

Expectation in motor planning and execution

Isobel Claire Weinberg

A dissertation submitted in partial fulfilment of the requirements for
the degree of Doctor of Philosophy (Neuroscience)

at



University College London
August 2017

I, Isobel Claire Weinberg confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Isobel Weinberg, August 2017

Abstract

Expectation has been studied extensively in the decision-making field and its possible implementation in influential decision-making models has been formulated. Decision-making has historically been studied separately to motor planning. However, recent data suggest decision-making and motor planning overlap in time, with competing action plans in motor cortex biased by the ongoing decision. There has therefore been increasing interest in studying the interplay between decision-making and motor planning.

Past experiments have typically studied the decision between two movements, when each movement is equally likely to be chosen. This aims to mimic the everyday situation in which we prepare a movement before knowing which it will be. However, a more common situation is that we expect to make one movement with a high likelihood, but also know there is a low likelihood of making a different movement. It is this uneven expectation across potential movements, and its effect on motor planning and execution, that is the focus of this thesis.

I first investigate expectation in motor planning. I propose expectation may play the same role in motor cortical excitability as it is proposed to in theoretical decision making models. A series of experiment did not support this hypothesis; I discuss possible reasons for this.

I next turn to an aspect of action execution: motor variability. There has been increasing interest in the idea that noise during motor planning is an important cause of motor variability. One theory has proposed that neural resources are divided when there are multiple motor plans, increasing motor variability. I propose that expectation interacts with this process by sharing these neural resources unevenly, so that variability is lower in the high-likelihood movement. I conduct two experiments to test

this idea, and, based on the results, propose that expectation interacts with the motor control policy to determine motor variability.

Table of contents

Abstract.....	3
Chapter One: Introduction	19
1.1. Introduction	19
1.2. Motor Planning	20
1.2.1. Actions and decisions occur in parallel, not series	20
1.2.2. Behavioural data in support of parallel motor planning.....	25
1.2.3. Two goals, one plan?	28
1.2.4. Summary and relevance.....	31
1.3. Prior Expectation	33
1.3.1. Decision model predictions about role of prior expectation....	34
1.3.2. High-level representations of expectation.....	40
1.3.3. Expectation in sensory areas	41
1.3.4. Expectation in motor representations.....	44
1.3.5. Differentiating attention and expectation.....	45
1.3.6. Summary and relevance.....	46
1.4. Motor variability	48
1.4.1. Motor variability as a limiting factor in motor control	48
1.4.2. The source of motor variability	51
1.4.3. Limited resources as a cause of motor variability	53
1.4.4. The advantages of motor variability	55
1.4.5. Summary and relevance.....	58
1.5. Transcranial Magnetic Stimulation.....	59
1.5.1. The motor evoked potential	59

1.5.2. The MEP as an assay of competing action plans	61
1.5.3. Inhibition of MEPs in the effector that moves.....	64
1.5.4. Summary and relevance.....	67
Chapter Two: The representation of expectation in corticospinal excitability during motor planning.....	68
2.1. Introduction	68
2.1.1. Assaying competition between motor plans	68
2.1.2. Expectation as a prestimulus bias	69
2.1.3. Hypothesis and key experimental features	70
2.1.4. What this study adds	71
2.1.5. An experimental challenge: minimising impulse supression ..	73
2.1.6. Three experiments	74
2.2. Methods.....	76
2.2.1. Participants	78
2.2.2. Experiment 1	78
2.2.3. Experiment 2	82
2.2.4. Experiment 3.....	85
2.2.5. Subject exclusions	85
2.2.6. Preprocessing of behavioural data.....	86
2.2.7. Preprocessing of MEPs	86
2.2.8. 'Too Slow' trials.....	87
2.2.9. Control trials.....	87
2.2.10. MEP normalisation	88
2.2.11. Primary and secondary analyses.....	89
2.2.12. Modelling analysis.....	89

2.2.13. Analysis of MEPs over time.....	93
2.2.14. Statistical analysis	94
2.3. Results.....	95
2.3.1. Experiment 1	95
2.3.2. Comparison of Experiment 2 and Experiment 3.....	99
2.3.3. Experiment 2 and Experiment 3 combined dataset	100
2.3.4. Comparison to baseline	122
2.4. Discussion.....	124
2.4.1. Was ‘impulse suppression’ in these experiments the reason for a lack of effect on MEPs at stimulus onset?	124
2.4.2. Was anatomical specificity of competitive representations the reason for a lack of effect on MEPs at stimulus onset?.....	127
2.4.3. Was the experiment underpowered?.....	128
2.4.4. Temporal variability of MEPs	129
2.4.5. Differences between conditions at stimulus onset when split by choice	129
2.4.6. Relationship between late MEPs and reaction speed	130
2.4.7. Suitability of TMS for experiments of this nature	131
2.4.8. Replication of previous studies.....	132
2.4.9. Alternative experimental approaches	133
2.4.10. Conclusion	135
2.5. Appendix I: ANOVA to compare MEP data from Experiments 2 and 3.....	137
2.6. Appendix II: ANOVA on stimulus-locked MEP data	138
Chapter Three: Does prior expectation distribute variability unevenly across motor plans?.....	139

3.1. Introduction	139
3.1.1. A hypothesis linking prior expectation and motor variability	141
3.1.2. Experimental design	142
3.2. Methods	143
3.2.1. Participants	143
3.2.2. Robotic apparatus	143
3.2.3. Trial protocol	144
3.2.4. Experimental protocol	147
3.2.5. Analysis	148
3.3. Results	154
3.3.1. Success rate	154
3.3.2. Reaction Time and Movement Time	156
3.3.3. Movement trajectories	157
3.3.4. Reach angle	159
3.3.5. Directional Bias	165
3.3.6. Differences between leftward and rightward movements ...	167
3.3.7. Reaction time as a determinant of variability	170
3.3.8. Learning	173
3.4. Discussion	176
3.4.1. Why does variability differ across expectation condition in leftward movements?	177
3.4.2. Why is there no variability difference in rightward movements?	177
3.4.3. Why does the variability difference between conditions diminish over time?	181

3.4.4. Why is there a difference in mean reach angle?.....	181
3.4.5. Variability effects are not a result of more changes of mind.	182
3.4.6. Variability effects are not driven by reaction time differences across conditions.....	183
3.4.7. Learning does not reduce motor variability in this experiment	184
3.4.8. Planning versus execution: the source of motor variability ..	185
3.4.9. Comparison to previous studies.....	186
3.4.10. Alternative metrics for analysing movement trajectories ..	186
3.4.11. Conclusion	187
3.5. Appendix I: ANOVA on learning effects within initial reach angle	189

Chapter Four: Do planning-related variability differences depend on a stringent control policy? 190

4.1. Introduction	190
4.1.1. Encouraging a non-stringent control policy	190
4.1.2. Target jumps as a paradigm to manipulate control policy.....	191
4.1.3. Relevance to Chapter Two	192
4.1.4. Relevance to target jump literature.....	193
4.2. Methods.....	195
4.2.1. Overview	195
4.2.2. Participants.....	195
4.2.3. Robotic apparatus.....	196
4.2.4. Trial protocol	196
4.2.5. Block types: jump and non-jump	199
4.2.6. Experimental design	200

4.2.7. Analysis.....	201
4.3. Results.....	205
4.3.1. Characteristics of jump trials.....	206
4.3.2. Reaction time and movement time	209
4.3.3. Replication of previous study	211
4.3.4. An interaction between block type and prior cue?	212
4.3.5. Early movement time and movement velocity	216
4.3.6. Variability as determined by success	218
4.4. Discussion.....	221
4.4.1. Effect of expectation on movement variability successfully replicated	221
4.4.2. Target jump manipulation increased movement variability.	221
4.4.3. Main analysis of interest did not support hypothesis	222
4.4.4. Was the hypothesis falsified?.....	225
4.4.5. Was the study underpowered?	227
4.4.6. Conclusion	227
4.5. Appendix I: ANOVAs on reaction time and movement time	228
4.5.1. Reaction time	228
4.5.2. Movement time	229
Chapter Five: General Discussion	230
5.1. Where next for the affordance competition hypothesis?	231
5.2. New competitors to the drift-diffusion model.....	235
5.3. Alternative theories of expectation in motor cortex.....	237
5.4. Appraisal of methods used in thesis.....	239
5.4.1. Paradigms for delivering probabilistic information.....	239

5.4.2. Measures of variability	240
5.5. Future work.....	241
Acknowledgements.....	245
References.....	246

List of figures

Figure 1.1 Action plans in monkey dorsal premotor cortex are biased by decision information.	21
Figure 1.2 PMd cells respond to the value of potential movements, but only when there is a choice.	23
Figure 1.3 Corticospinal excitability reflects dynamic competition between responses.	24
Figure 1.4 Nearby targets permit intermediate movements, whilst widely-spaced targets do not.	25
Figure 1.5 Intermediate trajectories in a 'go before you know' paradigm reflect target position	27
Figure 1.6 This ambiguous scene becomes easy to interpret with information about what it depicts; see text.	33
Figure 1.7 The drift-diffusion model implements a statistically optimal method for making a decision.	35
Figure 1.8 A random dot kinetogram task and corresponding activity in MT and LIP neurons.	36
Figure 1.9 Prior expectation could be incorporated into a drift-diffusion model in one of two ways: by reducing the distance between the two thresholds or increasing the gain on the integration of sensory evidence.	38
Figure 1.10 The integration of new information and existing beliefs, in a hierarchical network for predictive coding.	43
Figure 1.11 A visual perturbation is corrected for when it affects task performance, but not when it does not.	50
Figure 1.12 The waves evoked by PA TMS depend on the intensity of stimulation.	61
Figure 1.13 MEPs are different in the chosen versus unchosen effector during the decision time.	63

Figure 1.14 MEP suppression during the delay period affects all effectors and is strongest in the hand the subject expects to move.....	65
Figure 2.1 Tasks for the three experiments.	77
Figure 2.2 Parameters in the drift-diffusion model.....	91
Figure 2.3 Reaction time decreases with increasing rightward expectation.	96
Figure 2.4 Errors decrease with increasing prior expectation.....	97
Figure 2.5 Raw MEP size does not change with prior expectation.....	97
Figure 2.6 Relative MEP size does not change with prior expectation.	98
Figure 2.7 Right hand MEPs over time in the three prior cue conditions.	99
Figure 2.8 Subjects were fastest with the response they most expected to make.	101
Figure 2.9 Behavioural choice is modulated by prior cue.	102
Figure 2.10 No modulation of MEP by prior expectation cue at stimulus onset.	103
Figure 2.11 No modulation of raw MEPs by prior expectation cue at stimulus onset.	103
Figure 2.12 No effect of prior cue condition on MEP amplitude for trials in which the same prior cue had been seen on that trial and the preceding one.....	104
Figure 2.13 Post-hoc analysis shows that choice affects MEP size at stimulus onset in some conditions.	105
Figure 2.14 Modelled bias parameters by probability cue condition.....	107
Figure 2.15 No relationship between bias parameter slope and MEP slope.	108
Figure 2.16 No relationship between reaction speed slope and MEP slope.	110
Figure 2.17 MEPs plotted relative to stimulus onset are suppressed at the second timepoint and elevated at the third.....	111

Figure 2.18 Stimulus-locked MEPs do not vary by condition in when split by timepoint x prior cue x hand.	112
Figure 2.19 Response locked MEPs suggest activity may diverges later when prior expectation is stronger.	114
Figure 2.20 Response locked MEPs on trials in which subjects chose against the prior cue do not show statistically-significant divergence.	116
Figure 2.21 No relationship between bias parameter slopes and late MEP slopes.	118
Figure 2.22 Positive relationship between reaction speed slope and late MEP slope.	119
Figure 2.23 Control trial MEPs show separation by activity at latest timepoint, but not at stimulus onset or interim timepoints.	121
Figure 2.24 MEPs are elevated relative to pre-block baseline measure.	122
Figure 3.1 The robotic apparatus used allowed subjects to control a cursor with their hand movements.	144
Figure 3.2 Information on screen during a trial of the experiment.	145
Figure 3.3 Reach angle was calculated as the angle between position at sample and rest position, relative to the vertical midline.	149
Figure 3.4 Prior expectation influences success rate.	154
Figure 3.5 Reaction times are modulated by expectation but movement times are not.	156
Figure 3.6 Movement trajectories by expectation.	157
Figure 3.7 Cursor position at last sample before onset of forward movement, by cued target and prior expectation condition.	158
Figure 3.8 Left-cue trials show an early modulation of reach angle and reach angle variability by prior expectation.	160
Figure 3.9 Reach angle histograms in one subject.	163

Figure 3.10 Left-cue trials show a modulation of reach angle and reach angle variability by prior expectation when a speed threshold is applied.....	164
Figure 3.11 The effect of prior probability cue on reach angle variability in left-cue trials is not dependent on sample size.....	165
Figure 3.12 Prior probability cue and target affect starting direction....	166
Figure 3.13 Reach angle variability was higher for right-target movements than left-target movements early in the movement.	168
Figure 3.14 Early position is a stronger determinant of endpoint position in left-cue movements than right cue movements.	169
Figure 3.15 Reaction time does not predict endpoint error on a trial-by-trial basis.	171
Figure 3.16 Left-cue trials with faster reaction times have lower variability.	172
Figure 3.17 There is no effect of reaction time on movement variability when the prior cue seen is controlled for.....	173
Figure 3.18 Subjects improved over the course of the experiment.	174
Figure 3.19 Reach angle shows learning across experiment but reach angle variability does not.....	175
Figure 4.1 On-screen display for one trial.....	196
Figure 4.2 Success is lower in target jump blocks.....	206
Figure 4.3 Apparent target jump differed from the predetermined target jump shown in Figure 4.2B.	207
Figure 4.4 Mean reach angle and reach angle variability differ between non-jump and jump blocks.....	209
Figure 4.5 Reaction times are modulated by prior expectation but not block type.	210
Figure 4.6 Movement times are modulated by block type but not prior expectation.	211

Figure 4.7 As in the previous study, for trials in non-jump blocks, mean reach angle and reach angle variability depend on expectation in left-cue but not right-cue trials.	212
Figure 4.8 Effect of block type and expectation on reach angle mean and variability.	214
Figure 4.9 Block type affects time taken to reach occluder.	216
Figure 4.10 Block type affects movement velocity in movement prior to the occluder.	218
Figure 4.11 Variability is increased by lack of success for some trials only.	219

List of tables

Table 2.1 MEP data does not differ between Experiments 2 and 3.....	137
Table 2.2 No significant main effects in stimulus-locked MEP data.....	138
Table 3.1 Prior expectation does not affect block-by-block changes in reach angle.....	189
Table 4.1 Bayesian ANOVA shows anecdotal evidence against the interaction of interest, and extremely strong evidence for the two-main effect model.....	215
Table 4.2 Reaction time showed a significant main effect of prior expectation and a significant target x prior interaction.....	228
Table 4.3 Movement time showed a significant main effect of target cued and block type.	229

Abbreviations

BOLD	Blood-oxygenation-level dependent
EEG	Electro-encephalography
FEF	Frontal eye fields
FOF	Frontal orienting fields
fMRI	Functional magnetic resonance imaging
LIP	Lateral intraparietal cortex
M1	Primary motor cortex
MEG	Magneto-encephalography
MEP	Motor evoked potential
OFC	Optimal feedback control
PMd	Dorsal premotor cortex
PPC	Posterior parietal cortex
SD	Standard deviation
TMS	Transcranial magnetic stimulation

Chapter One: Introduction

1.1. Introduction

My PhD is about the interface between decision and actions. The fields studying these two processes have evolved separately, but there is increasing evidence that their brain implementation is not sequential, but concurrent and interactive.

In this introductory chapter, I review literature from both fields. On the side of movement, I begin by discussing literature from the field of action planning, with a focus on a theory that proposes multiple actions are planned in parallel. I explain how neurophysiological experiments in monkeys have proposed hypotheses that have been amenable to behavioural testing.

Turning to the decision literature, I focus on the role of expectation. Our beliefs about the world change the outcome of our decisions, and there has been much effort to place expectation into a theoretical framework.

The third section of this chapter discusses ideas about the variability present during repeated execution of the same movement. From where and why does this variability arise? Motor variability has been characterised alternately as a limiting factor on the motor system and an exploratory behaviour; as arising primarily from noise in action execution or primarily from noise in action planning. The experiments in my PhD applied ideas about expectation and action planning to the study of motor variability.

Transcranial magnetic stimulation has been one of the experimental tools used in my PhD. In the fourth section of this chapter, I review the methodological basis of transcranial magnetic stimulation and discuss its use as an assay of competing motor plans.

1.2. Motor Planning

1.2.1. Actions and decisions occur in parallel, not series

A common metaphor for the brain is that of a computer: they both seem to have inputs, execute programs, and produce outputs (Cisek, 1999). This concept of sequential computation (Donders, 1969) influenced neuroscience to view cognition and action planning as discrete modules, with cognitive processes executed before action planning begins. In such a view of action planning, we first resolve choice between competing targets – which item of food on the table to reach towards – and then motor cortex plans the kinematics of a reach towards the selected target. Sherrington saw response selection as the interface between anatomically distinct sensory and motor systems (Sherrington, 1910).

The assumptions underlying this way of thinking have been influential, and there has been a traditional separation between the study of neural decision making and of action planning. The decision making field has been shaped by classic experiments in monkeys that show parietal cortex neurons which are interested in the balance of evidence for competing hypotheses in a perceptual decision (Shadlen and Newsome, 2001a; Roitman and Shadlen, 2002). On the action planning side, primate experiments have shown there is preparatory activity prior to a movement in premotor cortex (Tanji and Evarts, 1976; Riehle and Requin, 1989; Alexander and Crutcher, 1990) and disrupting this activity delays movement onset (Churchland and Shenoy, 2007).




Figure 1.1 Action plans in monkey dorsal premotor cortex are biased by decision information. The spatial cues indicate two possible reach directions, and PMd responds by planning a potential movement to each cue. These are maintained during the memory period. When the colour cue indicating the correct reach direction appears, activity in the corresponding motor plan is amplified and the alternate plan is suppressed. The Go signal tells the monkey to execute the movement. Figure from Cisek and Kalaska (2010)

This theoretical separation has been challenged by neurophysiological data in which dorsal premotor cortex appears to reflect multiple possible actions plans (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011). Neural activity recorded from monkey premotor cortex is shown in Figure 1.1 for a task in which monkeys have to remember two possible locations to reach to, wait for a colour cue that indicates which is the correct target, and then make their reach (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011). In the memory period, dorsal premotor cortex (PMd) shows elevated activity in two sets of cells: those with receptive fields oriented to the two potential targets. When the colour cue appears, activity representing that target begins grows, whilst activity for the other target diminishes. So PMd maintains two potential motor plans whilst the direction of eventual movement is under consideration. The decision is reflected in this motor planning activity: multiple potential options are specified and gradually eliminated in response to relevant information. This is despite the fact that this strategy is not necessary for

success: monkeys could have withheld all motor planning processes until they had all the information needed.

These data have led to the ‘parallel processing’ idea: decision-making and motor-planning take place on parallel timescales, with motor plans being continually updated with new information (Cisek, 2005). The idea has been formalised into the “affordance competition hypothesis” (Cisek, 2007) in which an organism’s preoccupying challenge is to choose from the many competing potential actions (‘affordances’) offered by the environment, rather than specify the movement. As soon as potential targets for action are identified, movements begin to be planned. The decision about where to move to goes on in parallel, and, as it evolves towards completion, the winning hypothesis biases the corresponding action plan by amplifying its activity. The action plan corresponding to the losing hypothesis is suppressed. Multiple action plans thus compete for dominance in a similar way to that envisaged for conventional “decision-making” activity (Gold and Shadlen, 2007). The motor planning activity is biased by a range of relevant information, such as that from basal ganglia and prefrontal cortex. Choices emerge via a ‘distributed consensus’, in which relevant activity in one region propagates across the brain. Thus decisions based on stimulus features might emerge from sensory cortex and propagate forwards, whilst those based on abstract rules would emerge from frontal regions and propagate backwards, but both would ultimately bias movement plans.

A parallel processing scheme predicts that all factors that influence a choice between actions will be represented in motor cortex, not because they are explicitly encoded there, but because they are reflected in the competition between actions. Thus, when monkeys must hold targets in memory with potentially different reward values, the strength of activity for each target is biased by the value of the target (Pastor-Bernier and

Cisek, 2011). These biases are relative: activity is suppressed by a higher value for the opposing target. Importantly, when there is no choice to be made, value does not influence activity at all (

Figure 1.2). Decision variables are represented in so far as they subserve competition between actions.

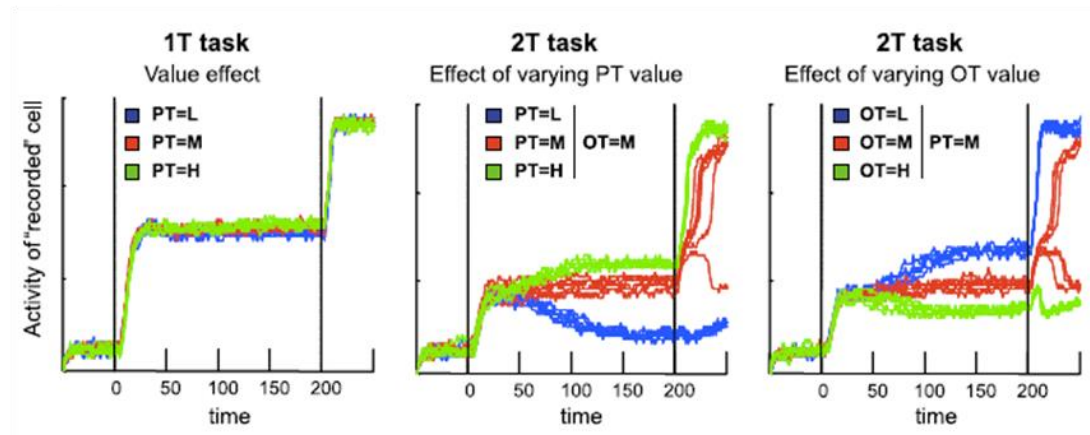


Figure 1.2 PMd cells respond to the value of potential movements, but only when there is a choice. Plots show activity of one cell in caudal PMd from onset of targets (first vertical black line) to go cue (second vertical black line). Left panel: in the one target version of the task, the cell shows no modulation by value. Middle panel: in the two target version of the task, the value of the cell's preferred target affects its activity so that high-value targets (green line) lead to more activity than medium (red line) or low value targets (blue line). This pattern is reversed when the value altered is that of the opposite target (right panel). Figure from Pastor-Bernier and Cisek (2011).

There is further evidence from other modalities that 'cognitive' factors are represented in motor cortex. In a value-based choice, the subjective value difference between the alternatives is reflected in corticospinal excitability (Klein-Flugge and Bestmann, 2012). During a perceptual decision, accumulating sensory evidence corresponds to the lateralisation of motor-selective MEG activity (Donner et al., 2009). In a task with confusing, distracting flankers, corticospinal excitability shows the dynamic competition between the two plans (Figure 1.3), with activity first growing in the hand that favours the flankers, before being superseded by activity on the correct side (Michelet et al., 2010).

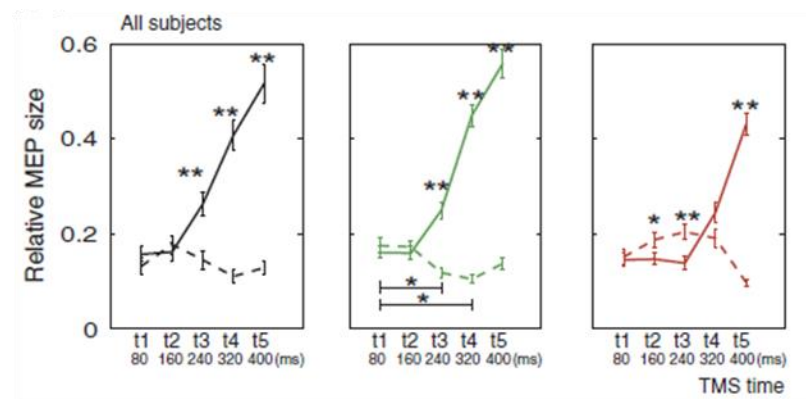


Figure 1.3 Corticospinal excitability reflects dynamic competition between responses. In this task, subjects had to flex or extend a finger, as cued by a central arrow. The figure shows the corticospinal excitability in the chosen muscle (solid line) and unchosen muscle (dashed line) when (1) only the central arrow was present (black lines), (2) the central arrow was flanked by ‘congruent’ side arrows that pointed in the same direction (green lines), and (3) the central arrow was flanked by ‘incongruent’ side arrows that confusingly pointed in the opposite direction (red lines). The red lines show an initial activation of the incorrect motor plan, followed by a change of mind. Figure from Michelet et al (2010).

Recent neuronal recording studies have shown how decision-making activity might transform from purely sensory to information suitable to subserve an action as it travels from sensory association cortex to motor association cortex. Posterior parietal cortex (PPC) is typically implicated in evidence accumulation, whilst frontal eye fields (FEF) are involved in saccade generation. In the rat, there is a specialisation in the tuning of these two areas, with PPC accumulating evidence in a graded manner, whilst the rat FEF homologue represents this information in a more categorical manner, analogous to, “If I had to go now, which way would I choose?” (Erlich et al., 2015; Hanks et al., 2015). With similar evidence of a progressive transformation from pure decision to decision-for-action information is a study of subjective decision-making in macaques, which find that the value of a choice is encoded in abstract form in orbitofrontal cortex and relative to action in lateral prefrontal cortex (Cai and Padoa-Schioppa, 2014).

1.2.2. Behavioural data in support of parallel motor planning

The neurophysiological data in Figure 1.1 have been replicated by modelling (Cisek, 2007). This draws on a literature which has seen premotor cortex as forming a population code over a movement-related metric, typically reach angle (Bastian et al., 2003; Georgopoulos and Carpenter, 2015). The two potential plans are modelled as peaks of activity existing within a continuous distribution of reach angles and thus represent a probability density function of potential movements. The peaks of activity interact with one another, competing when they are far apart, and mutually reinforcing when they are close together (Erlhagen and Schöner, 2002).

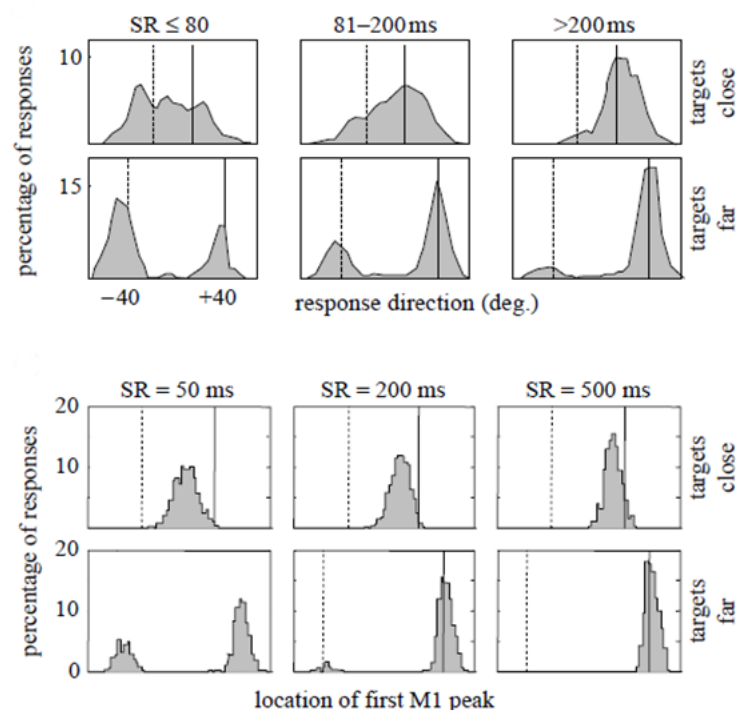


Figure 1.4 Nearby targets permit intermediate movements, whilst widely-spaced targets do not. Behavioural data (top panels; Ghez et al., 1997) and simulations of preferred direction of first M1 cell exceeding threshold (bottom panels). With longer intervals permitted for movement selection (left panels: shortest interval; right panels: longest interval), behavioural responses and modelled activity settle on the response orientation corresponding to rightward target. However, if a response is forced before this process is complete, responses tend to be between the two targets when they are close, but bimodally distributed to one or other of the targets when they are spaced far apart. Figure from Cisek (2007).

This makes a prediction: sometimes, the interaction between competing, similar action plans will determine the final movement, rather than the original plans. This idea is powerfully supported by the ability of Cisek's model (2007) and others similar (Erlhagen and Schöner, 2002) to reproduce behavioural data. If human subjects are cued to move before they know where to go, and the targets are far apart, they will select one or other target to move to (Ghez et al., 1997). However, if the targets are close together, a movement aiming between the two targets is made (Figure 1.4). This behaviour is explained by the models: positive feedback between similar actions leads to a single broad peak of activity, executed as an intermediate movement (Cisek, 2007). However, connections between actions which are spaced far apart in the population code are inhibitory and so only one action can win the competition. The model also explains why reaction time is longer when targets are far apart in space but not when there are more of them (Bock and Eversheim, 2000): interactions within a continuous distribution means angle subtended, rather than number of movements, is the important determinant of competition.

If intermediate movements are the result of interactions between competing plans, they offer a behavioural window into this competitive process (Ghez et al., 1997). A 'go before you know' paradigm forces subjects to begin moving before they know exactly which of several potential movements they will have to make. This paradigm has shown that intermediate reach trajectories (Chapman et al., 2010) and initial hand orientation (Stewart et al., 2013) do indeed reflect competition between the potential options: for instance, in a situation in which subjects are twice as likely to be cued to the right as to the left, the 'before you know' intermediate trajectory is pulled to the right (Figure 1.5). Saccade trajectories can also be made curved by the presence of distractors (McPeck et al., 2003).

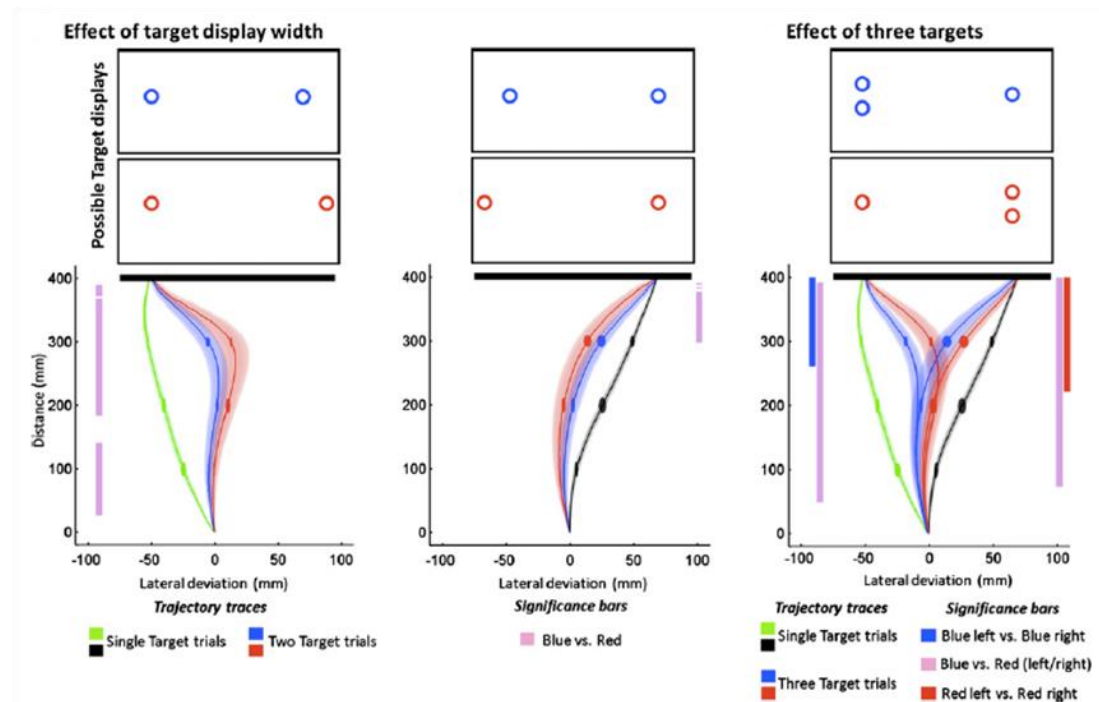


Figure 1.5 Intermediate trajectories in a ‘go before you know’ paradigm reflect target position and target likelihood. Red and blue lines show trials where both targets were present, whilst green and black lines show trials where only one target is present. On one target trials, trajectories aim towards the target. On two target trials, trajectories are initially oriented in between potential targets. Moving one target further to the side deviates this intermediate trajectory correspondingly (A). When one side of the display is more likely to be cued than the other, the intermediate trajectory is deviated towards the likely side (B). Figure from Chapman et al (2010).

Intermediate trajectories have even been found to reflect competition between high-level cognitive parameters during decisions (Song and Nakayama, 2009). For a task in which subjects have to make a movement to report whether one number is greater or less than another, there are more intermediate trajectories when the discrimination is difficult because the numbers are closer in value (Song and Nakayama, 2008). Similarly, relevant distractor words in a language-based task increase the proportion of intermediate trajectories (Finkbeiner et al., 2008).

The evidence above is compelling but it is possible that the intermediate trajectories are generated by averaging between the sensory representations of the two targets, with a trajectory then generated to this intermediate pseudo-target, rather than a true a motor averaging (Stewart et al., 2014; Gallivan et al., 2017). Arguing for a true motor

averaging, Gallivan et al (2015) found that when subjects were forced to move to a target that was ambiguous in the movement it specified, they tended to use motor elements of a plan to a previously-seen, non-ambiguous target. These results depended on interactions between specific movements and not the spatial organisation of the targets. This behaviour was observed despite the absence of a ‘go before you know’ paradigm: participants had all the necessary information before they began their movement.

In fact, interactions between movement plans appear to extend beyond movement specification to movement execution. Feedback gains are also averaged between targets (Gallivan et al., 2016b) and feedback gains in turn reflect the evolving decision (Selen et al., 2012). This suggests competition is at the level of feedback control policies, not just endpoints.

Behavioural data has not only been useful to understand the competitive processes entailed by parallel processing, but might hint at the reasons this system confers an adaptive advantage. Subjects tend to reuse elements of a competing plan in the executed plan, and when the two competing movements are compatible, reaction and movement times are shorter (Gallivan et al., 2015). In other words, sharing motor plan parameters might speed motor processing or reduce working memory load. Furthermore, intermediate movements may optimise reward or defer final decision-making in situations where evidence is uncertain (Haith et al., 2015; Wong and Haith, 2017).

1.2.3. Two goals, one plan?

Despite the evidence presented above, the affordance competition hypothesis is contested (Padoa-Schioppa, 2011). Wong and Haith (2017) have argued that intermediate trajectories represent a deliberate movement to an intermediate goal, rather than the execution of a

combination of plans. Similarly, intermediate feedback gains (Gallivan et al., 2016b) reflect selection of a halfway-house control policy which balances cost and accuracy. The findings of Stewart et al (2014) - that when there is an obstacle in the way of one plan, but not the intermediate movement, the intermediate movement is nevertheless deflected away from the obstructed side – is similarly explained as not a motor averaging phenomenon, but a unitary control policy, in which the obstructed movement is more effortful and thus costed against. Averaging is thus deliberate selection of, and planning for, an intermediate goal, because this is an optimal strategy (Haith et al., 2015), rather than because a combination of plans is being executed.

These ideas are drawn from a viewpoint in which motor planning is a process distinct from selecting a motor goal (Wong et al., 2015). What is the difference between a goal and a plan? Wong and Haith (2017) characterise the distinction as analogous to that between a cost function and a control policy in optimal control theory; in other words, one is strategic and the other is implementational. Whilst there are multiple potential goals, there is only one motor plan. The activity relating to multiple visual targets in premotor cortex (Cisek and Kalaska, 2005) is the representation of multiple possible motor goals rather than motor plans. Planning, unlike choosing a goal, is characterised as a typically quick process that has little influence on the reaction time (Wong et al., 2015, 2016). Planning follows serially from goal selection. This does not mean that planning must wait until goal selection is complete – a plan might be to an intermediate goal if the decision process is still ongoing – but a single plan is produced.

To support this argument, Wong and Haith (2017) argue that activity in frontal eye fields represents stimulus, potential goals, and the motor

plans in turn. Motor cortex might similarly represent motor goals which evolve to a single plan.

In this context, the finding that adding virtual barriers (Haith et al., 2015) or requiring speeded movements (Wong and Haith, 2017) abolishes intermediate movements is evidence for intermediate goals, on the grounds that the a motor averaging phenomenon in premotor cortex must be too low level to incorporate task rules.

Further evidence for the separation of plans and goals comes from the field of motor skill learning: learning is believed to bind together execution-related activity in the spinal cord or primary motor cortex ('motor primitives') and thus over time automate the process of response selection (Diedrichsen and Kornysheva, 2015).

These ideas have been conceptually linked to neurophysiological work describing motor cortex as a dynamical system for generating movement. According to this view, the goal of preparatory activity is to bring neural cortical activity to a particular state, from which it will evolve passively into movement activity without requiring additional input (Shenoy et al., 2013). This view, in which motor cortex acts as an internal pattern generator, is in contrast to traditional views which regard the motor cortex as coding distributions of movement-related variables, such as reach angle (Georgopoulos and Carpenter, 2015). There is experimental evidence in favour of the dynamical systems view: neural variability diminishes during movement preparation, as would be expected by a dynamical system bringing activity to a common starting point (Churchland et al., 2006, 2010). When projected into a lower dimensional space, neural responses tend to rotate with an amplitude and phase given by the preparatory activity (Churchland et al., 2012), suggesting oscillatory activity akin to that found in other motor systems (Grillner, 2006).

The dynamical systems view has the motor plan represented by a single population, with the firing rate of individual neurons only incidentally correlating with movement parameters such as speed and direction. This is argued to be incompatible with representing plans as two competing sets of neurons tuned to movement parameters (Cisek and Kalaska, 2005; Wong and Haith, 2017). Furthermore, the dynamical systems view argues that, because neural tuning to, for example, velocity or reach direction does not necessarily remain consistent from the preparatory period to the movement period, there is no need for preparatory neural activity to be a subthreshold version of activity in the movement period (Kaufman et al., 2013, 2014). This is a challenge to Cisek's model which encompasses rise-to-threshold behaviour (Cisek, 2007).

To summarise, the line of thinking which argues that the representations of plans and goals in motor cortex are conceptually distinct posits that cue-related activity in premotor cortex represents multiple goals rather than multiple plans. Spatial averaging behaviour arises from selection of an intermediate goal rather than an averaged plan.

1.2.4. Summary and relevance

Neurophysiological evidence shows that motor plans compete for dominance in premotor cortex and this has been used to understand human behaviour in which movements reflect the competing alternatives available. The sharing of neural resources implicit in parallel processing has implications for motor variability; this is the subject of an experiment in Chapter 3.

I have detailed the controversy over whether the activity patterns seen in premotor cortex truly reflect two executable motor plans. It is easy for the discussion to become semantic in the absence of clearly defined neural correlates for each putative process. Whether the goals hypothesis

proves to be true or not, it is not contested that there are signatures of value-based and perceptual decision-making in motor cortex. This has motivated increasing research interest in assaying decision variables in motor areas, which led to the experiment in Chapter 2. In the next section I turn my attention to one of these decision variables: prior expectation.

1.3. Prior Expectation

Our interpretation of the world around us depends on our prior beliefs about it. The scene in Figure 1.6 provides an intuitive example: interpretation is hard until we know the image shows a Dalmatian under a tree; after this, it is easy. Visual illusions highlight the extent to which perception is a process of applying existing beliefs to new data (Kersten and Yuille, 2003). Similarly, object recognition is speeded by a relevant visual scene that makes the object more likely (Bar, 2004; Enns and Lleras, 2008). Veridical expectations make our decisions faster and more accurate (Mulder et al., 2012).



Figure 1.6 This ambiguous scene becomes easy to interpret with information about what it depicts; see text. Image from Gregory (1970).

The combining of prior expectations and sensory evidence occurs in a Bayes-optimal manner (Kording and Wolpert, 2004). Our tendency to form predictions about the world lies at the heart of ideas about the brain as an inference machine (Helmholtz, 1867; Friston, 2012).

How is expectation represented in the brain? Neurons in superior colliculus (Basso and Wurtz, 1997, 1998) and lateral intraparietal cortex (LIP; Platt and Glimcher, 1999) fire more vigorously prior to a certain response being cued, the likelier that response is. These responses have been interpreted in the context of models of decision-making, and so I begin by discussing these.

1.3.1. Decision model predictions about role of prior expectation

Normative models of decision making were proposed to model the process of making a choice between two alternatives. The drift-diffusion model (DDM; Ratcliff, 1978) is the most prominent among other related models (Carpenter and Williams, 1995; Usher and McClelland, 2001; Brown and Heathcote, 2008). The drift-diffusion model represents decision making as an iterative process of repeated sampling from sensory evidence. The samples of sensory evidence are accumulated until there is enough evidence to select one of the two alternatives and make a decision.

For instance, consider a human subject faced with a popular psychophysical stimulus: a field of moving dots, with a small proportion moving either leftward or rightward amongst a majority moving randomly (Britten et al., 1992). These ‘random dot kinetograms’ (RDKs; Figure 1.8, top panel) require the subject to look at the stimulus and integrate evidence over time before making a left/right judgement. According to a drift-diffusion model, a decision variable represents the current state of accumulated evidence about the RDK. The decision variable starts at a level equidistant between two bounds which represent ‘decide right’ and ‘decide left’ (Figure 1.7). Every time sensory evidence is sampled, the decision variable changes, with sensory evidence consistent with leftward motion pushing the decision variable towards the ‘decide left’ bound and vice versa. Over time, the decision variable will

reach one of the two bounds, with the mean rate of the rise reflecting the strength of sensory evidence. The bound represents a criterion level of evidence which is regarded as sufficient to make a decision.

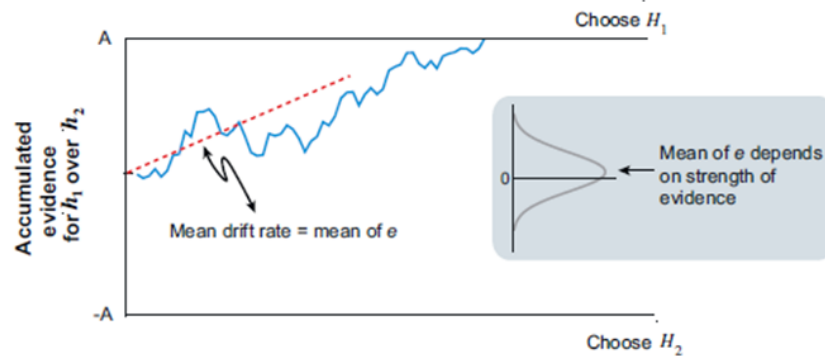


Figure 1.7 The drift-diffusion model implements a statistically optimal method for making a decision. The decision variable (blue line) starts in between the bounds representing the two hypotheses (H_1 and H_2). Evidence is repeatedly sampled and the decision variable is updated. In this example, the sensory evidence favours H_1 and so the activity of the decision variable rises to the H_1 bound, with mean rate of rise (red dashed line) representing the strength of the sensory evidence. When the blue line reaches a bound, the corresponding hypothesis is accepted as true. The time taken for this process, plus some non-decision time, equals the reaction time. Figure adapted from Gold and Shadlen (2007).

The drift-diffusion model is equivalent to the Sequential Probability Ratio Test (SPRT; Wald, 1945), the statistical test developed to make a decision at the desired accuracy with the fewest samples of evidence possible. The drift-diffusion model naturally replicates some features of human perceptual decisions e.g. that more difficult decisions – where the evidence is weaker – take longer. Indeed, the drift-diffusion model is very good at replicating reaction time distributions in a range of tasks (Ratcliff, 1978, 2002), including the RDK task (Ratcliff and McKoon, 2007).

There are other, competitor models to the drift-diffusion model, which share the basic feature of evidence-determined activity racing to a bound. In some of these models, rising activity reflect momentary rather than accumulating evidence (Carpenter and Williams, 1995; Cisek et al., 2009), in others separate units race to meet threshold (Carpenter et al., 2009), whilst others include a leak term so that integrated evidence does not

accumulate indefinitely (Usher and McClelland, 2001). All perform well at modelling reaction times, so I have focussed on the drift-diffusion model because it is the most widely used.

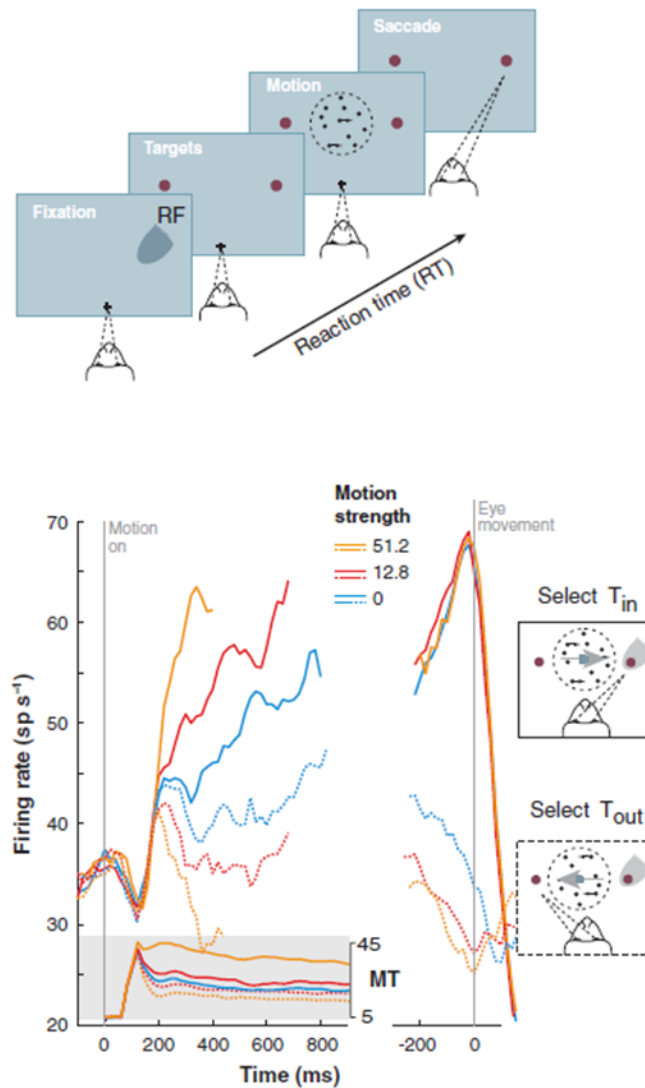


Figure 1.8 A random dot kinetogram task and corresponding activity in MT and LIP neurons. Top panel: the random dot kinetogram (RDK) is a field of noisy moving dots which moves either leftward or rightward. In this task, the direction of RDK motion indicates which of the two targets the monkey should make a saccade towards. RF indicates the neuron's receptive field corresponds to the right target. Bottom panel: the activity of the LIP neuron shows ramping activity in response to the RDK, with a steeper ramp when the signal-to-noise ratio of the RDK motion is higher (solid lines). When RDK motion favours the alternative target (dotted lines), the ramping is downwards, with, again, a rate depending on strength of stimulus motion. The grey shaded area shows the response of motion-sensitive neurons in visual cortex: these appear to represent the current strength of sensory evidence, in contrast to LIP neurons, which appear to represent evidence integrated over time. Figure from Gold and Shadlen (2007).

The drift-diffusion model made the leap from theoretical model to putative mechanism with the discovery of neurons in macaque lateral intraparietal cortex (LIP) that appeared to encode the decision variable during RDK tasks (Shadlen and Newsome, 2001a; Roitman and Shadlen, 2002). As in the drift-diffusion model, LIP neurons show ramping in activity level up to a fixed firing rate; their activity predicts behavioural choice (Figure 1.8; bottom panel). Stronger RDK motion leads to a steeper rise in activity and a shorter reaction time.

How can prior expectation be incorporated into decision models? Prior expectation makes reaction times faster and more accurate (e.g. Carpenter and Williams, 1995). There are two possible ways to alter a rise-to-threshold model to reflect this (Carpenter and Williams, 1995; Ratcliff and McKoon, 2007; Domenech and Dreher, 2010; de Lange et al., 2013; Summerfield and de Lange, 2014). Firstly, an increased expectation of one hypothesis in favour of another could bias the starting position (Figure 1.9). The decision variable would thus start closer to one bound than another, meaning less evidence would need to be accumulated to reach the decision. I refer to this as ‘baseline bias’. Alternatively, there could be a bias to the drift rate for the relevant hypothesis, so that accumulation progressed in bigger steps, making it easier to reach threshold. This is analogous to adjusting the gain on the sensory evidence for that hypothesis. I refer to this as ‘gain increase’.

Either modulation would have the effect of reducing reaction times, although they make different predictions about distributions (Carpenter and Williams, 1995) and error trials (Mulder et al., 2012). (There is also a third possibility: an increased prior expectation could reduce the height of the bound for the relevant hypothesis. This has the same effect as a baseline bias – less evidence need be acquired – and is behaviourally

indistinguishable; thus these two possibilities are usually considered together.)

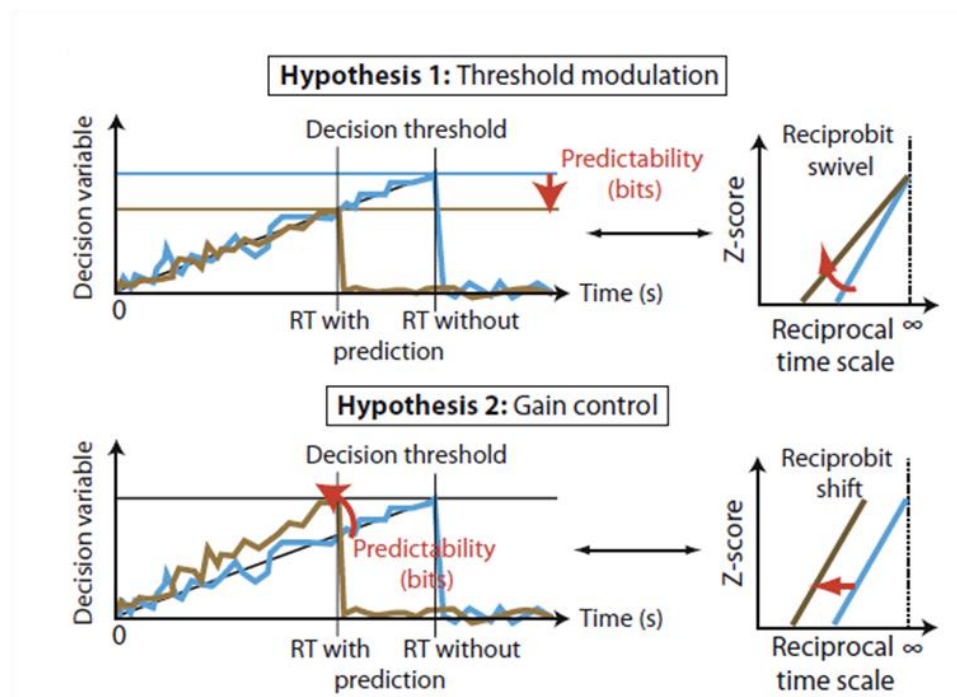


Figure 1.9 Prior expectation could be incorporated into a drift-diffusion model in one of two ways: by reducing the distance between the two thresholds or increasing the gain on the integration of sensory evidence. A reduced interthreshold distance is shown in the top panel; an increased gain on sensory evidence is shown in the bottom panel. Blue line: lower expectation; brown line: higher expectation. The reduced inter-threshold distance shown in the top panel could be achieved by a lowered top threshold or an elevated starting point; these have equivalent effects on reaction times. The plots on the right hand side show the corresponding effect on reaction time distributions, which are plotted as 'reciprobit' plots (Carpenter and Williams, 1995) in order to transform them to be linear. Figure from Domenech and Dreher (2010).

For a simpler model than the drift-diffusion model, analysis of reaction time distributions has repeatedly (Carpenter and Williams, 1995; Domenech and Dreher, 2010; Forstmann et al., 2010) favoured the first explanation: prior expectation increases the height of the starting bound and not the drift rate. For drift-diffusion models themselves, the results are considered to support a change to starting bound to the exclusion of (Simen et al., 2009; Mulder et al., 2012) or as a greater effect than (Ratcliff and McKoon, 2007) a drift rate change.

Thus, the quantities in rise-to-threshold models have intuitive correlates (Carpenter and Williams, 1995). The starting level reflects the level of belief about a hypothesis: the prior expectation. The rate of rise is a noisy representation of the level of evidence for a particular hypothesis, or for one hypothesis over the other, depending on the model (Reddi et al., 2003). The top threshold represents the amount of evidence for a decision, and modulations to this implement a natural speed-accuracy trade-off (Reddi and Carpenter, 2000). The models are Bayesian, in that they combine prior belief with new evidence (Summerfield and de Lange, 2014).

So behavioural modelling predicts prior expectation acts via a baseline bias and not gain increase. Is this prestimulus bias found in neural firing? This idea was lent credence by the finding that superior colliculus neurons (Basso and Wurtz, 1997, 1998) and LIP neurons (Platt and Glimcher, 1999; Churchland et al., 2008) fire at a lower level prior to stimulus if an alternative is less likely, but ultimately reach the same level prior to movement, exactly as would be predicted by the theory above. In a more direct test of the hypothesis, Rao et al (2012) found an effect of prior expectation that waned with time, exactly as might be expected for a baseline offset.

However, Hanks et al (2011) found the opposite: the bias to neural activity afforded by a prior expectation grew stronger over time. This might be due to task differences: Rao et al showed a stimulus of fixed duration, whilst Hanks used a reaction time task in which it might be more advantageous to weight the prior more strongly as time proceeds. In other words, taken together, these studies might imply the implementation of priors is dynamic according to task demands. A third study (Gold et al., 2008) used an alternative method to compare the

hypotheses by looking for neural correlates of trial-to-trial random biases. They failed to find any; this could have been due to lack of power.

1.3.2. High-level representations of expectation

Whilst neurophysiological studies on the coding of prior expectation in LIP have been inconclusive, imaging studies have sought to look more broadly for representations of expectation and identify ‘higher-level’ brain areas which might be the source of expectations induced by task context. A number of studies have used the decision modelling introduced in the previous section to identify relevant areas.

Domenech and Dreher (2010) found that, in a shape identification task, prior expectation based on a predictable order to stimuli modulated reaction times in a manner consistent with a threshold modulation in a decision model. Their fMRI analyses found that the anterior cingulate cortex was involved in this threshold modulation, whilst dorsolateral prefrontal cortex accumulated sensory evidence relating to the correct choice.

Forstmann et al (2010) induced an expectation about the direction of an RDK. Modelling the behavioural data using a simple decision model (Brown and Heathcote, 2008), they found that simply fitting the accumulator distance-to-threshold captured the data better than any other combination of parameter fitting. Using this parameter as a covariate in their fMRI analysis, they found that activity in the orbitofrontal cortex, hippocampus and putamen varied with the model-estimated distance-to-threshold. They suggested that a prior expectation led to activation of a selected corticostriatal loop, and thus began an action selection process.

Rahnev et al (2011) attempted to dissociate perceptual activity from the corticostriatal motor response observed by Forstmann et al (2010) by

withholding the stimulus-response mapping until after the decision in a similar RDK task. They used their behavioural data to compute signal detection theoretic quantities (Green and Swets, 1966) and then compared the bias from this computation to BOLD activity. They found increased activity in dorsolateral prefrontal cortex, and increased recurrent connections between this area and motion-sensitive area MT+.

Mulder et al (2012) used the drift-diffusion model to model the effect of prior expectation on an RDK task and again found that prior expectation changed the baseline of the evidence accumulation process. They found a variety of frontal and parietal regions covaried with this model baseline.

These studies and others (Shulman et al., 1999; Preusschhof et al., 2009) have found a wide variety of brain areas whose activity co-varies with levels of expectation and shows a prestimulus bias in response to expectation (Esterman and Yantis, 2010). This suggests expectation signals are represented widely around the brain in a task-specific manner. However, one finding that has been remarkably consistent across paradigms and models is that reaction time changes in response to expectation are consistent with a baseline alteration to the evidence accumulation process, in preference to the rate or gain of the evidence accumulation itself.

1.3.3. Expectation in sensory areas

Random dot kinetograms are designed to activate neurons in motion-sensitive visual cortex area MT (Britten et al., 1992). How does an expectation that a stimulus is coming affect this and other early sensory areas? The prediction made by the models outlined above is that early sensory areas should not show a modulation by expectation, because this is already represented as a baseline offset in the activity of the parietal, evidence-accumulating neurons. If it were true that stimulus likelihood

had been incorporated into the sensory representation, then expectation representation at the parietal level could not be by baseline offset alone – the expectation would be repeatedly incorporated with each sample of evidence and thus would also affect the drift rate.

Is this prediction – of expectation-invariant sensory representations – borne out? Rao et al (2012) found the pattern predicted by decision theory: a modulation by expectation in LIP but not MT. However, this finding is at odds with fMRI studies which have seen a modulation by expectation in early visual areas (Kok et al., 2012) and MT specifically (Shulman et al., 1999; Schlack and Albright, 2007). Indeed, visual context can drive activity in non-stimulated areas of V1 (Smith and Muckli, 2010).

Fitting these results into rise-to-threshold models would necessitate rejecting the hypothesis that, in biological implementation, expectation acts strictly on baseline alone. The Leaky Competing Accumulator (LCA; Usher and McClelland, 2001), a variant of the drift-diffusion model which is considered more biologically plausible, includes lateral inhibition between competing decision units. Lateral inhibition obliterates the distinction between baseline bias and drift rate, because, for example, a positive bias to one unit will inhibit the evidence accumulation of competing units. However, other lines of evidence also challenge the rise-to-threshold model account.

A consistent finding from the EEG literature is that unexpected stimuli evoke a ‘mismatch’ negativity; a larger amplitude event-related potential compared to expected stimuli (Näätänen et al., 1978, 1989). As well as attenuated responses to expected stimuli, responses to repeated stimuli are also suppressed (Todorovic and de Lange, 2012). Clearly, this dampening of expected stimuli does not fit well with the drift-diffusion model account, which hypothesises that a stronger expectation drives the amplitude of responses.

What is the explanation for the apparent contradiction? Predictive coding (Mumford, 1992; Rao and Ballard, 1999) offers a more complex and dynamic account of the role of expectations or Bayesian priors, rooted in the recurrent nature of neural connections (Figure 1.10). Expectations are fed down the hierarchy of brain areas whilst sensory input is passed upwards. At each level of processing, sensory areas compare their input and the prediction to compute sensory prediction errors and errors are progressively minimised to arrive at the best percept.

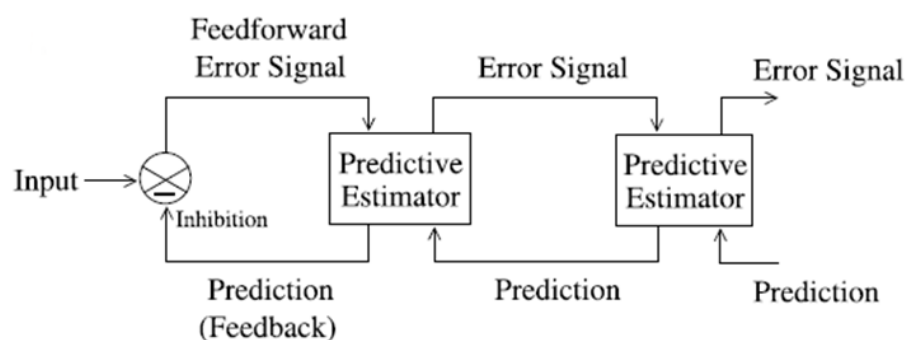


Figure 1.10 The integration of new information and existing beliefs, in a hierarchical network for predictive coding. Feedforward pathways carry prediction error signals whilst feedback pathways carry predictions. At each level of the hierarchy, a representation of the current input is maintained and compared to the prediction carried in the feedback connections to calculate a prediction error signal. The prediction error signal is then passed to the next level in the hierarchy via the feedforward connections. Figure from Rao and Ballard (1999).

Predictive coding has the power to explain previously puzzling ‘extra-classical’ receptive field effects (Rao and Ballard, 1999). Furthermore, a number of fMRI studies have found visual cortex activity consistent with predictive coding. Most strikingly, Kok and de Lange (2014) showed that when visual input is expected due to the context and also present, responses in V1 are suppressed, but responses are enhanced when there is the same expectation but no corresponding input. This fits well with the idea of a sensory prediction error. Furthermore, expectations increase top-down connectivity to visual areas (Summerfield et al., 2006), and when expected and observed perceptions diverge, bottom up connectivity increases (Summerfield and Koechlin, 2008).

Thus, there is now a body of evidence to suggest that expectations are represented in early sensory areas in a manner than cannot be accounted for by drift-diffusion modelling but is well captured by predictive coding. The rise-to-threshold model account, which influenced the design of experiments presented in this thesis, continues to be useful to understand the integration of evidence and existing beliefs in higher areas, but the purely feed-forward flow of information fails to capture the role of expectation in sensory areas.

1.3.4. Expectation in motor representations

Long-standing evidence has shown motor cortical neural activity in the period before having to make a movement (Wise et al., 1983; Riehle and Requin, 1989; Crammond and Kalaska, 2000). However, whether and how the motor cortex responds to information about the likelihood of having to make a particular movement is less well-studied. The affordance competition hypothesis (Cisek, 2007) would suggest that motor cortex is constantly being updated with decision-making signals and thus we might expect signals about expectation to move to be reflected in motor cortex. In the oculomotor system, superior collicular neurons reflect cue likelihood (Basso and Wurtz, 1997, 1998).

Partial information about an upcoming movement shapes population-tuned neural activity (Bastian et al., 2003). The likelihood of having to make a movement at a particular time also alters lateral field potentials (Roux et al., 2006). Corticospinal excitability as measured by transcranial magnetic stimulation appears to reflect quantities related to likelihood, such as surprise (Bestmann et al., 2008) and time elapsed without a cue (van Elswijk et al., 2007). Finally, lateralisation of MEG activity over motor cortex reflects a predictive cue (de Lange et al., 2013). Thus there is a variety of evidence from different modalities that the information carried by predictive cues may be reflected in motor cortex. A series of

experiments into the effect of the movement likelihood on corticospinal excitability were the subject of Chapter 2.

1.3.5. Differentiating attention and expectation

‘Everyone’, William James famously wrote, ‘knows what attention is’. A similar claim might be made for expectation: we have a rich and immediate introspective sense of what it means to expect a forthcoming stimulus, or for expectation to be violated (‘Surprise!’). – *Summerfield and Egner (2009)*

Despite everyone knowing what attention and expectation are, it is sometimes hard to differentiate them. For instance, the experimenter cues a subject that a target is likely to appear at a particular location, and they are faster and more accurate (e.g. Carpenter and Williams, 1995). Are neural changes a result of probabilistic expectation, spatial attention, or both?

There is almost exact overlap between the putative implementations of expectation and attention (gain control, baseline offset, etc; Summerfield and Egner, 2009), and they have often been elided in discussions of cueing paradigms, particularly in the field of visual perception (e.g. Esterman and Yantis, 2010). However, there is evidence that modulatory neural responses to predictive information persist even when attention is diminished by sleep (Nakano et al., 2008), anaesthesia (Yaron et al., 2012), or disorders of consciousness (Bekinschtein et al., 2009).

An experiment that aimed to orthogonalise the likelihood of a stimulus being relevant (as a proxy for attention) and the likelihood of a stimulus occurring (expectation; Wyart et al., 2012b), found that attention had effects consistent with a model which suppressed noise in signal processing. In other words, attention facilitated signal-processing most for stimuli in which the signal was present. By contrast, expectation

(likelihood) was best modelled by a baseline offset to signal processing. This would thus have the strongest effect on stimuli in which the signal is absent. Within a predictive coding framework, attention appears to increase the precision of prediction errors (Kok et al., 2012; Jiang et al., 2013). In short, expectation and attention seem to have distinct neural mechanisms.

1.3.6. Summary and relevance

Expectation signals are found in a huge variety of experiments looking at different brain areas and using different techniques. In particular, I have highlighted lines of evidence that have argued for expectation as a baseline bias, which was relevant to the design of experiments in Chapters 2, 3 and 4.

The literature on expectation is confused by an absence of a unifying definition or paradigm characterising expectation. For instance, some authors have considered it important that expectations should be implicitly developed (Kok et al., 2013) i.e. acquired unconsciously from task statistics rather than cued. Some studies have used a cue which is likely to come at one of several timepoints (Roux et al., 2006; van Elswijk et al., 2007) and assumed this temporal expectation represents expectations more generally. Some literature has used a task in which participants have to make an 'A/not A' judgement about each stimulus. The template A stimulus is assumed to represent an expectation, but, as test stimuli are equally likely to be A as to be not A, this is clearly not a representation of likelihood (Summerfield and Koechlin, 2008). Furthermore, some of the literature I have cited above used cues which were perfectly predictive of later stimuli (Shulman et al., 1999; Schlack and Albright, 2007); this could conceivably be represented differently to a probabilistic mapping. I have already discussed the common elision of expectation and visual attention in the scene perception literature.

In the experiments in this thesis, I used explicit cues that accurately informed subjects about the likelihood of one stimulus over another. This method is simple, clearly modulates subjects' belief about stimulus likelihood and, when trial by trial cues are used, allows the use of interleaved designs (see Chapter 2 for more).

1.4. Motor variability

In 1990, Larry Bird made 71 consecutive free throws. While a remarkable feat, one wonders why he missed the 72nd? Why could he not simply do what he had done the last 71 times? – *Churchland et al (2006)*

Movements are variable: it is rarely possible to perfectly repeat them. This is such a ubiquitous feature of human movement that we take it for granted, but why is it that we cannot do again what we have done before, even when, as the quote above illustrates, the potential gain is high?

In this section I review ideas about the source of motor variability, but first begin by describing models which have centred around controlling movement variability as the fundamental optimisation of the motor system.

1.4.1. Motor variability as a limiting factor in motor control

The control of motor variability has been a critically important concept in formulating hypotheses about online motor control. Under these hypotheses, variability is characterised as a limiting factor for the motor system. Optimal feedback control (Todorov and Jordan, 2002) was based on the experimental observation that movement variability is typically lower in task-relevant versus task-irrelevant dimensions: the ‘uncontrolled manifold’ (Scholz and Schöner, 1999). For instance, endpoint distributions in pointing movements reflect the shape of the target (Lametti et al., 2007; Berret et al., 2011; Nashed et al., 2012). Older models of motor control typically involved an ideal trajectory being planned and then movement execution system attempting to execute this as closely as possible (Flash and Hogan, 1985; Harris and Wolpert, 1998), but this fails to explain the existence of the uncontrolled manifold, which

implies that online mechanisms retain knowledge of the goal of the movement and not merely the trajectory.

According to optimal feedback control, the motor system generates a feedback control law. This responds to feedback (e.g. proprioceptive information) intelligently, according to task demands (Todorov and Jordan, 2002). For instance, a control law for a reaching movement might be defined in terms of endpoint error. If the control law predicts that noise or a perturbation during the movement will deviate the endpoint enough to interfere with task goals then it is corrected for, but otherwise it is not. Thus variability accumulates in task-irrelevant dimensions. The elegance of this system is that details of the movement are postponed until execution, so that trajectory planning and motor execution become one and the same.

In common with older models (Harris and Wolpert, 1998), optimal feedback control models include a term for noise at the neuromuscular junction which is dependent on the mean size of the motor-neuronal signal and thus a speed-accuracy trade-off arises naturally. Effort is also penalised.

Optimal feedback control predicts a 'minimal intervention' principle, where deviations from a mean trajectory are corrected for only when they prevent task goals being reached. This prediction has held over a huge variety of experimental paradigms. If subjects experience a visual perturbation (via an offset introduced into the position of a cursor representing hand position) when moving towards a rectangular target, they correct more to deviations along the short axis than the long axis (Knill et al., 2011). This is also true when the perturbation is mechanical (Nashed et al., 2012).

Feedback corrections seem to be highly flexible in reflecting complex task demands. For instance, corrections take into account how hard the subject needs to hit the target (Liu and Todorov, 2007). Corrections also intelligently reflect obstacles in the task environment (Nashed et al., 2012). If subjects learn that a perturbation is temporary, they do not correct for it, but if the perturbation is permanent, they do (Figure 1.11) (Franklin and Wolpert, 2008). Corrections to perturbations reflect knowledge about the dynamics of a task even when these were learnt in the absence of perturbations (Wagner and Smith, 2008; Cluff and Scott, 2013).

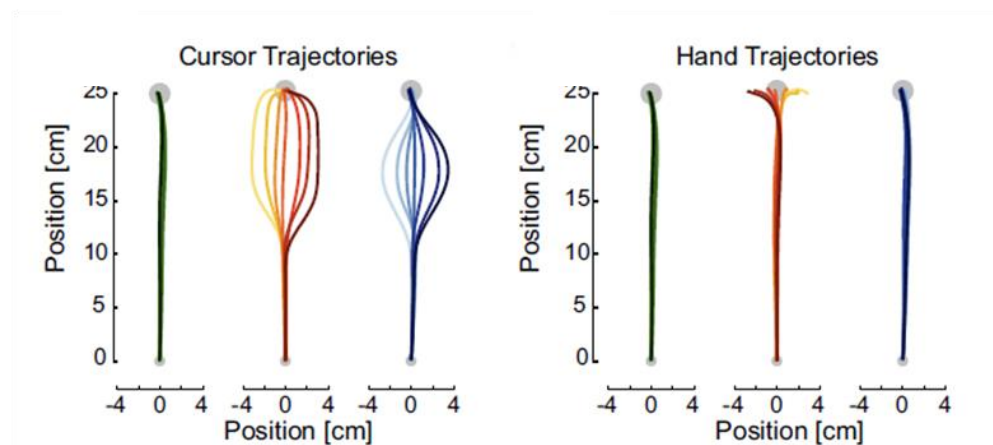


Figure 1.11 A visual perturbation is corrected for when it affects task performance, but not when it does not. Results from an experiment in which the subject is making a reaching movement with hand occluded and a cursor representing hand position. The experimenter sometimes induces a perturbation by falsely deviating the position of the cursor laterally (left panels). On some blocks, these perturbations are eventually reversed (blue lines) whilst on others they are not (yellow/orange lines). The hand trajectories (right panels) show that the subject adjusts their hand position to counteract the permanent perturbation (yellow/orange lines), but not the temporary one (blue lines). Figure from Franklin and Wolpert (2008).

Involuntary corrections to perturbation that are responsive to task demands begin within ~ 110 ms; in other words, within the class of response that is usually considered a 'reflex' (Pruszynski et al., 2008, 2011; Dimitriou et al., 2013). This applies even when the perturbation varies from trial to trial (Franklin and Wolpert, 2008). Similarly, if the task changes, feedback gains are updated within 100 ms (Dimitriou et al.,

2013), suggesting that optimal feedback laws are continuously recalculated. (However, in contradiction, another study failed to find trial-by-trial updating of feedback corrections (de Xivry, 2013)).

Corrections to perturbations are sensitive to task urgency (Crevecoeur et al., 2013) and are graded (Pruszynski et al., 2008) and present even for small perturbations (Crevecoeur et al., 2012). They peak in magnitude midway to the target (Liu and Todorov, 2007; Dimitriou et al., 2013), which is consistent with another optimal feedback control prediction: towards the end of the movement, there is an increased energetic cost to corrections, which begins to outweigh the benefit of a target hit. Optimal feedback control also accurately models coordination in two-handed movements (Diedrichsen and Dowling, 2009).

In further support of the idea that the motor system is optimised to deal with the detrimental effects of noise, subjects select their movements in a way that shows knowledge of their own endpoint distributions, and a rapid updating of this knowledge when this variability is increased (Trommershäuser et al., 2005).

In summary, optimal feedback control has been highly successful in accounting for feedback responses to experimentally-induced movement variability. However, the source of variability in natural movements is contested.

1.4.2. The source of motor variability

Optimal control theory (Todorov and Jordan, 2002) relies critically on the idea that motor noise is signal-dependent: that is, it scales with the size of the motor signal. Thus faster movements, requiring a larger signal, are more variable, imposing a natural speed-accuracy tradeoff (Harris and Wolpert, 1998). In these conceptions, the nervous system is critically limited by stochasticity in spiking and the motor system is optimised to

limit the impact of such noise on movement execution (Franklin and Wolpert, 2011). These theories have assumed that such noisy spiking is peripheral in origin (Jones et al., 2002) and has its detrimental effect primarily during action execution (van Beers et al., 2004).

However, the primary determinant of movement variability continues to be debated. Elsewhere, the importance of central sources of neuronal noise has been emphasised, whether sensory, motor, or computational (Churchland et al., 2006; Beck et al., 2012; Chaisanguanthum et al., 2014). Churchland and colleagues (2006) have argued that their dynamical systems view of motor cortex (Section 1.2.3) suggests motor planning involves complex optimisations, and variance in these would be expected to contribute to variance in motor execution. They found a trial-by-trial correspondence in variability in motor planning activity and reach velocity. They estimate that motor planning activity determines around half of movement variability.

Similarly, Beck and colleagues have argued that variability arises from approximations during inference: human limbs are too complex to be modelled accurately and our internal models of them are suboptimal (Beck et al., 2012). This would again predict that most variability arises during motor planning, when the models of limb dynamics are engaged. These ideas are in marked contrast to the optimal feedback control studies which have assumed that preparatory variability is negligible in a reaching task in humans (van Beers et al., 2004).

In short, the key determinant of movement variability continues to be debated, but there has been a renaissance in interest in the importance of motor planning in determining motor variability. Recent behavioural work, which I have built on in my PhD, has manipulated motor planning in order to determine the effect on motor variability and this is described in the next section.

1.4.3. Limited resources as a cause of motor variability

Multiple potential actions are planned in parallel (Section 1.2) and motor planning activity is considered a possible determinant of motor variability (Section 1.4.2). Does planning multiple actions have an effect on motor variability? When human subjects have to plan two possible movements rather than one, positional variability in the movement increases (Wijdenes et al., 2016). In this experiment, subjects had to move to one of two possible targets; the correct target was either cued early relative to the movement (i.e. only one plan need be maintained) or late (i.e. both plans need to be maintained). A control condition in which the cued target jumped showed that subjects did indeed maintain the expected number of plans, because there was a reaction time cost if the target jumped in the two-plan condition but not the one-plan condition.

Why should increasing the number of motor plans increase motor variability? The finding suggests that there is a limited resource in motor planning. Under a limited resource scheme, all motor plans must be represented by a shared neural resource. As the number of plans increases, the quality of each plan diminishes. Behaviourally, variability increases.

Limited resource models have had notable success in describing visual working memory. The precision of representations in visual working memory (as shown by line length discrimination between a presented stimulus and a remembered one (e.g. Palmer, 1990)) decreases as the number of items in working memory increases. Visual working memory was therefore posited to be item dependent, with acuity of representations dropping off sharply after three to four items (Luck and Vogel, 1997).

However, Bays and Hussain (2008) showed that when the discrimination task is made more difficult, there is no sudden drop in acuity after three or four items but instead a gradual decrease as the number of items increases. The relationship between number of items and precision of representation can be described by a power law. The results support an alternative hypothesis to a fixed item limit: that there is sharing of a limited resource which determines the precision with which items in working memory are internally represented. The limited resource is shared dynamically, with quick updating in response to saccadic gaze changes and attention-capturing cues (Bays and Husain, 2008).

Limited resource models have also been proposed to explain finite self-control (Muraven and Baumeister, 2000) and now motor variability (Wijdenes et al., 2016).

What could the biological substrate of a limited resource constraint be? The summed level of neuronal spiking in a population offers a candidate mechanism. Many neuronal populations have been observed to exhibit divisive normalisation, where neuronal inputs are scaled by the overall activity of the population such that the mean level of spiking is always similar (Carandini and Heeger, 2012). Divisive normalisation was originally proposed to explain receptive field responses in striate cortex (Heeger, 1992) and has since been used to explain receptive fields in area MT (Britten and Heuer, 1999), multiple object representation in inferotemporal cortex (Zoccolan et al., 2005), value coding in lateral intraparietal cortex (Louie et al., 2011), retinal photoreceptor adaptation (Normann and Perlman, 1979; Carandini and Heeger, 2012) and olfactory population codes in *Drosophila* (Olsen et al., 2010). Indeed, divisive normalisation has been proposed as a canonical computation in neural circuits (Carandini and Heeger, 2012). Divisive normalisation provides a mechanism for the context-dependence of value coding in LIP (Louie et

al., 2011) and has been proposed to explain rationality violations in economic choices (Louie et al., 2013). Divisive normalisation is the basis of the limited resource model which has been successful in modelling errors made in visual working memory tasks (Bays, 2014).

Divisive normalisation's ubiquity makes it a good candidate for a limited resource model of motor planning. In a normalised neuronal population, maintaining a representation of two possible reach targets rather than one leads to a reduced signal-to-noise ratio for each option, because the spiking in the population is shared across the plans. Indeed, normalisation is needed to explain many of the behavioural phenomena introduced in Section 1, and is incorporated into the affordance competition model (Pastor-Bernier and Cisek, 2011). Normalisation explains why there is a response time cost to preparing more than one movement (Cisek, 2007), and why there is an inverse relationship between the number of potential movements and the neural activity of each one (Basso and Wurtz, 1998; Cisek and Kalaska, 2005).

So, divisive normalisation might be the key to explaining many of the behavioural consequences of planning two movements, including an increased variability in execution. I have built on the limited resource ideas of Wijdenes and Bays (2016) in experiments in Chapters 3 and 4.

1.4.4. The advantages of motor variability

So far, I have discussed motor variability as a limitation; that is, in the context of hypotheses that posit that variability arises inevitably from noise in neuronal firing and is a limiting factor on the motor system (Harris and Wolpert, 1998). However, research in a different context has emphasised motor variability as a critical substrate of exploratory behaviour and motor learning (e.g. Therrien et al., 2016). There is

growing evidence that behavioural variability is sometimes desirable and gratuitously introduced.

In support of the link between variable behaviour and learning, it has been found that rats which are trained to be more variable are capable of learning a complicated motor sequence which less variable rats were not (Grunow and Neuringer, 2002). These findings were extended to humans in a study showing that subjects who are naturally more variable in a particular aspect of a baseline reach task were faster to learn to modify behaviour when a subsequent reach task rewarded that aspect, even though subjects were not aware exactly what was being rewarded (Wu et al., 2014). Furthermore, this finding was not limited to reward-based operant learning, which is believed to be driven by reward prediction errors in the basal ganglia (Glimcher, 2011), but also held in an error-based learning task; these are hypothesised to rely on cerebellar sensory prediction errors (Izawa et al., 2012).

An eloquent description of how variable behaviour might be generated deliberately as a substrate for learning comes from the literature on songbird song generation. Young songbirds have highly variable song production, but, as the birds age, they learn to produce a stereotyped song that mimics a tutor bird (Ölveczky et al., 2011). Interestingly, when a songbird is caged with a potential mate, it ceases to produce a variable song and produces its best attempt at a stereotyped song, but the variability in singing resumes once the mate is removed (Woolley et al., 2014). Neuronal recording studies have elucidated the spiking that drives this behaviour (Ölveczky et al., 2011) - the spiking itself drops in variability - and found that lesioning a basal ganglia-to-forebrain output tract obliterates the singing variability and produces a stereotyped song (Ölveczky et al., 2005). There is convergent evidence that dopamine is

also involved in producing behavioural variability in rats (Pesek-Cotton et al., 2011).

In songbirds, when there is a high possibility of reward – a mate – variability is temporarily halted. Is the same true in humans? Yes: human subjects increase their motor variability in response to an absence of reward (Pekny et al., 2015) or a punishment (Galea et al., 2013), or in a task that relies on reward-based rather than sensory feedback (Izawa and Shadmehr, 2011). Similarly, saccades in monkeys are less variable after a reward (Takikawa et al., 2002). These effects apply to both overall movement direction (Pekny et al., 2015) and low-level movement parameters (Takikawa et al., 2002; Galea et al., 2013). There is also a graded response: when the chance of getting a reward is lower, baseline variability increases.

Furthermore, the variability increase in response to absence of reward is obliterated in patients with Parkinson's Disease, suggesting a basal-ganglia dependent process, as in songbirds (Pekny et al., 2015). Similarly, dopamine antagonists obliterate the increased motor variability observed in response to financial punishment (Galea et al., 2013). The idea that variability increases in the absence of a reward fits into an idea from reinforcement learning of a so-called explore/exploit dilemma: organisms must trade-off between exploiting a current reward-giving option versus looking for a better one (e.g. Daw et al., 2006).

In addition to allowing learning and exploration, variability in itself might offer an advantage. Game theory says that randomness can be useful in outwitting opponents in predator-prey relationships (Neumann and Morgenstern, 1944). Randomness's role in exploration might even make it a substrate for higher-level processes like creativity (Carpenter, 1999).

1.4.5. Summary and relevance

The motor system is optimised to limit the effects of variability, and recent ideas suggest this variability might arise from limited neural resources during motor planning. In Chapters 3 and 4, I hypothesise that prior expectation plays a role in dividing up these resources, and test this idea experimentally.

1.5. Transcranial Magnetic Stimulation

Transcranial Magnetic Stimulation (TMS) is a neurostimulatory technique which can be used to assay action planning. I begin with a technical discussion of TMS before explaining its advantages in studying motor planning. In TMS, a coil-shaped electromagnet is placed on the scalp of an awake human. The electromagnet generates a magnetic field for a fraction of a millisecond, and so induces a current in and briefly excites underlying neural tissue (Barker et al., 1985). Whilst TMS can be used repetitively to study brain plasticity, or in paired pulse protocols to assay inhibitory circuits, I will focus the technical discussion on the type of TMS I used in the experiments in Chapter 2: single pulse TMS to hand area of motor cortex using a figure-of-eight coil in a postero-anterior (PA) orientation relative to the subject.

1.5.1. The motor evoked potential

A TMS pulse induces current flow parallel to the surface of the coil, with little radial current spread (Di Lazzaro and Rothwell, 2014). A figure-of-eight coil, in contrast to earlier circular coils, causes current summation across the two halves so that stimulation is strongest in the centre.

TMS excites the axons, rather than cell bodies, of cortical neurons. We know this because certain stimulation properties of TMS to motor cortex match those of large diameter myelinated peripheral fibres (Peterchev et al., 2013). Axons are excited by a potential difference along their length, and so TMS activates axons differentially according to their arrangement relative to the induced current (Salvador et al., 2011), with axons that bend out of a uniform electric field being most susceptible to stimulation (Amassian et al., 1992). Because of these properties, stimulation threshold varies with coil orientation, such that holding the coil in a PA orientation gives the lowest threshold for an evoked potential (Di Lazzaro

et al., 2001). This orientation is therefore commonly used, including for the experiments in this thesis.

TMS to motor cortex causes a muscular twitch contralaterally and an evoked potential which can be read out using electromyography (EMG) with surface electrodes. The amplitude of this motor evoked potential (MEP) is the metric of interest in many experiments using TMS.

Recordings from the surface of the spinal cord in awake humans have been used to understand the waves of neural discharge that sum, via motoneurones, to contribute to the MEP. Low intensity TMS evokes the I1 ('indirect 1') wave (Di Lazzaro et al., 1998a). This is a single discharge which is thought to originate from presynaptic activation of the pyramidal tract (hence the term 'indirect'). In support of this idea, I1 waves are sensitive to cortical excitability (Di Lazzaro et al., 1998b) suggesting a cortical origin. The latency of this wave is 1 ms longer than a direct wave; these will be discussed below.

Medium intensity TMS evokes further waves after the I1, this time as a volley, known as late I waves (Di Lazzaro et al., 1998a). They are believed to originate from neurons discharging at about 600 Hz. It is these waves that are depressed by activation of GABAergic circuits in intracortical inhibitory protocols (Di Lazzaro et al., 1998c).

The I1 wave may have a different origin to later I waves. Whilst I1 is believed to be monosynaptic, it has been proposed that late I waves arise from axons which synapse onto bursting neurons that in turn drive pyramidal tract neurons (Di Lazzaro and Rothwell, 2014).

High intensity TMS evokes further late I waves, and a D ('direct') wave (Di Lazzaro et al., 1998a). D waves are believed to originate from direct excitation of pyramidal tract neurons in the subcortical white matter,

some distance from the cell body. A given axon can produce both D and I waves (Patton and Amassian, 1954).

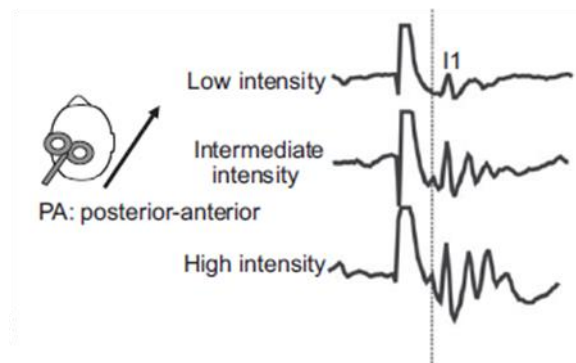


Figure 1.12 The waves evoked by PA TMS depend on the intensity of stimulation. Low intensity stimulation evokes an I1 wave. Intermediate intensity stimulation evokes an I1 wave and later I waves. High intensity stimulation evokes the I waves and a D wave (dashed line indicates timing of D wave). Figure from Di Lazzaro and Rothwell (2014).

Monophasic stimulation is typically used for TMS because it produces a simpler and more consistent pattern of cortical outputs than biphasic stimulation (Di Lazzaro et al., 2001). TMS-evoked responses are idiosyncratic: for example, it has been shown for AP TMS that the efficiency of late I wave recruitment varies between subjects (Hamada et al., 2013).

Importantly, voluntary contraction (both of relevant muscles (Di Lazzaro et al., 1998b) and non-relevant muscles (Andersen et al., 1999)) amplifies MEPs – the amplitude of evoked waves can be 50% higher than at rest. Thus MEP amplitude reflects the excitatory state of the pyramidal tract at the time of stimulation, and this is the basis of its use in studying action preparatory processes during the reaction time.

1.5.2. The MEP as an assay of competing action plans

TMS is highly temporally specific, allowing high-resolution sampling of corticospinal excitability within the human reaction time. This sampling can be directed to different effectors (index fingers, little finger, foot, etc). An early finding was that MEPs recorded from the effector during the

reaction time on a simple reaction time task (subjects respond to a cue with a fixed action) increase in amplitude, decrease in latency (Rossini et al., 1988) and increase in probability of being evoked (Starr et al., 1988) in the ~50 ms before EMG activity. Subsequently, it was found that the effector to move showed this pre-movement MEP facilitation in a variety of tasks whilst the corresponding, unchosen effector showed a suppression (Leocani et al., 2000). Thus it appeared that MEPs offered an insight into a suppression of the ipsilateral hemisphere prior to movement.

Inhibition of the ipsilateral motor cortex depends on the nature of the movement (Liepert et al., 2001) and the effectors involved (Sohn et al., 2003) and its mechanism is likely corticocortical (Weiss et al., 2003). Inhibition of the non-selected effector is context-dependent: it is stronger when movements are both upper limb than when one is upper limb and the other is lower limb (Labruna et al., 2013). Regarding excitation of the contralateral motor cortex, MEP excitability increases with an advance cue that fully or partially specifies the nature of the movement but not an uninformative cue (McMillan et al., 2004).

Thus it seems that MEPs offer a read-out of the competitive processes leading to one hemisphere being excited and the other suppressed. The idea that MEPs reflect dynamic competition was demonstrated particularly by the study by Michelet et al (2010) discussed in Section 1.2.1 (Figure 1.3) which showed the wrong effector initially being activated during a difficult decision. A number of studies have since used TMS to explore competition between action plans and find MEP correlates of factors which bias behaviour.

MEPs are facilitated by a biasing reward (Klein et al., 2012). Increasing temporal expectation also increases MEP amplitude (van Elswijk et al., 2007). MEPs reflect quantities related to the information signalled in the

task, like entropy and surprise (Bestmann et al., 2008). MEPs reflect the biomechanical cost of a decision: early MEPs are suppressed by high cost, perhaps reflecting a choice process, whilst later MEPs increase with high cost, perhaps reflecting a need for a greater EMG activation in more energetic movements (Cos et al., 2014).

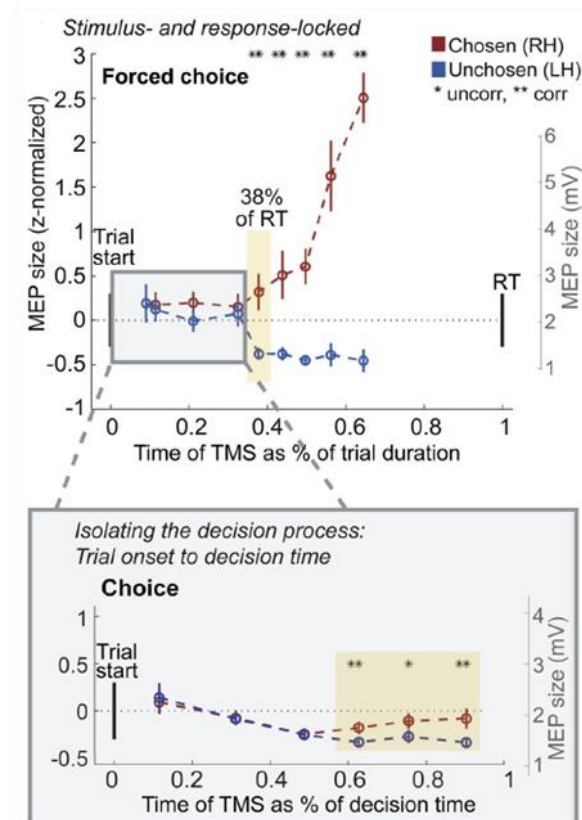


Figure 1.13 MEPs are different in the chosen versus unchosen effector during the decision time. In this experiment, the time for the decision was isolated using a comparison to control trials in which there was no choice process. Main panel shows the separation of MEPs in the chosen and unchosen effector across the whole reaction time. Inset shows this activity in just the decision time. The significant difference between MEP amplitude in the chosen and unchosen effectors in this period indicates that MEPs reflect the ongoing decision process. Figure from Klein-Flugge and Bestmann (2012).

In support of the idea that MEPs reflect competition whilst the decision is still ongoing (as in a parallel processing schema (Cisek, 2007)), MEPs in the chosen hand are differentiated from those in the unchosen hand at timepoints too early to reflect action preparation (Figure 1.13) (Klein-Flugge and Bestmann, 2012). Similarly, a more difficult face categorisation decision leads to divergence of MEP activity earlier relative

to the response, suggesting a continuous read out of the decision (Hadar et al., 2016).

1.5.3. Inhibition of MEPs in the effector that moves

Prior to movement, MEPs are facilitated in the selected effector and suppressed in the non-selected effector (Section 1.5.2). A consistent finding from the TMS literature thus seems counterintuitive: MEPs in *both* the effector and the antagonist are suppressed relative to an intertrial baseline prior to movement (Hasbroucq et al., 1997; Touge et al., 1998; van Elswijk et al., 2007; Tandonnet et al., 2010). Indeed, MEPs are *more* strongly suppressed in the effector that will move (Figure 1.14) (Duque and Ivry, 2009). This suppression is followed by the facilitation of the selected effector which was discussed above.

The phenomenon has classically been investigated using ‘instructed delay’ tasks in which the participant knows which effector they will have to move but must wait for a fixed delay period before beginning the movement (Hasbroucq et al., 1997, 1999; Davranche et al., 2007; Duque and Ivry, 2009). These tasks show MEPs progressively diminishing if the delay period is short (0.5 s) but not if it is long (2.5 s). So this type of delay was believed to be associated with placing a temporary brake on motor structures prior to movement.

However, even in a task without a delay, there is nevertheless MEP suppression after the imperative to move (Greenhouse et al., 2015b). Furthermore, simple reaction time tasks (in which the same response is always cued) also cause MEP suppression (Touge et al., 1998; Duque et al., 2014; Greenhouse et al., 2015b), suggesting the process is not critically dependent on choice selection. MEP suppression also occurs regardless of whether the precue is informative or not (Duque and Ivry, 2009), or how likely the precue is to lead to a movement (Sinclair and

Hammond, 2009). In other words, MEP suppression before the ramping up of MEPs in the chosen effector prior to movement is a widespread phenomenon which is relatively robust to particular task parameters.




Figure 1.14 MEP suppression during the delay period affects all effectors and is strongest in the hand the subject expects to move. The figure shows left index finger MEPs at intertrial baseline (black bars) compared to in the delay period before movement was cued (TMS_{Delay}) and in the reaction time after movement was cued ($TMS_{Movement}$). When the cue was informative (left panel), MEPs are suppressed most if the cue indicates that the relevant finger will be used in the response ('Bimanual' and 'Left'). However, MEPs are also suppressed in the non-relevant finger ('Right'). There is also suppression even if the cue was uninformative (right panel). Figure from Duque and Ivry (2009).

To explain the stronger suppression of MEPs in the chosen effector, a dual process account of MEP suppression has been suggested (Duque et al., 2010). According to this, one process reflects the outcome of the decision process (as described in Section 1.5.2), which might include the selected response inhibiting alternatives under a 'winner take all' scheme. This causes the suppression of non-selected effectors. The process is sometimes called 'competition resolution'. This type of inhibition is upregulated when conflicting stimuli which cause confusion over which movement to make are expected (Klein et al., 2014).

The second process is termed 'impulse suppression'. The winning representation is inhibited most strongly, perhaps as part of the preparation for movement or to prevent premature initiation. Evidence

for this ‘hold your horses’ signal includes the finding that inhibition is stronger when responses are more complex (Greenhouse et al., 2015a).

In support of the idea that there are two distinct processes, H-reflexes, which reflect spinal cord excitability, are diminished concurrently with MEP suppression in the chosen effector but not with suppression in the non-chosen (Hasbroucq et al., 1999; Duque et al., 2010). Furthermore, disrupting prefrontal areas with repetitive TMS reduces suppression of MEPs in the non-selected effectors (Duque et al., 2012, 2013), whilst doing the same to dorsal premotor cortex reduces suppression of MEPs in the selected effector (Duque et al., 2012), perhaps implicating separate cortical mechanisms in these phenomena.

Despite the dual process account, there is an alternative proposed explanation for MEP suppression: the spotlight account (Greenhouse et al., 2015b). This draws on the idea that MEP suppression could be a mechanism for increasing the gain of future signals (Hasbroucq et al., 1997). According to the spotlight account, a single process leads to broad inhibition of MEPs, with, as in a spotlight, inhibition strongest at the selected representation, but also extending less strongly onto non-selected representations. It is proposed that this broad inhibition is mediated by the basal ganglia, and there is focussed disinhibition of a relevant thalamocortical channel to initiate movement (Duque et al., 2017). Under this scheme, inhibition enhances future movement, rather than temporarily halting it. This idea is appealing because it dissociates MEP and behavioural inhibition, which have often been assumed to be linked, despite an absence of evidence to support this (Duque et al., 2017).

1.5.4. Summary and relevance

MEPs offer a temporally precise read-out of decision processes taking place in motor cortex and this was the basis for its use for the experiments in Chapter 2. The separate inhibitory processes that influence the MEP became important in interpreting these experiments; current theories about these processes are described above.

Chapter Two: The representation of expectation in corticospinal excitability during motor planning

2.1. Introduction

Neurophysiological evidence suggests that, prior to movement, multiple motor plans are prepared and a competitive process between these plans leads to selection of the final movement (Cisek, 2007; Pastor-Bernier and Cisek, 2011). This chapter uses Transcranial Magnetic Stimulation (TMS) to ask how beliefs ('prior expectations') are incorporated into this competition between motor plans. Expectation has been characterised as a prestimulus bias in literature concerning theoretical models of decision-making (Carpenter and Williams, 1995; Ratcliff and McKoon, 2007; Domenech and Dreher, 2010) and neural activity consistent with this has been recorded in superior colliculus (Basso and Wurtz, 1998) and parietal cortex (Platt and Glimcher, 1999; Rao et al., 2012). The question of this chapter is whether expectation is represented in corticospinal excitability as a prestimulus bias, as predicted by this work. Data from three experiments are presented.

2.1.1. Assaying competition between motor plans

The affordance competition hypothesis (Cisek, 2007) was formulated to explain neurophysiological data showing that activity in premotor cortex reflects sensory information (e.g. colour) only when it is relevant to the decision about which movement to make (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011). According to the hypothesis, motor plans are formulated whilst the decision about which movement to make is still ongoing. The motor plans are then biased by all signals relevant to the decision from elsewhere the brain (e.g. parietal cortex and basal ganglia). The motor plans are characterised as being in a continuous space representing a motor parameter, typically reach angle

(Georgopoulos and Carpenter, 2015). The plans can interact with one another in this space, so that executed movements are a product of the competitive interactions between the motor plans. Evidence for the affordance competition hypothesis is reviewed in detail in Chapter 1.

TMS can be used to assay corticospinal excitability for a particular effector (e.g. right index finger). Because of its temporal specificity, and the ability to compare competing effectors, over multiple trials, motor-evoked potentials (MEPs) can be used to construct a picture of ongoing competition between motor plans. TMS has been used in this way to study competition between motor plans in a value-based decision (Klein-Flugge and Bestmann, 2012). Apparent changes of mind when stimuli are highly conflicting have also been tracked using TMS (Michelet et al., 2010). TMS has also been used to demonstrate that corticospinal excitability varies with ‘cognitive’ parameters such as surprise (Bestmann et al., 2008) and action cost (Cos et al., 2014), as would be expected if competition in motor activity were the result of all decision information available. In short, TMS permits an assay in awake humans of the dynamic competition between motor plans that is central to the affordance competition hypothesis.

2.1.2. Expectation as a prestimulus bias

Prior expectation (which here is a belief about the likelihood of various events) has often been investigated in literature on decision-making because our expectations, along with new sensory evidence, combine to determine the outcomes of our decisions. This literature has focussed particularly on perceptual decision-making, studied in an established paradigm, the Random Dot Kinetogram (RDK). An RDK is a noisy field of moving dots from which the subject must determine coherent motion and thus must be viewed over time to determine the motion direction.

Modelling human subjects' performance on these tasks using the drift-diffusion model (Ratcliff and McKoon, 2007; Simen et al., 2009; Mulder et al., 2012) or race-to-threshold models (Carpenter and Williams, 1995; Domenech and Dreher, 2010; Forstmann et al., 2010) consistently finds that expectation is best modelled as a decision bias that occurs prior to the process of accumulating sensory evidence up to a decision bound. Neural recordings consistent with the idea of prior expectation as a prestimulus bias have been made in intraparietal cortex (Platt and Glimcher, 1999; Churchland et al., 2008; Rao et al., 2012) and superior colliculus (Basso and Wurtz, 1998). See Chapter 1 for more.

2.1.3. Hypothesis and key experimental features

The aim of this chapter was to measure the influence of prior expectation on corticospinal excitability. The particular hypothesis, informed the idea that decision information is continually biasing motor plans, and by past work on decision models, is that there would be a bias, prior to evidence accumulation, where corticospinal excitability for the expected effector should be higher than in the non-expected effector.

The experiments in the chapter involved human subjects receiving single-pulse TMS to left motor cortex during a choice between a left index finger movement and a right index finger movement, with various levels of prior expectation that one movement would be more likely than the other. Subjects were cued to expect, for instance, that on some trials, the probability that the upcoming stimulus would go left vs right was 10:90, whilst on others it was 50:50. The stimulus used to indicate to participants which way they should move was a Random Dot Kinetogram (RDK; Britten et al., 1992). These noisy moving dot stimuli involve a perceptual judgement that requires evidence accumulation over time (see Results). They have been widely used in decision-making studies, both behavioural (e.g. Donner et al., 2009; Rahnev et al., 2011) and

neurophysiological (e.g. Roitman and Shadlen, 2002; Hanks et al., 2011), and so were used here to align this study with past literature.

MEPs were measured in right-hand first dorsal interosseous (FDI). Due to the hypothesis that a bias would be present before evidence accumulation, the key MEP timepoint of interest was that at the appearance of the RDK: after the prior belief about its likely direction had been seen by subjects but before evidence accumulation had begun. I term this the ‘stimulus onset’ timepoint. The experiments also included some trials with later timepoints in order to build up a picture of ongoing excitability during the decision.

A further goal of the experiments was to use behavioural data recorded to model the bias parameter using a decision model (the drift diffusion model; Ratcliff, 1978) and to test for a correlation between this parameter and corticospinal excitability. The model allows us to access a metric of prestimulus bias, which is hypothesised to be dependent on prior expectation, in a way that simple behavioural measures do not.

2.1.4. What this study adds

If the hypothesis was borne out, it would add to a growing body of evidence for competing motor plans being biased by decision information (Cisek, 2007), and show that this happens not just during the evidence-accumulation part of the decision but also earlier. Corticospinal excitability has proved a fruitful way to investigate biased competition in humans because of its temporal and effector specificity. To date, TMS studies in this context have examined the Posner cueing task (Michelet et al., 2010) and a value-based task (Klein-Flugge and Bestmann, 2012), but there has not been any use of the evidence-integration task that perceptual decision-making studies have used.

A number of TMS studies have added credence to the idea that expectation can be detected in corticospinal excitability. Firstly, Bestmann et al (2008) have suggested that entropy and surprise are reflected in corticospinal excitability. As with the experiments in this chapter, they used a probabilistic cueing task. However, the experiment focussed on probabilities tracked over the many trials, incorporating the influence of both expectation and past events, and used information-theory to model the quantities. Here I wanted to look at the representation of expectation in isolation.

Secondly, Klein et al (2014) found MEP suppression at stimulus onset was higher in a context in which ‘conflict’ (a stimulus with confusing flankers) was expected. This informed the hypothesis here that expectation is incorporated at the timepoint of interest. However, the putative process studied here is a different one: Klein et al propose that MEPs are suppressed in this experiment as part of a stronger ‘impulse suppression’ process (discussed further below and in Chapter 1). This is considered distinct from the process of competition between motor plans investigated in this chapter (Duque et al., 2010).

Thirdly, van Elswijk et al (2007) show that MEPs increase with increasing temporal expectation. Again, this is promising for the present hypothesis, but this temporal expectation is about ‘when’ not ‘where’ to move. This experiment only involved one movement and so there was no competition between motor plans.

In other modalities, Bastian et al (2003) recorded primate motor cortex neurons and showed a population code for movement direction that is ‘preshaped’ by the number of precues (i.e. by prior information). More information about which was likely to be the direction of movement led to narrow peaks of population activity around the mean movement direction. This chapter aims to extend this to humans and link the

findings to parameters in the drift-diffusion model. De Lange et al (2013) use an RDK task and magneto-encephalography to find that prior expectation cause lateralisation of beta oscillations over motor cortex. This study is similar, but uses corticospinal excitability.

2.1.5. An experimental challenge: minimising impulse suppression

MEPs are increasingly understood to reflect multiple motor preparatory processes; they are the summed read-out of all the processes that affect corticospinal excitability (Bestmann and Duque, 2015). In particular, MEPs are known to have a phase of broad suppression prior to ramping up before movement (Greenhouse et al., 2015b), which is not easily understood in the context of biased competition. I will refer to this as ‘impulse suppression’, as it has often been hypothesised to play the functional role of a ‘hold your horses’ signal that prevents premature movement (Duque et al., 2010), even though this functional interpretation is now disputed (Greenhouse et al., 2015b; Duque et al., 2017).

The literature on impulse suppression was discussed in Chapter 1; to briefly summarise some key properties:

- MEP suppression has been shown to follow both a warning cue and an imperative cue giving the signal to move (Touge et al., 1998; Duque et al., 2014; Greenhouse et al., 2015b)
- It is robust to the particular properties of the cue (Duque and Ivry, 2009; Sinclair and Hammond, 2009)
- It affects all effectors and is strongest in the effector that will move (Duque and Ivry, 2009)
- In precue paradigms, MEP suppression is present if the precue-imperative gap is short (0.5 s) but obliterated when this delay is

long (2.0 - 2.5 s; Hasbroucq et al., 1997; Touge et al., 1998; Tandonnet et al., 2010)

I hypothesised that impulse suppression would mask the effect of prior expectation, because it is a non-effector specific signal. I believed prior expectation would affect motor plans by altering biased competition, which is considered a separate process. The three experiments presented in this chapter attempted various manipulations with the aim of minimising the impulse suppression process to allow measurement of the biased competition process.

2.1.6. Three experiments

The three experiments presented in this chapter all aimed to measure the effect of prior expectation on corticospinal excitability. The differences between the three were concerned with how to distinguish the biased competition process from the impulse suppression process.

Experiment 1 had a simple blocked design, in which the level of prior expectation (e.g. L:R 50:50) was only changed every block. On each block, a veridical cue indicated the left:right likelihood of the RDKs in that block. On each trial, participants had to make a finger movement (button press) to report the direction of an RDK and each trial involved one MEP measure via TMS. As is typical in decision-making experiments, each trial began with a fixation cross, which was designed to capture subjects' attention.

Experiment 1 showed no effect of prior expectation at stimulus onset and this indicated that prior expectation does not affect corticospinal excitability, or that its effect was potentially being masked by impulse suppression. I hypothesised that impulse suppression could have been induced by the fixation cross, which, because it was a fixed interval before the stimulus, acted as a warning cue.

Another possibility for the lack of an effect in Experiment 1 was that MEP fluctuations with time (see Discussion) were making a block design unsuccessful; an interleaved design was used in Experiment 2.

Experiment 2 aimed to answer the question, ‘how does prior expectation modify MEPs in the absence of a warning cue?’. It was a modification of Experiment 1 without a warning cue and other optimisations to the design, such as an interleaved structure (see Methods).

Experiment 2 also showed no effect of prior expectation at stimulus onset. Again, this raised the issue that impulse suppression was present. I hypothesised that the prior cue itself could have been acting as a warning cue, inducing MEP suppression. Experiment 3 induced a further modification to aim to reduce this.

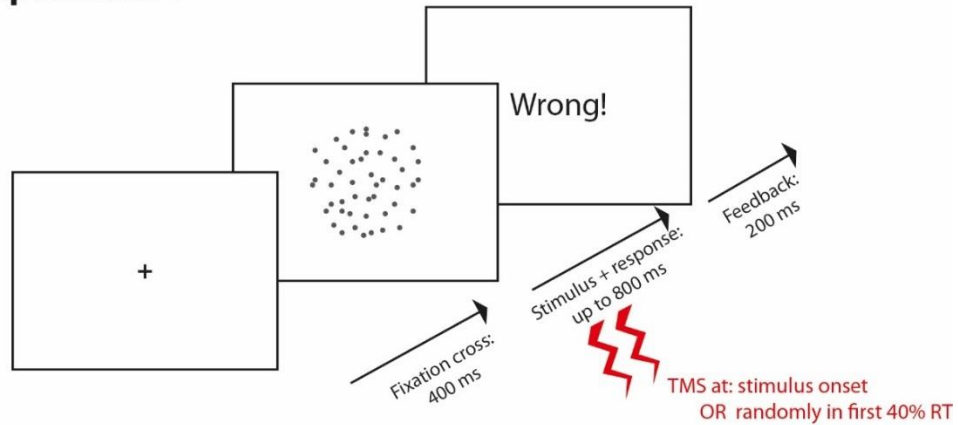
Experiment 3 aimed to answer the question, ‘how does prior expectation modify MEPs when the timing of stimulus onset is unpredictable?’. In this experiment, which was otherwise identical to Experiment 2, the length of prior cue (and thus the timing of stimulus onset) was jittered, to make the onset unpredictable. Although there is no evidence that unpredictable onset reduces impulse suppression, prominent theories (Duque et al., 2010) hypothesise that this suppression prevents the release of premature movements, and thus I hypothesised that making movement onset unpredictable might suppress this process. Furthermore, pilot data in three subjects indicated the hypothesised effect was strongly present in this experimental design, and this further motivated this experiment.

Experiment 2 and Experiment 3 were later combined into one dataset and analysed together (see Results).

2.2. Methods

All three experiments aimed to measure the effect of expectation to move on corticospinal excitability. In all three experiments, subjects watched a computer display. The instruction to move was given by a moving dot stimulus; subjects had to judge whether this was moving left or right and make a corresponding button press with either the left or right index finger. An expectation to move was induced by an explicit veridical cue which told subjects the likelihood the moving dot stimulus would be leftward versus be rightward. MEPs were measured during stimulus onset (i.e. when the moving dot stimulus appeared) and at various timepoints afterwards.


Experiment 1



3 conditions: Instruction at start of block gives likelihood of rightgoing stimulus for that block



Experiments 2 and 3



5 conditions: Veridical cue shows likelihood of rightgoing stimulus for that trial




Figure 2.1 Tasks for the three experiments. In Experiment 1, on each trial, a fixation cross was presented, followed by an RDK. Subjects had a deadline of 800 ms in which to judge the direction of motion of the RDK and make a corresponding response with an abduction of either the left or the right index finger. At the start of each block, subjects were shown a message (see text; figure shows simplified form) which instructed them as to the Left:Right probability of stimulus motion. In Experiments 2 and 3, on each trial, a probabilistic cue was presented, which instructed subjects as to

the probability of rightward stimulus motion. This was followed after a constant (Experiment 2) or jittered (Experiment 3) delay by an RDK. Again, subjects had to make a corresponding button press. In all experiments, TMS was delivered to left motor cortex to evoke an MEP in right FDI.

2.2.1. Participants

All participants were healthy human subjects who had been screened for contraindications to TMS and provided informed consent. Experiment 1 was carried out in 18 subjects (11 female; mean age = 26, $SD = 6$). Experiment 2 was carried out in 27 subjects (17 female; mean age = 21, $SD = 9$). Experiment 3 was carried out in 16 subjects (11 female; mean age = 23, $SD = 7$). The experiments were approved by a UCL ethics board.

2.2.2. Experiment 1

2.2.2.1. Task

On each trial, subjects viewed a Random Dot Kinetogram (RDK; (Britten et al., 1992)). RDKs were used because they require subjects to accumulate evidence over a longer time-period than e.g. a simple arrow stimulus, because they are widely-used for studying perceptual decision making (e.g. Shadlen and Newsome, 2001b), and because the difficulty level can be titrated to individual subjects.

RDKs use dots, replotted frame-by-frame, to give the illusion of a motion signal to left or right. This signal is corrupted by noise because a fraction of the dots move left or right whilst the remainder move randomly. In this experiment, the signal:noise ratio was fixed across the experiment at a difficulty level adjusted for individual subjects (see below on thresholding). Subjects were required to judge the direction of motion and respond with an index finger button press. Subjects rested the index fingers of both hands on either side of the arrow keys of a standard keyboard and pressed the left arrow key with left index finger or right arrow key with right finger to mark their response. This meant that the

button press involved a finger adduction (i.e. used the first dorsal interosseous).

To maximise subjects' behavioural change, a blocked design was used. The probability that the stimulus would move left vs right was not always equally balanced and this changed by block. The possible ratios were 10:90, 25:75, and 50:50 L:R probability. There were a total of six experimental blocks (each consisting of 66 trials), so each probability was repeated once. A random order of the three probabilities was generated for each participant and this was repeated twice, in order to avoid any participant repeating the same probability consecutively.

The probability of the block was explicitly cued at the start of the block. For instance, in a 50:50 block, subjects saw the following message: "For the next 66 trials, the probability the dot field will be moving LEFT is 50%, and the probability it will be moving RIGHT is 50%."

The RDK was displayed for 800 ms, or until subjects made a response. Subjects received feedback according to their response: 'Right!' if their button press correctly indicated the motion direction; 'Wrong!' if it indicated the opposite direction, 'Too Slow!' if subjects failed to make a response whilst the RDK was displayed; 'Too Early!' if subjects pressed in the period before stimulus appearance. This feedback was displayed for 200 ms. The interval between this and the next trial was jittered between 2,750 ms and 3,750 ms (randomly with a uniform distribution). After the intertrial interval, and prior to each stimulus, a fixation cross was presented for 400 ms.

2.2.2.2. Stimulus parameters

Subjects saw 300 dots of size 0.10 degrees of visual angle randomly plotted in a central circular aperture (diameter: 8.1 degrees of visual angle). A certain proportion of dots had coherent motion; that is, they

were replotted every frame with a constant offset along the x-axis to give an illusion of either leftward or rightward motion. Dots had an apparent speed of 4 degrees of visual angle/s. The remainder of dots were replotted with the same offset in a random direction.

Stimuli were viewed on an LCD monitor (Dell) of size 52 cm by 32.5 cm and 1920 by 1200 pixels resolution. Subject sat 90 cm away from the screen. The refresh rate was 60 Hz. Stimuli were presented in MATLAB, version 2014b (Mathworks, <https://www.mathworks.com/products/matlab.html>) and Psychtoolbox (Brainard, 1997; Pelli, 1997; Kleiner et al., 2007). The Magstim was triggered via the DAQ toolbox (Mathworks).

2.2.2.3. TMS

On every trial, a single TMS pulse was delivered to finger area of left motor cortex a 50 mm figure-of-eight-shaped coil connected to a monophasic Magstim 2002 stimulator (Magstim). The stimulation intensity had previously been adjusted for each subject to produce a Motor Evoked Potential (MEP) of approximately 1 mV in relaxed right first dorsal interosseous muscle (FDI). The coil was held throughout the experiment at a 45° angle relative to an anteroposterior axis by the experimenter. Optimal stimulation site was marked in pencil on subjects' scalps prior to the experiment. A block lasted approximately 5.5 minutes. The coil was repositioned on the head using the pencil marks at the start of each new block.

The time-point of TMS delivery was varied: half of pulses were delivered at 0 ms relative to stimulus presentation (i.e. stimulus onset) which was the key timepoint of interest. To build up a picture of corticospinal excitability after stimulus onset, the other half were delivered at a random point within the first 40% of individual subject reaction time

(measured prior to the main experiment; see Section 2.2.2.4). Before each block, and prior to subjects seeing the probability information, 14 TMS pulses were delivered. This was to determine a pre-block baseline measure of corticospinal excitability for the data analysis.

EMG was recorded from right hand FDI using disposable prewired electrodes (Henleys Medical Supplies Ltd). Electrodes were attached to muscle belly and thumb index finger metacarpophalangeal joint, and earthed at a bony wrist process.

2.2.2.4. Protocol

Prior to the main experiment, subjects completed a training, without TMS, of 30 trials which were similar to the main experiment but with a fixed RDK coherence of 70%.

Subjects next underwent a thresholding procedure which adjusted the difficulty of the RDK in order to standardise performance. They undertook 100 trials without probability information and the RDK coherence was modulated to individual ability such that subjects were getting 80% of trials correct. A Bayesian thresholding procedure was used (Quest; Watson and Pelli, 1983).

Subjects then undertook 3 blocks of 75 trials (one each at each level of prior cue) which were similar to the main experiment, and incorporated the individually thresholded coherence. These blocks were used to measure their mean reaction time, which determined the timing of the TMS.

Subjects then had their TMS 'hotspot' determined by adjusting the position of the coil on the left hemisphere until the point of maximum right FDI MEP amplitude was obtained. Stimulation intensity was

adjusted to produce an MEP of approximately 1mV at rest, stable over approximately 10 trials, with the right hand relaxed.

Prior to the main experiment, subjects were asked to relax their hands as much as possible between trials and ignore the TMS pulse. Subjects were shown the EMG trace at this point and practised relaxing their hand with this biofeedback. During the experiment, the experimenter monitored EMG activity and reminded the subject to relax if there was activity during the intertrial interval.

It was emphasised to subjects that they should try to make exclusive movements (to move the left finger **or** right finger). Each subject then carried out another brief training of 12 trials with TMS, so they could accustom themselves to doing the task with TMS pulses.

2.2.3. Experiment 2

As described above, this experiment was a modified version of Experiment 1, and so the method was similar, except in the following ways. The main changes to the method were, firstly, switching from a blocked to an interleaved design, to avoid problems with MEP fluctuations over time, and, secondly, removing the fixation cross, to avoid possible suppression of MEPs associated which has been reported for tasks with a cue a short interval before the imperative (Hasbroucq et al., 1997). The changes are detailed in full below:

Absence of fixation cross: As highlighted above and in the Introduction, in order to obliterate a possible warning cue, there was no fixation cross.

Trial types were interleaved across blocks: Prior cue levels were interleaved across blocks to reduce the possibility of MEP fluctuations over time affecting the results (see Discussion). Subjects underwent 400 trials in four blocks of 100 trials. They had a break between each trial

during which the coil was removed from the head. The trials were divided evenly between the five probability conditions (10/25/50/75/90% rightward probability) and interleaved randomly.

Forty trials, interleaved randomly amongst the remainder, were ‘control’ trials, meaning a large arrow which accurately indicated the direction of stimulus motion was displayed superimposed on the RDK. These trials were included to measure portion of reaction time that did not depend on RDK interpretation (Klein-Flugge and Bestmann, 2012) (see Section 2.2.9).

Prior cue on every trial: In order to permit an interleaved design, probability information (i.e. the prior expectation) was given to subjects on every trial, rather than at the start of a block. Each trial began, prior to RDK appearance, with the appearance of a bar which informed participants of the likelihood that the stimulus following would be moving rightward (see Figure 2.1). The bar gave a visual illustration of the likelihood, and this probability was also displayed as a percentage underneath the bar. The possible likelihoods of moving right were 10, 25, 50, 75 and 90%. The bar was displayed 2,950 ms.

Timing changes: The RDK was displayed for 700 ms rather than 800 ms to encourage a strong dependence on the prior cue.

Feedback was displayed for 600 ms, not 200 ms, for easier viewing.

TMS timepoints: Fixed timepoints were used to increase power in a stimulus-locked analysis (see Results).

TMS was delivered at one of three possible timepoints: at stimulus onset, 25% of subject’s mean reaction time, 45% of subject’s mean reaction time. Mean reaction time was in response to RDK stimuli and was

determined for each subject prior to the start of the experiment. It did not change over the course of the experiment.

Determination of mean reaction time: Prior to the main experiment, subjects undertook 100 trials which were similar to the main experiment, and incorporated the individually thresholded coherence, but excluded control trials. A subject's mean reaction time (across all trial types) was determined in these trials and used to determine timing of TMS delivery later.

Thresholded level: Prior to the main experiment, subjects were thresholded to a 70% correct performance, rather than 80%.

Subject instruction: Prior to the main experiment, subjects were told that they would see a bar indicating the likelihood that the next moving dot stimulus would be rightward. It was emphasised that this was probabilistic and a strong rightward probability could still be followed, on occasion, by a leftward moving stimulus.

RDK properties: An alteration was made to how RDK frames were calculated to make these stimuli consistent with other work using these stimuli. Rather than replotting coherent dots on consecutive frames, they were replotted every fourth frame with a constant offset along the x-axis to give an illusion of either leftward or rightward motion. The second frame had a separate random series of dots which were replotted as described for the first frame on the 5th, 8th, 11th frames and so on. Similarly, the third frame displayed another random distribution of dots which were replotted on the 6th, 9th, and 12th frames. This RDK protocol is in keeping with, e.g. Roitman and Shadlen (2002). It did not change the visual appearance of the stimuli.

In addition to right hand FDI, EMG was also recorded from right hand Abductor Pollicis Brevis, right hand Abductor Digiti Minimi, and left hand

FDI. This was to allow for comparison MEP size in non-target muscle if required during analysis. Ultimately, only FDI data was used; because of the null result in FDI data, the analysis was not extended to other muscles.

2.2.4. Experiment 3

The method for this experiment was the same as for Experiment 2 in every aspect, except that the duration of the probabilistic cue was jittered between 2,950 ms and 3,950 ms, which aimed to increase participants' uncertainty about when the RDK would appear on screen (Figure 2.1).

2.2.5. Subject exclusions

For **Experiment 1**, two subjects were excluded from analysis entirely: one had very few baseline MEPs successfully recorded; the other reported at debriefing deliberately slowing his behavioural responses 'to see what would happen'. This left sixteen subjects remaining in the analysis. Two further subjects had a single block removed in MEP analyses due to few baseline MEPs being recorded.

For **Experiment 2**, one subject was excluded because they had a mean reaction time of 389 ms ($SD = 65.1$ ms; mean reaction time across all subjects in this experiment = 521 ms, $SD = 40.7$ ms) and chose the more probable cue nearly 100% of the time in the 10%, 25%, 75% and 90% conditions, suggesting this subject was guessing.

For **Experiment 3**, one subject was excluded because 9.75% of their MEPs were under the MEP amplitude threshold, leading to concerns over a possible floor effect (mean across remaining subjects = 0.75%, $SD = 1.18\%$).

2.2.6. Preprocessing of behavioural data

For plots showing reaction time or success rates, trials in which subjects pressed before stimulus onset (Experiment 1: $M = 0\%$, $SD = 0\%$; Experiment 2: $M = 0.067\%$, $SD = 0.22\%$; Experiment 3: $M = 0.12\%$, $SD = 0.23\%$), and trials in which a key other than left or right was pressed were removed (Experiment 1: $M = 0.047\%$, $SD = 0.10\%$; Experiment 2: $M = 0.13\%$, $SD = 0.28\%$; Experiment 3: $M = 0.017\%$, $SD = 0.065\%$). Reaction time outliers were also removed (Experiment 1: $M = 0.078\%$, $SD = 0.15\%$; Experiment 2: $M = 0.087\%$, $SD = 0.22\%$; Experiment 3: $M = 0.13\%$, $SD = 0.23\%$). Reaction times were deemed outliers if they were more than three interquartile ranges above the third quartile or below the first quartile. Control trials were also excluded from these plots.

2.2.7. Preprocessing of MEPs

EMG activity was recorded in Signal (Cambridge Electronic Design Limited) and exported to MATLAB, version 2014b (Mathworks, <https://www.mathworks.com/products/matlab.html>). A 50 Hz notch filter was applied to EMG data either at time of recording or retrospectively if it was judged, on visual inspection, to be contaminated by 50 Hz noise. MEP amplitude was extracted by searching for the maxima and minima in a time window 0 to 60 ms after pulse delivery. Trials in which the rectified, baseline-corrected EMG value exceeded 0.1 mV at any point between 50 ms and 5 ms prior to the TMS pulse were discarded as muscle contraction is known to potentiate MEP amplitude (Hess et al., 1987; Rossini et al., 1988). 'MEPs' with an amplitude less than 0.05 mV were discarded; this threshold was to prevent the algorithm reporting background fluctuations as MEPs. The EMG traces were manually checked by an observer for approximately five subjects to ensure the algorithm described was producing plausible results. MEPs

were deemed outliers if they were more than three interquartile ranges above the third quartile or below the first quartile; these were also removed.

For **Experiment 1**, on average, 1.19% of MEPs were lost due to precontraction ($SD = 2.81\%$), 0.98% due to insufficient amplitude ($SD = 1.64\%$), and 0.53% due to outlier exclusion ($SD = 0.95\%$). For **Experiment 2**, on average, 1.37% of MEPs were lost due to precontraction ($SD = 1.38\%$), 0.42% due to insufficient amplitude ($SD = 0.84\%$), and 0.40% due to outlier exclusion ($SD = 0.94\%$). For **Experiment 3**, on average, 1.28% of MEPs were lost due to precontraction ($SD = 1.87\%$), 0.75% due to insufficient amplitude ($SD = 1.18\%$), and 0.20% due to outlier exclusion ($SD = 0.30\%$).

2.2.8. ‘Too Slow’ trials

On trials in which subjects did not make a response within the RDK time limit (800 ms or 700 ms), the stimulus disappeared and subjects saw a ‘Too Slow’ message (5.58% of trials, $SD = 3.44\%$). As no reaction time or response was recorded, these trials were not included in analyses, with the exception of analyses of MEP data at the first timepoint. These analyses did not rely on sorting MEPs by reaction time or response and so ‘Too Slow’ trials could be included.

2.2.9. Control trials

Experiments 2 and 3 included ‘control trials’ which showed an arrow in the same direction as the RDK. These trials are easier to respond to because they do not involve the evidence-accumulation stage of viewing the RDK. These trials were included because they could be used to measure the time needed to detect a stimulus and make the motor response, and identify changes in corticospinal excitability related to this (as used by Klein-Flugge and Bestmann, 2012). This then allows an

analysis of corticospinal excitability in non-control trials that focuses on the part of the trial prior to the timing of excitability changes in control trials. In other words, it allows an identification of purely decision-related activity.

The analysis of control trials here was based on that in Klein-Flugge and Bestmann (2012), although the method used was slightly different. The question was whether choice-selective activity is seen in a part of the reaction time known to be prior to ‘motor preparation’. The timing of this period was found using the control trials, which do not require an evidence accumulation phase.

MEPs in control trials were analysed by eventual choice and timepoint. It was found that there was no difference between activity by eventual choice at the first and second timepoints. Therefore, for each subject, the average timing of the second timepoint was calculated in milliseconds. This was the ‘decision epoch’, i.e. the period known not to invoke choice-selective activity on control trials. MEPs recorded on non-control trials in this epoch were then selected, and analysed by eventual choice, to see if choice biases were present in the decision epoch. Other than for this analysis, control trials were excluded from all analyses except for those of MEP data at the first timepoint. At this timepoint, subjects had not yet seen the stimulus and so heterogeneous trial types could be included.

2.2.10. MEP normalisation

In **Experiment 1**, MEPs were ‘normalised’ by in-block MEPs being divided by the median of preblock MEPs. There were 14 of these preblock MEPs on each block. Those with precontraction (as defined above) were excluded and the median of the remaining MEPs was used.

In **Experiments 2 and 3**, the interleaved design reduced the importance of normalisation and also permitted normalisation by z-scoring, which was the method used.

2.2.11. Primary and secondary analyses

Some of the analyses presented in this chapter were planned prior to data collection whilst others were planned post-hoc.

The primary MEP analysis planned was an analysis of normalised MEPs at stimulus onset plotted by prior cue (Figure 2.6 and Figure 2.10). Analyses of reaction times and error rates plotted by prior cue were also planned (e.g. Figure 2.3 and Figure 2.4). An analysis in which bias parameter was modelled (see Section 2.2.12) and related to MEP size was also planned (Figure 2.15).

Secondary analyses were: (1) analysing non-normalised MEPs at stimulus onset by prior cue (Figure 2.5 and Figure 2.11); (2) analysing MEPs at stimulus onset by later choice (Figure 2.13); (3) analysing the relationship between MEP size and reaction time (Figure 2.16 and Figure 2.22); and (4) extending the analysis of bias parameter to late MEPs (Figure 2.21). Furthermore, whilst an analysis of MEPs at timepoints post-stimulus onset was planned, specific decisions about how to bin or smooth these (Figure 2.7, Figure 2.19 and Figure 2.20) were made post-hoc.

2.2.12. Modelling analysis

The behavioural data in **Experiments 2 and 3** was used for drift-diffusion modelling. The drift-diffusion model is a decision model for modelling two-alternative forced choice tasks. In the drift-diffusion model, a decision variable representing current belief in a hypothesis

accumulates to one of two decision bounds. The model is described in more detail in Chapter 1.

As Chapter 1 explains, there is behavioural evidence that prior expectation provides a prestimulus bias to the activity of a drift-diffusion model. This line of thinking influenced us to look for a correlate in corticospinal activity: a prestimulus bias that represented prior expectation.

As Chapter 1 explains, there are many similar diffusion and race-to-threshold models. Here the drift-diffusion model was used because it has been widely used in past literature (e.g. Gold and Shadlen, 2007; Mulder et al., 2012; Rao et al., 2012).

Here I used the HDDM toolbox, version 0.6.0 (http://ski.clps.brown.edu/hddm_docs/, Wiecki et al., 2013) running in Python, version 2.7 (Python Software Foundation, <https://www.python.org/>), and. The HDDM allows for hierarchical Bayesian parameter estimation of the Drift Diffusion Model. Hierarchical here means parameters of the model are estimated for each subject individually under the assumption that there is a group-level distribution (i.e. participants are similar but not identical). This method is preferred when datasets are small (Ratcliff and Childers, 2015), and so is well suited to this dataset with <400 trials per subject.




Figure 2.2 Parameters in the drift-diffusion model. The figure shows the parameters in the evidence accumulation process. The decision variable must accumulate to one of the two response boundaries for a response to be selected. The distance between these two response boundaries is the threshold, a . The mean slope of the evidence accumulation signal is the drift rate, v . The parameter of interest in the analysis presented in this chapter is the bias, z , which allows the evidence accumulation to begin closer to one response boundary than the other and thus lower the evidence required for a particular response. There is also a portion of non-decision time, t , which represents a fixed time taken by non-decision processes. The dark red and blue lines show histograms of reaction times for each response generated by the model.

The model toolbox was used to estimate the parameters of the model in each subject under the assumption that the bias, z , of the model was different with different levels of prior expectation cue. All other parameters, including the drift rate, v , were assumed to be fixed across conditions (Figure 2.2). I did not include model comparison of alternative models, as the finding that prior expectation mainly or exclusively alters bias rather than drift rate has been consistent among many previous studies (Ratcliff and McKoon, 2007; Simen et al., 2009; Mulder et al., 2012), and was not the question motivating this study.

The model was operationalised as follows: (1) the *find_starting_values* function was used to find plausible starting parameters via optimisation; (2) Markov Chain Monte Carlo simulations were run to sample 20,000 times from the posterior distributions of the parameters. Four thousand samples were discarded as ‘burn-in’. This level of sampling was arrived at after visually inspecting plots of parameter estimates over time for evidence of Markov Chain convergence.

All trials in which the participant responded before RDK end were included in modelled data (i.e. previously-identified reaction time outliers were not excluded) but the model was instructed to assume a certain proportion of reaction times were outliers drawn from a separate distribution, and the starting proportion of these was set at 0.1%, which was informed by the outlier analysis reported above.

Due to the small numbers of trials, I did not attempt to model inter-trial parameters.

The existence of a relationship between modelled bias and MEP was tested. For each subject, the bias parameters derived from the model (one per prior cue level) were used to fit a linear regression (y: bias parameter, X: prior cue level and constant term). An equivalent linear regression was also fit for each subject for mean MEP at stimulus onset (y: mean MEP at stimulus onset, X: prior cue level and constant term). This gave, for each subject, an 'MEP slope' and a 'bias slope' (β coefficients from the linear regression). Across subjects a linear regression was fit to the bias slopes against the MEP slopes. This tested whether subjects who showed steeper changes in MEP size across prior conditions showed steeper changes in bias parameter.

Finally, this analysis was repeated to test a behavioural measure directly, in place of the modelled bias parameter: reaction speed for right-choice trials only. Reaction speed is the inverse of reaction time and it was used so the directionality would be the same as the bias parameter (stronger rightward expectation leading to both a larger rightward bias and a faster reaction speed in right-choice trials). As above, reaction speed was used to fit a linear regression within each subject (y: reaction speed, X: prior cue level and constant term). Again, the β coefficients of these relationship were regressed against the MEP slopes across subjects.

2.2.13. Analysis of MEPs over time

In both experiments, MEPs collected at timepoints after the stimulus onset timepoint were collated to create plots displaying excitability during the early part of the reaction time.

In **Experiment 1**, MEPs were recorded at random timepoints in the first 40% of the reaction time. The time of the TMS probe on each trial was divided by the reaction time on that trial to normalise MEP timings by variable trial reaction times. This is sometimes referred to as ‘stimulus-locked and response-locked’. Within each subject, the MEPs were then sorted into six equally-spaced time-bins between 0% and 40% of reaction time. A mean value was obtained for each bin within subjects, and this was then meaned across subjects. This analysis is shown in Figure 2.7.

MEPs in **Experiments 2 and 3** were recorded at discrete timepoints (0%, 25%, and 45% of mean reaction time) to give more power to ‘stimulus locked’ analyses. ‘Stimulus locked’ analyses with MEPs presented at their original timepoints are shown in Figure 2.17 and Figure 2.18.

In order to look at a continuous measure of MEPs over time, MEPs were also ‘response-locked’, by finding, on each trial, the time remaining between the time of the TMS probe and the reaction time. These data were used to calculate a moving average. In each subject, MEP data was averaged in bins 150 ms wide. The bin was applied to the data every 20 ms. Any bin that did not contain a minimum of three MEPs for the average was discarded. The bins were then averaged across subjects. Any bin that did not contain a minimum of two subjects’ data was discarded. The results of this moving average ‘response locked’ analysis are shown in Figure 2.19.

2.2.14. Statistical analysis

Reaction times and MEPs at stimulus onset were tested for statistical significance using one- or two-way repeated measures ANOVAs. Where assumptions of sphericity were violated (as determined by Mauchly's test), Greenhouse-Geisser corrections were applied, indicated in the text as non-whole number degrees of freedom. For Figure 2.19 and the analysis of control trials (Section 2.3.3.6), Multiple comparison t-tests were corrected for with a False Discovery Rate procedure (Storey, 2002) within each condition, with the False Discovery Rate set at 0.25 (Bonferroni correction was not used because the tests were not independent).

2.3. Results

Three experiments were conducted to measure the effect of an expectation to move on corticospinal excitability. The particular hypothesis was that an expectation to move a particular effector increases corticospinal excitability for that effector prior to stimulus onset, i.e. prior to the process of weighing evidence for the decision beginning. Subjects saw a moving dot stimulus (RDK) which instructed them whether to make a movement with either the left or right index finger. Prior to viewing the stimulus, they were biased by a cue that instructed them about the relative likelihood of a leftward vs a rightward stimulus. MEPs were measured at stimulus onset, the key timepoint to test the hypothesis, as well as at later timepoints during the ongoing decision.

2.3.1. Experiment 1

Experiment 1 tested the hypothesis above using a standard psychophysical design. The level of prior expectation cue (50:50, 25:75, or 10:90 L:R expectation) changed on a block basis.

I measured right hand MEPs and hypothesised that a stronger expectation of rightward movement (i.e. the 10:90 condition) would facilitate MEPs in this hand, and that this facilitation would be present at the time of stimulus onset.

2.3.1.1. Behavioural data

Behavioural data was analysed to test whether subjects had understood and used the prior cue given at the start of each block. Prior cue level modified reaction time (Figure 2.3). A more biasing cue (i.e. 10:90 versus 50:50) reduced the reaction time to respond to the RDK (repeated measures ANOVA, $F(2, 30) = 12.5$, $p < .001$, $\eta^2_p = .45$).




Figure 2.3 Reaction time decreases with increasing rightward expectation. Participants were slower to respond to a Random Dot Kinetogram stimulus when they were in a block in which 50% of stimuli were rightward compared to when they were in a block in which 90% of stimuli were rightward. So the experimental manipulation modified subjects' behaviour, $F(2, 30) = 12.5$, $p < .001$, $\eta^2_p = .45$. Error bars show standard error of mean.

Similarly, subjects made more errors when they had a less informative probability cue. They made more incorrect button presses in the unbiased condition (Figure 2.4B; repeated measures ANOVA, $F(1.18, 17.7) = 12.1$, $p = 0.0019$, $\eta^2_p = 0.447$). A similar, though non-significant, trend is present for no-response errors (i.e. did not respond to the RDK within 800 ms; Figure 2.4A; repeated measures ANOVA, $F(2, 30) = 2.36$, $p = 0.11$, $\eta^2_p = .14$).

To summarise, probability information given at the start of a block was used by subjects to improve their speed and accuracy. Did it also have an effect on corticospinal excitability? I first consider MEPs at stimulus onset ('0 ms'), the key timepoint of interest.




Figure 2.4 Errors decrease with increasing prior expectation. Both errors in which subjects did not press a button to respond to the stimulus before it timed out (A) and errors in which the button press did not match the direction of stimulus motion (B) decreased with increasing prior expectation. One-way ANOVAs showed a significant effect for ‘Too Slow’ errors, $F(1.18, 17.7) = 12.1$, $p = 0.0019$, $\eta^2_p = 0.447$, but not wrong response errors, $F(2, 30) = 2.36$, $p = 0.11$, $\eta^2_p = .14$. Participants made more errors when they were in a block in which 50% of stimuli were rightward compared to when they were in a block in which 90% of stimuli were rightward. Error bars show standard error of mean.

2.3.1.2. MEPs at stimulus onset




Figure 2.5 Raw MEP size does not change with prior expectation. MEPs measured in right FDI at stimulus onset (the appearance of the RDK) are shown. These were hypothesised be larger with increasing rightward expectation. In fact, there is no effect of level of prior expectation, $F(2, 30) = 0.67$, $p = .52$, $\eta^2_p = .043$. Error bars show standard error of the mean.

At stimulus onset, there was no change in raw MEP amplitude with probability condition (Figure 2.5; repeated measures ANOVA: $F(2, 30) =$

0.67, $p = .52$, $\eta^2_p = .043$). However, this analysis was always intended to look at a relative MEP measure, as the block design made the experiment susceptible to any baseline fluctuations in corticospinal excitability. A baseline excitability measure was taken at the start of each experimental block, before subjects had seen any probability information (see Methods). Here, within-block MEPs were divided by the median of the baseline MEPs measured at the start of the block (Figure 2.6). As with the raw MEPs, there was no change between prior cue levels (repeated measures ANOVA: $F(2, 30) = 2.24$, $p = .12$, $\eta^2_p = .13$).




Figure 2.6 Relative MEP size does not change with prior expectation. MEPs measured at stimulus onset were normalised by dividing by the median of baseline MEPs recorded prior to the start of the block. This was to reduce inter-subject variability in MEP size. MEPs were hypothesised to be larger with increasing rightward expectation. In fact, as with raw MEPs, there was no effect of prior expectation, $F(2, 30) = 2.24$, $p = .12$, $\eta^2_p = .13$. Error bars show standard error of the mean.

2.3.1.3. MEPs over time

In addition to the data at stimulus onset, MEPs were also sampled at later timepoints. Due to variation in reaction times, TMS delivered at timepoints beyond stimulus onset does not fall at a consistent point in the reaction time and so here the data has been scaled by percentage of reaction time elapsed ('stimulus and response locked') and binned (Figure 2.7).

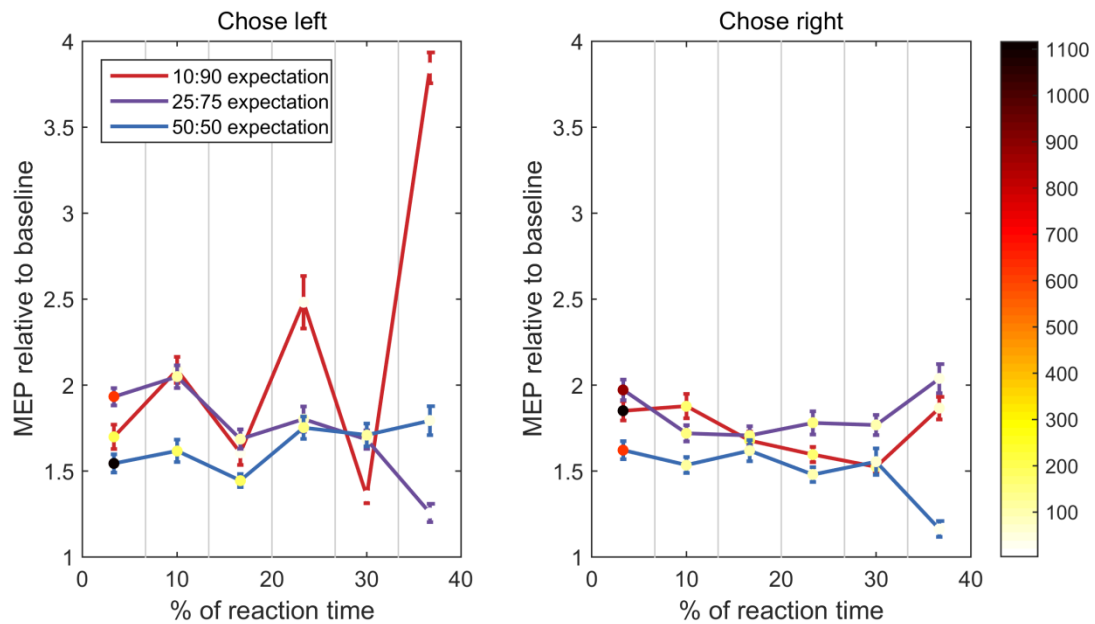


Figure 2.7 Right hand MEPs over time in the three prior cue conditions. Left panel shows trials on which the subject chose left; right panel shows trials on which the subject chose right. Timepoint of TMS delivery has been scaled by reaction time on the relevant trial, pooled across subjects, and sorted into six bins. Due to the scaling by reaction time, data becomes sparser near to 40% of reaction time. Colouring of points shows number of MEPs in relevant mean, according to colour scale on right hand side. Error bars show standard error of the mean.

2.3.2. Comparison of Experiment 2 and Experiment 3

Experiments 2 and 3 aimed to test the effect of prior expectation on corticospinal excitability under altered temporal predictability of stimulus onset. These were motivated by past literature suggesting temporally predictive cues can suppress MEPs (Duque et al., 2010). Experiment 2 involved stimulus onset at a fixed time, but without the fixation cross that had preceded stimulus onset in Experiment 1. Experiment 3 jittered stimulus onset to further reduce the temporal predictability.

Experiment 3 followed the same protocol as Experiment 2, except that the length of time the prior expectation cue was presented for was jittered, to reduce subjects' expectation about when the stimulus would appear.

A four-way ANOVA was used to compare MEPs in the two datasets, with Hand chosen, Probability, and Timepoint as within-subject factors, and Dataset as a between subject factor. This ANOVA revealed no main effects and no interactions. The results of this ANOVA are given in full in Appendix I.

A t-test was used to compare subject mean reaction times across the two datasets, collapsed across condition. As with the MEPs, this revealed no difference between datasets, $t(39) = 1.13$, $p = .27$, $d = .37$. The mean reaction time for Experiment 2 was 526 ms ($SD = 31.7$ ms) and the mean reaction time for Experiment 3 was 513 ms ($SD = 41.1$ ms).

Because no reaction time or MEP difference was found between these two datasets, they were combined for the remaining analyses, which are presented below.

2.3.3. Experiment 2 and Experiment 3 combined dataset

2.3.3.1. Behavioural data

Reaction times across the five conditions are shown in Figure 2.8A. A repeated measures ANOVA showed prior cue had a significant effect on reaction time $F(2.78, 111.1) = 15.76$, $p < .001$, $\eta^2_p = .283$. Subjects were slowest in the 50:50 condition, in which the prior cue did not offer any information about the likeliest direction of the stimulus. Biasing prior cues decreased reaction time, with a more strongly biasing cue increasing reaction time more. This suggests that subjects had understood the cues they were shown and were using them in the perceptual decision.

Separating reaction time by response as well as prior cue condition (Figure 2.8B) shows the advantage a biasing cue gives – a speeded reaction time – applies only when the response is that the subject has been primed to expect by the cue. (Note that the trials in which the cue moves in the opposite direction to that suggested by the prior are rarer,

and thus the bars in Figure 2.8A are not an average of the red and blue bars in Figure 2.8B.) These data show a main effect of probability cue condition, but no main effect of response direction on ANOVA. There is a significant interaction between these factors. Probability: $F(2.69, 104.8) = 3.18$, $p = .032$, $\eta^2_p = .075$; Hand: $F(1, 39) = 0.02$, $p = .902$, $\eta^2_p = .000$; Probability*Hand: $F(1.52, 59.2) = 22.2$, $p < .001$, $\eta^2_p = .36$.

In short, the biasing information given by the prior cues was used by subjects to make themselves faster, indicating that they had understood and used the information given.




Figure 2.8 Subjects were fastest with the response they most expected to make. Behavioural data for Experiments 2 and 3, showing: (A) Mean reaction time by prior cue condition, excluding control trials; (B) Mean reaction time by prior cue condition and response, excluding control trials. (B) shows that the U-shaped curve in A is composed of two linear trends in reaction time, so that subjects are fastest with right-hand responses with increasing expectation of a rightward stimulus, and vice versa. In (A), there is a significant effect of prior expectation, $F(2.78, 111.1) = 15.76$, $p < .001$, $\eta^2_p = .283$. In (B), there is a significant effect of prior expectation and a significant interaction, Probability: $F(2.69, 104.8) = 3.18$, $p = .032$, $\eta^2_p = .075$; Hand: $F(1, 39) = 0.02$, $p = .902$, $\eta^2_p = .000$; Probability*Hand: $F(1.52, 59.2) = 22.2$, $p < .001$, $\eta^2_p = .36$. In both panels, error bars show standard error of the mean.

Similarly, subjects scored more correct trials when the information was more strongly biasing (Figure 2.9A), $F(3.03, 121.1) = 8.91$, $p < .001$, $\eta^2_p = .18$. Figure 2.9B shows how choice varied by prior cue condition; subjects were more likely to choose in the direction of the bias with a more biasing cue.

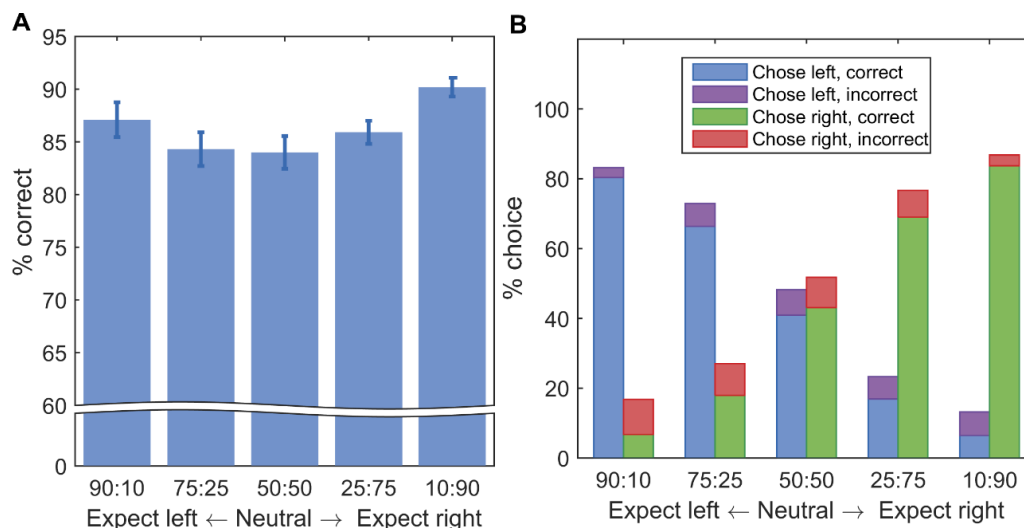


Figure 2.9 Behavioural choice is modulated by prior cue. Subjects' choice reflected the level of prior expectation. (A) Percentage of trials that were correct in the prior cue conditions. Error bars show standard error of the mean, $F(3.03, 121.1) = 8.91$, $p < .001$, $\eta^2_p = .18$. (B) Percentages of trials in each prior cue condition where left or right was chosen as the response, and where the chosen response was correct or incorrect.

2.3.3.2. MEPs at stimulus onset

With evidence that subjects had incorporated the prior cues in their decision-making, MEPs at the main timepoint of interest, stimulus onset ('0 ms'), were analysed. As detailed in the Introduction, I hypothesised that the prior cue shown would affect corticospinal excitability at this early timepoint, prior to subjects viewing the stimulus. In particular, I expected a stronger right-hand expectation to increase corticospinal excitability measured with a right FDI MEP.

Figure 2.10 shows MEPs at stimulus onset plotted by condition. As subjects had not seen the stimulus, these data are not sorted by eventual choice. Contrary to the hypothesis, there was no effect of prior cue on these MEPs, $F(3.09, 123.8) = 0.88$, $p = .45$, $\eta^2_p = .022$.

In case the system of z-scoring MEPs within blocks was obscuring a different trend in the raw data, raw MEPs were also plotted (Figure 2.11). The data show a similar trend to that after z-scoring, and, as with the z-

scored data, there was no effect of prior cue condition on MEP magnitude, $F(2.89, 115.6) = 0.59, p = .62, \eta^2_p = .015$.

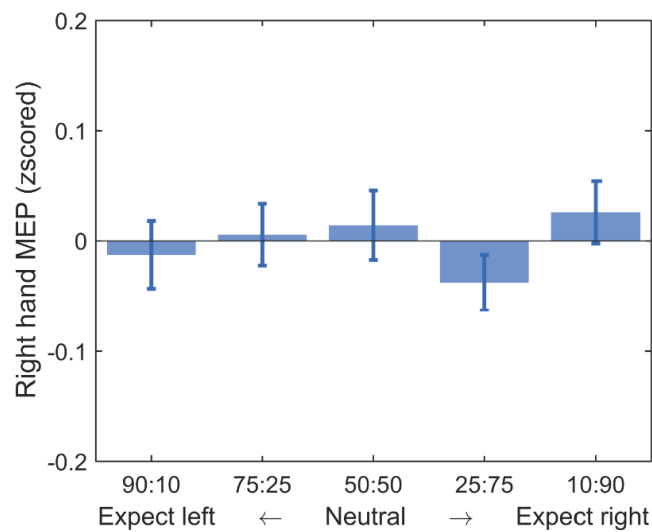


Figure 2.10 No modulation of MEP by prior expectation cue at stimulus onset. For Experiments 2 and 3, this figure shows MEPs recorded at stimulus onset, regardless of ultimate choice. It was hypothesised that a stronger rightward expectation would increase right hand MEP. In fact, there was no effect, $F(3.09, 123.8) = 0.88, p = .45, \eta^2_p = .022$. Within each subject, MEPs have been z-scored within blocks. Error bars show standard error of mean.

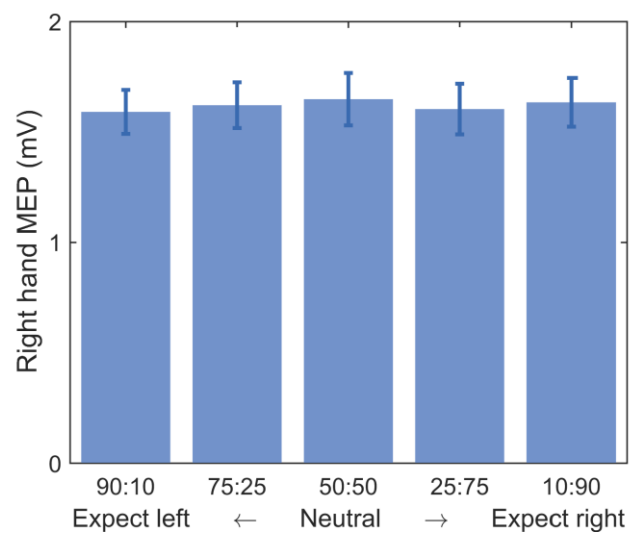


Figure 2.11 No modulation of raw MEPs by prior expectation cue at stimulus onset. To check the absence of effect in Figure 2.10 was not a problem with the normalisation method used (z-scoring), raw MEPs were also plotted. Once again, there was no effect of prior cue on MEP size, $F(2.89, 115.6) = 0.59, p = .62, \eta^2_p = .015$. Error bars show standard error of the mean.

Figure 2.12 once again shows MEPs measured at stimulus onset by prior cue condition, except that here only MEPs have been plotted for trials in

which both that trial and the preceding had the same prior cue. If there was a weak effect of prior cue on corticospinal excitability, this might be expected to be strengthened by successive identical cues. However, there was once again no effect of prior cue condition, $F(4, 160) = 1.50$, $p = .21$, $\eta^2_p = .036$.




Figure 2.12 No effect of prior cue condition on MEP amplitude for trials in which the same prior cue had been seen on that trial and the preceding one. Trials in which the same prior cue had been seen twice in a row were selected, to see if a stronger bias showed an effect of prior cue on MEP size. However, there was no effect of prior cue on MEPs in these trials, $F(4, 160) = 1.50$, $p = .21$, $\eta^2_p = .036$. MEPs have been z-scored within blocks within subjects. Error bars show standard error of the mean.

Figure 2.13 once again shows MEPs at stimulus onset, now sorted by eventual choice as well as prior cue condition. Subjects have not seen the stimulus at this point, so any differences in MEPs by right/left choice reflect a *cause* of that choice rather than a response to the stimulus. Testing this data via two-way ANOVA revealed no main effect of probability, $F(3.20, 118.6) = 2.18$, $p = .090$, $\eta^2_p = .056$, or choice, $F(1, 37) = 2.46$, $p = .13$, $\eta^2_p = .062$, and no significant interaction between these factors, $F(4, 148) = 0.67$, $p = .62$, $\eta^2_p = .018$.




Figure 2.13 Post-hoc analysis shows that choice affects MEP size at stimulus onset in some conditions. In this analysis of MEPs at stimulus onset split by eventual choice, only the stimulus has not yet been seen so any differences between MEPs according to eventual choice (left or right) are causal. There was no overall main effect of probability, $F(3.20, 118.6) = 2.18$, $p = .090$, $\eta^2_p = .056$, or choice, $F(1, 37) = 2.46$, $p = .13$, $\eta^2_p = .062$, and no significant interaction between these factors, $F(4, 148) = 0.67$, $p = .62$, $\eta^2_p = .018$. However, a post-hoc t-test showed that MEPs in the strongest ‘Expect left’ condition are reduced when the eventual choice is right, compared to when the eventual choice is left, $t(40) = 2.54$, $p = .015$, $d = 0.40$. Error bars show standard error of the mean. Annotated significance values show results of paired t-tests between indicated bars.

However, the plot of these data appears to show a trend in which MEPs in the strongest expectation conditions show a difference when the unexpected vs the expected choice is made (e.g. the ‘Chose right’ bar appears suppressed compared to the ‘Chose left’ bar in the 90:10 condition). Paired t-tests were used to test for a difference of left/right choice in the two extreme conditions. There was no effect of choice in the ‘Strongly expect right’ condition, $t(37) = 0.90$, $p = .37$, $d = 0.15$, but an effect in the ‘Strongly expect left’ condition, $t(40) = 2.54$, $p = .015$, $d = 0.40$.

To summarise, in the condition in which subjects most expected to move left, right hand MEPs were *smaller* in trials in which the eventual choice was left compared to when the eventual choice was right. Whilst there is no statistically significant effect in the complementary ‘Strongly expect

right' condition, there is a parallel trend. Note that the effect is contrary to what might be expected: we might guess that a higher right hand MEP would lead to a higher likelihood of moving right on that trial. This counterintuitive finding is examined in the Discussion.

2.3.3.3. MEP-behaviour correlations for stimulus onset MEPs

The hypothesis of this experiment – that an expectation to move would bias corticospinal excitability for the relevant effector – was influenced by ideas from the drift-diffusion and related models. Work with these models has suggested that expectation is incorporated as a prestimulus bias, altering model activity before evidence accumulation begins (Ratcliff and McKoon, 2007; Simen et al., 2009). The aim of these experiments was to look for a similar process in corticospinal excitability.

This motivated an analysis in which behavioural data (reaction time and selected response) was used to fit a drift-diffusion model with a bias parameter, z , allowed to vary between prior cue levels. From this an estimate of z for each prior cue level, within each subject, was obtained.

Figure 2.14 shows the estimated group-level (i.e. all subject) distributions for the bias parameter, z . The bias parameter is measured on a scale between 0 and 1, with 0.5 being unbiased, 0 indicating total certainty about choosing left, and 1 indicating total certainty about choosing right.

The figure shows the modelling produced the expected results. Bias towards the right-hand response is strongest in the most strongly expect right condition (10:90) and the remaining conditions rank in the expected order. The distributions do not overlap one another. The 50:50 condition has a mean bias slightly below 0.5, indicating, at the group level, a slight bias towards left-hand responses.




Figure 2.14 Modelled bias parameters by probability cue condition. A hierarchical Bayesian drift-diffusion model toolbox was used to fit the behavioural data. Plot shows group distributions of the parameter z , which represents a prestimulus bias to either a right-hand response (closer to 1) or a left-hand response (close to 0). The bias parameter was fit under the assumption drift rate and other parameters were fixed across prior cue conditions, but bias varied with condition. The resulting distributions reproduce the order of conditions with the strongest rightward expectation (10:90, blue line) leading to the strongest rightward bias.

The bias parameters from the modelling were used in a regression analysis to see if participants who showed steeper trends in bias parameters showed steeper trends in corticospinal excitability at stimulus onset. In other words, was there a relationship between bias in the model and MEP size?

For each participant, a regression was conducted for (a) mean bias value and (b) mean MEP at stimulus onset in the five prior cue levels. The slopes of these relationships gives the rate of change of (a) bias and (b) MEP by prior cue level. The bias slope and the MEP slope for each subject were then plotted against one another and a regression on these values conducted. This regression thus tested whether participants who showed greater changes in modelled bias parameter across conditions showed greater changes in MEP (Figure 2.15). The regression found no statistically significant relationship, $F(1, 39) = 0.007$, $p = .93$, $R^2 = .0002$. In summary, there was no relationship between trend in modelled response bias and trend in MEP size at stimulus onset.

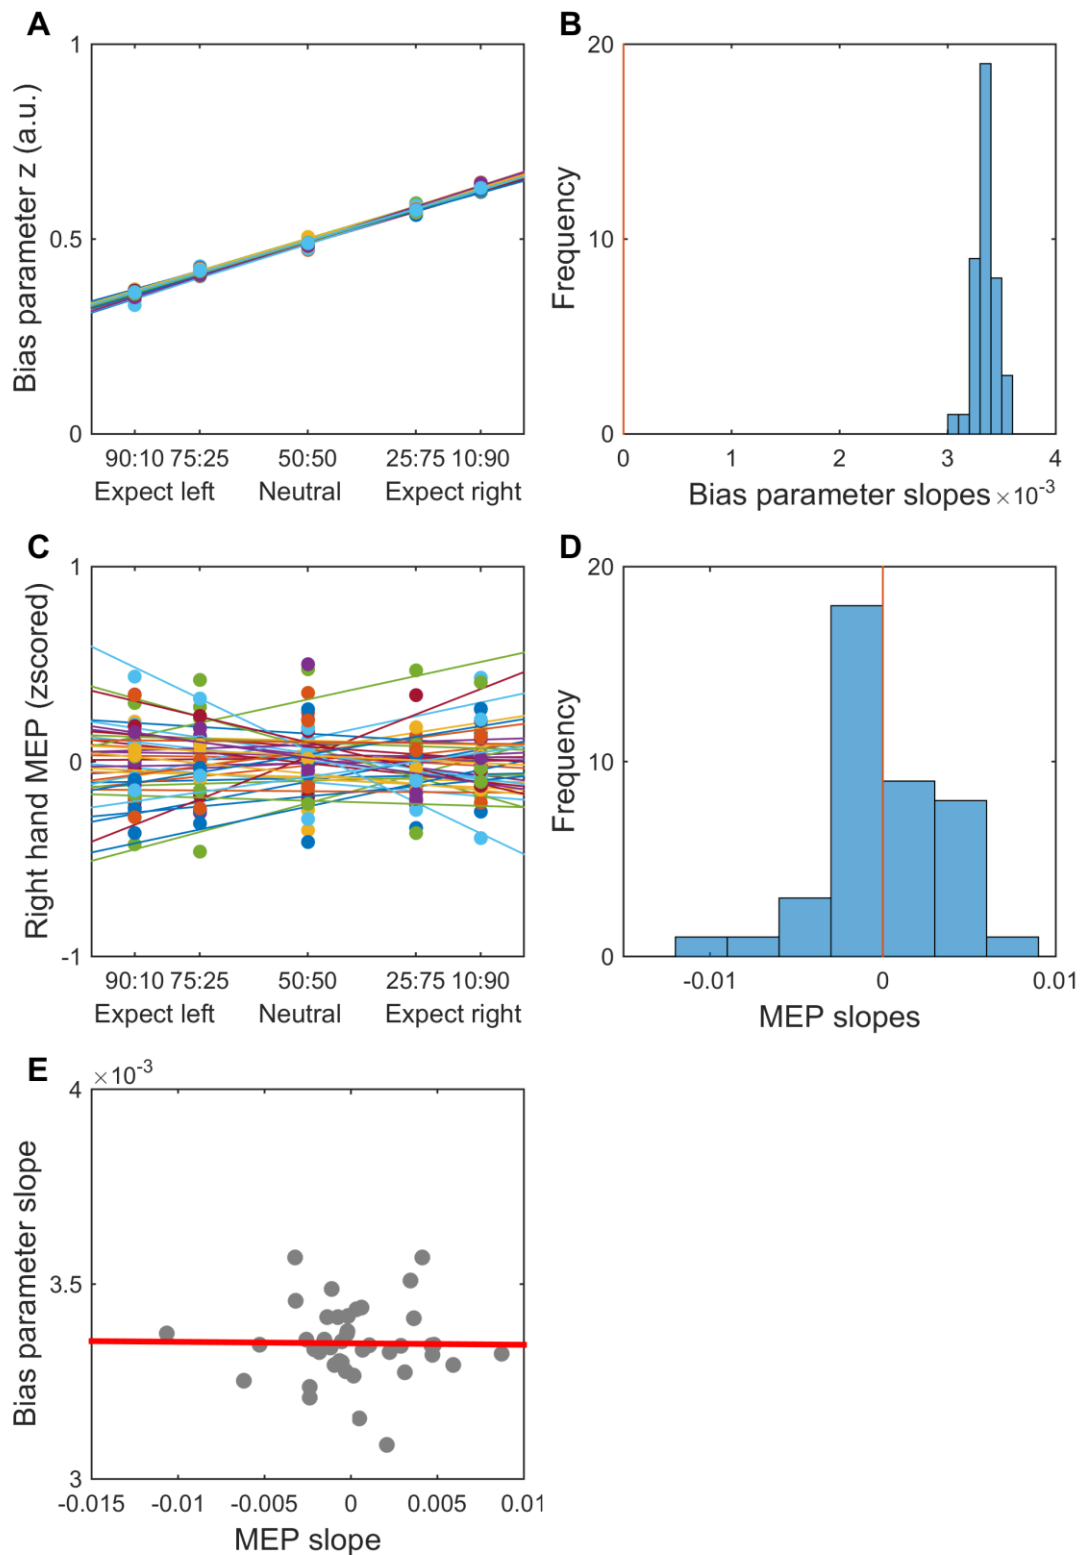


Figure 2.15 No relationship between bias parameter slope and MEP slope. (A) shows the bias parameters, z , produced by the drift diffusion model, for each condition and each subject. A simple linear regression was conducted on these points within each subject. (Each subject's data is shown by coloured points, with a line indicating the fits of the linear regression.) A histogram of the slopes (β) of these relationships is shown in (B). These slopes are significantly different from 0, $t(40) = 223$, $p < .001$, $d = 34.9$. The positive slopes indicate the model finds that subjects are, as expected, more

biased towards a right-hand response the more the prior cue indicates a rightward stimulus is likely. (C) shows an equivalent plot to (A) for MEP in the right FDI at stimulus onset. (D) shows the slopes (β) of within-subjects linear regressions on these data. These slopes are not significantly different from 0, $t(40) = 0.19$, $p = .85$, $d = 0.029$. In (E), for each subject, the bias parameter slope has been plotted against the MEP slope, to see whether subjects who show steeper relationships between experimental condition and modelled parameter also show steeper trends in MEPs. There was no significant relationship, $F(1, 39) = 0.007$, $p = .93$, $R^2 = .0002$.

In order to understand whether final reaction time, rather than modelled bias, might drive MEP differences, the above analysis was repeated using reaction speed (inverse of reaction time) on right-choice trials instead of bias parameter. As expected, most participants were faster to make a right-hand movement in strongly expect right conditions (Figure 2.16A and B). When the regression against MEP slopes was repeated (Figure 2.16C), the slope of the relationship was positive, but there was no statistically significant relationship, $F(1, 39) = 3.78$, $p = .059$, $R^2 = .088$.

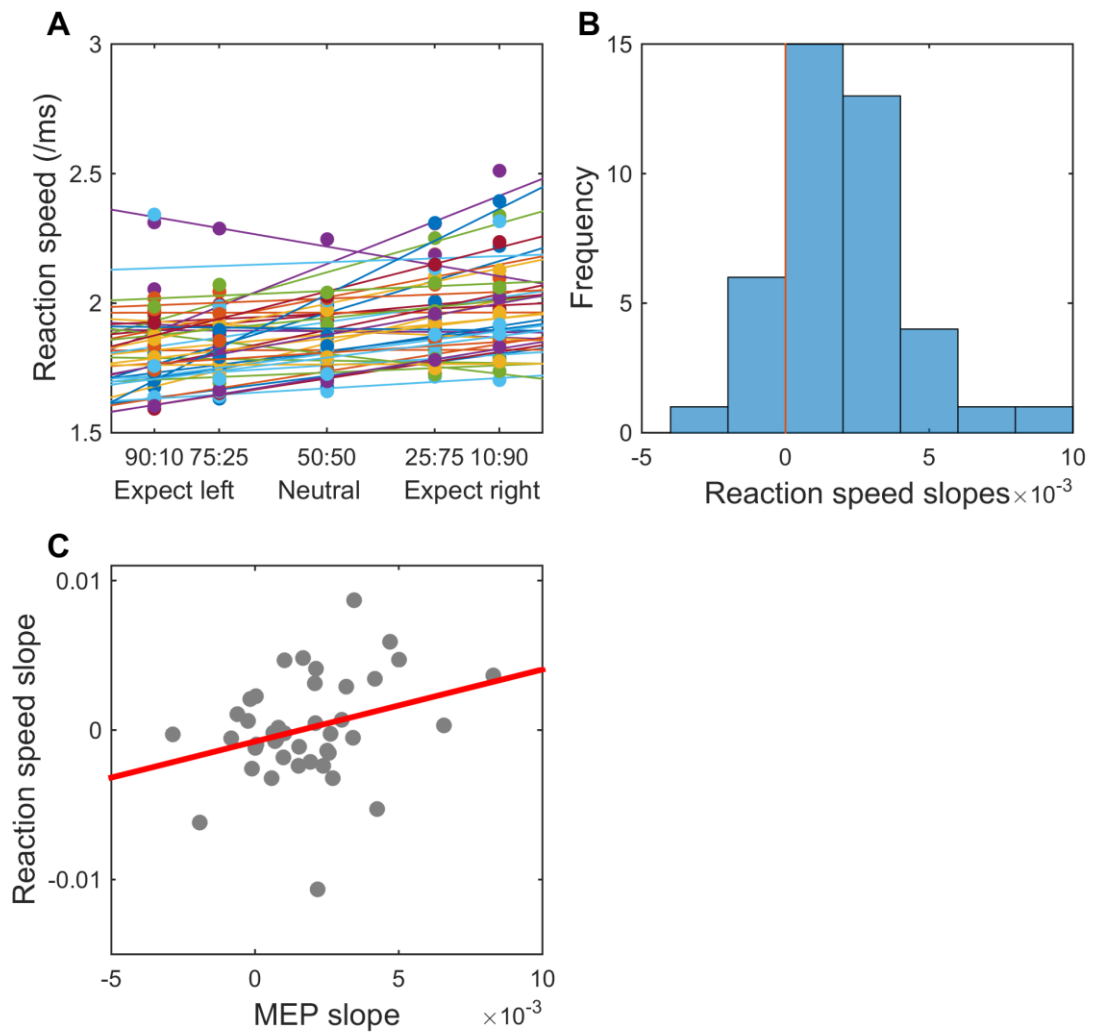


Figure 2.16 No relationship between reaction speed slope and MEP slope. (A) shows individual subject data for reaction speed (inverse of reaction time) in right-choice trials only. As expected, the trend is towards faster right-hand responses in conditions where there is a greater likelihood of a rightward stimulus. Linear regression models were fit to these data within each subject. (Each subject's data is shown by coloured points, with a line indicating the fits of the linear regression.) A histogram of the slopes (β) of these relationships is shown in (B). These slopes were significantly different from 0, $t(40) = 5.36$, $p < .001$, $d = 0.84$. Figure 2.15C and D show this analysis repeated for right FDI MEP at stimulus onset. (C) shows the slopes from plot (B) against the MEP slopes, and the line of the linear regression fit. The regression showed no significant relationship, $F(1, 39) = 3.78$, $p = .059$, $R^2 = .088$.

2.3.3.4. MEPs over time

In addition to recording MEPs at stimulus onset, MEPs were recorded at 25% and 45% of the subject's mean reaction time (as measured behaviourally before the first experimental block). This data is shown plotted by original timepoint (effectively aligned to stimulus onset;

‘stimulus locked’) and left/right choice in Figure 2.17 and further divided by prior cue condition in Figure 2.18.

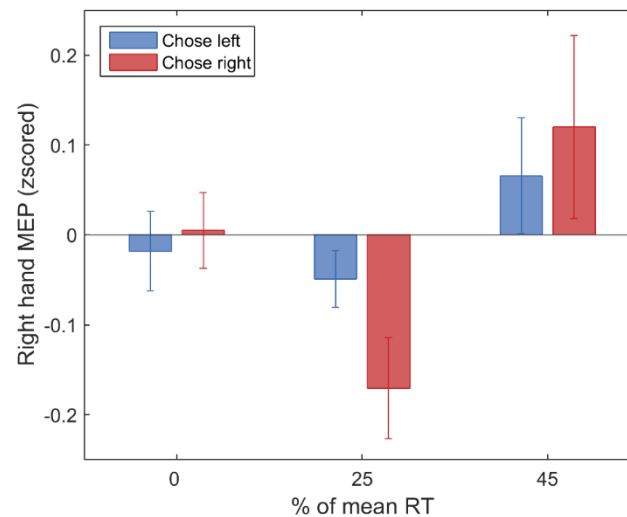


Figure 2.17 MEPs plotted relative to stimulus onset are suppressed at the second timepoint and elevated at the third. To look at broad patterns of MEP size over time, these were plotted according to timepoint of stimulation. There was a main effect of timepoint, $F(1.55, 62.0) = 4.68$, $p = .020$, $\eta^2_p = .11$, no main effect of choice, $F(1, 40) = 0.22$, $p = .64$, $\eta^2_p = .006$, and no interaction between these factors, $F(2, 80) = 2.00$, $p = .14$, $\eta^2_p = .048$. A post-hoc test showed that MEPs at the second timepoint were significantly lower than at the first timepoint, $t(40) = 2.23$, $p = .032$, $d = 0.35$. Although some studies have reported greater MEP suppression in the chosen effector, MEPs at the second timepoint did not differ by ultimate choice, $t(40) = -1.82$, $p = .61$, $d = 0.081$. Error bars show standard error of the mean.

A two-way timepoint x hand ANOVA on this data (i.e. reflecting the data as split in Figure 2.17) revealed a main effect of timepoint, $F(1.55, 62.0) = 4.68$, $p = .020$, $\eta^2_p = .11$, no main effect of choice, $F(1, 40) = 0.22$, $p = .64$, $\eta^2_p = .006$, and no interaction between these factors, $F(2, 80) = 2.00$, $p = .14$, $\eta^2_p = .048$. When the data are plotted locked to the stimulus, there is a trend for MEPs to be suppressed at the second timepoint relative to the first, and elevated at the third.

Suppression of MEPs shortly after stimulus onset has been recorded previously (Greenhouse et al., 2015b). To find whether this phenomenon was present in the data reported here, all data (i.e. collapsed across ‘Chose left’ and ‘Chose right’) at the first timepoint was compared to the second timepoint. MEPs at the second timepoint were indeed significantly

lower than at the first timepoint, $t(40) = 2.23$, $p = .032$, $d = 0.35$. Suppression of MEPs has been previously found to be strongest in the effector used compared to non-chosen effectors (Duque and Ivry, 2009). Again, to find if this applied in this dataset, ‘Chose left’ and ‘Chose right’ data at the second timepoint were compared. These were not significantly different, $t(40) = -1.82$, $p = .61$, $d = 0.081$.

A three-way timepoint x prior cue x hand ANOVA on the data (i.e. reflecting the data as split in Figure 2.18) showed no significant main effects or interactions. This ANOVA is given in full in Appendix II.

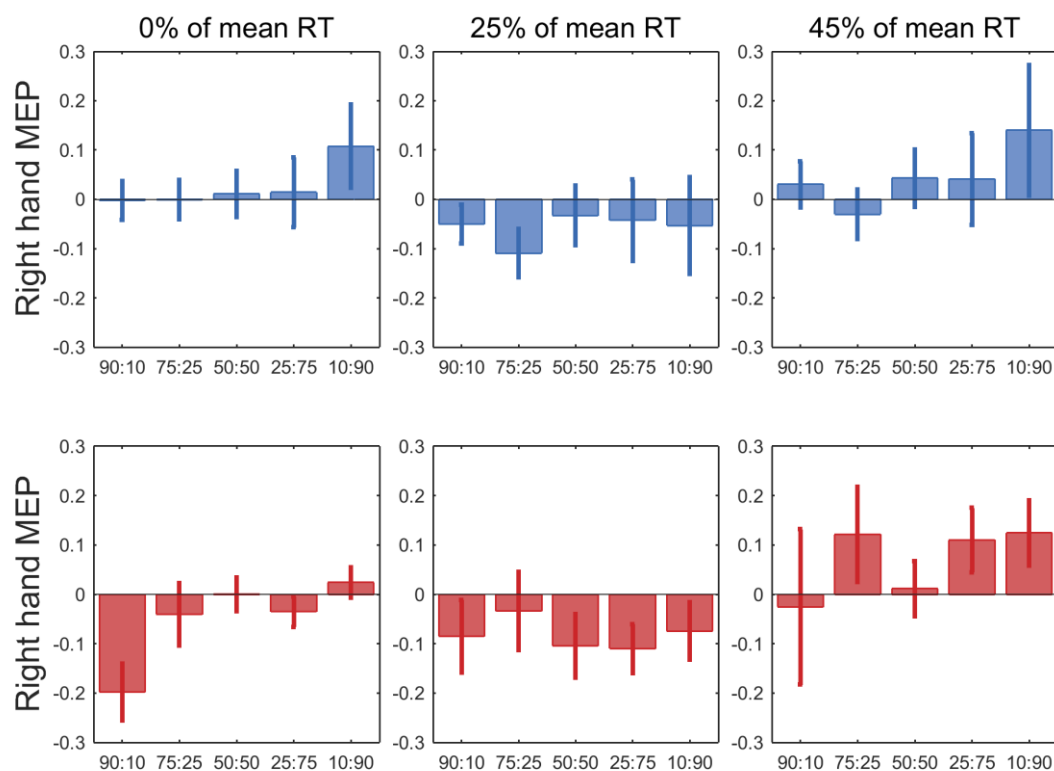


Figure 2.18 Stimulus-locked MEPs do not vary by condition in when split by timepoint x prior cue x hand. Here, data from Figure 2.17 is shown further split by prior cue condition. MEPs, aligned to stimulus onset, are plotted by choice made (upper panels: chose left; lower panels: chose right), timepoint (left panels: 0% of mean RT; middle panels: 25% of mean RT; right panels: 45% of mean RT), and prior cue condition. There are no significant main effects or interactions in an ANOVA of this data (see Appendix II). Error bars show standard error of the mean.

To understand better how MEPs evolved over the reaction time, they were ‘locked’ to the response (i.e. time of stimulation was recalculated to

be relative to reaction time on each trial). These data were then smoothed by calculating a moving average with a window width of 150ms, stepped across the data in steps of 20ms. These moving averages were meaned across subjects. These data are plotted in Figure 2.19 and Figure 2.20, split according to which hand was chosen on that trial. For each prior condition, the trials in which the right hand was chosen have been plotted against left-choice trials for the *opposite* prior category. For example, right-choice trials in the 10:90 condition have been paired with left-choice trials in the 90:10 condition. This pairs together ‘strongly expect right, chose right’ with ‘strongly expect left, chose left’ trials, and, as MEPs were always from right FDI, thus shows CSE in the chosen and unchosen effector for the strongest expectation condition. Thus, the most intuitive and useful comparison is produced by matching the data across conditions in this way. This produces three panels in which subjects had a neutral or biasing prior which they chose with (Figure 2.19), and two in which they had a biasing prior they chose against (Figure 2.20).

For each plot, a paired t-test was conducted on the data at each point (left-choice vs right-choice). These tests were corrected for multiple comparisons using a False Discovery Rate procedure (Benjamini and Hochberg, 1995). The first point at which the two traces are significantly different in a test surviving correction is shown with a grey line on the graph in Figure 2.19. This point comes progressively later as the cue becomes less biasing across conditions: 0.245, 0.205, and 0.185 ms prior to the response.

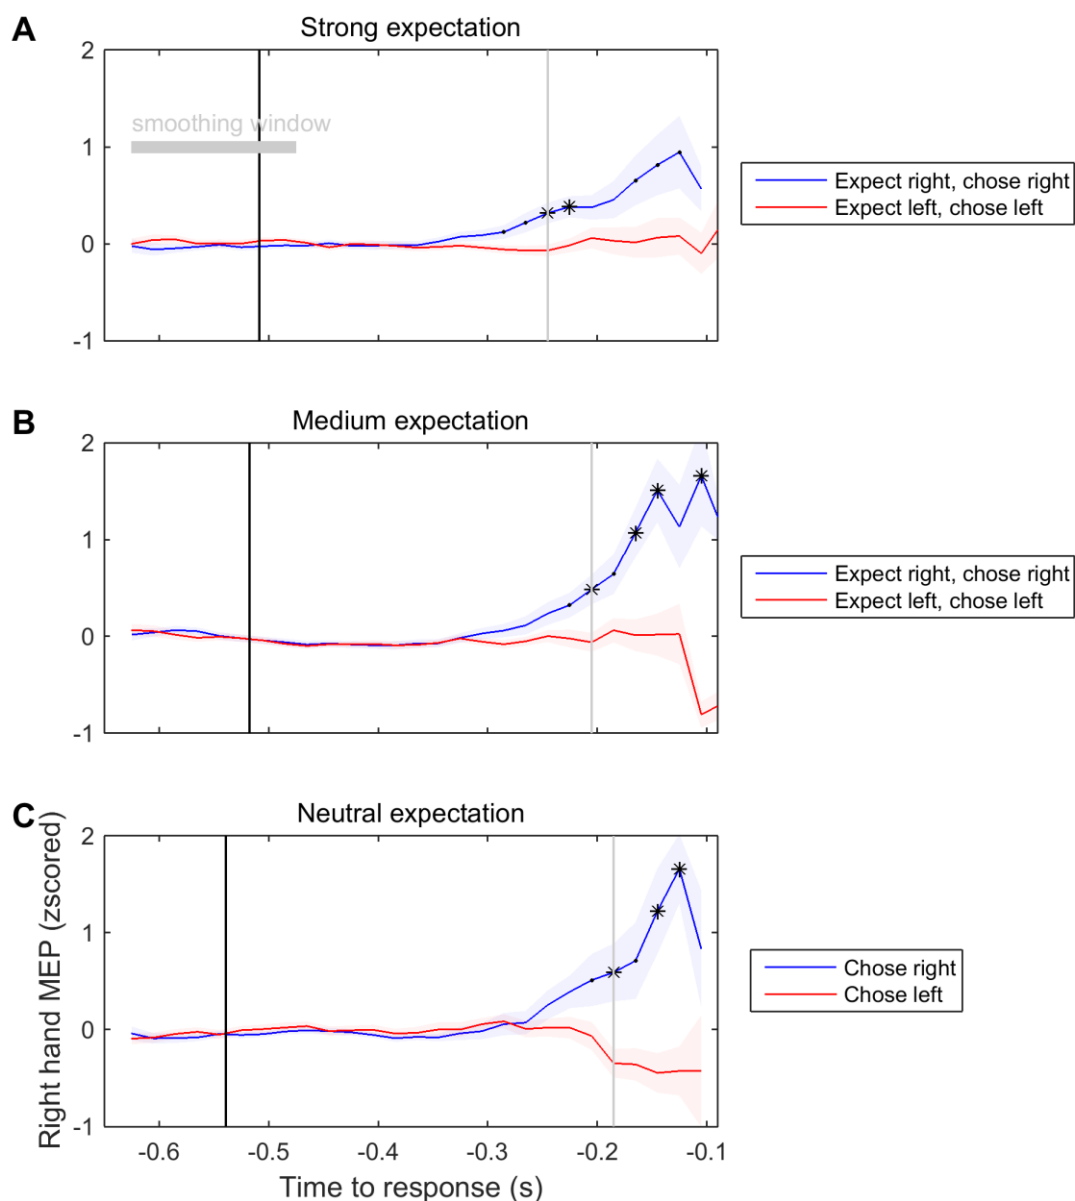


Figure 2.19 Response locked MEPs suggest activity may diverge later when prior expectation is stronger. Data have been re-organised so that MEPs are plotted relative to time of response rather than time of stimulus onset. MEPs were selected from trials in which subjects chose according to prior cue. Data were then smoothed by a moving average filter (window width: 150 ms; step size: 20 ms). Paired t-tests were conducted at each step. Significant t-tests are shown on the chose right trace with a black point (.); significant tests surviving FDR correction are shown with a black asterisk (*). Vertical grey lines show first significant point surviving correction. The first point of significant divergence is earliest relative to response in the neutral condition and latest in the strong expectation condition. See text for a discussion of the robustness of this result. Vertical black lines show mean reaction time for that condition relative to response; in other words, they show the average time of reaction start.

This suggested that there might be an effect whereby ‘Chose right’ MEP activity differentiates itself from ‘Chose left’ MEP activity earlier relative to the response when the prior expectation to move in that direction is

stronger. It is difficult to examine this hypothesis directly in this dataset, which was not designed to look in detail at MEPs over time. Because sampling occurred at two fixed timepoints after stimulus onset, the distribution of datapoints once expressed relative to response is determined by reaction time variability. The data are sparse close to the response, and some subjects do not have any data in this time-period (as shown by wider error bars close to response in Figure 2.19 and Figure 2.20). To attempt to test the hypothesis, the last timepoint at which there was data in all subjects in the response-locked analysis was used, which was 0.325 s prior to response. For each subject, for the strong, medium and neutral expectation conditions, the difference between 'Chose right' and 'Chose left' MEPs at this timepoint was calculated, and then ranked the resulting values within subjects. These ranks were tested to see if there was a greater tendency for e.g. the strong expectation condition to have a greater left-right difference; however, Friedman's test revealed no difference between the conditions, $X^2(2, 80) = 1.76$, $p = .42$. In sum, the data does not support the hypothesis of an earlier separation at a within subjects level, although this could be due to sparse data in the critical time-period rather than a genuine absence of effect.




Figure 2.20 Response locked MEPs on trials in which subjects chose against the prior cue do not show statistically-significant divergence. Data have been re-organised so that MEPs are plotted relative to time of response rather than time of stimulus onset. MEPs were selected from trials in which subjects chose in the opposite direction to that indicated by the prior cue. Data were then smoothed by a moving average filter (window width: 150 ms; step size: 20 ms). Paired t-tests were conducted at each step. Significant t-tests are shown on the chose right trace with a black point (.); significant tests surviving FDR correction are shown with a black asterix (*). Vertical black lines show mean reaction time for that condition relative to response; in other words, they show the average time of reaction start.

2.3.3.5. MEP-behaviour correlations for late MEPs

Analyses reported above failed to show a relationship between (a) modelled bias parameter or (b) reaction speed and MEP size at stimulus onset. Having failed to find this relationship with early MEPs at the time of RDK onset, the analysis was repeated with late MEPs (third timepoint; 45% of reaction time). This was to investigate whether late corticospinal activity predicts behaviour on this task, given that early corticospinal activity does not.

Thus the analysis described in Section 2.3.3.3 was repeated for MEPs recorded at the third timepoint. Firstly, slopes of bias parameter

regressions were regressed against MEP slopes (Figure 2.21). There was no relationship between these variables, $F(1, 39) = 0.36$, $p = .55$, $R^2 = .009$. Secondly, slopes of reaction speed regression were regressed against MEP slopes (Figure 2.22). Here, there was a significant relationship, $F(1, 39) = 8.61$, $p = .006$, $R^2 = .18$. This indicates that subjects who showed a steeper trend in late MEPs across prior conditions also showed a steeper trend in right-choice reaction speed. Note that the distribution of MEP slopes at the late timepoint was not significantly different from 0 ($t(40) = 1.09$, $p = .28$, $d = 0.17$), indicating no significant tendency for subjects to show a positive relationship between right FDI MEPs at the late timepoint and rightward expectation. Thus the MEP slope-reaction speed slope relationship exists despite this lack of a positive bias in the slopes themselves.

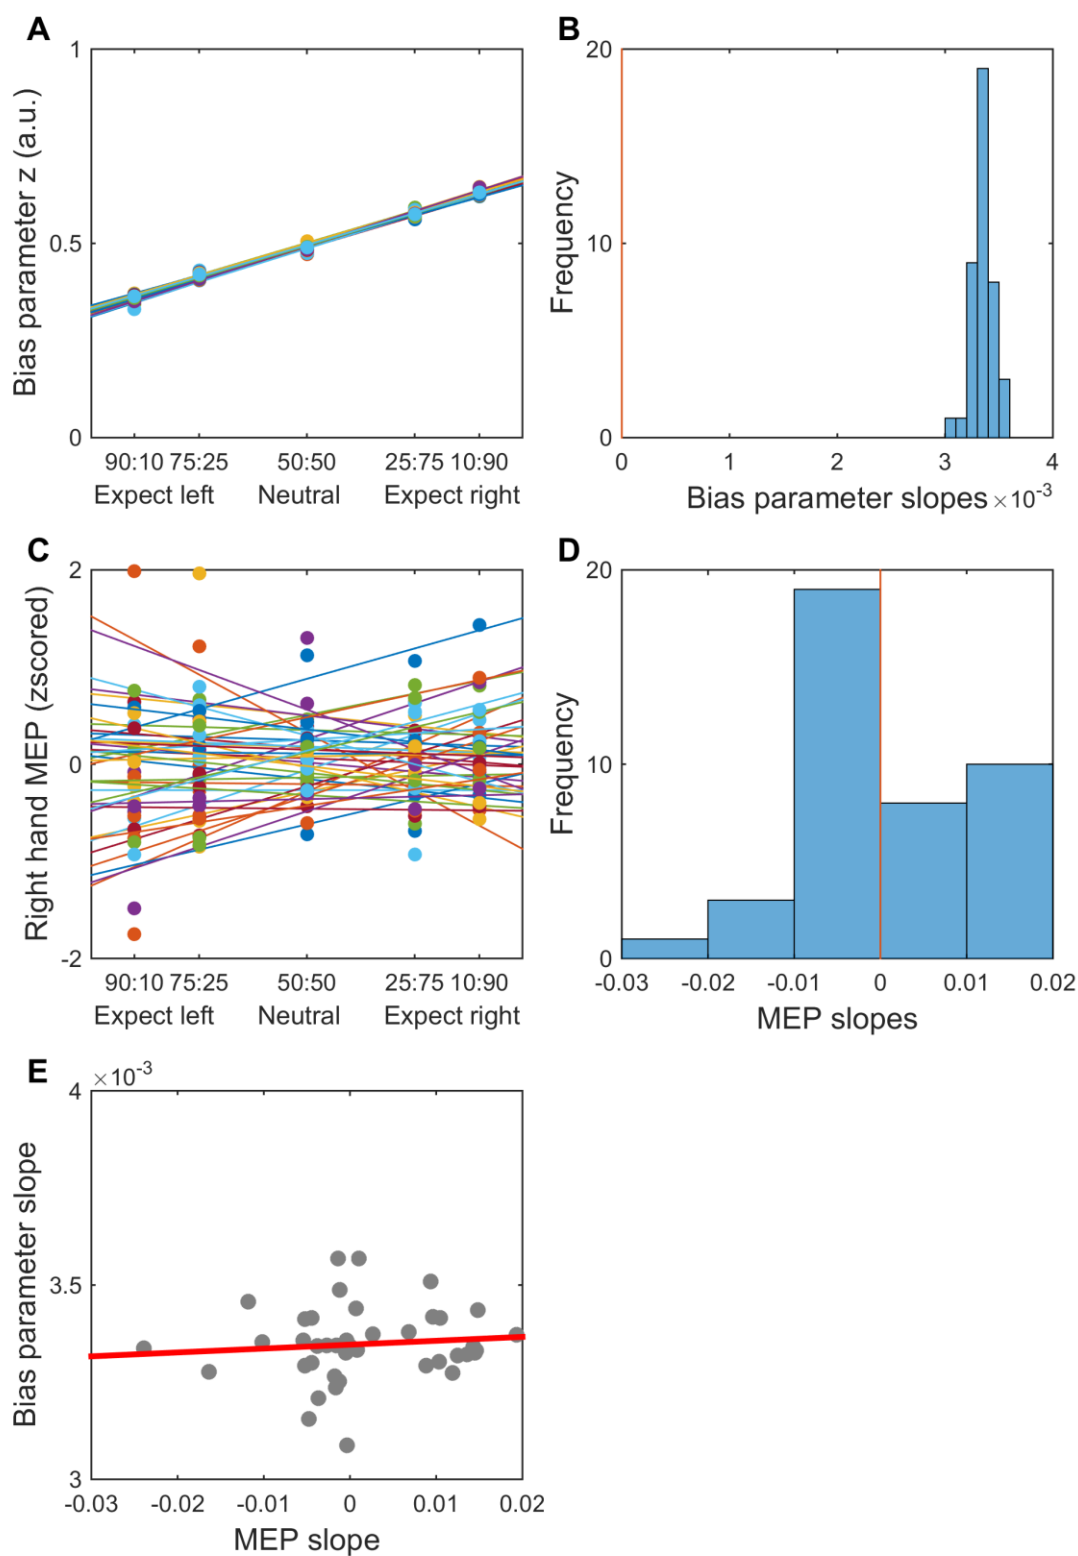


Figure 2.21 No relationship between bias parameter slopes and late MEP slopes. Plots (A) and (B) have previously been displayed in Figure 2.15 and are repeated here for visual convenience. (A) shows the bias parameters, z , produced by the drift diffusion model, for each condition and each subject. A simple linear regression was conducted on these points within each subject. (Each subject's data is shown by coloured points, with a line indicating the fits of the linear regression.) A histogram of the slopes (β) of these relationships is shown in (B). These slopes are significantly

different from 0, $t(40) = 223$, $p < .001$, $d = 34.9$. The positive slopes indicate the model finds that subjects are, as expected, more biased towards a right-hand response the more the prior cue indicates a right-hand response is likely. (C) shows an equivalent plot to (A) for MEP in the right FDI at the third timepoint. (D) shows the slopes (β) of within-subjects linear regressions on these data. These slopes are not significantly different from 0, $t(40) = 1.09$, $p = .28$, $d = 0.17$. In (E), for each subject, the bias parameter slope has been plotted against the MEP slope, to see whether subjects who show steeper relationships between experimental condition and modelled parameter also show steeper trends in MEPs. There was no significant relationship, $F(1, 39) = 0.36$, $p = .55$, $R^2 = .009$.

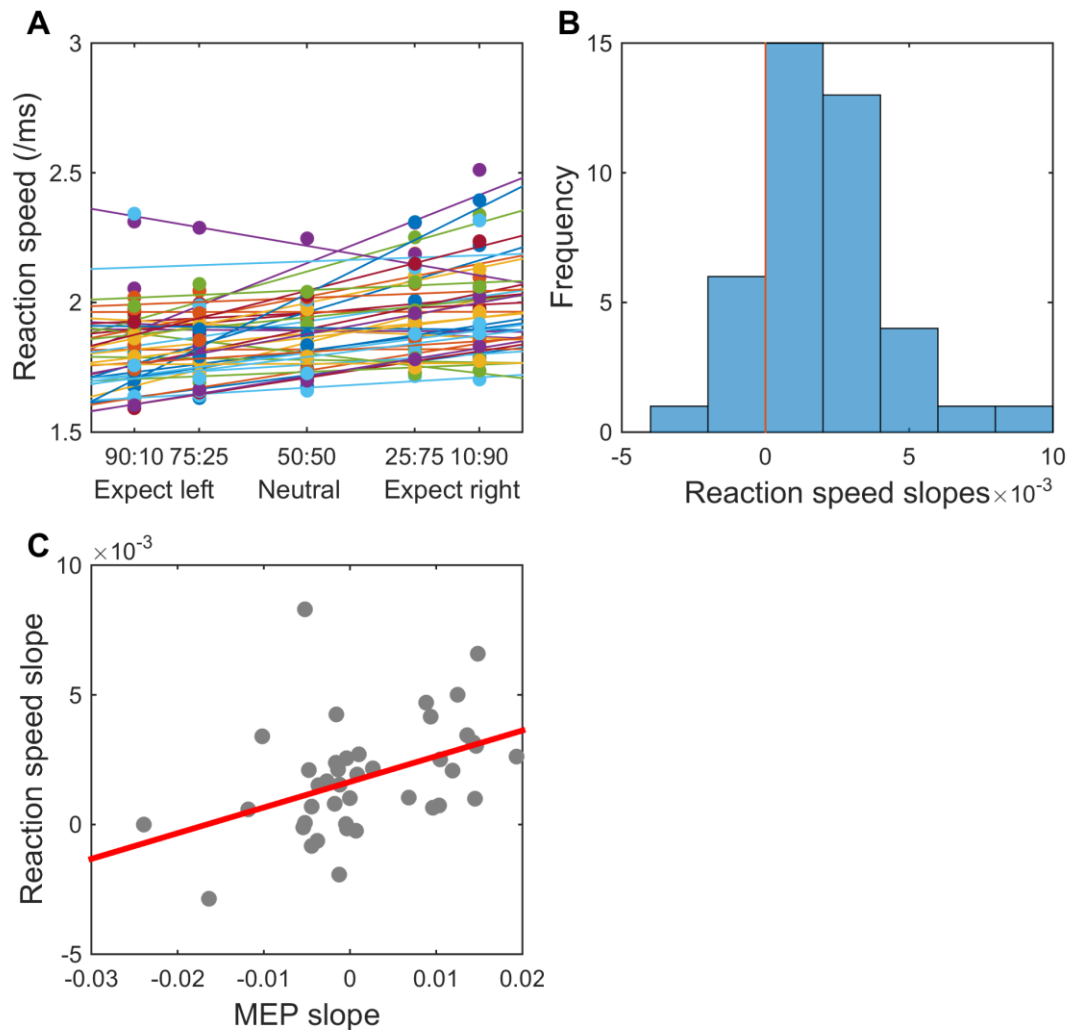


Figure 2.22 Positive relationship between reaction speed slope and late MEP slope. Plots (A) and (B) have previously been displayed in Figure 2.16 and are repeated here for visual convenience. (A) shows individual subject data for reaction speed (inverse of reaction time) in right-choice trials only. As expected, the trend is towards faster right-hand responses in conditions where there is a greater likelihood of a rightward stimulus. Linear regression models were fit to these data within each subject. (Each subject's data is shown by coloured points, with a line indicating the fits of the linear regression.) A histogram of the slopes (β) of these relationships is shown in (B). These slopes were significantly different from 0, $t(40) = 5.36$, $p < .001$, $d = 0.84$. (C) shows the slopes from plot (B) against the MEP slopes in the third timepoint (see Figure 2.21C and D), and the line of the linear regression fit. The regression showed a significant relationship, $F(1, 39) = 8.61$, $p = .006$, $R^2 = .18$.

2.3.3.6. MEPs in control trials

A previous study (Klein-Flugge and Bestmann, 2012) compared MEPs in simple and choice reaction time tasks, using the timing of separation of corticospinal excitability by eventual effector (left or right) in the simple reaction time task to measure a time for executing a simple motor response. The authors then analysed MEPs only prior to this ‘motor response’ epoch. This allowed them to show that MEP activity was differentiated by eventual choice prior to this epoch on choice trials, suggesting that this information leaks into corticospinal excitability whilst the decision is ongoing, substantiating ideas about parallel processing.

The experiments in this chapter included control trials (in which an arrow was superimposed on the RDK to indicate direction of dot motion, obliterating the need for an evidence accumulation phase) in order to allow for a similarly-motivated analysis.

MEPs are shown recorded in control trials only, plotted aligned to stimulus onset, separated by eventual choice (Figure 2.23). This showed a separation by eventual choice at the third timepoint ($t(27) = 2.57, p = .048, d = 0.49$), but not the first ($t(40) = 0.14, p = .89, d = 0.021$) or second ($t(40) = 1.48, p = .22, d = 0.23$).




Figure 2.23 Control trial MEPs show separation by activity at latest timepoint, but not at stimulus onset or interim timepoints. MEPs here were recorded during control trials, in which an arrow was displayed on the screen, removing the need for difficult RDK direction detection and making reaction times faster. MEPs have been plotted by eventual button press response (left or right). As indicated on the graph, right hand MEPs are significantly large in right-choice trials than left-choice trials by the third timepoint, at 45% of mean reaction time, but not at earlier timepoints. T-tests were corrected for multiple comparisons using the False Discovery Rate procedure (Storey, 2002).

Because there was no differentiation of activity by eventual choice at the first or second timepoint, it could be safely assumed that at least the epoch until the second timepoint could be considered ‘decision time’, in which MEP activity does not reflect motor preparation. For each subject, this decision time was calculated and MEPs in this epoch only on non-control trials were analysed for differences by eventual choice. In Klein-Flugge and Bestmann’s analysis (2012), the authors similarly defined a decision time epoch using forced choice trials and then found activity relating to choice in this epoch in the choice trials. However, in the experiment presented here, there was no difference by choice of MEPs in the decision time at either the first ($t(40) = 0.63, p = .53, d = 0.098$) or the second ($t(40) = 0.99, p = .49, d = 0.16$) timepoints. These t-tests were corrected for multiple comparisons using the False Discovery Rate procedure (Storey, 2002).

2.3.4. Comparison to baseline

In all experiments, in addition to the within-block MEPs presented so far, a pre-block, baseline MEP measure was also taken. It has been shown repeatedly that under some experimental conditions, MEPs can be suppressed prior to movement relative to a baseline measure (e.g. Hasbroucq et al., 1997). It has been hypothesised this is the result of a broad-based inhibition and is part of a process of preventing premature movement (Duque et al., 2010).




Figure 2.24 MEPs are elevated relative to pre-block baseline measure. This analysis sought a metric of whether MEPs recorded at the time of stimulus onset were suppressed relative to baseline for each experiment. For the three experiments, mean MEP at stimulus onset only was calculated relative to baseline measure on that block. In fact, MEPs on all three experiments were approximately double the baseline measure. The first two MEPs were excluded from the baseline measure. Error bars show standard error of the mean.

I hypothesised that MEPs at stimulus onset were suppressed in the experiment and this was responsible for the absence of modulation by prior cue level (see Discussion). I thus calculated MEPs relative to the pre-block baseline in the three experiments, as this measure may give some indication of the suppression (although note MEP suppression is typically calculated relative to a within-block, intertrial baseline (e.g. Duque and Ivry, 2009)).

Figure 2.24 shows MEPs at stimulus onset relative to the appropriate block baseline measure, collapsed across all prior cue conditions. In all three experiments, MEPs were higher than at baseline. The implications of this are examined in the discussion.

2.4. Discussion

This chapter investigated whether an expectation about which direction to move modulates corticospinal excitability in the early period prior to the evidence-accumulation phase of the decision. Three datasets investigating this question were presented. Contrary to the hypothesis, there was no effect of prior expectation on corticospinal excitability prior to stimulus onset. There was some indication in the data that a stronger prior might lead to earlier right-left separation in activity relative to response, but this was not statistically robust; the experiment was not designed to look at this time period in detail and data here was sparse.

This discussion will consider a number of possible explanations for the null result. The first is that there is no representation of prior expectation in motor cortex; in other words the proposed hypothesis is false. A second explanation is that MEP suppression, which is known to occur either just before or just after a movement is cued (Greenhouse et al., 2015b), affected the expectation conditions indiscriminately and thus obscured an effect. A third explanation is that MEPs are too variable to assess an effect of this magnitude. I will examine these explanations in further detail and then discuss the positive findings of this experiment.

2.4.1. Was ‘impulse suppression’ in these experiments the reason for a lack of effect on MEPs at stimulus onset?

What is the relevance of the ‘impulse suppression’ to the experiments presented in this chapter? Because impulse suppression affects all effectors that might be involved in the upcoming movement (Duque et al., 2010; Greenhouse et al., 2015b), it is considered to be a broad inhibitory signal, and thus could have obscured differences between conditions in the data.

The Introduction to this chapter detailed the three experimental designs aimed at minimising the effect of MEP suppression. In Experiment 1, a fixation cross was always presented 400 ms before stimulus onset. A fixation cross is common in both perceptual tasks (e.g. Shadlen and Newsome, 2001b) and instructed delay TMS studies (e.g. Labruna et al., 2013). However, the fixation cross may have acted as a warning cue and thus lead to MEP suppression by the time of stimulus onset. Experiment 2 therefore had no fixation cross. However, it is possible in this experiment that the presentation of prior cues - which were presented a fixed 2,950 ms prior to stimulus onset - themselves acted as a warning cue (Sinclair and Hammond, 2009; Duque et al., 2010). Experiment 3 was similar to Experiment 2, except that the duration between the prior cue and the stimulus was jittered. There is no evidence in the literature that jittering the length of the delay period reduces suppression; however, pilot data in three subjects (not presented separately in chapter; incorporated into dataset for Experiment 3) which showed an effect of prior cue at stimulus onset, indicating that reducing temporal expectation by jittering the length of the delay might unmask an effect of prior expectation. Furthermore, it has repeatedly been found that long pre-movement delays (which are considered less 'predictable' than short pre-movement delays) obliterate MEP suppression (Touge et al., 1998; Hasbroucq et al., 1999, 1999; Tandonnet et al., 2010).

So was there MEP suppression in the experiment which obscured any effect of prior cue? There are several reasons to think MEP suppression was not present in Experiments 2 and 3.

Firstly, MEPs are high relative to a pre-block baseline (Figure 2.24). This is unsatisfactory as a direct measurement of suppression, because, in studies that investigate MEP suppression, MEP amplitude is typically compared to a baseline measured during the intertrial interval. This was

not part of the experimental design (although with hindsight, this would have been a useful metric). However, some studies have additionally compared to a pre-block baseline (Greenhouse et al., 2015b). Repeating the process on this data shows that MEPs at stimulus onset are facilitated, not suppressed.

Secondly, there was a delay of at least 2,950 ms between prior cue and stimulus onset. It has been repeatedly demonstrated that long delays (2 – 2.5 s compared against 0.5 s) in instructed delay tasks obliterate impulse suppression (Hasbroucq et al., 1997; Touge et al., 1998; Tandonnet et al., 2010). This suggests that the delay used was long enough to obliterate MEP suppression in the delay period. Note however that the papers cited above used simpler tasks and the cues prior to the delay period fully specified the movement that would be required at the time of the imperative; thus, they are not directly comparable.

Thirdly, in the combined dataset for Experiments 2 and 3, MEPs are suppressed at the second timepoint relative to the first (Figure 2.17). This has been observed after the imperative in no-delay reaction time tasks (Greenhouse et al., 2015b), and is interpreted as being the same MEP suppression phenomenon as seen in instructed delay tasks, with a shifted timing in response to the specifics of the task. The presence of a post-stimulus suppression in the data suggests that MEPs were not suppressed at stimulus onset. (It is possible that there are two successive stages of suppression, but this has not been described in the literature.)

In sum, there is no compelling reason to think that MEP suppression was the reason for the null result. It is possible that it was present, but a number of lines of reasoning point towards its absence.

2.4.2. Was anatomical specificity of competitive representations the reason for a lack of effect on MEPs at stimulus onset?

If broad suppression of MEPs was not the cause of the null result, then why was there an absence of hypothesised activity in response to the prior cue? One possibility is that stimulus information dynamically biases motor activity in dorsal premotor cortex but not primary motor cortex. The key studies demonstrating biasing of motor plans by cognitive information when this cognitive information was relevant for decision found this activity in PMd (Pastor-Bernier and Cisek, 2011; Cisek and Pastor-Bernier, 2014); the authors did not find the same phenomenon in M1 (Cisek; personal communication).

However, other studies have successfully used TMS to demonstrate response competition (Michelet et al., 2010) and the impact of cognitive variables such as expected value (Klein-Flugge and Bestmann, 2012) on corticospinal excitability. Furthermore, the specificity of TMS with a figure-of-eight coil for M1 remains unclear. MEPs are composed of D ('direct') and I ('indirect') waves (Di Lazzaro et al., 1998a). D waves are generated directly from the pyramidal cell outflow tract (Patton and Amassian, 1954), but I waves are believed to be dependent on corticocortical circuits (Di Lazzaro et al., 1998b). Invasive electrical stimulation of premotor cells alone elicits I waves and an EMG response, but more weakly than when M1 is stimulated (Cerri, 2003; Shimazu et al., 2004; Schmidlin et al., 2008), and mini-coil TMS to PMd will not evoke MEPs (Groppa et al., 2012). However, TMS or electrical stimulation of PMd enhances later I waves and thus MEPs evoked from M1 (Cerri, 2003; Shimazu et al., 2004; Groppa et al., 2012). Thus a figure-of-eight coil activating both PMd and M1 could result in MEPs being evoked which depend on PMd activation. The hand motor hotspot (the site at which maximum MEP size is elicited) has been found to be located in M1 in just

under half of subjects, with the remaining half having a more anterior premotor location (Ahdab et al., 2016).

Hence it is not possible to say definitively that an effect of prior on motor plans in PMd but not M1 would have led to the results observed. A paired pulse protocol looking at the strength of PMd to M1 facilitation (Tokimura et al., 1996; Ziemann et al., 1998) could be used to test directly the idea that prior information is represented in PMd.

It is possible that there was no effect of prior information on corticospinal excitability at stimulus onset because prior information is not transmitted to any motor areas until action preparation begins. However, a previous magneto-encephalography study with a similar motivation to the experiments presented in this chapter found expectation was represented in motor areas from around one second before the cue to move (de Lange et al., 2013).

2.4.3. Was the experiment underpowered?

A final possible reason for the null result is that the study was underpowered. However, the sample size of 41 subjects (Experiments 2 and 3 combined) is approximately double what similar studies have used in the past. This sample size gives an 80% power to detect a difference between means of 0.45 standard deviations, using a paired t-test. It's possible there was an effect present that was smaller than this, which would be very hard to measure using a TMS experiment. The differences observed in MEP size at stimulus onset between conditions were less than 0.1 standard deviations; if these reflected genuine differences, the study would not have been powered to detect them.

2.4.4. Temporal variability of MEPs

Between Experiment 1 and Experiment 2, alongside other changes to the design highlighted elsewhere, I switched from a design where the experimental conditions were varied by block to an ‘interleaved’ design where they were varied by trial. This was as a result of gaining experimental experience of the variability of MEPs. There is also evidence in the literature that MEPs are highly variable, with exact variability depending on factors such as coil and stimulation intensity (Kiers et al., 1993). Variation in physiological factors might also play a role; post-exercise fatigue decreases MEPs (Brasil-Neto et al., 1993). More recent studies have begun to attempt to explain the determinants of this variability (Klein-Flügge et al., 2013; Goetz et al., 2014). Importantly, MEP variability has a strong temporal component, with the first few recorded MEPs being much larger than subsequent ones (Brasil-Neto et al., 1993; Schmidt et al., 2009). One paper found that early MEPs are approximately log-normally distributed, whilst subsequent MEPs are approximately normal (Schmidt et al., 2009). The authors recommend excluding the first 20 MEPs when analysing experiments; however, this practice has not been adopted in the literature cited in this chapter. I believe these findings support the idea of using interleaved designs in preference to blocked designs in cognitive MEP experiments.

2.4.5. Differences between conditions at stimulus onset when split by choice

Figure 2.13 shows MEPs at stimulus onset split by choice. Unexpectedly, even though subjects had not seen the stimulus at this stage, in the 90:10 (strongly expect left) condition, MEPs were lower when subjects chose right (i.e. against the expectation) compared to when they chose left. The data pattern in the opposite condition appears to reverse this pattern, although the difference is not significant here. (Note that there are very

few trials in these conditions as the stimulus only moves against the prior in 10% of cases.) These findings are counterintuitive because one might expect a higher right hand MEP to predict right hand choice, but the right hand MEP was lower when subjects managed to overcome a strong expectation that the stimulus would be leftward and move the right finger. Why is this? One speculative explanation is that these trials reflect a delayed or failed impulse suppression process. Under this explanation, differences between conditions are usually obliterated by impulse suppression by the time of stimulus onset, but occasionally this process fails or is late, reflecting a delay to the preparation process. This delayed preparation means the action is further from being released and thus there is more time to overcome a bias on the basis of evidence presented by the stimulus. The effect only shows up in the most extreme conditions because these are the hardest to overcome and are only overcome in trials with delayed impulse suppression. Note that this explanation relies critically on the presence of MEP suppression at stimulus onset, and above I have presented reasons why this might not be present in these data. Thus this explanation is presented as speculation.

2.4.6. Relationship between late MEPs and reaction speed

The results showed that subjects who showed stronger trends (steeper regression slopes) in right FDI MEPs measured at the third timepoint showed stronger trends in right-hand reaction speed across prior cue conditions. There was no relationship (regression was non-significant) for the same analysis using modelled bias parameter rather than the direct behavioural measure of reaction speed. Interestingly, this relationship was present, even though the MEPs slopes themselves were not mostly positive (i.e. there was no tendency for most subjects to show a positive relationship between third-timepoint MEP and prior cue). This highlights the value of analyses which link individual subject behavioural

measures to corticospinal excitability (Klein-Flugge and Bestmann, 2012). The analyses indicate that MEP size at stimulus onset does not determine (or is not determined by) reaction time, but MEP size at later timepoints does. (The same analysis with stimulus-onset MEPs showed no effect).

It is interesting that modelling a bias parameter in each subject weakened rather than strengthened these analyses. An fMRI study on prior expectation by Forstmann et al (2010) found results only when a similar modelled bias parameter was included in the fMRI analysis as a covariate. This highlighted the value of modelling in these experiments and influenced the choice of analyses for this experiment. Why was the modelling less valuable in this case? Visual inspection of the modelled bias parameters shows that inter-subject variability is very low and every subject shows a steep trend in bias parameters across prior cue conditions. This contrasts with the reaction speed measure, which is much noisier. This is an effect of the hierarchical modelling, which assumes individual subject parameters are drawn from a group distribution. This noise-reduction effect might sometimes be advantageous, but I speculate that this reduced inter-subject variability to the point where an analysis based on individual differences was not useful.

2.4.7. Suitability of TMS for experiments of this nature

At the time of designing the experiments presented in this chapter, there were no other studies looking at evidence-accumulation based decisions. In the interim, Hadar et al (2016) published a TMS study looking at how MEP size over the reaction time depends on stimulus strength. In short, their study was similarly motivated by studying the correlates of biased competition in a perceptual decision using TMS, but they chose to focus on stimulus strength, which is considered to affect the rate of evidence

accumulation, rather than prior expectation, which is considered to cause a biased starting point.

Similarly to the analysis presented in this chapter, they used smoothing to generate a continuous MEP trace. They find an earlier divergence of activity for hard over easy decisions in a response locked analysis. However, this result was achieved only after a post-hoc exclusion of trials which did not produce strong reaction time differences. Furthermore, each of their eight participants undertook an unusually high number of trials (1,920, compared to 400 in this study). Thus this study also struggled to find hypothesised effects even with large statistical power. It is interesting to consider whether, despite its advantages (temporal and effector specificity), high MEP variability and the complexity of unpicking multiple component processes which contribute to the MEP, mean TMS is not, in fact, a good method for this kind of experiment.

2.4.8. Replication of previous studies

How do the results presented in this chapter compare to those in previous studies? Firstly, the canonical finding that corticospinal excitability in the chosen effector increases prior to movement was replicated (e.g. Leocani et al., 2000). Secondly, a suppression of MEPs shortly after stimulus onset (Greenhouse et al., 2015b), strongest in the chosen effector (Duque and Ivry, 2009), has been shown previously and this pattern was found in the data from Experiments 2 and 3 (Figure 2.17). Thirdly, in the behavioural measures, a more biasing prior expectation was shown to make subjects faster and more accurate, which is another standard result (e.g. Carpenter and Williams, 1995).

A finding this study failed to replicate was that decision-related activity is present in corticospinal excitability in the ‘decision epoch’, which was previously shown in a value-based choice task (Klein-Flugge and

Bestmann, 2012). However, the experiments in this chapter had less precision in this analysis, as only three timepoints were tested, whilst the previous study tested six. This was necessary to the design of the experiments in this chapter as there were more kinds of decision type (five levels of prior cue) and thus needed fewer timepoints to give statistical power. So loss of temporal resolution, or a difference in decision type (value-based vs perceptual decision) could have been reasons for the discrepancy in findings.

Bestmann et al (2008) conducted a study in which the main focus was the effect of entropy and surprise on corticospinal excitability, but they also briefly report that corticospinal excitability was higher in blocks in which the conditional stimulus had a higher likelihood of validly cueing the imperative stimulus. This is analogous to the experiments in this chapter, in which an effect of expectation at stimulus onset was not found. It is difficult to say why this is. One possibility is that it is due to a timing difference, as these authors timed TMS at 200 ms prior to the imperative stimulus to move.

2.4.9. Alternative experimental approaches

The experiments in this chapter attempted to measure corticospinal excitability in order to investigate how it was modulated by expectation. A similar question could have been investigated with other experimental approaches.

M1 excitability can also be measured via M/EEG. The advantage of these methods is a temporally precise signal. Furthermore, data recording is continuous through the reaction time, unlike with an MEP, and thus data about temporal dynamics is richer for the same number of trials. These advantages make this a promising method for this kind of experiment.

However, unlike with an MEP, it is not possible to tell from the cortical recordings of M/EEG, which activity is translated to the corticospinal tract. Furthermore, increased activity seen on M/EEG could come from populations of inhibitory neurons, and thus it is difficult to disentangle motor facilitation and suppression.

A similar experiment to the ones in this chapter was performed using MEG by de Lange et al (2013). In contrast to the findings in this chapter, this study did find a modulation of M1 oscillatory activity by prestimulus bias. However, the conflict with the results presented here suggests the oscillatory activity they found may not translate into corticospinal activity.

Another alternative is a TMS paired pulse protocol (Reis et al., 2008). In such protocols, a pair of MEP pulses are delivered, either at the same cortical site (through the same coil) or at a different cortical site (using a second coil). Specific paired pulse paradigms are known to differentially modulate the MEP. For instance, a subthreshold first stimulus (the Test Stimulus, TS) applied to M1 6 – 25 ms before a suprathreshold second stimulus (Conditioning Stimulus, CS) will produce a facilitation (Intracortical Facilitation, ICF) of the MEP (Kujirai et al., 1993). Similarly, when the protocol is modified to have shorter intervals of 1 – 6 ms between the stimuli, the MEP is instead inhibited (Kujirai et al., 1993); this process is called short interval intracortical inhibition (SICI).

These paired pulse protocols could be combined with the behavioural paradigm described in this chapter to test whether different levels of expectation modulate SICI and ICF. Relatedly, increased temporal expectation given by a warning stimulus has been shown to decrease SICI and increase ICF (Tandonnet et al., 2010). However, the literature the experiments in this chapter was based on does not suggest an a priori hypothesis about how SICI and ICF are modulated by expectation.

A more hypothesis-driven alternative would be to use TMS to test whether expectation is acting at PMd rather than M1. There are two methods for testing PMd-M1 connections using TMS.

Firstly, repetitive TMS can be applied to PMd by a protocol known to either facilitate or suppress PMd activity. Single pulse TMS to M1 can then be used to test whether the MEP is in turn facilitated/suppressed by the PMd modulation (Rizzo et al., 2004). For instance, repetitive TMS to PMd (used to inactivate PMd) decreases inhibition during motor preparation, suggesting PMd is responsible for that inhibition (Duque et al., 2012).

Alternatively, paired pulse TMS can be used with a coil over PMd. For instance, a subthreshold CS over ipsilateral PMd with a CS-TS interval of 6 ms reduces the MEP evoked from M1 (Civardi et al., 2001). A suprathreshold CS will cause facilitation. Contralateral PMd-M1 connections can also be tested (Mochizuki et al., 2004).

To test whether expectation modifies PMd-M1 connections, an ipsilateral paired pulse protocol would be optimal. The experimental question would be whether MEP suppression or facilitation by PMd is altered by the level of prior expectation.

2.4.10. Conclusion

In conclusion, three experiments found no effect of prior expectation on corticospinal excitability. This seems unlikely to have been because MEP suppression masked a result. The null result could be: (1) because prior expectation is, in fact, not represented in motor areas before stimulus appearance; or (2) because MEP variability renders TMS an unsuitable technique for studying subtle activity changes in a perceptual decision, even with a large number of subjects.

2.5. Appendix I: ANOVA to compare MEP data from Experiments 2 and 3

Within Subjects Effects

	Sphericity Correction	Sum of Squares	df	Mean Square	F	p
Hand	None	0.018	1.000	0.018	0.086	0.773
Hand * Dataset	None	0.113	1.000	0.113	0.551	0.467
Residual	None	3.701	18.000	0.206		
Prob	None	0.560	4.000	0.140	0.606	0.659
Prob * Dataset	None	0.770	4.000	0.192	0.834	0.508
Residual	None	16.622	72.000	0.231		
Time	None	0.738	2.000	0.369	1.003	0.377
Time * Dataset	None	0.150	2.000	0.075	0.203	0.817
Residual	None	13.234	36.000	0.368		
Hand * Prob	None	0.357	4.000	0.089	0.624	0.647
Hand * Prob * Dataset	None	0.878	4.000	0.219	1.536	0.201
Residual	None	10.287	72.000	0.143		
Hand * Time	None	0.776	2.000	0.388	1.718	0.194
Hand * Time * Dataset	None	0.852	2.000	0.426	1.886	0.166
Residual	None	8.133	36.000	0.226		
Prob * Time	None	1.973	8.000	0.247	1.317	0.239
Prob * Time * Dataset	None	1.680	8.000	0.210	1.121	0.352
Residual	None	26.967	144.000	0.187		
Hand * Prob * Time	None	1.154	8.000	0.144	0.756	0.642
Hand * Prob * Time * Dataset	None	1.058	8.000	0.132	0.693	0.697
Residual	None	27.476	144.000	0.191		

Between Subjects Effects

	Sum of Squares	df	Mean Square	F	p
Dataset	0.028	1	0.028	0.237	0.632
Residual	2.130	18	0.118		

Table 2.1 MEP data does not differ between Experiments 2 and 3. These tables report a four-way repeated measures ANOVA on the MEP data from Experiments 2 and 3, with hand chosen (L/R), probability condition (10:90/25:75/50:50/75:35/90:10), and timepoint (0%/25%/45%) as within subjects effects, and dataset (Experiment 2/Experiment 3) as between subjects effects.

2.6. Appendix II: ANOVA on stimulus-locked MEP data

Within Subjects Effects

	Sphericity Correction	Sum of Squares	df	Mean Square	F	p	η^2_p
Time	None	1.129	2.000	0.564	1.661	0.205	0.089
	Greenhouse-Geisser	1.129	1.724	0.655	1.661	0.209	0.089
Residual	None	11.552	34.000	0.340			
	Greenhouse-Geisser	11.552	29.308	0.394			
Hand	None	1.59e -5	1.000	1.59e -5	6.4e-5	0.994	0.000
	Greenhouse-Geisser	1.59e -5	1.000	1.59e -5	6.4e -5	0.994	0.000
Residual	None	4.222	17.000	0.248			
	Greenhouse-Geisser	4.222	17.000	0.248			
Prob	None	1.789	4.000	0.447	1.893	0.122	0.100
	Greenhouse-Geisser	1.789	3.072	0.582	1.893	0.141	0.100
Residual	None	16.067	68.000	0.236			
	Greenhouse-Geisser	16.067	52.221	0.308			
Time * Hand	None	0.309	2.000	0.155	0.693	0.507	0.039
	Greenhouse-Geisser	0.309	1.772	0.175	0.693	0.491	0.039
Residual	None	7.582	34.000	0.223			
	Greenhouse-Geisser	7.582	30.119	0.252			
Time * Prob	None	1.408	^a 8.000	^a 0.176	0.827	^a 0.580	0.046
	Greenhouse-Geisser	1.408	^a 4.349	^a 0.324	0.827	^a 0.520	0.046
Residual	None	28.923	136.000	0.213			
	Greenhouse-Geisser	28.923	73.927	0.391			
Hand * Prob	None	0.696	4.000	0.174	1.078	0.374	0.060
	Greenhouse-Geisser	0.696	2.804	0.248	1.078	0.364	0.060
Residual	None	10.972	68.000	0.161			
	Greenhouse-Geisser	10.972	47.668	0.230			
Time * Hand * Prob	None	0.753	8.000	0.094	0.481	0.868	0.028
	Greenhouse-Geisser	0.753	4.491	0.168	0.481	0.770	0.028
Residual	None	26.616	136.000	0.196			
	Greenhouse-Geisser	26.616	76.351	0.349			

^a Mauchly's test of sphericity indicates that the assumption of sphericity is violated ($p < .05$).

Table 2.2 No significant main effects in stimulus-locked MEP data. This table reports the results of a three-way repeated measures ANOVA on the MEP data from Experiments 2 and 3 (combined). The data is 'stimulus-locked' i.e. categorised by timepoint of stimulation. ANOVA factors were timepoint (stimulus onset/25% of reaction time/45% of reaction time) x probability condition (10:90/25:75/50:50/75:35/90:10) x hand chosen for response (left/right).

Chapter Three: Does prior expectation distribute variability unevenly across motor plans?

3.1. Introduction

There is mounting evidence that multiple alternative motor plans are created and maintained in motor cortex whilst the decision about which movement to make is ongoing. Both neurophysiological (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011) and behavioural (Chapman et al., 2010; Stewart et al., 2014; Gallivan et al., 2015, 2016b, 2017) studies suggest that offering experimental subjects two potential targets leads to two motor plans being maintained, and these can interact, leading to a movement trajectory which is initially intermediate between the two potential task goals. Experiments which excluded simple spatial relationships between targets (Gallivan et al., 2015, 2017) have suggested that the multiple plans are true motor representations and not dependent on sensory averaging. The ‘affordance competition hypothesis’ (Cisek, 2007) is discussed in more detail in Chapter One.

What is the effect of maintaining multiple motor plans on the quality of the executed movements? When human subjects are required to plan two pointing movements rather than one, movement variability increases (Wijdenes et al., 2016). This effect is also present when subjects must prepare motor plans with both hands vs one hand to a single target, showing it is not dependent on competition for resources at the level of visuospatial representation.

The finding that more motor plans means more variability is consistent with a limited resource model, in which there is a finite neural resource that must represent various options; in this case, the various possible targets of the reach. Divisive normalisation provides a potential biological substrate for this: it limits the cumulative level of neuronal spiking and

thus makes it a finite resource. Many neuronal populations have been observed to exhibit divisive normalisation, where neuronal inputs are scaled by the overall activity of the population such that the mean level of spiking in the population is always similar (Olsen et al., 2010; Carandini and Heeger, 2012). Adaptation by photoreceptors to light of different intensities is an example of normalisation (Normann and Perlman, 1979; Carandini and Heeger, 2012).

In a normalised neuronal population, maintaining representations of two motor plans rather than one would lead to a reduced signal-to-noise ratio for each plan. Limited resource models have been successful in modelling errors made when subjects need to remember multiple items in visual working memory tasks (Bays, 2014), and a similar mechanism has been proposed to underlie the motor variability finding (Wijdenes et al., 2016).

The concept of a limited resource in motor planning is attractive because it imposes a biological disadvantage to implementing multiple motor plans, and thus proposes a mechanism by which perceptual or decision uncertainty could translate to higher movement variability. Additionally, it places a constraint on the number of motor plans, which seems necessary to counterbalance the optimisation that multiple motor planning brings.

For the idea of a limited resource determining motor variability to be correct, noisier representations during motor planning must translate into variability during the movement. The primary source of motor variability has been a matter for debate. The pre-eminent theory of motor control relies on noise generated at the periphery during motor execution being the main source of variability during movement (Harris and Wolpert, 1998; Todorov and Jordan, 2002). In this theory, noise scales with the size of the signal, such that faster movements are more variable, and the nervous system is optimised to limit the impact of such

noise on movement execution. However, elsewhere, the importance of central sources of neuronal noise has been emphasised, whether sensory, motor, or computational (Churchland et al., 2006; Beck et al., 2012; Chaisanguanthum et al., 2014). Churchland and colleagues (2006) have argued that, in macaques, variable spiking during motor planning explains more than half of the variability in executed reaches. The experiment by Wijdenes and colleagues (2016) found a measurable difference in variability due to planning two movements rather than one, suggesting that it is reasonable to study the effect of motor planning on motor variability in behavioural tasks.

3.1.1. A hypothesis linking prior expectation and motor variability

This chapter examines the situation in which a subject must plan two possible movements, but expects to make one with a much stronger likelihood than the other. This is analogous to many real-world situations: when reaching out to take a cup of hot tea from a friend, we must plan, with a strong expectation of execution, a movement to grasp the cup, but also, with a much lower likelihood, a movement to quickly pull our hand out of the way, in case they spill the tea.

It is known that a stronger expectation of making a movement increases neuronal spiking for cells interested in that movements prior to evidence accumulation in lateral intraparietal cortex (Rao et al., 2012) and superior colliculus (Basso and Wurtz, 1997, 1998). In this chapter, I hypothesise that uneven spiking rates translate into an effect on motor variability.

Motivated by the limited resource theories of Bays (Bays, 2014, 2015; Wijdenes et al., 2016), I propose that the effect of prior expectation in motor cortex is to allocate the limited resource to motor plans unevenly, so that the signal-to-noise ratio of the expected motor plan is boosted at

the expense of the other. Assuming that the impact of limited resources in motor planning can be measured in motor variability (Wijdenes et al., 2016), this predicts that the variability of the expected movement is reduced whilst the variability of the unexpected movement is increased. As the variability increase derives from the motor planning stage, I expect it to predominate in the first part of the movement, and thus will analyse motor variability throughout the movement. (Other studies have typically focussed only on endpoint variability (e.g. Pekny et al., 2015)).

3.1.2. Experimental design

In this chapter, I test the hypothesis that an increased prior expectation to move reduces motor variability. I describe a behavioural experiment in which human subjects make two-dimensional reaching movements and a cue before movement gives subjects a trial-by-trial expectation about which movement they are likely to have to make.

Subjects made right-handed reaching movements to either a target oriented at either 45° to the left or to the right of the starting position. On each trial, before seeing a cue which instructed subjects which direction to move in, subjects saw a probabilistic cue which informed them one of three possibilities: the left target was more likely to be cued (with an 80% probability); the right target was more likely to be cued (80% probability), or both targets were equally likely to be cued. These cues were veridical. After this, either the left or the right target was cued, and subjects had one second in which to respond with a speeded reaching movement, aiming to win points by hitting the cued target.

3.2. Methods

The aim of the experiment was to measure movement variability under various levels of expectation about making the movement. This was tested with an experiment in which human subjects were required to make fast reaching movements, using an onscreen cursor, to one of two targets, each positioned at 45° to the horizontal either side of the starting position. On each trial, subjects were shown probabilistic information (a ‘prior cue’) which indicated whether the cued target was likely to be the left one, the right one, or either target. After the target was cued, subjects had a brief period to execute a movement to the target and were given feedback on whether their movement had been successful or not.

3.2.1. Participants

Fourteen participants were tested (3 female, mean age = 24, $SD = 5.36$) recruited through a university subject pool. The choice of sample size was based on previous behavioural experiments studying variability (Pekny et al., 2015; Wijdenes et al., 2016). Participants gave written informed consent. The experiment was approved by the research ethics committee of University College London (United Kingdom). Subjects were naïve to the purpose of the experiment.

3.2.2. Robotic apparatus

Subjects rested their right forearm in a plastic support and grasped a manipulandum in their semipronated right hand at approximately chest height (Figure 3.1). Their forehead was supported on a rest. The manipulandum moves with two degrees of freedom. Subjects viewed reach targets and the cursor via a screen of size 40 cm by 64.5 cm projected through a mirror of 30 cm by 36 cm which obscured a direct view of the hand and arm. The display appeared to be in the same plane as the hand. The apparatus was controlled by custom C++ code

(Microsoft). The position of the manipulandum was recorded at 200 Hz and this position was Kalman filtered and drawn to the screen as a circular outline cursor of diameter 0.3 cm. Screen information was updated at 200 Hz and the screen display was refreshed at 60Hz. Objects on the screen were drawn in white on a black background, except where stated.



Figure 3.1 The robotic apparatus used allowed subjects to control a cursor with their hand movements. Subjects grasped a manipulandum in their semi-pronated right hand. Their view of their own arm was obscured by a screen (translucent in the diagram). The screen contained a cursor which reflected the position of the manipulandum. Targets which indicated to the subjects which movement they would have to make were displayed on the screen. Figure modified from one by J Galea.

3.2.3. Trial protocol

The starting box and two targets were displayed on the screen at all times. The starting box was a square outline of length 0.7 cm centered at 0 cm (x-dimension) and -12 cm (y-dimension) relative to the centre of the screen. The two targets were circular outlines of diameter 1 cm positioned 20 cm away from the starting box at 45° angles.

To begin a trial, the participant had to move the cursor inside the starting box. Once the subject was within 0.25 cm of the centre of the starting box, the cursor would disappear and a probability cue would appear (Figure 3.2).

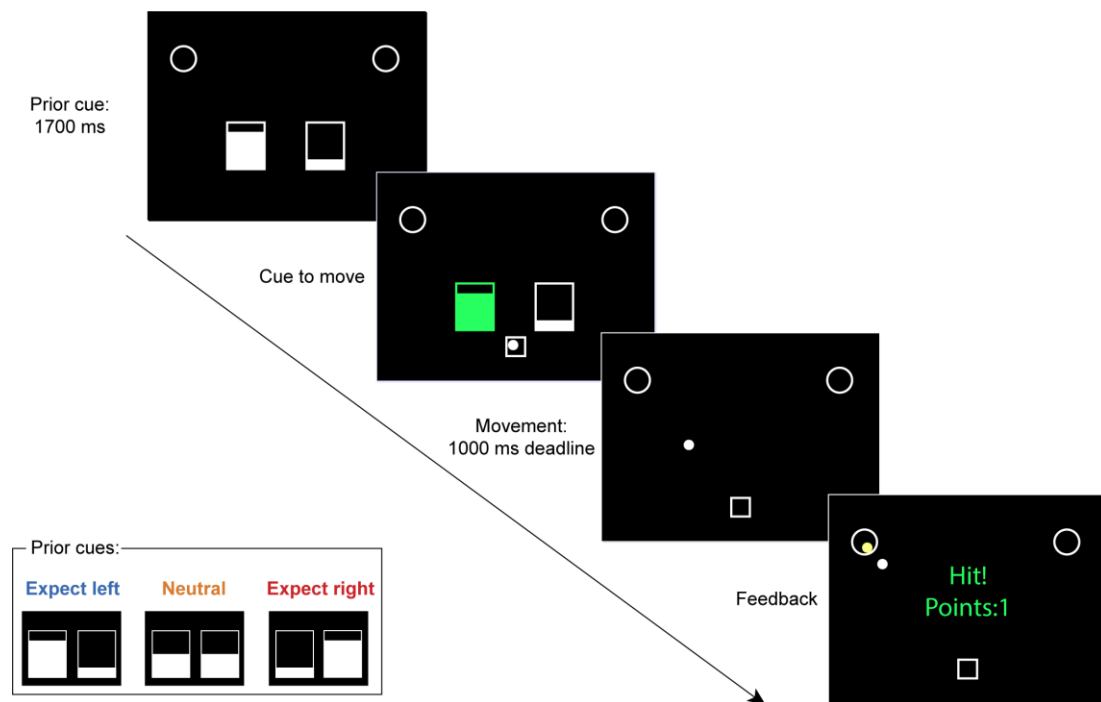


Figure 3.2 Information on screen during a trial of the experiment. Subjects saw a prior cue (1st panel; here the cue is 80:20) followed by a target cue which also indicated the movement should begin (2nd panel). During movement, subjects saw their hand position represented with a cursor (3rd panel). After movement end, subjects were given endpoint feedback (yellow dot) and scored points according to the accuracy of the reach (4th panel). On the trial displayed, the prior cue indicates an 80% probability the reach will be to the left, and the left target is cued. Inset shows the three possible prior cues. Not to scale. Cursor depicted as filled rather than outline for visibility.

On each trial the prior probability cue could be 20/80 ('20% probability the left target will be cued; 80% probability the right target will be cued'), 50/50 or 80/20, and these probabilities were displayed to subjects visually. The prior cue was composed of two rectangle outlines (position: -3 cm or 3cm (x-dimension), -8.5 cm (y-dimension) relative to the centre of the screen; size: 2 cm x 5 cm). Each rectangle was filled with white in proportion to the relative likelihood of the corresponding target being cued (see Figure 3.2).

Once the subject had held the required starting position for 1700 ms, one of the probability cue rectangles turned green and the cursor reappeared. (If the subject left the starting box prematurely, the rectangle failed to turn green, a motorised robotic manipulandum pushed them back towards the starting position, and the timer for holding an acceptable start position was reset.) The probability cue rectangle turning green signalled which target to move to (left target if the left rectangle turned green, right target if the right rectangle turned green). This was also the cue to move.

The rectangles turned green rather than the targets, to encourage subjects to focus their attention on the two (close together) rectangles rather than the two targets (which were far apart from one another), reducing a possible confound in which a subject would be quicker to detect a particular cue if she happened to be attending to the corresponding target.

After this cue to move, subjects had to complete their response within 1000 ms (i.e. a combined reaction time and movement time criterion of 1000 ms). Movement start was operationalised as a movement speed greater than 3.5 cm/s. Once the subjects began moving, the probability cue disappeared. The movement was considered ended if the displacement was greater than or equal to 20 cm (the distance of the targets from the starting box) or the reaction time plus movement time was greater than the deadline. If the subject's speed fell below 3.5 cm/s during the movement, this was considered a stop and if movement did not resume within 40 ms, this was also considered the movement end.

When the movement had ended, feedback about the movement was displayed. Subjects saw a static yellow cursor of same size as the movement cursor at the position where the movement had ended. If the cursor position at the end of movement was within the cued target,

subjects saw 'Hit! Points: 1' displayed in green. If a subject ended the movement inside the non-cued target they saw 'Wrong target! Points: 0'. (In practice this did not happen a single time, for any subject.) If a subject made a movement of sufficient amplitude but did not land inside a target, they saw 'Miss! Points: 0' displayed in red. If a subject failed to make a movement of sufficient amplitude within the deadline, 'Too Slow' was displayed in white.

Once the movement was over, the motorised robotic manipulandum pushed subjects back towards the start position. If subjects' movement amplitude was greater than 23 cm, they were also pushed back towards the starting position, in order to prevent them hitting the back of the apparatus.

3.2.4. Experimental protocol

Each subject undertook one training block and six experimental blocks. A block consisted of 90 trials. There were three levels of the probability cue factor (20/80, 50/50, 80/20). These conditions were interleaved across blocks to avoid block effects. There were 30 trials of each level in a block. The order of trials was randomised within blocks. The schedule of cued targets was determined such that the likelihood levels subjects experienced in a block were equal to the probability cue for that level (e.g. for the 30 trials in a block with a 50/50 cue, 15 were cued for the left target and 15 for the right target).

The nature of the probability cues and scoring system was explained to participants and they were coached through the first approximately 10 trials by the experimenter. Subjects were paid £5 for their participation and an additional £6-10 based on the score they accumulated during the experimental blocks, to increase motivation. (Score being the summed total of all the hit trials.) Subjects were aware that the training block did

not count towards their score and the experimental blocks did. Participants could choose whether to take a break between blocks.

3.2.5. Analysis

Data were analysed in MATLAB, version 2014b (Mathworks, <https://www.mathworks.com/products/matlab.html>) and JASP, version 0.8.1.2 (Jasp Team, <https://jasp-stats.org/>).

3.2.5.1. Processing of position data

During the experiments, the measurement sampling rate was 200 Hz, giving an average number of samples per movement trajectory (i.e. after reaching a movement speed threshold 3.5 cm/s) of 95.4 ($SD = 8.86$). To facilitate comparison across trials, data was interpolated so that each trial comprised a fixed number of samples, irrespective of movement speed on that trial. Each measured trajectory was interpolated to 100 points using spline interpolation based on cubic interpolation of neighbouring values (`interp1` function; Matlab). This effectively normalised all trajectories with respect to time. Interpolated trajectories were used in all analyses.

The interpolated trajectories were used to compute a reach angle over the course of the movement. Past studies involving reaching movements have also used reach angle (Pekny et al., 2015; Wijdenes et al., 2016); it has the advantage of collapsing x - and y -data into a single metric. Prior to angle calculation, rest position was first subtracted; this effectively equalises the starting point across trials. Reach angle was calculated for each interpolated sample by finding the angle between the position at that sample and the rest position, relative to the vertical midline (see Figure 3.3).




Figure 3.3 Reach angle was calculated as the angle between position at sample and rest position, relative to the vertical midline. Leftward reach angles are negative.

3.2.5.2. Defining movement onset

The apparatus generated velocity data in addition to position, which was calculated by a Kalman filtering process during the experiment. These velocities were used to define the start of movement. For the majority of the analyses, movement onset was operationalised in the same way as during the experiment: reaching a 3.5 cm/s speed criterion. Thus reaction time was defined as the time between the green cue to move being shown and a movement reaching the 3.5 cm/s speed criterion. Movement time was defined as the time between the 3.5 cm/s speed criterion being hit and the movement end.

The above analyses included all data after the speed criterion was reached. Some analyses of trajectory and reach angle which attempted to capture earlier data were also conducted. For these, the start of the trajectory was defined as the point at which a zero or negative y velocity became a positive y velocity (but only in runs of samples where the velocity continued increasing until it hit the 3.5 cm/s speed criterion; i.e. earlier samples for the same movement were selected). These start on average 53.2 ms earlier ($SD = 8.98$), representing on average 0.095 cm of movement ($SD = 0.007$ cm). The sample before this start was defined as the rest position, and this rest position was used in every analysis, regardless of whether the speed criterion was used.

3.2.5.3. Defining trial success

Success was defined as an endpoint within the displayed target confines within the time limit.

3.2.5.4. Trial exclusion

When sorting trials by experimental condition (target or prior cue level), trials which were 'too slow' and 'change of mind' trials were excluded. These trials were excluded in order not to artificially inflate measures of variability with trials that might have markedly different trajectories because they were erroneous.

It was important that measures of reach angle variability were not contaminated by trials in which the subject moved towards non-cued target, because more of these trials might be expected in low-expectation conditions. For this reason, trials in which the reach angle indicated that the subject moved towards the incorrect target, however briefly, were excluded. On almost all these trials, the subject switched during the movement to aiming towards the cued target during the movement (mean number of trials in which subjects did not switch = 0.079% of trials, $SD = 0.14\%$). For this reason, they were termed 'change of mind' trials.

Change of mind trials were defined as trials in which the sign of the reach angle (with negative reach angles indicating a movement in the leftward direction and positive reach angles indicating a movement in the rightward direction) was opposite to the vertical half of the screen in which the movement ended (i.e. started with a leftgoing reach angle and finished in the right side of the screen, or vice versa). The reach angle used in this change of mind exclusion was that at the very start of the movement (i.e. not post the speed threshold).

Too slow trials were defined as trials in which the participant had not made a movement of sufficient amplitude with the time limit. These were also included in case they selectively increased variability in some experimental conditions.

3.2.5.5. Analysis of positional correlations

To test the idea that early position determines final position more strongly in leftward movements than rightward movements, an analysis of how position correlation with endpoint was performed. Within each subject, for each trial, across every interpolated position sample, squared deviation from the subject's mean position at that sample (with sign preserved, such that negative deviations remained negative after squaring) was calculated. At each sample, these deviations were regressed against the equivalent deviations in the final sample i.e. the endpoint. This produced, for each subject, at each sample, a slope (beta) of the regression, which could be used to assess how strongly position at that sample determined final position. By definition, the slope at the final sample was one.

3.2.5.6. Analysis of reaction time as a determinant of performance

Two analyses were performed to determine whether reaction time predicted (1) endpoint performance or (2) reach angle variability during movement.

For the endpoint analysis, for each subject, on each trial, endpoint error was calculated as the Pythagorean distance between actual endpoint and endpoint at the centre of the target cued at that trial (a point 20 cm away and oriented at $\pm 45^\circ$ from the centre of the starting box). Across trials, these endpoint errors were regressed against the reaction times on the corresponding trial (using the MATLAB `robustfit` function). For each subject, this produced a slope (β) of this relationship. These slopes were

then t-tested against 0 to look for evidence of a systematic tendency in reaction time-error relationships (distribution of β s significantly different to 0).

For the analysis of reach angle variability, for each subject, trials were sorted according to the trial reaction time, based on that subject's reaction time quartiles. The variability (standard deviation) of the reach angle in trials across the four groups was then calculated.

3.2.5.7. Statistical analysis

Success rates, movement times, reaction times, rest positions, initial reach angle and standard deviation of initial reach angle were analysed by two or three-way repeated measures ANOVAs, as indicated in the text. Where assumptions of sphericity were violated (as determined by Mauchly's test), Greenhouse-Geisser corrections were applied, indicated in the text as non-whole number degrees of freedom. When conducting the two- and three-way ANOVAs on initial reach angle, reach angles for the left target were multiplied by -1 to make them positive and avoid trivially significant interactions.

For the analysis of positional correlation (Figure 3.14), multiple t-tests were used to compare data across left-cue and right-cue movements, and these were corrected for using the False Discovery Rate (FDR) procedure (Benjamini and Hochberg, 1995).

For analyses of continuous data, paired t-tests and one-way repeated measures ANOVAs were conducted using the `spm1d` Matlab toolbox (Pataky, 2012) for statistical parametric mapping, which uses random field theory (Worsley et al., 1996) to make statistical inferences about one-dimensional continua. Statistical tests under random field theory are suitable for continuous data which violate the assumption of independence, making them unsuitable for standard multiple

comparisons corrections. The toolbox returns the number of clusters in the data which exceed an F-statistic or t-statistic threshold required for an alpha value (significance level) of 0.05.

Figure 3.18 and Figure 3.19 show analyses conducted to look at learning over the course of the experiment. For these, the reach angle used was that of the first interpolated sample post-speed threshold.

3.2.5.8. Bootstrapping analysis

In order to exclude the possibility that systematic changes in number of trials in different conditions were driving differences in standard deviation of reach angle (e.g. there are four times as many right-cue trials than left-cue trials in the 20:80 condition), a bootstrapping analysis was conducted. For each subject, at each level of each factor, for each sample, reach angles were sampled with replacement from the full pool of reach angle datapoints, such that the number of datapoints equalled the number of datapoints in the most sparse condition. The standard deviation of this new sample was found. This process was repeated 100 times and the mean of the bootstrapped standard deviation was calculated. Figure labelling

Figures are labelled with the prior expectation conditions as 'Expect left' (which corresponds to the 80:20 left:right expectation), 'Neutral' (50:50 expectation) and 'Expect right' (20:80 expectation).

3.3. Results

Fourteen participants were tested on a simple reaching task. Subjects used a robotic manipulandum to control a cursor and make a reaching movement to either a target oriented either 45° to the left or 45° to the right of the starting position. Prior to each reach, they saw a probabilistic cue ('prior cue') informing them of one of three conditions: that they were likely to have to move to the left (80% probability), that they were equally likely to have to move to either target, or that they were likely to have to move to the right (80% probability). Following the probability cue, the target for the reach was cued ('target cue') and subjects had one second to respond. The hypothesis under test was that expectation about where to reach affected variability in the subsequent reach.

3.3.1. Success rate

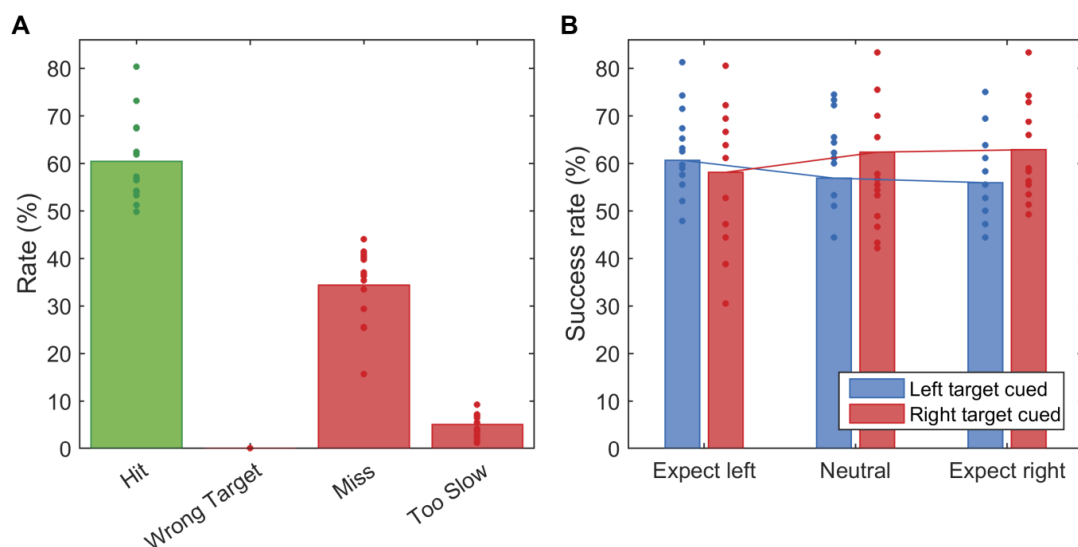


Figure 3.4 Prior expectation influences success rate. A successful trial was one in which the subject hit the cued target within the 1000 ms deadline. (A) Global success rate and breakdown of non-scoring trials. On 'Wrong Target' trials, subjects successfully hit the non-cued target. On 'Miss' trials, subjects completed a movement of sufficient amplitude within the deadline but did not finish inside the target. On 'Too Slow' trials, subjects did not complete a movement of sufficient amplitude within the deadline. (B) Success rate by cued target and prior expectation condition. Filled circles show individual subject data.

Task performance, and whether this differed by prior cue and movement made, were first analysed. Subjects were successful in finishing the

movement within the boundaries of the cued target on 60.4% of trials ($SD = 8.90\%$; Figure 3.4A). They failed to make a reach of sufficient amplitude within the 1000 ms deadline on 5.09% of trials ($SD = 2.68\%$). Subjects never successfully hit the non-cued target ($M = 0.00\%$ of trials; $SD = 0.00\%$). Furthermore, reaches in which subjects finished the movement in the incorrect vertical half of the screen were counted (i.e. on the left side of the screen when the cue was for the right target and vice versa). Most subjects also did not have a single trial in which they made this error ($M = 0.079\%$ of trials, $SD = 0.14\%$). This indicates that subjects had time following the cue to make a reach to the cued target, and miss trials ($M = 34.4\%$ of trials, $SD = 7.85\%$) were primarily caused by inaccurate execution to the correct target rather than an attempt to move to the wrong target.

Figure 3.4B shows success rate analysed by probability cue condition and cued target. A 2×3 ANOVA revealed no main effect of target, $F(1, 13) = 2.49$, $p = .14$, $\eta^2_p = .16$, or prior probability cue, $F(1.84, 24.0) = 0.010$, $p = .99$, $\eta^2_p = .001$, but a significant target \times prior interaction, $F(1.84, 24.0) = 3.92$, $p = .037$, $\eta^2_p = 0.23$. This interaction shows that subjects are more successful on the reach they expect to carry out. To test whether subjects were more successful at right-cue movements in the neutral condition, as suggested by the plot, a post-hoc t -test was conducted. This showed no success rate difference between left-cue and right-cue movements in the neutral condition, $t(13) = -1.98$, $p = .069$, $d = 0.53$.

This analysis suggests participants find right target reaches easier, as shown by increased success for right target cued movements in the neutral expectation condition. This advantage is exacerbated by the expect right condition, but when subjects expected to go left, the trend reversed and they were more successful in left target cued reaches.

3.3.2. Reaction Time and Movement Time

During the experiment subjects had 1000 ms between movement cue and timeout in which to react to the cue and complete the movement. Reaction time and movement time data were analysed by prior cue and movement made, in order to understand whether different expectations induced by the prior cue had conferred speed advantages in reaction and and/or movement.

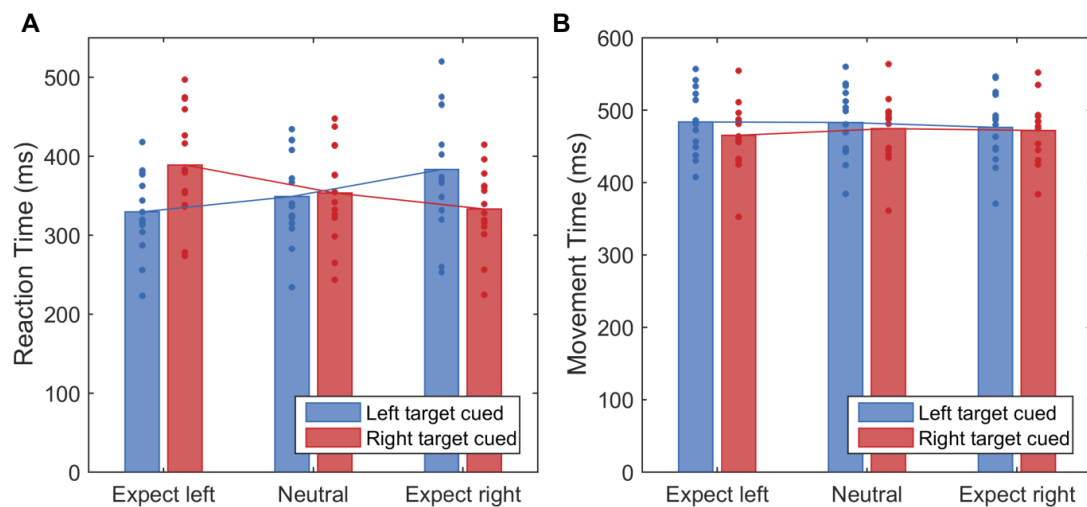


Figure 3.5 Reaction times are modulated by expectation but movement times are not. Mean (A) reaction time and (B) movement time by cued target and prior expectation condition. Filled circles show individual subject data.

The mean reaction time during the experiment was 345 ms ($SD = 57.1$ ms). An expectation to move in a particular direction conferred a speed advantage when the eventual cue was in the expected direction (Figure 3.5A), but a speed penalty when the non-expected target was cued. This was shown by a significant effect of prior, $F(1.43, 18.5) = 10.168$, $p = .002$, $\eta^2_p = .44$, and target x prior interaction, $F(1.14, 14.8) = 51.8$, $p = <.001$, $\eta^2_p = .80$, in a 2x3 ANOVA. There was no main effect of target, $F(1, 13) = 1.14$, $p = .26$, $\eta^2_p = .098$.

The mean movement time was 477 ms ($SD = 44.3$ ms). There was no difference in movement time by target or prior probability cue (Figure

3.5B). A 2 x 3 ANOVA showed no effect of target, $(1, 13) = 3.03, p = .11$, $\eta^2_p = .19$, no effect of prior, $F(1.29, 16.8) = 2.43, p = .13$, $\eta^2_p = .16$, and no interaction between these factors, $F(1.46, 19.0) = 2.47, p = .12$, $\eta^2_p = .16$

3.3.3. Movement trajectories




Figure 3.6 Movement trajectories by expectation. Mean interpolated trajectory across subjects to left and right targets in the 'Expect left' (left panel), 'Neutral' (centre panel) and 'Expect right' (right panel) conditions. Shaded error bars show standard deviation. In this and subsequent analyses, 'Too Slow' and 'Change of Mind' trials were excluded.

Positional data during the reach was recorded at 200 Hz. To facilitate comparison across trials, data was interpolated so that each trial comprised a fixed number of samples, irrespective of movement speed on that trial.

Figure 3.6 shows mean trajectories across all subjects to the two targets in the three conditions. To compute these trajectories, every sample in which the subjects had a positive velocity in the y direction which was part of a full movement was used (i.e. the subject did not become still again before speed threshold reached; see Methods). Trials which were labelled 'too slow' (did not make a movement of sufficient amplitude in one second) or 'change of mind' (starting direction did not match finishing half; see Methods) were excluded in these and subsequent analyses.

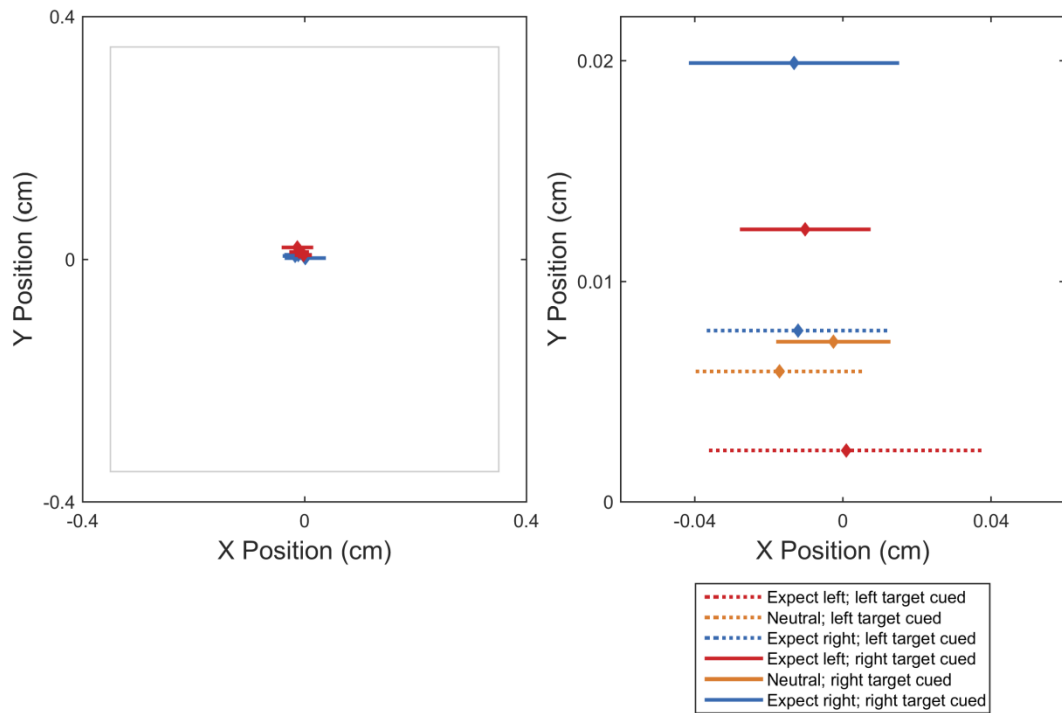


Figure 3.7 Cursor position at last sample before onset of forward movement, by cued target and prior expectation condition. Left panel shows cursor position relative to the 0.7 x 0.7 cm starting box. Starting box shown as grey outline. Starting position for left cue movements shown in blue; starting positions for right cue movements shown in red. Right panel shows an enlargement of the same data, colour-coded according to prior expectation condition.

The last sample before the first used in the trajectory was defined as the rest position. Rest position was analysed before calculating reach angle (see below) to check there was no systematic difference in starting position across conditions. Figure 3.7 shows the rest positions in the conditions, in relation to the starting box in which subjects were required to maintain a position within whilst the prior probability cue was displayed. A 2x3 ANOVA on the x coordinate revealed no main effect of target, $F(1, 13) = 0.029$, $p = .87$, $\eta^2_p = .002$, or of prior probability cue, $F(2, 26) = 1.74$, $p = .20$, $\eta^2_p = .12$, but a significant target x prior interaction, $F(2, 26) = 8.33$, $p = .002$, $\eta^2_p = .39$. Visual inspection of these x -coordinates did not show any systematic ordering of the conditions and suggested the interaction was spurious. A 2x3 ANOVA on the y coordinate revealed no significant effects of target, $F(1, 13) = 4.42$, $p = .055$, $\eta^2_p = .25$, or prior

probability cue, $F(2, 26) = 1.70$, $p = .20$, $\eta^2_p = .12$, and no significant interaction between these factors, $F(1.29, 16.8) = 1.07$, $p = .34$, $\eta^2_p = .076$.

3.3.4. Reach angle

The interpolated trajectories were used to compute a reach angle over the course of the movement (Figure 3.3). Past studies involving reaching movements have also used reach angle (Pekny et al., 2015; Wijdenes et al., 2016); it has the advantage of collapsing x - and y -data into a single metric. Prior to angle calculation, rest position was first subtracted; this effectively equalises the starting point across trials.

The key hypothesis of this chapter was that when expectation to make a particular movement was higher, variability (in reach angle) would be lower. On these analyses (Figure 3.8), and analyses of continuous data throughout this chapter and the next, random field theory statistical tests were used (see Methods). These are suited to continua because random field theory designed to deal with correlated data, whereas conducting repeated t-tests would have violated the assumption of independence.

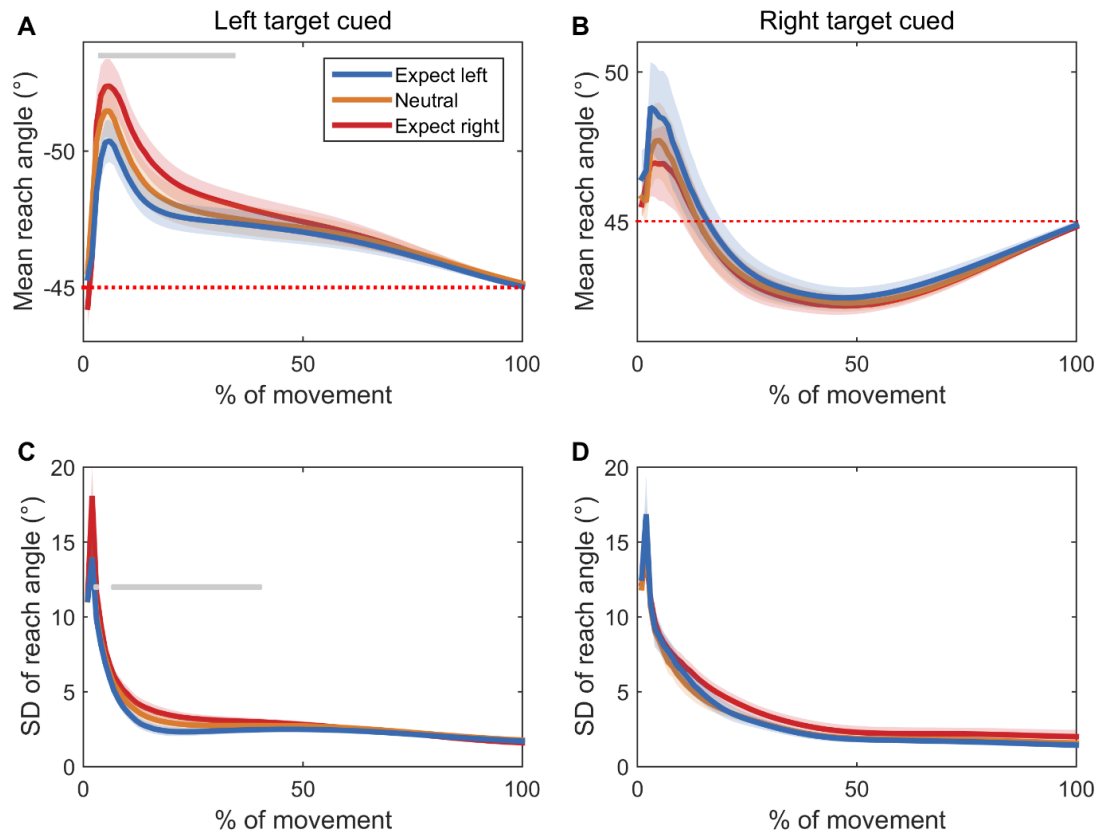


Figure 3.8 Left-cue trials show an early modulation of reach angle and reach angle variability by prior expectation. (A) Mean reach angle across course of movement for interpolated trajectories, in left-cue trials (left panel) and right-cue trials (right panel). (B) Mean standard deviation of reach angle in left-cue trials (left panel) and right-cue trials (right panel). Grey bars indicate clusters of the movement for which there is a significant effect of prior expectation condition. Red dashed line indicates orientation of targets. Shaded error bars show standard error of the mean.

For completeness, mean reach angle was analysed in addition to reach angle variability. Figure 3.8A and B show reach angle, divided by prior probability cue condition. In both left target cued and right target cued trials, subjects have a tendency to end the reach at a reach angle of 45° (the target orientation) relative to their rest position. In left target cued trials, this is preceded by reach angles that are wider than 45° (larger angle between the cursor and the vertical midline), whereas for the right target cued reaches, subjects are initially wider than 45° but then show reach angles that are narrower.

For the left target cue trials only, reach angle differed by probability condition in the first third of the reach (one-way ANOVA using random

field theory showed one suprathreshold cluster, from 2.15% to 33.5% of movement, $F(2, 26)$ threshold = 6.71, $p = <.001$). Subjects followed a trajectory that was more wide of the target when they expected to go right compared to expecting to go left. This modulation of reach angle by prior condition did not apply for right target cued reaches, where a one-way ANOVA using random field theory showed no suprathreshold clusters, $F(2, 26)$ threshold = 6.53.

Figure 3.8C and D show the standard deviation in reach angle, averaged across subjects. Reach angle variance shows a difference by condition in the early part of the movement for the left cued trials but not the right cued trials. One-way ANOVA for left cued trials showed two suprathreshold clusters, from 1.86% to 2.34% of movement and from 5.71% to 39.9% of movement, $F(2, 26)$ threshold = 7.32, p of first cluster = 0.050, p of second cluster = $<.001$. One-way ANOVA for right cued trials showed no suprathreshold clusters, $F(2, 26)$ threshold = 6.74. In summary, the predicted effect was seen in left cued trials but not right cued trials.

This analysis of reach angle variability partially supports the hypothesis, in that a variability difference between prior cue conditions was seen in left cued reaches. For this to be a valid finding, it is important that only variability across reaches in which the subject was aiming in the same direction (i.e. to the same target) are included. Including reaches in which the subject aimed to the opposite target has the potential to artificially increase variability in low expectation conditions, in which an increase in the number of trials to the wrong target might be expected. For this reason, trials in which the subject aimed in the wrong direction (termed 'change of mind' trials because an initial aim in the wrong direction was readjusted over the movement) were excluded in the above analyses. However, because some trials showed small amounts of stop-start

movement prior to the definition of movement start used, it is possible that variability could be increased by mind-changes that had happened in this period, without the trials being excluded. Therefore, as an even more stringent test, later portions of the movement were selected and used, but with an exclusion criterion based on the full movement shown in Figure 3.8. Replicating the pattern in this analysis would give reassurance that early changes of mind are not driving later variability differences.

Raw position data were re-interpolated, now using only samples in which the subject had achieved a threshold speed of 3.5 cm/s. (This was the same way movement was operationalised by the apparatus during the experiment, when it was needed to determine whether subjects had e.g. moved too early.) This meant movement ‘start’ began on average 53.2 ms later ($SD = 8.98$ ms), in which time subjects had moved on average 0.095 cm ($SD = 0.007$ cm) – in other words, they were still within the starting box. The same change of mind exclusion criterion from before was used. This leads to higher change of mind rates than when using a change of mind criterion calculated based on post-speed threshold data only. (Stringent criterion: $M = 4.55\%$ of trials, $SD = 4.29\%$ of trials; other criterion: $M = 1.71\%$ of trials, $SD = 2.58\%$ of trials.) Because trials were excluded based on a very early reach angle and only later reach angles were plotted, this excluded even trials with a direction change very early in the movement, when the cursor was moving within the start box with very low speed.

To aid visualisation of change of mind trials, Figure 3.9 shows histograms of reach angle over the course of the movement for one example subject who had a high change of mind rate. Figure 3.9A shows reach angle in all trials, and shows a gradual narrowing of distribution of reach angles as the movement progresses. Figure 3.9B shows only change of mind trials,

demonstrating how incorrect reach angles are substituted with correct ones over the course of the movement in these trials.

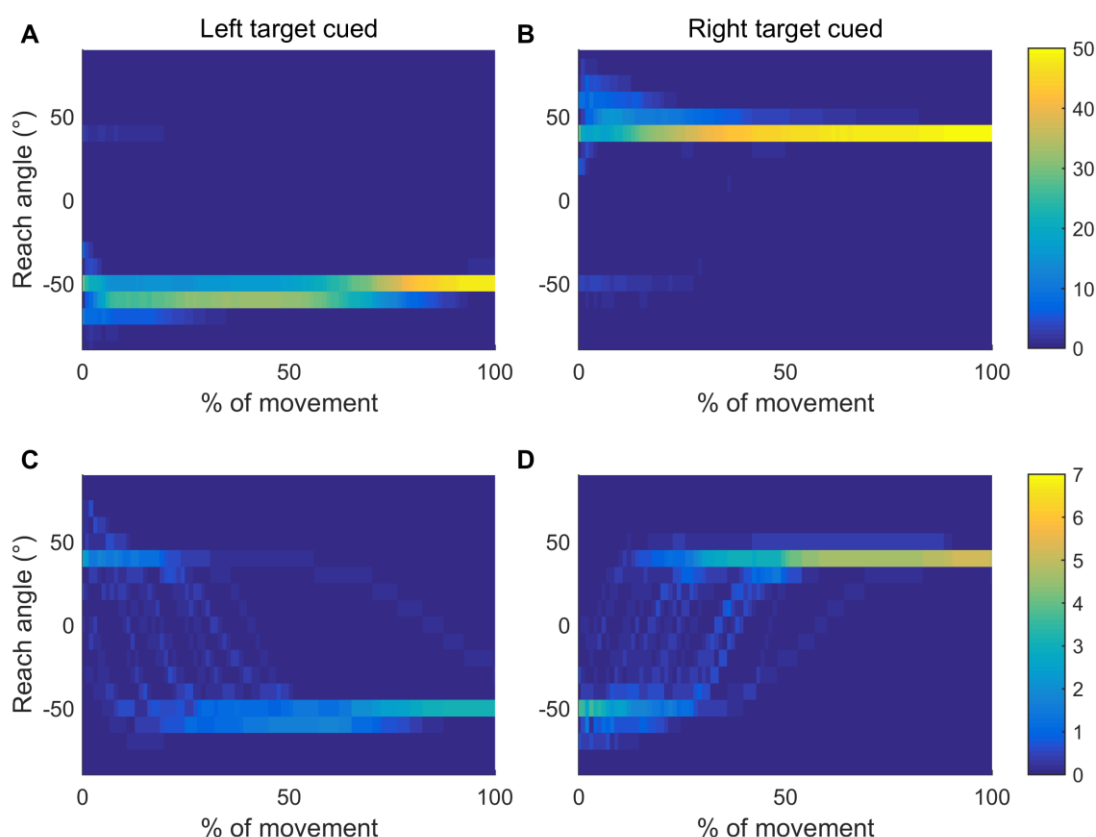


Figure 3.9 Reach angle histograms in one subject. Plots show reach angle histograms through the reach, with reach angles for that point in the trial grouped into bins of width 10°; brighter colours indicate higher number of trials in a bin. (A) All left cue trials; (B) All right cue trials; (C) Change of mind left cue trials; (D) Change of mind right cue trials. These plots show reach angle for the entirety of the reach (i.e. including pre-speed threshold data). The colourbar gives the scale: number of trials expressed as a percentage of all trials across the experiment.

Figure 3.10 shows the reach angle and reach angle variability, with only the later parts of the movement interpolated but using the stringent early change of mind exclusion. This figure displays a reduced range of reach angles and lower reach angle variability than Figure 3.8, due to elimination of the high variability, early portion of the movement. However, the pattern is the same as for Figure 3.9: for left target cued reach angle, there is an early effect of prior probability cue (one-way ANOVA using random field theory showed one suprathreshold cluster, 0 – 29.1% of movement, $F(2, 26)$ threshold = 5.86, $p = .007$). There is also an

early effect of prior probability cue on left target cued reach angle variability (one-way ANOVA showed one suprathreshold cluster, 0 – 39.0% of movement, $F(2, 26)$ threshold = 5.73, $p = .003$). For right target cued trials, there was no effect on reach angle (one-way ANOVA showed no suprathreshold clusters, $F(2, 26)$ threshold = 5.89) or reach angle variability (one-way ANOVA showed no suprathreshold clusters, $F(2, 26)$ threshold = 5.41). Thus the same results were found in the more stringent analysis.

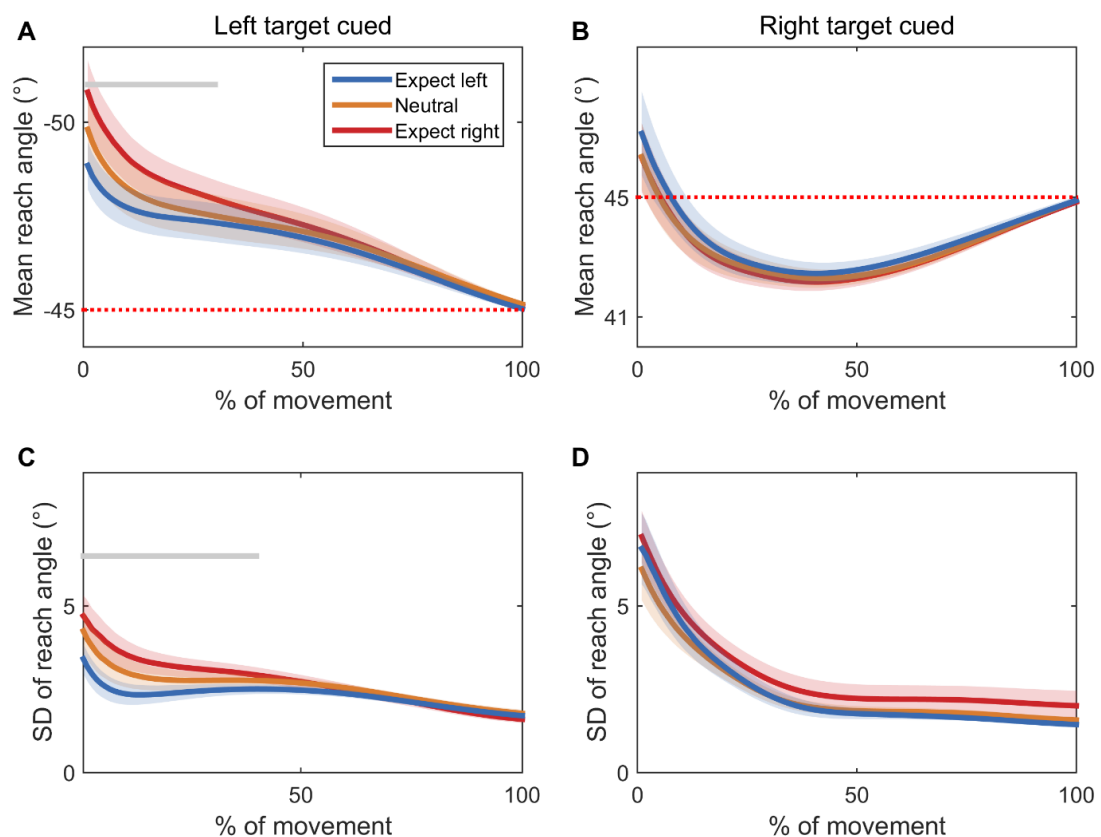


Figure 3.10 Left-cue trials show a modulation of reach angle and reach angle variability by prior expectation when a speed threshold is applied. Details as for Figure 3.9 except that here only samples in which a criterion speed threshold was met have been interpolated to produce the reach angles. Mean reach angle is shown by prior expectation condition for trials in which the left target was cued (A) or the right target was cued (B). Red dashed lines indicate target orientation. SD of the reach angle is shown by prior expectation condition for trials in which the left target was cued (A) or the right target was cued (B). Grey bars indicate a significant difference between conditions.

A constraint of using veridical prior probability cues is that the experimental conditions were unbalanced with regard to trial number. (For instance, in the ‘Expect right’ condition, 80% of trials had a right

target cue, whilst only 20% had a left target cue.) To establish that the lower number of trials was not driving higher variability in reach angle, a bootstrapped version of the reach angle variability analysis was conducted, in which the data set was sampled with replacement in such a way as to balance the number of trials across conditions (Figure 3.11). The same pattern as in the full dataset was found: for the left target cued trials, a one-way ANOVA showed one suprathreshold cluster, from 0% to 39.0% of movement, $F(2, 26)$ threshold = 5.73, $p = .003$. For the right target cued trials, a one-way ANOVA revealed no suprathreshold clusters, $F(2, 26)$ threshold = 5.41.

Similarly, a version of the analysis in which the data were not interpolated but averaged by the original sample number once again revealed the same pattern on statistical testing (not shown).

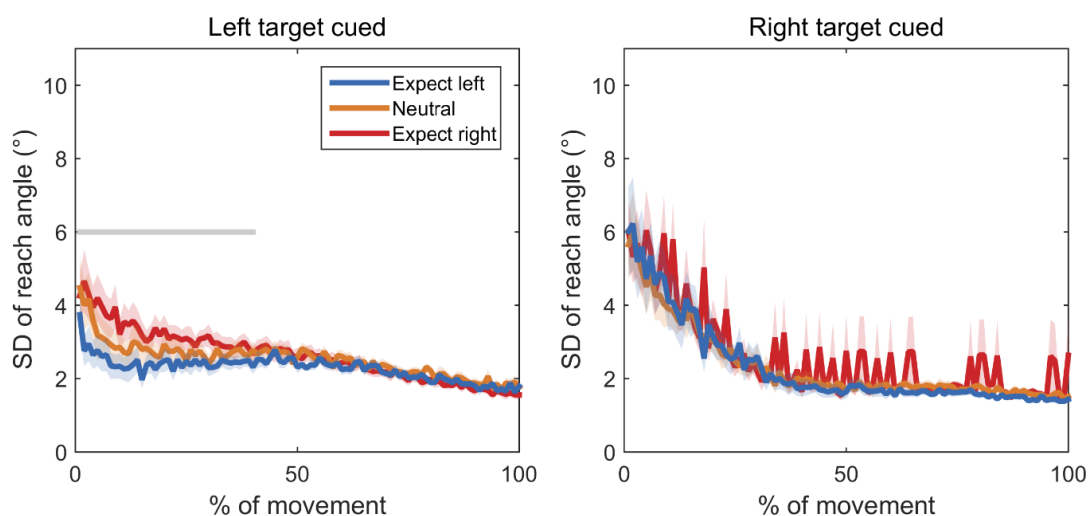


Figure 3.11 The effect of prior probability cue on reach angle variability in left-cue trials is not dependent on sample size. A bootstrapping analysis sampled reach angles with replacement to match sample size across conditions. For each new sample generated, the standard deviation was calculated, and the average of this bootstrapped SD is plotted. Grey bar indicates a difference between conditions.

3.3.5. Directional Bias

Having excluded from the reach angle analyses some trials on which subjects moved in the wrong (non-cued) direction, the rates of these

excluded trials were analysed to understand if they were distributed unevenly across experimental conditions and so build up a picture of participants' strategy. Whilst change of mind trials are shown separately to trials in which subjects started in the wrong direction, in fact these are almost exactly the same trials. Participants almost always attempted to correct a trial in which they had begun in the wrong direction by switching direction ('change of mind'), and so almost never finished the reaches still oriented in the wrong direction ($M = 0.079\%$ of trials, $SD = 0.14\%$). In other words, change of mind trials are caused by an initial start in the wrong direction and a correction of this.

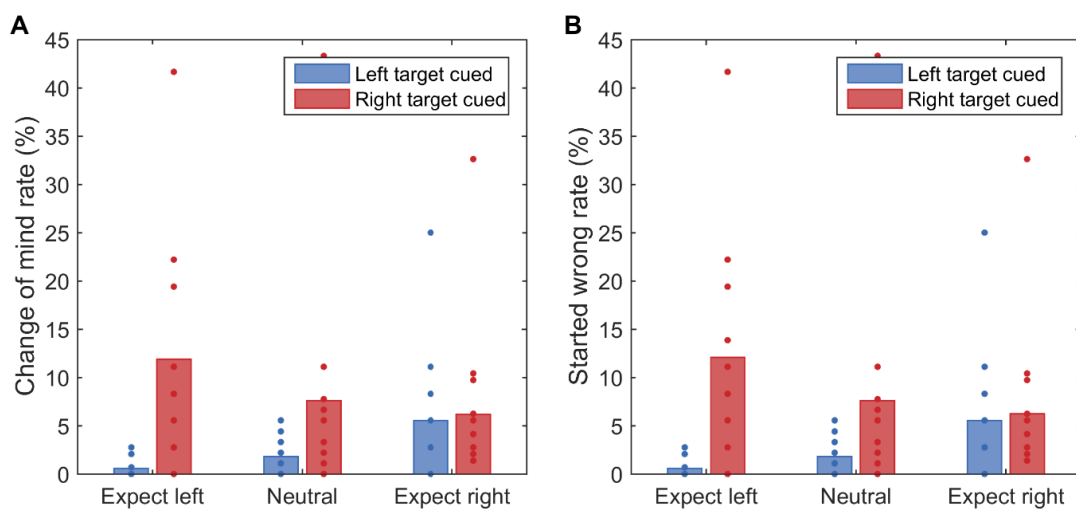


Figure 3.12 Prior probability cue and target affect starting direction. Mean (A) change of mind rate and (B) rate of trials in which the starting direction did not match the cue direction, by cued direction and target. Rates here are as a percentage of trials in that condition, not all trials in the experiment. Filled circles show individual subject means.

An analysis of change of mind trials by condition (Figure 3.12A) shows that subjects were more likely to change their mind on right-cue trials: a 2x3 ANOVA showed a significant effect of target, $F(1, 13) = 6.93, p = .021, \eta^2_p = .35$. There was no effect of prior probability cue, $F(1.36, 17.6) = 0.90, p = .39, \eta^2_p = .065$, but there was a significant target x prior interaction, $F(1.16, 15.0) = 5.22, p = .033, \eta^2_p = .29$. This interaction reflects the

discrepancy between change of mind rates in left and right target trials being strongest when subjects expected to go left.

A similar analysis of the rate at which subjects initially reached in the wrong direction (Figure 3.12B) is visually indistinguishable from the analysis of change of mind trials. As with the change of mind trial analysis, a 2x3 analysis of started wrong trials found an effect of target, $F(1, 13) = 7.16, p = .019, \eta^2_p = .36$, no main effect of prior probability cue, $F(1.36, 17.7) = 1.00, p = .36, \eta^2_p = .072$, and a target x prior interaction, $F(1.16, 15.1) = 5.37, p = .031, \eta^2_p = .29$.

So, the higher rate of change of mind in right cue trials indicates an increased tendency to reach towards the left and later correct this. This effect is strongest when the leftward expectation is highest.

(In Figure 3.12, one subject has markedly higher rates of change of mind in the right-cue trials than the remainder. This was because Subject 14 had a tendency to make rightward reaches with an initial very strong curvature, leading to an initial leftward angle. This idiosyncratic behaviour is not driving the effect – excluding Subject 14 from Figure 3.12 does not change the appearance of the trend. Excluding Subject 14 from the ANOVAs leads to an additional significant main effect of prior probability cue in both cases but does not obliterate any existing significant effects.)

3.3.6. Differences between leftward and rightward movements

The analyses of mean reach angle and reach angle variability showed a difference of prior cue condition for left-cue but not right-cue reaches. This motivated an investigation into the differences between these movements. The analyses of reach angle variability appeared to show that right cued reaches had higher initial variability, despite equivalent success rates.

To test this, variability was examined for all leftward reaches compared to all rightward reaches. Indeed, in the early portion of the movement, right-hand reaches are more variable than left-hand reaches (Figure 3.13); t-test showed one suprathreshold cluster, from 1.00% to 10.4% of movement, $t(13)$ threshold = 2.91, $p = .047$.




Figure 3.13 Reach angle variability was higher for right-target movements than left-target movements early in the movement. To test the hypothesis that the control policy in the early part of the movement was less stringent in rightward movements compared to leftward movements, standard deviation of reach angle in left-cue movements and right-cue movements were calculated. Grey bar shows where right variability was different to left variability in a random field theory t-test. Shaded error bars show standard error of the mean.

On the basis of this variability difference, I hypothesised this was because these movements had non-symmetrical control policies. In particular, I hypothesised that there was less stringent control of variability early on in rightward reaches.

In addition to a gross variability difference, a second test of this hypothesis would be to determine whether, in rightward reaches, early position was a weaker predictor of final position compared to in leftward reaches. Finding a weaker relationship would again be supportive of the hypothesis that early control policy in rightward reaches was less stringent than in leftward reaches.

To test this, an analysis was conducted where, for each sample of each trajectory, the deviation of the position from the mean position at that

sample was calculated. Thus each trial was expressed in terms of deviation from the mean position at every point along the trajectory. In each subject, at each sample, these deviations on the trials were regressed against the endpoint deviation. The slope (beta) of this relationship gave a measure, in each subject, at each sample, of the strength of the relationship between current position and end position. Figure 3.14A and B shows how, as might be expected, this relationship strengthened over the course of the movement, so that deviations closer to the end of the movement were better predictors of final endpoint position than deviations earlier in the movement.

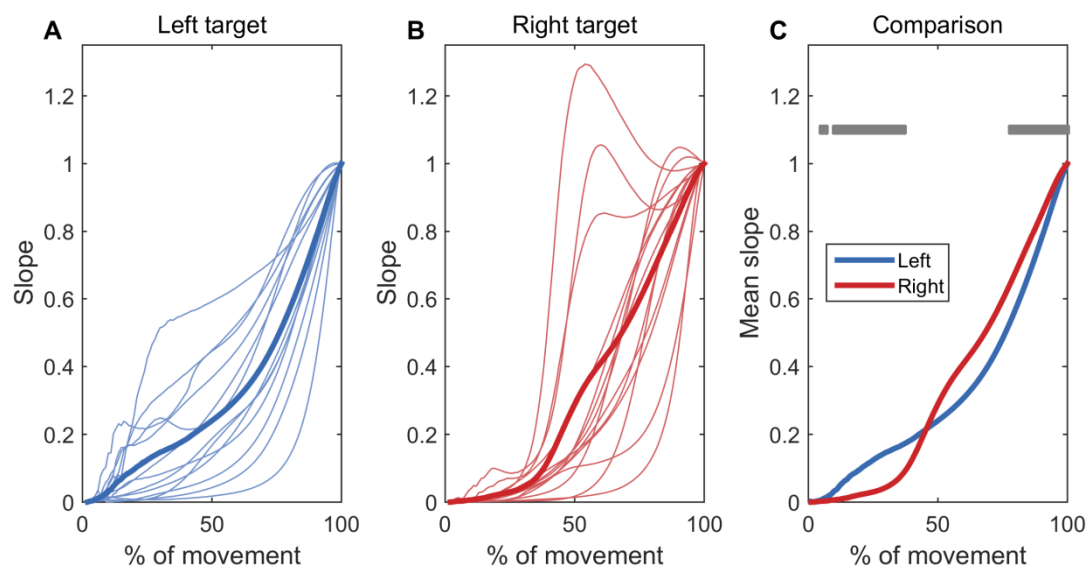


Figure 3.14 Early position is a stronger determinant of endpoint position in left-cue movements than right cue movements. To test the idea that early position determines final position more strongly in leftward movements than rightward movements, the following analysis was performed: at each sample, for each trial, squared deviation (with sign preserved) from the subject's mean position at that sample was calculated. At each sample, these deviations were regressed against the equivalent deviations from the mean endpoint. This produced, for each subject, at each sample, a slope (beta) of the regression. These are shown plotted for left-cue movements in (A) and for right-cue movements in (B). Thin line shows individual subject betas; thick line shows mean. The means are shown again in (C), with grey bars to indicate regions of significant difference (t-tests with False Discovery Rate correction for multiple comparison).

My question was whether, early on in the movement, deviation in rightward movements was a weaker predictor of deviation in endpoint position than in leftward movements. To test this, the distributions of

slopes for right and left were compared to one another. Indeed, right slopes were lower (less strong relationship with endpoint deviation) for right than for left early in the movement. In the last part of the movement, the trend was reversed (Figure 3.14C). This analysis supports a suggestion that position on rightward reaches was less stringently controlled early on in the movement than on leftward reaches, and that this is compensated for by more stringent control later in the movement.

3.3.7. Reaction time as a determinant of variability

Figure 3.5 shows that prior cue affects reaction time, with an unexpected cue being associated with a longer reaction time. It is important to know if performance of the reaching movement differed by reaction time. In particular, did variability during the movement vary with reaction time? If so, it needs to be considered whether reaction time differences could be driving the variability difference seen across conditions (Section 3.3.4).

To look for a relationship between reaction time and task performance, two analyses were performed.

Firstly, for a trial-by-trial measure of whether reaction time predicts performance, endpoint error (distance of endpoint from 'ideal' endpoint, the centre of the target) was regressed against reaction time within each subject (Figure 3.15A). The betas for this relationship were then t-tested against zero, to find out whether there was a systematic tendency for these betas to be positive or negative, which would indicate a relationship between reaction time and endpoint error. In fact, the betas were not different from zero, $t(13) = 0.58$, $p = .57$, $d = 0.15$ (Figure 3.15B).

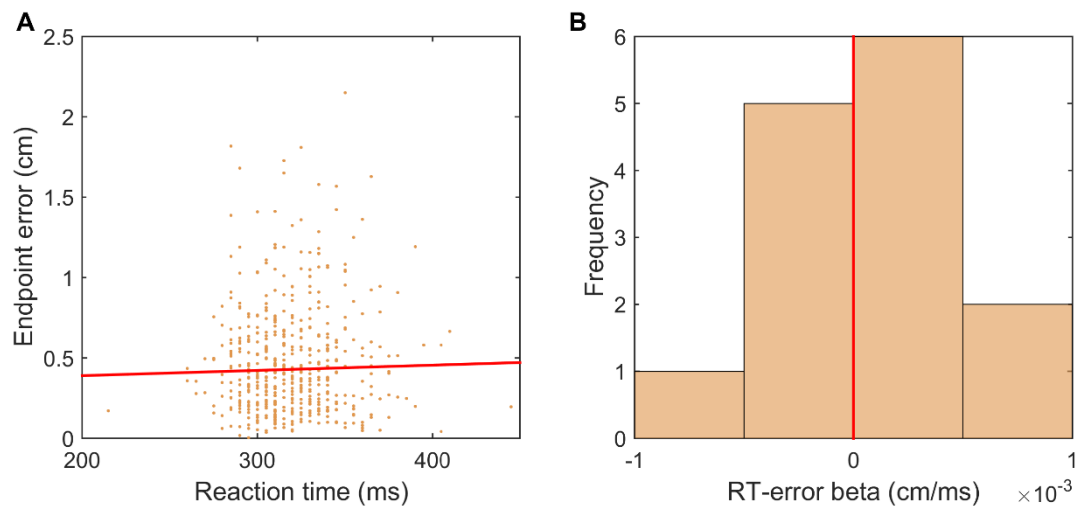


Figure 3.15 Reaction time does not predict endpoint error on a trial-by-trial basis. For each subject, reaction time on each trial was plotted against endpoint error on that trial (Pythagorean distance of endpoint from centre of cued target). This is shown for one example subject in (A). A regression was performed on this data (indicated by red line in (A)). The slopes (betas) of these regressions for all subjects are shown in (B). The distribution of these betas was found not to be significantly different from zero (red line).

(This analysis excludes change of mind trials, as do the main analyses of variability (Section 3.3.4). In case these trials were critical to revealing a relationship between reaction time and endpoint error, the analysis was repeated with these trials included. This again showed no difference of the betas from zero, $t(13) = 0.88$, $p = .40$, $d = 0.23$.)

The goal of the next analysis was to test for an effect of reaction time on reach angle variability across the movement, similar to the analyses of reach angle variability shown above (Section 3.3.4). To achieve this, trials were sorted by their reaction time quartile (assessed within subjects), and variability across trials in each quartile was calculated (Figure 3.16).

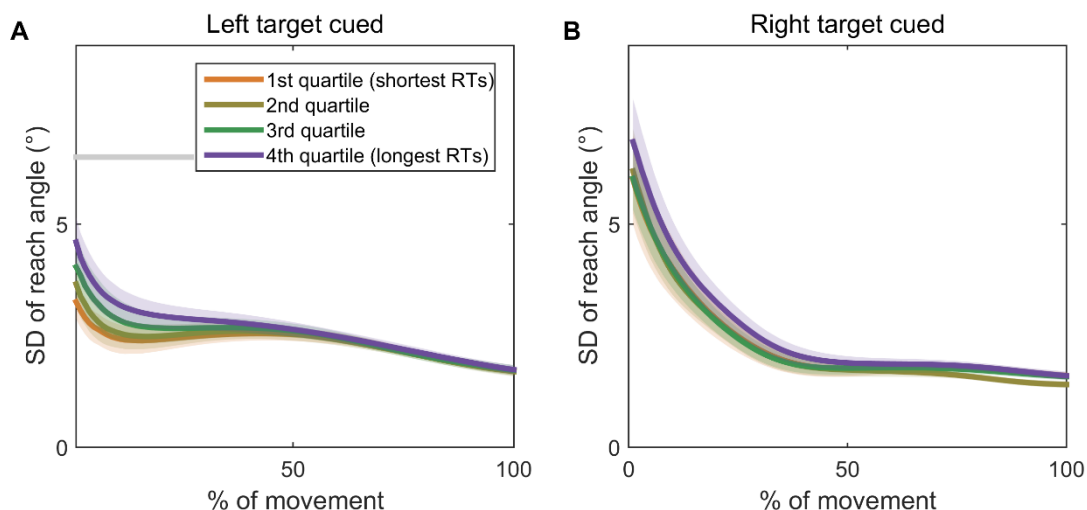


Figure 3.16 Left-cue trials with faster reaction times have lower variability. Trials were sorted by the reaction time quartile and variability in reach angle across the movement was calculated, for left cue movements (A) and right cue movements (B). Grey bars show areas of the movement with a significant effect of reaction time quartile. There was a significant effect of reaction time quartile on reach angle variability in the early part of the movement for left-cue trials. Shaded error bars show standard error of the mean.

As Figure 3.16A shows, for the left target cued movements, there is a difference in reach angle variability by reaction time quartile (a one-way ANOVA using random field theory showed one suprathreshold cluster, from 0.00% to 26.7% of movement, $F(3, 39)$ threshold = 4.51, $p = .011$). This is such that the shortest reaction times (fastest reaction speeds) lead to the lowest variability. There is no such relationship for the right-cue movements (Figure 3.16B; a one-way ANOVA using random field theory showed no suprathreshold clusters, $F(3, 39)$ threshold = 4.70).

However, the analysis above includes data from all prior cue conditions and is thus potentially confounded by any effects of prior cue. In other words, an effect of prior cue on variability might be driving the effect of reaction time on variability. Therefore a more stringent test of the idea that reaction time affects variability is to repeat the analysis, including only trials with a neutral (50:50) cue. All trials will then have the same expectation and so any effect of reaction time on variability will be a direct effect.

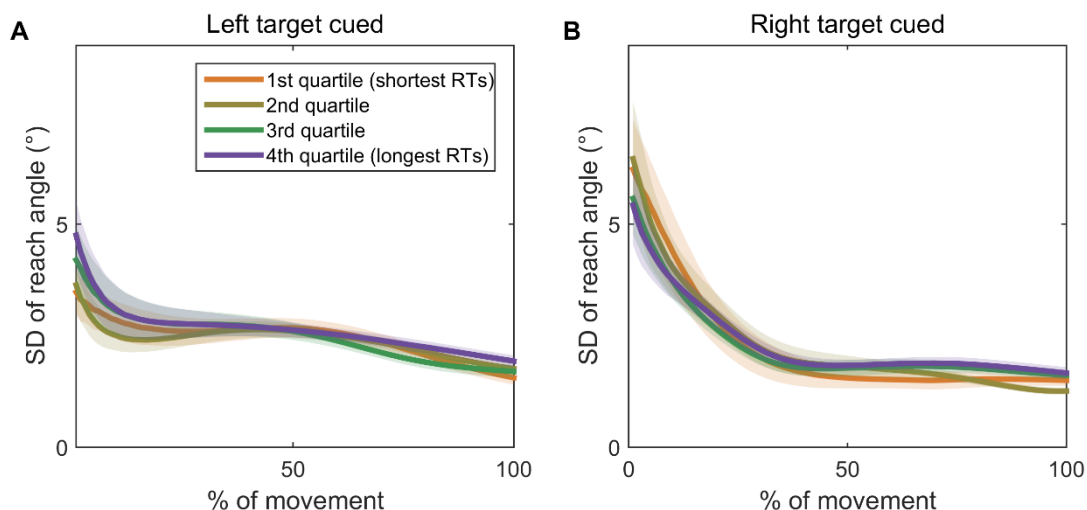


Figure 3.17 There is no effect of reaction time on movement variability when the prior cue seen is controlled for. Figure details as for Figure 3.16, except only trials in the Neutral (50:50) expectation condition have been included. This was done to determine whether the effect of reaction time on movement variability was driven by an underlying difference in prior cue seen.

When this is done (Figure 3.17), no effect of reaction time on variability is found for either left-cue movements (a one-way ANOVA using random field theory showed no suprathreshold clusters, $F(3, 39)$ threshold = 4.50) or right-cue movements (a one-way ANOVA using random field theory showed no suprathreshold clusters, $F(3, 39)$ threshold = 4.64). This indicates the effect in Figure 3.16 was driven by differences in prior cue.

3.3.8. Learning

Experimentally, participants appeared to improve performance over the course of the experiment. Success rates were analysed to understand whether this was indeed occurring and what changes were occurring in reach parameters to drive this learning effect.

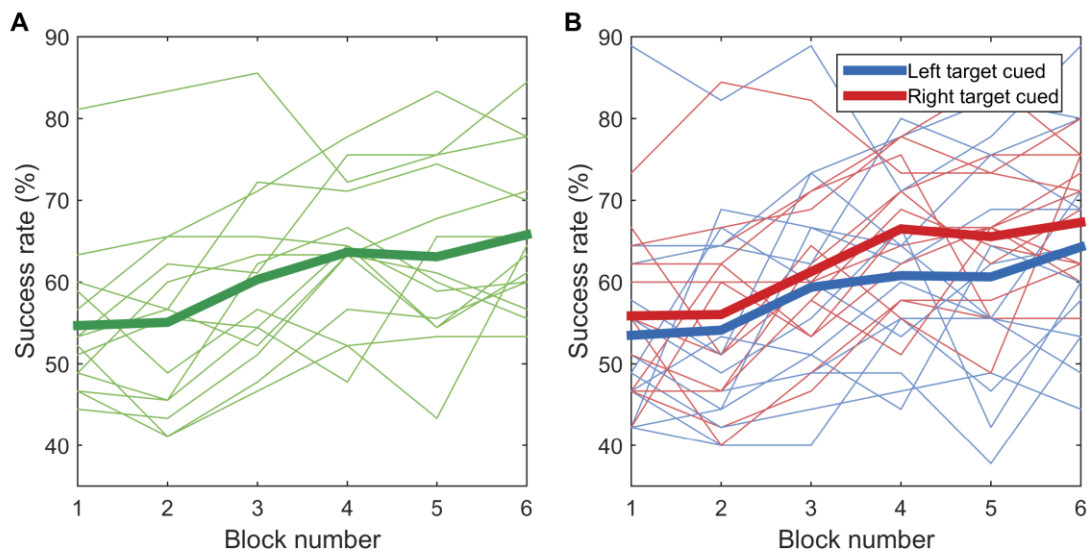


Figure 3.18 Subjects improved over the course of the experiment. Success rate was calculated by block number to test for performance improvement. (A) shows overall success rate by block. (B) shows success rate by block and cued target. Thick lines show mean; thin lines show individual subject data.

Participants were rewarded financially for their cumulative score across the experiment and an analysis of success rate by block shows participants tended to increase their scores across the duration of the experiment (Figure 3.18). A 6x2 ANOVA showed a main effect of block, $F(5, 65) = 9.37, p = <.001, \eta^2_p = .42$, but no effect of right/left target, $F(1, 13) = 4.12, p = .063, \eta^2_p = .24$, or block x target interaction, $F(5, 65) = 0.55, p = .74, \eta^2_p = .041$. In other words, the pattern of improvement did not differ between left-cue and right-cue trials.

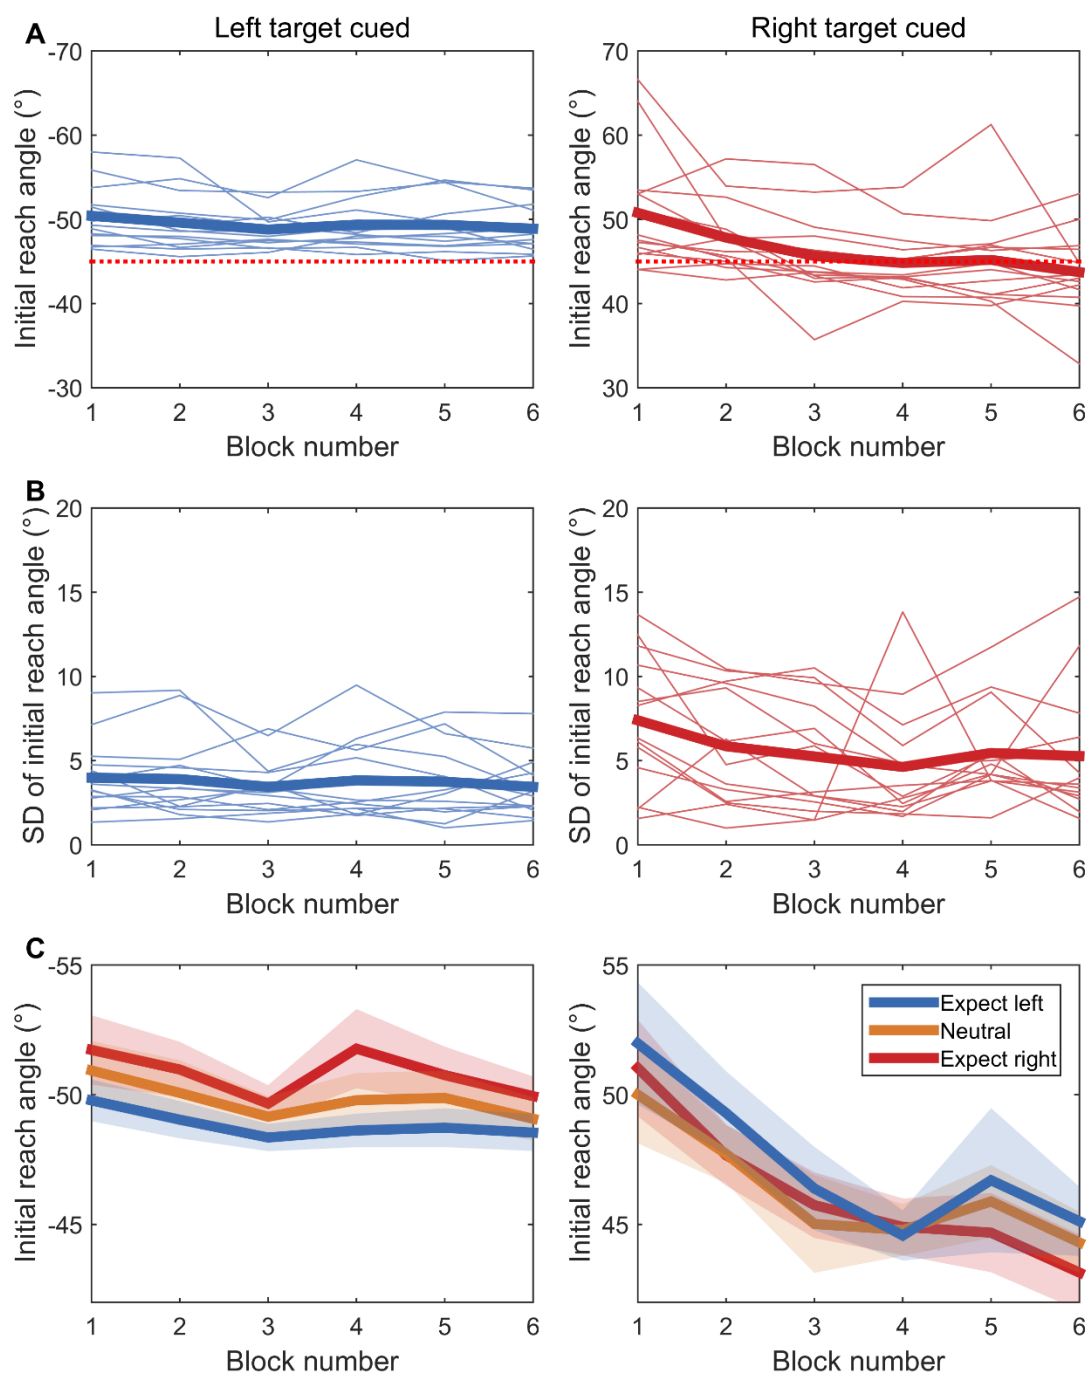


Figure 3.19 Reach angle shows learning across experiment but reach angle variability does not. These analyses were conducted to understand whether movement parameters showed trends over blocks (suggesting learning) to parallel the increase in success rate over blocks. (A) Mean initial (i.e. first post-speed threshold sample) reach angle across blocks for left-cue (left panel) and right-cue (right panel) trials. (B) Mean standard deviation by initial reach angle for left-cue (left panel) and right-cue (right panel) trials. (C) Data in (A) presented by prior expectation condition. In (A) and (B), thick lines show mean and thin lines show individual subject data. In (C), thick lines show mean and shaded error regions indicate standard error of the mean.

In order to investigate what might be driving subjects' improvements in hitting the cued target over the course of the analysis, initial reach angle

and initial reach angle variability were analysed by block (Figure 3.19A & B). Initial reach angle varied with block as shown by a 2 x 6 ANOVA, $F(5, 65) = 12.2$, $p = <.001$, $\eta^2_p = .48$, and target cued, $F(1, 13) = 11.5$, $p = .005$, $\eta^2_p = .47$, and the interaction between these factors was significant, $F(2.58, 33.5) = 4.87$, $p = .009$, $\eta^2_p = .27$, indicating that the pattern of improvement differed with reach direction. The magnitude of the change in reach angle is larger in the right hand than in the left hand. (Note that in this analysis the leftward reach angle was multiplied by -1 to prevent a trivially significant interaction).

In contrast, the variability in reach angle showed no modulation by block, $F(5, 65) = 2.05$, $p = .083$, $\eta^2_p = .014$, although there was a main effect of target, $F(1, 13) = 6.45$, $p = .025$, $\eta^2_p = .33$. There was no interaction between these factors, $F(2.85, 37.0) = 1.31$, $p = .29$, $\eta^2_p = .092$. In short, there was a learning effect in the initial reach angle itself but not in the variability in this reach angle.

To investigate whether the learning effect in initial reach angle differed across prior probability conditions (Figure 3.19C), a 2 x 3 x 6 ANOVA was conducted. As expected from analyses described above, this revealed main effects of block number, $F(5, 65) = 10.1$, $p <.001$, $\eta^2_p = .44$, and of target, $F(1, 13) = 12.0$, $p = 0.004$, $\eta^2_p = .48$, a block x target interaction, $F(2.49, 32.3) = 4.49$, $p = 0.013$, $\eta^2_p = .26$, and a target x prior interaction, $F(2, 26) = 10.6$, $p < 0.001$, $\eta^2_p = 0.45$. However, there was no block x prior interaction, $F(4.26, 55.4) = 0.71$, $p = .60$, $\eta^2_p = .052$, indicating that prior probability level did not modulate reach angle learning. The results of this ANOVA are given in full in Appendix I.

3.4. Discussion

This experiment asked how an expectation of making one movement over another affects the variability of those movements. This was tested in an

experiment comparing reaching movements to two targets under unequal levels of expectation. My key hypothesis was that a lower prior expectation of having to reach towards a particular target would lead to a higher variability in movements towards this target.

3.4.1. Why does variability differ across expectation condition in leftward movements?

A variability difference across conditions was observed in leftward movements, with the lowest variability present when participants had expected to make that movement, and the highest present when participants had not expected to make that movement. This is in keeping with the prediction made: that making a movement, having expected to make the alternative movement, would increase variability.

So, low expectation of making a leftward movement reduces leftward movement quality. This is echoed by the findings in other metrics: reaction times are lengthened, and success rate decreased, in lower expectation conditions. Whilst these findings of a reaction time effect are common in the literature (e.g. Carpenter and Williams, 1995), to my knowledge this is the first study to find an effect of expectation on movement quality.

3.4.2. Why is there no variability difference in rightward movements?

Do the findings fit with a limited resource model of motor planning, as proposed in the Introduction? The findings in left-cue trials, where low expectation increases variability, fit with the idea proposed: low expectation to make a movement reduces the quality of the associated motor plan.

However, the hypothesised variability effect was not present in right-cue trials, in which there was no effect of prior cue. This appears to be a

challenge to a limited resource hypothesis. I proposed that prior expectation distributes limited resources unevenly, such that the quality of one motor plan increases at the cost of the other. If the improved variability in one movement does not come at a cost to the other, how can a limited resource be the cause? To answer this, I propose that if the two motor plans are understood as two differing feedback control policies, the findings can be successfully interpreted under a limited resource model.

According to this line of reasoning, the two movements (leftward and rightward) have distinct, non-symmetrical control policies. We can understand the limited resource as determining the quality of the control policies. In other words, the distribution of the limited resource affects the fidelity of corrections to deviations from the optimal trajectory. If corrections are inaccurate or incomplete, movement variability increases.

This understanding requires characterising motor planning as the planning of two control policies, rather than two movement trajectories. This is in keeping with modern understandings of motor planning (Gallivan et al., 2016b).

Here I propose that worse motor planning leads to a less successful control policy which increases variability. However, this increased variability would only be observed in portions of the movement in which there was active control of motor variability, which I have termed a 'stringent' control policy. If the feedback policy was not stringent in a particular portion of the movement, then worse motor planning would not be expected to affect variability. Hence variability changes arise through an interaction between the limited resource in motor planning and the control policy.

According to this explanation, in this experiment, rightward movements had a control policy which did not control variability stringently in the

early part of the movement and so there was no effect of expectation here.

What is the evidence that leftward and rightward movements had non-symmetrical control policies in this experiment? Firstly, variability in rightward movements was higher than in leftward ones in the early portion of the movement (Figure 3.13). Secondly, early position showed a weaker relationship with final position in rightward movements compared to leftward ones (Figure 3.14). For late position, the relationship was reversed, suggesting that, for rightward movements, control is initially less stringent, and more stringent later in the movement.

In the context of these findings, I interpret the bias to starting with leftward movements (Figure 3.12) as a strategy to begin the execution of these more demanding movements early on.

Why should leftward and rightward movements have had non-symmetrical control policies? The difference arises from the biomechanics of the two movements. Right handed movements towards a rightward target oriented at 45° can be achieved by elbow extension only, whilst right handed movements towards a leftward target involve shoulder and elbow flexion. Involving two joints over one might be expected to increase movement variability, or conversely, increase the amount of control of variability needed for a successful movement. Van Beers et al (2004) studied centre-out reaching movements in all directions, and found marked differences in the endpoint distributions of movements at 45° and 225° (the single-joint movements) from movements in other directions.

3.4.2.1. Feedforward not feedback

Previous studies have shown that control policies are highly flexible. Fast correction occurs within around 150 ms of external perturbation and, although involuntary, these corrections are responsive to sophisticated task demands (Franklin and Wolpert, 2008). More generally, perturbations are only corrected when they interfere with achieving the task goal (Liu and Todorov, 2007; Knill et al., 2011; Nashed et al., 2012; Dimitriou et al., 2013).

Whilst these studies have focussed on the corrections that occur in response to an experimenter-induced perturbation i.e. feedback corrections, I propose corrections that drive the effect in this experiment are feedforward ones. This is because the differences between conditions are present most strongly very early in the movement, which is too soon even for fast feedback corrections.

How does a motor control policy allow for feedforward corrections? The motor system contains a forward model which allows it to predict the sensory consequences of future states (Miall and Wolpert, 1996; Bhushan and Shadmehr, 1999), and learning studies suggest that feedback control processes have access to the current forward model (Wagner and Smith, 2008). One study found that including feedforward strategies in their optimal control model was necessary to describe experimental data sufficiently (Yeo et al., 2016). Here, there are differences between conditions, and I propose these differences rely on a difference in control model. Because these differences are present from the start of the movement, I propose they rely on feedforward error correction, rather than feedback.

3.4.3. Why does the variability difference between conditions diminish over time?

The variability difference in leftward movements caused by varying levels of expectation is greatest at the start of the movement and diminishes through the movement (Figure 3.10). It loses statistical significance at around 40% of the movement. Why does the effect progressively diminish? I propose this is because over the course of the movement, resources are reallocated to the motor plan being executed, and the effect of planning becomes less important.

The study that posited a limited resource theory for visual working memory (Bays and Husain, 2008) found that the distribution of the limited resource was highly dynamic, with rapid reallocation based on the forthcoming saccade, attention, and so on. Here, I propose that the processes underlying distribution of the limited resource in motor planning are similarly dynamic. Indeed, Gallivan et al (2016a) have shown that motor plans are rapidly restructured when a second possible movement target is presented subsequent to the first one.

3.4.4. Why is there a difference in mean reach angle?

As well as a difference of prior cue on reach angle variability in leftward movements, there is also a difference in the reach angle mean for these movements, with narrower movements in the high-expectation condition. I interpret these, once again, as an effect of differing control policies. I propose that the difference in mean reach angle indicates variability accumulates or is controlled more in one direction than the other. It may be that the dynamics of the movement lead to a tendency to reach wider (further from the vertical midline) unless actively corrected. When control of this tendency is weaker, the mean reach angle becomes larger (further from the midline). In support of the idea that the same process

that is driving variability differences is also driving mean differences, these effects are mirrored: they are present over the same period and both only present in leftward movements.

As an alternative explanation, could spatial averaging explain the difference in mean reach angle? It is known that when subjects have to make movements to multiple targets in go-before-you-know paradigms (Chapman et al., 2010), they initially make trajectories that are intermediate between the targets, weighted by the likelihood of having to move in either direction. There was no corollary of this in phenomenon in these results. This is perhaps unsurprising, as subjects had all the information necessary for target selection before they needed to move. There was little tendency for subjects to choose intermediate reach angles, even very early in the reach and on trials where they changed their mind (Figure 3.9). Furthermore, there was a significant effect of target and prior on reaction time but not movement time (Figure 3.5), supporting the idea that a single movement was selected before the movement was initiated. Crucially, the differences in reach angle found actually run counter to the ‘spatial averaging’ phenomenon observed by Chapman et al (2010) – when subjects expected to go left, their reach angle was narrower i.e. the trajectory was closer to the rightward target in coordinate space.

3.4.5. Variability effects are not a result of more changes of mind

An alternative explanation for the variability effects seen in this experiment is that change of mind trials – trials in which the subject initially reaches towards one target but finishes the movement on the other side of the screen – are driving a difference between prior conditions. Under this argument, a lower expectation to move to a particular target would lead to an increased tendency to initially make the wrong movement and switch to the correct target, thus increasing

variability, either directly by increasing the spread of data points, or indirectly, with the hurried switch leading the subject to execute the replacement movement more sloppily. A highly conservative change of mind exclusion was used in which all data points that had a forward velocity that ramped into movement were included to look for changes of mind – in other words, data recorded at very low speeds early in the movement were used, rather than just data points that had reached the speed threshold. Yet it is hard to exclude switching in motor plans at very low velocities at which the movement is still starting and stopping, and it is not possible to exclude decision switching that happens before movement initiation.

However, there are a number of reasons I do not think a change of mind phenomenon could be driving the observed effect. Firstly and most importantly, a change of mind effect would be predicted to lead to mean reach angles which show the opposite effect to the one they do: switching from the right to the left target when right was expected should lead to reaches which are on average spatially nearer to the right target, but these reaches are on average closer to the left target. Secondly, any explanation which relies on cancelling one plan and replanning another would be at odds with the current literature, in which there is neurophysiological (Cisek and Kalaska, 2005) and behavioural (Gallivan et al., 2015) evidence for multiple motor plans maintained in concert.

3.4.6. Variability effects are not driven by reaction time differences across conditions

The results show a reaction time difference between prior cue conditions (Figure 3.5); expected cues lead to faster responses. If movement variability varied systematically with reaction time, it is possible that movement variability changes across conditions could be driven indirectly by reaction time differences rather than by a direct effect on

variability itself. There are a number of reasons to think this was not the case, based on analyses of reaction time and variability presented in Section 3.3.7.

When data from all prior cue conditions was included, there was an effect of reaction time on movement variability in left-cue movements (Figure 3.16). Speed-accuracy tradeoffs are well established in a variety of paradigms in behavioural neuroscience (Heitz, 2014). For this relationship to fit with a speed-accuracy tradeoff, we would expect the fastest responses to lead to the most variable movements. However, the opposite direction of effect is found: faster reactions lead to less variable movements. This is more consistent with an effect of prior cue, with stronger expectation making movements both faster and less variable (i.e. the extra information improves movement speed and quality). Consistent with the idea that the expectation-variability relationship drives the reaction time-variability relationship and not vice versa, there is a reaction time effect only when there is a variability effect (left cue trials, early part of the movement).

A simple way to test whether the prior cue is driving variability differences across reaction time quartiles is to exclude the effect of prior cue by selecting only trials on which a neutral (50:50) cue was seen. When this is done (Figure 3.17), no relationship between reaction time and variability is seen.

3.4.7. Learning does not reduce motor variability in this experiment

Subjects improved their performance over the experiment (Figure 3.18). There is a parallel change in initial reach angle, for both rightward and leftward movements, but not in initial reach angle variability. The absence of change in initial variability suggests that this variability is a constraint that cannot be overcome by learning, or at least not in the 540-

trial timescale of the experiment. The prior expectation cues did not alter the rate of reach angle change. So the effect of expectation on motor variability, which I propose as one of limited resources, is not modifiable by learning, indicating it is a fundamental effect.

3.4.8. Planning versus execution: the source of motor variability

In the interpretation of the results proposed above, it must be planning of a control policy and not merely trajectory planning alone that is under a limited resource constraint. This is needed to explain why rightward movements do not show an effect of prior cue trajectory. A limited resource effect acting purely via motor planning would affect both movements. So I believe that an interplay between the effect during motor planning and execution noise during the movement must be important.

As discussed in Chapter 1, there has been debate about whether noise during motor planning or execution noise is the main determinant of motor variability (Harris and Wolpert, 1998; van Beers et al., 2004; Churchland et al., 2006; Chaisanguanthum et al., 2014). The experiment described in this chapter was designed to examine a putative determinant of motor variability that has its effect at the motor planning stage. However, the explanation for the results that has been proposed in this discussion relies on the control of execution noise. In other words, in this paradigm, noise in motor planning manifests as worse control of noise during motor execution. This is an important distinction that may resolve some of the apparently contradictory experimental data that exist. The planning versus execution debate may not be meaningful because motor planning is concerned with reducing execution noise.

3.4.9. Comparison to previous studies

The affordance competition hypothesis has been extended to suggest that it is not merely multiple movements that are planned in parallel, but multiple task-specific feedback control policies. Gallivan et al (2016b) perturbed subjects in a portion of a movement in which the subjects were still uncertain whether they would have to finish within the bounds of a narrow, stringent target or a wide, lenient target. The gain of the correction subjects made was intermediate between the gain when reaching to just the narrow or just the wide target. In other words, just as there is spatial averaging between competing motor plans when final target location is uncertain (Chapman et al., 2010), there appears to be feedback control averaging when the final accuracy demand of the task is uncertain. This finding, and the flexible nature of feedback control policies, are the foundation for the interpretation of results proposed here.

I have proposed that the control of variability (in other words, the gain on corrections to noise) is weaker if the subject was not expecting to reach to a particular target. In Gallivan et al's paper (2016b), they find that feedback gains are higher with higher uncertainty about a movement. However, I attribute this differences in task: in their task, in which both targets are in the same location, success can always be maximised by increasing feedback gains, whereas this task has an either/or nature, in that subjects must choose the right target to be rewarded.

3.4.10. Alternative metrics for analysing movement trajectories

In the analyses presented here, reach angle relative to starting position was calculated and used as the primary metric for calculating e.g. movement variability. The primary reason for using reach angle was that this experiment was based on a study by Wijdenes et al (2016), which used variability in reach angle at moment of peak speed as the main

dependent variable. In order to make the analysis comparable to that of Wijdenes at all, I also used variability in reach angle, but extended this to be a continuous measure rather than choosing a single point during the movement.

An often-used alternative to reach angle is movement jerk (third derivative of position). It is used because an early model of reach trajectories operated on a principle of minimising jerk across the movement i.e. assumed that the motor system sought to execute the smoothest movement possible (Flash and Hogan, 1985). However, there has been no standard measure for summarising jerk across a movement (e.g. mean jerk, cumulative jerk, mean jerk normalised by peak speed and many other variants have all been used) and this has caused inconsistencies in results (Hogan and Sternad, 2009). One alternative, similar to the continuous reach angle measure I used, would be to deal with continuous metrics of jerk across the movement, rather than calculating a summary. Investigating variability in movement jerk, and whether this varies by prior expectation cue, could be a possible future line of analysis for the data in this chapter.

3.4.11. Conclusion

To conclude, the hypothesised effect of expectation on movement variability was present, but only in leftward reaches. I propose this is because leftward movements have a more stringent control of early variability. I support this with supplementary analyses showing a stronger relationship between early position and end position for leftward reaches.

I believe this is the first study to demonstrate an effect of prior expectation on motor variability. The expectation affected motor execution even though subjects were fully cued as to the correct target

before the movement began. I propose the effect of expectation on variability arises because planning movements with uneven expectations divides the resources of the motor system unevenly.

In sum, expectation determines limited resource allocation, and this interacts with motor control policy to produce the final effect.

3.5. Appendix I: ANOVA on learning effects within initial reach angle

Within Subjects ANOVA

	Sphericity Correction	Sum of Squares	df	Mean Square	F	p	η^2_p
Target	None	1301.20	1.000	1301.201	11.996	0.004	0.480
	Greenhouse-Geisser	1301.20	1.000	1301.201	11.996	0.004	0.480
Residual	None	1410.07	13.000	108.467			
	Greenhouse-Geisser	1410.07	13.000	108.467			
Block	None	1004.87	5.000	200.973	10.078	<.001	0.437
	Greenhouse-Geisser	1004.87	3.203	313.716	10.078	<.001	0.437
Residual	None	1296.21	65.000	19.942			
	Greenhouse-Geisser	1296.21	41.640	31.129			
Prior	None	20.74	2.000	10.372	2.585	0.095	0.166
	Greenhouse-Geisser	20.74	1.538	13.485	2.585	0.111	0.166
Residual	None	104.31	26.000	4.012			
	Greenhouse-Geisser	104.31	19.996	5.216			
Target * Block	None	458.57	^a 5.000	^a 91.714	^a 4.494	^a 0.001	^a 0.257
	Greenhouse-Geisser	458.57	^a 2.486	^a 184.441	^a 4.494	^a 0.013	^a 0.257
Residual	None	1326.59	65.000	20.409			
	Greenhouse-Geisser	1326.59	32.321	41.044			
Target * Prior	None	209.34	2.000	104.670	10.561	<.001	0.448
	Greenhouse-Geisser	209.34	1.573	133.072	10.561	0.001	0.448
Residual	None	257.69	26.000	9.911			
	Greenhouse-Geisser	257.69	20.451	12.601			
Block * Prior	None	43.66	^a 10.000	^a 4.366	^a 0.709	^a 0.715	^a 0.052
	Greenhouse-Geisser	43.66	^a 4.259	^a 10.251	^a 0.709	^a 0.598	^a 0.052
Residual	None	800.45	130.000	6.157			
	Greenhouse-Geisser	800.45	55.361	14.459			
Target * Block * Prior	None	27.48	^a 10.000	^a 2.748	^a 0.488	^a 0.895	^a 0.036
	Greenhouse-Geisser	27.48	^a 4.034	^a 6.813	^a 0.488	^a 0.746	^a 0.036
Residual	None	731.43	130.000	5.626			
	Greenhouse-Geisser	731.43	52.440	13.948			

^a Mauchly's test of sphericity indicates that the assumption of sphericity is violated ($p < .05$).

Table 3.1 Prior expectation does not affect block-by-block changes in reach angle. This table reports a repeated measures ANOVA which was conducted on initial reach angle to understand whether prior expectation condition drives learning-related changes in reach angle. Factors in the three-way ANOVA were target (left-cue/right-cue) x block number (1/2/3/4/5/6) x prior expectation condition (Expect left/Neutral/Expect right).

Chapter Four: Do planning-related variability differences depend on a stringent control policy?

4.1. Introduction

In the previous chapter, I reported that a stronger expectation to make a particular movement reduced variability in reach angle, for a leftward reaching movement but not a rightward reaching movement. To explain this finding in the context of a limited resource model, I hypothesised this was due to differing control policies in leftward versus rightward movements. Specifically, I suggested that there is stronger control of position early in leftward movements than rightward movements. This arises because these movements have non-symmetrical biomechanics: a 45° rightward movement with the right hand involves flexion at a single joint, whilst moving leftward with the right hand involves movement at two joints (van Beers et al., 2004).

I proposed that prior expectation acts at the motor planning stage to affect the quality of the planned control policy. It therefore only has an effect on variability if positional corrections are made under a stringent control policy. The aim of the present chapter was to test this proposal.

4.1.1. Encouraging a non-stringent control policy

In this chapter, I aimed to design an experimental condition in which there was no incentive to strictly control position in the early part of the movement. This was in order to test the idea that this would cause an effect of prior expectation on motor variability to be reduced or obliterated. Effectively, this would mimic the control policy I hypothesised that subjects used in rightward movements in the previous experiment.

The manipulation used to encourage subjects to adopt a non-stringent control policy was to make the target jump during the movement. The experiment presented in this chapter is similar to that in the previous chapter, but introduces a target jump manipulation in half the experimental blocks. In these blocks, as the subject moved the cursor through an occluder a fixed distance from the start position, the target would change position unpredictably. This required subjects to change their trajectory. Importantly, the target jumped relative to the cursor position at the occluder, rather than its own position, so there was no obvious optimal strategy for the initial reach.

In other words, the early movement in the jump blocks is similar to that in the non-jump blocks; but, in the jump blocks, unlike in the non-jump blocks, this part of the movement is unimportant to final success. I used a blocked design (jump vs non-jump blocks) so there was no uncertainty about whether the target would jump or not.

4.1.2. Target jumps as a paradigm to manipulate control policy

Why should the target jump encourage a non-stringent control policy prior to the jump? I hypothesised that, as aiming for a particular position in the early part of the movement was not advantageous, subjects should adopt a policy of correcting less stringently for deviations from the trajectory than in the equivalent portion of the non-jump block movements.

To my knowledge, a target jump manipulation has not been previously used for the reason it was here: to make subjects' movement less stringent and more variable prior to a jump they were aware was coming. However, a large literature on Optimal Feedback Control (Harris and Wolpert, 1998; Todorov and Jordan, 2002) has suggested that subjects

should allow variability to accumulate more when there is no advantage to accuracy.

Multiple experiments find that subjects correct to perturbations most when they interfere with task success and only partially otherwise. This was reviewed in Chapter 1; to recap: subjects correct more to a deviation that will prevent them from hitting the target (Knill et al., 2011; Nashed et al., 2012), and more when they know they are reaching in a context in which perturbations will not be corrected externally (Franklin and Wolpert, 2008). Corrections to perturbations reflect knowledge about obstacles in the environment (Nashed et al., 2012) and the learnt dynamics of a task (Wagner and Smith, 2008; Cluff and Scott, 2013). These studies support the view that feedback policies are highly flexible and reflect ‘high-level’ parameters such as knowledge of task environment. This drove the assumption here that subjects would, on target jump blocks, learn to adopt a lax control policy prior to the jump.

4.1.3. Relevance to Chapter Two

The jump trials used here are analogous to rightward movements in the previous experiment. In the previous chapter it was shown that rightward variability was higher than leftward variability in the early parts of the movement. Furthermore, deviations from the mean position early in the movement predicted endpoint deviations better for leftward than rightward movements. I therefore postulated that the control policy in leftward movements involved a more stringent control of early variability, whilst in rightward movements, more stringent control was applied later in the movements.

By introducing a target jump, I expected that the control policy in both leftward and rightward trials would become more like the control policy in rightward trials in the previous experiment. Thus I expected to

obliterate the effect of prior expectation on reach angle and reach angle variability in the leftward direction of target jump blocks only. This would suggest that the effect of prior was dependent on a control policy that stringently corrected for feedforward deviations in the early part of the movement.

4.1.4. Relevance to target jump literature

Target jump paradigms are common in motor experiments, but have typically been used to investigate feedback processes (e.g. Desmurget et al., 1999; Brenner and Smeets, 2003; Wijdenes et al., 2011). For instance, target jump paradigms have shown that corrective responses are multi-component, with a fast, mostly automatic response at 100 - 150 ms (Prablanc and Martin, 1992) which is likely subcortical (Day and Brown, 2001), and a second response showing a greater degree of voluntary control at 200 – 300 ms (Day and Lyon, 2000). Corrections can be initiated quickly even if two corrections are required in quick succession (Wijdenes et al., 2011). Corrections can be initiated without conscious awareness of the jump (Goodale et al., 1986; Prablanc and Martin, 1992). Disrupting the posterior parietal cortex with TMS interferes with corrections (Desmurget et al., 1999).

Thus much work has gone into describing the corrective responses *after* target jump, which represent replanning processes. These studies have not examined whether knowledge of a target jump affects movement *prior* to the jump. Here, target jumps were used with the aim of inducing in subjects a movement policy where the initial portion of the movement was unimportant to final success. Only the first portion of movement was of interest in the analyses.

The corrections that subjects made in response to the target jump were not studied because they did not relate to the hypothesis posed here.

Furthermore, in order to maximise unpredictability for the subjects, the target jump was to a position that changed on every trial (see Methods), rather than to one of e.g. two fixed positions, as is more usual in previous experiments (e.g. Day and Lyon, 2000). This would have made the trajectories of corrections to jumps difficult to analyse without approximations such as binning trials.

4.2. Methods

4.2.1. Overview

The experiment was designed to (1) introduce a target jump manipulation and see whether this made subjects more variable in the early part of the movement, suggesting a less stringent control policy, (2) discover whether adoption of this less stringent control policy reduced or obliterated an effect of prior expectation on variability.

As in the previous chapter, subjects used a robotic apparatus to control a cursor and make a rapid reaching movement from a starting position to one of two possible targets, each one positioned at 45 degrees to the horizontal either side of the starting position. On each trial, one of the two targets was cued, which indicated which to move towards. Before this ‘target cue’, probabilistic information (‘prior cue’) was displayed to subjects, indicating the relative probability of each target being cued (Figure 4.1). After the target was cued, subjects had 1200 ms to execute a movement to the target and were given feedback on whether their movement had been successful or not.

Unlike in the previous experiment, the movement involved moving the cursor through an occluder at 25% of the distance through the movement. There were two types of blocks. On trials in a non-jump block, nothing happened when cursor passed through the occluder. On jump blocks, the target position changed.

4.2.2. Participants

Twenty-two subjects were tested in total (15 female, mean age = 25, $SD = 4.40$) recruited through a university subject pool. Participants gave written informed consent. The experiment was approved by the research


ethics committee of University College London (United Kingdom). Subjects were naïve to the purpose of the experiment.

4.2.3. Robotic apparatus

The robotic apparatus and screen were as described in the previous chapter (Section 3.2.2).

4.2.4. Trial protocol

Non-jump trials:



Jump trials:

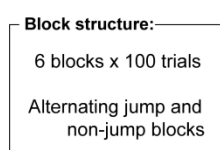




Figure 4.1 On-screen display for one trial. Top series shows a trial in a non-jump block. First, the prior cue is displayed (here: 80:20, i.e. ‘Expect left’). Secondly, one of the rectangles turns green, which is the cue to move (here: left cue). Thirdly, the subject makes a movement, with cursor feedback of hand position. The subject has 1200 ms to respond and to complete the movement. Fourthly, the subject sees endpoint feedback (yellow dot) and written feedback; they gain a point if the trial was successful. Bottom series shows a trial in a jump block. These trials are the same as in the non-jump blocks, except that when the cursor passes through the arc-shaped occluder (3rd panel), the target the subject is reaching towards changes position (‘jumps’).

Screen objects are not displayed to scale; dimensions given in text. For visibility, cursor has been shown as solid rather than an outline.

The starting box and two targets were displayed on the screen at all times. The starting box was a square outline of length 0.7 cm centered at 0 cm (x-dimension) and -12 cm (y-dimension) relative to the centre of the screen. An occluder was also displayed on screen at all times. This occluder was an arc located at 5 cm from the starting box and spanned 180° (Figure 4.1).

The two targets were arc-shaped. The change from circular targets (in the previous experiment) to arc targets (in this experiment) was made to facilitate calculation of target-jumps in reach angle space rather than coordinate space. The targets were each 5° wide, making them roughly equivalent in size to the circular targets of the previous experiment. At the start of the trial they were always oriented at 45° to the left and right of the starting box, at a distance of 20 cm.

As with the previous experiment, to begin a trial, the participant had to move the cursor inside the starting box. Once the subject was within 0.25 cm of the centre of the starting box, the cursor would disappear and a probability cue would appear.

The prior probability cues were visually as described in the previous chapter. In this experiment, two prior levels were used: 20:80 and 80:20. This was a reduction from the three prior levels of the previous chapter, to increase power, as the target jump introduced another condition to the experiment.

Once the subject had held an acceptable starting position for 1700 ms, one of the probability cue rectangles turned green and the cursor reappeared. (If the subject left the starting box prematurely, the rectangle failed to turn green, a motorised robotic manipulandum pushed them back towards the starting position, and the timer for holding an acceptable start position was reset.) The probability cue rectangle turning

green signalled which target to move to (left target if the left rectangle turned green, right target if the right rectangle turned green). This was also the cue to move.

After this cue to move, subjects had to try and hit the cued target and complete their response within 1200 ms (i.e. a combined reaction time and movement time criterion of 1200 ms). This was lengthened from 1000 ms in the previous experiment because the jump trials (see below) increased difficulty significantly.

As in the previous chapter, movement onset was operationalised as a movement speed greater than 3.5 cm/s. Once the subjects began moving, the probability cue disappeared. The movement was considered ended if the displacement was greater than or equal to 20 cm (the distance of the targets from the starting box) or the reaction time plus movement time was greater than the deadline.

When the movement had ended, feedback about the movement was displayed. Subjects saw a static yellow cursor of same size as the movement cursor at the position where the movement had ended. If the cursor position at the end of movement was within the cued target, subjects saw 'Hit! Points: 1' displayed in green. If a subject ended the movement inside the non-cued target they saw 'Wrong target! Points: 0'. If a subject made a movement of sufficient amplitude but did not land on the target, they saw 'Miss! Points: 0' displayed in red. If a subject failed to make a movement of sufficient amplitude within the deadline, 'Too Slow' was displayed in white.

As in the previous experiment, once the movement was over, the motorised robotic manipulandum pushed subjects back towards the start position. If subjects' movement amplitude was greater than 23 cm, they

were also pushed back towards the starting position, to prevent them hitting the back of the apparatus.

4.2.5. Block types: jump and non-jump

There were two block types: jump and non-jump. Trials in non-jump blocks were the same as in the previous experiment: the target stayed in a fixed position. Although the arc shaped occluder was visible, there was no change when the cursor passed through this.

For trials in jump blocks, the target jumped when the cursor passed through the occluder (i.e. at a movement amplitude of 5 cm). The aim of the target jump manipulation was that there would be no optimal position that subjects could maximise performance by aiming towards. The target jumped relative to the participant's reach angle, which was calculated online. The amplitudes of the possible target jumps were: 5°, 7.5°, 10°, 12.5°, or 15° either leftward or rightward. The jump was selected randomly from these each trial and added to the participant's reach angle at the moment of passing the occluder in order to calculate the new target position. Only the cued target moved.

Because it was the cued target that jumped, and it jumped relative to the participant's reach angle, on trials in which the subject was aiming towards the wrong target at the time of reaching the occluder, the correct target would jump to the opposite side so that both targets were displayed on the same side of the screen. This occurred in 0.14% ($SD = 0.24\%$) of trials.

Furthermore, on trials in which subjects made a reach at a wide angle, it was possible for targets to jump below the horizontal. This involved 0.15% ($SD = 0.50\%$) of trials.

The behaviour of the jump manipulation was not changed to prevent these two types of trials (target jumps to opposite side; target jumps below horizontal) as I did not want to introduce any kind of fixed mapping which participants could exploit, aiming at a particular angle in order to produce a predictable post-jump target position. As the totals given above show, the number of trials affected was very small. As movement analysed in this chapter was pre-target jump, whilst these trials only produced idiosyncratic behaviour post-target jump, trials in which the target jumped below the horizontal were not excluded from the analysis. Trials in which the target jumped to the opposite side were excluded for a different reason (see 'Analysis', below).

4.2.6. Experimental design

Each subject undertook six experimental blocks, of 100 trials each. Three blocks were target jump blocks and three blocks were non-jump blocks. Each subject began with either a target jump or a non-jump block, chosen at random, after which the block type alternated.

There were two levels of prior probability cue (20:80, 80:20). Thus there were 50 trials of each level in a block. The order of trials was randomised within blocks. The schedule of cued targets was determined such that the likelihood levels subjects experienced in a block were equal to the probability cue for that level (e.g. for the 50 trials in a block with a 80:20 cue, 40 were cued for the left target and 10 for the right target).

Prior the main blocks, subjects completed two training blocks. They practised a non-jump block for 50 trials, followed by a jump block for 80 trials.

The nature of the probability cues and scoring system was explained to participants and they were coached through the first approximately 10 trials by the experimenter. Subjects were paid £5 for their participation

and an additional £6-10 based on the score they accumulated during the experimental blocks, to increase motivation. (Score being the summed total of all the hit trials.) Subjects were aware that the training block did not count towards their score and the experimental blocks did. Participants could choose whether to take a break between blocks.

4.2.7. Analysis

4.2.7.1. Data processing

Data were analysed in MATLAB, version 2014b (Mathworks, <https://www.mathworks.com/products/matlab.html>) and JASP, version 0.8.1.2 (Jasp Team, <https://jasp-stats.org/>). ‘Movement’ samples were selected between the start of the trial and the last sample prior to reaching the occluder (i.e. movement amplitude 5cm; 25% of total required movement amplitude). Note therefore that the parts of the trajectory beyond the occluder were discarded; the question of interest here was the effect of anticipating a target jump on the early movement variability, rather than the corrections made in response to the jump itself.

Movement was defined in the same way as during the experiment (filtered velocity > 3.5 cm/s). Trials in which there was no movement or the movement did not reach the arc were discarded (0.030% of trials, $SD = 0.066\%$).

Each measured trajectory was interpolated to 100 points using spline interpolation based on cubic interpolation of neighbouring values (interp1 function; Matlab). This effectively normalised all trajectories with respect to time. Interpolated trajectories were used in all analyses.

As in the previous experiment, interpolated trajectories were used to calculate a reach angle relative to the vertical midline. Before angle was calculated, the position one sample prior to the start of the trajectory was

subtracted from each trajectory; this effectively normalised the start position of each trajectory to $(x, y) = (0, 0)$.

4.2.7.2. Analyses of reaction time and movement time

These analyses included all trials i.e. unsuccessful trials were not excluded. As in the previous chapter, reaction time was defined as time between cue to move and movement speed greater than 3.5 cm/s. Movement time was the time between reaching a movement speed greater than 3.5 cm/s and movement end.

4.2.7.3. Analysis of reach angle and reach angle variability

As in the previous experiment, an emphasis was placed on a stringent exclusion of trials in which the subject aimed initially for the opposite target, even momentarily, in order to avoid artificially inflating measures of variability. In the previous experiment, this was achieved by using change of mind by the time of reaching the final target. Since in this experiment, the movement of interest was that up to the occluder, a similar exclusion was achieved instead by discarding trials in which the initial reach angle was not in the same half of the screen (i.e. aiming left vs aiming right) as the cued target. These were 0.65% of trials, $SD = 0.78\%$. The first trial of each block was also excluded, to reduce the impact of the switch from jump to non-jump blocks.

Note that, unlike in the previous experiment, ‘too slow’ trials in which the subject did not manage to hit the target were not excluded. Because the pre-occluder portion of the movement was the movement of interest, there was no longer a clear reason to exclude these trials. Trials in which the subject was too slow to reach the occluder had already been removed at the data processing stage (see above).

After these exclusions, trials were sorted by experimental condition and the mean and standard deviation of this reach angle were calculated within conditions.

4.2.7.4. Analysis of X and Y velocity

The robotic apparatus used a Kalman filtering technique to calculate a smooth estimate of velocity during the trial. These estimates were used to produce the velocity plots in Figure 4.10. The samples were selected and interpolated in exactly the same way as for positional trajectories.

4.2.7.5. Analysis of effect of success or failure on previous trial

The aim of this analysis was to understand how positive or negative feedback altered reach angle when performing the same movement on a subsequent trial (Galea et al., 2013; Pekny et al., 2015). Therefore, for this analysis (Figure 4.11), trials were included only if the same direction was cued as on the previous trial. Excluding trials in which the direction cued was not the same as on the previous trial left 49.5% of trials remaining ($SD = 1.89\%$).

4.2.7.6. Statistical analysis

Success rates, reaction times, and movement times were analysed using paired t-tests and two or three-way repeated measures ANOVAs, as indicated in the text. For the ANOVAs, as all factors had only two levels, assumptions of sphericity were not violated.

For the continuous measures of reach angle, reach angle standard deviation, and velocity, t-tests and ANOVAs were conducted using the spm1d Matlab toolbox (Pataky, 2012) for statistical parametric mapping, which uses random field theory to make statistical inferences about one-dimensional continua (Worsley et al., 1996). Statistical tests under random field theory are suitable for continuous data which violate the

assumption of independence, making them unsuitable for standard multiple comparisons corrections. The toolbox returns the number and positions of clusters which exceed the F-statistic threshold at an alpha value (significance level) of 0.05. All t-tests were paired and two-tailed, and all ANOVAs were repeated measures.

As indicated in the Results section, one Bayesian repeated measures ANOVA (Rouder et al., 2012) was carried out on the leftward reach angle variability data. A single measure of variability was obtained by averaging reach angle standard deviation across the movement epoch. The prior used was that each model, and the null model, were equally likely i.e. a prior of 0.2.

4.3. Results

In Chapter 3, I presented evidence that a higher expectation of having to move to a particular target decreases the variability associated with that movement, and argued that this effect is dependent on a less stringent or less effective online control policy. This experiment aimed to test that assertion.

This experiment compared the effect of an expectation about having to make a movement to a particular location in two types of movement: firstly, a simple reach to a target orientated at 45° , and, secondly, the same reach in which the target jumped when the cursor passed through an occluder 25% of the way through the movement, requiring a change of direction.

Subjects alternated between jump and non-jump blocks. On the jump blocks, the target jumped relative to the participant's reach angle at the time of jump so there was no advantage to aiming accurately in any particular direction. The aim of these target jump trials in this experiment was to reduce the value of strict control of variability early in the movement, as the forced change of direction meant this would no longer contribute to success.

Movement trajectories were analysed in the early part of the movement; up until the occluder was reached. If the assertion that the effect of expectation on variability depends on control policy was correct, it would be expected that, on jump trials (in which control policy did not need to be stringent), there would be no effect of prior expectation on variability. Because an effect of prior expectation was previously only seen on leftward trials, this difference is only expected in leftward trials. Thus the critical test of the hypothesis is an interaction between prior expectation and block type (non-jump vs jump) in leftward trials.

4.3.1. Characteristics of jump trials

Experimental blocks were composed alternately of all non-jump trials, or all target jump trials. Subjects completed six blocks in total; whether they began on a standard or a jump block was randomised for each subject.

As anticipated, the target jump made the task more difficult and subjects were significantly more successful on the non-jump blocks, $t(21) = 8.46$, $p < .001$, $d = 1.80$ (Figure 4.2A).

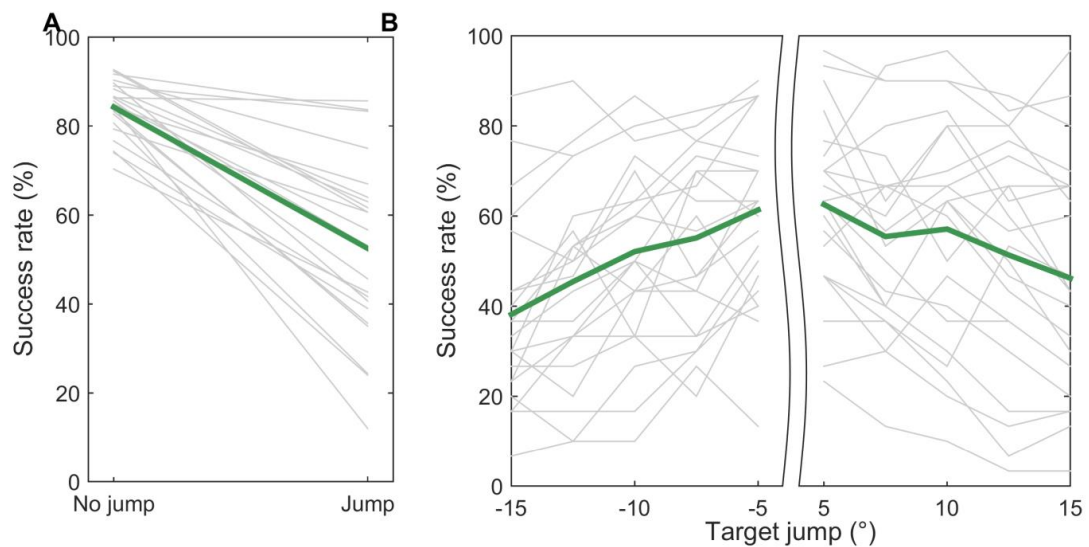


Figure 4.2 Success is lower in target jump blocks. The jump blocks required subject to change direction under time pressure at 25% of the movement, and this is reflected in the success rate. (A) Success rate in non-jump vs jump blocks. (B) For jump blocks only, success rate as a function of angular target jump. Green line shows mean; grey lines show individual subject data.

The target jump was one of a range of fixed sizes between 5 and 15° (target width: 5°) either clockwise or counterclockwise relative to the subject's reach angle at the time of crossing the occluder. Subjects were less successful with increasing size of target jump (Figure 4.2B).

Because the target jumped relative to the subject's own position, it is also possible to plot the target jump relative to the original target position (i.e. the apparent target jump) rather than the preprogrammed target jump. Figure 4.3A shows histograms of these apparent target jumps in all subjects. As the histograms show, a consequence of the target jumping

relative to the subject's position was that in a proportion of trials there was no apparent target jump ($M = 4.49\%$ of target jump trials where apparent jump was less than 1° magnitude, $SD = 1.82\%$). However, as the target could jump with equal likelihood in either direction, subjects had no way to predict when these would occur.

Plotting success relative to apparent jump shows, as with programmed jump, that larger apparent jumps decrease success (Figure 4.3B).

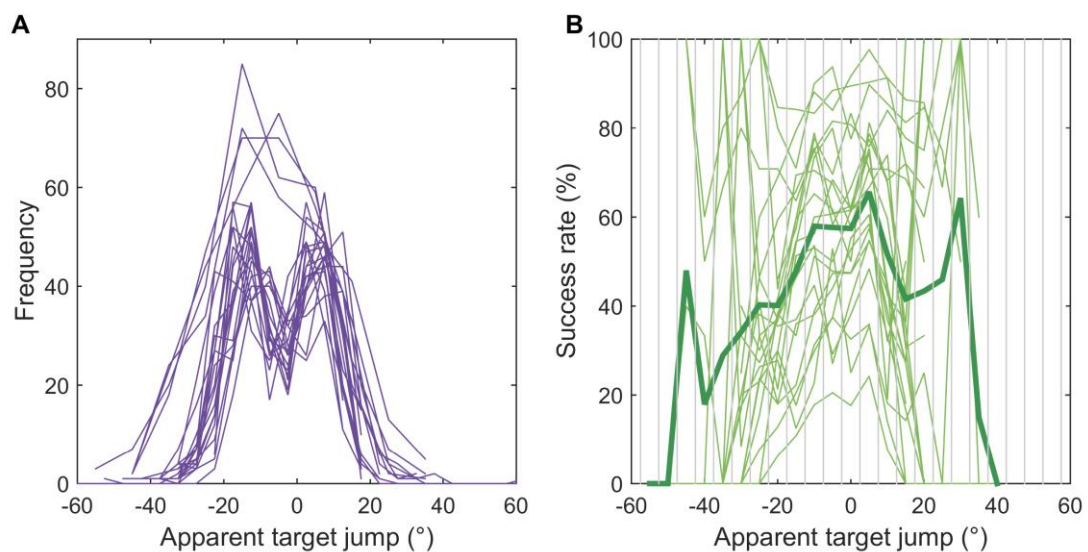


Figure 4.3 Apparent target jump differed from the predetermined target jump shown in Figure 4.2B. Because target jump was calculated relative to subject's reach angle, the visual appearance of the jump was not the same magnitude as the pre-determined magnitude. Here, apparent jumps are shown. (A) Histogram to show distribution of apparent target jumps; each line shows data for one subject. (B) Success rate by apparent target jump; thin lines show individual subject data; thick line shows mean. Grey vertical lines indicate bin edges. Both graphs show data only between $\pm 60^\circ$ for scaling reasons although there were very low numbers of trials beyond this.

The aim of the jump trial blocks was to produce a context in which subjects did not need to use a stringent control policy early in the movement. If this was achieved, pre-jump variability would be expected to be higher in jump trials than the equivalent portion of the movement in standard trials, because there would be low levels of correction of positional error.

Data was recorded as x and y positions; these trajectories were truncated at the occluder, interpolated to give the same number of data points on each trial (see Methods), and the angle on each trial relative to the midline was calculated, after subtracting that trial's rest position. These analyses are equivalent to those in Chapter 3, except that trajectories were not truncated in Chapter 3 i.e. these analyses only look at an early 25% of the movement.

Figure 4.4A and B show that mean reach angle was different between jump and no-jump trials in leftward movements, but not rightward movements. For left-cue movements, a t-test using random field theory showed one suprathreshold cluster, from 0.00% to 99.0% of movement, $t(21)$ threshold = 2.50, $p < .001$. For right-cue movements, a t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.49. Figure 4.4C and D show that the prediction of higher variability in jump trials was borne out. Reach angle variability is higher for jump trials in both leftward and rightward movements. For left-cue movements, a t-test using random field theory showed one suprathreshold cluster, from 19.8% to 99.0% of movement, $t(21)$ threshold = 2.54, $p = .012$. For right-cue movements, a t-test using random field theory showed one suprathreshold cluster, from 0.00% to 99.0% of movement, $t(21)$ threshold = 2.62, $p < .001$.

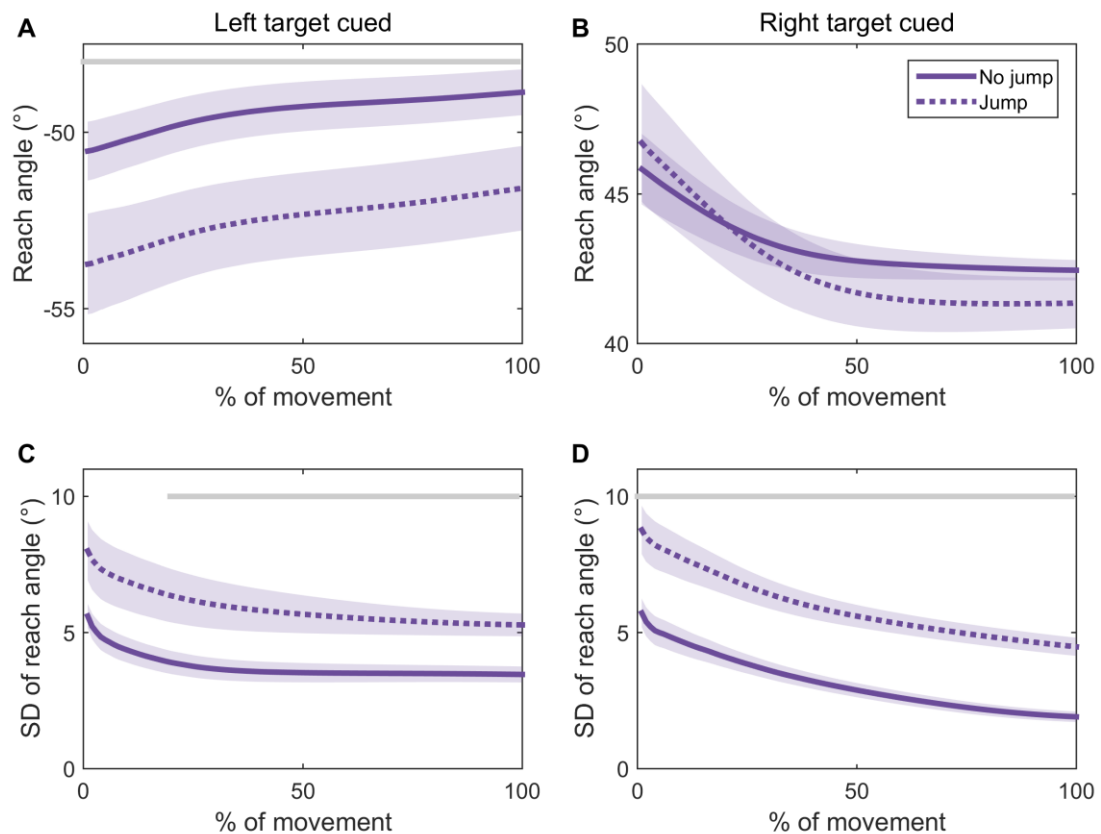


Figure 4.4 Mean reach angle and reach angle variability differ between non-jump and jump blocks. (A) and (B) show reach angle in interpolated trajectories for start of movement until occluder. (C) and (D) show the standard deviation of the reach angle. Left panel shows left-cue trials and right panel shows right-cue trials. Grey bars indicate clusters of the movement for which there is a significant effect of block type (jump vs non-jump). Shaded area shows standard error of the mean.

4.3.2. Reaction time and movement time

The mean reaction time during the experiment was 384 ms ($SD = 50.2$ ms). A three-way prior \times block type \times target cued ANOVA showed a main effect of prior expectation on reaction time ($F(1, 21) = 8.46, p = .008, \eta^2_p = .29$), but not block type ($F(1, 21) = 1.67, p = .21, \eta^2_p = .074$) or target cued ($F(1, 21) = 3.02, p = .097, \eta^2_p = .13$). There was also a target \times prior interaction ($F(1, 21) = 158, p < .001, \eta^2_p = 0.88$), but no other significant interactions. This ANOVA is shown in full in Appendix 1.

These results replicate the findings in Chapter 3, where an expectation to move to a particular target was found to decrease the reaction time for

movements to that target but increase the reaction time for movements to the other target.

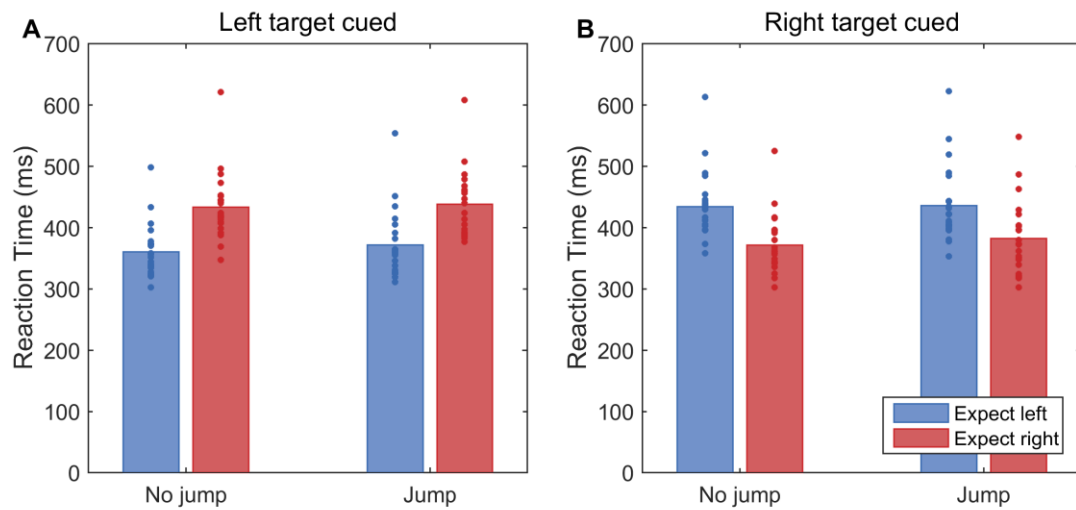


Figure 4.5 Reaction times are modulated by prior expectation but not block type. Mean reaction time by expectation and block type in (A) left cue and (B) right-cue trials. Filled circles show individual subject data

The mean movement time during the experiment was 616 ms ($SD = 59.1$ ms). Movements were slower in jump blocks than non-jump blocks. A three-way prior \times block type \times target cued ANOVA showed a main effect of target cued ($F(1, 21) = 7.20, p = .014, \eta^2_p = .26$) and of block type ($F(1, 21) = 39.9, p < .001, \eta^2_p = 0.66$). There was no main effect of prior expectation ($F(1, 21) = 0.11, p = .75, \eta^2_p = .005$), and no significant interactions between these factors. This ANOVA is given in full in the appendix.




Figure 4.6 Movement times are modulated by block type but not prior expectation. Mean movement time by expectation and block type in (a) left-cue and (B) right-cue trials. Filled circles show individual subject data.

4.3.3. Replication of previous study

The non-jump blocks were highly similar to the experiment presented in Chapter 3, although with only two prior levels ('Expect left' and 'Expect right'; there was no 'Neutral' condition). This allows us to check for a replication by using trials in standard blocks only (original analysis: Figure 3.10).

For reach angle, the same pattern was found as in the previous experiment: there was a mean difference in reach angle for leftward movements but not rightward movements (Figure 4.7A and B). For left-cue movements, a t-test using random field theory showed one suprathreshold cluster, from 0.00% to 99.0% of movement, $t(21)$ threshold = 2.63, $p < .001$. For right-cue movements, a t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.65.

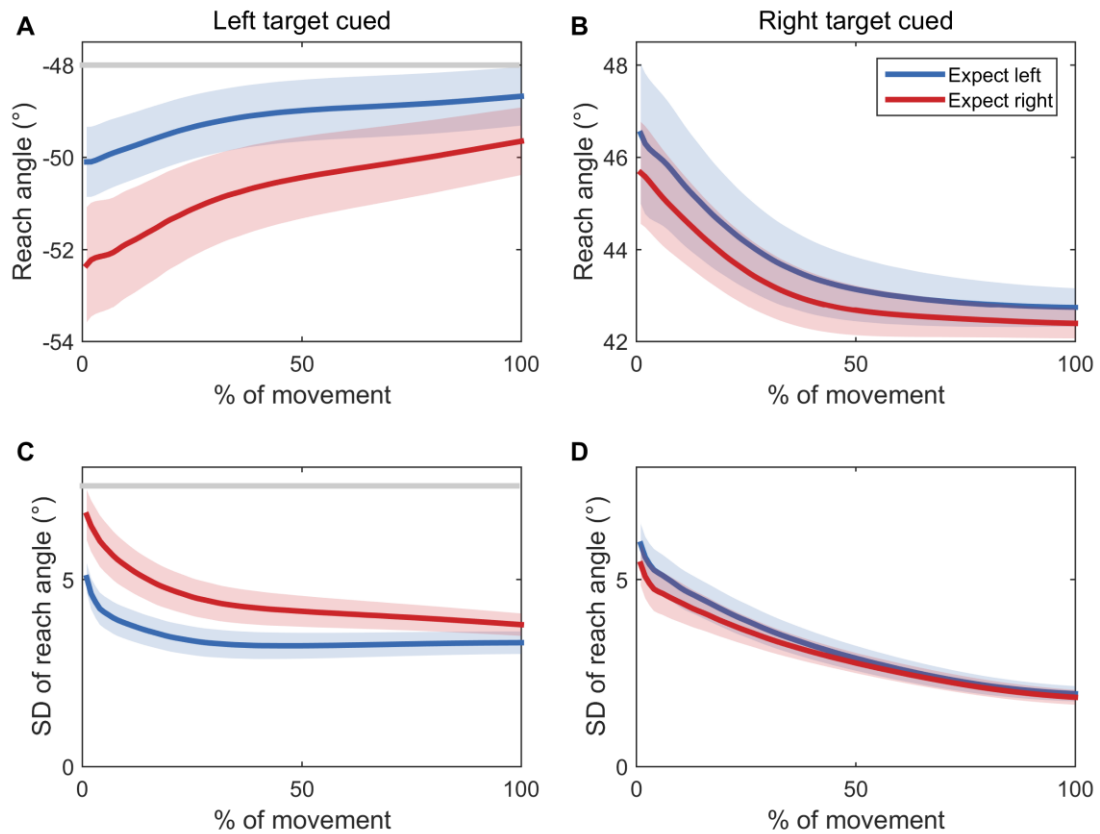


Figure 4.7 As in the previous study, for trials in non-jump blocks, mean reach angle and reach angle variability depend on expectation in left-cue but not right-cue trials. The aim of this analysis was to check for a replication of the results of Chapter 3. Top panel: mean reach angle for movement until occluder for interpolated trajectories in left cue (A) and right-cue (B) trials. Bottom panel: mean standard deviation of reach angle in left-cue trials (C) and right-cue trials (D). Grey bars indicate clusters of movement for which there is a significant effect of prior expectation condition. Shaded error bars show standard error of the mean.

Similarly, for reach angle variability (Figure 4.7C and D), a difference in reach angle variability in left-cue movements, but not right-cue movements, was found. For left-cue movements, a t-test using random field theory showed one suprathreshold cluster, from 0.00% to 99.0% of movement, $t(21)$ threshold = 2.76, $p < .001$. For right-cue movements, a t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.76. Again, this replicates the previous study.

4.3.4. An interaction between block type and prior cue?

Having replicated the previous study, jump trials were now included in the analysis. This allows us to test whether the effect of expectation on

reach angle mean and variability was reduced or obliterated in jump trials, as hypothesised.

Figure 4.8A and B show reach angle between movement start and reaching the occluder. For left-cue movements, a two-way ANOVA using random field theory showed an effect of prior (one suprathreshold cluster, from 0.00% to 99.0% of movement, $F(1, 21)$ threshold = 5.84, $p < .001$), an effect of block type (one suprathreshold cluster, from 0.00% to 99.0% of movement, $F(1, 21)$ threshold = 5.84, $p = .005$), but no interaction ($F(1, 21)$ threshold = 5.84).

For right-cue movements, a two-way ANOVA using random field theory showed an effect of prior in the later part of the movement (one suprathreshold cluster, from 43.6% to 99.0% of movement, $F(1, 21)$ threshold = 6.37, $p = .027$), no effect of block type ($F(1, 21)$ threshold = 6.37) and no interaction ($F(1, 21)$ threshold = 6.37).

Figure 4.8C and D show reach angle variability between movement start and reaching the occluder. The litmus test of the hypothesis proposed is, in leftward movements, the presence of an interaction between prior and trial type. In fact, this interaction was absent. Instead, for leftward movements, there was a main effect of both prior expectation and block type. A two-way ANOVA using random field theory showed an effect of prior (one suprathreshold cluster, from 0.00% to 99.0% of movement, $F(1, 21)$ threshold = 6.58, $p = .002$), a later effect of block type (one suprathreshold cluster, from 30.2% to 99.0% of movement, $F(1, 21)$ threshold = 6.58, $p = .014$) and no interaction ($F(1, 21)$ threshold = 6.58).

For right-cue movements, a two-way ANOVA using random field theory showed no effect of prior ($F(1, 21)$ threshold = 6.99), an effect of block type (one suprathreshold cluster, from 0.00% to 99.0% of movement, $F(1, 21)$ threshold = 6.99, $p < .001$) and an early interaction (one

suprathreshold cluster, from 0.00% to 30.6% of movement, $F(1, 21)$ threshold = 6.99, $p = .033$).

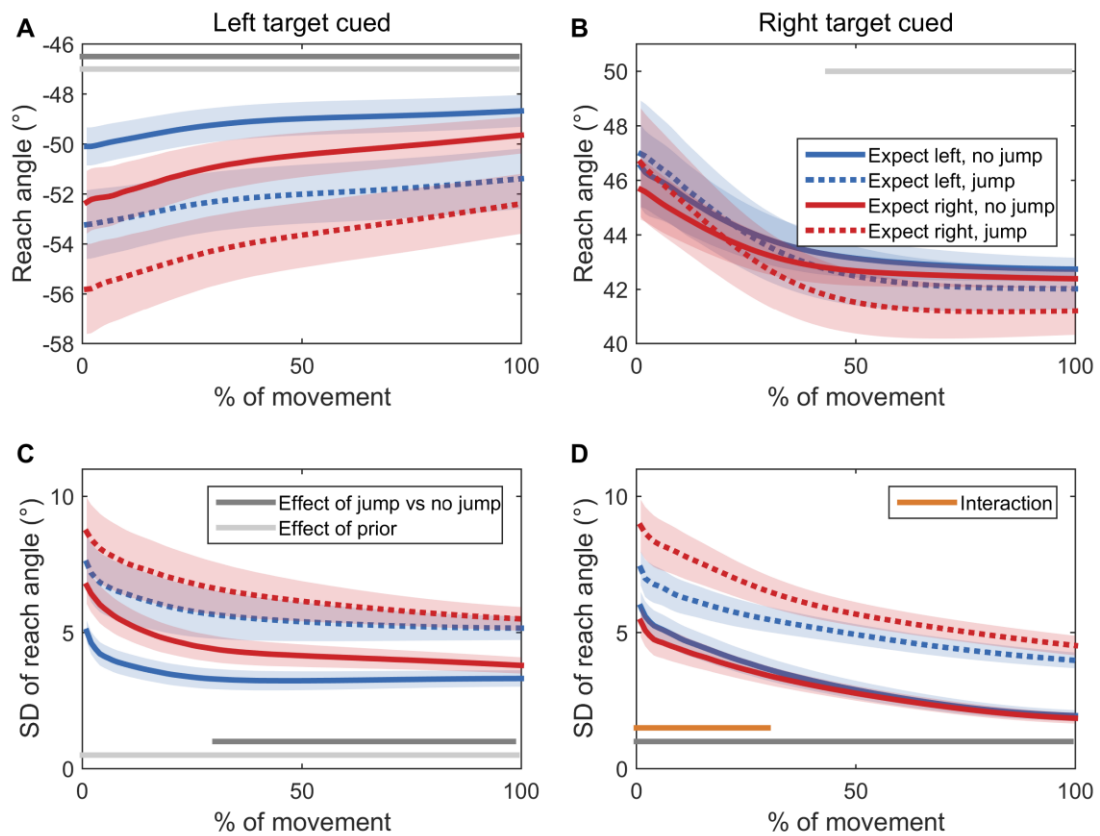


Figure 4.8 Effect of block type and expectation on reach angle mean and variability. This figure shows the key analysis of interest in this experiment. A significant interaction between block type (jump vs no jump) and prior expectation in panel C would have confirmed the hypothesis by indicating that variability is only modulated by prior expectation in the presence of a stringent control policy (i.e. non-jump blocks). Top panel: mean reach angle for interpolated trajectories until occluder in left-cue trials (A) and right-cue trials (B). Bottom panel: mean standard deviation of reach angle in left-cue trials (C) and right-cue trials (D). Light grey bars show a significant effect of prior expectation condition. Dark grey bars show a significant effect of block type (jump vs non-jump). Orange bars show a significant interaction between these two factors. Shaded error bars show standard error of the mean.

For the data shown in Figure 4.8D, two post-hoc t-tests were conducted in order to understand the trends underlying the interaction. For non-jump trials, variability was compared in the ‘Expect left’ and ‘Expect right’ conditions, and found no effect here: a t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.76. Variability for the ‘Expect left’ and ‘Expect right’ conditions in the jump trials was also compared, and again found no effect. A t-test using random field theory

showed no suprathreshold clusters, $t(21)$ threshold = 2.70. Thus there was no significant difference of prior expectation in rightward movements in either jump or non-jump trials considered in isolation.

4.3.4.1. Was the lack of effect due to lack of power?

I wanted an indication of whether the absence of the expected interaction was due to lack of power or not. Therefore movement variability was averaged across the sample (0% to 100% of movement) for 'left target cued' data only and a Bayesian ANOVA was conducted, which compared the five models given in the Table 4.1.

Models	P(M)	P(M data)	BF _M	BF ₁₀	% error
Null model (incl. subject)	0.200	3.900e -4	0.002	1.000	
Prior	0.200	3.047e -4	0.001	0.781	0.922
Block type	0.200	0.400	2.668	1025.923	2.026
Prior + Block type	0.200	0.470	3.550	1205.767	3.772
Prior + Block type + Prior * Block type	0.200	0.129	0.593	330.844	1.634

Note. All models include subject.

Table 4.1 Bayesian ANOVA shows anecdotal evidence against the interaction of interest, and extremely strong evidence for the two-main effect model. This ANOVA was conducted to give an indication as to whether the study was underpowered to detect a block type (jump/non-jump) x prior (Expect left/Expect right) interaction. The ANOVA showed anecdotal evidence against the interaction i.e. tends away from supporting the idea that the study was underpowered.

In keeping with the random field theory ANOVA, the most likely model was that with a main effect of prior and a main effect of trial type. This model had a Bayes Factor against the null hypothesis of 1,206, indicating these data are 1,206 times more likely to be observed under this model. A Bayes Factor of > 100 is usually interpreted as 'Extreme evidence' (Wagenmakers et al., 2011).

The comparison of interest was the model with the interaction over the model without. How strong was the evidence against the presence of an interaction in addition to the two main effects? This is given by finding the ratio of the Bayes Factors for the model with the interaction against the model with the main effects alone. This gives a Bayes Factor of 0.274.

This can be interpreted as ‘anecdotal’ or ‘barely worth mentioning’ evidence against the interaction (Kass and Raftery, 1995).

To summarise, contrary to the hypothesis, an effect of prior on reach angle variability is present in jump trials. This could be because higher variability in the jump trials arises not due to a genuinely less stringent control policy, but for other reasons. I investigated two possible confounds that could explain variability differences between jump and non-jump trials in the following section. These are: (1) movement velocity: faster movement in jump trials could lead to higher variability; (2) lower success: exploration in space in jump trials following negative feedback could lead to higher variability.

4.3.5. Early movement time and movement velocity




Figure 4.9 Block type affects time taken to reach occluder. The aim of this analysis was to discover whether pre-occluder movement time differed by block type. Mean movement time to reach occluder by block type (jump vs non-jump) and target cued. Filled circles show individual subject data.

Above, an analysis of total movement time was presented, finding that subjects were slower on jump blocks. However, this includes time taken to respond to the target jump and does not give an indication of their movement speed prior to the occluder, in the portion of the movement that has been considered in the other analyses.

Analysing the movement time to the occluder showed that movements to the occluder were faster in jump blocks compared to non-jump blocks. A two-way block type x target cued ANOVA showed a main effect of block type, $F(1, 21) = 6.59, p = .018, \eta^2_p = .24$, no effect of cue direction (left vs right), $F(1, 21) = 3.89, p = .062, \eta^2_p = .16$, and a significant interaction, $F(1, 21) = 8.84, p = .007, \eta^2_p = .30$.

This result shows that movements to the occluder were faster in non-jump trials than in jump trials. This could have been a source of the variability difference observed between these two trial types (see Discussion). This raised the question of whether similar differences in early movement time exist between movements at different levels of prior expectation, which have been demonstrated to show different levels of variability. I wanted to check whether this was the case, to check this difference did not drive variability differences seen in the previous chapter.

I have already shown no difference of prior expectation on overall movement time (see previous chapter). To further investigate this, the analysis above was repeated (movement time to reach a movement of amplitude 5 cm, which was the location of the occluder) on the dataset in Chapter 3. I found no effect of prior expectation of this metric in leftward movements, $F(2, 26) = 0.78, p = 0.47, \eta^2_p = 0.056$, and no effect for rightward movements, $F(2, 26) = 0.31, p = 0.73, \eta^2_p = 0.024$. This confirmed that early movement time differences were not the cause of variability differences in the previous experiment.

I next analysed velocity during the movement prior to the occluder. X velocity in leftward movements showed a significant effect of block type for most of the movement: a t-test using random field theory showed one suprathreshold cluster, from 0.00% to 84.6% of movement, $t(21)$ threshold = 2.90, $p < .001$. For x velocity in rightward movements, a t-test

using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.84.

Y velocity in leftward movements showed no effect of block type: a t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.85. For y velocity in rightward movements, a t-test using random field theory showed one suprathreshold cluster, from 22.7% to 72.6% of movement, $t(21)$ threshold = 2.91, $p < .001$.

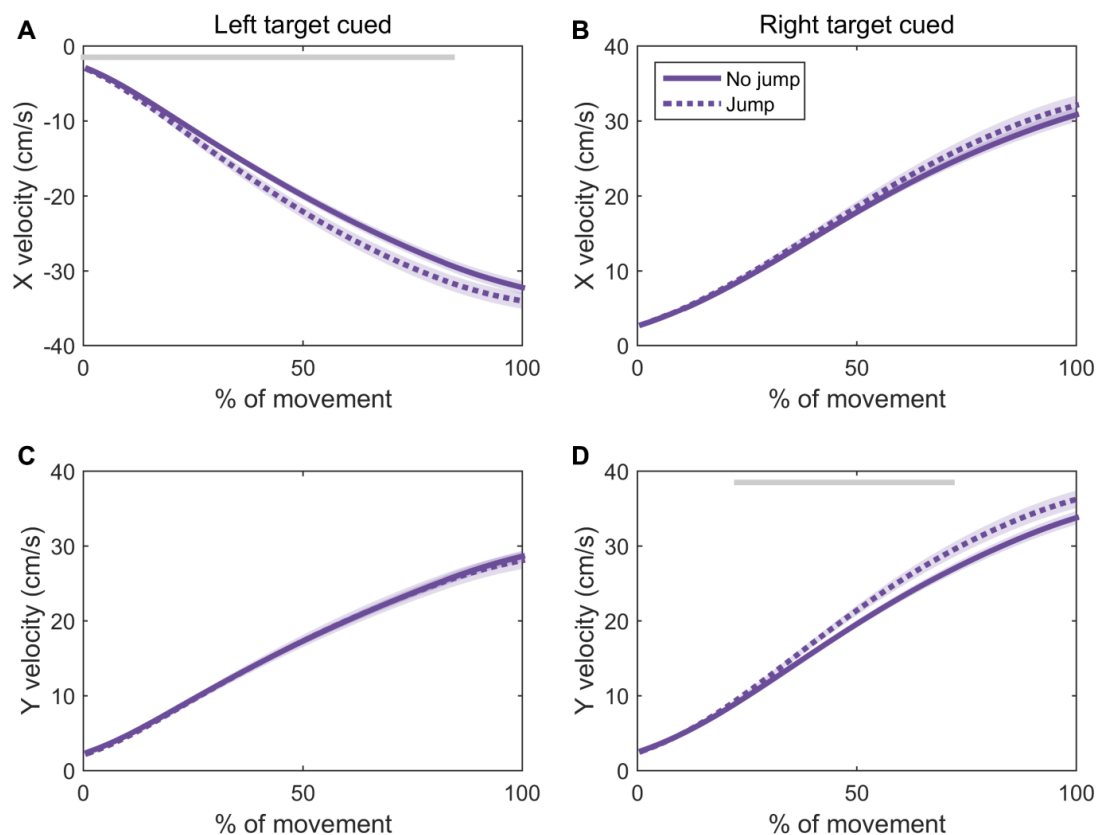


Figure 4.10 Block type affects movement velocity in movement prior to the occluder. Top panel: mean X velocity in left-cue (A) and right-cue (B) trials. Bottom panel: mean Y velocity in left-cue (C) and right-cue (D) trials. Grey bars show a significant effect of block type (jump vs non-jump). Shaded error bars show standard error of the mean.

4.3.6. Variability as determined by success

Previous studies have indicated that human movement variability increases in response to negative feedback (Galea et al., 2013; Pekny et al., 2015). Specifically, variability depends strongly on whether the

previous trial was successful or not (Pekny et al., 2015). Given that success rates differed greatly between non-jump and jump blocks (see above), I wanted to explore whether I also observed this effect of success on variability in subsequent trials. If so, this could indicate that differences in variability between the two block types were driven by differences in success, rather than in movement strategy.

Trials were divided by the block type (jump or non-jump) and whether the previous trial had ended in a target hit (and positive feedback) or a miss (with 'Points: 0' feedback). Since the experiment involved two separate reaching movements (leftward or rightward) and since I was interested in the effect of success or failure in amending the same movement in future, I used only trials in which the same movement was cued twice in a row (i.e. leftward, leftward or rightward, rightward).

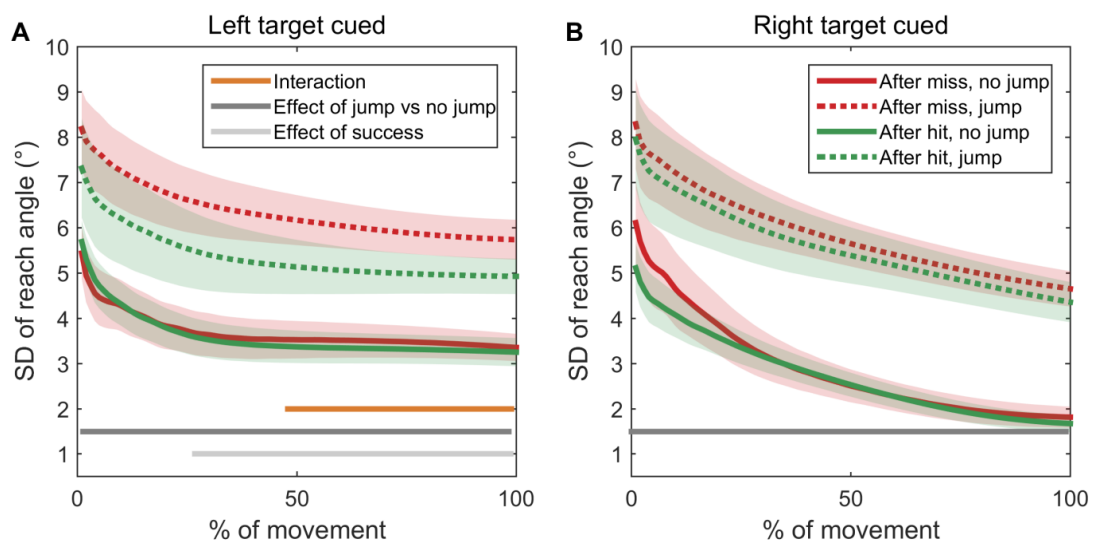


Figure 4.11 Variability is increased by lack of success for some trials only. The aim of this analysis was to discover whether unsuccessful trials drove an increase in variability on the subsequent trial. The figure shows reach angle variability on trials sorted by whether the previous trial was a hit or a miss, and by block type. (A) shows left-cue trials and (B) shows right-cue trials. Only trials where the same direction was cued twice in a row have been included. Dark grey bar shows significant effect of block type (jump vs non-jump). Light grey bar shows significant effect of success (hit on previous trial vs miss on previous trial). Orange bar shows significant interaction. Shaded error bars show standard error of the mean.

The variability in these groups of trials showed, as in Figure 4.4, a modulation by block type. Only leftward movements showed the effect of interest: higher variability following unsuccessful trials over successful ones (Figure 4.11).

For leftward movements, there were significant main effects and an interaction. A two-way ANOVA using random field theory showed an effect of success (one suprathreshold cluster, from 26.7% to 99.0% of movement, $F(1, 21)$ threshold = 6.78, $p = .008$), an effect of block type (one suprathreshold cluster, from 1.29% to 99.0% of movement, $F(1, 21)$ threshold = 6.78, $p = .002$), and a significant interaction (one suprathreshold cluster, from 47.9% to 99.0% of movement, $F(1, 21)$ threshold = 6.78, $p = .020$).

For rightward movements, the only significant finding was a main effect of block type (jump vs no jump). A two-way ANOVA using random field theory showed no effect of success ($F(1, 21)$ threshold = 6.90), an effect of block type (one suprathreshold cluster, from 0.00% to 99.0% of movement, $F(1, 21)$ threshold = 6.90, $p < .001$), and no interaction ($F(1, 21)$ threshold = 6.90).

For the leftward movements, visual inspection and the significant interaction suggested there was no effect of success in the no jump trials (Figure 4.11A; solid lines). To test this, a post-hoc t-test, was conducted which confirmed there was no effect of success (hit vs miss) for these trials. The t-test using random field theory showed no suprathreshold clusters, $t(21)$ threshold = 2.84. In other words, the variability difference for hit vs miss trials in leftwards movements was driven purely by jump trials.

4.4. Discussion

I wanted to test whether a previously-described (Chapter 3) effect of prior expectation on motor variability in a reaching movement was dependent on the control policy used in the movement. The experiment presented in this chapter examined the effect of prior expectation on motor variability, in two contexts: (1) a simple 45° reaching movement; (2) the same reaching movement but involving a target jump 25% of the way through the movement.

4.4.1. Effect of expectation on movement variability successfully replicated

Because the reaching movements that did not involve a jump are equivalent to those in the experiment presented in Chapter 3, I was able to test whether we replicate the results of this experiment. Indeed, the findings are replicated faithfully (Figure 4.7). As in the previous experiment, prior expectation to move modulated mean reach angle in leftward but not rightward movements. An expectation to move left reduced movement variability in leftward movements, but this was not mirrored in rightward movements. Again, this matches with the findings of the previous experiment.

The reaction time and movement time findings also matched the previous study: again there was an effect of prior expectation cue on reaction time but not movement time.

4.4.2. Target jump manipulation increased movement variability

The ‘target jump’ trials were included in this experiment with the aim of introducing a type of trial in which the early part of the movement did not contribute to movement success. I surmised this would induce subjects to

adopt a lax control policy where there was little correction for deviations from the mean trajectory.

If subjects had indeed adopted such a control policy, I would expect to see higher variability in the movements in target jump blocks. Indeed, this difference was present (Figure 4.4). The target jump blocks had a longer total movement time but there was no reaction time effect.

4.4.3. Main analysis of interest did not support hypothesis

The key question of this experiment was whether the effect already observed - a lower expectation to make a leftward movement increasing variability in that movement - was obliterated during the target jump blocks. I theorised that, during these blocks, control policy would be less stringent. I had hypothesised in the previous chapter that an effect of prior expectation is only seen in a portion of the movement in which there is active control of variability, in order to explain why the effect of prior expectation is seen in leftward but not rightward movements. So I hypothesised, on blocks where there was a target jump, there would be a reduced or no difference between variability across different levels of prior expectation in leftward movements.

However, a lower expectation to move increased variability in jump trials just as it did in non-jump trials. There was a main effect of prior expectation and of block type (jump versus non-jump) in leftward movements, but no interaction between these factors.

Thus the findings provide no evidence for the hypothesis. This could be because the hypothesis is false. However, a second possibility is that the target jump manipulation did not, in fact, induce a lax control policy in participants. Under this interpretation of the data, subjects did in fact use a stringent control policy on jump blocks, perhaps because this somehow conferred an advantage in the process of making an adjustment in

response to the target jump. If this was indeed the case, how can the higher variability on jump blocks be explained?

I identified two factors, not relating to control policy, which could have been driving variability differences between the target jump and non-jump trials: (1) movement velocity, and (2) success rate. These are discussed in turn; I then discuss whether the hypothesis was falsified.

4.4.3.1. Movement velocity was higher in jump blocks

Despite an overall longer movement time in jump blocks (

Figure 4.6), movement times for the portion of the movement which was analysed were shorter in jump blocks (Figure 4.9). This is a sensible strategy for these trials: moving faster in the early portion of the trial buys as much time as possible in which to redirect the movement in response to the target jump. In keeping with this difference in movement time, velocities were also faster in jump trials, although whether this is evident in the x or y component of velocity differs between leftward and rightward movements (Figure 4.10).

So a higher velocity in target jump trials could be the cause of the increased variability seen in these trials. If this was the case, it would explain why the effect of prior was present in these trials. This would suggest that there was still a control policy which stringently controlled deviations, but a different speed-accuracy tradeoff (Heitz, 2014) had been used.

If it were true that these movement time differences were driving variability, it would mean the variability metric used (standard deviation in reach angle) is highly sensitive to small movement time differences of ~10 ms over a 5 cm (~190 ms) movement. I therefore checked that

differences in early movement time were not present in the previous experiment, and confirmed that there were none.

One way to try and minimise speed-accuracy tradeoff differences in a future version of this experiment could be to attempt to fix movement speed by training and instructing subjects appropriately. The most constrained way to do this would be to attempt to standardise the timing of both the first (pre-jump) and second (post-target jump) parts of the movement. If a subject reached either the jump position or the final target too early or too late, visual cues would indicate to a subject that that trial had been discarded, and whether this was due to being too fast or too slow. Through training subjects would learn to time their movements within the bounds chosen.

However, the disadvantages of this approach are that it would likely lead to lots of lost trials due to incorrect timing and/or an extensive training requirement. Furthermore, it seems likely that these techniques do not represent a truly standardised movement speed, and instead lead to speed being restricted to a narrow range, which may not fully obliterate differences between conditions.

4.4.3.2. Did success rate drive variability differences?

There was a marked reduction in success rate in the jump blocks compared to the non-jump blocks (Figure 4.2A). This was in contrast to the previous experiment where success differences between conditions were much smaller. I considered whether this strong difference in success could have been the source of the variability difference between jump and non-jump blocks, confounding the results.

Reward information has been shown to alter motor variability (Takikawa et al., 2002; Izawa and Shadmehr, 2011; Galea et al., 2013), including on a trial-by-trial basis (Pekny et al., 2015). Here I drew on a study by Pekny et

al (2015). This study found that variability in a reaching task increased in response to failure on the preceding trial. Furthermore, when reward probability was lower, global reach variability was higher.

In a similar analysis, I found, in leftward movements only, a main effect of success on previous trial for the later two thirds of the movement. There was an interaction with block type (jump vs non-jump). Post-hoc analysis confirmed the effect of success was not present in non-jump blocks (Figure 4.11).

It is interesting that I replicated the results of Pekny et al in just one movement, for one trial type only, rather than more widely. Perhaps the effect was present in jump trials because the difficult nature of these trials and lower success rate encouraged exploration. Pekny et al's study encouraged exploration by, unknown to participants, shifting the rewarded zone periodically. This might have prompted participants to adopt a 'when you don't succeed, explore in space' strategy. In other words, their experimental finding might be specific to tasks with a benefit to spatial exploration. However, my study was not designed for these analyses and selecting only trials where the same target was cued twice consecutively reduced the amount of data available and might have meant the study was underpowered for these particular analyses.

The key finding from this analysis of success is that analysing by success does not obliterate the effect of block type (jump vs non-jump) on variability. In other words, there is a variability increase due to jump trials beyond the one driven by success. This suggests that differences in success rate are not sufficient to explain the variability increase caused by the experimental manipulation.

4.4.4. Was the hypothesis falsified?

A velocity difference between the conditions could have driven the increased variability, which would mean this was not evidence that subjects adopted a laxer control policy for the jumping blocks. An alternative explanation is that subjects did indeed adopt this lax control policy and the lack of the predicted result indicates the hypothesis – that an effect of prior on variability depends on a stringent control policy being planned – was false.

If this was the case, this would indicate the key finding was that, even when subjects had a lax control policy because it was not important for the early movement to be accurate, a lower prior expectation still increased movement variability. In short, the effect of prior expectation does not depend on movement control policy, and instead only on planning.

However, the problem with this explanation is that it struggles to reconcile the results of the previous chapter. These were that rightward movements were not susceptible to a change in variability as a consequence of prior expectation. Indeed, I replicated this effect again in the most recent experiment and also found it applied to jump blocks. It is difficult to find an explanation for this that does not rely on differing dynamics of the two movements (i.e. different control policies).

Explanations that are based on the idea that there is a floor effect for variability of rightward movements are unsatisfactory as I showed in the previous chapter that initially rightward variability is higher than left; this informed the idea that control of variability in the rightward movement is less stringent.

Therefore, the explanation for the findings based on measured differences in movement velocity between jump and non-jump blocks

(Section 4.4.3.1) seems easiest to reconcile with the data across the two chapters.

4.4.5. Was the study underpowered?

One possibility is the sample size lacked the power to demonstrate a larger effect of prior in non-jump trials than in jump trials. To try and assess this, I conducted a Bayesian ANOVA. This found anecdotal evidence against the presence of the interaction. In other words, this analysis tended to suggest the hypothesis was false rather than the study underpowered, but only weakly.

The sample size was 22 subjects, which is larger than typical amongst similar studies (Pekny et al., 2015; Wijdenes et al., 2016). So, rather than increasing the sample size, reducing measurement noise via the use of a different manipulation seems a more fruitful way to attempt to resolve this problem.

4.4.6. Conclusion

To conclude, the evidence did not support the hypothesis proposed. It could be that the hypothesis is incorrect, but this poses a challenge in understanding the experiment in the previous chapter. It could also be that the study is underpowered, but the evidence from the Bayesian ANOVA leans towards suggesting it was not.

An alternative possibility is that subjects did not adopt the strategy of a lax control policy. Instead, greater variability was caused by other differences between the target jump and non-jump blocks. Specifically, a faster velocity in the early part of the movement for jump blocks could have driven higher variability.

4.5. Appendix I: ANOVAs on reaction time and movement time

4.5.1. Reaction time

Within Subjects Effects						
	Sum of Squares	df	Mean Square	F	p	η^2_p
Target	1128.08	1	1128.08	3.018	0.097	0.126
Residual	7850.22	21	373.82			
Prior	1450.61	1	1450.61	8.461	0.008	0.287
Residual	3600.47	21	171.45			
Block Type	2263.57	1	2263.57	1.671	0.210	0.074
Residual	28448.23	21	1354.68			
Target * Prior	179593.51	1	179593.51	157.988	< .001	0.883
Residual	23871.89	21	1136.76			
Target * Block Type	35.42	1	35.42	0.359	0.556	0.017
Residual	2073.90	21	98.76			
Prior * Block Type	17.86	1	17.86	0.181	0.675	0.009
Residual	2069.37	21	98.54			
Target * Prior * Block Type	670.18	1	670.18	4.009	0.058	0.160
Residual	3510.72	21	167.18			

Table 4.2 Reaction time showed a significant main effect of prior expectation and a significant target x prior interaction. Table reports a three-way repeated measures ANOVA on reaction times with factors target cued (left-cue/right-cue) x block type (jump/non-jump) x prior expectation cue (Expect left/Expect right).

4.5.2. Movement time

Within Subjects Effects						
	Sum of Squares	df	Mean Square	F	p	η^2_p
Target	3344.207	1	3344.207	7.200	0.014	0.255
Residual	9754.293	21	464.490			
Prior	26.111	1	26.111	0.105	0.749	0.005
Residual	5221.674	21	248.651			
Block Type	314608.017	1	314608.017	39.860	< .001	0.655
Residual	165750.445	21	7892.878			
Target * Prior	870.198	1	870.198	1.885	0.184	0.082
Residual	9697.001	21	461.762			
Target * Block Type	224.350	1	224.350	0.929	0.346	0.042
Residual	5070.148	21	241.436			
Prior * Block Type	8.951	1	8.951	0.089	0.768	0.004
Residual	2112.275	21	100.585			
Target * Prior * Block Type	239.875	1	239.875	1.351	0.258	0.060
Residual	3727.578	21	177.504			

Table 4.3 Movement time showed a significant main effect of target cued and block type. Table reports a three-way repeated measures ANOVA on movement times with factors target cued (left-cue/right-cue) x block type (jump/non-jump) x prior expectation cue (Expect left/Expect right).

Chapter Five: General Discussion

This thesis looked at how prior expectation acts on motor processes. Firstly, I studied prior expectation in motor planning. A series of TMS experiments were carried out to attempt to substantiate a hypothesis that prior expectation biases motor system excitability prior to a decision (Chapter 2). The results of the experiment did not support the hypothesis. It is unclear whether the hypothesis has been falsified or whether TMS is not a suitable method for studying small behavioural biases, despite a number of previous studies which suggested it might be promising (Bestmann et al., 2008; Michelet et al., 2010; Klein-Flugge and Bestmann, 2012).

The second part of the thesis studied prior expectation in motor planning, specifically in the context of theories which have proposed how a limited resource in motor planning influences variability during movement (Wijdenes et al., 2016). I proposed a new hypothesis: that prior expectation unevenly weights the allocation of resources to motor plans, and thus will lead to unequal variability between two planned movements.

In Chapter 3, I found evidence to support this hypothesis, with the expected effect present in one movement but not the other. I proposed this was due to an interaction with the specific control policies of the movement and tested this idea in Chapter 4. I did not find support for this idea in Chapter 4; I proposed that this might be because the experiment was confounded by differing movement velocities across conditions.

In this general discussion, I will first revisit and broaden the themes presented in Chapter 1 which have been pertinent to the design of the experiments in this thesis. In particular, I will outline what might be the next stages in investigating these theories, and give some of the more

novel and contested evidence. As this thesis has used human experiments, I will ask to what extent these developments can be tested in humans.

I will then return to the experiments presented in this thesis, specifically those on motor variability, and present some ideas for the next steps in this line of research.

5.1. Where next for the affordance competition hypothesis?

The idea that multiple motor plans are formed during a decision about movement is now prominent, and heavily informed the experiments in this thesis. Evidence for this idea comes from neurophysiology (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek, 2011), behavioural studies (Chapman et al., 2010; Stewart et al., 2013; Gallivan et al., 2015, 2016b, 2017) and TMS studies (Michelet et al., 2010; Klein-Flugge and Bestmann, 2012).

As detailed in Chapter 1, there has been a shift away from the idea that motor cortex is a population code representing movement parameters such as reach direction (Georgopoulos and Carpenter, 2015), which is intertwined with the modelled instantiations of the affordance competition hypothesis (Cisek, 2007). In particular, a new line of evidence supports the idea that motor cortex acts as a dynamical system, and so the role of preparatory activity is to bring the neural population to a particular state, from which it can passively evolve into movement activity (Churchland et al., 2006, 2012; Churchland and Shenoy, 2007; Shenoy et al., 2013; Kaufman et al., 2014). It has therefore been argued that the multiple motor plans of the affordance competition hypothesis are more accurately multiple motor goals (Wong and Haith, 2017). (See Chapter 1 for more.) Leaving aside this debate, the core findings of (1) representation of ‘cognitive’ variables in motor cortex when they are relevant to a decision (Cisek and Kalaska, 2005; Pastor-Bernier and Cisek,

2011) and (2) competition and interaction between motor representations (Ghez et al., 1997; Chapman et al., 2010; Gallivan et al., 2015) still stand.

Having established this hypothesis of competition between motor plans, what is the next step in investigation of these theories? One avenue could be to move beyond the recognition that activity relating to a decision is represented in multiple brain areas to studying the distinct role of these brain areas in an ongoing decision to move. The affordance competition hypothesis talks about a ‘distributed consensus’ of decision activity; what is the unique contribution of each part of the consensus?

Work in rats has elucidated distinct representations of accumulating evidence during a decision to move. These studies focussed on a perceptual decision requiring evidence accumulation. While posterior parietal cortex appears to encode accumulating evidence in a graded manner, the tuning curve in the FOF (Frontal Orienting Fields; rat premotor homologue) is steeper; analogous to a categorical representation of information encoding the current best choice (Hanks et al., 2015). Furthermore, brief optogenetic inactivation of the FOF only induces an ipsilateral choice bias if it occurs near the end of the decision time – supporting the idea that it is involved in committing to or remembering (Piet et al., 2017) a choice, or categorising accumulated evidence into a binary choice, but not the evidence accumulation itself (Erlich et al., 2015). This is an elegant way to show how representation of decisions in motor areas are not a straightforward translation of the evolving decision and involve modification of the information to be more suitable to subserve action.

The idea of multiple representations of similar decision information raises an interesting question: if multiple regions show evidence-accumulator activity, which is the underlying evidence accumulator?

Erlich et al (2015) found that pharmacological FOF inactivation (i.e. over a longer time period than optogenetic inactivation) introduced a strong behavioural bias consistent with a role as the output pathway of the evidence accumulator. By contrast, inactivation of the posterior parietal cortex did not cause a choice bias, except on trials in which the animal was free to give a random response of its choice. In short, the finding is consistent with the posterior parietal cortex being a weak alternative evidence accumulator that only plays a role when there is no incoming sensory evidence. The finding has been replicated in primates where LIP inactivation again only biased decisions on the free choice trials (Katz et al., 2016). This is interesting because seminal neurophysiological decision-making work referenced in Chapters 1 and 2 has focussed on lateral intraparietal cortex (LIP), under the assumption that this area is the site of the evidence accumulator (reviewed in Huk et al., 2017). Yet it seems that an activity correlation with the ongoing decision does not equal causation.

So which is the crucial region for evidence accumulation? Brody and Hanks (2016) report that currently unpublished data (Yartsev et al) suggest the causal circuitry in the rat is subcortical, in a region not previously studied in the context of evidence accumulation.

Thus the next stage of research on the affordance competition hypothesis might be interest in the specific roles different brain areas play in a single decision for action. Do human studies, such as those presented in this thesis, have a role to play in extending this work? The low temporal resolution of the BOLD response makes it difficult to extract the accumulating evidence signal in fMRI studies (discussed in Hanks and Summerfield, 2017). However, M/EEG studies have more potential for measuring these time-critical signals. There has been a focus on work to find neural correlates of evidence accumulation signals, similar to those

in LIP. Accumulating signals which correlate with behavioural measures have been found in a centroparietal positivity (Kelly and O'Connell, 2013, 2015), in lateralised beta band power over motor cortex (Donner et al., 2009; de Lange et al., 2013), and in theta band power (van Vugt et al., 2012).

These studies open up the possibility of more complex analyses to decompose how the evidence accumulation signal evolves over brain areas, similar to the rat studies above. This is facilitated by new paradigms, which ask participants to make a decision about a stream of evidence presented as a series of discrete stimuli (e.g. judging whether the average axis of a series of Gabor patches is cardinal or diagonal; Wyart et al., 2012a). Because these use quantised streams of evidence, the experimenter can measure the contribution of individual pieces of evidence to the decision, as in the rat studies (Hanks et al., 2015).

Wyart et al (2012a) used this task to characterise how each stimulus added to perceptual, decision, or response information. These parameters could then be regressed against brain activity. They found distinct representations of each type of information: perceptual information encoding peaked at occipital electrodes, whilst decision information encoding was more broadly distributed and peaked later at parietal electrodes. Response information was encoded in beta-band activity over motor electrodes.

Using this analysis, the authors were able to detect that decision information was multiplicatively weighted by the phase of delta oscillations over parietal cortex, and then integrated additively over motor cortex. In other words, the accumulation signal in parietal cortex is not yet a pure response signal. This provides further evidence for distinct decision and motor stages, with distinct computations. Another approach comes from Hunt et al (2012), who correlated MEG activity with a

network attractor model to find brain areas that actively contribute to value-based choice.

These studies illustrate how it is possible for paradigms to orthogonalise distinct processes within a single motor response, and so decompose the representation of these processes.

5.2. New competitors to the drift-diffusion model

Evidence-accumulation and race-to-threshold models were introduced in Chapter 1; the hypothesis that was the basis of Chapter 2 was influenced by them. The drift-diffusion model in particular is canonical and widely used across human and animal neuroscience (e.g. Shadlen and Newsome, 2001a; Roitman and Shadlen, 2002; Mulder et al., 2012). Despite this, the drift-diffusion model as a putative neural computation continues to be debated. It is clear that the drift-diffusion model is effective in modelling reaction time distributions under various task manipulations, such as evidence level, urgency or expectation (Ratcliff, 2002; Ratcliff and McKoon, 2007). The debate has centred on whether activity in lateral intraparietal cortex is effectively a neural implementation of evidence accumulation, as has often been argued, implicitly or explicitly (Gold and Shadlen, 2007).

These classic studies average neural activity across trials, which has the potential to obscure trial-to-trial dynamics; in particular, it has the potential to make abrupt transitions look graded when averaged over many trials (Churchland and Kiani, 2016). Similarly, population heterogeneity can be obscured by averaging across a population of neurons. Some authors have analysed spike trains in single-neurons to seek evidence for models in which there is a ‘step’ between states rather than ramping activity (Durstewitz and Deco, 2007; Miller and Katz, 2010) and found evidence that spike trains in LIP during a decision-making task

indeed step from one fixed firing rate to another (Latimer et al., 2015). Under this model, the ramping activity appears only when neural responses are stimulus-locked and averaged. Stronger stimuli lead to steeper ramping, not because of a change within single neuron responses, but because of a greater proportion of up-steps over down-steps in the population. However, this has been strongly contested (Shadlen et al., 2016), and the dynamics of neural responses may depend on the particular task (Latimer et al., 2016).

A second challenge to the drift-diffusion model has come from a theoretical background. Cisek has argued that the ramping activity seen in neural activity or inferred behaviourally has always been assumed to track accumulating evidence, but there is no evidence that it is not driven by elapsed time (Cisek et al., 2009). The urgency-gating model proposes that neural activity is composed of two combined signals. Firstly, an ‘urgency’ signal that grows with time and pushes a decision towards completion (and thus action) whether or not there is decisive evidence. This can be seen as a motor signal. This is combined with the second signal, which represents momentary sensory evidence (and is thus larger if sensory evidence is strong). The group has produced behavioural and neural studies supporting the model (Cisek et al., 2009; Thura et al., 2012; Thura and Cisek, 2014).

In sum, the extent to which the DDM can be used to understand neural responses is still contested. One of the advantages of the DDM is that parameters in the model can be assigned intuitive correlates (Carpenter and Williams, 1995). The top threshold can be understood as representing urgency; the drift rate as evidence; and the starting bias as prior expectation. The latter pair influenced this thesis. It is an open question how quantities such as prior expectation would be represented in any competitor models.

Is there a way to distinguish between various decision models in humans? Recent TMS studies have had success linking MEPs to modelled quantities (Bestmann et al., 2008; Klein-Flugge and Bestmann, 2012; Hadar et al., 2016) and raised the prospect that TMS could be used to distinguish between models (Hadar et al., 2016). The idea that ongoing decisions are present in motor cortex (Cisek, 2007) suggests that TMS should be able to access decision computations. On the other hand, in Chapter 1, I did not find any evidence to support the proposed hypothesis and an analysis attempting to link modelled bias to early corticospinal excitability was unsuccessful. I speculated one reason for these null results could be MEP variability (Kiers et al., 1993; Schmidt et al., 2009). Furthermore, TMS produces an aggregate population signal which makes it unsuitable for testing models based on single-neuron responses (Latimer et al., 2015). Any studies that were to attempt to differentiate between models using MEPs would need many more trials than is conventional.

5.3. Alternative theories of expectation in motor cortex

This thesis examined some aspects of the role of prior expectation in motor processes. I stayed within a framework influenced by past neurophysiological recordings (Basso and Wurtz, 1998; Platt and Glimcher, 1999; Bastian et al., 2003) and the models used to understand them (Ratcliff, 1978; Gold and Shadlen, 2007). However, there are other frameworks in which to understand prediction. As detailed in Chapter 1, the drift-diffusion model cannot explain the diverse nature of expectation-related signals in sensory areas, but predictive coding can (Summerfield and de Lange, 2014).

To date most work on predictive coding has been in the sensory domain (e.g. Summerfield et al., 2006; Summerfield and Koechlin, 2008). Can this theory be extended to explain the role of prediction in the motor domain?

Active inference, a framework that extends predictive coding (Friston et al., 2011), includes a theory of motor function.

According to active inference, motor planning should be understood not as generating a motor command but as the process of generating a sensory (specifically, a proprioceptive) prediction. The movement is then the process of fulfilling this prediction. Movement initiation involves increasing the uncertainty on current sensory information to make the predicted state more certain than the current one; this results in movement in order to attain the predicted state. There is an obvious analogy to the predictive coding theory of sensory cortex, where predictions are made about sensory inputs and prediction errors are conveyed back.

According to a mapping of predictive coding onto the laminar architecture of the cortex, this explains why the motor cortex is agranular (has rudimentary or no Layer IV; Shipp et al., 2013). It is theorised that this layer is concerned with conveying prediction errors and thus is not needed in motor cortex, where predictions are fulfilled by movement.

According to active inference, rather than the motor system planning a cost function (as in optimal control theories referred to in this thesis; Todorov and Jordan, 2002), a prior belief is generated about the desired trajectory (Friston et al., 2011; Adams et al., 2013).

The theory of active inference in the motor domain and the mapping onto laminar architecture are clearly exciting steps towards a deeper understanding of expectation in motor areas. Whether the hypothesis can be tested in humans remains to be seen. Recently, fMRI paradigms have started to arise for testing the predictions of predictive coding in perceptual areas (Summerfield and Koechlin, 2008; Kok and de Lange, 2014), so perhaps motor paradigms will follow. It has been proposed that

beta-band oscillations over motor cortex reflect uncertainty computations (Tan et al., 2016), and therefore track active inference processes (Palmer et al., 2016). So beta oscillations might be a way to test active inference predictions about the role of expectation in motor control.

5.4. Appraisal of methods used in thesis

In this section I discuss the benefits and drawbacks of some of the experimental methods used in this thesis. Further discussions are located elsewhere: in Chapter 1, I consider whether TMS was a suitable method for assessing decisions in motor cortex. An alternative to the use of single-handed leftward and rightward movements (as in Chapters 3 and 4) is discussed in the context of future work in Section 5.5.

5.4.1. Paradigms for delivering probabilistic information

In the experiments of this thesis, I used simple, explicit ‘prior cues’ to induce a probabilistic expectation in subjects about the direction of an upcoming movement. As discussed in Chapter 1, a variety of paradigms have been proposed to access ‘expectation’. The method used in this thesis has the advantage that it is clear on any trial what the participant’s belief is, and the explicit nature of the cues allows an interleaved design, because expectations do not need to be acquired slowly by learning. This was particularly useful in the TMS experiments of Chapter 1, where there was concern that temporal trends within MEPs (Schmidt et al., 2009) could compromise a blocked design. Despite being interleaved, the tasks used produced robust reaction time effects.

The disadvantage of this task is that it only allows the tracking of very simple types of probabilistic information. Recently-developed tasks are more complex (Galea et al., 2012; Marshall et al., 2016), and allow inferences about subjects’ beliefs about multiple kinds of information e.g.

both beliefs about individual trials and the wider context. These tasks, however, require model-based analyses to infer the quantities subjects are tracking.

5.4.2. Measures of variability

There has been no standard metric of movement variability in past experiments. Past measurements of variability have included standard deviation in maximal acceleration (Galea et al., 2013), standard deviation of the deviation of movement direction from a straight line at the moment of peak speed on that trial (Wijdenes et al., 2016), standard deviation of trial-to-trial change in overall movement reach angle (Pekny et al., 2015), and analyses of endpoint distributions as described by 95% confidence ellipses (van Beers et al., 2004). In short, it has been common to focus either on variability at movement endpoint, or variability at a single point during the trial.

To my knowledge, the analyses presented in this thesis are the first to plot variability continuously over the movement. The motivation for this was that the hypotheses concerned factors affected motor planning, and so any variability effects were expected to be strongest at the start of the movement. Indeed, this was borne out in the results. I believe the continuous metric is a strength of the analyses as it gives a more complete picture of variability changes over the movement than previous studies.

One limitation of this method is it is not suitable for analysis by conventional statistical methods. Instead, a variant of random field theory (Worsley et al., 1996) optimised for one-dimensional data (Pataky, 2012) was used to perform t-tests and ANOVAs, although it was necessary to collapse the data into a single variability metric for some analyses.

5.5. Future work

What should follow from the experiments presented in this thesis? In this thesis, I drew on recently proposed ideas that suggest that a limited resource model has implications for motor variability because it places a constraint on the planning resources allocated to each motor plan (Wijdenes et al., 2016). To me, an interesting aspect of the affordance competition hypothesis (Cisek, 2007) is the tension between an optimal strategy and the biological implementation of this. Planning multiple movements is argued to be necessary for dynamic behaviour in complex environments (Cisek and Kalaska, 2010). Planning many movements is clearly advantageous; a shared resource framework provides the balance to this by limiting the number of movements that can be planned well (i.e. with low execution variability). Extending and testing the idea of a limited resource framework in motor areas has the potential to offer new insights into motor variability, which has always been an area of debate (Harris and Wolpert, 1998; van Beers et al., 2004; Churchland et al., 2006)

In Chapter 3 I proposed and tested a hypothesis about how expectation would interact with a limited resource model. Specifically, I proposed that prior expectation leads to uneven distribution of resources across two motor plans, making the more expected movement less variable in execution.

Testing this idea, I found the expected effect in leftward movements but not rightward movements. I proposed this was due to an interaction with the different control policies for leftward and rightward movement and in Chapter 4 tested the idea that a stringent control policy is needed to see the effect. The experiment in Chapter 4 failed to support this hypothesis.

I discussed in Chapter 4 whether there were confounding factors in the paradigm used to test this idea. One option for future work would be to

attempt to design an improved paradigm to test the hypothesis and conduct another experiment. However, a continuing complicating factor in these experiments is the comparison between two different movements (right-handed movements to a target 45° left or 45° right) involving non-symmetrical patterns of muscle recruitment which are known to produce different early variabilities and endpoint distributions (van Beers et al., 2004).

Thus the next experiment I propose to pursue would be to repeat the experiment in Chapter 3 again, this time using two-handed movements (i.e. handle of robotic manipulandum is grasped with both hands, rather than right-hand only). This would have the effect of making the 45° leftward and rightward movements much more similar to one another, as they would both involve use of both hands. Handedness means that the two movements are not completely equivalent, but it seems likely this difference is less than the difference when both movements are completed using the right hand, and one solution to this would be to recruit a group with equal proportions left and right-handers, for separate analysis if necessary.

One possibility is that the explanation proposed for the results of Chapter 3 – that an effect of expectation on motor variability in leftward but not rightward movements was due to differing control policies between these movements – was not due to differences in control policy. This would explain why the experiment in Chapter 4 failed to find any evidence in support of this idea. What then is the reason for the result in Chapter 3? The result could be caused by switching. Under this explanation, subjects plan a (perhaps lower-effort, single-joint) rightward movement and plan the leftward movement only when it is cued. Thus increased variability in lower expectation conditions is as a result of lower expectation to switch plans, not lower expectation at the planning stage.

The idea that a single movement is planned contradicts a wide literature on multiple motor planning (Cisek, 2007; Chapman et al., 2010), detailed in Chapter 1. It is also unable to explain the results of a previous experiment on motor variability in choice vs forced-choice movements (Wijdenes et al., 2016). On the other hand, it explains the null result of Chapter 4, because it does not predict that inducing a lax control policy (as in this experiment) would change the modulation of variability.

How could the idea that the result of Chapter 3 is due to plan-switching be tested? The strength of a two-handed experiment is that it would also test the idea that the results observed are due to plan-switching. If switching were occurring, I would expect a two-handed version of the experiment to either show the same result as Chapter 4, with an effect of variability in one direction but not the other (because one movement is still planned in preference to the other, with later switching) or a complete null result (because with two equivalent movements, neither one is planned until advance, and planning is withheld until the movement is fully specified).

Assuming a two-handed version of the experiment replicated the finding in Chapter 3, I would next design experiments to bring the idea I have advanced - that limited resources for motor plans can be shared unevenly - away from the specific idea of prior expectation and into a more general framework. There are lots of factors that influence the relative importance of one motor plan over another and might similarly be hypothesised to weight motor plans unevenly, leading to a variability difference between the two movements. The investigation of these could be amenable to the analyses developed (specifically, continuous measures of variability along the movement) in Chapters 3 and 4.

One experiment in this vein would be to test whether, when multiple movements are planned in a value-based choice, movements towards the

lower-value item are more variable. Value encoding is already known to be relative, and the proposed encoding is via divisive normalisation (Louie et al., 2011; Pastor-Bernier and Cisek, 2011). Showing a similar effect of value on variability to that for prior expectation would be the first step towards extending these experiments.

If these experiments were successful, a next step could be to attempt to prove that divisive normalisation is the underlying mechanism, as proposed. This would be a challenge in humans as divisive normalisation is a cellular mechanism, which, to date, has been studied in animal models (Olsen et al., 2010; Louie et al., 2011; Carandini and Heeger, 2012). One alternative to animal work is to show results match with the pattern predicted by modelling a divisive normalisation processing. This has been done successfully for visual working memory errors by modelling spike patterns (Bays, 2014, 2015), and there are also preliminary results using whole-population models of divisive normalisation to make inferences about motor adaptation (Kato et al., 2017). Thus modelling approaches could be one option for linking more definitively to divisive normalisation.

An alternative approach to attempting to link to divisive normalisation mechanisms could be to use neuroimaging in humans. M/EEG would be more suitable than fMRI because temporal resolution would be more critical than spatial. A possible experiment would be to measure lateralised population activity, and, in particular, peak strength in activity, during a task where multiple movements have to be planned with different likelihoods of being executed, and show that (1) the activity associated with a particular movement is larger in amplitude when that movement is more expected; (2) the overall summed population activity is relatively constant; as predicted by a divisive normalisation hypothesis.

Acknowledgements

My primary supervisor, Sven Bestmann, has been unfailingly patient, supportive and fun to work with during a steep learning curve. The transition from medical school was an abrupt one; Sven deserves credit for maintaining an impressive poker face whilst I presented some truly wacky ideas to him and still allowing me to pursue my own areas of interest, which is what made this PhD fun. He was calm and encouraging through many set-backs. Thanks also to Steve Kennerley, my subsidiary supervisor, who has also been immensely supportive and instrumental in keeping the project on track and on time.

The MBPhD programme facilitated this PhD and were generous in giving me total freedom to pursue a project of scientific interest. Thanks to all involved in enabling this. In particular, Gordon Stewart made it possible for me to apply for the programme and encouraged me to do so. Daniel Marks has been a huge source of support in both the clinical and research aspects of the programme and was a reassuring voice in times of crisis.

The rotating members of the Bestmann lab group and the entire fourth floor office have made the Sobell Department a fun and friendly place to work. Every person in the office has been generous with their time scientifically, wonderful to spend time with non-scientifically, and, by and large, obsessed with running. All of this has been to my benefit, even the running.

My family have been encouraging throughout. This PhD prolongs my university years to last a decade; they have been kind in (mostly) not mentioning this. Last and not least, Jan Freyberg has acted as proof reader, sounding board, statistics obsessive, counsellor, MATLAB helpdesk, critic, cheerer-upper and much more: thank you.

References

- Adams RA, Shipp S, Friston KJ (2013) Predictions not commands: active inference in the motor system. *Brain Struct Funct* 218:611–643.
- Ahdab R, Ayache SS, Brugières P, Farhat WH, Lefaucheur J-P (2016) The Hand Motor Hotspot is not Always Located in the Hand Knob: A Neuronavigated Transcranial Magnetic Stimulation Study. *Brain Topogr* 29:590–597.
- Alexander GE, Crutcher MD (1990) Preparation for movement: neural representations of intended direction in three motor areas of the monkey. *J Neurophysiol* 64:133–150.
- Amassian VE, Eberle L, Maccabee PJ, Cracco RQ (1992) Modelling magnetic coil excitation of human cerebral cortex with a peripheral nerve immersed in a brain-shaped volume conductor: the significance of fiber bending in excitation. *Electroencephalogr Clin Neurophysiol Potentials Sect* 85:291–301.
- Andersen B, Rösler KM, Lauritzen M (1999) Nonspecific facilitation of responses to transcranial magnetic stimulation. *Muscle Nerve* 22:857–863.
- Bar M (2004) Visual objects in context. *Nat Rev Neurosci* 5:617–629.
- Barker AT, Jalinous R, Freeston IL (1985) Non-invasive magnetic stimulation of human motor cortex. *The Lancet* 325:1106–1107.
- Basso MA, Wurtz RH (1997) Modulation of neuronal activity by target uncertainty. *Nature* 389:66–69.
- Basso MA, Wurtz RH (1998) Modulation of Neuronal Activity in Superior Colliculus by Changes in Target Probability. *J Neurosci* 18:7519–7534.
- Bastian A, Schöner G, Riehle A (2003) Preshaping and continuous evolution of motor cortical representations during movement preparation. *Eur J Neurosci* 18:2047–2058.
- Bays PM (2014) Noise in Neural Populations Accounts for Errors in Working Memory. *J Neurosci* 34:3632–3645.
- Bays PM (2015) Spikes not slots: noise in neural populations limits working memory. *Trends Cogn Sci* 19:431–438.

- Bays PM, Husain M (2008) Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science* 321:851–854.
- Beck JM, Ma WJ, Pitkow X, Latham PE, Pouget A (2012) Not Noisy, Just Wrong: The Role of Suboptimal Inference in Behavioral Variability. *Neuron* 74:30–39.
- Bekinschtein TA, Dehaene S, Rohaut B, Tadel F, Cohen L, Naccache L (2009) Neural signature of the conscious processing of auditory regularities. *Proc Natl Acad Sci* 106:1672–1677.
- Benjamini Y, Hochberg Y (1995) Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *J R Stat Soc Ser B Methodol* 57:289–300.
- Berret B, Chiovetto E, Nori F, Pozzo T (2011) Manifold reaching paradigm: how do we handle target redundancy? *J Neurophysiol* 106:2086–2102.
- Bestmann S, Duque J (2015) Transcranial Magnetic Stimulation: Decomposing the Processes Underlying Action Preparation. *The Neuroscientist*:1073858415592594.
- Bestmann S, Harrison LM, Blankenburg F, Mars RB, Haggard P, Friston KJ, Rothwell JC (2008) Influence of Uncertainty and Surprise on Human Corticospinal Excitability during Preparation for Action. *Curr Biol* 18:775–780.
- Bhushan N, Shadmehr R (1999) Evidence for a forward dynamics model in human adaptive motor control. *Adv Neural Inf Process Syst*:3–9.
- Bock O, Eversheim U (2000) The mechanisms of movement preparation: a precuing study. *Behav Brain Res* 108:85–90.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433–436.
- Brasil-Neto JP, Pascual-Leone A, Valls-Solé J, Cammarota A, Cohen LG, Hallett M (1993) Postexercise depression of motor evoked potentials: a measure of central nervous system fatigue. *Exp Brain Res* 93:181–184.
- Brenner E, Smeets JBJ (2003) Fast corrections of movements with a computer mouse. *Spat Vis* 16:365–376.
- Britten KH, Heuer HW (1999) Spatial Summation in the Receptive Fields of MT Neurons. *J Neurosci* 19:5074–5084.

- Britten KH, Shadlen MN, Newsome WT, Movshon JA (1992) The analysis of visual motion: a comparison of neuronal and psychophysical performance. *J Neurosci* 12:4745–4765.
- Brody CD, Hanks TD (2016) Neural underpinnings of the evidence accumulator. *Curr Opin Neurobiol* 37:149–157.
- Brown SD, Heathcote A (2008) The simplest complete model of choice response time: Linear ballistic accumulation. *Cognit Psychol* 57:153–178.
- Cai X, Padoa-Schioppa C (2014) Contributions of Orbitofrontal and Lateral Prefrontal Cortices to Economic Choice and the Good-to-Action Transformation. *Neuron* 81:1140–1151.
- Carandini M, Heeger DJ (2012) Normalization as a canonical neural computation. *Nat Rev Neurosci* 13:51–62.
- Carpenter RHS (1999) A neural mechanism that randomises behaviour. *J Conscious Stud* 6:13–22.
- Carpenter RHS, Reddi BAJ, Anderson AJ (2009) A simple two-stage model predicts response time distributions. *J Physiol* 587:4051–4062.
- Carpenter RHS, Williams MLL (1995) Neural computation of log likelihood in control of saccadic eye movements. *Nature* 377:59–62.
- Cerri G (2003) Facilitation From Ventral Premotor Cortex of Primary Motor Cortex Outputs to Macaque Hand Muscles. *J Neurophysiol* 90:832–842.
- Chaisanguanthum KS, Shen HH, Sabes PN (2014) Motor Variability Arises from a Slow Random Walk in Neural State. *J Neurosci* 34:12071–12080.
- Chapman CS, Gallivan JP, Wood DK, Milne JL, Culham JC, Goodale MA (2010) Reaching for the unknown: Multiple target encoding and real-time decision-making in a rapid reach task. *Cognition* 116:168–176.
- Churchland AK, Kiani R (2016) Three challenges for connecting model to mechanism in decision-making. *Curr Opin Behav Sci* 11:74–80.
- Churchland AK, Kiani R, Shadlen MN (2008) Decision-making with multiple alternatives. *Nat Neurosci* 11:693–702.

- Churchland MM et al. (2010) Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nat Neurosci* 13:369–378.
- Churchland MM, Afshar A, Shenoy KV (2006) A central source of movement variability. *Neuron* 52:1085–1096.
- Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, Shenoy KV (2012) Neural population dynamics during reaching. *Nature* 487:51–56.
- Churchland MM, Shenoy KV (2007) Delay of Movement Caused by Disruption of Cortical Preparatory Activity. *J Neurophysiol* 97:348–359.
- Cisek P (1999) Beyond the Computer Metaphor: Behaviour as Interaction. *J Conscious Stud* 6:11–12.
- Cisek P (2005) Neural representations of motor plans, desired trajectories, and controlled objects. *Cogn Process* 6:15–24.
- Cisek P (2007) Cortical mechanisms of action selection: the affordance competition hypothesis. *Philos Trans R Soc Lond B Biol Sci* 362:1585–1599.
- Cisek P, Kalaska JF (2005) Neural Correlates of Reaching Decisions in Dorsal Premotor Cortex: Specification of Multiple Direction Choices and Final Selection of Action. *Neuron* 45:801–814.
- Cisek P, Kalaska JF (2010) Neural Mechanisms for Interacting with a World Full of Action Choices. *Annu Rev Neurosci* 33:269–298.
- Cisek P, Pastor-Bernier A (2014) On the challenges and mechanisms of embodied decisions. *Philos Trans R Soc B Biol Sci* 369:20130479–20130479.
- Cisek P, Puskas GA, El-Murr S (2009) Decisions in Changing Conditions: The Urgency-Gating Model. *J Neurosci* 29:11560–11571.
- Civardi C, Cantello R, Asselman P, Rothwell JC (2001) Transcranial Magnetic Stimulation Can Be Used to Test Connections to Primary Motor Areas from Frontal and Medial Cortex in Humans. *NeuroImage* 14:1444–1453.
- Cluff T, Scott SH (2013) Rapid Feedback Responses Correlate with Reach Adaptation and Properties of Novel Upper Limb Loads. *J Neurosci* 33:15903–15914.

- Cos I, Duque J, Cisek P (2014) Rapid prediction of biomechanical costs during action decisions. *J Neurophysiol* 112:1256–1266.
- Crammond DJ, Kalaska JF (2000) Prior Information in Motor and Premotor Cortex: Activity During the Delay Period and Effect on Pre-Movement Activity. *J Neurophysiol* 84:986–1005.
- Crevecoeur F, Kurtzer I, Bourke T, Scott SH (2013) Feedback responses rapidly scale with the urgency to correct for external perturbations. *J Neurophysiol* 110:1323–1332.
- Crevecoeur F, Kurtzer I, Scott SH (2012) Fast corrective responses are evoked by perturbations approaching the natural variability of posture and movement tasks. *J Neurophysiol* 107:2821–2832.
- Davranche K, Tandonnet C, Burle B, Meynier C, Vidal F, Hasbroucq T (2007) The dual nature of time preparation: neural activation and suppression revealed by transcranial magnetic stimulation of the motor cortex. *Eur J Neurosci* 25:3766–3774.
- Daw ND, O'Doherty JP, Dayan P, Seymour B, Dolan RJ (2006) Cortical substrates for exploratory decisions in humans. *Nature* 441:876–879.
- Day BL, Brown P (2001) Evidence for subcortical involvement in the visual control of human reaching. *Brain* 124:1832–1840.
- Day BL, Lyon IN (2000) Voluntary modification of automatic arm movements evoked by motion of a visual target. *Exp Brain Res* 130:159–168.
- de Lange FP, Rahnev DA, Donner TH, Lau H (2013) Prestimulus Oscillatory Activity over Motor Cortex Reflects Perceptual Expectations. *J Neurosci* 33:1400–1410.
- de Xivry J-JO (2013) Trial-to-Trial Reoptimization of Motor Behavior Due to Changes in Task Demands Is Limited. *PLoS ONE* 8:e66013.
- Desmurget M, Epstein CM, Turner RS, Prablanc C, Alexander GE, Grafton ST (1999) Role of the posterior parietal cortex in updating reaching movements to a visual target. *Nat Neurosci* 2:563–567.
- Di Lazzaro V, Oliviero A, Mazzone P, Insola A, Pilato F, Saturno E, Accurso A, Tonali P, Rothwell J (2001) Comparison of descending volleys evoked by monophasic and biphasic magnetic stimulation of the motor cortex in conscious humans. *Exp Brain Res* 141:121–127.

- Di Lazzaro V, Oliviero A, Profice P, Saturno E, Pilato F, Insola A, Mazzone P, Tonali P, Rothwell JC (1998a) Comparison of descending volleys evoked by transcranial magnetic and electric stimulation in conscious humans. *Electroencephalogr Clin Neurophysiol Mot Control* 109:397–401.
- Di Lazzaro V, Restuccia D, Oliviero A, Profice P, Ferrara L, Insola A, Mazzone P, Tonali P, Rothwell JC (1998b) Effects of voluntary contraction on descending volleys evoked by transcranial stimulation in conscious humans. *J Physiol* 508:625–633.
- Di Lazzaro V, Restuccia D, Oliviero A, Profice P, Ferrara L, Insola A, Mazzone P, Tonali P, Rothwell JC (1998c) Magnetic transcranial stimulation at intensities below active motor threshold activates intracortical inhibitory circuits. *Exp Brain Res* 119:265–268.
- Di Lazzaro V, Rothwell JC (2014) Corticospinal activity evoked and modulated by non-invasive stimulation of the intact human motor cortex. *J Physiol* 592:4115–4128.
- Diedrichsen J, Dowling N (2009) Bimanual coordination as task-dependent linear control policies. *Hum Mov Sci* 28:334–347.
- Diedrichsen J, Kornysheva K (2015) Motor skill learning between selection and execution. *Trends Cogn Sci* 19:227–233.
- Dimitriou M, Wolpert DM, Franklin DW (2013) The Temporal Evolution of Feedback Gains Rapidly Update to Task Demands. *J Neurosci* 33:10898–10909.
- Domenech P, Dreher J-C (2010) Decision Threshold Modulation in the Human Brain. *J Neurosci* 30:14305–14317.
- Donders FC (1969) On the speed of mental processes. *Acta Psychol (Amst)* 30:412–431.
- Donner TH, Siegel M, Fries P, Engel AK (2009) Buildup of Choice-Predictive Activity in Human Motor Cortex during Perceptual Decision Making. *Curr Biol* 19:1581–1585.
- Duque J, Greenhouse I, Labruna L, Ivry RB (2017) Physiological Markers of Motor Inhibition during Human Behavior. *Trends Neurosci* 40:219–236.
- Duque J, Ivry RB (2009) Role of Corticospinal Suppression during Motor Preparation. *Cereb Cortex* 19:2013–2024.

- Duque J, Labruna L, Cazaes C, Ivry RB (2014) Dissociating the influence of response selection and task anticipation on corticospinal suppression during response preparation. *Neuropsychologia* 65:287–296.
- Duque J, Labruna L, Verset S, Olivier E, Ivry RB (2012) Dissociating the Role of Prefrontal and Premotor Cortices in Controlling Inhibitory Mechanisms during Motor Preparation. *J Neurosci* 32:806–816.
- Duque J, Lew D, Mazzocchio R, Olivier E, Ivry RB (2010) Evidence for Two Concurrent Inhibitory Mechanisms during Response Preparation. *J Neurosci* 30:3793–3802.
- Duque J, Olivier E, Rushworth M (2013) Top-Down Inhibitory Control Exerted by the Medial Frontal Cortex during Action Selection under Conflict. *J Cogn Neurosci* 25:1634–1648.
- Durstewitz D, Deco G (2007) Computational significance of transient dynamics in cortical networks: Computational significance of transient dynamics. *Eur J Neurosci* 27:217–227.
- Enns JT, Lleras A (2008) What's next? New evidence for prediction in human vision. *Trends Cogn Sci* 12:327–333.
- Erlhagen W, Schöner G (2002) Dynamic field theory of movement preparation. *Psychol Rev* 109:545–572.
- Erlich JC, Brunton BW, Duan CA, Hanks TD, Brody CD (2015) Distinct effects of prefrontal and parietal cortex inactivations on an accumulation of evidence task in the rat. *eLife* 4:e05457.
- Esterman M, Yantis S (2010) Perceptual Expectation Evokes Category-Selective Cortical Activity. *Cereb Cortex* 20:1245–1253.
- Finkbeiner M, Song J-H, Nakayama K, Caramazza A (2008) Engaging the Motor System with Masked Orthographic Primes: A Kinematic Analysis. *Vis Cogn* 16:11–22.
- Flash T, Hogan N (1985) The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci* 5:1688–1703.
- Forstmann BU, Anwander A, Schäfer A, Neumann J, Brown S, Wagenmakers E-J, Bogacz R, Turner R (2010) Cortico-striatal connections predict control over speed and accuracy in perceptual decision making. *Proc Natl Acad Sci* 107:15916–15920.

- Franklin DW, Wolpert DM (2008) Specificity of Reflex Adaptation for Task-Relevant Variability. *J Neurosci* 28:14165–14175.
- Franklin DW, Wolpert DM (2011) Computational Mechanisms of Sensorimotor Control. *Neuron* 72:425–442.
- Friston K (2012) Prediction, perception and agency. *Int J Psychophysiol* 83:248–252.
- Friston K, Mattout J, Kilner J (2011) Action understanding and active inference. *Biol Cybern* 104:137–160.
- Galea JM, Bestmann S, Beigi M, Jahanshahi M, Rothwell JC (2012) Action Reprogramming in Parkinson's Disease: Response to Prediction Error Is Modulated by Levels of Dopamine. *J Neurosci* 32:542–550.
- Galea JM, Ruge D, Buijink A, Bestmann S, Rothwell JC (2013) Punishment-Induced Behavioral and Neurophysiological Variability Reveals Dopamine-Dependent Selection of Kinematic Movement Parameters. *J Neurosci* 33:3981–3988.
- Gallivan JP, Barton KS, Chapman CS, Wolpert DM, Randall Flanagan J (2015) Action plan co-optimization reveals the parallel encoding of competing reach movements. *Nat Commun* 6:7428.
- Gallivan JP, Bowman NAR, Chapman CS, Wolpert DM, Flanagan JR (2016a) The sequential encoding of competing action goals involves dynamic restructuring of motor plans in working memory. *J Neurophysiol* 115:3113–3122.
- Gallivan JP, Logan L, Wolpert DM, Flanagan JR (2016b) Parallel specification of competing sensorimotor control policies for alternative action options. *Nat Neurosci* 19:320–326.
- Gallivan JP, Stewart BM, Baugh LA, Wolpert DM, Flanagan JR (2017) Rapid Automatic Motor Encoding of Competing Reach Options. *Cell Rep* 18:1619–1626.
- Georgopoulos AP, Carpenter AF (2015) Coding of movements in the motor cortex. *Curr Opin Neurobiol* 33:34–39.
- Ghez C, Favilla M, Ghilardi MF, Gordon J, Bermejo R, Pullman S (1997) Discrete and continuous planning of hand movements and isometric force trajectories. *Exp Brain Res* 115:217–233.
- Glimcher PW (2011) Understanding dopamine and reinforcement learning: the dopamine reward prediction error hypothesis. *Proc Natl Acad Sci U S A* 108 Suppl 3:15647–15654.

- Goetz SM, Luber B, Lisanby SH, Peterchev AV (2014) A novel model incorporating two variability sources for describing motor evoked potentials. *Brain Stimulat* 7:541–552.
- Gold JI, Law C-T, Connolly P, Bennur S (2008) The Relative Influences of Priors and Sensory Evidence on an Oculomotor Decision Variable During Perceptual Learning. *J Neurophysiol* 100:2653–2668.
- Gold JI, Shadlen MN (2007) The Neural Basis of Decision Making. *Annu Rev Neurosci* 30:535–574.
- Goodale MA, Pelisson D, Prablanc C (1986) Large adjustments in visually guided reaching do not depend on vision of the hand or perception of target displacement. *Nature* 320:748–750.
- Green DM, Swets JA (1966) *Signal Detection Theory and Psychophysics*. New York: Wiley.
- Greenhouse I, Saks D, Hoang T, Ivry RB (2015a) Inhibition during response preparation is sensitive to response complexity. *J Neurophysiol* 113:2792–2800.
- Greenhouse I, Sias A, Labruna L, Ivry RB (2015b) Nonspecific Inhibition of the Motor System during Response Preparation. *J Neurosci* 35:10675–10684.
- Gregory RL (1970) *The intelligent eye*. London: Weidenfeld & Nicolson.
- Grillner S (2006) Biological Pattern Generation: The Cellular and Computational Logic of Networks in Motion. *Neuron* 52:751–766.
- Groppa S, Schlaak BH, Münchau A, Werner-Petroll N, Dünneberger J, Bäumer T, van Nuenen BFL, Siebner HR (2012) The human dorsal premotor cortex facilitates the excitability of ipsilateral primary motor cortex via a short latency cortico-cortical route. *Hum Brain Mapp* 33:419–430.
- Grunow A, Neuringer A (2002) Learning to vary and varying to learn. *Psychon Bull Rev* 9:250–258.
- Hadar AA, Rowe P, Di Costa S, Jones A, Yarrow K (2016) Motor-evoked potentials reveal a motor-cortical readout of evidence accumulation for sensorimotor decisions. *Psychophysiology* 53:1721–1731.
- Haith AM, Huberdeau DM, Krakauer JW (2015) Hedging Your Bets: Intermediate Movements as Optimal Behavior in the Context of an Incomplete Decision. *PLOS Comput Biol* 11:e1004171.

- Hamada M, Murase N, Hasan A, Balaratnam M, Rothwell JC (2013) The role of interneuron networks in driving human motor cortical plasticity. *Cereb Cortex* 23:1593–1605.
- Hanks TD, Kopec CD, Brunton BW, Duan CA, Erlich JC, Brody CD (2015) Distinct relationships of parietal and prefrontal cortices to evidence accumulation. *Nature* 520:220–223.
- Hanks TD, Mazurek ME, Kiani R, Hopp E, Shadlen MN (2011) Elapsed Decision Time Affects the Weighting of Prior Probability in a Perceptual Decision Task. *J Neurosci* 31:6339–6352.
- Hanks TD, Summerfield C (2017) Perceptual Decision Making in Rodents, Monkeys, and Humans. *Neuron* 93:15–31.
- Harris CM, Wolpert DM (1998) Signal-dependent noise determines motor planning. *Nature* 394:780–784.
- Hasbroucq T, Kaneko H, Akamatsu M, Possamaï C-A (1999) The time-course of preparatory spinal and cortico-spinal inhibition: an H-reflex and transcranial magnetic stimulation study in man. *Exp Brain Res* 124:33–41.
- Hasbroucq T, Kaneko H, Akamatsu M, Possamaï C-A (1997) Preparatory inhibition of cortico-spinal excitability: a transcranial magnetic stimulation study in man. *Cogn Brain Res* 5:185–192.
- Heeger D (1992) Half-squaring in responses of cat striate cells.
- Heitz RP (2014) The speed-accuracy tradeoff: history, physiology, methodology, and behavior. *Front Neurosci* 8:150.
- Helmholtz H von (1867) Concerning the perceptions in general. In: *Readings in the history of psychology*, pp 214–230 Century psychology series. East Norwalk, CT, US: Appleton-Century-Crofts.
- Hess CW, Mills KR, Murray NM (1987) Responses in small hand muscles from magnetic stimulation of the human brain. *J Physiol* 388:397–419.
- Hogan N, Sternad D (2009) Sensitivity of Smoothness Measures to Movement Duration, Amplitude, and Arrests. *J Mot Behav* 41:529–534.
- Huk AC, Katz LN, Yates JL (2017) The Role of the Lateral Intraparietal Area in (the Study of) Decision Making. *Annu Rev Neurosci* 40:349–372.

- Hunt LT, Kolling N, Soltani A, Woolrich MW, Rushworth MFS, Behrens TEJ (2012) Mechanisms underlying cortical activity during value-guided choice. *Nat Neurosci* 15:470–476.
- Izawa J, Criscimagna-Hemminger SE, Shadmehr R (2012) Cerebellar Contributions to Reach Adaptation and Learning Sensory Consequences of Action. *J Neurosci* 32:4230–4239.
- Izawa J, Shadmehr R (2011) Learning from Sensory and Reward Prediction Errors during Motor Adaptation. *PLoS Comput Biol* 7:e1002012.
- Jiang J, Summerfield C, Egnér T (2013) Attention Sharpens the Distinction between Expected and Unexpected Percepts in the Visual Brain. *J Neurosci* 33:18438–18447.
- Jones KE, Hamilton AF de C, Wolpert DM (2002) Sources of Signal-Dependent Noise During Isometric Force Production. *J Neurophysiol* 88:1533–1544.
- Kass RE, Raftery AE (1995) Bayes Factors. *J Am Stat Assoc* 90:773–795.
- Kato Y, Hayashi T, Nozaki D (2017) Divisively normalized integration of visual and proprioceptive information for motor adaptation. Available at: <http://www.abstractsonline.com/pp8/#!/4376/presentation/6810> [Accessed April 9, 2018].
- Katz LN, Yates JL, Pillow JW, Huk AC (2016) Dissociated functional significance of decision-related activity in the primate dorsal stream. *Nature* 535:285–288.
- Kaufman MT, Churchland MM, Ryu SI, Shenoy KV (2014) Cortical activity in the null space: permitting preparation without movement. *Nat Neurosci* 17:440–448.
- Kaufman MT, Churchland MM, Shenoy KV (2013) The roles of monkey M1 neuron classes in movement preparation and execution. *J Neurophysiol* 110:817–825.
- Kelly SP, O’Connell RG (2013) Internal and External Influences on the Rate of Sensory Evidence Accumulation in the Human Brain. *J Neurosci* 33:19434–19441.
- Kelly SP, O’Connell RG (2015) The neural processes underlying perceptual decision making in humans: Recent progress and future directions. *J Physiol-Paris* 109:27–37.

- Kersten D, Yuille A (2003) Bayesian models of object perception. *Curr Opin Neurobiol* 13:150–158.
- Kiers L, Cros D, Chiappa KH, Fang J (1993) Variability of motor potentials evoked by transcranial magnetic stimulation. *Electroencephalogr Clin Neurophysiol Potentials Sect* 89:415–423.
- Klein P-A, Olivier E, Duque J (2012) Influence of Reward on Corticospinal Excitability during Movement Preparation. *J Neurosci* 32:18124–18136.
- Klein P-A, Petitjean C, Olivier E, Duque J (2014) Top-down suppression of incompatible motor activations during response selection under conflict. *NeuroImage* 86:138–149.
- Kleiner M, Brainard D, Pelli D, Ingling A, Murray R, Broussard C (2007) What's new in Psychtoolbox-3. *Perception* 36:1.
- Klein-Flugge MC, Bestmann S (2012) Time-Dependent Changes in Human Corticospinal Excitability Reveal Value-Based Competition for Action during Decision Processing. *J Neurosci* 32:8373–8382.
- Klein-Flügge MC, Nobbs D, Pitcher JB, Bestmann S (2013) Variability of Human Corticospinal Excitability Tracks the State of Action Preparation. *J Neurosci* 33:5564–5572.
- Knill DC, Bondada A, Chhabra M (2011) Flexible, Task-Dependent Use of Sensory Feedback to Control Hand Movements. *J Neurosci* 31:1219–1237.
- Kok P, Brouwer GJ, Gerven MAJ van, Lange FP de (2013) Prior Expectations Bias Sensory Representations in Visual Cortex. *J Neurosci* 33:16275–16284.
- Kok P, de Lange FP (2014) Shape Perception Simultaneously Up- and Downregulates Neural Activity in the Primary Visual Cortex. *Curr Biol* 24:1531–1535.
- Kok P, Rahnev D, Jehee JFM, Lau HC, Lange D, P F (2012) Attention Reverses the Effect of Prediction in Silencing Sensory Signals. *Cereb Cortex* 22:2197–2206.
- Kording KP, Wolpert DM (2004) Bayesian integration in sensorimotor learning. *Nature* 427:244–247.
- Kujirai T, Caramia MD, Rothwell JC, Day BL, Thompson PD, Ferbert A, Wroe S, Asselman P, Marsden CD (1993) Corticocortical inhibition in human motor cortex. *J Physiol* 471:501–519.

- Labruna L, Lebon F, Duque J, Klein P-A, Cazares C, Ivry RB (2013) Generic Inhibition of the Selected Movement and Constrained Inhibition of Nonselected Movements during Response Preparation. *J Cogn Neurosci* 26:269–278.
- Lametti DR, Houle G, Ostry DJ (2007) Control of Movement Variability and the Regulation of Limb Impedance. *J Neurophysiol* 98:3516–3524.
- Latimer KW, Yates JL, Meister ML, Huk AC, Pillow JW (2015) Single-trial spike trains in parietal cortex reveal discrete steps during decision-making. *Science* 349:184–187.
- Latimer KW, Yates JL, Meister MLR, Huk AC, Pillow JW (2016) Response to Comment on “Single-trial spike trains in parietal cortex reveal discrete steps during decision-making.” *Science* 351:1406–1406.
- Leocani L, Cohen LG, Wassermann EM, Ikoma K, Hallett M (2000) Human corticospinal excitability evaluated with transcranial magnetic stimulation during different reaction time paradigms. *Brain* 123:1161–1173.
- Liepert J, Dettmers C, Terborg C, Weiller C (2001) Inhibition of ipsilateral motor cortex during phasic generation of low force. *Clin Neurophysiol* 112:114–121.
- Liu D, Todorov E (2007) Evidence for the Flexible Sensorimotor Strategies Predicted by Optimal Feedback Control. *J Neurosci* 27:9354–9368.
- Louie K, Grattan LE, Glimcher PW (2011) Reward Value-Based Gain Control: Divisive Normalization in Parietal Cortex. *J Neurosci* 31:10627–10639.
- Louie K, Khaw MW, Glimcher PW (2013) Normalization is a general neural mechanism for context-dependent decision making. *Proc Natl Acad Sci* 110:6139–6144.
- Luck SJ, Vogel EK (1997) The capacity of visual working memory for features and conjunctions. *Nature* 390:279–281.
- Marshall L, Mathys C, Ruge D, Berker AO de, Dayan P, Stephan KE, Bestmann S (2016) Pharmacological Fingerprints of Contextual Uncertainty. *PLOS Biol* 14:e1002575.
- McMillan S, Nougier V, Byblow WD (2004) Human corticospinal excitability during a precued reaction time paradigm. *Exp Brain Res* 156:80–87.

- McPeck RM, Han JH, Keller EL (2003) Competition Between Saccade Goals in the Superior Colliculus Produces Saccade Curvature. *J Neurophysiol* 89:2577–2590.
- Miall RC, Wolpert DM (1996) Forward Models for Physiological Motor Control. *Neural Netw* 9:1265–1279.
- Michelet T, Duncan GH, Cisek P (2010) Response Competition in the Primary Motor Cortex: Corticospinal Excitability Reflects Response Replacement During Simple Decisions. *J Neurophysiol* 104:119–127.
- Miller P, Katz DB (2010) Stochastic Transitions between Neural States in Taste Processing and Decision-Making. *J Neurosci* 30:2559–2570.
- Mochizuki H, Huang Y-Z, Rothwell JC (2004) Interhemispheric interaction between human dorsal premotor and contralateral primary motor cortex. *J Physiol* 561:331–338.
- Mulder MJ, Wagenmakers E-J, Ratcliff R, Boekel W, Forstmann BU (2012) Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *J Neurosci* 32:2335–2343.
- Mumford D (1992) On the computational architecture of the neocortex. *Biol Cybern* 66:241–251.
- Muraven M, Baumeister RF (2000) Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychol Bull* 126:247–259.
- Näätänen R, Gaillard AWK, Mäntysalo S (1978) Early selective-attention effect on evoked potential reinterpreted. *Acta Psychol (Amst)* 42:313–329.
- Näätänen R, Paavilainen P, Alho K, Reinikainen K, Sams M (1989) Do event-related potentials reveal the mechanism of the auditory sensory memory in the human brain? *Neurosci Lett* 98:217–221.
- Nakano T, Homae F, Watanabe H, Taga G (2008) Anticipatory Cortical Activation Precedes Auditory Events in Sleeping Infants. *PLOS ONE* 3:e3912.
- Nashed JY, Crevecoeur F, Scott SH (2012) Influence of the behavioral goal and environmental obstacles on rapid feedback responses. *J Neurophysiol* 108:999–1009.
- Neumann JV, Morgenstern O (1944) *Theory of Games and Economic Behavior*. Princeton University Press.

- Normann RA, Perlman I (1979) The effects of background illumination on the photoresponses of red and green cones. *J Physiol* 286:491–507.
- Olsen SR, Bhandawat V, Wilson RI (2010) Divisive Normalization in Olfactory Population Codes. *Neuron* 66:287–299.
- Ölveczky BP, Andalman AS, Fee MS (2005) Vocal Experimentation in the Juvenile Songbird Requires a Basal Ganglia Circuit. *PLoS Biol* 3:e153.
- Ölveczky BP, Otchy TM, Goldberg JH, Aronov D, Fee MS (2011) Changes in the neural control of a complex motor sequence during learning. *J Neurophysiol* 106:386–397.
- Padoa-Schioppa C (2011) Neurobiology of Economic Choice: A Good-Based Model. *Annu Rev Neurosci* 34:333–359.
- Palmer C, Zapparoli L, Kilner JM (2016) A New Framework to Explain Sensorimotor Beta Oscillations. *Trends Cogn Sci* 20:321–323.
- Palmer J (1990) Attentional limits on the perception and memory of visual information. *J Exp Psychol Hum Percept Perform* 16:332–350.
- Pastor-Bernier A, Cisek P (2011) Neural Correlates of Biased Competition in Premotor Cortex. *J Neurosci* 31:7083–7088.
- Pataky TC (2012) One-dimensional statistical parametric mapping in Python. *Comput Methods Biomech Biomed Engin* 15:295–301.
- Patton HD, Amassian VE (1954) Single- and Multiple-Unit Analysis of Cortical Stage of Pyramidal Tract Activation. *J Neurophysiol* 17:345–363.
- Pekny SE, Izawa J, Shadmehr R (2015) Reward-Dependent Modulation of Movement Variability. *J Neurosci* 35:4015–4024.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat Vis* 10:437–442.
- Pesek-Cotton EF, Johnson JE, Newland MC (2011) Reinforcing behavioral variability: an analysis of dopamine-receptor subtypes and intermittent reinforcement. *Pharmacol Biochem Behav* 97:551–559.
- Peterchev AV, Goetz SM, Westin GG, Luber B, Lisanby SH (2013) Pulse width dependence of motor threshold and input–output curve characterized with controllable pulse parameter transcranial

magnetic stimulation. *Clin Neurophysiol Off J Int Fed Clin Neurophysiol* 124:1364–1372.

Piet AT, Erlich JC, Kopec CD, Brody CD (2017) Rat Prefrontal Cortex Inactivations during Decision Making Are Explained by Bistable Attractor Dynamics. *Neural Comput*:1–26.

Platt ML, Glimcher PW (1999) Neural correlates of decision variables in parietal cortex. *Nature* 400:233–238.

Prablanc C, Martin O (1992) Automatic control during hand reaching at undetected two-dimensional target displacements. *J Neurophysiol* 67:455–469.

Preuschhof C, Schubert T, Villringer A, Heekeren HR (2009) Prior Information Biases Stimulus Representations during Vibrotactile Decision Making. *J Cogn Neurosci* 22:875–887.

Pruszynski JA, Kurtzer I, Scott SH (2008) Rapid Motor Responses Are Appropriately Tuned to the Metrics of a Visuospatial Task. *J Neurophysiol* 100:224–238.

Pruszynski JA, Kurtzer I, Scott SH (2011) The long-latency reflex is composed of at least two functionally independent processes. *J Neurophysiol* 106:449–459.

Rahnev D, Lau H, Lange FP de (2011) Prior Expectation Modulates the Interaction between Sensory and Prefrontal Regions in the Human Brain. *J Neurosci* 31:10741–10748.

Rao RPN, Ballard DH (1999) Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nat Neurosci* 2:79–87.

Rao V, DeAngelis GC, Snyder LH (2012) Neural Correlates of Prior Expectations of Motion in the Lateral Intraparietal and Middle Temporal Areas. *J Neurosci* 32:10063–10074.

Ratcliff R (1978) A theory of memory retrieval. *Psychol Rev* 85:59–108.

Ratcliff R (2002) A diffusion model account of response time and accuracy in a brightness discrimination task: Fitting real data and failing to fit fake but plausible data. *Psychon Bull Rev* 9:278–291.

Ratcliff R, Childers R (2015) Individual Differences and Fitting Methods for the Two-Choice Diffusion Model of Decision Making. *Decis Wash DC* 2015:10.1037/dec0000030.

- Ratcliff R, McKoon G (2007) The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Comput* 20:873–922.
- Reddi B a. J, Asrress KN, Carpenter RHS (2003) Accuracy, Information, and Response Time in a Saccadic Decision Task. *J Neurophysiol* 90:3538–3546.
- Reddi BA, Carpenter RHS (2000) The influence of urgency on decision time. *Nat Neurosci* 3:827–830.
- Reis J, Swayne OB, Vandermeeren Y, Camus M, Dimyan MA, Harris-Love M, Perez MA, Ragert P, Rothwell JC, Cohen LG (2008) Contribution of transcranial magnetic stimulation to the understanding of cortical mechanisms involved in motor control: TMS and motor control. *J Physiol* 586:325–351.
- Riehle A, Requin J (1989) Monkey primary motor and premotor cortex: single-cell activity related to prior information about direction and extent of an intended movement. *J Neurophysiol* 61:534–549.
- Rizzo V, Siebner HR, Modugno N, Pesenti A, Münchau A, Gerschlagel W, Webb RM, Rothwell JC (2004) Shaping the excitability of human motor cortex with premotor rTMS. *J Physiol* 554:483–495.
- Roitman JD, Shadlen MN (2002) Response of Neurons in the Lateral Intraparietal Area During a Combined Visual Discrimination Reaction Time Task. *J Neurosci* 22:9475–9489.
- Rossini P, Zarola F, Stalberg E, Caramia M (1988) Pre-movement facilitation of motor-evoked potentials in man during transcranial stimulation of the central motor pathways. *Brain Res* 458:20–30.
- Rouder JN, Morey RD, Speckman PL, Province JM (2012) Default Bayes factors for ANOVA designs. *J Math Psychol* 56:356–374.
- Roux S, Mackay WA, Riehle A (2006) The pre-movement component of motor cortical local field potentials reflects the level of expectancy. *Behav Brain Res* 169:335–351.
- Salvador R, Silva S, Basser PJ, Miranda PC (2011) Determining which mechanisms lead to activation in the motor cortex: a modeling study of transcranial magnetic stimulation using realistic stimulus waveforms and sulcal geometry. *Clin Neurophysiol Off J Int Fed Clin Neurophysiol* 122:748–758.
- Schlack A, Albright TD (2007) Remembering Visual Motion: Neural Correlates of Associative Plasticity and Motion Recall in Cortical Area MT. *Neuron* 53:881–890.

- Schmidlin E, Brochier T, Maier MA, Kirkwood PA, Lemon RN (2008) Pronounced Reduction of Digit Motor Responses Evoked from Macaque Ventral Premotor Cortex after Reversible Inactivation of the Primary Motor Cortex Hand Area. *J Neurosci* 28:5772–5783.
- Schmidt S, Cichy RM, Kraft A, Brocke J, Irlbacher K, Brandt SA (2009) An initial transient-state and reliable measures of corticospinal excitability in TMS studies. *Clin Neurophysiol* 120:987–993.
- Scholz JP, Schöner G (1999) The uncontrolled manifold concept: identifying control variables for a functional task. *Exp Brain Res* 126:289–306.
- Selen LPJ, Shadlen MN, Wolpert DM (2012) Deliberation in the Motor System: Reflex Gains Track Evolving Evidence Leading to a Decision. *J Neurosci* 32:2276–2286.
- Shadlen MN, Kiani R, Newsome WT, Gold JI, Wolpert DM, Zylberberg A, Ditterich J, Lafuente V de, Yang T, Roitman J (2016) Comment on “Single-trial spike trains in parietal cortex reveal discrete steps during decision-making.” *Science* 351:1406–1406.
- Shadlen MN, Newsome WT (2001a) Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *J Neurophysiol* 86:1916–1936.
- Shadlen MN, Newsome WT (2001b) Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *J Neurophysiol* 86:1916–1936.
- Shenoy KV, Sahani M, Churchland MM (2013) Cortical Control of Arm Movements: A Dynamical Systems Perspective. *Annu Rev Neurosci* 36:337–359.
- Sherrington C (1910) *The Integrative Action of the Nervous System*. CUP Archive.
- Shimazu H, Maier MA, Cerri G, Kirkwood PA, Lemon RN (2004) Macaque ventral premotor cortex exerts powerful facilitation of motor cortex outputs to upper limb motoneurons. *J Neurosci* 24:1200–1211.
- Shipp S, Adams RA, Friston KJ (2013) Reflections on agranular architecture: predictive coding in the motor cortex. *Trends Neurosci* 36:706–716.
- Shulman GL, Ollinger JM, Akbudak E, Conturo TE, Snyder AZ, Petersen SE, Corbetta M (1999) Areas Involved in Encoding and Applying

- Directional Expectations to Moving Objects. *J Neurosci* 19:9480–9496.
- Simen P, Contreras D, Buck C, Hu P, Holmes P, Cohen J (2009) Reward rate optimization in two-alternative decision making: empirical tests of theoretical predictions. *J Exp Psychol Hum Percept Perform* 35:1865–1897.
- Sinclair C, Hammond GR (2009) Excitatory and inhibitory processes in primary motor cortex during the foreperiod of a warned reaction time task are unrelated to response expectancy. *Exp Brain Res* 194:103–113.
- Smith FW, Muckli L (2010) Nonstimulated early visual areas carry information about surrounding context. *Proc Natl Acad Sci* 107:20099–20103.
- Sohn YH, Jung HY, Kaelin-Lang A, Hallett M (2003) Excitability of the ipsilateral motor cortex during phasic voluntary hand movement. *Exp Brain Res* 148:176–185.
- Song J-H, Nakayama K (2008) Numeric comparison in a visually-guided manual reaching task. *Cognition* 106:994–1003.
- Song J-H, Nakayama K (2009) Hidden cognitive states revealed in choice reaching tasks. *Trends Cogn Sci* 13:360–366.
- Starr A, Caramia M, Zarola F, Rossini P (1988) Enhancement of motor cortical excitability in humans by non-invasive electrical stimulation appears prior to voluntary movement. *Electroencephalogr Clin Neurophysiol* 70:26–32.
- Stewart BM, Baugh LA, Gallivan JP, Flanagan JR (2013) Simultaneous encoding of the direction and orientation of potential targets during reach planning: evidence of multiple competing reach plans. *J Neurophysiol* 110:807–816.
- Stewart BM, Gallivan JP, Baugh LA, Flanagan JR (2014) Motor, not visual, encoding of potential reach targets. *Curr Biol* 24:R953–R954.
- Storey JD (2002) A direct approach to false discovery rates. *J R Stat Soc Ser B Stat Methodol* 64:479–498.
- Summerfield C, de Lange FP (2014) Expectation in perceptual decision making: neural and computational mechanisms. *Nat Rev Neurosci* 15:745–756.

- Summerfield C, Egner T (2009) Expectation (and attention) in visual cognition. *Trends Cogn Sci* 13:403–409.
- Summerfield C, Egner T, Greene M, Koechlin E, Mangels J, Hirsch J (2006) Predictive Codes for Forthcoming Perception in the Frontal Cortex. *Science* 314:1311–1314.
- Summerfield C, Koechlin E (2008) A Neural Representation of Prior Information during Perceptual Inference. *Neuron* 59:336–347.
- Takikawa Y, Kawagoe R, Itoh H, Nakahara H, Hikosaka O (2002) Modulation of saccadic eye movements by predicted reward outcome. *Exp Brain Res* 142:284–291.
- Tan H, Wade C, Brown P (2016) Post-Movement Beta Activity in Sensorimotor Cortex Indexes Confidence in the Estimations from Internal Models. *J Neurosci* 36:1516–1528.
- Tandonnet C, Garry MI, Summers JJ (2010) Cortical activation during temporal preparation assessed by transcranial magnetic stimulation. *Biol Psychol* 85:481–486.
- Tanji J, Evarts EV (1976) Anticipatory activity of motor cortex neurons in relation to direction of an intended movement. *J Neurophysiol* 39:1062–1068.
- Therrien AS, Wolpert DM, Bastian AJ (2016) Effective reinforcement learning following cerebellar damage requires a balance between exploration and motor noise. *Brain* 139:101–114.
- Thura D, Beauregard-Racine J, Fradet C-W, Cisek P (2012) Decision making by urgency gating: theory and experimental support. *J Neurophysiol* 108:2912–2930.
- Thura D, Cisek P (2014) Deliberation and Commitment in the Premotor and Primary Motor Cortex during Dynamic Decision Making. *Neuron* 81:1401–1416.
- Todorov E, Jordan MI (2002) Optimal feedback control as a theory of motor coordination. *Nat Neurosci* 5:1226–1235.
- Todorovic A, de Lange FP (2012) Repetition Suppression and Expectation Suppression Are Dissociable in Time in Early Auditory Evoked Fields. *J Neurosci* 32:13389–13395.
- Tokimura H, Ridding MC, Tokimura Y, Amassian VE, Rothwell JC (1996) Short latency facilitation between pairs of threshold magnetic

stimuli applied to human motor cortex. *Electroencephalogr Clin Neurophysiol Mot Control* 101:263–272.

Touge T, Taylor JL, Rothwell JC (1998) Reduced excitability of the cortico-spinal system during the warning period of a reaction time task. *Electroencephalogr Clin Neurophysiol Mot Control* 109:489–495.

Trommershäuser J, Gepshtein S, Maloney LT, Landy MS, Banks MS (2005) Optimal Compensation for Changes in Task-Relevant Movement Variability. *J Neurosci* 25:7169–7178.

Usher M, McClelland JL (2001) The time course of perceptual choice: The leaky, competing accumulator model. *Psychol Rev* 108:550–592.

van Beers RJ, Haggard P, Wolpert DM (2004) The Role of Execution Noise in Movement Variability. *J Neurophysiol* 91:1050–1063.

van Elswijk G, Kleine BU, Overeem S, Stegeman DF (2007) Expectancy induces dynamic modulation of corticospinal excitability. *J Cogn Neurosci* 19:121–131.

van Vugt M, Karlijn M, Simen P, Nystrom LE, Holmes P, Cohen JD (2012) EEG Oscillations Reveal Neural Correlates of Evidence Accumulation. *Front Neurosci* 6:106.

Wagenmakers E, Wetzels R, Borsboom D, van der Maas HLJ (2011) Why psychologists must change the way they analyze their data: The case of psi: Comment on Bem (2011). *J Pers Soc Psychol* 100:426–432.

Wagner MJ, Smith MA (2008) Shared Internal Models for Feedforward and Feedback Control. *J Neurosci* 28:10663–10673.

Wald A (1945) Sequential Tests of Statistical Hypotheses. *Ann Math Stat* 16:117–186.

Watson AB, Pelli DG (1983) Quest: A Bayesian adaptive psychometric method. *Percept Psychophys* 33:113–120.

Weiss AC, Weiller C, Liepert J (2003) Pre-movement motor excitability is reduced ipsilateral to low force pinch grips. *J Neural Transm* 110:201–208.

Wiecki TV, Sofer I, Frank MJ (2013) HDDM: Hierarchical Bayesian estimation of the Drift-Diffusion Model in Python. *Front Neuroinformatics* 7:14.

- Wijdenes LO, Brenner E, Smeets JBJ (2011) Fast and fine-tuned corrections when the target of a hand movement is displaced. *Exp Brain Res* 214:453–462.
- Wijdenes LO, Ivry RB, Bays PM (2016) Competition between movement plans increases motor variability: evidence of a shared resource for movement planning. *J Neurophysiol*:jn.00113.2016.
- Wise SP, Weinrich M, Mauritz K-H (1983) Motor aspects of cue-related neuronal activity in premotor cortex of the rhesus monkey. *Brain Res* 260:301–305.
- Wong AL, Goldsmith J, Krakauer JW (2016) A motor planning stage represents the shape of upcoming movement trajectories. *J Neurophysiol*:jn.01064.2015.
- Wong AL, Haith AM (2017) Motor planning flexibly optimizes performance under uncertainty about task goals. *Nat Commun* 8:14624.
- Wong AL, Haith AM, Krakauer JW (2015) Motor Planning. *The Neuroscientist* 21:385–398.
- Woolley SC, Rajan R, Joshua M, Doupe AJ (2014) Emergence of Context-Dependent Variability across a Basal Ganglia Network. *Neuron* 82:208–223.
- Worsley KJ, Marrett S, Neelin P, Vandal AC, Friston KJ, Evans AC (1996) A unified statistical approach for determining significant signals in images of cerebral activation. *Hum Brain Mapp* 4:58–73.
- Wu HG, Miyamoto YR, Castro LNG, Ölveczky BP, Smith MA (2014) Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nat Neurosci* 17:312–321.
- Wyart V, de Gardelle V, Scholl J, Summerfield C (2012a) Rhythmic Fluctuations in Evidence Accumulation during Decision Making in the Human Brain. *Neuron* 76:847–858.
- Wyart V, Nobre AC, Summerfield C (2012b) Dissociable prior influences of signal probability and relevance on visual contrast sensitivity. *Proc Natl Acad Sci U S A* 109:3593–3598.
- Yaron A, Hershenhoren I, Nelken I (2012) Sensitivity to Complex Statistical Regularities in Rat Auditory Cortex. *Neuron* 76:603–615.
- Yeo S-H, Franklin DW, Wolpert DM (2016) When Optimal Feedback Control Is Not Enough: Feedforward Strategies Are Required for

Optimal Control with Active Sensing. PLOS Comput Biol 12:e1005190.

Ziemann U, Tergau F, Wassermann EM, Wischer S, Hildebrandt J, Paulus W (1998) Demonstration of facilitatory I wave interaction in the human motor cortex by paired transcranial magnetic stimulation. J Physiol 511:181–190.

Zoccolan D, Cox DD, DiCarlo JJ (2005) Multiple Object Response Normalization in Monkey Inferotemporal Cortex. J Neurosci 25:8150–8164.