, obtained with fMRI in the same individuals. The results of this study demonstrated that the beamformer and MNE approaches were most suitable for interrogating the activity of early visual areas with MEG, and could localize signals into the quater-field retintopic subregions of early visual areas. This result informed the the final study of this thesis, in which MEG brain recordings were taken from human participants whilst they made perceptual decisions in the cylinder under different reward conditions. Data were analyzed with the beamformer approach. Firstly, we confirmed the role of area hMT+ complex, the human analogue to macaque area V5/MT (Huk *et al.*, 2002), in responding to cylinder stimuli approximately 150 to 200 ms after stimulus onset, after the activation of primary visual cortex. Then, the effect of a response-biasing cue for unequal pay-offs, versus a neutral cue of equal pay-offs, was investigated, as well as the brain response to the onset of a highly-rewarded, versus a low-rewarded, cylinder stimulus. The results imply that reward affects visual processing in the human brain prior to and during perceptual judgements, and revealed for the first time an effect on the visual precuneus in these perceptual decisions, but further subjects are required to fully confirm these effects.

Taken together, the results presented in Chapters 2, 3, 4 and 7 provide strong, consistent evidence that reward influences the processing of sensory information that underlies visual perceptual decisions in monkeys and humans. However, the results fall short of providing a direct observation of exactly how and where reward influence perceptual decision-making. Further work is required to provide direct proof.

8.2. Future directions

8.2.1. Further analysis of MEG data & the potential of connectivity studies

The MEG data presented in Chapter 7 of this thesis suggests that the onset of a reward cue which indicates unequal pay-offs, as opposed to equal pay-offs, and the onset of a highly-rewarded compared to a low-rewarded cylinder stimulus, result in differential activity in early visual areas, the visual precuneus and prefrontal cortex. However, no direct effect of reward was found on broad-band MEG activity recorded from visual area hMT+, which is thought to contribute sensory evidence in the perceptual task in humans as well as in monkeys (Huk *et al.*, 2002). These data merit further investigation and analysis. Firstly, only seven subjects participated in the study, due to practical limitations (uncertain helium provision to the MEG scanner). This

small sample limits the sensitivity of group level inferences because of the relatively large variability of effects between individual subjects. The experiment needs to be repeated on at least 3-5 more subjects, to increase the power of the analyses. Non-parametric methods of 4D cluster analysis must then be developed and applied to the data to infer the significance of the effects (p-value, see Hunt *et al.* (2012)). Secondly, MEG signals are complex and multi-dimensional, and many analysis approaches are still in development. Further analysis of the MEG data using new analysis techniques that incorporate time-frequency decomposition and connectivity modeling should be done, as these may reveal effects of reward on sensory processing that lie in the interactions between signals in different areas and which cannot be measured in the basic, broadband analysis approach carried out in Chatper 7.

There is growing evidence that when sensory evidence is processed for perceptual judgements, bottom-up information about the visual stimulus is propagated forward on a gamma frequency channel, whilst top-down priors are propagated backwards to sensory regions on a beta frequency channel (Wang, 2010). In a recent human MEG study, (Arnal *et al.*, 2011) showed that when visual speech input was incongruent with auditory signals, the degree of audio-visual incongruence scaled the coupling between low-beta and high-gamma frequency band activity in a multisensory region of the superior temporal sulcus (area STS). An effect of reward on sensory processing in the cylinder task might similarly be observed in time-frequency decompositions to certain frequency channels in visual and sensorimotor brain regions. Specifically, compared to the neutral reward condition where rewards for both perceptual outcomes are equal, it might be expected that when there is a reward bias toward a particular perceptual outcome, more top-down information is transmitted to early visual areas, especially those associated with larger pay-off. According to the hypothesis, this would be expected to be revealed in a beta-frequency band, posited to carry such top-down information. On the other hand, when stimulus information contradicts the prior reward bias, there might be an increase in gamma-frequency band activity.

Effects of reward on visual sensory processing imply functional connectivity between cognitive brain regions, which represent reward predictions and behavioural strategies, and early visual areas that represent sensory evidence (Siegel *et al.*, 2011). Dynamic causal modeling (DCM) is a method of analysis of neuroimaging data which uses a neuronal mass model to infer functional connectivity between different brain regions (David *et al.*, 2004; Kiebel *et al.*, 2009; Friston & Dolan, 2010). Although inferences about causality based upon DCM analysis must be interpreted with caution (Lohmann *et al.*, 2012), DCM has strongly implicated the contribution of

feedback components in brain responses evoked by auditory stimuli, both in health (Garrido *et al.*, 2007) and disease (Boly *et al.*, 2011). Summerfield & Koechlin (2008) applied DCM to fMRI data recorded whilst humans judged simple visual stimuli with choices biased, through task instructions, towards a particular perceptual interpretation. They showed that forward connectivity was increased when expected and observed perceptions diverged and feedback connectivity increased with prior expectations. From these findings, we can posit particular hypotheses for perceptual decision-making under reward. For example, when the visual stimulus disagrees with the perceptual outcome that has the greater reward value, it might be that forward connectivity is strengthened. An application of DCM to the MEG would be necessary to evaluate these effects. Alternatively, functional brain networks during perceptual decision-making could be evaluated in the future by using recently-developed MEG beamformer connecitivity analysis techniques (Schoffelen & Gross, 2009), such as the nulling beamformer (Hui *et al.*, 2010).

8.2.2. Single cell electrophysiology: reward and sensory neurons

The most direct demonstration of an effect of reward on the sensory evidence that underlies perceptual decisions would be to record from single neurons in sensory brain regions, and measure how the information that they contain about the visual stimulus is affected by reward conditions. The most appropriate animal model would be the macaque monkey, as this species has a very similar visual system to that of the human and has the ability to learn the complex perceptual decision-making tasks under different reward conditions (Feng *et al.*, 2009; Teichert & Ferrera, 2010). Animals should be trained to make judgements in the cylinder task whilst single neuron recordings are made in area V5/MT, which contains the neural correlates of perceptual decisions in this task (Dodd *et al.* (2001); Chapter 2). Reward would be a droplet of fluid for a correct choices, and reward size for a correct choice should be pseudorandomly varied for each perceptual outcome and appropriately cued to the monkey. This would allow an investigation into the effect of decision value (reward) on mean firing rate, variance, and timing of the spike train from single neurons in area V5/MT. Since this neuronal activity correlates strongly with subjective perceptual judgements in the cylinder task (choice probability = 0.67, Dodd *et al.*, 2001), the effect of reward on the size of this correlation could be measured. In an unbalanced reward condition, these analyses would reveal any link between reward effects on early visual areas and the effect of reward on the drift rate of the diffusion model reported in this thesis. Furthermore, zero-disparity, bistable cylinder stimuli contain less information about the perceptual categorization of the stimulus compared to unambiguous stimuli. There might be a differential effect of reward on the neuronal signals that represent the bistable stimuli compared

to unambiguous stimuli, and this should be investigated.

Recordings of neuronal activity in parietal area LIP, which correlates with the accumulation of sensory evidence (Roitman & Shadlen, 2002), have already been made under different reward conditions (Rorie *et al.*, 2010). However, that study had some limitations, as the neuronal results from the two subjects varied considerably (Rorie *et al.*, 2010). This, along with the new evidence presented in this thesis, makes the case for a replication of that experiment to be carried out, measuring again the effects of reward on the accumulation of sensory evidence in sensorimotor regions. Moreover, simultaneous recordings should be made between single neurons visual area V5/MT and area LIP under different reward conditions, to directly probe the interactions between these two stages of the perceptual decision process and how these may be affected by reward condition. These studies should carefully control for potential effects of attention as manipulated in previous attentional paradigms (for discussion see Chapter 1).

8.2.3. Computational modeling of reward and sensory information processing

Finally, neural network models of sensory processing and perception must be designed which fully incorporate effects of reward on perceptual decision-making. The drift-diffusion model has already been extended to incorporate cortico-basal ganglia circuit mechanisms, which have been posited to implement the decision threshold, or decision bound, during decision-making (Lo & Wang, 2006). This has been further extended to incorporate effects of reward on synapses in LIP and basal ganglia, and can account for learning and behaviour of monkeys under different reward conditions in some non-perceptual value-based decision tasks (Soltani & Wang, 2006, 2010). As more data become available regarding exactly how reward affects the information held in neuronal activity in sensory and other brain regions, these findings can be incorporated into current models of perceptual decision-making.

There is an ultimate dichotomy inherent in perceptual processes, namely the tension between the need to use prior knowledge about the world to aid (or *bias*) the interpretation of noisy and incomplete sensory information as efficiently as possible, and the need to keep some record of what is/was *really there*, bias-free; not least so that we can modulate the extent of our perceptual biases in the future through trial and error (e.g. Gregory, 1997). Further research into how reward, along with other contextual or cognitive factors, affects the sensory evidence that underlies perceptual judgements will take us a long way forward into understanding this fascinating aspect of cognitive neuroscience.

A. Mathematical details of a drift-diffusion process applied to microstimulation data

The goal of this model was to understand the probability of a perceptual choice in a particular direction as a function of cylinder disparity (stimulus strength), microstimulation condition and reward condition, for the data presented in Chapter 2. The aim of the modeling analysis was to indicate where in the decision-making pathway the effects of reward might be acting. The model adaptation was based upon that presented by (Palmer *et al.*, 2005).

A.1. Basic drift-diffusion model

A.1.1. Stimulus strength C and response direction R

In the simple 2-alternative forced choice (2AFC) visual discrimination task used in the study presented in Chapter 2, in each trial the animal had to discriminate the direction of rotation of a binocular disparity-rendered, structure-from-motion cylinder stimulus (Figure 2.1). Two possible directions of rotation, Cyl_a and Cyl_b have the respective correct responses R_a and R_b. The difference between stimuli of rotation directions Cyl_a and Cyl_b consists of a change in the sign of the stimulus' binocular disparity value, C. By arbitrary assignation, $C > 0$ produces a stimulus from category Cyl_a and $C < 0$ produces a stimulus from category Cyl_b. The absolute value of C corresponds to the absolute stimulus strength. When $C = 0$, the stimulus is perceptually bistable at the short viewing durations used here such that without a response bias this stimulus is as likely to be identified as either Cyl_a or Cyl_b. In Chapters 2 and 3, I use the convention that a positive stimulus disparity value C corresponds to a cylinder stimulus rotating in a clockwise (CCW) direction as viewed from above, such that Cyl_a, Cyl_b, R_a and R_b represent Cyl_{CCW}, Cyl_{CW}, R_{CCW}, and R_{CW}, respectively.

A.1.2. The diffusion process: accumulating evidence $U(t)$

The crucial aspect of the drift-diffusion model of perceptual decision-making is the accumulation of evidence over time towards one outcome or the other by a diffusion process (Figure 3.1 A). In the particular mathematical instantiation of the drift-diffusion model adapted here, $U(t)$ represents the accumulated evidence at any stage of its path over time t. $U(t)$ is known as the *decision variable* (DV). When a neutral starting point for the diffusion process is assumed (equidistant from the two decision bounds), at $t = 0$, $U(0) = 0$. In the continuous version of the diffusion process, at t the expected value $E(U(t))$ and variance $Var(U(t))$ of the decision variable are defined by (Karlin & Taylor, 1975):

$$E(U(t)) = \mu t \tag{A.1}$$

$$Var(U(t)) = \sigma^2 t \tag{A.2}$$

Therefore, μ describes the expected rate (referred to as the *drift rate*) at which evidence accumulates into the decision variable. The squared diffusion coefficient σ^2 describes the expected rate of increase of the variance. In turn, these parameters are dependent upon the stimulus strength C, such that weak stimuli have values of μ near zero and strong stimuli have larger values of μ.

The continuous version of the drift-diffusion process described above can be simplified by approximation to a discrete-time version, where discrete increments of evidence are accumulated over n steps of time Δt. The change in the accumulated evidence over Δt is:

$$\Delta U_t = U(t + \Delta t) - U(t) \tag{A.3}$$

where ΔU_t is a normally-distributed random variable (Palmer et al., 2005). The subscript t can be dropped because the evidence accumulation is assumed to be the same over each time increment i.e. stationary. The distribution of $\Delta U(t)$ has mean and variance of $\mu\Delta t$ and $\sigma^2\Delta t$ respectively, and at time $t = n\Delta t$:

$$U(n\Delta t) = \sum_{i=1}^{n} \Delta U_i \tag{A.4}$$

Therefore, the continuous drift-diffusion process can be simplified by an approximation to a

Guassian random walk in discrete time (Brownian motion), where ΔU_i represents the momentary evidence for choice R_a or R_b at each time step Δt (Palmer *et al.*, 2005).

A.1.3. The diffusion process with two bounds, A and B

The diffusion model for choice behaviour has two decision bounds, which correspond to the two possible decisions/choices. The decision at the end of the trial corresponds to the decision bound reached by the accumulated evidence $U(t)$. Specifically, the following decision rule is applied at each time increment:
- if $-B < U(t) < A$, then continue sampling,
- if $U(t) > A$, then respond R_a,
- if $U(t) < -B$, then respond R_b.

As an example, Figure 3.1 A (Chapter 3) depicts a single trial, showing the drift rate of the decision variable accumulating the evidence for choice A over B as a function of time. The mean drift rate μ corresponds to the mean of the strength of the evidence. In the example, the distribution of sensory evidence is shifted away from μ=0 and towards the decision bound that corresponds to choice A. Indeed, in the example, the drift-diffusion path of the decision variable eventually reaches decision bound A. This drift-diffusion schema with a single decision variable is equivalent to race accumulator models, where the evidence for each alternative is integrated separately, under the simplifying assumption that the inputs to each accumulator are perfectly negatively correlated (Palmer *et al.*, 2005).

Importantly, this allows us to describe perceptual choices in the decision task using the model. For any given stimulus strength C, we can derive the probability ($P_a(C)$) of stopping at the upper bound A and from that, the complementary probability of stopping at lower bound B ($P_b(C) = 1 - P_a(C)$) (Luce, 1986; Ross, 1996). When drift rate $\mu \neq 0$:

$$P_a(C) = \frac{e^{2\mu B/\sigma^2} - 1}{e^{2\mu B/\sigma^2} + e^{-2\mu A/\sigma^2}} \tag{A.5}$$

and for the special case when drift rate $\mu = 0$, the limiting expression is:

$$P_a(C) = \frac{A}{A + B} \tag{A.6}$$

Equations A.5 & A.6 therefore predict the probability of each perceptual report as a function of

stimulus strength C.

A.1.3.1. Special case of symmetrical bounds B and constant variance σ^2

Intrinsic bias towards one choice or the other can be represented by the relative distances between the decision variable starting point and the locations of decision bounds A and B. However, in the microstimulation study presented in Chapters 2 & 3, there was no reason to assume a response bias. Therefore a simplifying assumption was applied such that $A = B$, setting both bounds equidistant from the starting point of the decision variable's diffusion path (as shown in Figure 3.1 A). Parameter B is henceforth used to indicate the distance of the decision bounds.

A further simplifying assumption is that the drift diffusion coefficient σ^2 is constant for all values of stimulus strength C. This means that model parameters B and μ can both be normalized by division by σ^2, and parameter σ^2 can henceforth be cancelled out from the model. Under these assumptions Equations A.5 & A.6 can be simplified and rearranged to the following logistic form (Link, 1992; Palmer *et al.*, 2005):

$$P_a(C) = \frac{1}{1 + e^{-(B\mu')}} \tag{A.7}$$

where B now represents the normalized decision bound distance and μ' represents the normalized drift rate.

A.1.4. Linear scaling of drift rate μ' by stimulus strenght C

A final simplifying assumption is made that the normalized drift rate μ' of the decision variable is linearly dependent upon the stimulus disparity strength C ($\mu' = kC$, for some k) rather than non-linearly dependent ($\mu' = sign(x)|kx|^{\lambda}$, for some k and λ). The former model, known as the *proportional-rate diffusion model*, has been shown to be a good general approximation for the scaling relationship between stimulus strength and drift rate in simple 2AFC visual discrimination tasks (Palmer *et al.*, 2005). Furthermore, the expected difference in firing rates between direction-selective neurons in area V5/MT is 0 when C = 0 and varies linearly, on average, as a function of motion strength (Britten *et al.*, 1993). Therefore, the effect of stimulus strength C on the rate at which evidence accrues into the decision variable is assumed to scale linearly by parameter k, such that k represents the sensitivity of the decision model to stimulus strength C. Perceptual choices in the task are then predicted as a function of stimulus strength C by:

$$P_{CW}(C) = \frac{1}{1 + e^{-(BkC)}} \tag{A.8}$$

Therefore, if μ' corresponds to the mean strength of the evidence towards one perceptual choice or the other, then the proportional-rate diffusion model indicates that the stimulus strength C is converted linearly into this neuronal evidence.

A.2. Adaptation of the drift-diffusion model to microstimulation experiment

A.2.1. Representation of electrical microstimulation S

To recap, neuronal activity in extra-striate visual area V5/MT represents evidence, evoked by the visual stimulus, which contributes to the perceptual decision via an influence on the drift rate (μ') of the decision variable (Figure 3.1). Parameter k linearly scales the sensitivity of this evidence to the strength of the visual stimulus, such that $\mu' = kC$. The effect of electrical microstimulation applied to a small selective group of neurons in V5/MT was conceived as an additional evidence (S) added to the visual stimulus-dependent evidence (kC) represented in the region. When microstimulation is applied, the drift rate is dependent on both the microstimulation and the visual stimulus-dependent evidence, such that $\mu' = kC + S$. This is consistent with the main reported effect of V5/MT microstimulation being an additive effect on reported percept in the direction of the stimulus selectivity of the microstimulated neurons (Salzman *et al.*, 1990, 1992; Uka & DeAngelis, 2006; Hanks *et al.*, 2006). The logistic function becomes:

$$P_{CW}(C) = \frac{1}{1 + e^{-(B(\kappa C+S))}} \tag{A.9}$$

As can be seen, changes to the stimulus scaling sensitivity k and the normalized distance-to-bound value B have different effects on the output of Equation A.9. Normally, a full distribution of chronometric (reaction time) data is required to differentially fit these two parameters to behavioural decision data (Ratcliff & McKoon, 2008). In the fixed duration (FD) task design used in Chapter 2, reaction time data was not available. However, the microstimulation intervention S allows an important distinction to be made between any changes to, or effects on, the stimulus sensitivity k of visuals stimulus-depedent evidence in area V5/MT and overall changes/effects

on the integration of all combined evidence up to the decision-bound B. Parameter k now specifically represents the sensitivity of the decisions to the stimulus strength, which can be understood as the gain between the stimulus strength and the aspect of the decision variable drift that is driven by visually evoked sensory evidence alone. Parameter B represents the distance between the starting point of the drift-diffusion process to the bound(s) and the overall gain of the accumulation rate: aspects that effect both visually- and electrically-evoked evidence. For simplicity, B will be referred to as the bound parameter and k as the stimulus sensitivity.

A.2.2. Representation of the null bias x

Finally, a well-described behavioural effect, observed in microstimulation experiments in which stimulated and non-stimulated trials are interleaved, is the *null bias*. This refers to the animals' tendency to have a choice bias towards the non-preferred (NULL) direction on non-stimulated trials (Salzman *et al.*, 1992). For simplicity, in this model the null bias is represented by the additive inclusion of a parameter x:

$$P_{CW}(C) = \frac{1}{1 + e^{-(B(\kappa C + S) + x)}} \tag{A.10}$$

Parameter x corresponds to the y-intercept of the logistic regression function.

Therefore, in its final form, the logistic function describing the simplified drift-diffusion model contains three parameters that require fitting to the data: k, B and x.

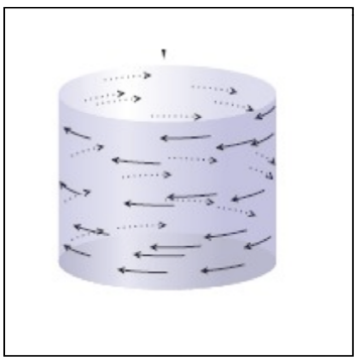
B. Cylinder task: participant instruction sheets

Figure B.1.: Instruction sheet for Chapter 4 training task:

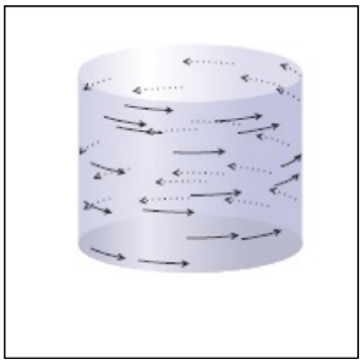
Single Rotating Cylinder: Training Task

A single cylinder is shown in the centre of the screen – it looks a bit like an upright drinks can. The cylinder is transparent and made up of many dots on its surface.

The dots on the front surface of the cylinder are moving either to the left or to the right:



front going **LEFT (or clockwise)**



front going **RIGHT (or counterclockwise)**

Task:

- Always look at the dot in the centre of the screen.

- A cylinder will appear on the screen. Click the left mouse button if the dots on the front surface are going to the left. Click the right mouse button if the dots on the front surface are going to the right.

- If you are not sure, make your best guess. Always try to be as accurate and as speedy as you can.

- We'll do 3 or 4 blocks of about 150 trials each, to get you practiced on the cylinder task.

Figure B.2.: Example instruction sheet for psychophysics study reward task (first page):

white = 1 point / black = 8 points

Rotating Cylinder, with Points

The cylinder is shown in the centre of the screen. The dots on the front surface of the cylinder are moving either to the left or the right:



front going **LEFT (or clockwise)**



front going **RIGHT (or counterclockwise)**

Task:

- Always look at the dot in the centre of the screen.

- A cylinder will appear. Click the left mouse button if the dots on the front surface are going left. Click the right mouse button if the dots on the front surface are going right.

- You will earn points if you click the correct mouse button.

- Before each cylinder, a points cue will appear on the screen to show you how many points you can get for a **correct** answer:

 1 point for a correct click, in either direction.

 1 point for a correct left click; 8 points for a correct right click.

 8 points for a correct left click; 1 point for a correct right click.

 8 points for a correct click, in either direction.

Aim: earn as many points as you can! Points are converted to money at the end of the experiment (250 points ~ £1.00).

Figure B.3.: Example instruction sheet for psychophysics study reward task (second page):

- Use the reward cue to help you get as many points as you can on each trial.

- After you've clicked a button, the cylinder will disappear and the number of points you earned will be shown on the screen.

- Each block has about 150-200 trials. Every 20 trials, an update will appear, telling you how many points you've earned in the block so far.

- The entire experiment consists of about 15 blocks, done over 3 or 4 sessions.

Figure B.4.: Example instruction sheet for MEG scan reward task (first page):

white = 8 points / black = 1 point

Rotating Cylinder, with Points

The cylinder is shown in the centre of the screen. The dots on the front surface of the cylinder are moving either to the left or the right:



front going **LEFT (or clockwise)**



front going **RIGHT (or counterclockwise)**

TASK

- **Always look at the dot in the centre of the screen.**

- A cylinder will appear. Lift your left thumb from the response button if the dots on the front surface are going left. Lift your right thumb from the response button if the dots on the front surface are going right.

- (Keep your thumbs covering the response buttons at other times.)

- You will earn points if you give the correct response.

- Before each cylinder, a **points cue** will appear on the screen to show you how many points you can get for a **correct** answer:

 **5 points** for a correct response, in either direction.

 **8 points** for a correct left response; **1 point** for a correct right response.

 **1 point** for a correct left response; **8 points** for a correct right response.

Figure B.5.: Example instruction sheet for MEG scan reward task (second page):

Aim: earn as many points as you can! Points are converted to money at the end of the experiment (100 points = £1.00).

- Use the **points cue** to help you get as many points as you can on each trial.

- After you've clicked a mouse button, the cylinder will disappear and the number of points you earned will be shown on the screen.

- Each block has **100 trials**. Every 20 trials, an update will appear, telling you the number of points you have earned in the block so far.

- The entire experiment consists of 8 blocks, and takes about 1 hour 30 mins in the MEG scanner.

Special Instructions For MEG Scan:

- **Metal**: Please make sure all metal is removed before entering the MEG scanner room.

- **Eye-blinks:** Please try to avoid blinking during the cylinder, or whilst giving the responses, as blinks interfere with the signal measurements. The best time to blink is after the number of points has been shown at the end of the trial.

Bibliography

ABBOTT, L.F. & CHANCE, F.S. (2005). Drivers and modulators from push-pull and balanced synaptic input. *Progress in Brain Research*, 149:147–155.

ABI-DARGHAM, A., RODENHISER, J., PRINTZ, D., ZEA-PONCE, Y., GIL, R., KEGELES, L.S., WEISS, R., COOPER, T.B., MANN, J.J., VAN HEERTUM, R.L., GORMAN, J.M. & LARUELLE, M. (2000). Increased baseline occupancy of D2 receptors by dopamine in schizophrenia. *Proceedings of the National Academy of Sciences*, 97(14):8104–8109.

AFRAZ, S.R., KIANI, R. & ESTEKY, H. (2006). Microstimulation of inferotemporal cortex influences face catergorization. *Nature*, 442(10):692–695.

ALBRIGHT, T.D. & DESIMONE, R. (1987). Local precision of visuotopic organization in the middle temporal area (MT) of the macaque. *Experimental Brain Research*, 65:582–592.

ARNAL, L.H., WYART, V. & GIRAUD, A.L. (2011). Transitions in neural oscillations reflect prediction errors generated in audiovisual speech. *Nature Neuroscience*, 14(6):797–803.

ARON, A.R., ROBBINS, T.W. & POLDRACK, R.A. (2004). Inhibition and the right inferior frontal cortex. *TRENDS in Cognitive Sciences*, 8(4):170–177.

BAINES, S., RUZ, M., RAO, A., DENISON, R. & NOBRE, A.C. (2011). Modulation of neural activity by motivational and spatial biases. *Neuropsychologia*, 49:2489–2497.

BALCI, F., SIMEN, P., NIYOGI, R., SAXE, A., HUGHES, J.A., HOLMES, P.J. & COHEN, J.D. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics*, 73:640–657.

BAR-GAD, I. & BERGMAN, H. (2001). Stepping out of the box: information processing in the neural networks of the basal ganglia. *Current Opinion in Neurobiology*, 11(6):689–695.

BARTON, J.J.S., SHARPE, J.A. & RAYMOND, J.E. (1996). Directional deficits in pursuit and motion perception in humans with unilateral cortical lesions. *Brain*, 119:1535–1550.

BELOVA, M.A., PATON, J.J., MORRISON, S.E. & SALZMAN, C.D. (2007). Expectation modulates

neural responses to pleasant and aversive stimuli in primate amygdala. *Neuron*, 55(6):970–984.

BERGER, J. (1972). *Ways of Seeing*. London: British Broadcasting Corporation and Penguin Books.

BICHOT, N.P., ROSSI, A.F. & DESIMONE, R. (2005). Parallel and serial neural mechanisms for visual search in macaque area v4. *Science*, 308:529–534.

BOGACZ, R., BROWN, E., MOEHLIS, J., HOLMES, P.J. & COHEN, J.D. (2006). The physics of optimal decision making: a formal analysis of models of performance in two-alternative forced-choice tasks. *Psychological Review*, 113(4):700–765.

BOGACZ, R., HU, P.T., HOLMES, P.J. & COHEN, J.D. (2010). Do humans produce the speed-accuracy tradeoff that maximizes reward rate? *Quarterly Journal of Experimental Psychology (Colchester)*, 63(5):863–891.

BOLY, M., GARRIDO, M.I., GOSSERIES, O., BRUNO, M.A., BOVEROUX, P., SCHNAKERS, C., MASSIMINI, M., LITVAK, V., LAUREYS, S. & FRISTON, K.J. (2011). Preserved feedforward but impaired top-down processes in the vegetative state. *Science*, 332:858–862.

BOYNTON, G.M. (2009). A framework for describing the effects of attention on visual responses. *Vision Research*, 49:1129–1143.

BRADLEY, D.C., CHANG, G.C. & ANDERSEN, R.A. (1998). Encoding of three-dimensional structure-from-motion by primate area MT neurons. *Nature*, 392(16):714–717.

BRADLEY, D.C., QIAN, N. & ANDERSEN, R.A. (1995). Integration of motion and stereopsis in middle temporal cortical area of macaques. *Nature*, 373(16):609–611.

BRETT, M., PENNY, W. & KIEBEL, S.J. (2003). An introduction to Random Field Theory. In R.S.J. Frackowiak, editor, *Human brain function*. Elsevier Science (USA).

BRIDGE, H. & CLARE, S. (2006). High-resolution MRI: in vivo histology? *Philosophical Transactions of the Royal Society*, 361:137–146.

BRIDGE, H., CLARE, S., JENKINSON, M., JEZZARD, P., PARKER, A.J. & MATTHEWS, P.M. (2005). Independent anatomical and functional measures of the V1/V2 boundary in human visual cortex. *Journal of Vision*, 5:93–102.

BRITTEN, K.H., NEWSOME, W.T., SHADLEN, M.N., CELEBRINI, S. & MOVSHON, J.A. (1996). A relationship between behavioural choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*, 13:87–100.

BRITTEN, K.H., SHADLEN, M.N., NEWSOME, W.T. & MOVSHON, J.A. (1992). The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of Neuroscience*, 12(12):4745–4765.

BRITTEN, K.H., SHADLEN, M.N., NEWSOME, W.T. & MOVSHON, J.A. (1993). Responses of neurons in macaque MT to stochastic motion signals. *Visual Neuroscience*, 10:1157–1169.

BROOKES, M.J., ZUMER, J.M., STEVENSON, C.M., HALE, J.R., BARNES, G.R., VRBA, J. & MORRIS, P.G. (2010). Investigating spatial specificity and data averaging in MEG. *NeuroImage*, 49:525–538.

BROUWER, G.J. & VAN EE, R. (2007). Visual cortex allows prediction of perceptual states during ambiguous structure-from-motion. *The Journal of Neuroscience*, 27(5):1015–1023.

BUSEMEYER, J.R. & TOWNSEND, J.T. (1993). Decision field theory: a dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100:432–459.

CALLICOTT, J.H., BERTOLINO, A., MATTAY, V.S., LANGHEIM, F.J., DUYN, J., COPPOLA, R., GOLDBERG, T.E. & WEINBERGER, D.R. (2000). Physiological dysfunction of the dorsolateral prefrontal cortex in schizophrenia revisited. *Cerebral Cortex*, 10(11):1078–1092.

CALLICOTT, J.H., MATTAY, V.S., VERCHINISKI, B.A., MARENCO, S., EGAN, M.F. & WEINBERGER, D.R. (2003). Complexity of prefrontal cortical dysfunction in schizophrenia: more than up or down. *American Journal of Psychiatry*, 160(12):2209–2215.

CAVANNA, A.E. & TRIMBLE, M.R. (2006). The precuneus: a review of its functional anatomy and behavioural correlates. *Brain*, 129:564–583.

CELEBRINI, S. & NEWSOME, W.T. (1994). Neuronal and psychophysical sensitivity to motion signals in extrastriate area MST of the macaque monkey. *The Journal of Neuroscience*, 14(7):4109–4124.

CELEBRINI, S. & NEWSOME, W.T. (1995). Microstimulation of extrastriate area MST influences performance on a direction discrimination task. *Journal of Neurophysiology*, 73(2):437–448.

CHANCE, F.S., ABBOTT, L.F. & REYES, A.D. (2002). Gain modulation from background synaptic input. *Neuron*, 35:773–782.

CHANDRA, M. (2001). *Statistical Quality Control*. New York: CRC Press.

CHEN, W.T., KO, Y.C., LIAO, K.K., HSIEH, J.C., YEH, T.C., WU, Z.A., HO, L.T. & LIN, Y.Y. (2005). Optimal check size and reversal rate to elicit pattern reversal MEG responses. *Canadian Journal of Neurologial Science*, 32:218–224.

CHURCHLAND, A.K., KIANI, R., CHAUDHURI, R., WANG, X.J., POUGET, A. & SHADLEN, M.N. (2011). Variance as a signature of neural computations during decision making. *Neuron*, 69:818–831.

CISEK, P., PUSKAS, G.A. & EL-MURR, S. (2009). Decisions in changing conditions: the urgency-gating model. *The Journal of Neuroscience*, 29:11560–11571.

COHEN, M.R. & MAUNSELL, J.H.R. (2011). Using neuronal populations to study the mechanisms underlying spatial and feature attention. *Neuron*, 70:1192–1204.

COLBY, C.L. & GOLDBERG, M.E. (1999). Space and attention in parietal cortex. *Annual Reviews of Neuroscience*, 22:319–349.

CORBETTA, M. & SHULMAN, G.L. (2002). Control of goal-directed and stimulus-driven attention in the brain. *Nature Reviews Neuroscience*, 3:201–215.

COTTEREAU, B., LORENCEAU, J., GRAMFORT, A., CLERC, M., THIRION, B. & BAILLET, S. (2011). Phase delays within visual cortex shape the response to steady-state visual stimulation. *NeuroImage*, 54:1919–1929.

DALE, A.M., FISCHL, B.R. & SERENO, M.I. (1999). Cortical surface-based analysis I: Segmentation and surface reconstruction. *NeuroImage*, 9:179–194.

DALE, A.M., LIU, A.K., FISCHL, B.R., BUCKNER, R.L., BELLIVEAU, J.W., LEWINE, J.D. & HALGREN, E. (2000). Dynamic statistical parametric mapping: combining fMRI and MEG for high-resolution imaging of cortical activity. *Neuron*, 26:55–67.

DALE, A.M. & SERENO, M.I. (1993). Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction. *Journal of Cognitive Neuroscience*, 5:162–176.

DASSONVILLE, P., ZHU, X.H., UGURBIL, K., KIM, S.G. & ASHE, J. (1997). Functional activation in motor cortex reflects the direction and degree of handedness. *Proceedings of the National Academy of Sciences USA*, 94:14015–14018.

DAVID, O., COSMELLI, D. & FRISTON, K.J. (2004). Evaluation of different measures of functional connectivity using a neural mass model. *NeuroImage*, 21:659–673.

DEANGELIS, G.C., CUMMING, B.G. & NEWSOME, W.T. (1998). Cortical area MT and the perception of stereoscopic depth. *Nature*, 394(13):677–670.

DEANGELIS, G.C. & NEWSOME, W.T. (1999). Organization of Disparity-Selective Neurons in Macaque Area MT. *The Journal of Neuroscience*, 19(4):1398–1415.

DEANGELIS, G.C. & NEWSOME, W.T. (2007). Perceptual 'read-out' of conjoined direction and disparity maps in extrastriate area MT. *PLoS Biology*, 2(3):394–404.

DESIMONE, R. & UNGERLEIDER, L.G. (1986). Multiple visual areas in the caudal superior temporal sulcus of the macaque. *The Journal of Comparative Neurology*, 248:164–189.

DEYOE, E.A., BANDETTINI, P., NEITZ, J., MILLER, D. & WINANS, P. (1994). Functional magnetic resonance imaging (FMRI) of the human brain. *Journal of Neuroscience Methods*, 54:171–187.

DIEDERICH, A. (2008). A further test of sequential-sampling models that account for payoff effects on response bias in perceptal decision tasks. *Perception & Psychophysics*, 70(2):229–256.

DIEDERICH, A. & BUSEMEYER, J.R. (2006). Modeling the effects of payoff on response bias in a perceptual discrimination task: bound-change, drift-rate-change, or two-stage-processing hypothesis. *Perception & Psychophysics*, 68(2):194–207.

DODD, J.V., KRUG, K., CUMMING, B.G. & PARKER, A.J. (2001). Perceptually Bistable Three-Dimensional Figures Evoke High Choice Probabilities in Cortical Area MT. *The Journal of Neuroscience*, 21(13):4809–4821.

DONNER, T.H., SIEGEL, M., FRIES, P. & ENGEL, A.K. (2009). Buildup of choice-predictive activity in human motor cortex during perceptual decision-making. *Current Biology*, 19:1481–1585.

ECKHOFF, P., HOLMES, P.J., LAW, C., CONNOLLY, P.M. & GOLD, J.I. (2008). On diffusion processes with variable drift rates as models for decision making during learning. *New Journal of Physics*, 10:1367–2630.

EICKHOFF, S.B., PAUS, T., CASPERS, S., GROSBRAS, M.H., EVANS, A.C., ZILLES, K. & AMUNTS, K. (2007). Assignment of functional activations to probabilistic cytoarchitectonic areas revisited. *NeuroImage*, 36(3):511–521.

EICKHOFF, S.B., STEPHAN, K.E., MOHLBERG, H., GREFKES, C., FINK, G.R., AMUNTS, K. & ZILLES, K. (2005). A new SPM toolbox for combining probabilistic cytoarchitectonic maps and functional imaging data. *NeuroImage*, 25:1325–1335.

ENGEL, S.A., GLOVER, G.H. & WANDELL, B.A. (1997). Retinotopic organization in human visual cortex and the spatial precision of functional MRI. *Cerebral Cortex*, 7:181–192.

ENGEL, S.A., RUMELHART, D.E., WANDELL, B.A., LEE, A.T., GLOVER, G.H., CHICHILNISKY, E.J. & SHADLEN, M.N. (1994). fMRI of human visual cortex. *Nature*, 369:525.

FENG, S., HOLMES, P.J., RORIE, A.E. & NEWSOME, W.T. (2009). Can monkeys choose optimally when faced with noisy stimuli and unequal rewards? *PLoS Computational Biology*, 5(2):1–15.

FISCHL, B.R. (2012). Freesurfer. In press.

FISCHL, B.R., SERENO, M.I. & DALE, A.M. (1999). Cortical surface-based analysis II: Inflation, flattening and a surface-based coordinate system. *NeuroImage*, 9:195–207.

FLEMING, S.M., ADN OLIVER J. HULME, L.W., SAHANI, M. & DOLAN, R.J. (2010). Effects of category-specific costs on neural systems for perceptual decision-making. *Journal of Neurophysiology*, 103:3238–3247.

FRISTON, K.J. (2003). Introduction: experimental design and statistical parametric mapping. In R.S.J. Frackowiak, editor, *Human brain function*. Elsevier Science (USA).

FRISTON, K.J. (2008). Hierarchical models in the brain. *PLoS Computational Biology*, 4(11):1–24.

FRISTON, K.J. (2009). *SPM8 Manual*. The FIL Methods Group (and honorary members), Institute of Neurology, University College London, 2009 edition.

FRISTON, K.J. & DOLAN, R.J. (2010). Computational and dynamic models in neuroimaging. *NeuroImage*, 52:752–765.

FRISTON, K.J., HARRISON, L., DAUNIZEAU, J., KIEBEL, S.J., PHILLIPS, C., TRUJILLO-BARRETO, N., HENSON, R.N., FLANDIN, G. & MATTOUT, J. (2008). Multiple sparse priors for the M/EEG inverse problem. *NeuroImage*, 39:1104–1120.

FRISTON, K.J., MATTOUT, J., TRUJILLO-BARRETO, N., ASHBURNER, J. & PENNY, W. (2006). Variational free energy and the Laplace approximation. *NeuroImage*, 34:220–234.

FRITSCH, G. & HITZIG, E. (1870). On the electrical excitability of the cerebrum. *Arch. f. Anat., Physiol. und wissenschaftl. Mediz.*, 37:300–32.

GAO, J., TORTELL, R. & MCCLELLAND, J.L. (2011). Dynamic integration of reward and stimulus information in perceptual decision-making. *PLoS One*, 6(3):e16749.

GARRIDO, M.I., KILNER, J., KIEBEL, S.J. & FRISTON, K.J. (2007). Evoked brain responses are generated by feedback loops. *Proceedings of the National Academy of Sciences*, 104(52):20961–20966.

GOLD, J.I. & SHADLEN, M.N. (2000). Representation of a perceptual decision in developing oculomotor commands. *Nature*, 404(23):390–394.

GOLD, J.I. & SHADLEN, M.N. (2001). Neural computations that underlie decisions about sensory stimuli. *TRENDS in Cognitive Sciences*, 5(1):10–16.

GOLD, J.I. & SHADLEN, M.N. (2002). Banburismus and the brain: decoding the relationship between sensory stimuli, decisions and reward. *Neuron*, 36:299–308.

GOLD, J.I. & SHADLEN, M.N. (2007). The neural basis of decision making. *Annual Reviews of Neuroscience*, 30:535–574.

GOODALE, M.A. & MILNER, A.D. (1992). Separate visual pathways for perception and action. *TRENDS in Neurosciences*, 15(1):20–25.

GOTTLIEB, J. (2007). From thought to action: the parietal cortex as a bridge between perception, action and cognition. *Neuron*, 53:9–16.

GOTTLIEB, J. & BALAN, P. (2010). Attention as a decision in information space. *TRENDS in Cognitive Sciences*, 14:240–248.

GREEN, D.M. & SWETS, J.A. (1966). *Signal Detection Theory and Psychophysics*. New York: Wiley.

GREGORY, R. (1997). Knowledge in perception and illusion. *Philosophical Transactions of the Royal Society*, 352:1121–1128.

GROSSBERG, S. & PILLY, P.K. (2008). Temporal dynamics of decision-making during motion perception in the visual cortex. *Vision Research*, 48:1345–1373.

GRUNBAUM, A.S.F. & SHERRINGTON, C.S. (1902). Observations on the physiology of the cerebral cortex of some of the higher apes (preliminary communication). *Proceedings of the Royal Society*, 69:206–9.

HAGLER, D.J., HALGREN, E., MARTINEZ, A., HUANG, M., HILLYARD, S.A. & DALE, A.M. (2009). Source estimates for MEG/EEG visual evoked responses constrained by multiple, retinotopically-mapped stimulus locations. *Human Brain Mapping*, 30(4):1290–1309.

HAMALAINEN, M., HARI, R., ILMOIEMI, R.J., KNUUTILA, J. & LOUNASMAA, O.V. (1993). Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, 65:413–503.

HAMALAINEN, M. & ILMOIEMI, R.J. (1994). Interpreting magnetic fields of the brain: minimum norm estimates. *Medical and Biological Engineering and Computing*, 32:35–42.

HANES, D.P. & SCHALL, J.D. (1996). Neural control of voluntary movement initiation. *Science*, 274:427–430.

HANKS, T.D., DITTERICH, J. & SHADLEN, M.N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, 9(5):682–689.

HANSEN, P.C., KRINGELBACH, M.L. & HARI, R. (2010). *MEG: An introduction to methods*. Oxford University Press, New York.

HEEKEREN, H.R., MARRETT, S., BANDETTI, P.A. & UNGERLEIDER, L.G. (2004). A general mechanism for perceptual decision-making in the human brain. *Nature*, 431(14):859–862.

HEEKEREN, H.R., MARRETT, S. & UNGERLEIDER, L.G. (2008). The neural systems that mediate human perceptual decision-making. *Nature Reviews Neuroscience*, 9:467–479.

HENSON, R.N., MATTOUT, J., PHILIPS, C. & FRISTON, K.J. (2009). Selecting forward models for MEG source reconstruction using model-evidence. *NeuroImage*, 46:168–176.

HERRNSTEIN, R.J. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behaviour*, 4:267–72.

HILLEBRAND, A. & BARNES, G.R. (2011). Practical constrains on estimation of source extent with MEG beamformers. *NeuroImage*, 54:2732–2740.

HOLMES, G. (1918). Disturbance of vision by cerebral lesions. *British Journal of Ophthalmology*, 2:353–384.

HOLMES, G. (1944). The organization of the visual cortex in man. *Proceedings of the Royal Society*, 132:348–361.

HORTON, J.C. & HOYT, W.F. (1991). Quadrantic visual field defects: a hallmark of lesions in extrastriate (V2/V3) cortex. *Brain*, 114:1703–1718.

HOUWELING, A.R. & BRECHT, M. (2008). Behavioural report of single neuron stimulation in somatosensory cortex. *Nature*, 451(3):65–69.

HUI, H.B., PANTAZIS, D., BRESSLER, S.L. & LEAHY, R.M. (2010). Identifying true cortical interactions in MEG using the nulling beamformer. *NeuroImage*, 49:3161–3174.

HUK, A.C., DOUGHERTY, R.F. & HEEGER, D.J. (2002). Retinotopy and functional subdivision of human areas MT and MST. *The Journal of Neuroscience*, 22(16):7195–7205.

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

HUNT, L., KOLLING, N., SOLTANI, A., WOOLRICH, M., RUSHWORTH, M.F.S. & BEHRENS, T.E.J. (2012). Mechanisms underlying cortical activity during value-guided choice. *Nature Neuroscience*, 15:470–476.

HURVICH, C.M. & TSAI, C.L. (1989). Regression and time series model selection in small samples. *Biometrika*, 76:297–307.

JIANG, Y., BOEHLER, C.N., NONNIG, N., DUZEL, E., HOPF, J.M., HEINZE, H.J. & SCHOENFELD,

M.A. (2008). Binding 3-D object perception in the human visual cortex. *Journal of Cognitive Neuroscience*, 20(4):553–562.

JOHANSSON, R.S. & VALLBO, A.B. (1979). Detection of tactile stimuli thresholds of afferent units related to psychophysical thresholds in the human hand. *Journal of Physiology*, 297:405–422.

KARLIN, S. & TAYLOR, H.M. (1975). *A first course in stochastic processes*. New York: Academic Press.

KENNERLEY, S.W., WALTON, M.E., BEHRENS, T.E.J., BUCKLEY, M.J. & RUSHWORTH, M.F.S. (2006). Optimal decision making and the anterior cingulate cortex. *Nature Neuroscience*, 9(7):940–947.

KIANI, R., ESTEKY, H., MIRPOUR, K. & TANAKA, K. (2007). Object category structure in response patterns of neuronal population in monkey inferior temporal cortex. *Journal of Neurophysiology*, 97:4296–4309.

KIANI, R., HANKS, T.D. & SHADLEN, M.N. (2008). Bounded integration in parietal cortex underlies decisions even when viewing duration is dictated by the environment. *The Journal of Neuroscience*, 28(12):3017–3029.

KIEBEL, S.J., GARRIDO, M.I., MORAN, R., CHEN, C.C. & FRISTON, K.J. (2009). Dynamic Causal Modelling for EEG and MEG. *Human Brain Mapping*, 30:1866–1876.

KNILL, D.C. & RICHARDS, W., editors (1996). *Perception as Bayesian Inference*. Cambridge University Press, 2 edition.

KOUNEIHER, F., CHARRON, S. & KOECHLIN, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience*, 12(7):939–947.

KRUG, K., CUMMING, B.G. & PARKER, A.J. (2004). Comparing perceptual signals of single V5/MT neurons in two binocular depth tasks. *Journal of Neurophysiology*, 92:1586–1596.

LANGERS, D.R.M. & VAN DIJK, P. (2011). Mapping the tonotopic organization in human auditory cortex with minimally salient acoustic stimulation. *Cerebral Cortex*, 34(1):264–73.

LARUELLE, M., ABI-DARGHAM, A., GIL, R., KEGELES, L.S. & INNIS, R. (1999). Increased dopamine transmission in schizophrenia: relationship to illness phases. *Biological Psychiatry*, 46(1):56–72.

LAUWERS, K., SAUNDERS, R., VOGELS, R., VANDENBUSSCHE, E. & ORBAN, G.A. (2000). Impairment in motion discrimination tasks is unrelated to amount of damage to superior temporal sulcus in motion areas. *The Journal of Comparative Neurology*, 420(4):539–557.

LIAO, C.H., WORSLEY, K.J., POLINE, J.B., ASTON, J.A.D., DUNCAN, G.H. & EVANS, A.C. (2002). Estimating the delay of the fMRI response. *NeuroImage*, 16:593–606.

LIN, F.H., WITZEL, T., AHLFORS, S.P., STUFFLEBEAM, S.M., BELLIVEAU, J.W. & HAMALAINEN, M. (2006). Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates. *NeuroImage*, 31:160–171.

LINK, S.W. (1992). *The wave theory of difference and similarity*. Hillsdale, NJ:Erlbaum.

LISTON, D.B. & STONE, L.S. (2008). Effects of prior information and reward on oculomotor and perceptual choices. *The Journal of Neuroscience*, 28(51):13866–13875.

LITVAK, V., MATTOUT, J., KIEBEL, S.J., PHILLIPS, C., HENSON, R., KILNER, J., BARNES, G.R., OOSTENVALD, R., DAUNIZEAU, J., FLANDIN, G., PENNY, W. & FRISTON, K.J. (2011). EEG and MEG data analysis in SPM8. *Computational Intelligence and Neuroscience*, 2011(852961):1–32.

LO, C.C. & WANG, X.J. (2006). Cortico-basal ganglia circuit mechanism for a decision threshold in reaction time tasks. *Nature Neuroscience*, 9(7):956–963.

LOGOTHETIS, N.K., PAULS, J., AUGATH, M., TRINATH, T. & OELTERMANN, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412:150–158.

LOGOTHETIS, N.K. & WANDELL, B.A. (2004). Interpreting the BOLD signal. *Annual Reviews of Physiology*, 66:735–769.

LOHMANN, G., ERFURTH, K., MULLER, K. & TURNER, R. (2012). Critical comments on dynamic causal modelling. *NeuroImage*, 59:2322–2329.

LYNN, C. & CURRAN, W. (2010). A comparison of monkey and human motion processing mechanisms. *Vision Research*, 50:2137–2141.

MADDOX, W.T. & BOHIL, C.J. (2003). A theorical framework for understanding the simulatneous base-rate and payoff manipulations on decision criterion learning in perceptual categorization. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 29:307–320.

MALONEY, L.T. & WANDELL, B.A. (1984). A model of a single visual channel's response to weak test lights. *Vision Research*, 24:633–640.

MANTINI, D., HASSON, U., BETTI, V., PERRUCCI, M.G., ROMANI, G.L., CORBETTA, M., ORBAN, G.A. & VANDUFFEL, W. (2012). Interspecies activity correlations reveal functional correspondance between monkey and human brain areas. *Nature Methods*, 9(3):276–283.

MARGULIES, D.S., VINCENT, J.L., KELLY, C., LOHMANN, G., UDDIN, L.Q., BISWAL, B.B., VILL-

RINGER, A., CASTELLANOS, F.X., MILHAM, M.P. & PETRIDES, M. (2009). Precuneus shares intrinsic functional architecture in humans and monkeys. *Proceedings of the National Academy of Sciences*, 106(47):20069–20074.

MATTOUT, J., HENSON, R. & FRISTON, K.J. (2007). Canonical source reconstruction for MEG. *Computational Intelligence and Neuroscience*, 2007(67613).

MAUNSELL, J.H.R. (2004). Neuronal representations of cognitive state: reward or attention? *TRENDS in Cognitive Sciences*, 8(6):261–265.

MAUNSELL, J.H.R. & TREUE, S. (2006). Feature-based attentiona in visual cortex. *TRENDS in Neurosciences*, 29(6):317–322.

MAUNSELL, J.H.R. & VAN ESSEN, D.C. (1983). Functional Properties of Neurons in Middle Temporal Visual Area of the Macaque Monkey. II. Binocular Interactions and Sensitivity to Binocular Disparity. *Journal of Neurophysiology*, 49(5):1148–1167.

MAZUREK, M.E., ROITMAN, J.E., DITTERICH, J. & SHADLEN, M.N. (2003). A role for neural integrators in perceptual decision making. *Cerebral Cortex*, 13:1257–1269.

MCADAMS, C.J. & MAUNSELL, J.H.R. (2000). Attention to both space and feature modulates neuronal responses in macaque area V4. *Journal of Neurophysiology*, 83:1751–1755.

MENON, R.S., GATI, J.S., GOODYEAR, B.G., LUKNOWSKY, D.C. & THOMAS, C.G. (1998). Spatial and temporal resolution of functional magnetic resonance imaging. *The International Journal of Biochemistry & Cell Biology*, 76(2-3):560–71.

MILOSAVLJEVIC, M., MALMAUD, J., HUTH, A., KOCH, C. & RANGEL, A. (2010). The drift diffusion model can account for the accuracy and rection time of value-based choices under high and low time pressure. *Judgement and Decision Making*, 5(6):437–449.

MININI, L., PARKER, A.J. & BRIDGE, H. (2010). Neural modulation by binocular disparity greatest in human dorsal visual stream. *Journal of Neurophysiology*, 104:169–178.

MORADI, F., LIU, L.C., CHENG, K., WAGGONER, R.A., TANAKA, K. & IONNIDES, A.A. (2003). Consistant and precise localization of brain activity in human primary visual cortex by MEG and fMRI. *NeuroImage*, 18:595–609.

MORRISON, S.E. & SALZMAN, C.D. (2010). Re-valuing the amygdala. *Current Opinion in Neurobiology*, 20(2):221–230.

MORRISON, S.E. & SALZMAN, C.D. (2011). Representations of appetitive and aversive information in the primate orbitofrontal cortex. *Annals of the New York Academy of Sciences*, 1239:59–70.

MULDER, M.J., WAGENMAKERS, E.J., RATCLIFF, R., BOEKEL, W. & FORSTMANN, B.U. (2012). Bias in the brain: A diffusion model analysis of prior probability and potential payoff. *The Journal of Neuroscience*, 32(7):2335–2343.

MURAKIAMI, S. & OKADA, Y.C. (2006). Contributions of principal neocortical neurons to magnetoencephalography and electoencephalography signals. *Journal of Physiology*, 575(3):925–936.

NAKAMURA, A., KAKIGI, R., HOSHIYAMA, M., KOYAMA, S., KITAMURA, Y. & SHIMOJO, M. (1997). Visual evoked cortical magnetic fields to pattern reversal stimulation. *Cognitive Brain Research*, 6:9–22.

NAWROT, M. & BLAKE, R. (1993). On the perceptual identity of dynamic stereopsis and kinetic depth. *Vision Research*.

NERI, P., BRIDGE, H. & HEEGER, D.J. (2004). Stereoscopic processing of absolute and relative disparity in human visual cortex. *Journal of Neurophysiology*, 92:1880–1891.

NEWSOME, W.T., BRITTEN, K.H. & MOVSHON, J.A. (1989a). Neuronal correlates of a perceptual decision. *Nature*, 341(6237):52–54.

NEWSOME, W.T., BRITTEN, K.H., MOVSHON, J.A. & SHADLEN, M.N. (1989b). Single neurons and the perception of visual motion. In D.K. Lam & D. Gilbert, editors, *Neural Mechanisms of of Visual Perception*. Portfolio Publishing, The Woodlands, Texas.

NIENBORG, H. & CUMMING, B.G. (2006). Macaque V2 neurons, but not V1 neurons, show choice-related activity. *The Journal of Neuroscience*, 26(37):9567–9578.

NIENBORG, H. & CUMMING, B.G. (2009). Decision-related activity in sensory neurons reflects more than a neuron's causal effect. *Nature*, 459:89–92.

NISHIYAMA, T., OHDE, H., HARUTA, Y., MASHIMA, Y. & OGUCHI, Y. (2004). Multifocal magnetoencephalogram applied to objective visual field analysis. *Japanese Journal of Opthalmology*, 48:115–122.

OKADA, Y.C., WU, J. & KYUHOU, S. (1997). Genesis of MEG signals in a mammalian CNS structure. *Electroencephalography and clinical Neurophysiology*, 103:474–485.

OOSTENVALD, R., FRIES, P., MARIS, E. & SCHOFFELEN, J.M. (2011). FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. *Computational Intelligence and Neuroscience*, 2011(156869):1–9.

PACK, C.C., BEREZOVSKII, V.K. & BORN, R.T. (2001). Dynamic properties of neurons in cortical area MT in alert and anaesthetized macaque monkeys. *Nature*, 414:905–908.

PADOA-SCHIOPPA, C. & ASSAD, J.A. (2008). The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nature Neuroscience*, 11(1):95–102.

PALMER, J., HUK, A.C. & SHADLEN, M.N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, 5:376–404.

PANTOJA, J., RIBEIRO, S., WIEST, M., SOARES, E., GERVASONI, D., LEMOS, N.A.M. & NICOLELIS, M.A.L. (2007). Neuronal activity in the primary somatosensory thalamocortical loop is modulated by reward contingency during tactile discrimination. *The Journal of Neuroscience*, 27(39):10608–10620.

PARKER, A.J. & NEWSOME, W.T. (1998). Sense and the single neuron: Probing the physiology of perception. *Annual Reviews of Neuroscience*, 21:227–77.

PASTERNAK, T. & MERIGAN, W.H. (1994). Motion perception following lesions of the superior temporal sulcus in the monkey. *Cerebral Cortex*, 4:247–259.

PATON, J.J., BELOVA, M.A., MORRISON, S.E. & SALZMAN, C.D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature*, 439(7078):865–870.

PENFIELD, W. & BOLDREY, E. (1937). Somatic motor and sensory representation in the cerebral cortex of man as studied by electrical microstimulation. *Brain*, 60:389–443.

PENNY, W. & HOLMES, A.J. (2003). Random-effects analysis. In R.S.J. Frackowiak, editor, *Human brain function*. Elsevier Science (USA).

PERRY, G., ADJAMIN, P., THAI, N.J., HOLLIDAY, I., HILLEBRAND, A. & BARNES, G.R. (2011). Retinotopic mapping of the primary visual cortex - a challenge for MEG imaging of the human cortex. *European Journal of Neuroscience*, 34:652–661.

PLATT, M.L. & GLIMCHER, P.W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400:233–238.

PLEGER, B., BLANKENBURG, F., RUFF, C.C., DRIVER, J. & DOLAN, R.J. (2008). Reward facilitates tactile judgements and modulates haemodynamic responses in human primary somatosensory cortex. *The Journal of Neuroscience*, 28(33):8161–8168.

PLEGER, B., RUFF, C.C., BLANKENBURG, F., KLOPPEL, S., DRIVER, J. & DOLAN, R.J. (2009). Influence of dopaminergically mediated reward on somatosensory decision-making. *PLoS Biology*, 7(7):e1000164–e100173.

POGHOSYAN, V. & IONNIDES, A.A. (2007). Precise mapping of early visual responses in space and time. *NeuroImage*, 35:759–770.

RATCLIFF, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2):59–108.

RATCLIFF, R. (1979). Group reaction time distributions and an analysis of distribution statistics. *Psychological Bulletin*, 86:446–461.

RATCLIFF, R. (1993). Methods for dealing with reaction time outliers. *Psychonomic Bulletin*, 114(3):510–532.

RATCLIFF, R. & MCKOON, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, 20:873–922.

RATCLIFF, R. & SMITH, P.L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111:333–367.

RATCLIFF, R., THAPAR, A., SMITH, P.L. & MCKOON, G. (2005). Aging and response times: a comparison of sequential sampling models. In J. Duncan, P. McLeod & L. Phillips, editors, *Speed, control and age*. New York: Oxford University Press.

RATCLIFF, R. & TUERLINCKX, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, 9(3):438–481.

REDDI, B.A., N., A.K. & CARPENTER, R.H. (2003). Accuracy, information, and response time in a saccadic decision task. *Journal of Neurophysiology*, 90:3538–3546.

REYNOLDS, J.H. & HEEGER, D.J. (2009). The normalization model of attention. *Neuron*, 61(2):168–185.

ROBERTS, S. (1959). Control chart tests based on geomatric moving averages. *Technometrics*, 1:239–250.

ROBINSON, D.A. (1972). Eye movements evoked by collicular stimulation in the alert monkey. *Vision Research*, 12:1795–1808.

ROBINSON, D.A. & FUCHS, A.F. (1969). Eye movements evoked by stimulation of frontal eye fields. *Journal of Neurophysiology*, 32(5):637–648.

ROBINSON, S.E. & VRBA, J. (1999). *Functional neuroimaging by synthetic aperture magnetometry (SAM)*, pages 302–305. Tohoku University Press, Sendai.

ROITMAN, J.D. & SHADLEN, M.N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *The Journal of Neuroscience*, 22(21):9475–9489.

ROLLS, E.T. & GRABENHORST, F. (2008). The orbitofrontal cortex and beyond: from affect to decision-making. *Progress in Neurobiology*, 86:216–244.

ROMO, R., HERNANDEZ, A., ZAINOS, A., BRODY, C.D. & LEMUS, L. (2000). Sensing without touching: psychophysical performance based on cortical microstimulation. *Neuron*, 26:273–278.

ROMO, R., HERNANDEZ, A., ZAINOS, A., LEMUS, L. & BRODY, C.D. (2002). Neuronal correlates of decision-making in secondary somatosensory cortex. *Nature Neuroscience*, 5:1217–1225.

RORIE, A.E., GAO, J., MCCLELLAND, J.L. & NEWSOME, W.T. (2010). Integration of Sensory and Reward Information during Perceptual Decision-Making in Lateral Intraparietal Cortex (LIP) of the Macaque Monkey. *PLoS One*, 5(2):1–21.

RUSHWORTH, M.F.S., MARS, R.B. & SUMMERFIELD, C. (2009). General mechanisms for making decisions? *Current Opinion in Neurobiology*, 19:75–83.

SALZMAN, C.D., BELOVA, M.A. & PATON, J.J. (2005). Beetles, boxes and brain cells: neural mechanisms underlying valuation and learning. *Current Opinion in Neurobiology*, 15:721–729.

SALZMAN, C.D., BRITTEN, K.H. & NEWSOME, W.T. (1990). Cortical microstimulation influences perceptual judgements of motion direction. *Nature*, 346:174–177.

SALZMAN, C.D., MURASUGI, C.M., BRITTEN, K.H. & NEWSOME, W.T. (1992). Microstimulation in Visual Area MT: Effects on Direction Discrimination Performance. *The Journal of Neuroscience*, 12(6):2331–2355.

SCHOFFELEN, J.M. & GROSS, J. (2009). Source connectivity analysis with MEG and EEG. *Human Brain Mapping*, 30:1857–1865.

SCHULTZ, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80:1–27.

SCHULTZ, W. (2001). Reward signalling by dopamine neurons. *Neuroscientist*, 7(4):293–304.

SCHWARZ, G. (1978). Estimating the dimension of a model. *Annals of Statistics*, 6:461–464.

SERENCES, J.T. (2008). Value-based modulations in human visual cortex. *Neuron*, 60:1169–1181.

SERENCES, J.T. & SAPROO, S. (2010). Population response profiles in early visual cortex are biased in favour of more valuable stimuli. *Journal of Neurophysiology*, 104:76–87.

SERENO, M.I., DALE, A.M., REPPAS, J.B., KWONG, K.K., BELLIVEAU, J.W., BRADY, T.J., ROSEN, B.R. & TOOTELL, R.B.H. (1995). Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science*, 268:889 – 893.

SHADLEN, M.N., BRITTEN, K.H., NEWSOME, W.T. & MOVSHON, J.A. (1996). A computational analysis of the relationship between neuronal and behavioural responses to visual motion. *The Journal of Neuroscience*, 16(4):1486–1510.

SHADLEN, M.N. & NEWSOME, W.T. (2001). Neural basis of a perceptual decision in the parietal cortex (area LIP) of the rhesus monkey. *Journal of Physiology*, 86:1916–1936.

SHARON, D., HAMALAINEN, M., TOOTELL, R.B.H., HALGREN, E. & BELLIVEAU, J.W. (2007). The advantage of combining MEG and EEG: Comparison to fMRI in focally stimulated visual cortex. *NeuroImage*, 36:1225–1235.

SHULER, M.G. & BEAR, M.F. (2006). Reward timing in the primary visual cortex. *Science*, 311:1606–1610.

SIEGEL, M., ENGEL, A.K. & DONNER, T.H. (2011). Cortical network dynamics of perceptual decision-making in the human brain. *Frontiers in Human Neuroscience*, 5(21):1–12.

SMITH, P.L. & VICKERS, D. (1988). The accumulator model of two-choice discriminations. *Journal of Mathematical Psychology*, 32:135–168.

SOLTANI, A. & WANG, X.J. (2006). A biophysically based neural model of matching law behavioura: melioration by stochastic synapses. *The Journal of Neuroscience*, 26(14):3731–3744.

SOLTANI, A. & WANG, X.J. (2010). Synaptic computation underlying probabilistic inference. *Nature Neuroscience*, 13(1):112–121.

SPECKMAN, P.L. & ROUDER, J.N. (2004). A comment on Heathcote, Brown, and Mewhort's QMLE method for response time distributions. *Psychonomic Bulletin & Review*, 11(3):574–576.

STEFFEN, A., ROCKSTROH, B., WIENBRUCH, C. & MILLER, G.A. (2011). Distinct cognitive mechanisms in a gambling task share neural mechanisms. *Psychophysiology*, 48:1037–1046.

STEPHEN, J.M., AINE, C.J., CHRISTNER, R.F., RANKEN, D., HUANG, M. & BEST, E. (2002). Central versus peripheral visual field stimulation results in timing differences in dorsal stream sources as measured with MEG. *Vision Research*, 42:3059–3074.

SUGRUE, L.P., CORRADO, G.S. & NEWSOME, W.T. (2004). Matching behaviour and the representation of value in the parietal cortex. *Science*, 304:1782–1788.

SUGRUE, L.P., CORRADO, G.S. & NEWSOME, W.T. (2005). Choosing the greater of two goods: neural currencies for valuation and decision-making. *Nature Reviews Neuroscience*, 6:363–375.

SUMMERFIELD, C. & KOECHLIN, E. (2008). A neural representation of prior information during perceptual inference. *Neuron*, 59:336–347.

SUMMERFIELD, C. & KOECHLIN, E. (2010). Economic value biases uncertain perceptual choices in the parietal and prefrontal cortices. *Frontiers in Human Neuroscience*, 4:1–12.

SUPER, H., SPEKREIJSE, H. & LAMME, V.A.F. (2001). Two distinct modes of sensory processing observed in monkey primary visual cortex (V1). *Nature Neuroscience*, 4(3):304–310.

TEHOVNIK, E.G., SLOCUM, W.M. & SCHILLER, P.H. (1999). Behavioural conditions affecting saccadic eye movements elicted electrically from the frontal lobes of primates. *European Journal of Neuroscience*, 11:2431–2443.

TEICHERT, T. & FERRERA, V.P. (2010). Suboptimal integration of reward magnitude and prior reward likelihood in categorical decisions by monkeys. *Frontiers in Neuroscience*, 4(186):1–13.

TREUE, S. (2003). Visual attention: the where, what, how and why of saliency. *Current Opinion in Neurobiology*, 13:428–432.

TREUE, S., HUSAIN, M. & ANDERSEN, R.A. (1991). Human perception of structure from motion. *Vision Research*, 31(1):59–75.

TREUE, S. & MAUNSELL, J.H.R. (1999). Effects of attention on the processing of motion in macaque middle temporal and medial superior temporal visual cortical areas. *The Journal of Neuroscience*, 19(17):7591–7602.

TUERLINCKX, F. (2004). The efficient computation of the cumulative distribution and probability density functions in the diffusion model. *Behaviour Research Methods, Instruments & Computers*, 36(4):702–716.

UKA, T. & DEANGELIS, G.C. (2006). Linking neural representation to function in stereoscopic depth perception: roles of MT in course vs. fine disparity discrimination. *The Journal of Neuroscience*, 26(25):6791.

ULLMAN, S. (1979). *The Interpretation of Visual Motion*. MIT Press.

USHER, M. & MCCLELLAND, J.L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*, 108:550–592.

VAN ESSEN, D.C., MAUNSELL, J.H.R. & BIXBY, J.L. (1981). The middle temporal visual area in the macaque: myeloarchitecture, connections, functional properties and topographic organization. *The Journal of Comparative Neurology*, 199:293–326.

VAN VEEN, B.D. & BUCKLEY, K.M. (1988). Beamforming: A versatile approach to spatial filtering. *ASSP Magazine*, 5(2):4–24.

VAN VEEN, B.D., VAN DRONGELEN, W., YUCHTMAN, M. & SUZUKI, A. (1997). Localization of brain electrical acitvity via linearly constrained minimum variance spatial filtering. *IEEE Transactions on Biomedical Enginerring*, 44(9):867–880.

VANDEKERCKHOVE, J. & TUERLINCKX, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, 14(6):1101–1026.

VANDEKERCKHOVE, J. & TUERLINCKX, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behaviour Research Methods*, 40(1):61–72.

VOSS, A., ROTHERMUND, K. & VOSS, J. (2004). Interpreting the parameters of the diffusion model: An empirical investigation. *Memory and Cognition*, 32(7):1206–1220.

VOSS, A., ROTHERMUND, K. & VOSS, J. (2008). Interpreting ambiguous stimuli: separating perceptual and judgemental biases. *Journal of Experimental Social Psychology*, 44:1048–1056.

WALLACH, H. & O'CONNELL, D.N. (1953). The kinetic depth effect. *Journal of Experimental Psychology*, 45(4):205–217.

WANDELL, B.A. (1999). Computational neuroimaging of human visual cortex. *Annual Reviews of Neuroscience*, 22:145–73.

WANDELL, B.A., BREWER, A.A. & DOUGHERTY, R.F. (2005). Visual field map clusters in human cortex. *Philosophical Transactions of the Royal Society*, 360:693–707.

WANG, X.J. (2010). Neurophysiological and computational principles of cortical rhythms in cognition. *Psychological Review*, 90:1195–1268.

WEIL, R.S., FURL, N., SYMMONDS, M., FLANDIN, G., DOLAN, R.J., DRIVER, J. & REES, G. (2010). Rewarding feedback after correct visual discrimination has both general and specific influences on visual cortex. *Journal of Neurophysiology*, 104(3):1746–1757.

WEISKRANTZ, L. (1986). *Blindsight: A case study and implications*. Oxford University Press, Oxford.

WENZLAFF, H., BAUER, M., MAESS, B. & HEEKEREN, H.R. (2011). Neural characterization of the speed-accuracy tradeoff in a perceptual decision-making task. *The Journal of Neuroscience*, 31(4):1254–1266.

WHEATSTONE, C. (1838). Conributions to the Physiology of Vision. - Part the First. On some

remarkable, and hitherto unobserved, Phenomena of Binocular Vision. *Philosophical Transactions of the Royal Society*, 128:371–394.

WINTERER, G. & WEINBERGER, D.R. (2004). Genes, dopamine and cortical signal-to-noise ratio in schizophrenia. *TRENDS in Neurosciences*, 27(11):683–690.

WOOLRICH, M., HUNT, L., GROVES, A. & BARNES, G.R. (2011). MEG beamforming using Bayesian PCA for adaptive data covariance matrix regularization. *NeuroImage*, 57:1466–1479.

YARBUS, A.L. (1967). *Eye movements and vision*. New York: Plenum Press.

YOSHIOKA, T., TUYAMA, K., KAWATO, M., YAMASHITA, O., NISHINA, S., YAMAGISHI, N. & AKI SATO, M. (2008). Evaluation of hierarchical Bayesian method through retinotopic brain activities reconstruction from fMRI and MEG signals. *NeuroImage*, 42:1397–1413.

ZEKI, S. (1990). A century of cerebral achromatopsia. *Brain*, 113:1721–1777.

ZEKI, S.M. (1974). Functional organization of a visual area in the posterior bank of the superior temporal sulcus of the rhesus monkey. *Journal of Physiology*, 236:549–573.

ZIHL, J., VON CRAMON, D. & MAI, N. (1983). Selective disturbance of movement vision after bilateral brain damage. *Brain*, 106:313–340.