

Intentional inhibition of human action

Erman Misirlisoy

Institute of Cognitive Neuroscience
Division of Psychology and Language Sciences
University College London

PhD Thesis

2014

I, Erman Misirlisoy, confirm that the work presented in this thesis is my own. Any information derived from other sources is fully cited and referenced in the thesis.

Signed,

A handwritten signature in black ink, appearing to read 'Erman', written in a cursive style.

Erman Misirlisoy

Abstract

A crucial component of human behavioural flexibility is the capacity to inhibit actions at the last moment before action execution. This behavioural inhibition is often not an immediate reaction to external stimuli, but rather an endogenous ‘free’ decision. Knowledge about such ‘intentional inhibition’ is currently limited, with most research focused on stimulus-driven inhibition. This thesis will examine intentional inhibition, using several different experimental approaches. The behavioural experiments reported in the initial chapters found that intentional inhibition directly alters sensory processing during decision-making. In addition, there were unique effects of prior event sequences on subsequent decisions to either act or inhibit. Brain imaging methods using EEG and fMRI showed distinct neural mechanisms associated with intentional inhibition, which did not apply to rule-based inhibition. Work with Tourette syndrome patients indicated that the intentional inhibition of involuntary motor tics affects brain activity associated with voluntary actions. Furthermore, attentional manipulation strategies were shown to be highly effective in reducing tics, which may open up alternative behavioural treatment approaches for tic disorders. This thesis concludes by demonstrating that intentional inhibition is a bona fide cognitive function that can be studied using behavioural and neuroscientific methods. It also develops a cognitive model in which behavioural inhibition varies along a continuum from ‘instructed inhibition’ to ‘intentional inhibition’. This model may be useful as a guide for future work.

Acknowledgements

I am enormously grateful to my supervisor Patrick Haggard for his stimulating discussions and astute comments over the last three years. He has played a major role in all of the work within this thesis. I have improved considerably as a researcher and writer thanks to his input. I would also like to thank the ESRC for funding the work (Grant RES-062-23-2183).

Research projects with my collaborators have been a great source of enjoyment and development. First, I offer my thanks to Max-Philipp Stenner for working with me on experiment 3.4. For experiments 3.3 and 3.4, I am also grateful to Hun Choi and Ksenia Vinogradova who helped with data collection.

For experiment 5.2, I would like to thank Marcel Brass, Jelle Demanet, and Wouter De Baene for allowing me to work in the Ghent lab. Wouter also contributed significantly to programming the experiment and helping with data analysis.

For experiment 6.1, I would like to thank Alexander Munchau, Valerie Brandt, and Jennifer Tubing, who were hugely welcoming during the project in Hamburg. In particular, I greatly appreciate their efforts in recruiting Tourette syndrome patients for the study. I owe additional thanks to the patients themselves for their time and cooperation as participants in the study. For experiment 6.2 of the thesis, I would like to thank Giovanna Moretto, Christos Ganos, Elisa Filevich, Jan Feldheim, Ursula Kahl, Friedhelm Hummel, Alexander Munchau, and Patrick Haggard, for designing the experiment and collecting the data that I analysed and reported.

Further thanks go out to the people I have had the pleasure of sharing an office with during my PhD: Mathew Salvaris, Christina Fuentes, Esther Kuhn, Nura Sidarus, and Steve Di Costa. They created an environment that was both fun and productive.

I am profoundly indebted to Metin, Yesim, and Ertan Misirlisoy, whose unconditional love and guidance have been a constant inspiration. I would also like to thank Elisabeth Engl and her family for their support and encouragement. My friends, Tom, Nick, and Jon, have also provided invaluable companionship when it was most needed.

Final thanks go to the examiners of this thesis: Prof. Richard Ridderinkhof and Dr. Mark Edwards. The exam was challenging but also highly rewarding. Their incisive criticisms and suggestions improved the thesis enormously.

Contents

Chapter 1. Intentional inhibition: An introduction.....	10
1.1. Action inhibition in the real world.....	14
1.1.1. The importance of inhibition for everyday function.....	14
1.1.2. Intentional inhibition: a definition.....	14
1.1.3. Inhibition in a social context.....	14
1.1.4. Inhibition in law and morality.....	15
1.2. The scientific study of instructed inhibition.....	15
1.2.1. The stop signal paradigm.....	16
1.2.2. The go/nogo paradigm.....	17
1.2.3. The neural correlates of instructed inhibition.....	17
1.2.3.1. EEG.....	17
1.2.3.2. Functional magnetic resonance imaging (fMRI).....	18
1.2.3.3. Brain stimulation.....	18
1.3. The scientific study of intentional inhibition.....	19
1.3.1. Relevance and importance for scientific research.....	19
1.3.2. Similarities to existing concepts.....	20
1.3.2.1. Self-control.....	20
1.3.2.2. Delayed gratification/delay discounting.....	21
1.3.2.3. Decision-making.....	22
1.3.3. Challenges in intentional inhibition research.....	23
1.3.3.1. 'Free' choices.....	23
1.3.3.2. Prepotent action and pre-decision.....	23
1.3.3.3. A lack of behavioural output.....	24
1.3.4. Past research in intentional inhibition.....	25
1.3.4.1. Behavioural investigations.....	25
1.3.4.2. fMRI.....	26
1.3.4.3. EEG.....	27
1.4. Disorders of action and inhibition.....	27
1.4.1. Psychiatric disorders.....	28
1.4.1.1. Attention deficit hyperactivity disorder (ADHD).....	28
1.4.1.2. Obsessive-compulsive disorder (OCD).....	28
1.4.1.3. Drug abuse.....	31
1.4.2. Tourette syndrome.....	31
1.4.3. Neuropsychological disorders and brain lesions.....	32

1.4.3.1. <i>Anarchic hand syndrome</i>	32
1.5. <i>Knowledge gaps in the study of intentional inhibition</i>	33
1.6. <i>Aims of the current thesis</i>	34
Chapter 2. <i>Perceptual processing and intentional inhibition</i>	36
2.1. <i>Introduction</i>	36
2.2. <i>Experiment 2.1: Evidence accumulation for action and inhibition decisions</i>	37
2.2.1. <i>Introduction</i>	37
2.2.2. <i>Methods</i>	39
2.2.2.1. <i>Participants</i>	39
2.2.2.2. <i>Design</i>	39
2.2.2.3. <i>Procedure</i>	40
2.2.3. <i>Results</i>	42
2.2.4. <i>Discussion</i>	46
2.3. <i>Experiment 2.2: Manipulating perceptual outcomes of action and inhibition decisions</i> ..	49
2.3.1. <i>Introduction</i>	49
2.3.2. <i>Methods</i>	49
2.3.2.1. <i>Participants</i>	49
2.3.2.2. <i>Design and Procedure</i>	49
2.3.3. <i>Results</i>	51
2.3.4. <i>Discussion</i>	54
2.4. <i>Experiment 2.3: Inhibition and agency over sensory outcomes</i>	55
2.4.1. <i>Introduction</i>	55
2.4.2. <i>Methods</i>	55
2.4.2.1. <i>Participants</i>	55
2.4.2.2. <i>Design</i>	56
2.4.2.3. <i>Procedure</i>	56
2.4.3. <i>Results</i>	58
2.4.4. <i>Discussion</i>	61
2.5. <i>General discussion</i>	63
Chapter 3. <i>Effects of sequential dependence and the gambler’s fallacy</i>	66
3.1. <i>Introduction</i>	66
3.2. <i>Experiment 3.1: The gambler’s fallacy effect on inhibition decisions – an initial</i>	
<i>exploration</i>	69
3.2.1. <i>Introduction</i>	69
3.2.2. <i>Methods</i>	70
3.2.2.1. <i>Participants</i>	70

2.1.2. Design	70
2.1.3. Procedure.....	72
3.2.3. Results	73
3.2.3.1 Reaction times.....	73
3.2.3.2 Free choices.....	74
3.2.4. Discussion	75
3.3. Experiment 3.2: The gambler’s fallacy effect on inhibition decisions – reducing automatic activation.....	76
3.3.1. Introduction.....	76
3.3.2. Methods.....	76
3.3.2.1 Participants.....	76
3.3.2.2 Design and Procedure	77
3.3.3. Results	77
3.3.3.1 Reaction times.....	77
3.3.3.2 Free choices.....	78
3.3.3.3 Between-experiment comparisons of free choices	79
3.2.4. Discussion	79
3.4. Experiment 3.3: The gambler’s fallacy effect on inhibition decisions – controlling for action frequency.....	80
3.4.1. Introduction.....	80
3.4.2. Methods.....	80
3.4.2.1 Participants.....	80
3.4.2.2 Design and Procedure	81
3.4.3 Results	81
3.4.3.1 Reaction times.....	81
3.4.3.2 Free choices.....	82
3.4.3.3 Between-experiment comparisons of free choices	82
3.4.4. Discussion	83
3.5. Experiment 3.4: Priming sequential choices to act or inhibit.....	85
3.5.1. Introduction.....	85
3.5.2. Methods.....	85
3.5.2.1. Participants.....	85
3.5.2.2. Design.....	86
3.5.2.3. Procedure.....	87
3.5.3. Results	88
3.5.4. Discussion	90

3.6.	<i>General discussion</i>	92
Chapter 4.	A natural experiment: Kickers and goalkeepers in football penalty shoot-out	96
4.1.	<i>Introduction</i>	96
4.2.	<i>Methods</i>	98
4.3.	<i>Results</i>	99
4.4.	<i>Discussion</i>	106
Chapter 5.	The neural correlates of intentional inhibition	110
5.1.	<i>Introduction</i>	110
5.2.	<i>Experiment 5.1: Omissions in action sequences – EEG</i>	110
5.2.1.	<i>Introduction</i>	110
5.2.2.	<i>Methods</i>	112
5.2.2.1.	<i>Participants</i>	112
5.2.2.2.	<i>Design</i>	112
5.2.2.3.	<i>Procedure</i>	113
5.2.2.4.	<i>EEG</i>	113
5.2.3.	<i>Results</i>	114
5.2.4.	<i>Discussion</i>	118
5.3.	<i>Experiment 5.2. Omissions in action sequences – fMRI</i>	124
5.3.1.	<i>Introduction</i>	124
5.3.2.	<i>Methods</i>	126
5.3.2.1.	<i>Participants</i>	126
5.3.2.2.	<i>Design</i>	126
5.3.2.3.	<i>Procedure</i>	127
5.3.2.4.	<i>Scanning procedure</i>	129
5.3.2.5.	<i>Data preprocessing and analysis</i>	129
5.3.3.	<i>Results</i>	130
5.3.4.	<i>Discussion</i>	134
5.4.	<i>General discussion</i>	137
Chapter 6.	Tourette syndrome and tic inhibition	139
6.1.	<i>Introduction</i>	139
6.2.	<i>Experiment 6.1: The role of attention in tics</i>	140
6.2.1.	<i>Introduction</i>	140
6.2.2.	<i>Methods</i>	140
6.2.2.1.	<i>Participants</i>	140
6.2.2.2.	<i>Apparatus and materials</i>	141
6.2.2.3.	<i>Design and procedure</i>	141

6.2.2.4. Tic counting	144
6.2.3. Results	144
6.2.4. Discussion	147
6.3. Experiment 6.2: Intentional action and tic inhibition: an EEG and psychophysical study in Tourette syndrome	152
6.3.1. Introduction.....	152
6.3.2. Methods.....	152
6.3.2.1. Participants.....	152
6.3.2.2. Design and procedure.....	152
6.3.2.3. Measurements and EEG processing	153
6.3.3. Results	154
6.3.4. Discussion	161
6.4. General discussion	162
Chapter 7. General discussion.....	164
7.1. Intentional inhibition as a distinct construct	164
7.1.1. Intentional action vs intentional inhibition	164
7.1.2. Intentional inhibition vs instructed inhibition.....	165
7.2. Remaining problems for intentional inhibition	166
7.3. A continuum of instructed to intentional inhibition	167
7.4. Conclusions	171
References.....	173

List of figures

Figure 1.1. Key principles underlying the stop-signal task.....	16
Figure 2.1. Schematic illustration of varying perceptual evidence strength for action and inhibition decisions	38
Figure 2.2. Basic sequence of events in a typical trial (Experiment 2.1).....	41
Figure 2.3. Reaction times (ms) for actions and inhibitions in response to green and red instruction stimuli (Experiment 2.1)	42
Figure 2.4. Response times (ms) for decisions about the direction of moving dots (Experiment 2.1)	43
Figure 2.5. Drift rate values for actions and inhibitions across each motion coherence condition (Experiment 2.1).	45
Figure 2.6. Boundary values for actions and inhibitions across each motion coherence condition (Experiment 2.1).	46
Figure 2.7. Basic sequence of events in an example trial (Experiment 2.2).....	50
Figure 2.8. Response times (ms) for decisions about the direction of moving dots (Experiment 2.2).	51
Figure 2.9. Drift rate values for actions and inhibitions across each motion coherence condition (Experiment 2.2).	52
Figure 2.9.1 Boundary parameter values for actions and inhibitions across each motion coherence condition (Experiment 2.2)	53
Figure 2.9.2. Basic structure of trials in action and inhibition blocks (Experiment 2.3).	58
Figure 2.9.3. Proportion of ‘yes’ responses when asked whether a delay was present between the executed response (action or inhibition) and the subsequently appearing illumination stimulus (Experiment 2.3).....	59
Figure 2.9.4. d' values for detecting delays between response and subsequent illumination stimulus, and the difference in d' between actions and inhibitions (Experiment 2.3).....	60
Figure 2.9.5. Subjective ratings of agency towards illumination stimulus (left), and difference in agency ratings between actions and inhibition (right) (Experiment 2.3).	61
Figure 3.1. Schematic illustration of the AAC and HAA hypotheses.	69
Figure 3.2. Experimental paradigm showing different trial types (Experiment 3.1).	71
Figure 3.3. Mean reaction times to instructed action stimuli as a function of number of preceding action or inhibition responses (Experiment 3.1).	74

Figure 3.4. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions (Experiment 3.1).....	75
Figure 3.5. Mean reaction times to instructed action stimuli as a function of number of preceding consecutive action/inhibition responses (Experiment 3.2).....	78
Figure 3.6. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions (Experiment 3.2).....	79
Figure 3.7. Mean reaction times to instructed action stimuli as a function of number of preceding consecutive action/inhibition responses (Experiment 3.3).....	81
Figure 3.8. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions (Experiment 3.3).....	82
Figure 3.9. Basic stimulus sequence for a typical up-arrow trial with compatible priming (Experiment 3.4).	87
Figure 3.9.1. Reaction times to an action stimulus following 1, 2, 3, or 4, preceding consecutive action or inhibition targets (Experiment 3.4).....	88
Figure 3.9.2. Expectancy judgements for an action target stimulus in the next trial following 1, 2, 3, or 4 consecutive preceding action or inhibition targets (Experiment 3.4).....	89
Figure 3.9.3. Proportion of action responses on free choice trials following 1, 2, 3, or 4 action or inhibition responses (Experiment 3.4).....	90
Figure 4.1. Percentage of dives in the opposite direction to last dive direction following runs of 1, 2, or 3 repeated dives in the same direction.	101
Figure 4.2. Percentage of shots in the opposite direction to last dive direction following runs of 1, 2, or 3 repeated dives in the same direction.	101
Figure 4.3. Four consecutive shots by Portuguese kickers against the English goalkeeper during the Euro 2004 semifinal penalty shootout.	102
Figure 4.4. Percentage of goalkeeper dives in the opposite direction to the last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction.	103
Figure 4.5. Percentage of kicker shots in the opposite direction to the last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction.	103
Figure 4.6. Percentage of dives in the opposite direction to last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction.	106
Figure 4.7. Percentage of kicks in the opposite direction to last kick direction, following runs of 1, 2, or 3 repeated kicks in the same direction.	106
Figure 5.1. ERP data (n = 18) for standard, omission – 1, and omission + 1 actions at electrode C3 (Experiment 5.1).....	116

Figure 5.2. ERP data (n = 18) for omissions at electrode Cz (left) and across the scalp (right) (Experiment 5.1).	118
Figure 5.3. Basic stimulus sequence for voluntary-W' and rule-based conditions (Experiment 5.2).	128
Figure 5.4. M' and W' judgements relative to time of actual inhibition (Experiment 5.2)..	131
Figure 5.5. Brain activations for action > baseline contrast (Experiment 5.2).	132
Figure 5.6. Brain activations for rule-based and voluntary inhibition contrasts (Experiment 5.2).	133
Figure 6.1. Key features of the experimental design (Experiment 6.1).	143
Figure 6.2. Inter-rater reliability between rater counts used for analysis and independent rater naïve to all aspects of the experiment (Experiment 6.1).	145
Figure 6.3. Mean number of tics per trial for each task condition and baseline (Experiment 6.1).	146
Figure 6.4. Judgements of intention and action in TS patients and volunteers (Experiment 6.2)	155
Figure 6.5. Readiness potentials for volunteers and patients in M and W judgement conditions (Experiment 6.2).	156
Figure 6.6. Event-related spectral power (ERSP) for volunteers and patients in M and W judgement conditions, shown as frequency x time plots (Experiment 6.2).	157
Figure 6.7. Effects of voluntary tic suppression on readiness potentials in M and W judgement conditions (Experiment 6.2).	160
Figure 6.8. Effects of voluntary tic suppression on event-related spectral power (ERSP) in M and W judgement conditions, shown as frequency x time plots (Experiment 6.2).	160
Figure 7.1. Two dimensions of an instructed-intentional continuum for behavioural inhibition	168

List of tables

Table 1.1. Behavioural performance and brain activity during stop-signal/Go-NoGo tasks for various patient groups.	30
Table 3.1. Distribution of runs within each block.	72
Table 5.1. Rewards (i.e., appropriately timed actions, see text) and mean temporal intervals in rule-based and voluntary conditions	115
Table 6.1. Percentage of correctly recalled items in the finger and colour tasks during freely tic and tic inhibition conditions.....	145
Table 6.2. Temporal judgements (ms) and tic counts for TS patients in M and W judgement conditions during suppression and no suppression of tics.	159

Chapter 1. Intentional inhibition: An introduction

1.1. Action inhibition in the real world

1.1.1. The importance of inhibition for everyday function

In everyday life, we navigate and interact with the external world through action. These actions can directly change the environment. In many cases, we may plan an action but then decide that a predicted change in the environment will be inappropriate. In such situations, we are often able to inhibit the planned action, thus preventing its execution and any of its causal external effects. This capacity to stop action up until “the last moment” is a major part of human behavioural flexibility. Without inhibition, actions would be impulsive and reactive, with little room for correcting maladaptive decisions.

1.1.2. Intentional inhibition: a definition

Inhibition can arise from two possible processes. One can be thought of as ‘instructed inhibition’, in which we inhibit our actions in response to some external signal, and the other is ‘intentional inhibition’ in which we make an endogenous self-generated decision to halt our action (Filevich et al., 2012). Intentional inhibition is therefore the ‘free’ decision to inhibit an action we have planned, without an immediate external signal instructing us to do so. This is arguably the more common form of inhibition in everyday life. We rarely receive an external stimulus telling us to withhold action. Instead, we usually must make our own decision about whether or not to act. In sports such as tennis for example, we may see an approaching ball, which we prepare to hit back. We may realise at the last moment that there is a chance the opponent’s ball will land ‘out’. We could then decide to inhibit our action and let the ball pass us with the hope it lands beyond the court lines. During this sequence of events, there is no obvious or explicit external stop signal. The decision to inhibit is made endogenously based on ambiguous evidence in the environment.

1.1.3. Inhibition in a social context

Intentional inhibition is frequently used in social interactions. When communicating or competing with others, we are in a continuous state of planning what to say or do next. In

everyday conversation, there is typically no immediate external stimulus to instruct us to inhibit our next action. Instead, inhibition in this case must be an endogenous self-generated decision based on our broader perception of the environment. Often we may be inclined to say something inappropriate, for example a taboo word, which we then feel must be stopped (Severens et al., 2011a, 2011b). The neural correlates of inhibiting speech show significant overlap with those involved in inhibiting manual actions (Xue et al. 2008). We regulate our interactions with others by acting through speech and gesture, but also through being able to inhibit each of these. An effective balance between action and inhibition may be very important for factors such as turn-taking during successful communication (Stivers et al., 2009).

1.1.4. Inhibition in law and morality

In discussions about moral and legal responsibility, we will often hear “they were not forced to commit the act; they could have stopped themselves”. These statements strongly highlight the importance of our capacity for intentional inhibition. Many crimes involve an impulsive act and often a failure to inhibit a regrettable action at the last moment. Studying the mechanisms underlying intentional inhibition could be very useful in understanding the risk factors for failed inhibition, and how such risk factors can be minimised. Research on instructed stimulus-driven inhibition is useful in revealing basic processes and neural correlates of inhibitory processes (Verbruggen & Logan, 2008), but it lacks the ‘real-world’ quality of everyday action control which intentional inhibition tries to capture. Studying intentional inhibition gives a unique insight into how we make *decisions* to inhibit action, rather than simply how we employ a generic inhibitory process.

1.2. The scientific study of instructed inhibition

Most of what we know about inhibitory function comes from experiments in which participants are instructed to inhibit an action when they receive an immediate sensory stimulus. The participant does not make an endogenous decision to inhibit an action, but rather reacts to an unambiguous signal in the environment which directs them to inhibit their action. The two most popular paradigms for examining this have been the stop signal and go/nogo paradigms (Verbruggen & Logan, 2008).

1.2.1. The stop signal paradigm

In the stop-signal task, participants prepare to execute an action following a Go cue. They then receive a 'stop' stimulus instructing them to withhold the action at the last moment (Verbruggen & Logan, 2008). The time between the go signal and stop signal is referred to as the stop signal delay (SSD). The probability of successfully inhibiting a response can be manipulated by varying this delay. Errors of commission are more likely with a longer SSD, as the action execution process is interrupted later. Therefore, stopping becomes more difficult.

The SSD can be adjusted to the point at which the probability of a successful inhibition is 50%. When this SSD is subtracted from the participant's median reaction time for actions, it yields the stop signal reaction time (SSRT) (Logan, 1994; Logan & Cowan, 1984). This can be seen below in figure 1.1. The SSRT provides information on how quickly the inhibition process can be applied successfully before action execution. A shorter SSRT means more effective inhibition. Studies using fMRI have shown that areas such as the inferior frontal gyrus (IFG), pre-supplementary motor area (preSMA) and subthalamic nucleus (STN) become active during successful stopping (Aron & Poldrack, 2006; Duann et al., 2009; Garavan, Ross, & Stein, 1999; Li et al., 2006; Menon et al., 2001; Rubia et al., 2003). Activity in inhibitory motor areas is greater for participants with shorter SSRTs and therefore a better ability to inhibit their responses (Aron & Poldrack, 2006; Li et al., 2006).




Figure 1.1. Key principles underlying the stop-signal task. SSD is set to produce 50% correct inhibitions and 50% errors of commission following the stop signal. The SSD can then be subtracted from the median of the reaction time distribution for actions (dotted line) to give the SSRT (reproduced from Aron & Poldrack, 2006).

1.2.2. The go/nogo paradigm

In the Go/NoGo task, participants are exposed to one of two cues that instruct them to either act or withhold action. This differs from the stop signal task in that there is not always an action cue present on every trial. Instead of forcing participants to initiate action responses on every trial before presenting a stop stimulus, the go/nogo task manipulates action prepotency by varying the probability of go vs nogo cues (Eimer, 1993). Highly frequent go cues in a block of trials for example, would force people into being more ready to act. This would consequently increase the level of inhibition required to withhold the action when a NoGo cue appears.

1.2.3. The neural correlates of instructed inhibition

1.2.3.1. EEG

EEG is frequently used to study action inhibition. Event-related potential (ERP) correlates of inhibition are well-known. In the go/nogo task, nogo cues elicit greater N2 and P3 components than Go cues (Bokura, Yamaguchi, & Kobayashi, 2001; Eimer, 1993).

Participants with high false alarm rates (i.e. acting when they are required to inhibit their action) have smaller nogo-N2 components than those with lower false alarm rates, suggesting that the amplitude of the N2 component may reflect the strength of the inhibition process between participants (Falkenstein, Hoorman, & Hohnsbein, 1999).

Evidence from the stop-signal tasks provides convergent evidence that N2-P3 components are important in action inhibition. Trials which include stop signals produce higher-amplitude components than trials in which there is no stop signal (Kok et al., 2004). N2 and P3 components in response to stop signals are also known to be larger when those stop signals are infrequent rather than frequent: due to the dominance of action responses, action preparation is more prominent and requires stronger inhibition to be stopped (Ramautar et al., 2004). Patients with attention deficit hyperactivity disorder (ADHD) who exhibit problems with action inhibition show smaller N2 components than healthy controls (Dimoska et al., 2003). Action inhibition paradigms using EEG strongly suggest that the N2 component is closely related to inhibitory processes.

Other EEG evidence examines time-frequency oscillation data rather than ERPs. Neural oscillations in the alpha and beta band show strong desynchronisation patterns (reduced power) over motor areas during voluntary actions (Pfurtscheller & Lopes da Silva, 1999). Increases in alpha and beta power from baseline power levels have been linked to top-

down action inhibition processes (Klimesch et al., 2007). Intracranial EEG measurements in patient groups undergoing surgery also show greater beta activity in brain areas related to inhibition such as the right inferior frontal gyrus (IFG), when patients successfully stop an action compared to when they are unsuccessful (Swann et al., 2009). In these stop trials, patients also show reduced action-related beta desynchronisation in primary motor areas, suggesting that synchronising beta oscillations in IFG may directly inhibit action at the final brain areas for motor output.

1.2.3.2. Functional magnetic resonance imaging (fMRI)

Evidence from fMRI gives detailed spatial information about the neural correlates of inhibition. In the stop-signal task, action inhibition is associated with activation in the right IFG and subthalamic nucleus (STN). Activation likelihood estimate (ALE) meta-analyses of go-nogo tasks also suggest an important role for the inferior frontal gyrus in action inhibition (Simmonds et al., 2008). Networks involving the IFG, STN, and pre-supplementary motor area appear to be critical for inhibitory function but there is disagreement about the nature of the functional connections between each node. Some suggest direct connections between IFG and STN during action inhibition (Aron & Poldrack, 2006), while others propose direct connections only between IFG and pre-SMA with projections from the pre-SMA then reaching the STN (Duann et al., 2009). The pre-SMA may have a less specific role in inhibition, in particular given that SSRT correlates with IFG activity, but *not* with pre-SMA activity (Aron et al., 2007). Its role in action stopping could be a more general function such as action monitoring and conflict-detection.

1.2.3.3. Brain stimulation

Studies using transcranial magnetic stimulation (TMS) and transcranial direct current stimulation (tDCS) are informative about the causal roles of particular brain areas for inhibition. Repetitive TMS at 1Hz has been used to disrupt activity at various anatomical sites in the brain during a stop signal task. Disruption of the right IFG produced significant impairments of inhibition performance, with longer SSRTs and more inhibition errors (Chambers et al., 2006). Stimulation of the angular gyrus or middle frontal gyrus produced no such changes in behaviour, suggesting a critical role specifically for the IFG in action inhibition.

Stimulation with tDCS over the pre-SMA also significantly affects inhibitory performance in a stop-signal task. Anodal stimulation – thought to enhance neural activity – reduces the number of errors of commission made when a stop signal is encountered. Conversely, cathodal stimulation – thought to reduce neural activity – shows a trend towards increasing the number of errors compared to baseline in which there is no stimulation (Hsu et al., 2011). Therefore, enhancement and impairment of pre-SMA function seem to improve and impair inhibitory function respectively. Similarly, anodal stimulation of the right IFG improves inhibition in the stop-signal task by improving SSRTs compared to control sham stimulation (Jacobson et al., 2011). Direct electrical stimulation of the right IFG slows down actions, particularly when participants are additionally recruiting task-related behavioural inhibition (Wessel et al., 2013). These brain stimulation findings support neuroimaging evidence in identifying inhibitory roles for right IFG and pre-SMA areas (see above).

1.3. The scientific study of intentional inhibition

Research into instructed inhibition has been very informative, particularly in exposing important neural networks underlying successful stopping of actions. Intentional inhibition has been relatively neglected in research however. This could be for two reasons: 1) it is assumed that intentional inhibition is not distinct from instructed inhibition; 2) it is difficult to design good experiments to address intentional inhibition. This section will highlight the importance of studying intentional inhibition as a distinct process. It will then go on to identify existing work in the area and the challenges to relevant experimental design.

1.3.1. Relevance and importance for scientific research

We know a lot about the ‘inhibition’ part of intentional inhibition but not much about the ‘intentional’ part. A similar trend can be found in the literature on action. The most influential experiments in action-related ‘free will’ began relatively recently (e.g. Libet et al., 1983) even though research on the basic neuroscience of action such as the readiness potential started a long time before this (e.g. Kornhuber & Deecke, 1965). Research on ‘free’ intentional actions has advanced a long way and we now have a strong understanding about how it differs from instructed action (Fleming et al., 2009), and the areas of the brain that are important in driving our choices to act (Soon et al., 2008).

Intentional inhibition has yet to reach this milestone. Our understanding of its neural precursors and how it differs from instructed inhibition is rather limited. Improving our knowledge is important however, particularly because intentional inhibition is a major part of our everyday lives. If it is a unique and independent function from instructed inhibition, then there may be a large gap in our understanding of everyday inhibitory behaviour. In addition to examining how we react to external stop signals, we need to study how we make our own endogenous decisions to cancel planned actions at the last moment.

1.3.2. Similarities to existing concepts

1.3.2.1. Self-control

One well-studied concept in experimental psychology is very relevant to intentional inhibition, and is referred to as ‘self-control’. A key theory in this area is ‘ego depletion theory’. This explains that our capacity for self-control is limited and can be exhausted (Baumeister, Vohs, & Tice, 2007; Muraven & Baumeister, 2000). The capacity for self-control on a task reduces after having exerted it on a previous task, even when the subject of the tasks is entirely different. For example exerting self-control in one task involving suppressing forbidden thoughts, impairs self-control performance in a subsequent task involving persistence at unsolvable anagrams (Muraven, Tice, & Baumeister, 1998). This work illustrates that the inhibition of action is an active effortful process which depletes a central resource (Muraven, Tice, & Baumeister, 1998). Depletion of this resource may explain many failures of self-control in the real world – for example in attempts to quit smoking or persistence at dieting.

Self-control usually refers to our capacity to refrain from giving in to some temptation over extended periods of time. The benefits of self-control in society are far-reaching and emphasis is placed on having this quality throughout life. Many have argued that the decline in violence – particularly homicide – over the last few centuries in many parts of the world including Europe, has been due to increasing levels of self-control (Eisner, 2001; Pinker, 2011). This increase may be explained by a number of factors, including improving governments and judicial institutions with a state monopoly on violence, and expanding schooling and literacy (Eisner, 2001). Self-control at a young age as measured by questionnaires, reports, and performance on delay-of-gratification tasks has been found to predict later academic success (Duckworth & Seligman, 2005; Shoda, Mischel, & Peake,

1990). Self-control also correlates with better psychological adjustment, including fewer alcohol abuse problems and anxiety disorders (Tangney, Baumeister, & Boone, 2004).

Our understanding of self-control is still relatively limited. Continued study with a wider set of experimental approaches is necessary. Ego depletion theory for example is a useful tool for comprehending the limits of self-control but lacks any detailed account of how the inhibition mechanism itself functions. Many other constructs in psychology have been associated with limited resources, including memory (Awh, Barton, & Vogel, 2007; Bays & Husain, 2008), and other attentional capacities (Marois & Ivanoff, 2005). The ‘limited resource’ idea itself is therefore not novel. To improve our understanding of self-control, we need to study both the higher-level decision processes associated with it, and the more specific intentional inhibition mechanisms that underlie it.

1.3.2.2. Delayed gratification/delay discounting

The phenomenon of self-control is well-captured in studies that use ‘delay-of-gratification’ tasks (McClure et al., 2004, 2007; Metcalfe & Mischel, 1999; Mischel, Shoda, & Rodriguez, 1989). In these studies, participants are offered the option of obtaining a small immediate reward, or waiting a length of time in order to obtain a larger or preferred reward. Humans are known to value immediate rewards greater than delayed rewards (e.g. “I would rather have £100 now than £100 in a month”) and preferences follow a function in which the future is discounted hyperbolically (Green & Myerson, 2004; Kirby & Herrnstein, 1995). However, everyday life is fraught with many competing motives such as drug addiction, health states, retirement planning, and other economic factors which make building an effective model with a single discounting rate - and successfully applying it to the real world - a difficult task (Frederick, Loewenstein, & O’Donoghue, 2002). When rewards involve a drug to which a participant is addicted for example, the future is discounted far more steeply than normal (Bickel & Marsch, 2001; Bickel et al., 2007; Kirby & Petry, 2004).

Several studies have investigated the delay-of-gratification effect in children (Mischel & Ebbesen, 1970; Mischel, Ebbesen, & Zeiss, 1972). Participants from 3-5 years old were presented with cookies and pretzels and told that they could immediately have their less preferred item, or wait for an extended period of time and receive their favoured item (Mischel & Ebbesen, 1970). When the objects were out of view and not attended to, children waited on average 11 minutes before settling for the immediate reward. Self-control was

weakened when rewards were attended to during the waiting period, with waiting times dropping as low as 1 minute on average. An inability to distract attention away from desired objects can diminish our ability to inhibit actions towards them.

1.3.2.3. Decision-making

At the core of intentional inhibition is the need to decide between alternatives. Rather than deciding between two actions, intentional inhibition usually involves the choice between acting and inhibiting action. The decision-making literature is vast, with most work focusing on how humans and non-human animals decide between two movements. A commonly used task is the random dot motion task (Gold & Shadlen, 2007), in which moving dots are presented on a screen, and participants must decide whether most of them are moving left or right. In monkeys, the lateral intraparietal area (LIP) predictively codes for the decision to move the eyes left or right following the accumulation of sensory evidence from the dot motion stimuli (Shadlen & Newsome, 2001). This happens even when all dots are moving entirely randomly and therefore there is no coherent motion to the left or right – in this case there is no strong external evidence to guide decisions and the monkeys may endogenously generate a decision.

Another important factor which may contribute to decision-making is spontaneous neural activity (Fox & Raichle, 2007; Sadaghiani et al., 2010). When making ambiguous decisions, spontaneous noise in sensory or motor neural systems can often determine what we do. For example, during perceptual decisions about ambiguous vase-face figures, levels of pre-stimulus activity in the fusiform face area can predict whether a face is perceived in the subsequent figure (Hesselmann et al., 2008). Random spontaneous activity or variability in biological systems is an adaptive advantage as it reduces predictability against predators (Brembs, 2011). It is therefore no surprise that such variability contributes to decision-making.

1.3.3. *Challenges in intentional inhibition research*

1.3.3.1. *'Free' choices*

Intentional inhibition involves a 'free' choice to stop action. The term 'free' typically refers to the human ability to make or withhold actions at will in everyday life. This is a rather difficult context to recreate in the laboratory, as flexibility in behaviour must be minimised. Experimental tasks introduce numerous constraints, both on behaviour itself and on the general context for decision-making. An important criticism for all research in the area is therefore that free choices cannot be studied when freedom is being minimised in the experimental setup.

Although there are necessary constraints on human behaviour in behavioural experiments, the term 'free choice' can still be usefully and objectively defined for study. Free choices in the lab refer to action and inhibition decisions that are not instructed by immediate identifiable external signals (Haggard, 2008). They are usually contrasted with actions in which one particular response must be made following an external stimulus.

Different elements of a free choice can be isolated and studied in experiments. Participants can choose *what* action to make, *when* to make an action, or *whether* to make an action (Brass & Haggard, 2008). For example, participants could decide between using their left hand or right hand for a response (Soon et al., 2008); this would be an example of a *what* decision. Participants could also decide on the exact point in time at which they want to press a button (Libet et al., 1983); here they would be executing a *when* decision. Finally, they could decide whether to execute the action at all (Brass & Haggard, 2007); this is the *whether* decision and frequently relates to intentional inhibition. Each of these refers to a different aspect of free choice, but all of them involve the presence of a choice between multiple available options – a choice which is not instructed by any obvious or immediate signal. Studying voluntary responses in this way and comparing them with actions or inhibitions in response to a specific informational stimulus, allows us to gain useful insights into the nature of endogenous decisions.

1.3.3.2. *Prepotent action and pre-decision*

Inhibition is assumed to be an active process as it involves the suppression of a prepotent action impulse. If there is no prepotent action impulse to suppress, then in principle there cannot be any inhibition. This is a challenge when it comes to intentional inhibition research,

as experiments must somehow ensure that participants are preparing an action before cancelling it. In typical instructed inhibition paradigms, this can be done quite easily by forcing participants to act often and presenting them with unexpected stop signals late in the action preparation process on some actions. However, stop signals cannot be used in intentional inhibition experiments.

Although stop signals cannot be used, participants can still be forced to act frequently enough that they develop a tendency to prepare action in intentional inhibition paradigms. Instead of stop signals, participants can then be presented with choice stimuli, which encourage them to decide at that moment whether to continue with a prepared action or choose to suppress it (Parkinson & Haggard, 2014). Another way to encourage action preparation is to provide a motivation for action responses, by for example presenting unpleasant sensory outcomes when actions fail (Kuhn et al., 2009). Participants can also be given strict time limits for making action responses, so that they must always be ready to act before seeing action or choice targets.

1.3.3.3. A lack of behavioural output

One other characteristic challenge to studying intentional inhibition is that there is no directly measurable behavioural output. This naturally makes behavioural experiments difficult because there are no reaction times or movement dynamics to measure. However, it also makes imaging experiments difficult, because there is no way to precisely time-lock to the inhibition response in the recorded data. There can also be no explicit stop signal to link to inhibition itself. Although free choice stimuli can be presented, it can be difficult to disentangle the intentional inhibition process from the decision-making process in imaging data. Often the choice stimulus itself can also act as an explicit stop signal in forcing participants to inhibit their action preparation while they decide on what they want to do next. This means that any activity time-locked to the choice stimulus may not be distinct from instructed inhibition.

This problem requires careful experimental design with appropriate control conditions. For example, although time-locking to the choice stimulus could elicit decision-making as well as intentional inhibition, these could be disentangled by comparing inhibition choices with action (Kuhn et al., 2009). Decision-making should be common to both of these choices, so contrasting them allows direct comparison of the effects of intentional inhibition

itself with intentional action. Attempts to imprecisely time-lock to an inhibition response itself are also possible, for example by time-locking to mean action reaction times during an inhibition. This imprecision could be accounted for by comparing inferred times of intentional inhibition with inferred times of instructed inhibition (Misirlisoy & Haggard, 2014). These should suffer from the same imprecision and so any differences between them in imaging data should be explained by differences in the intentionality of the inhibition process, rather than the imprecise time-locking.

1.3.4. Past research in intentional inhibition

Despite the significant challenges to intentional inhibition research, several studies have attempted to tackle the difficult problem. Below is a description and evaluation of the existing literature on intentional inhibition.

1.3.4.1. Behavioural investigations

Standard behavioural paradigms for studying intentional inhibition are lacking. One approach is to use existing instructed inhibition paradigms, and adapt them to include a ‘free choice’ component. Parkinson & Haggard (2014) used a go/nogo paradigm with subliminal priming in order to investigate whether proportions of free choices to inhibit action can be influenced by priming. In addition to standard go and nogo response targets, free choice targets were included which instructed participants to freely decide between acting and inhibiting action. Two types of priming were used: positive compatibility effect (PCE) in which compatible primes facilitate responding compared to neutral primes, and negative compatibility effect (NCE) in which compatible primes inhibit responding compared to neutral primes (Eimer & Schlaghecken, 1998). With PCE priming, no overall significant effect could be found on free choices. With NCE priming on the other hand, go and nogo primes increased the probability that the opposite behavioural response would be freely chosen. This inhibitory priming which facilitates opposing responses is therefore capable of subconsciously influencing free choices to inhibit action.

Another important influence on our capacity to choose to inhibit actions is higher-level beliefs. Participants who have been exposed to anti-free will messages, and therefore have been induced to reject the concept of human free will, are less likely to exert self-control

and inhibit actions when given the choice (Rigoni et al., 2012). These influences on our free decisions to inhibit action have only become apparent through studying intentional inhibition. Instructed inhibition paradigms are not able to investigate such effects on human choice.

Intentional inhibition also differs to action in its affinity to external events that follow its occurrence. For example, the ‘binding effect’ in agency experiments shows that subjective judgements of the time of a tone are earlier when that tone follows a voluntary action than when it does not (Haggard et al., 2002). This suggests that the feeling of having caused an external event temporally attracts that event towards our action. However, when such an event follows inhibition of action, there is in fact a trend in the opposite direction towards ‘repulsion’ of the tone (Haggard et al., 2009). There may therefore be a natural human assumption that inhibitions do *not* cause external events.

1.3.4.2. *fMRI*

The spatial neural correlates of intentional inhibition have been illuminated by fMRI studies. Brass & Haggard (2007) adapted the Libet paradigm to study intentional inhibition. In the standard Libet task, participants watch a rotating clock and press a button at a time that they choose (Libet et al., 1983). They then report the time on the clock at which they first had the ‘urge’ or intention to press the button. In the adapted task, participants could also choose to inhibit their urge to press the button at the last moment before executing the intended action. When actions were inhibited, significant activation was found in the dorsal fronto-median cortex (dFMC), in contrast to executed actions (Brass & Haggard, 2007). This study provided the first indication that the dFMC area may be recruited in intentional inhibition.

Subsequent studies have supported this conclusion. In the ‘marble task’ (Kuhn et al., 2009), participants watch a marble rolling down a ramp. When this marble reaches the bottom of the ramp, it smashes and creates an unpleasant sound. In some cases when this marble changes colour from white to green, participants are instructed to press a key as fast as possible – if successful, the marble is intercepted and is prevented from smashing. When the marble remains white, participants must instead freely choose whether or not to act and prevent the marble rolling down the ramp. When participants freely decide to prevent action compared to executing action, activation is found in the dFMC (Kuhn et al., 2009). Increased effective connectivity is also found between the dFMC and pre-SMA areas, when participants decide to inhibit rather than execute action.

Other studies using the marble task suggest strong overlap in brain activations between intentional and instructed inhibition, but a unique contribution from the dFMC depending on previous trial history. With a greater number of preceding consecutive green marble (forced action) trials, the dFMC is less recruited for intentional inhibition during choice trials (Schel et al., 2014). This suggests that greater automatization of action reduces the contribution of dFMC to choices, and consequently reduces the likelihood that participants will choose to inhibit action.

1.3.4.3. EEG

EEG experiments contribute further information about the neural correlates of intentional inhibition. The adapted Libet task used for studying intentional inhibition in fMRI (Brass & Haggard, 2007) has also been used with EEG. When participants choose to inhibit their action, there is an increase in alpha/lower-beta oscillation power, just after the initial intention to act (Walsh et al., 2010). When looking at neural activity *preceding* choice stimuli, as represented by contingent negative variation (CNV), greater amplitudes can be found when participants subsequently choose to act than when participants choose to inhibit action (Filevich et al., 2013).

1.4. Disorders of action and inhibition

As with many psychological processes, inhibitory mechanisms can become dysfunctional. Studies on disorders of inhibition are helpful in the development of new treatments and also in providing alternative perspectives on intentional inhibition in healthy people. It is often difficult to reliably infer what the core deficits or causes of a neuropsychiatric disorder are. The following section will attempt to identify and describe disorders that include inhibitory deficits as one of the major symptomatic categories.

1.4.1. Psychiatric disorders

1.4.1.1. Attention deficit hyperactivity disorder (ADHD)

Attention deficit hyperactivity disorder (ADHD) is a disorder characterised by impulsive behaviours and a lack of ability to sustain attention. It has been associated with deficient response inhibition mechanisms (Barkley, 1997; Scheres et al., 2004), and with altered function in IFG (Rubia et al., 2005). In a meta-analysis using ALE with 16 neuroimaging studies, both the IFG and striatum showed significantly reduced activity in ADHD patients relative to healthy controls (Dickstein et al., 2006). Another meta-analysis of 8 studies between 1990 and 1997 investigated children with ADHD, and children with conduct disorder. It found very similar deficits of inhibitory function across the two groups, with both showing slower SSRTs than healthy controls (Oosterlaan, Logan, & Sergeant, 1998).

A review of the available literature on twin studies across the world suggests an average heritability estimate of 76% for ADHD (Faraone et al., 2005). Assuming that response inhibition is one of the critical impairments in ADHD, this suggests that the individual capacity to inhibit actions could be largely innate. Candidate genes associated with the disorder include those within dopamine systems (e.g. DAT1, DRD4, DRD5) and serotonin systems (e.g. 5HTT), and those associated with synaptic plasticity or synaptic vesicle function in the regulation of neurotransmitter release (e.g. SNAP-25) (Gizer, Ficks, & Waldman, 2009). Data from twin studies in the wider healthy population also suggest that genetic factors explain high amounts of variance in levels of self-control (Beaver et al., 2008; Beaver, Ratchford, & Ferguson, 2009; Beaver et al., 2009). Experiments looking more directly at inhibitory performance itself also suggest that the individual capacity to exert self-control and inhibit action may be a trait of genetic origin (Aron & Poldrack, 2005; Congdon et al., 2008).

1.4.1.2. Obsessive-compulsive disorder (OCD)

Several behavioural studies suggest inhibitory problems may characterise obsessive-compulsive disorder (OCD). In an oculomotor antisaccade task where participants are instructed to move their eyes in the opposite direction to which a target is presented, more inhibition failures are observed in OCD patients compared to healthy controls, because patients reflexively look in the direction of the stimulus (Rosenberg et al., 1997). Similarly, a lower percentage of successful inhibition responses have been reported in Go/NoGo tasks

compared to healthy controls (Penades et al., 2007). One study compared OCD patients with patients diagnosed with panic disorder and found more errors of commission to NoGo stimuli amongst OCD patients (Bannon et al., 2002).

Trichotillomania is a disorder that has been related to OCD and the clinical presentation involves compulsive hair pulling from the body (Ferrao et al., 2009; Swedo & Leonard, 1992). High scores on scales of novelty seeking suggest impulsivity may be one characteristic associated with this disorder (Lochner et al., 2005). Patients with trichotillomania have been found to exhibit longer SSRTs and therefore worse inhibition in the stop-signal task compared to OCD patients, with symptom severity significantly correlating with SSRT (Chamberlain et al., 2006).

Table 1.1 below displays behavioural performance and fMRI activity during standard action inhibition tasks for patients with ADHD, OCD, and other disorders relative to healthy controls. It is interesting to note that in many of these fMRI studies, no difference could be found in behavioural performance between healthy controls and patients. It may be that when external stop signals instruct inhibition, impairment is less clear and consistent. Given the differences in inhibitory network activity during these tasks – particularly in the IFG, pre-SMA and striatum – it would be interesting to test these patients on more endogenous types of inhibition task to see whether behavioural impairment becomes more apparent.

Disorder	Performance in a stop signal/go nogo task (patients relative to controls)	Neural basis – fMRI activation (*reduced relative to control ** increased relative to control)	Relevance to action inhibition
ADHD	<i>Rubia et al. (2005)</i> No difference	<i>Rubia et al. (2005)</i> * IFG, insula, precentral gyrus, STG	Inability to sustain attention and inhibit action impulses
	<i>Tamm et al. (2004)</i> More errors of omission and commission	<i>Tamm et al. (2004)</i> *Cingulate cortex, SMA, SFG, middle frontal	
Drug and alcohol abuse	<i>Kaufman et al. (2003)</i> More errors of omission and commission	<i>Kaufman et al. (2003)</i> *ACC, insula, pre-SMA, IFG	Inability to inhibit impulse to consume drugs/alcohol
	<i>Li et al. (2008)</i> No difference	<i>Li et al. (2008)</i> *ACC	
	<i>Li et al. (2009)</i> Slower RT, higher stop success	<i>Li et al. (2009)</i> *Dorsolateral prefrontal cortex, cingulate gyrus	
Obsessive-compulsive disorder	<i>Woolley et al. (2008)</i> No difference	<i>Woolley et al. (2008)</i> *OFC, thalamus, striatum, insula, STG	Inability to inhibit perseverative thoughts and behavioural compulsions
	<i>Maltby et al. (2005)</i> No difference	<i>Maltby et al. (2005)</i> ** ACC, OFC, lateral prefrontal, thalamus, striatum (caudate)	
Schizophrenia	<i>Rubia et al. (2001)</i> No difference	<i>Rubia et al. (2001)</i> *ACC, SFG ** Thalamus, striatum (putamen)	Inability to inhibit inappropriate responses and irrelevant task sets
	<i>Kaladjian et al. (2007)</i> No difference in error – slower RTs	<i>Kaladjian et al. (2007)</i> *IFG	

Table 1.1. Behavioural performance and brain activity during stop-signal/Go-NoGo tasks for various patient groups. ACC – anterior cingulate cortex, IFG – inferior frontal gyrus, OFC – orbitofrontal cortex, STG – superior temporal gyrus, SFG - superior frontal gyrus

1.4.1.3. Drug abuse

Chronic drug users may also show inhibitory deficits. Impulsivity and failures to inhibit action are highly relevant to drug abuse and play a critical role in the various phases of addiction (Perry & Carroll, 2008). Methamphetamine abusers show abnormalities in inhibitory neural networks and show significantly longer SSRTs than healthy controls (Monterosso et al., 2005). Similarly, cocaine users also show behavioural impairment in inhibition, with longer SSRTs than controls and a reduced probability of successfully inhibiting responses (Fillmore & Rush, 2002). Action performance on the other hand is often preserved in these groups, suggesting that impairment is specific to inhibitory function rather than more global behaviour. Inhibitory deficits in drug abuse may be better understood with intentional inhibition research. A detailed and comprehensive account of inhibition in drug-related disorders requires a stronger understanding of failures to inhibit habitual actions, especially in the context of everyday repeated decision-making.

1.4.2. Tourette syndrome

Tourette syndrome (TS) appears to be a special case amongst the disorders related to inhibition, in that the areas typically related to tic inhibition such as the striatum, prefrontal, and medial frontal areas are increased rather than reduced in their activity (Baym et al., 2008; Bohlhalter et al., 2006). Rather than TS patients having an impaired inhibitory system, it seems likely that their inhibitory system could be normal (Roessner et al., 2008) and potentially working excessively in attempting to deal with overwhelming motor urges. Immediately prior to tic onset, activation in areas such as the SMA and anterior cingulate cortex (ACC) reflect the upcoming urge to move (Bohlhalter et al., 2006), and these same areas have been associated with the building urge to scratch an itch following histamine injections (Hsieh et al., 1994). Levels of activity in areas including the IFG and STN positively correlate with tic severity suggesting that this inhibitory network may act as a compensatory system for TS motor urges and symptoms (Baym et al., 2008). Coherence in neural activity between sensorimotor and frontomedian areas seem to be important in the suppression of involuntary twitches. Furthermore, this network shows increased activity relative to healthy controls during inhibition of voluntary movement in go/nogo tasks despite no behavioural differences in performance (Serrien et al., 2005).

1.4.3. Neuropsychological disorders and brain lesions

Inhibitory control can be impaired by acquired damage to specific areas of the brain. Lesions to the orbitofrontal cortex in rats lead to slower SSRTs, and STN lesions produce a more general impairment in stopping accuracy (Eagle et al., 2008). Human patients with lesions to the IFG show profound impairment in inhibition with significantly slower SSRTs that correlate with the amount of damage in the area (Aron et al., 2003). Deactivating the IFG with 1Hz transcranial magnetic stimulation (TMS) for 15 minutes also impairs inhibitory control by reducing the number of successful inhibitions in a stop-signal task (Chambers et al., 2006). Lesions to medial frontal areas, in particular the supplementary motor area (SMA), affect inhibitory control in a Go/Nogo task, by increasing the number of errors of commission to an infrequent NoGo signal (Picton et al., 2007).

1.4.3.1. Anarchic hand syndrome

Anarchic Hand Syndrome (AHS) leads to the production of unwanted and involuntary movements of a limb (Marchetti & Della Sala, 1998). These movements cannot be voluntarily suppressed or inhibited by the patient and often will result in attempts of physical restraint by the alternative healthy limb. A common behavioural presentation is utilisation behaviour in which the presence of tools in the environment elicits automatic and uncontrollable actions directed at them by the affected limb. In a clinical review of 20 reports of AHS based on anatomical damage in the brain, separate types of the disorder could be identified (Feinberg et al., 1992). A 'frontal' type of AHS involved left hemisphere lesions of the SMA, anterior cingulate, and medial prefrontal cortex and resulted in frequent compulsive movements to environmental stimuli by the dominant limb. Callosal AHS occurred following damage restricted to the corpus callosum, and this typically resulted in AHS of the nondominant limb, primarily involving behavioural patterns of intermanual conflict in which interhemispheric inhibition of the limb is absent. In these cases, less frequent utilisation behaviour and compulsive grasping movements were reported than those found in frontal AHS where damage to critical inhibition areas is more prominent.

An experiment conducted on a frontal AHS patient examined whether endogenously generated actions differed from exogenously generated actions in the affected limb (Kritikos, Breen, & Mattingley, 2005). In this experiment, the patient sat in front of a response board

with 23 buttons containing LEDs. The patient was instructed to press a sequence of buttons as indicated by the appearing lights on the buttons as quickly as possible. In the exogenous condition, after a key was pressed down, the next light appeared at the new response location, requiring a quick and externally determined action. In the endogenous condition however, the light only appeared at the new location after the previous button was released. The authors reasoned that this condition required a new internally generated action. The difference in response times between exogenous and endogenous responses was significantly larger when responses were made using the affected anarchic hand rather than the normal hand. AHS significantly altered the capacity to endogenously generate actions with the affected limb. No such pattern was found between hands in the performance of a healthy control participant.

AHS is a disorder with particular relevance to action inhibition. It suggests that actions are continuously suppressed in healthy people, preventing them from engaging in compulsive behaviours driven by external stimuli. Several theories suggest that the perception of objects in the environment automatically triggers related action representations and motor planning (Grèzes & Decety, 2002; Grèzes et al., 2003). Without an inhibitory system in place to put a brake on these automatic motor plans, one would be in a constant state of utilisation behaviour.

The experiment by Kritikos et al. (2005) highlights differences between exogenous and endogenous action initiation, and shows that endogenous actions are particularly affected when the capacity to exert inhibitory control is damaged. The endogenous condition in their experimental design still involved acting towards a stimulus in the environment however. An investigation of action and inhibition performance in purely self-generated endogenous circumstances would be interesting in disorders such as AHS, and in many of the other disorders described in this section.

1.5. Knowledge gaps in the study of intentional inhibition

Intentional inhibition remains a relatively unexplored area. Initial questioning should aim to clearly identify whether intentional inhibition is a bona fide cognitive function worth studying in its own right. It could be argued that inhibition is simply a symmetrical response alternative to action, with no special qualities as an option in decision-making (Mostofsky &

Simmonds, 2008). Similarly, it may not be distinct or unique in comparison to instructed inhibition, and so typical go/nogo and stop signal paradigms may be sufficient for our understanding of everyday behavioural inhibition.

If initial research suggests that intentional inhibition is in fact a distinct cognitive function, there could be significant benefits to studying it. Further research could contribute to clinical psychology and our basic knowledge of underlying inhibitory mechanisms in the brain. Inhibitory failures and deficits characterise several psychiatric disorders, and yet patient performance in instructed inhibition tasks often does not differ from healthy participants. Using experiments that are capable of reflecting everyday impairment experienced by patients is an important task. Similarly, the focus on instructed inhibition tasks in academic research leaves a large gap in our knowledge of intentional inhibition which may be more relevant to real-world behaviour.

As described above in section 1.3.4., basic research on the behavioural qualities and neural correlates of intentional inhibition has already started. Progress is still necessary however, in particular with developing a more detailed account of the factors that ultimately influence or determine our decision to inhibit action. An understanding of the interaction between intentional inhibition and perceptual decision-making is also currently lacking. Our everyday decisions to act or inhibit are often based on weighing up the external evidence available to us, and then assessing the perceptual outcomes that follow the overt expression of our decision. This cycle of monitoring the world and using the information in our action/inhibition choices is crucial to human behaviour, and should therefore be studied experimentally.

1.6. Aims of the current thesis

This thesis will aim to address some of the important gaps in our understanding of intentional inhibition. It will attempt to develop and present a coherent account of how we make our decisions to inhibit action, and the influences and biases that contribute to that decision.

Chapter 2 will focus on the interaction between perception and intentional inhibition decision-making. It will begin by looking at the accumulation of perceptual evidence leading up to our decisions to act or inhibit. It will investigate whether the evidence accumulation

process between action and inhibition is identical, or whether there are biases introduced by the response systems that alter how that evidence is used. It will also report an experiment which examines how perception of post-decision sensory outcomes is affected by action and inhibition responses.

Chapter 3 will examine the effects of past events on our future decisions to inhibit action. It will look at both behavioural and neural precursors of our inhibition decisions. It will investigate how cognitive biases affect our choices, and how unconscious external information relates to these biases. Chapter 4 will further present a natural experiment studying the effects of cognitive biases in elite sports, and describing how differences in action/inhibition states between opponents can dramatically alter competitive decisions.

Chapter 5 will look at the neural correlates of intentional inhibition using two different methods: EEG and fMRI. It will explore how spontaneous states in the brain can contribute to our final decisions, and the specific areas of the brain that may underlie intentional inhibition.

Chapter 6 will focus on patients with Tourette syndrome and will explore the unusual extent of intentional inhibition that they must engage in. It will describe an experiment in which tic frequency reductions due to attentional manipulation are compared with those due to tic inhibition. It will also investigate the EEG neural correlates of tic inhibition.

The thesis will end with a general discussion. This will use the evidence presented through the preceding experimental chapters to propose a model of inhibitory decision-making in which intentional and instructed inhibition lie on extremes of a continuum, rather than being entirely independent processes.

Chapter 2. Perceptual processing and intentional inhibition

This chapter reports three experiments that explore the effects of intentional inhibition on perceptual processing and decision-making. Experiments 2.1 and 2.2 investigated how people accumulate external evidence when making decisions to act or inhibit. The data showed that decisions to act are made significantly more cautiously than decisions to inhibit, even when external evidence for each response is identical. Experiment 2.3 investigated the perceptual processing of sensory outcomes following action and inhibition decisions. Action and inhibition responses produced asymmetric detection abilities and feelings of agency over delayed sensory outcomes. Overall, the chapter demonstrates that decision-making parameters and perceptual processing capacities are adjusted depending on whether intentional action or inhibition responses are made. The data suggests that intentional inhibition can usefully be studied using the methods and frameworks of decision-making and action-perception interactions.

2.1. Introduction

Action and perception systems are tightly coupled in everyday behaviour. In their bidirectional interaction, perceptions of objects directly adjust the action system to allow for efficient motor behaviours (Makris et al., 2011), and actions prime the perceptual system to focus on the most relevant information in the environment (Schutz-Bosbach & Prinz, 2007). However, little is known about how perception could interact with inhibition decisions. It is possible that these interactions are identical to those with the action system given that action and inhibition responses can be considered as symmetrical response alternatives (Mostofsky & Simmonds, 2008). However it is also possible that the inhibition system is unique in the way it guides or treats perceptual information. For example, inhibitions do not involve the direct physical manipulation of objects in the same way that actions do, and therefore the inhibitory system may not require perceptual monitoring with high temporal precision. The evaluation of perceptual evidence may then be driven in different ways by action and inhibition systems.

This chapter will address this question by presenting three experiments. The first will compare how we accumulate external evidence in order to make action vs inhibition decisions. The second experiment will test the possibility that differences between action and

inhibition systems in evidence accumulation could be driven by inherent differences in their capacity to directly and immediately affect events in the external world. The final experiment will examine how we perceive these external events and sensory outcomes caused by actions and inhibitions.

2.2. *Experiment 2.1: Evidence accumulation for action and inhibition decisions*

2.2.1. *Introduction*

Deciding between alternative responses in the real world involves weighing up the external evidence in favour of each option, before committing to a decision when we feel the evidence is strong enough. Decisions are often between whether to act or inhibit action, for example when deciding whether to cross the road immediately or to inhibit this impulse and wait longer. Inhibitions can be thought of as response options similar to actions (Mostofsky & Simmonds, 2008), and so decision-makers may accumulate evidence in favour of both action and inhibition alternatives. When we see a car approaching, we must evaluate several pieces of information in deciding whether or not to cross. Such evidence includes how fast the car is travelling, how far away it is, how fast we can travel across the road to the other side, and how much of a hurry we are in. After accumulating and evaluating this information, we can make our final decision and either act and cross immediately, or inhibit and cross in the future.

Decisions between action alternatives are well studied, with most experiments using a perceptual discrimination paradigm in which responses are leftward or rightward actions (e.g. Huk & Shadlen, 2005; Churchland et al., 2008; Hanks et al., 2006). Decisions between voluntary choices to act or inhibit on the other hand are less studied. The amount of external evidence available in making a decision to inhibit could arguably determine how ‘intentional’ inhibition is. With entirely ambiguous evidence, the decision to inhibit is more intentional and endogenous, while a clear instructional stimulus produces more stimulus-driven inhibition.

Although action and inhibition decisions can depend on accumulation of the same type of evidence, their respective effects in the external world can be considerably different. Actions immediately alter the world in some way: for example objects may move or the environment may change as a direct result of having performed the action. Inhibitions on the other hand are characterised by a lack of action and therefore a lack of any immediate change

in the world. Given this strong asymmetry in the outcomes of each response in the external world, there may also be differences in the decision-making processes that culminate in each response. For example, there may be a greater degree of caution associated with decision-making that will lead to action compared to inhibition. Regretting the choice to inhibit may only mean that a subsequent action is executed later than it could have been, but regretting action may be more problematic because it can produce an irreversible effect in the external world. More conservative decision-making may be expected when associated with action compared to inhibition.

Here, we directly compare action and inhibition in a dot motion discrimination task (Hanks et al., 2006). Participants view moving dots and must decide whether the dots are moving primarily left or right. While observing the dots, participants continuously move a mouse cursor around a screen and then respond by producing an action response (move faster) or an inhibition response (stop moving) depending on the perceived direction of motion. The number of dots moving coherently is manipulated between 40% and 0%. We therefore place action and inhibition on a continuum of instructed to voluntary: more reliable external evidence allows more instructed responses while more ambiguous evidence allows more voluntary responses. Unlike much previous work (e.g. Karch et al., 2009), we do not treat instructed and voluntary choice responses as independent processes, but rather consider them as points on a continuous scale of externally to internally driven behaviour. However, the presence of no consistent sensory evidence at all (i.e. 0% coherence) is formally equivalent to the classic ‘free choice’ condition in experiments that compare externally and internally generated actions (Jahanshahi et al., 1995). Figure 2.1. below illustrates this instructed-to-intentional continuum.

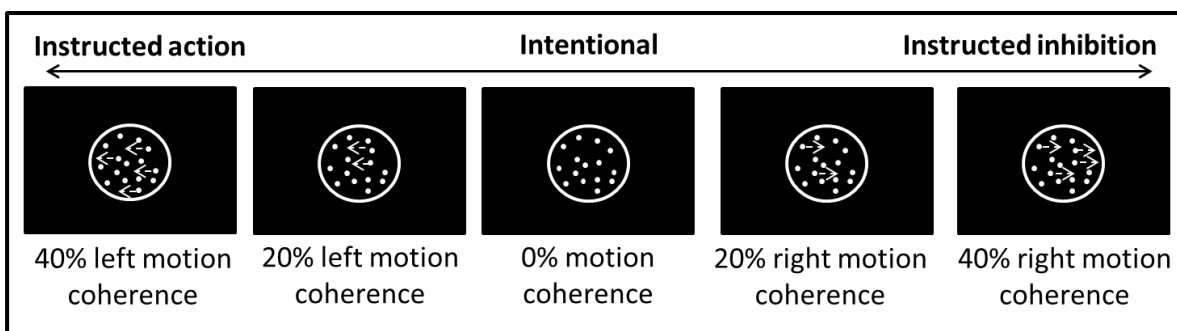


Figure 2.1. Schematic illustration of varying perceptual evidence strength for action and inhibition decisions. Perception of leftward motion indicates an action response and rightward motion indicates inhibition. Dots move randomly except for a possible subset in each trial between 40% and 0%. This dictates how ‘instructed’ a response is.

We examine the time taken to make a decision and investigate whether decision time is differently modulated for action and inhibition as external evidence becomes more ambiguous. If greater caution is associated with actions, the extent to which participants slow down their decisions for more ambiguous evidence should be greater than for inhibition.

2.2.2. Methods

2.2.2.1. Participants

Fifteen right-handed participants (7 female) took part in the experiment. One participant (female) could not maintain smooth cursor motion and therefore could not execute recordable responses in the task. Her participation was stopped early and no data was recorded. All participants were 18-40 years old. None had any history of psychological or neurological disorders. Ethical permission for the study was granted by the local ethics committee at the Institute of Cognitive Neuroscience.

2.2.2.2. Design

This experiment used the dot motion paradigm (e.g. Hanks et al., 2006). Moving dots were presented inside an enclosed circle in the centre of a computer screen. On each trial, the majority of these dots moved in entirely random directions. Depending on the condition, some dots in each trial could also coherently move horizontally left or horizontally right. There were three motion coherence conditions in each block of trials. 70 trials in each block were 0% coherence, in which 100% of dots moved in random directions. 70 trials in each block were 20% coherence, in which 20% of the dots on the screen moved in a coherent direction horizontally left (35 trials) or right (35 trials) while 80% were random. Finally, 70 trials in each block were 40% coherence, in which 40% of the dots on the screen moved in a coherent direction horizontally left (35 trials) or right (35 trials) while 60% were random. Conditions were randomised in order within each block.

2.2.2.3. Procedure

Participants were asked to begin moving a cursor around the screen using the mouse in a circular motion when they saw the word ‘START’. This start signal appeared in the centre of the screen for 1 second. Following this, a hollow circle appeared in the centre of the screen. Participants continued to move the mouse cursor around this circle. They were asked to maintain a steady cursor speed, which they learned during an initial training phase. If they moved faster than this pre-trained speed, a signal appeared saying ‘TOO FAST’, and the trial ended early. If they moved slower than this pre-trained speed, a signal appeared saying ‘TOO SLOW’ and the trial ended early. Speed was recorded by continuously sampling x and y coordinates for mouse position in a loop approximately every 16ms. When the average change in mouse position across the last 4 samples exceeded 60 pixels, this was considered too fast. When the average change in mouse position across the last 4 samples was below 10 pixels, this was considered too slow.

If participants maintained steady cursor speed and motion during the initial mouse motion period, moving dots were presented inside the hollow circle on the screen, after a random interval between 2 and 4 seconds. Dots remained on-screen for 2 seconds and participants were asked to decide on whether they were moving left or right. Participants performed two types of block. In one, they were asked to immediately inhibit moving the cursor and remain still if they thought the dots were coherently moving to the left (‘inhibition’ response), and to immediately initiate faster cursor motion if they thought the dots were coherently moving to the right (‘action’ response). The other block instructed the opposite responses for each direction. Block order was counterbalanced across participants.

Participants were instructed to make all responses within the 2 second interval while the dots were on the screen. Reaction times for actions were recorded as the time between the onset of the dots on the screen, and the immediate initiation of the speeded cursor motion (average change in mouse position across the last 4 samples exceeding 60 pixels). Reaction times for inhibitions were recorded as the time between the onset of the dots and the immediate stopping of cursor motion. After the dots disappeared, participants had to remain still or continue moving the mouse cursor rapidly depending on their response. After a random interval between 0.5 and 1 second, the words ‘End of trial’ appeared and the next trial began. If participants made no response the words ‘NO RESPONSE’ were presented, and if they accidentally performed both responses in the same trial, the words ‘BOTH RESPONSES MADE’ were presented. All stimuli were presented white on a black

background. The basic sequence of events in a trial can be seen below in figure 2.2. In order to motivate participants to perform the task, they were told there would be an additional £3 reward on top of their £12 payment if they performed well in the motion direction discrimination. All participants were paid £15 at the end of the experiment.

Participants also performed brief pre-experimental and post-experimental blocks in which there were no moving dots, and instead participants simply responded to a red (35 trials) or green (35 trials) signal appearing inside the centre circle. Trials were randomised in order. As with the main block trials, participants initiated steady and circular cursor motion around the screen upon seeing the ‘START’ signal. If the circle then filled with green, participants immediately initiated the action response with rapid cursor motion. If the circle filled with red, they immediately stopped moving the cursor and held it still for the inhibition response. These brief blocks were used in order to see whether there is a general difference in reaction times for actions and inhibitions in this particular experimental paradigm.

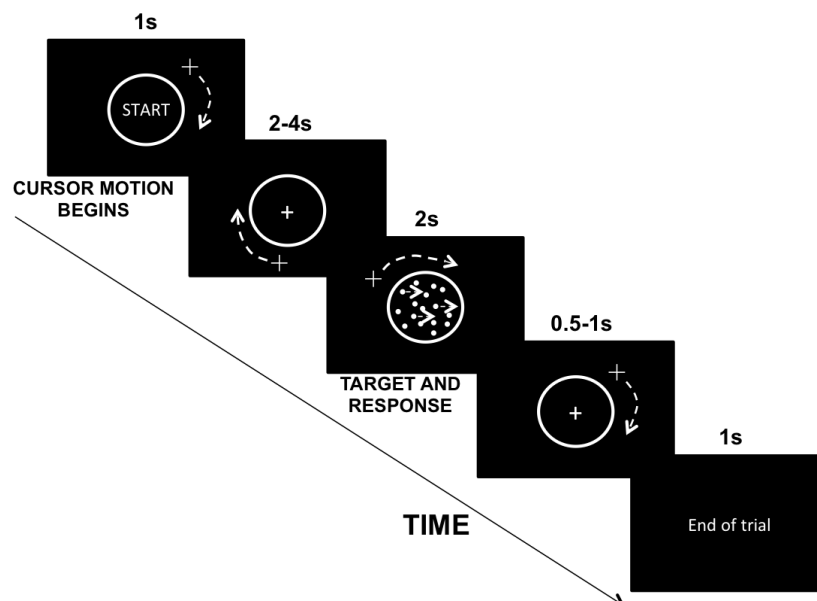


Figure 2.2. Basic sequence of events in a typical trial. After the dots appear, participants must decide whether there is any coherent motion to the left or right and execute the corresponding inhibition or action response. In this example they execute an action response by beginning to move the cursor rapidly and continuing this until the words ‘End of trial’.

2.2.3. Results

Figure 2.3 below shows reaction times for actions and inhibitions in the pre-experimental and post-experimental tasks. Action reaction times show how long it takes for the participant to initiate speeded cursor motion after the onset of the green signal. Inhibition reaction times show how long it takes for the participant to stop cursor motion after the onset of the red signal.




Figure 2.3. Reaction times (ms) for actions and inhibitions in response to green and red instruction stimuli respectively

Action reaction times were significantly slower than inhibition reaction times in the pre-experimental and post-experimental task blocks, $t(13) = 5.14$, $p < 0.001$. This suggests that inhibition responses in this particular paradigm may be slightly easier to initiate than action responses in terms of reaction time. Results in the experimental decision-making tasks should be interpreted in the context of this general difference between the responses.

Figure 2.4 shows response times in the experimental conditions during decision-making about dots with varying motion coherence.

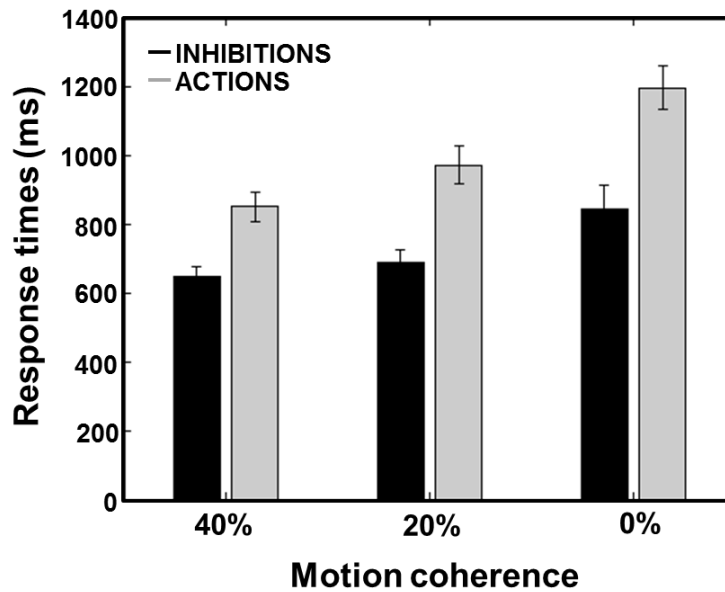


Figure 2.4. Response times (ms) for decisions about the direction of moving dots. Dots had 40%, 20%, or 0% motion coherence to the left or right.

Unsurprisingly, a 3x2 within-subjects ANOVA showed a significant main effect of motion coherence on reaction times, with participants taking significantly longer to respond to lower coherence – and therefore more ambiguous – dot motion stimuli, $F(2,26) = 37.24$, $p < 0.001$ (identical F and p -values with a Greenhouse-Geisser correction). There was also a main effect of response type, with actions being significantly slower than inhibitions, $F(1,13) = 44.5$, $p < 0.001$. This was also expected given that the pre and post-experimental tasks (see figure 2.3) showed a general responding advantage for inhibitions, suggesting that they may be easier to execute in the context of the present experiment. Interestingly however, a significant interaction was also found, $F(2,26) = 12.74$, $p < 0.001$. The slowing in response time as motion coherence decreased was much greater for actions than it was for inhibitions.

The greater caution in action-related responding was also reflected in the choice frequencies for 0% motion coherence data. Participants were significantly more likely to choose to inhibit (58.68%) than choose to act (41.32%), $t(13) = 2.95$, $p = 0.01$. Therefore, in addition to being slower to act than inhibit with more ambiguous evidence, participants were also less likely to choose to act.

We performed an additional diffusion-to-bound model analysis (Palmer et al., 2005; Ratcliff & McKoon, 2008) to investigate which decision-making parameters were most relevant for explaining the differences between action and inhibition. We used the DMATOOBOX (Vandekerckhove & Tuerlinckx, 2008), to study the relevance of two key parameters: drift rate and boundary (Ratcliff & McKoon, 2008). Our interest was to identify the independent contributions of these two parameters in explaining the behavioural effect from Figure 2.4. Therefore, rather than running a single model with both parameters varying simultaneously, we ran separate analyses in which each parameter was allowed to vary freely on its own while all other parameters were kept constant. We then compared how each individual parameter varied, and how closely this variation matched the behavioural response time data.

We started by investigating whether a difference in drift rate can explain the varying reaction time patterns for actions and inhibitions. Drift rate refers to the rate at which evidence is accumulated towards a decision boundary at which point a response is made, and it generally varies with stimulus strength or motion coherence (Palmer et al., 2005; Ratcliff & McKoon, 2008). We ran 2 models: a control model in which there were no free parameters available to vary, and a second model in which drift rate was allowed to vary (Vandekerckhove & Tuerlinckx, 2008). We modelled only successful responses and excluded any errors. Models were based on a 3x2 design (motion coherence x response type), and so we were able to obtain 6 separate drift rate parameter values for analysis. The Bayesian Information Criterion (BIC) was significantly better for model 2 (547.35) than model 1 (620.53) ($p = 0.003$), so we used drift rate values from this model in which drift rate could vary.

Figure 2.5 below shows drift rate values in each condition.




Figure 2.5. Drift rate values for actions and inhibitions across each motion coherence condition.

A 3x2 ANOVA returned a significant main effect of motion coherence, $F(2,26) = 20.78$, $p < 0.001$, and a significant main effect of response type, $F(1,13) = 46.26$, $p < 0.001$. Drift rate reduced with lower motion coherence and was consistently lower for evidence accumulation towards an inhibition response than an action response. However, there was no significant interaction, $F(2,26) = 0.2$, $p = 0.82$. This lack of an interaction suggests that drift rate cannot explain the interaction found for the behavioural response time data.

There is some evidence to suggest that decision boundaries may also vary over time within a trial, depending on motion coherence or uncertainty (Churchland et al., 2008; Wenzlaff et al., 2011). Therefore we ran a further diffusion model analysis, in an identical fashion to the previous drift rate analysis, except that the second model now used a free boundary parameter instead of a free drift rate parameter. Again, model 2 produced a significantly better BIC (559.82) than model 1 (620.53) ($p = 0.0003$), so we adopted this model for analysis of the free boundary parameter across conditions. Figure 2.6 below shows the boundary parameter value in each condition.




Figure 2.6. Boundary values for actions and inhibitions across each motion coherence condition.

A 3x2 ANOVA on the boundary values returned a significant main effect of motion coherence, $F(2,26) = 20.80$, $p < 0.001$ ($p < 0.001$ Greenhouse-Geisser corrected), and a significant main effect of response type, $F(1,13) = 56.29$, $p < 0.001$. However, unlike with the drift rate parameter values, the boundary values also returned a significant interaction between motion coherence and response type, $F(2,26) = 8.19$, $p = 0.002$ ($p = 0.007$ Greenhouse-Geisser corrected). The pattern of results looks strikingly similar to the response time patterns (see figure 2.4).

2.2.4. Discussion

This experiment demonstrated a clear difference between the decision to act and the decision to inhibit. Participants were consistently slower to choose to act than inhibit. This may have been an artefact of the way in which people expressed actions and inhibitions in the current task – perhaps the threshold used to define a speeded action response meant that it took longer for a response to be recorded. However, a more interpretable finding is that the difference in response times between actions and inhibitions increased as external evidence became weaker and more ambiguous. This interaction cannot be explained simply by a

difference in how the responses are recorded, but rather must be interpreted as an increasing difference in decision time and caution involved in making a decision to act or inhibit. Although both actions and inhibitions slowed across the 40% to 0% motion coherence conditions, actions slowed to a much greater extent than inhibitions. This suggests that actions may be associated with a greater level of caution in real-world behaviour, perhaps because actions directly and often irreversibly change the external world, while inhibitions do not. An action that is regretted after execution is likely to have more dramatically negative consequences than a regretted inhibition, so decisions to act are made more carefully and require a greater level of certainty.

The diffusion-to-bound analysis aimed to identify a decision-making parameter that could explain the response time difference found between actions and inhibitions. Drift rate, which reflects the rate of evidence accumulation, typically increases with greater stimulus strength or motion coherence (Ratcliff & McKoon, 2008). Clearer stimuli in the environment allow evidence to be gathered at a higher rate towards a faster decision. Expectedly, we found that drift rate decreased with weaker motion coherence (Ratcliff & McKoon, 2008). However, the drift rate values did not show an interaction between motion coherence and response type, and therefore could not explain the response time effects. We then explored the boundary parameter, and found that the pattern this produced was very consistent with the response time data. The boundary increased across motion coherence levels for both action and inhibition, but the increase for actions was significantly greater. The boundary parameter is typically not considered to vary with motion coherence but some previous models have proposed conditions under which the boundary could vary within a trial (Churchland et al., 2008; Wenzlaff et al., 2008).

Our data suggests that the boundary on a decision may be sensitive to the interaction between environmental uncertainty and the type of response being executed. When external evidence is highly ambiguous (e.g. 0% motion coherence), the boundary on an action response may be set very high to ensure that a large amount of evidence is accumulated before an action decision is expressed. This minimises the chance that an error is made (Ratcliff & Rouder, 1998). With inhibitions on the other hand, errors may be less costly, and so the boundary is less sensitive to change and does not shift up much with weaker evidence. The boundary on a decision is similar to the ‘criterion’ in signal detection theory (Green & Swets, 1974). When attempting to detect a noisy stimulus, the criterion can be adjusted to determine the number of hits and false alarms in detection. With a low criterion and therefore

many ‘yes’ responses when detecting stimuli, there are a large number of hits when the stimulus is present but also a large number of false alarms when the stimulus is absent. With a more conservative criterion, false alarms are less frequent as a stimulus must be detected far more clearly before a ‘yes’ response is made, but weaker stimuli are also frequently missed. This is analogous to the decision boundary in the current experiment in which a lower boundary means faster responses but a higher risk of errors. The criterion may also be important for conscious perception and determines whether a stimulus is reported as consciously seen or simply guessed (Lau, 2007; Lau & Passingham, 2006). The boundary on a decision may also vary with caution and uncertainty, as decisions require more or less confidence about conscious awareness of sensory evidence before a decision is expressed.

Experiment 1 suggests that one of the major differences between intentional action and intentional inhibition is the caution with which decisions and responses are made. We propose that the greater caution associated with action than inhibition is driven by the greater tendency for actions to directly change the environment. When these changes caused by action happen to be disadvantageous, it is often impossible to reverse them. It therefore serves us well to have an action system which is a conservative decision-maker, and an inhibition system which is capable of making faster but lower risk decisions. Probabilities of action choices are indeed very sensitive to perceived risk in a decision-making task (Dror et al., 1999). Increased neural activity in the nucleus accumbens underlies excessively risky action decisions, while increase anterior insula activity predicts excessively risk-averse action decisions (Kuhnen & Knutson, 2005). This kind of risk variability is typically not associated with inhibition decisions and therefore inhibition decisions may be less sensitive to potential errors in decision-making.

2.3. *Experiment 2.2: Manipulating perceptual outcomes of action and inhibition decisions*

2.3.1. *Introduction*

Experiment 1 proposed that the differences between action and inhibition in decision-making were due to their intrinsic differences in producing direct and immediate effects in the external world. In the following experiment, this interpretation is tested more directly, by manipulating the external effects that follow from actions and inhibitions. By creating artificial situations in which actions and inhibitions unusually produce similar external consequences, we can test the prediction that differences in decision-making between action and inhibition will disappear when the effects generated by each response are matched.

2.3.2. *Methods*

2.3.2.1. *Participants*

Sixteen right-handed participants (8 female) took part in the experiment. One participant (female) was a significant outlier in their reaction time data pattern for the difference between actions and inhibitions across motion coherence levels, as identified by a Grubbs test ($p < 0.05$). They were also >2 standard deviations away from the mean across participants and were therefore removed leaving fifteen participants. All participants were 18-40 years old. None had any history of psychological or neurological disorders. Ethical permission for the study was granted by the local ethics committee at the Institute of Cognitive Neuroscience.

2.3.2.2. *Design and Procedure*

The design and procedure were broadly identical to that in Experiment 1 (see Section 2.2.2.2.). The first critical difference was that in this experiment there were 4 blocks in total instead of 2. Each block contained 50 trials that were 0% coherence, 50 trials that were 20% coherence, and 50 trials that were 40% coherence. In the 40% and 20% coherence trials, the dots on the screen were moving left in 25 trials and right in 25 trials. In 2 of the 4 blocks, participants performed the action response for leftward motion (leftward-action blocks) and inhibition response for rightward motion. The opposite was true for the remaining blocks (rightward-action blocks). Motion coherence conditions were randomised in order within each block. The order of the blocks was counterbalanced across participants. There were no

additional experimental blocks with red and green signals (see figure 2.3), but there was an initial practice block to get participants familiar with the task.

The additional manipulation in this experiment was that either all actions or all inhibitions in each block were followed by an immediate visual cue. In 1 of the 2 rightward-action blocks, action responses were immediately followed by a blue signal, which filled the circle in the centre of the screen as soon as the action response was detected. In the other rightward-action block, inhibition responses were immediately followed by the blue signal. The same applied for the leftward-action blocks. The blue signal remained on-screen until the end of the trial, and was designed to represent an immediate external consequence produced by participants' responses. This stimulus had no instructional or informational content and did not provide positive or negative feedback about whether a response was correct. It was entirely task-irrelevant and acted as an immediate 'response recorded' confirmation signal. The new stimulus sequence can be seen in figure 2.7 below.

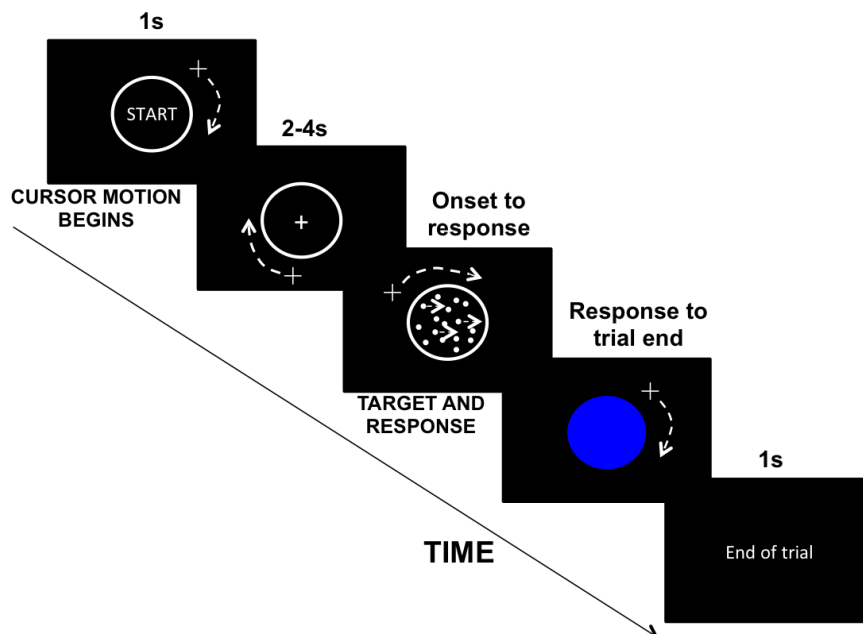


Figure 2.7. Basic sequence of events in an example trial. After the dots appear, participants must decide whether there is any coherent motion to the left or right and execute the corresponding inhibition or action response. In this example they execute an action response by beginning to move the cursor rapidly. A blue stimulus appears immediately following their response. They continue the response until the words 'End of trial'.

2.3.3. Results

Figure 2.8 below shows response times in all conditions for the responses that were *not* followed by a blue signal in each block (left), and for the responses that were followed by an immediate blue signal in each block (right). The results on the left should be directly comparable to those in figure 2.4 as the responses were identical.

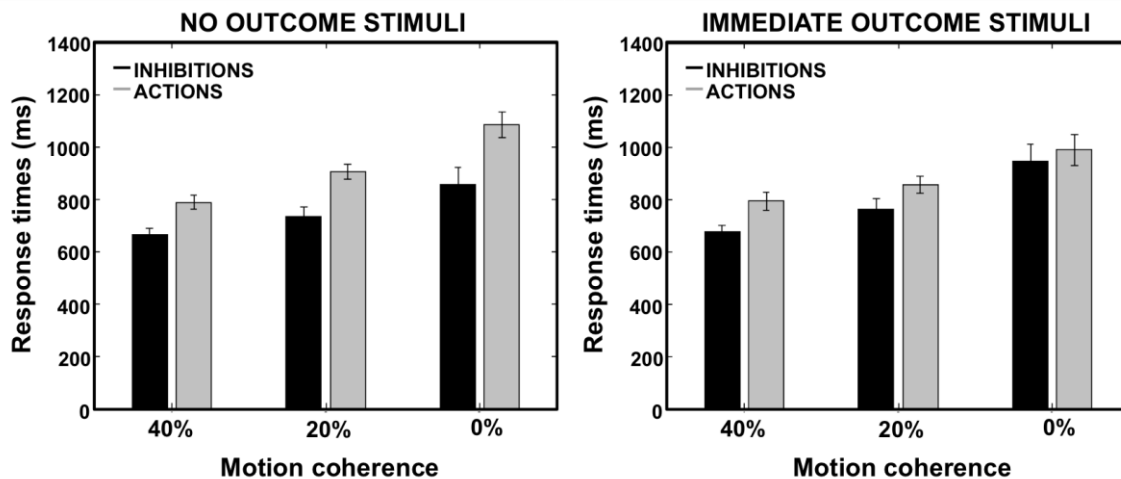


Figure 2.8. Response times (ms) for decisions about the direction of moving dots. Dots had 40%, 20%, or 0% motion coherence to the left or right. Responses for the graph on the left did not result in the appearance of any external stimulus, while responses on the right produced an immediate outcome stimulus.

A 3x2 ANOVA on the data with no outcome stimuli (left graph in figure 2.8) showed a significant main effect of motion coherence on response times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 22.81$, $p < 0.001$ ($p < 0.001$ Greenhouse-Geisser corrected). There was also a main effect of response type, with slower actions than inhibitions, $F(1,14) = 26.26$, $p < 0.001$. There was also a significant interaction, $F(2,28) = 6.56$, $p < 0.001$, with actions slowing more dramatically than inhibitions with decreasing motion coherence. These results are statistically identical to those from Experiment 1 (see section 2.2).

An identical analysis on the responses that produced an immediate outcome stimulus (right graph in figure 2.8) showed a significant main effect of motion coherence on response

times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 31.18$, $p < 0.001$ ($p < 0.001$ Greenhouse-Geisser corrected). There was also a main effect of response type, with slower actions than inhibitions, $F(1,14) = 15.32$, $p = 0.002$. No significant interaction was found, $F(2,28) = 2.20$, $p = 0.13$.

Choice biases were also interesting. In blocks where actions were followed by immediate outcome stimuli, participants were significantly more likely to choose to act (56.97%), than choose to inhibit (43.03%), $t(14) = 2.57$, $p = 0.02$. However, when inhibitions were followed by the outcome stimuli, there was no significant choice bias as participants were equally likely to choose to act (52.26%) and inhibit (47.74%), $t(14) = 0.63$, $p = 0.54$. It therefore seems that there is a reward value to the sense of agency over an external sensory outcome, but *only* when this relates to our actions. This is consistent with the idea that actions have stronger effect-producing qualities in the real world than inhibitions.

We further analysed the data using the same diffusion-to-bound modelling procedures as experiment 1. Figure 2.9 below shows the results for the drift rate parameter.

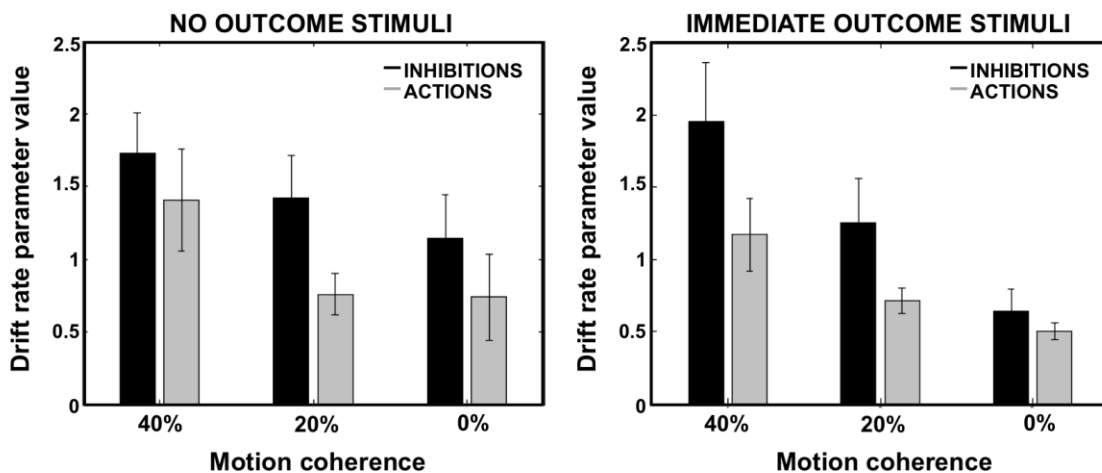


Figure 2.9. Drift rate values for actions and inhibitions across each motion coherence condition. Responses for the graph on the left did not result in the appearance of any external stimulus, while responses on the right produced an immediate outcome stimulus.

A 3x2 ANOVA on the drift rates from responses with no outcome stimuli (left graph in figure 2.9) showed a significant main effect of motion coherence on response times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 7.35$, $p = 0.003$. There was also a main effect of response type, with slower actions than inhibitions, $F(1,14) = 7.43$, $p = 0.02$. However, there was no significant interaction, $F(2,28) = 0.56$, $p = 0.58$.

Drift rates from responses producing an immediate outcome stimulus (right graph in figure 2.9) showed a significant main effect of motion coherence on response times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 12.23$, $p < 0.001$ ($p < 0.001$ Greenhouse-Geisser corrected). There was also a main effect of response type, with slower actions than inhibitions, $F(1,14) = 7.47$, $p = 0.02$. No significant interaction was found, $F(2,28) = 2.01$, $p = 0.15$.

Figure 2.9.1 displays the data for the boundary parameter.

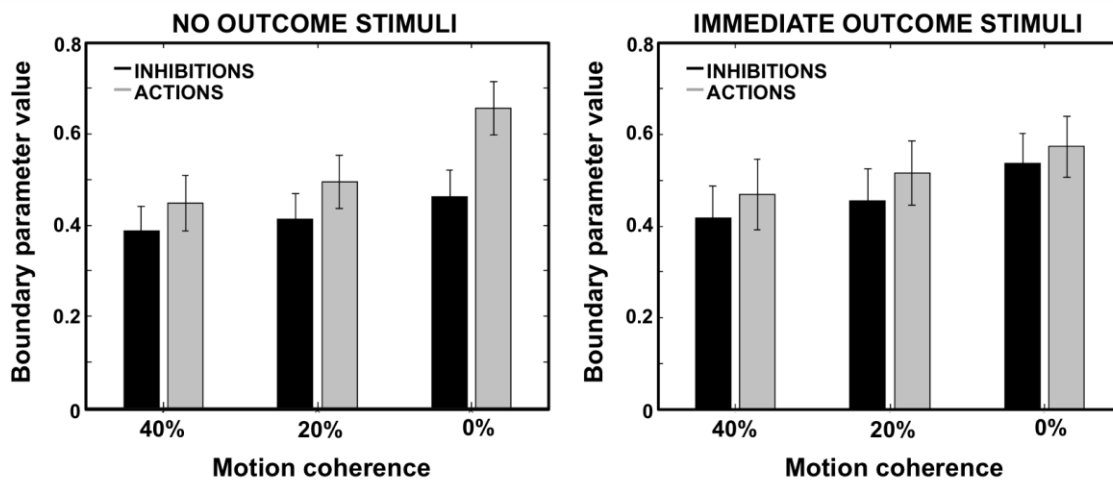


Figure 2.9.1 Boundary parameter values for actions and inhibitions across each motion coherence condition. Responses for the graph on the left did not result in the appearance of any external stimulus, while responses on the right produced an immediate outcome stimulus.

A 3x2 ANOVA on the boundary values from responses with no outcome stimuli (left graph in figure 2.9.1) showed a significant main effect of motion coherence on response times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 14.22$, $p < 0.001$ ($p = 0.002$, Greenhouse-Geisser corrected). There was a significant main

effect of response type, with slower actions than inhibitions, $F(1,14) = 17.30$, $p < 0.001$. A significant interaction was also found, $F(2,28) = 4.43$, $p = 0.02$ ($p = 0.047$, Greenhouse-Geisser corrected).

Boundary values from responses producing an immediate outcome stimulus (right graph in figure 2.9.1) showed a significant main effect of motion coherence on response times, with participants taking significantly longer to respond to lower motion coherence, $F(2,28) = 19.78$, $p < 0.001$ ($p < 0.001$ Greenhouse-Geisser corrected). There was also a main effect of response type, with slower actions than inhibitions, $F(1,14) = 18.17$, $p < 0.001$. No significant interaction was found, $F(2,28) = 0.44$, $p = 0.65$.

When comparing the outcome and no outcome responses directly in a $3 \times 2 \times 2$ ANOVA, a significant three-way interaction is found, $F(2,28) = 5.05$, $p = 0.005$ ($p = 0.03$ Greenhouse-Geisser corrected). This confirms that the motion coherence \times response type two-way interaction patterns independently found in the two outcome conditions differ significantly.

2.3.4. Discussion

Experiment 2 replicated the effect from experiment 1. ‘Normal’ action and inhibition responses without an immediate external outcome added, resulted in greater slowing of action decisions than inhibition decisions across decreasing motion coherence levels. We wanted to test whether an important factor underlying this asymmetry was the natural associations of action and inhibition responses with immediate sensory consequences in the external world. Actions tend to produce immediate external consequences while inhibitions prevent action and therefore generally do not alter the world. We introduced immediate external outcomes following each response, and found that this removed the asymmetry in response times. When comparing action responses that produced a sensory consequence with inhibition responses that produced the same sensory consequence, the response times for each response slowed by the same amount during more difficult decisions.

Consistent with experiment 1, we also found that the boundary parameter rather than the drift rate parameter in the diffusion-to-bound model was capable of explaining the relationship between action and inhibition response times. As external evidence became more ambiguous, the boundary parameter was shifted higher for action decisions than inhibition

decisions, when the outcomes of each were not artificially altered. When we introduced the blue external stimulus as an immediate causal consequence of each response, both the differences in response time patterns and the differences in boundary parameter patterns disappeared. This suggests that a key factor driving the asymmetry in decision-making between action and inhibition is the capacity of each response to produce direct and immediate changes in the environment.

The findings from experiment 2 are directly relevant to the concept of agency, which describes the feeling of being in control of our actions and the consequences they produce in the world (Haggard & Tsakiris, 2009). The results suggest that action and inhibition differ in how they affect external events, which may mean that they also produce different feelings of agency. Although there is a strong understanding of the relationship between action and agency (Chambon et al., 2013; Moore et al., 2010; Wenke et al., 2010), little is known about how inhibition relates to agency. Therefore, experiment 3 below aims to examine the agentic qualities of inhibition in more detail, in order to identify how it is unique from action.

2.4. Experiment 2.3: Inhibition and agency over sensory outcomes

2.4.1. Introduction

Experiment 2 investigated how the manipulation of sensory outcomes following actions and inhibitions can directly influence how people use external evidence in their decision-making. It supported the idea that under natural circumstances, action and inhibition are inherently asymmetric in their capacity to affect the external world, and that this information biases decision-making. Experiment 3 will explore this asymmetry in more detail, by directly measuring temporal judgements of external events that are caused by actions and inhibitions. It will also compare the subjective feelings of agency associated with the two responses.

2.4.2. Methods

2.4.2.1. Participants

Sixteen right-handed participants (10 female) took part in the experiment. One participant (female) could not maintain steady cursor motion to perform the task and therefore no data was collected. This left fifteen participants for analysis. All participants were 18-40 years old.

None had any history of psychological or neurological disorders. Ethical permission for the study was granted by the local ethics committee at the Institute of Cognitive Neuroscience.

2.4.2.2. Design

Participants performed two experimental blocks. In one block, they executed action responses while in the other they executed inhibition responses. For action responses, participants gave a computer mouse a single push. In the inhibition block, participants began circular cursor motion around the screen before the trial started and then abruptly stopped mouse motion to respond during the trial. Following action and inhibition responses, participants received a visual stimulus on the screen that could be immediate or delayed. We measured participants' detection abilities for these delays, as well as their subjective feelings of control over the onset of the stimulus. We compared this data across the action and inhibition responses.

2.4.2.3. Procedure

In action blocks, participants rested their right index finger on the right side of the mouse upon seeing the word 'START' on the screen. When they were ready, they pressed the spacebar on the keyboard with their left hand to initiate the trial. When they did this, a hollow white circle appeared in the centre of the screen against the black background. Participants were instructed to wait at least 1 second after initiation of the trial, and then to give the mouse a single brief push away from their hand with the index finger whenever they wanted to. If they did this too early, they received an 'ERROR' signal, and the trial was replaced at the end of the block. This was to avoid accidental or unprepared responses. Following a successful response, the circle in the centre of the screen was illuminated with a white fill, either immediately or delayed between 50ms to 500ms (in 50ms increments) after the response. This illuminated circle remained on-screen for between 1 and 2 seconds. Participants were asked to remain relaxed during this time and ensure they did not move the mouse again. If they did move the mouse, this also produced an 'ERROR' signal and replacement of the trial. After successful responses, an 'End of trial' signal was presented for 1 second. Participants were then asked two questions. They were first asked "How much did you feel in control of the light?". Participants responded by moving a slider along a continuous scale from 0 – 100% using their right hand on the mouse. They confirmed their answer by clicking the mouse button. A second question then appeared: "Was the light delayed?". They answered

this by responding with '1' on the keyboard number pad for 'yes', and '2' on the number pad for 'no'. After both responses were made, the next trial began.

Inhibition blocks were broadly identical to action blocks, except for the way in which the response was made. Rather than remaining at rest, participants began moving the mouse in a circular motion upon seeing the word 'START'. They then pressed the spacebar with their left hand when they were ready while continuing mouse motion. After initiating the trial, they were instructed to continue moving for at least 1 second, before abruptly stopping mouse motion whenever they wanted after that (responding too early also produced an 'ERROR' signal just as with action responses). This abrupt stop produced the illumination stimulus either immediately or delayed as described above for actions. During this time, participants were asked to ensure that their hand remained perfectly still to avoid any kind of motion of the mouse. If any motion was detected before the end of the trial, participants received an 'ERROR' signal on the screen. After successful responses, participants received the same questions as described above for actions.

Each of the possible illumination stimulus delays between 0ms and 500ms (0ms, 50ms, 100ms, 150ms, 200ms, 250ms, 300ms, 350ms, 400ms, 450ms, 500ms) were equally represented in each block. There were 18 trials for each delay. This meant there were 198 trials per action and inhibition block (11 delays with 18 trials each). Block order was counterbalanced across participants and order of delay presentation was randomised within each block. Figure 2.9.2. below displays the basic structure of a trial in each block.

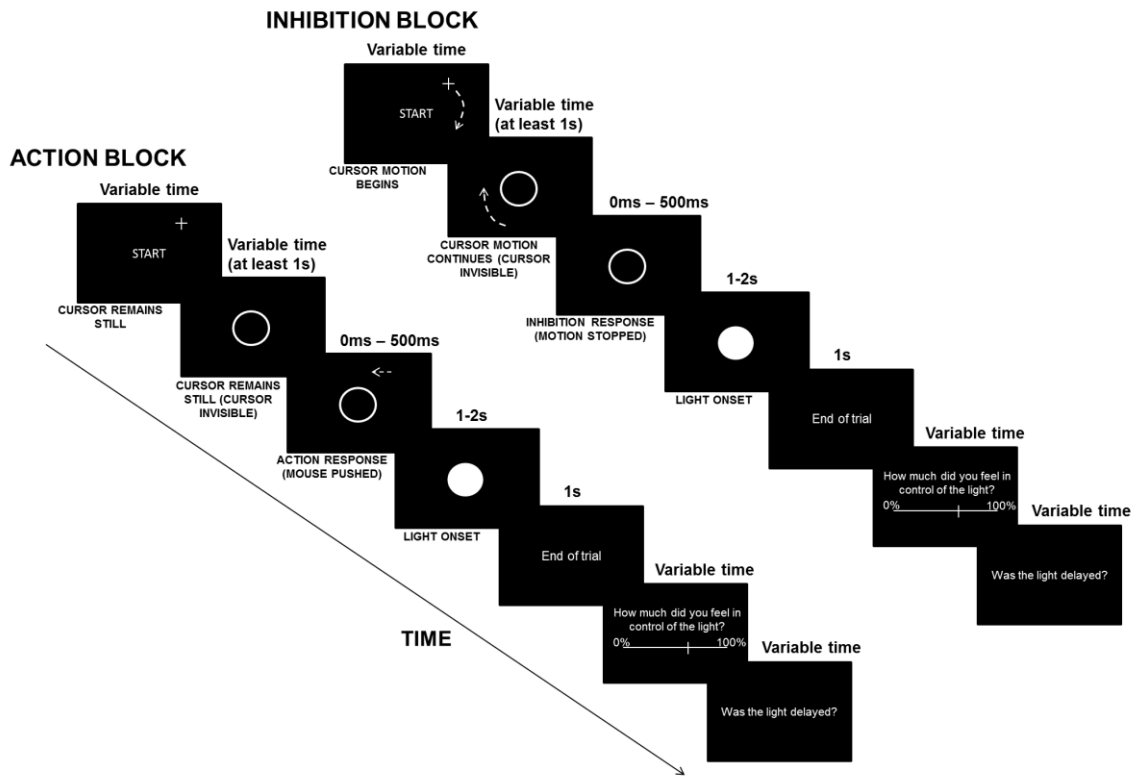


Figure 2.9.2. Basic structure of trials in action and inhibition blocks.

2.4.3. Results

Figure 2.9.3 below shows the data for participants' perceptions of whether a delay was present in the appearance of the illumination stimulus following responses.

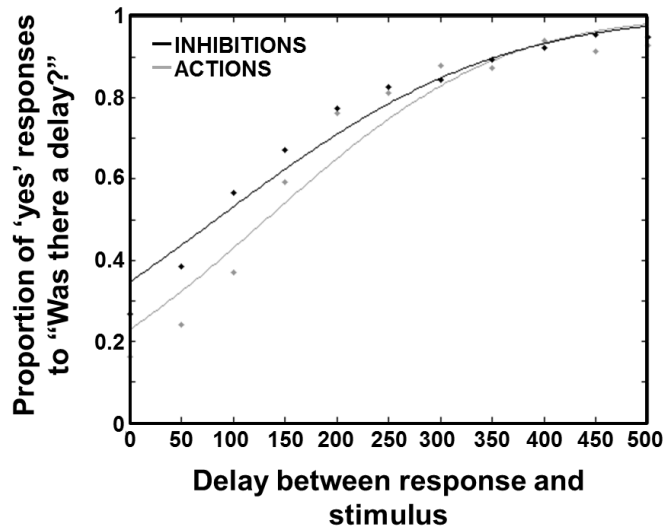


Figure 2.9.3. Proportion of 'yes' responses when asked whether a delay was present between the executed response (action or inhibition) and the subsequently appearing illumination stimulus. Psychometric curves (probit) are fitted to the data.

Analysis of the psychometric curves from figure 2.9.3 suggested no significant difference in the slopes between action and inhibition, $t(14) = 1.43$, $p = 0.17$. However a significant difference in the delay at which 50% of responses were reported as 'yes' was found, $t(14) = 3.15$, $p = 0.007$. Participants were more likely to report a delay for inhibitions than actions at the shorter delays between response and stimulus. In particular, the false alarm rate for inhibitions (28%) (i.e. the proportion of 'yes' responses at 0ms delay) was significantly higher than the false alarm rate for actions (17%), $t(14) = 2.45$, $p = 0.03$.

D-prime (d') measures were also analysed and compared between action and inhibition. Any false alarm/hit rates with values of 0 or 1 were altered to 0.06 ($1/N$) and 0.94 ($(N-1)/N$) respectively to allow z-transforms for d' calculations (N being the total number of trials in each delay condition). Figure 2.9.4 below displays these d' values.

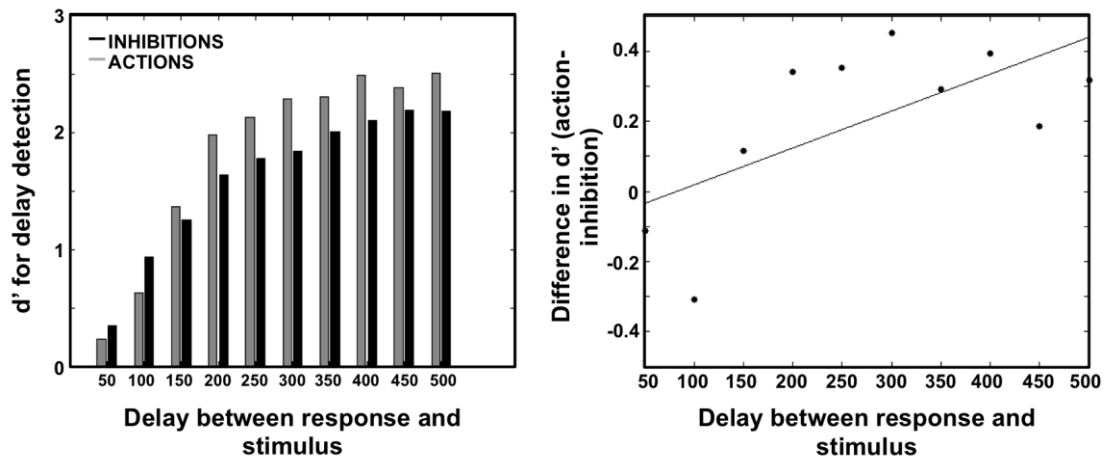


Figure 2.9.4. *d'* values for detecting delays between response and subsequent illumination stimulus (left), and the difference in *d'* between actions and inhibitions (right).

Unsurprisingly, detection performance for delays improved as delays got longer. The difference in *d'* between actions and inhibitions was more interesting however. The right panel in figure 2.9.4 shows the difference between actions and inhibitions increasing across the delay intervals. A t-test on the slope values for each participant shows that the positive slope is significantly greater than zero, $t(14) = 2.43$, $p = 0.03$. As delay intervals got longer, *d'* for actions improved more rapidly and to a greater extent than for inhibition. The *d'* values also show some sign of a plateau at the largest intervals, and perhaps an inverted U-shape function developing for the difference between actions and inhibitions, as *d'* reaches ceiling for both responses.

Subjective judgements of agency or 'control' over the illumination stimulus onset were also analysed. Figure 2.9.5. below displays mean agency ratings and the difference in ratings between actions and inhibitions.

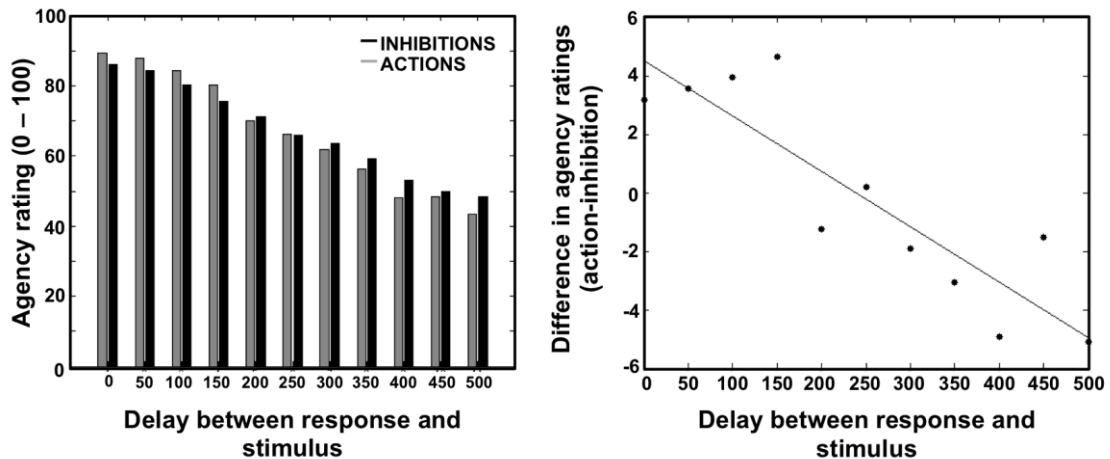


Figure 2.9.5. Subjective ratings of agency towards illumination stimulus (left), and difference in agency ratings between actions and inhibition (right).

Feelings of agency or control over the illumination stimulus reduced as the delay between response and stimulus grew longer. The differences between actions and inhibitions showed an interesting pattern: at smaller delays, participants felt they had more agency over stimuli following actions, but at longer delays, this pattern reversed and inhibitions produced stronger judgements of agency. The negative slope in this pattern of differences (right panel in figure 2.9.5) was significantly different to zero, $t(14) = -2.20$, $p = 0.045$. The reversal in agency advantage between action and inhibition responses seemed to happen at a delay of over 250ms (left panel in figure 2.9.5).

2.4.4. Discussion

Experiment 3 explored the concept of inhibition-related agency in more detail. It demonstrated that there are clear differences in how external events are perceived when they are a product of our actions compared to our inhibitions. First, delays in the appearance of expected external events are generally more precisely detected when they follow our actions than our inhibitions. In the present experiment, this was primarily due to a high false alarm rate following inhibitions, in which participants believed sensory outcomes were delayed when in fact they were not. Second, our feelings of control over external events display an interesting pattern: stronger agency ratings are reported for actions than inhibitions at small delays, but weaker agency ratings are reported at long delays.

In agreement with experiments 1 and 2, the reversal in agency ratings for actions and inhibitions could be explained by a natural asymmetry in how each system relates to sensory outcomes in the external world. Actions tend to produce immediate effects without delays in sensory consequences. Therefore, the action system may be particularly sensitive to detecting immediate delays in external events, and may also have an enhanced sense of control for events that occur immediately or very soon after action. Inhibitions on the other hand tend not to produce immediate effects in the external world. The inhibition system may therefore be less precise than the action system in detecting whether external events are delayed or not. However, it may also have an enhanced sense of control over events that are not immediate, due to its natural association with longer-term consequences rather than immediate consequences in the external world.

The findings also suggest an interesting relationship between delay detection abilities and subjective agency judgements towards external objects. Although actions showed a relatively consistent advantage over inhibitions in detection of objective delays, there was no consistent advantage in agency judgements. This suggests that although delay detection and agency are likely to be related (agency decreases as delay detection increases), delay detection alone cannot explain the difference in agency judgements between actions and inhibitions. There are other factors implicit to action and inhibition systems that drive asymmetric feelings of agency over external events. We propose that these factors arise from the everyday interactions of each system with the external world; namely the ability of the action system to cause immediate external effects, and the inhibition system's longer-term focus of effect.

The temporal characteristics of agency can be studied in relation to action with what is known as 'intentional binding' (Haggard, Clark, & Kalogeras., 2002). This effect shows that the perceived time between an action and a subsequent sensory event is shorter when that sensory event is caused by a participant's intentional action than when it is not. This temporal attraction between the events is a strong and robust marker of agency (Humphreys & Buehner, 2009). Interestingly, there is a trend towards temporal repulsion of a sensory event that follows inhibition rather than action (Haggard, Poonian, & Walsh, 2009). This emphasises that the agency associated with inhibition is very different to the agency associated with action. While action responses have a natural affinity to immediate external events that are perceived as being caused by action, the inhibition system fails to produce any subjective linkage between an inhibition response and subsequent sensory events. In fact,

inhibitions may have a natural association with *not* producing sensory consequences due to their prevention of action. This could explain the temporal repulsion rather than binding effect between inhibitions and events that follow inhibitions (Haggard et al., 2009).

2.5. *General discussion*

The experiments in this chapter have explored the relationship between perceptual judgements and intentional inhibition. Experiments 1 and 2 used a dot motion decision-making paradigm to investigate differences in how action and inhibition systems accumulate perceptual evidence in order to make a decision. They showed an asymmetric process, in which actions accumulate evidence more cautiously than inhibitions, and set a higher boundary on the accumulation process when perceptual evidence is weak. This was due to a general difference between the two systems in the types of outcome that they generate in the real world. Actions immediately and often irreversibly change the world, and therefore require a greater amount of certainty before a decision is expressed.

Experiment 3 investigated the processing of action and inhibition outcomes in more detail. It contributed additional support for the idea that action and inhibition systems are asymmetric in the types of outcomes that they naturally produce in the external world. It showed distinct processing of external outcomes following action and inhibition, in which delayed action-outcomes were more precisely detected overall. It additionally demonstrated a stronger feeling of control over action outcomes than inhibition outcomes with very small delays in outcomes, but a stronger feeling of control over inhibition outcomes at longer delays.

The experiments in this chapter are relevant to two key areas of existing research. The first is decision-making. The dot-motion task is typically used to investigate decisions between two action alternatives (Gold & Shadlen, 2007; Hanks et al., 2006; Roitman & Shadlen, 2002). In most cases, participants make leftward or rightward saccades or manual responses depending on the perceived direction of moving dots. This creates a balanced decision process between the two alternatives and each response is considered symmetrical. In experiments 1 and 2 in this chapter, two asymmetrical responses were competing with each other: intentional action vs intentional inhibition. This allowed us to identify critical differences between action and inhibition systems in relation to decision-making processes. Identical external evidence was accumulated in two contrasting ways, depending on whether

the ultimate response was action or inhibition. Action decisions required a greater level of evidence accumulation than inhibition decisions.

The second relevant area of research is in agency. Agency is usually studied in the context of action (David et al., 2008; Chambon et al., 2013; Farrer & Frith, 2002; Farrer et al., 2003; Wenke et al., 2010). Experiment 3 in this chapter considered the possibility of inhibition rather than action creating a sense of agency over an outcome. The feeling of agency that accompanied inhibition was shown to be very distinct from that which accompanies action. There are real-world examples of inhibition causing an immediate sensory outcome. One example might be when we inhibit hitting a cancel key on a keyboard and see an immediate download bar begin. Experiment 3 suggests that this outcome would be processed with less temporal precision than if it had followed action. An immediate inhibition outcome also returns a lower sense of control than if we had pressed a key to initiate the download. Overall however, immediate sensory outcomes following inhibition are a much less familiar event in the natural world than action-driven outcomes, because inhibitions usually prevent an immediate outcome from happening. This is likely to explain the distinct perceptual processing of external events found for each response.

The results presented in this chapter are relevant to real world behaviour and therefore present some practical implications. There are many situations in which we must decide between acting now or inhibiting action. When we cross the road, we decide between running immediately or waiting for an approaching car to pass first. When students revise for exams, they often need to decide between revising now or inhibiting this intention to act and postponing until later. This chapter has demonstrated that our decision-making is naturally asymmetrical for these two responses, and therefore we may be implicitly biased in decisions between action and inhibition. Given that actions are typically made more cautiously than inhibitions (see experiments 1 and 2), we may often be biased towards inhibition. In the case of crossing the road, this may be a beneficial bias as it prevents the possibility of a premature action, but in other cases it can lead to a cycle of procrastination. In these disadvantageous situations, it could be useful to reframe a problem as action vs action to provide a more balanced decision process: for example, rather than asking “should I revise now or not?”, students should ask “should I revise for module A or module B?”. In law, criminal responsibility can be assigned for both actions and omissions of actions that result in illegal activity (Fincham & Jaspars, 1980). However, the asymmetry between action and inhibition

responses as identified in the present experiment should be considered in ascribing such responsibility.

Chapter 3. Effects of sequential dependence and the gambler's fallacy

This chapter investigates whether 'free' action and inhibition decisions are influenced by prior sequences of events. Experiments 3.1-3.3 studied the competing effects of the 'gambler's fallacy' and habitual 'automatic activation' on sequence-embedded free choices to act or inhibit. Participants were vulnerable to a 'switching bias' similar to the gambler's fallacy in their cognitive decision-making. Following longer runs of repeated events, they were more likely to select an opposite novel response when given a free choice. Interestingly however, runs of repeated actions also produced a habitual bias to continue acting when given the choice. The same habitual responding bias was not present following runs of inhibitions. Experiment 3.4 investigated the effects of subliminal priming on these sequential dependence effects. It found that priming can affect habitual automatic activation effects, but not the higher level cognitive effects of switching biases. Overall, chapter 3 demonstrates that intentional inhibition decisions are not necessarily 'one-off' decisions independent of prior events. It highlights the important role that sequential effects may play in free decisions.

3.1. Introduction

Chapter 2 examined aspects of perceptual decision-making in the context of intentional inhibition. This chapter will continue to look at decision-making but from a different perspective. Rather than looking at the accumulation of immediate low-level perceptual evidence for inhibition decisions, the experiments presented here investigate the role of previous historical events in influencing a current decision. All of the experiments in this chapter will look at the biasing influence of a cognitive fallacy known as the 'gambler's fallacy'. By studying this, we can begin to understand the relevance of higher-level cognitive factors in driving decisions to intentionally inhibit actions.

The cognitive factors that influence choices to inhibit have been relatively little studied. 'Ego depletion' (Baumeister, Vohs, & Tice, 2007; Muraven & Baumeister, 2000) is one theory that attempts to explain failures of self-control. However, there are likely to be many other cognitive factors that are relevant. In the real world, failures to withhold action can simply be the product of bad decision-making, without any obvious prior depletion of self-control resources.

Behavioural choices naturally depend on previous action history. Repeated actions may become habits, triggered automatically by context or other Pavlovian stimuli, even when they do not fulfill current goals. Recent theories of action control distinguish between iterative flow of successive actions triggered by sensorimotor contexts, and control of discrete goal-directed actions involving volition and planning (Redgrave et al., 2010). Moreover, with repetition, the same action could shift from being voluntary and goal-directed, to being habitual and contextually-triggered (Skinner, 1971). Thus, voluntary action may involve a form of internal generation, which escapes from repetitive routine behaviour.

Sequential dependence between repeated actions has recently been investigated using a laboratory paradigm designed by Perruchet et al. (2006). They gave participants an auditory tone, followed by a cue, which required action on a random subset of trials. Participants were additionally asked how likely they thought an action cue was to appear following the next tone. The data were analysed according to the length of the preceding run of cues requiring action. As run length increased, reaction times became faster. However, conscious expectation of the cue simultaneously *decreased*. Perruchet et al. (2006) attributed the speeding of reaction time to ‘automatic activation’ caused by repeated recent reinforcement of the association between the warning stimulus and the cue. In contrast, they attributed the dwindling expectancy of the cue to a ‘gambler’s fallacy’ (Tune, 1964; Tversky & Kahneman, 1971): after a run of several trials in which the cue to act was presented, participants believed that the next trial must surely have no cue. Interestingly, the gambler’s fallacy is an anticipatory strategy based on a conscious model of events. In contrast, the automatic activation process might operate entirely outside consciousness.

This chapter shows that the sequential flow of voluntary decisions to act or inhibit involves the same influences as Perruchet previously found for expectancy of a sensory cue. Consider a person asked to freely choose whether to act or inhibit action. Automatic activation would then correspond to a routinised repetition of previous choices, while gambler’s fallacy would correspond to a conscious, high-level decision to try a new choice. Crucially, the contributions of these two processes to any individual action decision can be investigated by considering influences of recent action history.

In addition to gambler’s fallacy, the decision to try a new behaviour in the current experiments can also be considered as a decision to explore new behavioural options, rather than continuing to exploit current options (Cohen et al., 2007). This account takes a more

internally-focused view of voluntary choices, while the gambler's fallacy account focusses more on inferences from recent external events. In this experiment, voluntary choices depended both on monitoring external event sequences, and on internal motivations for choosing to act. These two influences are usually treated separately in psychological research but they clearly relate to each other in everyday human behaviour. Imagine the case of a goalkeeper in the penalty shoot-out stage of a football (soccer) match. Firstly, the goalkeeper can monitor the ongoing sequence of ball locations targeted by the strikers and anticipatorily adapt their behaviour. For example, after the ball has repeatedly been kicked to the left, the goalkeeper may assume the next one will go right. Secondly, the goalkeeper will have internal motivations for diving one way or the other in order to save the ball – perhaps they have dived to the left very often, and feel their actions are becoming too predictable. They should therefore consider a novel response of diving to the right, to avoid leaving an easy option for the striker. This behavioural landscape therefore includes some elements recalling the gambler's fallacy, but also a wider range of causes and effects. Therefore, the experiments in this chapter use the term Anticipatory Action Control (AAC) to refer to the conscious strategising that underlies voluntary decisions in sequential choice behaviours. Our choice of terminology aims to bridge the gap between externally-driven strategy such as gambler's fallacy, and an internally-motivated strategy such as exploration/exploitation. However, the behavioural sign of AAC in the current experiment recalls that of gambler's fallacy: an increased probability of choosing to switch responses following longer runs of the same response.

Experiment 1 is an investigation into whether choices to act or inhibit action are influenced by previous sequential history of behaviour along AAC lines. A common observation regarding motor action in both experimental and phenomenological research is the “prepotency” of responding (e.g. Simpson et al., 2012). This refers to the intensity of a motor impulse at the moment of action execution. Habitual or automatic action patterns may underlie several mental health problems, including obsessive-compulsive behaviour (Graybiel & Rauch, 2000), Tourette's syndrome (Leckman et al., 2006), eating disorders (Davis et al., 1999; Favaro & Santonastaso, 1998), and addiction (Everitt & Robbins, 2005). To examine experimentally how habitual action patterns influence current choices to act or inhibit, we instructed healthy participants to execute or inhibit actions in response to targets paired with auditory tones. We investigated the extent to which runs of preceding action or inhibition responses would influence free choice on a *subsequent* trial either to repeat the

same option (hypothesis of automatic activation - HAA) or to switch to the other option (anticipatory action control - AAC) (see figure 3.1 below). In a final natural experiment, we explore effects of the gambler's fallacy and inhibitory states in real-life professional football penalty shoot-outs.

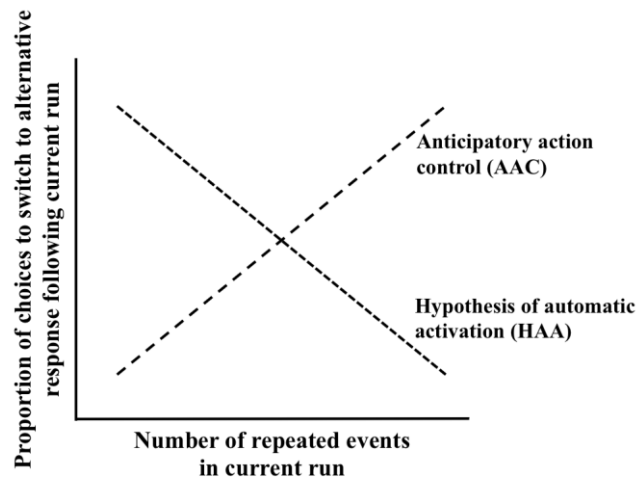


Figure 3.1. Schematic illustration of the AAC and HAA hypotheses. With consecutively repeating events, AAC would encourage selecting the opposite response when given the choice, while HAA would encourage continuation with the current response.

3.2. Experiment 3.1: The gambler's fallacy effect on inhibition decisions – an initial exploration

3.2.1. Introduction

This experiment aimed to compare HAA and AAC processes for free choices to act or inhibit. The HAA mechanism predicts that long runs of action or of inhibition should induce prepotency. A behavioural sign of prepotency of actions might be progressively faster responding due to HAA (Perruchet et al., 2006). A behavioural sign of prepotency of inhibitions might be a slowing of ensuing action responses. Moreover, free choice might also be affected. Long runs of action or of inhibition should make that option more likely to be selected in subsequent free-choice trials, if HAA is present.

The AAC mechanism makes quite different predictions. Specifically, long runs of a single option should promote a switch to the alternative option. Thus, an increasing run of repeated actions should produce a bias to choose inhibition on subsequent free-choice trials, and vice versa.

3.2.2. *Methods*

3.2.2.1. *Participants*

Twenty participants (ages 18-40 years, 13 females) were recruited for the experiment. None had any history of neurological or psychiatric disorders and all gave informed consent.

2.1.2. *Design*

The design was based on a soccer goalkeeper either diving to save a ball, or inhibiting this action and remaining in the goal centre. In each trial, participants saw 3 hollow circular placeholders in the centre of the screen. In a given block, the central circle and either the left or right circle were illuminated to mark them as the two active response options. The third circle was disabled. There were an even number of right and left blocks and their order was randomised. Thus, at any time participants could choose between one action possibility, and inhibition: they chose whether to act, but not what action to perform (Brass & Haggard, 2008).

Participants acted as the goalkeeper, trying to save a ball in a penalty situation. There were 8 blocks of 106 trials each. In 90 trials ('instructed' trials, 84.9%), the ball appeared randomly in one of the two active locations (i.e., centre and left, or centre and right, according to the block; see figure 3.2). An auditory tone (500Hz, 100ms duration) preceded ball onset by 600ms. If the ball appeared in the left or right circle (blocked factor), participants had to press a key as fast as possible with the left or right hand respectively ('instructed action' trial). Reaction times beyond 2 standard deviations from the mean within a condition were removed. Left and right hand responses were blocked. This meant that action responses were specified in advance to maximise motor preparation (Frith & Done, 1986). However, if the ball appeared in the central location, participants were instructed to inhibit their response ('instructed inhibition' trial).

In 16 of the 106 trials per block, a question mark appeared above the two active locations, and no green ball was shown. On these 'free choice' trials, participants were told

that the ball was hidden from view and they should therefore guess its location, by either pressing the key with the designated hand for that block, or inhibiting their keypress response. The actual position of the ball was randomly assigned to one of these positions. No immediate feedback was given on choice trials, but overall feedback on choice performance was given at the end of each block (see *procedure* for rationale).

The sequence of trials on each block was fixed to ensure appropriate runs of trials for testing hypotheses related to HAA and AAC-type mechanisms (Perruchet et al., 2006). Extended runs of several instructed action and instructed inhibition trials. There were 12 independent runs of 1 action, 6 runs of 2 actions, 3 runs of 3 actions, and 3 runs of 4 actions in each block. The same distribution was applied for runs of inhibition trials, creating 8 types of run defined by a 4 (length) x 2 (action/inhibition) factorial arrangement (see table 3.1). Two runs from each of these 8 conditions were followed by a choice trial, giving a total of 16 choice trials. The runs were concatenated in a random order (Perruchet et al., 2006) with the restriction that runs of action were always followed by runs of inhibition and vice versa.

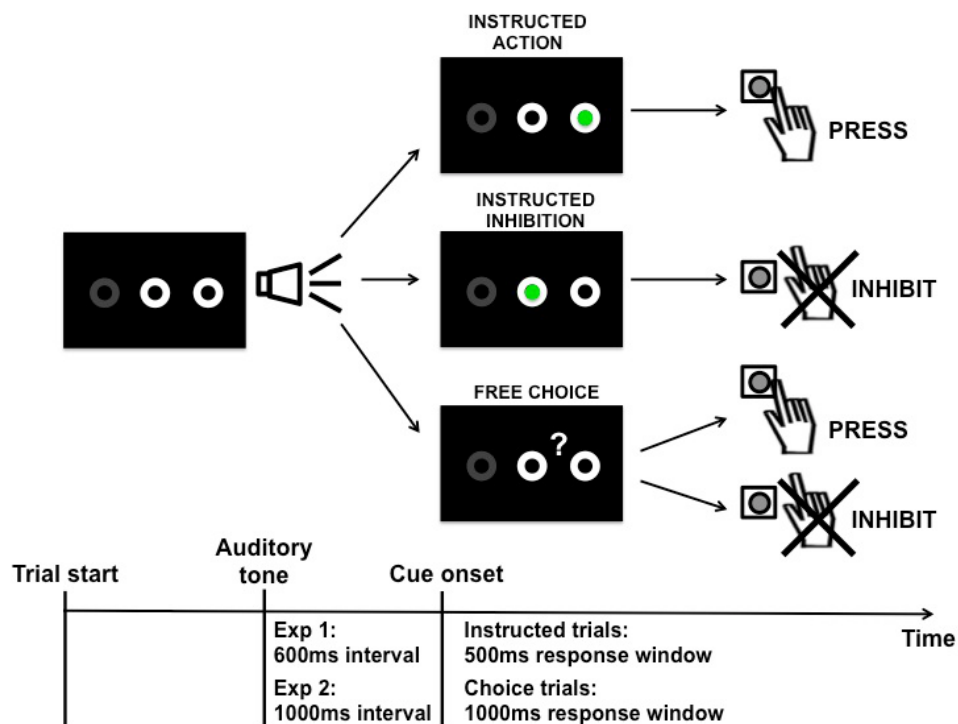


Figure 3.2. Experimental paradigm showing different trial types. The illustration shows a block in which all responses were made with the right hand.

	<i>Number of consecutive trials in a run</i>			
	<i>One</i>	<i>Two</i>	<i>Three</i>	<i>Four</i>
Runs of instructed actions	12 (2)	6 (2)	3 (2)	3 (2)
Runs of instructed inhibitions	12 (2)	6 (2)	3 (2)	3 (2)

Table 3.1. Distribution of runs within each block. The total number of runs in each group that are followed by a choice trial are shown in parentheses.

2.1.3. Procedure

Participants were told that the location of the ball in all trials was determined at random with a 50% chance of it appearing in the two possible locations available in each block. The structure of the trials was in fact pre-determined (see Design), but this instruction was intended to induce the belief that the sequence was random and thus allow a gambler's fallacy-like effect similar to Perruchet et al. (2006). Participants performed a 2-minute practice session of the task before starting the first experimental block.

Each block was preceded by the on-screen instruction 'Left block' or 'Right block'. Participants were told to be highly prepared to ensure a rapid response, corresponding to the left or right ball location. Participants started the block by pressing the spacebar, which revealed the placeholders. An auditory tone was then heard followed by the ball stimulus after a 600ms interval. If participants responded within a 500ms window to a ball in the left/right position, the word 'SAVE' was presented, while slower responses (errors of omission) returned the word 'MISS'. When a central ball appeared, 'SAVE' was presented when participants successfully inhibited their action, while errors of commission also produced a 'MISS', as would be the case in a real soccer game. The very short 500ms response window was chosen to ensure that action responses were always prepared in advance, and were highly prepotent. The inter-trial interval varied randomly between 1-3 seconds.

In choice trials, participants had up to 1000ms to make their decision. Later responses produced a 'MISS'. If they pressed the designated key on the keyboard, the corresponding left/right placeholder became red to indicate their response. If they inhibited action, the central placeholder became illuminated in red to indicate guessing the centre location.

At the end of each block participants received feedback about overall performance. They were rewarded an extra 40 pence per block if they saved the ball in at least 90% of instructed trials, and in more than 50% of choice trials (feedback given separately). Participants were told that making the same response in every choice trial would result in only 50% choice accuracy and therefore no reward. Given the random 50/50 distribution of ball locations in choice trials (see Design), this was in fact true. This encouraged them to try to adjust their responses according to their belief about the actual ball location.

3.2.3. Results

3.2.3.1 Reaction times

Reaction times on instructed trials decreased with longer preceding runs of instructed action responses (see figure 3.3 below). This pattern was analysed using a statistical model established previously (Perruchet et al., 2006). Briefly, trend analysis was used to investigate whether this reduction was linearly related to the number of preceding consecutive actions, by coding run lengths of 1-4 actions with coefficients -3 -1 1 3 for each participant (Hays, 1994). The trend, which corresponds to the average slope of the black line in figure 3.3 for each participant, was significantly less than 0, $t(19) = -2.96$, $p < 0.01$. That is, reaction times decreased linearly with number of preceding action responses, suggesting that responding was facilitated following repeated tone-action stimulus pairings (cf Perruchet et al., 2006).

We repeated this analysis for reaction times for instructed actions following runs of inhibitions. In contrast to the action run slope, this had a positive trend which approached the boundary of statistical significance: $t(19) = 1.95$, $p = 0.066$. The slopes were significantly different from each other, $t(19) = -3.88$, $p < 0.001$.

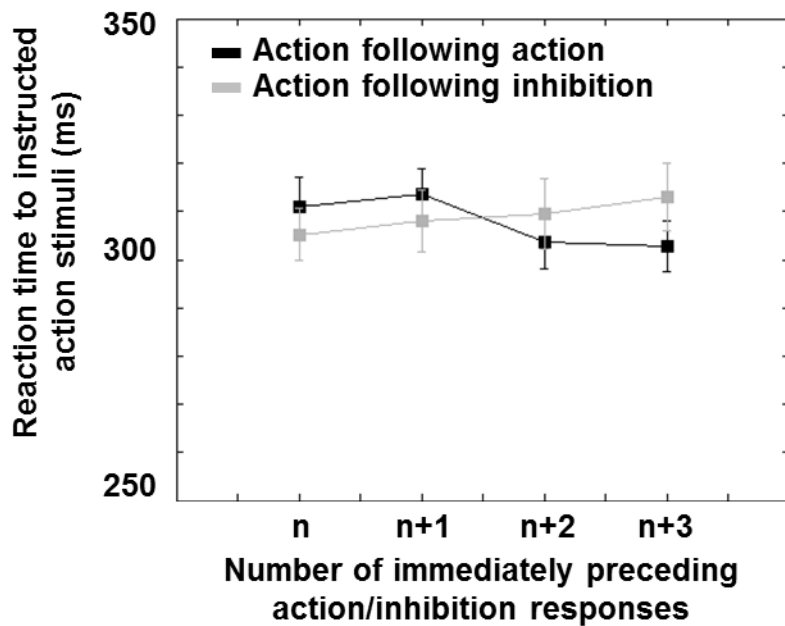


Figure 3.3. Mean reaction times to instructed action stimuli as a function of number of preceding action or inhibition responses. Error bars show standard error across participants.

3.2.3.2 Free choices

We investigated whether long runs of a behaviour might bias ‘free’ choice to favour the alternative behaviour. The percentage of free choices to act increased with number of preceding instructed inhibition trials (linear trend, $t(19) = 4.5$, $p < 0.001$). In contrast, the percentage of free choices to inhibit action did not increase with the number of preceding instructed actions, and in fact showed a non-significant decrease (linear trend, $t(19) = -0.75$, $p > 0.05$). These slopes differed significantly from each other $t(19) = 4.85$, $p < 0.001$.




Figure 3.4. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions. Error bars show standard error across participants.

3.2.4. Discussion

Experiment 1 found that an AAC effect exists for increasingly long runs of inhibitions. With repeated inhibition, people became more likely to switch to action when given a free choice. However, this was not the case for runs of actions: the probability of freely choosing to inhibit did not vary with preceding run length, and no AAC effect was observed. Thus, whereas repeated instructed inhibition trials supported a free decision to switch to an action response, repeated instructed actions did not support the corresponding free decision to switch to inhibition. This asymmetry is surprising, given that mechanisms such as AAC might be expected to work similarly in both cases. The automatic activation mechanism for routinised action appears to over-ride any AAC effect. Interestingly a smaller effect of slowing reaction times with longer runs of preceding inhibitions was also found, suggesting some build-up of ‘automatic inhibition’ is possible with repeated co-exposure to a tone and inhibition target. However unlike the automatic activation for actions, this did not seem to interfere with voluntary choices.

Caution is required in interpreting the lack of an AAC effect following runs of actions. In the Perruchet et al. (2006) study, the influences of gambler's fallacy and automatic activation loaded on different dependent variables – conscious expectancy and reaction time. However, our dependent variable of free choice might be sensitive to *both* influences of HAA and AAC. Indeed, free choices between alternatives appear to depend both on conscious strategy (Jahanshahi et al., 2000), and also on unconscious processes (Wenke et al., 2010; Parkinson and Haggard, 2014). In this case, any AAC effect encouraging a choice to inhibit after runs of repeated actions, would be cancelled by effects of automatic activation. If so, any factor that could reduce automatic activation could potentially unmask an AAC effect for switching from action to inhibition. This was investigated in a second experiment.

3.3. Experiment 3.2: The gambler's fallacy effect on inhibition decisions – reducing automatic activation

3.3.1. Introduction

This experiment aimed to reduce the effects of automatic activation on free choices following repeated runs of actions. Automatic activation depends on the association between the warning tone and action cue, built up over runs of tone-cue pairings (Perruchet et al., 2006). Associative learning theories emphasise that association strength depends not only on the number of pairings experienced, but also on the temporal proximity of the two events associated (Schneiderman & Gormezano, 1964; Smith, 1968). This suggests that increasing the time interval between tone and action stimuli should weaken the association, thus reducing automatic activation effects, and potentially unmasking an AAC effect following runs of actions.

3.3.2. Methods

3.3.2.1 Participants

Twenty-one participants (ages 18-40 years, 16 females) were recruited for the experiment. One was discarded because her highly stereotyped pattern of free choices produced a statistical imbalance in the experimental design (she never chose to inhibit on free-choice trials). None had any history of neurological or psychiatric disorders and all gave informed consent.

3.3.2.2 Design and Procedure

The task was entirely identical to that in experiment 1, except that the tone-target interval was extended to 1000ms instead of 600ms. This value was based on the suggestion that intervals longer than 750ms interval may not support automatic activation mechanisms (Perruchet et al., 2006). The feedback 'SAVE' was now also only given on action trials if participants responded within 320ms rather than 500ms as in Experiment 1. This lower value was designed to increase the likelihood that actions would be prepared on each trial, producing a prepotent action that would require inhibition on the inhibition trials.

3.3.3. Results

3.3.3.1 Reaction times

As in experiment 1, longer runs of consecutive instructed actions produced faster responses (linear trend, $t(19) = -2.39$, $p < 0.05$; see figure 3.5 below), and there was again a positive slope for reaction times following inhibition runs indicating a slowing down (this time highly significant), $t(19) = 3.6$, $p < 0.01$. These two slopes differed significantly from each other, $t(19) = -4.02$, $p < 0.001$. The slope for action runs did not differ from that in experiment 1 (see figure 3.3), $t(38) = -0.13$, $p > 0.05$, but the inhibition slope was significantly more positive than the positive trend found for inhibition runs in experiment 1, $t(38) = -2.33$, $p < 0.05$.




Figure 3.5. Mean reaction times to instructed action stimuli as a function of number of preceding consecutive action/inhibition responses. Error bars show standard error.

3.3.3.2 Free choices

As with experiment 1, increasingly long runs of instructed inhibitions again increased the probability of a free choice to act: linear trend, $t(19) = 4.61$, $p < 0.001$. However, analysis of free choice outcomes following actions showed a very different pattern to that of experiment 1. Unlike experiment 1, choices to inhibit now increased significantly with the number of preceding instructed actions: linear trend, $t(19) = 3.68$, $p < 0.01$.

The positive values of both free choice slopes (see figure 3.6) suggest that both acting and inhibiting induce an AAC effect. The trend to choose action did not now differ from the trend to choose inhibition, $t(19) = -0.9$, $p > 0.05$.




Figure 3.6. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions. Error bars show standard error.

3.3.3.3 Between-experiment comparisons of free choices

We next subjected the individual trend values to a mixed ANOVA with a between-subjects factor of experiment and a within-subjects factor of choice type (action following runs of inhibition vs inhibition following runs of action). This showed no main effect of experiment, $F(1,38) = 1.66, p > 0.05$, but a main effect of choice type, $F(1,38) = 17.3, p < 0.001$. More importantly a significant interaction was present between the factors, $F(1,38) = 8.55, p < 0.01$. Simple effects t-tests were used to explore the interaction. These showed that the trend for choosing to act following inhibition runs did not differ across the two experiments, $t(38) = 0.68, p > 0.05$, but the trend for choosing to inhibit following action runs was stronger in experiment 2 than in experiment 1, $t(38) = 2.84, p < 0.01$.

3.2.4. Discussion

Experiment 2 showed the same speeding of reaction times following longer runs of actions as experiment 1. However, choice outcomes showed the predicted pattern, with both action and

inhibition runs producing AAC effects. Comparison with experiment 1 showed that increasing the temporal delay between warning tone and action stimulus attenuated the way that a sequence of repeated behaviours influenced subsequent free choices. In particular, extending the temporal interval between warning signal and target unmasked an AAC effect favouring choice to inhibit following repeated actions.

3.4. Experiment 3.3: The gambler's fallacy effect on inhibition decisions – controlling for action frequency

3.4.1. Introduction

Two possible explanations could account for the differences between experiment 1 and experiment 2. First, the emergence of the AAC pattern following actions might be due to the extended tone-target interval in experiment 2. Alternatively, the slight overall increase in the intervals between successive actions, due to the longer tone-target interval, might be responsible. To compare these two possibilities, we ran a third experiment which used an identical tone-target interval to experiment 1 (600ms) but extended inter-trial intervals, which reduced the overall rate of repeated actions to even lower than the level of experiment 2. If the results replicate our experiment 1, this would suggest that the tone-target interval is itself responsible for modulating free choices. However, if the results replicate experiment 2, this would suggest that the overall action rate is the key factor underlying asymmetric sequentiality of free choices to act or inhibit.

3.4.2. Methods

3.4.2.1 Participants

Twenty-one participants (ages 18-40 years, 11 females) were recruited for the experiment. None had any history of neurological or psychiatric disorders and all gave informed consent.

3.4.2.2 Design and Procedure

The task was identical to that in Experiment 2. However, the tone-target interval was now set at 600ms (identical to experiment 1) and the inter-trial interval was increased to 2-4 seconds.

3.4.3 Results

3.4.3.1 Reaction times

As in both previous experiments, longer consecutive runs of instructed actions produced faster responses (linear trend, $t(20) = -4.13$, $p < 0.001$; see figure 3.7 below), and longer runs of instructed inhibitions produced slower responses (linear trend, $t(20) = 2.26$, $p < 0.05$). These slopes were significantly different to each other, $t(20) = -4.01$, $p < 0.001$. The slope for action runs did not differ from that in experiment 1 (see figure 3.2), $t(39) = 0.13$, $p > 0.05$, and also did not differ from that in experiment 2 (see figure 3.4), $t(39) = -0.03$, $p > 0.05$. The slope for inhibition runs did not differ from experiment 1, $t(39) = -0.29$, $p > 0.05$, but was significantly less positive than that in experiment 2, $t(39) = -2.15$, $p < 0.05$.




Figure 3.7. Mean reaction times to instructed action stimuli as a function of number of preceding consecutive action/inhibition responses. Error bars show standard error.

3.4.3.2 Free choices

As with experiments 1 and 2, increasingly long runs of instructed inhibitions again increased the probability of a free choice to act: linear trend, $t(20) = 3.2$, $p < 0.01$. Analysis of free choice outcomes following actions showed a pattern very much like experiment 1, but unlike experiment 2. In particular, choices to inhibit did not change with increasing runs of preceding actions: linear trend, recalling experiment 1: $t(20) = 0.54$, $p > 0.05$. The two slopes also differed significantly from each other (see figure 3.8), $t(20) = 2.39$, $p < 0.05$.

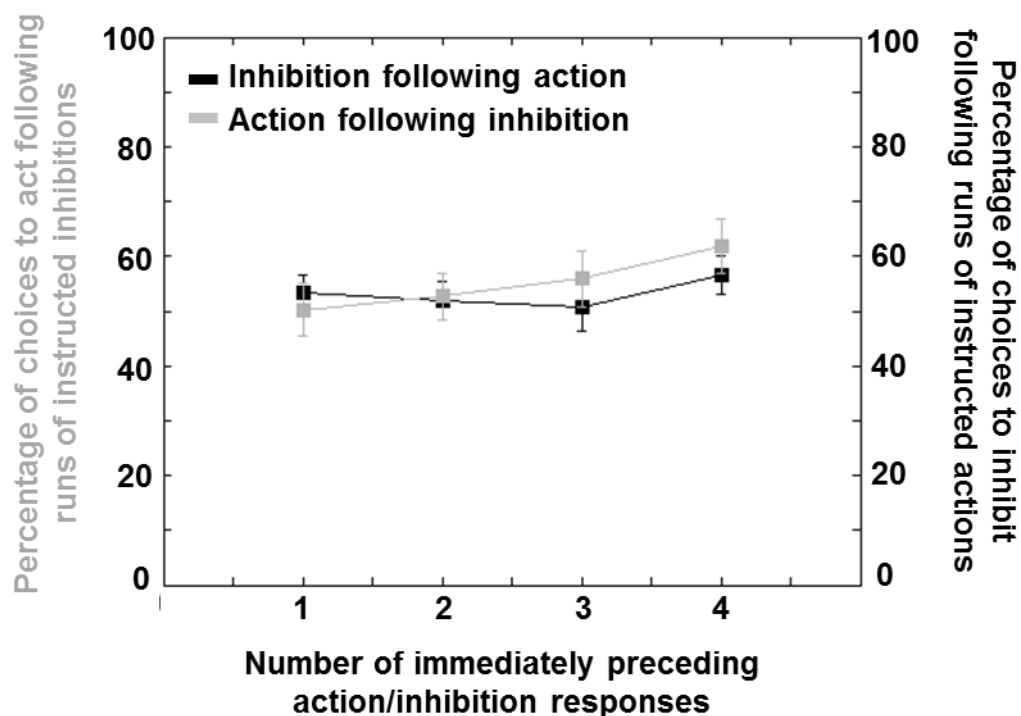


Figure 3.8. Percentage of choices to switch response as a function of number of preceding instructed actions/inhibitions. Error bars show standard error.

3.4.3.3 Between-experiment comparisons of free choices

We ran mixed ANOVAs with a between-subjects factor of experiment and a within-subjects factor of choice type (action following runs of inhibition vs inhibition following runs of action). When comparing experiment 1 with experiment 3, there was no main effect of experiment, $F(1,39) = 3.41$, $p > 0.05$. There was a main effect of choice type, $F(1,39) = 8.08$, $p < 0.01$ but no significant interaction between the factors, $F(1,39) = 3.2$, $p > 0.05$.

When comparing experiment 2 with experiment 3, there was again no main effect of experiment, $F(1,39) = 2.89$, $p > 0.05$. However as well as a main effect of choice type, $F(1,39) = 5.81$, $p < 0.05$, there was now also a significant interaction between the factors, $F(1,39) = 4.95$, $p < 0.05$. T-tests exploring the interaction showed a very similar pattern to the differences found between experiments 1 and 2. The trend for choosing to act following inhibition runs did not differ across the two experiments, $t(39) = 0.74$, $p > 0.05$, but the trend for choosing to inhibit following action runs was more strongly positive in experiment 2 than in experiment 3, $t(39) = 2.72$, $p < 0.01$.

3.4.4. Discussion

Experiment 3 showed a very similar pattern of results to experiment 1. Again, an AAC effect was found for action behaviours following runs of inhibitions but not for inhibition behaviours following runs of actions. We attempted to answer whether the absence of the AAC effect following runs of actions in experiment 1 is due to a strong tone-target association or whether it is due to simply a high frequency of action in recent behavioural history. We maintained the tone-target interval from experiment 1 but extended the inter-trial intervals to reduce action frequency and found that the resulting effect was identical to experiment 1. This strongly suggests that the repeated tight temporal association between tone and target is responsible for the prepotency to continue acting. Reducing this association by extending the tone-target interval (as in experiment 2) restores the AAC effect following runs of actions. Put colloquially, an automatic impulse to act appears to be the enemy of trying to inhibit.

We have investigated how repeated instances of action or inhibition influence subsequent free choices either to act or to inhibit. We applied Perruchet's (2006) dual process model to voluntary choices. Specifically, we proposed that the choice to act or inhibit would depend on the balance in recent behavioural history between automatic activation processes, and anticipatory action control processes (e.g. conscious expectancy). In all three experiments, reaction times decreased with increasing runs of consecutive actions, suggesting a prepotent action tendency, and growing automatization of responding. In addition, reaction times to action targets slowed down following longer runs of preceding inhibitions, suggesting that the processes related to inhibition of action may also show some automatization with repetition.

Experiment 1 found that recent behavioural history had asymmetric effects on decisions to act and on decisions to inhibit. Specifically, long runs of inhibition biased free choice in favour of action, but long runs of action did not comparably bias free choice in favour of inhibition. Experiments 2 and 3 showed that this asymmetry was specifically dependent on the time interval during which action preparation develops. Overall, this pattern of results broadly suggests that rapid, prepotent action leaves little opportunity for inhibition to emerge.

Free choice is often held to depend on conscious thought (Jahanshahi et al., 2000). In the present experiments, an AAC pattern was found for free choices to act or inhibit. The conscious expectancy and cognitive reasoning which influences AAC also depends on a conscious model of external events (Burns & Corpus, 2004; Johnson-Laird, 1994). However, the choice to inhibit action following runs of actions in experiments 1 and 3, showed a reduced AAC pattern compared to the choice to act following inhibitions, as a result of stronger associations between external stimuli and actions. This suggests that strong automatic activation for actions can hinder the contribution of conscious cognitive control when a choice is available. The resulting pattern of behaviour consequently fails to reflect patterns of conscious reasoning and expectancy. In this sense, habitual action and cognitive control are opposing influences on voluntary behaviour.

Perruchet et al. (2006) describe automatic activation as an increasing strength of association between tone and motor response following consecutive tone-action pairings. Following sufficient repetition, the warning tone itself elicits motor activation, facilitating reaction times. We found speeding reaction times with longer action runs in all of the above experiments. Interestingly, experiment 2 still showed this speeding RT effect from automatic activation despite losing the effect of action perseverance in choices following action runs. This suggests that automatic activation may primarily influence reaction times, and only secondarily influence free choices, if its effects are strong enough to persist. Indeed, automatic activation seems most relevant to single choice reaction time tasks, as the added complexity of two-choice decision tasks appears to reduce the effect of speeding (Livesey & Costa, 2014). Automatic activation may occur most strongly when an action can be fully prepared in advance. When an action must be chosen without preparation in advance, either on the basis of an instructional cue, or endogenously, this mechanism may be less effective.

Our study provides novel laboratory evidence that repeated pre-prepared actions directly reduce inhibition, in addition to facilitating action. Successive pairings of auditory tones and actions in close temporal proximity removed a preference for inhibition in subsequent free choice following repeated actions. Frequent co-occurrences of tone and action appear to elicit a positive reinforcement mechanism, so that voluntarily controlled inhibition of the action then becomes unlikely. This mechanism could potentially explain difficulties in inhibiting actions that have become routinised.

3.5. Experiment 3.4: Priming sequential choices to act or inhibit

3.5.1. Introduction

In a fourth experiment we investigated whether unconscious primes can influence the HAA and AAC processes from the previous three experiments. Priming experiments can distinguish between processes that are modulated by unconscious information and those that depend only on conscious information. This is useful for developing a more detailed understanding about the inhibitory systems that underlie the free choice and automatic activation systems. We tested whether the gambler's fallacy and automatic activation effects can be reduced or enhanced by priming, and whether these changes also influence free choices. We also obtained a more explicit measure of conscious expectancy for upcoming trials. This allowed us to more directly investigate whether conscious expectancy or automatic activation drives free choices.

3.5.2. Methods

3.5.2.1. Participants

Twenty participants (ages 18-40 years, 13 females) were recruited for the experiment. None had any history of neurological or psychiatric disorders and all gave informed consent.

3.5.2.2. *Design*

Participants were required to act and press a key or inhibit pressing a key depending on whether an arrow on the screen pointed up or down. Half of all participants acted when the arrow pointed up, and inhibited when the arrow pointed down. The other half did the reverse. Occasionally the arrow pointed both ways (free choice trial), in which case participants had to freely decide whether to act or inhibit. All up and down target arrows were preceded by subliminal metacontrast masked primes which pointed in the same direction (compatible) or in the opposite direction (incompatible) to the target arrows. All free choice targets were preceded by neutral primes in which the arrow pointed in both directions. Primes subtended a visual angle of 1.18 degrees x 2.84 degrees, metacontrast masks subtended an angle of 1.66 degrees x 3.18 degrees, and targets subtended an angle of 1.95 degrees x 5.3 degrees. All stimuli were presented in the centre of the screen.

Trials were grouped into various run lengths, in a similar fashion to experiments 1-3 of this chapter. A run was defined as a particular trial type repeated a number of times before switching to another trial type. There could be a run with a single up-arrow trial (run length 1), 2 repeated up-arrow trials (run length 2), 3 repeated up-arrow trials (run length 3), or 4 repeated up-arrow trials (run length 4). The same applied for down-arrow trials. All runs used *only* compatible priming or *only* incompatible priming. Priming compatibility was never mixed within a run. For each condition (i.e. 2x2 design with arrow direction (up/down) and prime compatibility (compatible/incompatible)), there were 80 runs with a run length of 1, 40 runs with a run length of 2, 20 runs with a run length of 3, and 20 trials with a run length of 4. The sequence of runs was randomised in order, with the restriction that runs using up and down target arrows always alternated. 10 of the runs within each group of runs was always followed by a free choice trial. 3 of these free choice trials in each group were followed by a 'filler' trial in which the a single trial of the same type as the preceding run was repeated, and 2 of the free choice trials in each group were followed by 2 filler trials in which the same trial type was repeated twice. This was to prevent any strong predictability of trial types switching immediately after free choice trials. Filler trials were not used for analysis. Participants were given a break every 100 or 101 runs.

Every trial was followed by a visual analogue scale, which participants used to indicate their expectancy for the next trial. The scale stretched horizontally across the screen from -8 degrees of visual angle to 8 degrees. To the left of the scaled was the word 'Up' and to the right of the scale was the word 'Down'. A cursor was always presented in the centre of

this horizontal scale subtending a vertical length of 1.5 degrees. Participants moved this cursor along the scale left or right, depending on their level of expectancy for either an up or a down arrow on the next trial.

3.5.2.3. Procedure

Each trial started with an inter-trial interval (jittered between 500 and 1000 ms, drawn from a uniform random distribution). Following this, an 800Hz auditory warning beep was presented for 100ms. 500ms after this beep, the prime was presented for a single frame at a 60Hz screen refresh rate. The mask then followed 3 frames after prime onset and remained on-screen for 7 frames. Target onset was 2 frames after mask onset with a duration of 7 frames. Participants had a response window of 500 ms after target onset. Any error in responses or delayed responses resulted in a red 'x' (font size of 1 degree) presented on-screen for 300ms. Following the response window interval, the conscious expectancy scale appeared on the screen. Participants then moved the cursor along the scale using the computer mouse and clicked the left mouse button to indicate how much they expected an up or down arrow target arrow to appear in the next trial. The basic stimulus sequence can be seen in figure 3.9 below.

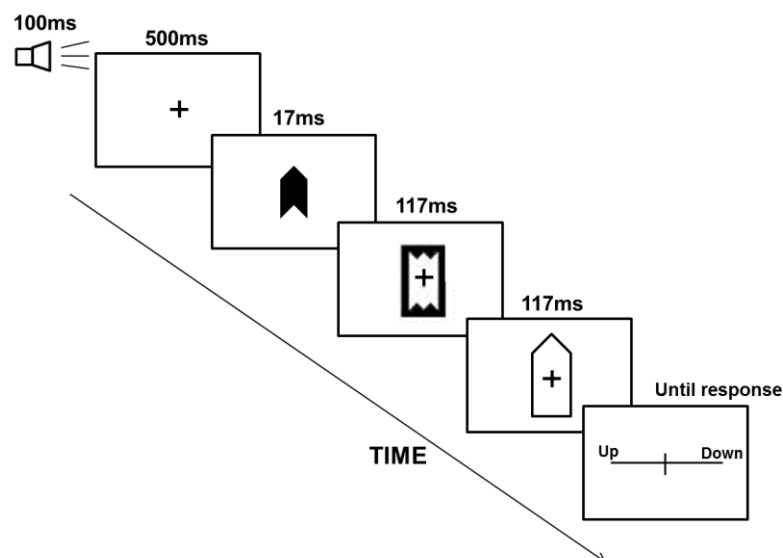


Figure 3.9. Basic stimulus sequence for a typical up-arrow trial with compatible priming. Participants hear a tone then see the masked prime and target. They then indicate which way they think the target on the next trial will point.

3.5.3. Results

Reaction times

Figure 3.9.1 below shows the reaction times for each response and prime compatibility condition across the various run lengths.




Figure 3.9.1. Reaction times to an action stimulus following 1, 2, 3, or 4, preceding consecutive action or inhibition targets. Priming compatibility refers to both the current trial and all preceding trials in each run.

We calculated linear trends for each of the conditions in figure 3.9.1 and then ran a 2x2 ANOVA. A significant main effect of response type was found, $F(1,19) = 10.48$, $p = 0.004$, with actions producing more negative slopes and inhibitions more positive slopes. No main effect of compatibility could be found, $F(1,19) = 2.66$, $p = 0.12$, but a significant interaction was present, $F(1,19) = 5.66$, $p = 0.03$. This was driven by a bigger difference between the action compatibility conditions than the inhibition compatibility conditions. Compatibly primed actions produced a significantly steeper slope than incompatibly primed actions, $t(19) = -2.91$, $p = 0.009$. No difference was found between priming conditions for inhibitions, $t(19) = -0.30$, $p = 0.76$.

Conscious expectancy

We coded expectancy ratings into binary judgements to make the analysis comparable to that in experiments 1-3. We compared the number of ‘action next’ expectancy judgements to ‘inhibition next’ expectancy judgements across the various run lengths. Figure 3.9.2 displays participants’ expectancy judgements for an action target appearing in the next trial.




Figure 3.9.2. Expectancy judgements for an action target stimulus in the next trial following 1, 2, 3, or 4 consecutive preceding action or inhibition targets. Priming compatibility refers to both the current trial and all preceding trials in each run.

A 2x2 ANOVA was again used on the linear trend values. A main effect of response type was found, $F(1,19) = 4.89$, $p = 0.04$. With more frequent consecutive actions, participants believed an action target was less likely to appear on the next trial, while more frequent consecutive inhibitions made participants believe an action target was more likely to appear on the next trial. Conscious expectancy judgements were therefore in line with the gambler’s fallacy. No main effect of prime compatibility was present, $F(1,19) = 1.35$, $p = 0.26$, and no significant interaction was found, $F(1,19) = 1.28$, $p = 0.27$.

Free choices

Figure 3.9.3 below shows the proportion of action choices made on free choice trials.

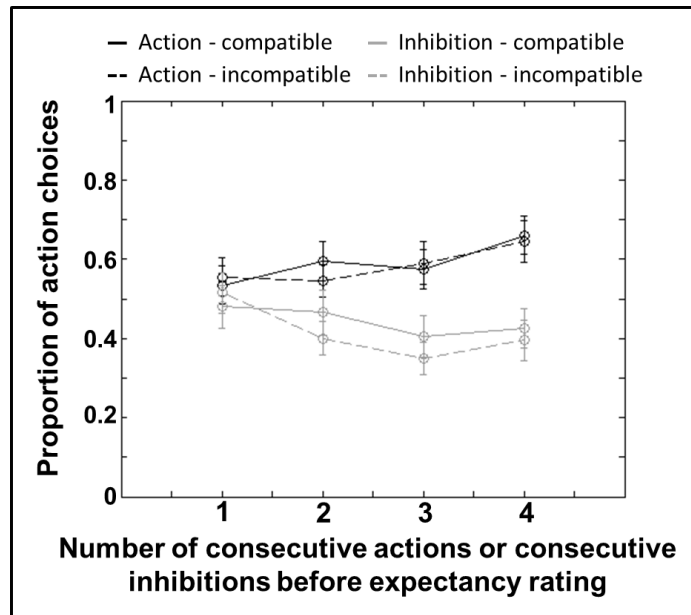


Figure 3.9.3. Proportion of action responses on free choice trials following 1, 2, 3, or 4 action or inhibition responses. Priming compatibility refers to both the current trial and all preceding trials in each run.

A 2x2 ANOVA returned a significant main effect of response type, $F(1,19) = 11.16$, $p = 0.003$. The proportion of action choices increased as the number of preceding consecutive actions increased, while the proportion of action choices decreased as the number of preceding inhibitions increased. This is the opposite pattern to the gambler's fallacy and is in line with the automatic activation hypothesis. No effect of prime compatibility was found, $F(1,19) = 0.44$, $p = 0.52$, and no significant interaction was found, $F(1,19) = 0.40$, $p = 0.54$.

3.5.4. Discussion

Consistent with experiments 1-3, we found speeding reaction times with consecutively repeated action responses. As predicted by HAA, repeated associations between the auditory tone signal and action response elicit automatic motor preparation when the tone is heard

again. This allows participants to produce faster responses. Interestingly, compatible priming for actions not only reduced reaction times overall, but also enhanced the extent of the speeding effect across longer action runs. This suggests that compatible priming increased the level of automatic motor preparation elicited by the warning cue when cues and action targets were repeatedly paired. In other words, it facilitated the strength of the association between auditory signals and action responses.

Conscious expectancy judgements showed the typical gambler's fallacy effect in which repeated occurrences of the same binary event encourage the belief that a different event will occur next. This was true for both actions and inhibitions. As previously demonstrated, the gambler's fallacy effect and the automatic activation effect coexist despite their conflicting influences (Perruchet et al., 2006).

We investigated whether free choices were driven more by automatic activation or more by the gambler's fallacy. Participants were more likely to act when they had acted frequently in the immediately preceding trials, and more likely to inhibit they had frequently inhibited. This is consistent with the automatic activation process rather than the gambler's fallacy. Interestingly, unconscious processes were therefore a stronger influence on free choices than conscious processes. Indeed, priming of the motor system has previously been shown to be insensitive to conscious expectancy (Cressman et al., 2013). Although participants believed that the next target stimulus would be an inhibition target following several action targets, they were more likely to repeat action rather than inhibit when given an entirely free choice. This recalls an analogy with habitual drug-taking behaviours where people cannot help but continue drug-taking actions even though they are consciously aware that it has become a dysfunctional behaviour (Everitt & Robbins, 2005; Robinson & Berridge, 1993).

Unconscious priming has been shown to directly influence free choices to act or inhibit (Parkinson & Haggard, 2014). While this may be true for priming of immediate choices, our data shows that priming effects in behavioural history do not accumulate to affect a subsequent free choice. The results from experiment 2 in this chapter showed that automatic activation can be reduced to a level at which it stops influencing free choices but continues to influence reaction times (see figures 3.5 and 3.6). The priming results from this experiment are consistent with this finding. Priming strongly affected the speeding effect of reaction times across accumulating runs, but free choices were immune to this priming effect.

This supports a model in which implicit or unconscious processes first affect a lower-level motor system, before feeding forward to influence higher-level decisions if their effects are strong enough.

This experiment has shown that entirely free decisions to act or inhibit are influenced by prior event history. Free choices seem to be driven more by implicit automatic activation effects than conscious expectancy effects. Our capacity to inhibit action may be diminished by frequent action in our immediate history (see also experiments 3.1-3.3). Frequent inhibition in behavioural history seems to enhance the probability that inhibition will be subsequently chosen, even when there is an explicit judgement of expectancy for action. Therefore, in contrast to the theory of ego-depletion (Baumeister et al., 2007; Muraven & Baumeister, 2000), frequent inhibition may in fact increase our capacity to subsequently choose inhibition over action.

The priming manipulations in this experiment have shown that unconscious sensory inputs affect actions more than inhibitions. The action system may therefore be more open to external influences than the inhibition system. This is perhaps due to the dynamic nature of action system: continuous actions can be rapidly adjusted online during execution based on updated sensory input. Inhibitions on the other hand are ‘one-shot’ responses that halt motor output. Sensitive adjustments in reaction to external signals are therefore less useful for inhibition. Although unconscious priming may affect inhibition to some extent (van Gaal et al., 2008, 2009, 2010), the present experiment shows that the inhibition system is less sensitive to these influences than the action system.

3.6. *General discussion*

This chapter has highlighted the importance of sequential dependence in action and inhibition decisions. When making a decision about whether or not to act, our preceding history of behaviours influences the outcome of our current decision. We are vulnerable to effects of both habitual continuation of action, and effects of the gambler’s fallacy (Perruchet et al., 2006). Free choices to act or inhibit action depend strongly on our perceptions of immediately preceding events.

These findings are relevant to common real-world scenarios including addiction, and voluntary control of involuntary habits. Consider the example of taking a “recreational”

drug. Inhibiting drug-taking behaviour may become more difficult as the number of recent drug-taking behaviours increases, due to the effects of automatic activation (see experiments 1-3). A repeated association of the urge to take the drug with the immediate action of taking it may make it much harder to subsequently stop when the urge arises. Experiment 2 in this chapter hinted that the timing between warning events (i.e. the auditory tones) and actions may modify the problematic balance between action and inhibition: when the two events are spaced some time apart, a capacity to terminate runs of action by inhibition may emerge. Therefore delaying a drug-taking action following the drug-taking urge may provide a means to weaken the urge-action association. This could in turn ease the strong desire to continue the drug-taking behaviour when the urge next arises.

The example of trying to stop smoking offers some interesting real-world evidence relevant to this point. Smokers often try to quit by gradually reducing smoking frequency (Falba et al., 2004). This technique is said to work even with those not motivated to quit in the first place (Fagerström, 2005). That is, gradually weakening the association between the urge to smoke and the action of smoking should restore inhibitory control over smoking to some extent. Interestingly, substitutes such as e-cigarettes reduce the desire to smoke (Bullen et al., 2010), despite lack of any correlation between the nicotine inhaled from them, and the overall reduction in smoking frequency. This suggests that the success of such devices in helping to stop smoking may be attributable to the maintenance of ritual smoking gestures and actions (Polosa et al., 2011). Part of the difficulty with quitting smoking may lie in the difficulty of inhibiting the highly habitual action of smoking, and not only the withdrawal of the drug. Our method provides a useful laboratory analog of this process. Several accounts of addiction have emphasized behavioural reinforcement and habit-formation underlying drug-taking actions (Everitt & Robbins, 2005; Robinson & Berridge, 1993). Our results suggest an additional, novel element, based on the interval of time between wanting a reinforcer and acting to obtain it. Extending this interval may allow sufficient time for residual capacity to inhibit to be expressed.

Experiments 1 and 3 identified a relative difficulty in inhibiting an urge to act following repeated tone-action pairings. This situation recalls the building impulse to tic in movement disorders such as Tourette's syndrome. Neuropsychiatric conditions offer some insights into the neural mechanisms that may underlie the interplay between action and inhibition, since the pathways underlying tic generation and inhibition have been extensively

studied. Areas such as the SMA and anterior cingulate cortex (ACC) increase their activity prior to tic onset, reflecting the upcoming urge to move (Bohlhalter et al., 2006). Areas associated with inhibitory control, including the inferior frontal gyrus and subthalamic nucleus correlate with tic severity (Baym et al., 2008), suggesting these areas act as a compensatory mechanism for the presence of overwhelming motor urges. The building urge of a desire to tic (Ganos et al., 2012) may be somewhat analogous to a building impulse to act following repetitive actions in healthy individuals. In both cases, the drive to act may be difficult to contain voluntarily. Treatments which encourage patients to control tic actions following the urge to tic (e.g. habit reversal) seem to show some efficacy (Bate et al., 2011). This may be due to the fact that the association between urge and action is weakened over time, eventually making it easier to stop tics.

The method used in experiments 1-3 may also provide an experimental model of self-control. In these experiments, choice behaviour following long runs of inhibitions appeared largely consistent with ego-depletion theories (Baumeister, Vohs, & Tice, 2007; Muraven & Baumeister, 2000). Participants generally chose to act rather than inhibit, after longer preceding runs of inhibition. This suggests that in addition to the effects of AAC, the preceding run of inhibition may have depleted the capacity to inhibit. However, the ego-depletion theory does not seem consistent with our data on the reaction time of instructed actions. We found that reaction times to action signals following progressively longer runs of previous inhibitions were slower. If inhibitory capacity is depleted following repeated inhibition, then faster subsequent reaction times should be expected. Instead inhibitory capacity appears to build with repeated tone-inhibition pairings causing a slowing of reaction times. Interestingly, the effect on reaction times appears dissociated from the effects on choice. Although participants slowed in their reaction times, they showed an increasing tendency to act following long runs of inhibition across all experiments, in line with AAC.

The intentional inhibition system may therefore have some special qualities in comparison to the intentional action system. First, choices to inhibit can be readily linked to the conscious reasoning and expectancy that underlie AAC. Choices to inhibit appear to be less vulnerable than choices to act to the effects of automatic activation from linkage to external stimuli – even when such activation affects motor reaction times. Second, there may be a difference in the optimum temporal interval for linkage between an external stimulus and inhibitory behaviour compared to that for action behaviour. Unlike the highly similar

speeding reaction time effects for runs of action across all experiments, slowing reaction time effects for runs of inhibition were much stronger in experiment 2 than in experiments 1 and 3. Experiment 2 used a tone-target interval of 1000ms rather than the 600ms used in the other two experiments. This suggests that association of tone and inhibition with repeated co-exposure occurs most effectively at a longer latency than that found for actions. Actions are effectively linked to warning tones at shorter intervals, but inhibitions require longer intervals. This could be due to a natural asymmetry between action and inhibition in linking to external events (Haggard, Poonian, & Walsh, 2009). While actions have a tendency to perceptually bind with related events, inhibition may show an active lack of binding to events temporally close in time. A longer interval is required between external stimulus and inhibition, than between external stimulus and action, for any association to form between them.

The AAC and gambler's fallacy effects throughout this chapter relate to the distinction between exploratory and exploitative behaviour (Cohen et al., 2007). The tendency to switch to a novel behaviour after a repeated run can be interpreted as the result of an exploratory decision. The results of experiments 1 and 3 suggest that this exploration favours action following inhibition rather than inhibition following action. Nevertheless, neuroeconomic studies of temporal discounting or delayed gratification (Green & Myerson, 2004; McClure et al., 2004; Mischel & Ebbesen, 1970) suggest that exploratory inhibition may often produce significant benefits. Experiment 2 suggests that exploratory inhibition becomes more likely when external stimuli are less able to drive voluntary choices.

Repeated and habitual action appears to be the enemy of inhibitory self-control. In cases where we believe that inhibiting our actions may be the best thing to do, recent behavioural history can make this difficult. When forced to act repetitively in response to a signal, our subsequent choices may be contaminated by an automatic bias to continue acting when the signal again appears. Weakening associations between such signals and actions by extending the temporal interval between them may allow re-emergence of a level of voluntary inhibitory control.

Chapter 4. A natural experiment: Kickers and goalkeepers in football penalty shoot-out

This chapter reports data on sequential effects in action control from the unusual but natural setting of a penalty shoot-out in professional football. This experiment examined the action and inhibition choices of goalkeepers and kickers in a highly competitive context. It found that kickers were significantly more likely to kick to the centre of the goal than goalkeepers were to remain in the centre of the goal. This was likely due to the asymmetric requirement for inhibition during central choices for kickers and goalkeepers. The experiment additionally found that goalkeepers were more vulnerable than kickers to switching biases following runs of repeated events. Chapter 4 provides an important bridge between intentional inhibition research and real-world behaviour. It shows that performance in high-profile competition may depend on the balance between action and inhibition response alternatives during decision-making.

4.1. Introduction

The experiments in this chapter so far have reported experiments from a controlled laboratory setting. Often, it can be difficult to know whether findings from the lab apply to real world behaviour. This final experiment of chapter 3, reports the findings of a natural experiment in which we studied professional football (soccer) goalkeepers and kickers in penalty shoot-outs of international elite tournaments. We aimed to identify whether the gambler's fallacy applies to sequences of behaviour in this natural setting. We also examined whether there are differences in the performance of goalkeepers and kickers during penalties, given the different roles that intentional inhibition plays in their behaviour.

Sports provide powerful demonstrations of cognitive strategies underlying competitive behaviour (Yarrow, Brown, & Krakauer, 2009). Penalty shoot-outs in football (soccer) involve direct competition between elite players and absorb the attention of millions. The penalty shootout between Germany and England in the 1990 World Cup semifinal was viewed by an estimated 46.49% of the UK population (BBC News, 2004). In a penalty shootout, a goalkeeper must defend their goal without teammate assistance while an opposing series of kickers aim to kick the ball past them into the net. As in many sports (Land & McLeod, 2000), the ball during a penalty kick often approaches too quickly for the goalkeeper to react to its direction of motion; instead, the goal-keeper must guess the likely

direction of the kick, and dive in anticipation, if they are to have a chance of saving the shot (Savelsbergh et al., 2002, 2005; McMorris & Hauxwell, 1997).

There is an interesting behavioural asymmetry between goalkeepers and kickers during the penalty shoot-out. When kickers shoot the ball towards the centre of the goal, their kick is an action in the same way that kicking to their left or right is an action. A kicking action must always be executed. For the goalkeeper however, remaining in the centre of the goal requires the inhibition of a prepotent and frequent action to dive towards the sides of the goal (Chiappori, Levitt, & Groseclose, 2002; Palacios-Huerta, 2003). To attempt to save a kick towards the centre of the goal, the goalkeeper must prevent the urge to perform a diving action. To test whether this inhibitory asymmetry between goalkeepers and kickers has an effect on actual behavioural decisions in the competition, we compared the frequency with which goalkeepers remained in the centre and the frequency with which kickers kicked to centre. Centre kicks should be more frequent than centre goalkeeper positions, if inhibition provides an extra element of difficulty for the goalkeeper.

Previous game-theoretic models have viewed football penalty kicks as a zero-sum, simultaneous-move game. Analyses of penalty kicks during match play have suggested that goalkeepers and kickers both use a mixed strategy, randomly choosing between left and right dives or kicks (Azar & Bar-Eli, 2011; Chiappori, Levitt, & Groseclose, 2002; Palacios-Huerta, 2003). In this case, the penalty kick situation resembles a mixed-strategy Nash equilibrium. Crucially, this equilibrium can hold only if both goalkeepers' and kickers' choices are serially independent and uncorrelated, because any predictable pattern of sequential behaviour could be exploited by the other party. Remarkably, choices in match-play penalty kicks were indeed found to be random and serially independent (Chiappori, Levitt, & Groseclose, 2002; Palacios-Huerta, 2003), even though human attempts to "act randomly" are generally poor (Tune, 1964; Wagenaar, 1972). However, penalty kicks in match play are infrequent and are generally separated by long periods of time, so serial independence may be unsurprising.

In contrast, the penalty shootout is based on a series of multiple kicks taken in rapid succession, after normal match play has failed to produce a clear winner. Each party could easily monitor for sequential regularities in the other's behaviour and could exploit any departure from randomness. The goalkeeper might anticipate the direction of the kick, for example by second-guessing the kicker's behavioural choices (Gallagher & Frith, 2003;

Siegal & Varley, 2002). In particular, sequences of human choices exhibit a ‘‘gambler’s fallacy’’ (Tune, 1964; Tversky & Kahneman, 1971). This involves the belief that in a series of independent random binary events, such as a coin toss, the alternative event becomes increasingly likely to occur after progressively longer runs of one outcome. Following repeated kicks of the ball in the same direction, a goalkeeper subject to the gambler’s fallacy will anticipate the ball going in the opposite direction on the next kick and will therefore dive accordingly. Importantly, the gambler’s fallacy is not based on the value of any single previous outcome, but on the run length over which an outcome has already occurred. If goalkeepers exhibit this regularity, kickers could potentially exploit it to gain advantage. Previous game-theoretic studies mostly considered match-play penalty kicks (Palacios-Huerta, 2003). Although some examined penalty shootouts (Apesteguia & Palacios-Huerta, 2010), they focused primarily on motivational factors predicting performance, rather than on runs of decisions by individual players.

Here, we show that centre kicks and centre goalkeeper positions during penalty shoot-outs are imbalanced. We propose that this is due to a difference between goalkeepers and kickers in the inhibition they require to select the centre of the goal. We also show that goalkeepers display a clear sequential bias. Following repeated kicks in the same direction, goalkeepers are increasingly likely to dive in the opposite direction on the next kick. Surprisingly, kickers fail to exploit these goalkeeper biases. Our findings highlight the importance of monitoring and predicting sequential behaviour in real-world competition, and the importance of intentional inhibition in driving performance differences between opponents in elite competition.

4.2. *Methods*

We examined online videos and statistics of all penalty shoot-outs from FIFA World Cup and UEFA Euro finals tournaments between 1976 (when penalty shoot-outs were introduced into the Euro finals tournament) and 2012. There were 361 penalty kicks in 37 shoot-outs. We recorded the direction each ball was kicked (left/right/centre), and also the direction in which the goalkeeper moved. We first compared the frequency of centre choices for kickers and goalkeepers. Then, when analysing behavioural sequences, centre kicks and dives were

removed to leave only left/right binary choices for analysis.

We analysed the data based on repetition lengths of kick directions from consecutive kickers on a single team. When a left or right kick was the first in a sequence (e.g. the very first penalty taken or a new direction chosen later in the sequence), we recorded whether the goalkeeper opponent on the subsequent penalty dived in this same direction or in the opposite direction. The same direction was coded as a 'stay' while the opposite direction was coded as a 'switch'. This was repeated for kickers – if the next kicker kicked in the same direction on the subsequent penalty, this was coded as a 'stay' while the opposite direction was coded as a 'switch'. The same process was repeated for a run of two repeated kick directions, in which the kickers on a team consecutively kicked left twice in a row or right twice in row. The same again applied for runs of 3 in which the kickers consecutively kicked in the same direction three times in a row. The percentage of switches relative to stays was used in analysis. Analyses used bootstrap resampling (Efron & Tibshirani, 1994; Wasserman & Bockenholt, 1989) to compare the actual data, to results obtained from the same data shuffled 10,000 times.

4.3. *Results*

Centre positions

Overall, kicks to the goal centre were rare (9.14%), and goalkeepers also rarely remained in the centre (2.49% of dives). Kickers were significantly more likely to kick the ball to the centre of the goal than goalkeepers were to remain in the centre of the goal ($p < 0.001$ based on a binomial test). Of all the kicks to the centre of the goal, goalkeepers only successfully remained in the centre of the goal on 6.06% of penalties. When the goalkeeper did choose to remain in the centre of the goal, they correctly anticipated a central ball direction only 22.22% of the time. In most cases, kickers kicked to the sides of the goal. If the goalkeeper is actively deciding not to dive when remaining in the centre, then this decision appears to be maladaptive. Goalkeepers therefore produced dysfunctional centre behaviour relative to kickers. This is likely due to an asymmetric requirement for inhibition in expressing these decisions.

Sequential effects and the gambler's fallacy

We had no strong prior assumptions about the effects of centre choices on behavioural sequences. Therefore, given the rarity of centre responses, we removed all penalties involving a centre kick or a centre goalkeeper position and analysed the remaining set of 321 penalty kicks with clear left/right choices.

Goalkeepers were approximately equally likely to dive left or right, showing no significant difference from 50% (46.73% left; $p = 0.27$, binomial test). Similarly, kickers were equally likely to shoot left or right (46.42% left; $p = 0.22$, binomial test). A chi-square test showed no significant difference between these proportions ($p = 0.94$). Importantly, goalkeepers were no more likely to dive in the direction of the kick than would be expected by chance (53.58%; $p = 0.22$, binomial distribution). This suggests that goalkeepers are indeed unable to react to kicker behaviour and endogenously choose each dive direction.

We next investigated possible exploitable regularities, by analysing effects of repeated ball direction across consecutive kicks on goalkeeper behaviour and determining whether goalkeepers showed a gambler's fallacy. We also examined whether sequences of kick direction did indeed depend on previous kick history. We thus measured the direction of the kick, and of the goalkeeper's dive, after runs of one ($n = 159$), two ($n = 66$), or three ($n = 16$) consecutive kicks in the same direction during a single penalty shootout. We also considered the possibility that goalkeepers' previous dive directions, rather than previous kick directions, could predict future behaviour. In other words, we examined how often goalkeepers dived in the opposite direction following 1, 2, or 3 previous consecutive dives in the same direction. Similarly we examined how often kickers shot in the opposite direction following 1, 2, or 3 previous consecutive dives in the same direction. We used bootstrap resampling (Efron & Tibshirani, 1994; Wasserman & Bockenholt, 1989) to compare the actual data to the same data randomly shuffled 10,000 times. The data can be seen below in figures 4.1 and 4.2. Neither goalkeepers nor kickers showed a pattern of behaviour significantly linked with runs of consecutive goalkeeper behaviour, as none of the actual data crossed 95% confidence intervals for random data. We found no evidence that dive history was relevant for kickers or goalkeepers.




Figure 4.1. Percentage of dives in the opposite direction to last dive direction following runs of 1, 2, or 3 repeated dives in the same direction. Blue bars indicate 10,000 iterations of random data across 20 bins. Dotted black lines are one-tailed 95% confidence intervals of the shuffled distribution. Red lines indicate actual goalkeeper data.

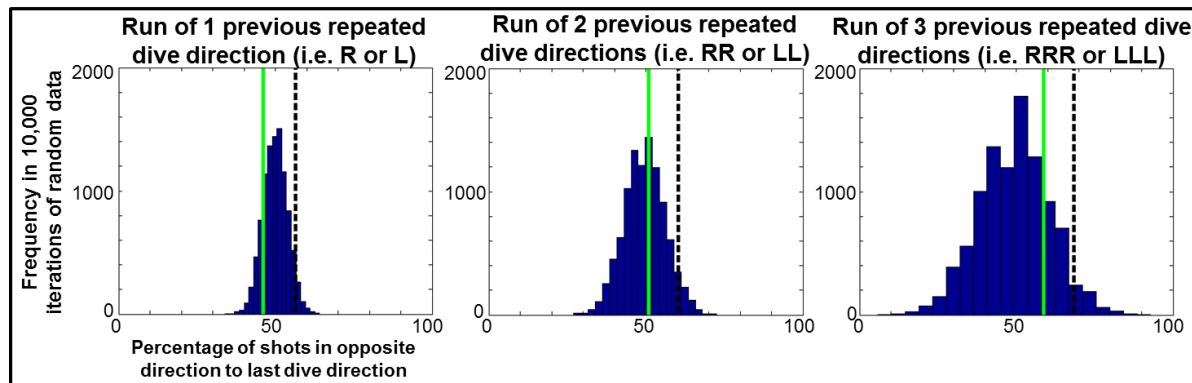


Figure 4.2. Percentage of shots in the opposite direction to last dive direction following runs of 1, 2, or 3 repeated dives in the same direction. Blue bars indicate 10,000 iterations of random data across 20 bins. Dotted black lines are one-tailed 95% confidence intervals of the shuffled distribution. Green lines indicate actual kicker data.

The main analysis here presents the results based on runs of kick directions rather than dive directions. We tested whether goalkeepers and kickers were more likely than chance to switch direction in their behaviour following long repetitions of kicks in the same direction.

When kickers repeatedly kicked in the same direction, goalkeepers became progressively more likely to dive in the opposite direction on the next kick, confirming a gambler's fallacy. Figure 4.3 shows an example in which the goalkeeper dived to the right

after three successive kicks to the left, thus failing to save the ball that the fourth kicker again directed to the left.




Figure 4.3. Four consecutive shots by Portuguese kickers against the English goalkeeper during the Euro 2004 semifinal penalty shootout. The Portuguese kickers are shown in red and the English goalkeeper in blue; the green circle is the ball. All four kicks were to the left of the goal. The goalkeeper made a rightward dive following three consecutive kicks to the left.

Figure 4.4 shows overall goalkeeper data (red line) for percentage of dives in the opposite direction to the last kick. This percentage increased monotonically as run length increased and then crossed the bootstrapped 95% confidence interval after three kicks in the same direction. Goalkeepers' action choices therefore show a pattern similar to the gambler's fallacy (Tune, 1964; Tversky & Kahneman, 1971).

Figure 4.5 indicates that kickers showed a rather different sequential pattern. Kickers showed a trend to switch to a new direction after a run of two kicks in the same direction ($p = 0.07$). Importantly, this trend did not escalate further as run lengths became longer: kickers were no more likely to switch directions after a run of three than would be expected by chance. Overall, kickers seem to show less predictable behavioural sequences than goalkeepers, with a less obvious gambler's fallacy pattern. One important reason for this difference could be that the kicks in penalty shootouts are generated by multiple agents, who act relatively independently. In contrast, the behavioural sequence of dives is generated by the goalkeeper alone, so that cognitive limitations on an individual's random generation figure prominently in his action decisions.




Figure 4.4. Percentage of goalkeeper dives in the opposite direction to the last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction. Blue bars indicate 10,000 iterations of random data across 20 bins. Dashed black lines are one-tailed 95% confidence intervals of the shuffled distribution. Red lines indicate actual goalkeeper data.




Figure 4.5. Percentage of kicker shots in the opposite direction to the last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction. Blue bars indicate 10,000 iterations of random data across 20 bins. Dashed black lines are one-tailed 95% confidence intervals of the shuffled distribution. Green lines indicate actual kicker data.

The goalkeeper’s choices could arise for two reasons. They may believe that the sequence of kick directions is genuinely random and simply be subject to the pervasive gambler’s fallacy (Tune, 1964; Tversky & Kahneman, 1971). Alternatively, the goalkeeper’s behaviour could reflect a “cognitive hierarchy”: the goalkeeper may believe that the kickers will display a gambler’s fallacy (Camerer, Ho & Chong, 2004), though in fact kickers do not. In any case, goalkeepers appear to behave nonoptimally, since they show a sequential bias not aligned with kicker behaviour. The goalkeepers’ sequential bias was not functional, since the dive direction that they chose after runs of three repeated kick directions was no more

likely than chance to coincide with the actual direction of the kick on that attempt (same direction 56.25%; $p = 0.79$, binomial distribution).

Analyses of gambler's fallacy assume that incrementing runs of repeated events affect behavioural choices (Ayton & Fischer, 2004; Perruchet, Cleeremans, & Destrebecqz, 2006). However, other forms of serial dependence could also occur. Lagged regression offers a general framework for considering independent contributions of multiple previous events in predicting the current state. We therefore also applied discrete-choice regression for unbalanced panel data to our data set. This analysis approach has been popular in econometrics (e.g. Arellano & Honore, 2001) and also in behavioural economics, including game-theoretic analyses of tournaments (Apesteguia & Palacios-Huerta, 2010; Chiappori et al., 2002). We investigated whether current dive direction depended on kicking direction, and on dive direction one to three kicks previously. First, we used a random-effects model, as recommended for short panels (Apesteguia & Palacios-Huerta, 2010; Chiappori et al., 2002). The overall model did not provide significant evidence for serial dependence ($F(6,95) = 1.63$, $p = 0.15$), consistent with a previous report (Palacios-Huerta, 2003). Interestingly, the partial coefficients nevertheless showed a significant negative relation between current dive direction and the kick direction three kicks previously ($p = 0.04$), after the contribution of other events was taken into account. We also fitted an additional model using first-difference estimators; these have been preferred because they can be shown to provide unbiased and consistent modelling that can account for unobserved heterogeneity (Wooldridge, 2002). The overall first-difference model was significant ($p < 0.001$). There was again a significant negative effect of the kick direction three kicks before the current dive ($p = 0.03$). However the goalkeeper's behaviour now showed a negative relation to their own previous dive direction at lags one to three (all $p < 0.01$). Recent studies have pointed to additional factors that may affect penalty shootout scores as a result of their motivational significance, notably kicking order (first versus second kicking team in the shootout) and difference between the two teams' scores immediately before each kick (Apesteguia & Palacios-Huerta, 2010). Adding these motivational predictors did not change the pattern of significance for other terms in either regression model.

Such lagged regressions do not consider effects of repetitions within sequences and therefore provide a different perspective from our runs analysis. However, the results are in broad agreement with the runs analysis, in showing a serial dependence on kick direction three kicks before the current dive. They also suggest that goalkeepers may have a general

tendency to avoid repetitions in their sequential diving direction behaviour, consistent with human behaviour during random sequence generation (Spatt & Goldenberg, 1993).

Kickers tend to use the inside of their dominant foot to make contact with the ball during a penalty kick. This means that right-footed and left-footed players find it easier to kick to their left and right, respectively (Chiappori et al., 2002; Palacios-Huerta, 2003). A goalkeeper may use information about the kicker's "natural side" in deciding which way to dive. In our data set, goalkeepers dived to the kicker's natural side 178 times, against 143 dives to the nonnatural side ($p = 0.06$, binomial test), suggesting that this information may be important. We therefore repeated our analyses after recoding all kick directions as "natural" or "nonnatural" for that kicker. Interestingly, this analysis showed that goalkeepers were significantly more likely than chance to dive in the opposite direction to the last kick after a run of one repeated kick direction ($p = 0.03$) (i.e., they were more likely to dive to the nonnatural side of kicker 2 if kicker 1 had just kicked to their natural side, and vice versa). However, goalkeepers did not use information about natural kicking direction for longer runs. In fact, the probability of switching dive direction relative to current kicker preference (i.e., natural side) did not increase with run length (see figure 4.6). Goalkeepers therefore may use information about the kicker's preferences at the start of a run, but they do not appear to keep track of kicker-specific preferences over extended sequences. To do so, the goalkeeper would need reliable information about each preceding kicker's preferred foot and the corresponding actual ball direction. This information would rapidly approach memory capacity. Instead, goalkeepers initially take kicker preference into account in choosing dive direction but thereafter simply track sequences of kicks to left and right. The gambler's fallacy effect identified in our main analysis is therefore based on accumulating information about left and right ball directions without comprehensive modelling of kicker preference. Figure 4.7 further shows that kickers showed no runs-based bias for kicking towards natural or non-natural sides.




Figure 4.6. Percentage of dives in the opposite direction to last ball direction following runs of 1, 2, or 3 repeated kicks in the same direction. Direction is coded in terms of ‘natural’ and ‘non-natural’ kicking directions depending on which foot the kicker uses. Blue bars indicate 10,000 iterations of random data across 20 bins. Dotted black lines are one-tailed 95% confidence intervals of the shuffled distribution. Red lines indicate actual goalkeeper data.




Figure 4.7. Percentage of kicks in the opposite direction to last kick direction, following runs of 1, 2, or 3 repeated kicks in the same direction. Direction is coded in terms of ‘natural’ and ‘non-natural’ kicking directions depending on which foot the kicker uses. Blue bars indicate 10,000 iterations of random data across 20 bins. Dotted black lines are one-tailed 95% confidence intervals of the shuffled distribution. Green lines indicate actual kicker data.

4.4. Discussion

This experiment provided two key findings: 1) Goalkeepers stay in the centre of the goal far more rarely than kickers kick to the centre; 2) Goalkeepers show a gambler’s fallacy in their sequential diving patterns while kickers do not show such a pattern in their kicking patterns.

The greater probability of a centre kick compared to a centre goalkeeper has been previously reported (Chiappori et al., 2002). Some have proposed it arises because kicks to the side have a high chance of scoring when the goalkeeper stays centre, while kicks to centre when the goalkeeper dives to the side have a relatively lower chance of scoring (Chiappori et al., 2002). Therefore, it is riskier for the goalkeeper to wrongly stay in the centre of the goal than to wrongly dive to the side. However, a goalkeeper should be able to monitor and notice that kickers kick to the centre far more frequently than they themselves remain there. Their lack of adjustment to this pattern suggests a potential imbalance in the behavioural states of goalkeeper and kicker. We propose that there is an additional factor to consider when explaining why goalkeepers do not remain centre more often to prevent excessive exploitation of the centre position by kickers. For a kicker, kicks to centre are equivalent to kicks to the side in terms of the action required. For a goalkeeper however, remaining in the centre requires intentional inhibition of a diving action. Given that kicks and dives to the side are so frequent, intentional inhibition of a highly prepotent diving action will be very difficult. This is in line with experimental work showing that frequent action reduces the success of inhibitions (Nieuwenhuis et al., 2003). It is therefore not surprising that goalkeepers remain in the centre less frequently than kickers kick to the centre. They are in fact less able to avoid a prepotent diving action due to the demands of intentional inhibition.

Perruchet and colleagues link the gambler's fallacy to conscious expectancy about the likely pattern of events (Perruchet, Cleeremans, & Destrebecqz, 2006), while others suggest that the gambler's fallacy arises because of a cognitive default based on sampling without replacement (Ayton & Fischer, 2004). Paradoxically, in simple laboratory reaction tasks, the conscious expectancy that a long run will shortly end seems to coexist with increasingly fast reaction times (Perruchet, Cleeremans, & Destrebecqz, 2006). Extending this idea to penalty shootouts, the goalkeeper shown in figure 4.3 might have dived very rapidly, and thus had a particularly high chance of saving the fourth kick, if the kick had in fact gone in the direction of his dive. In addition to this, experiments 1 and 3 in this chapter showed that the capacity to choose to inhibit action can be reduced with frequent preceding actions. This could mean that although the goalkeeper might dive faster with consecutive dives, they are also less able to inhibit action to remain in the centre of the goal when necessary.

Previous game-theoretic studies of elite sports have not specifically considered the gambler's fallacy but have investigated serial dependence of choices in general. In-match

penalty kicks did not show serial dependence (Chiappori, Levitt, & Groseclose, 2002; Palacios-Huerta, 2003). Kovash and Levitt (2009) found clear negative serial correlations in baseball pitches and NFL football passes, though their data sets were much larger than that used here. The gambler's fallacy is conventionally measured as increasing perceptual expectancy of a novel event with increasing run length (Perruchet, Cleeremans, & Destrebecqz, 2006). Kovash and Levitt instead reported lag-one correlations in production of events. Nevertheless, their data are consistent with the broad view that the gambler's fallacy may be an important factor limiting minimax play (i.e. an optimum strategy that minimises worst possible losses) in elite sports (Palacios-Huerta, 2003).

Since goalkeepers are carefully selected and highly trained, one may ask why their vulnerability of regular behavioural sequences persists. Kickers may simply fail to detect, or other- wise fail to exploit, the fundamental limitation in goalkeepers' random generation, which means that the vulnerability can continue without being penalized. Goalkeepers may not be aware of the vulnerability, perhaps because it has not been prioritized in selection and training. In fact, penalty shootouts are relatively rare—our sample contained 168 matches that could potentially have produced a penalty shootout, but only 37 were eventually settled in this way. Moreover, in-match penalty kicks do not show regular patterns of kicking direction (Chiappori, Levitt, & Groseclose, 2002; Palacios-Huerta, 2003). Therefore, the occurrence and implications of the gambler's fallacy in penalty shootouts may have gone unnoticed.

The penalty shootout pits the will of one goalkeeper against the will of many kickers. The cognitive functions of monitoring past performance and generating novel actions are highly relevant to both parties. These cognitive functions have been extensively studied in individuals, but little is known about how they are coordinated across individuals within a group (Adolphs, 2003). We suggest that the goalkeeper's and kickers' performance can be understood by considering the asymmetries in competition between a single individual (the goalkeeper) and a group of individuals (the kickers). First, individuals within a group each produce one-shot behaviours, and therefore effects of preceding action history have less effect on their behaviour. For example, even if most kicks are to the sides of the goal, kickers should be equally able to kick to the centre when they want to on their individual turn. Goalkeepers on the other hand are more vulnerable to the preceding history of actions. Frequent dives create a strong impulse to continue diving, which must be inhibited to remain in the centre. Another asymmetry between a single individual and a group of many

individuals depends on behaviour-monitoring and behaviour-generating processes. For example, goalkeepers may have complete autobiographical memory of the sequence of kick directions, and of their own diving behaviour, because they face every kick. In contrast, kickers have autobiographical memory for their own kick only and may rely on less direct experience for their teammates' kicks. Kickers may therefore fail to monitor the goalkeeper's diving behaviour. They may then fail to detect and exploit regular patterns of goalkeeper dives.

The goalkeeper should have a monitoring advantage over kickers, since self-related actions and material are better retained than other classes of material (Macrae et al., 2004; Symons & Johnson, 1997). However, goalkeepers have the disadvantage of producing predictable sequential behaviour, due to limitations in generating random sequences. In contrast, the group of kickers may collectively produce a more random sequence than any single individual, because each kicker chooses their kick direction only once. Distributed cognition across the group of kickers may increase randomness. Similarly, other animals can produce stochastic or random behaviour at the group level (Bednekoff & Lima, 1998; Perony et al., 2012). Thus, in summary, the individual goalkeeper has an advantage in the monitoring function but a disadvantage in random generation. The group of kickers, in contrast, has a disadvantage in monitoring, but an advantage in random generation.

In the penalty shootout, both fairness and success may depend on the balance between the individual goalkeeper's ability in random generation and kickers' collective ability to predict the goalkeeper's intention by monitoring sequential behavioural patterns. In terms of the explore/exploit dichotomy (Cohen et al., 2007), this chapter suggests that in football penalty shoot-outs, exploration of the centre position of the goal is much less likely when inhibition rather than action is required to remain there. Goalkeepers continue exploiting the sides of the goal, seemingly with the assumption that kickers will kick there more often than they actually do. Our research might motivate cognitive training in random generation for goalkeepers, and in inter- communicative monitoring and prediction for kickers. These could help teams better prepare for nail-biting penalty shoot- outs in future football tournaments.

Chapter 5. The neural correlates of intentional inhibition

This chapter explores the neural mechanisms underlying intentional inhibition decisions. Experiment 5.1 used EEG to examine the effects of continuously fluctuating neural activity on 'free' decisions to act or inhibit. It showed that reduced readiness potential amplitudes were associated with the decision to inhibit action. Such amplitude reductions were likely driven by spontaneous neural fluctuations. Experiment 5.2 examined the spatial coordinates of neural activity related to intentional inhibition. It found an important role for the pre-SMA in voluntary inhibition, in contrast to the commonly reported right IFG activation for instructed inhibition. Chapter 5 reveals distinct neural mechanisms for intentional and instructed inhibition, and shows that our free decisions to inhibit action may depend on spontaneously fluctuating states of the motor system.

5.1. Introduction

Behavioural experiments reveal important facts about how intentional inhibition is used and how it is influenced in everyday life. However, the information they provide about the neural processes that govern inhibitory function is limited. This chapter will present two experiments that directly examine brain activity related to the decision to inhibit action. The first uses EEG to investigate the neural activity that precedes and biases the decision to inhibit. The second uses fMRI to identify areas of the brain that are important for intentional inhibition.

5.2. Experiment 5.1: Omissions in action sequences – EEG

5.2.1. Introduction

The brain mechanisms involved in intentional inhibition remain unclear. Most neurocomputational models of voluntary action have focused on action generation rather than inhibition. For example, several models of frontal cortex are based on hierarchies (e.g., Kouneiher, Charron, & Koechlin, 2009; Koechlin, Ody, & Kouneiher, 2003; Brass & Von Cramon, 2002) with anterior areas generating abstract aspects of a plan and posterior motor areas executing them or generating stimulus-driven responses. However, the mechanisms that initiate plans at the highest level areas remain unexplained.

One could instead think of the voluntary motor system as a loop, in which each action depends on a preceding action, rather than as a linear process with an unexplained initiation. Frontal cortico- BG loops (Alexander & Crutcher, 1990; Alexander, DeLong, & Strick, 1986) could produce internally generated action sequences (Boecker et al., 1998; Tanji & Shima, 1994; Brothie, Iansek, & Horne, 1991) by chaining each action to the one before. On this view, there need not be any obvious hierarchical starting point for voluntary behaviors. Thinking of voluntary action as an iterating loop with a characteristic activation level may also explain the strong relation between cognitive resources and action inhibition. Inhibition of action becomes difficult and self-control may break down when a drive to act is sustainedly present (Baumeister, Vohs, & Tice, 2007; Mischel & Ebbesen, 1970). Here we examine the opposite possibility, whether action inhibition emerges at points where the level of activation in a repeated behavior is momentarily reduced.

We have considered how voluntary decisions to act or not to act arise within a continuous action sequence. Participants were asked to omit one particular action within a repetitive manual motor sequence, based either on an endogenous voluntary choice or an external instruction. Incorporating intentional inhibition into a regular action sequence provides a background of prepotent action. As a result, it may be necessary to truly inhibit an action, as opposed to merely failing to initiate it (Filevich et al., 2012; Kühn et al., 2009). Furthermore, studying a regular sequence of actions helps to fix action timing, allowing us to infer when inhibition of action should occur, if it is present.

We investigated the relation between preparation of the next element in a continuous motor sequence and the voluntary decision to omit an action by examining the late, lateralized component of the readiness potential (RP; Shibasaki & Hallett, 2006). Because this component immediately precedes voluntary actions (Matsushashi & Hallett, 2008; Sirigu et al., 2004; Libet, Gleason, Wright, & Pearl, 1983), its magnitude offers a valuable neural signature of the generation versus inhibition of voluntary action. We described above a nonhierarchical loop model for voluntary chaining of sequences of motor actions. Such models predict that activation states are passed around successive iterations of the loop. On this view, the generation or inhibition of a current action may depend on the level of system activity associated with a previous action. For example, a falling level of activation over successive loop iterations may eventually lead to an inhibition or failure to generate the next action.

5.2.2. *Methods*

5.2.2.1. *Participants*

Twenty-six right-handed participants (15 men, 11 women) were tested. Eight participants were excluded (three had excessive blink/ EOG artifacts, one could not produce a regular action sequence, one could not avoid tapping their foot in addition to their finger, one made excessive finger movements between actions, and two made too few voluntary omissions for ERP analysis to be possible), leaving 18 participants with usable data. All had normal or corrected-to-normal vision, and none had a history of neurological or psychiatric disorders.

5.2.2.2. *Design*

In all tasks, participants pressed the space button on a keyboard with their right index finger once every 2 sec in a self-paced manner. In the rule-based omission task, participants used a rule given by the experimenter to omit every fourth or sixth keypress in the sequence. In the voluntary omission task, participants were instructed to omit a keypress when they themselves chose. Participants were asked to be as spontaneous as possible in voluntary omissions and to decide at the very last moment prior to action. They were asked to avoid preplanning omissions or following a rule. Rule-based and voluntary tasks had 20 trials each. The order of these two tasks was alternated across participants. Individual trials ended once 30 keypresses were made, resulting in trial durations of approximately 1 min.

For the first nine participants, the instruction for the rule-based condition was to omit every fourth keypress. However, preliminary inspection of the data showed that voluntary omissions tended to occur more rarely than this, producing a confound between condition and motor activity. Therefore, the remaining nine participants were instructed to omit every sixth keypress. This successfully balanced overall action and omission frequency across conditions.

5.2.2.3. *Procedure*

Participants first practiced synchronizing keypresses to a 0.5-Hz auditory metronome, allowing them to learn the required rhythm. In the subsequent tasks, no external pacing stimulus was given. Participants were asked to avoid counting seconds in their timing and simply to follow the rhythm they had learned in the synchronization phase. In all trials, participants were first presented with an instruction to “Press space to begin the trial.” Following the keypress, a white fixation cross on a black background appeared centrally on the monitor, which participants were asked to remain fixated on while continuing keypresses at an internally paced rate of approximately 0.5 Hz.

In the practice task, participants produced a sequence of 30 keypresses at 0.5 Hz in each trial. Then the experimental tasks began. In the rule-based omission condition, participants were instructed to omit every fourth or every sixth keypress (see Design). In the voluntary omission condition, participants freely chose when to omit keypresses. Precise timing between keypresses was incentivized in the experimental session. Participants received feedback about their performance at the end of each trial. The mean interval between their actions had to be within 1700–2300 msec, and the standard deviation of these intervals had to be below 200. Omitted actions required an extended average interval of 3500–4500 msec between consecutive actions and standard deviation below 500. Fifteen pence per trial was gained for passing all criteria. Optimal performance could potentially increase participant income by £6.

5.2.2.4. *EEG*

Twenty-seven EEG channels (FT8, FC6, FC4, FC2, FCz, FC1, FC3, FC5, FT7, T8, C6, C4, C2, Cz, C1, C3, C5, T7, TP8, CP6, CP4, CP2, CPz, CP1, CP3, CP5, TP7) were recorded from sensorimotor areas. The ground electrode was at scalp position AFz, and the reference electrode was attached to the right earlobe. EOG electrodes were attached to the external canthi of each eye and the supra and suborbital areas of the right eye. A bandpass filter between 0.1 and 30 Hz and a notch filter of 50 Hz were applied, and the sampling rate was 256 Hz. Data were preprocessed and analyzed in EEGLAB v10.2.5.6b running in Matlab 7.10.

Epochs were defined from 1000 msec before to 500 msec after each keypress and baseline-corrected at –1000 to –800 msec. The baseline correction served to remove effects

of very slow EEG drifts and isolate the component of the EEG related specifically to each action. Omission events were inferred using the temporal interval between immediately preceding actions. If the omitted action is labeled n and the preceding actions $n - 1$ and $n - 2$, then whenever an extended interkeypress interval indicated an omission, the temporal interval between $n - 1$ and $n - 2$ was repeated to insert an omission event in the data. Although this was only an estimate of the omitted event timing, its accuracy was assumed to be similar for rule-based and voluntary omissions.

Epochs were discarded if the potential from any EOG electrode fell outside $+80$ to $-80 \mu\text{V}$ or any other electrode fell outside $+100$ to $-100 \mu\text{V}$. Improbable epochs that contained EEG signal amplitudes exceeding five standard deviations of the mean probability distribution or the mean kurtosis value were removed (Delorme, Sejnowski, & Makeig, 2007). Linear detrending was applied over the recording period to identify and remove drift (Matsushashi & Hallett, 2008). Participants were discarded if the number of events for any condition fell below 50 after artifact rejection (see Participants).

To investigate action preparation, the EEG signal for the last 500 msec before action onset was averaged from the C3 electrode (the late *bereitschaftspotential* (or late RP) over the contralateral motor area; Shibasaki & Hallett, 2006). Given the short intervals of approximately 2 sec between each action event, earlier components risked contamination from previous action related activity. Therefore, only the late RP found at C3 was analyzed. The same procedure was used for inferred omission onsets.

5.2.3. *Results*

Inspection of EEG traces and statistical comparisons showed no differences between the groups that omitted every fourth and every sixth action in the rule-based condition. Therefore, the data were pooled across groups. We defined “omission $- 1$ ” actions as those that immediately precede omissions and “omission $+ 1$ ” actions as those that follow omissions. All other actions were classified as “standard actions.”

Behavioral data

Behavioral data are shown in table 5.1.

	<i>Rule-based</i>	<i>Voluntary</i>
Rewards (%)	70 (23.3)	68.3 (24.4)
Action intervals (msec)	1977.2 (109.0)	2045.5 (110.7)
Inhibition intervals (msec)	3870.5 (198.5)	3935.9 (206.8)

Table 5.1. Rewards (i.e., appropriately timed actions, see text) and mean temporal intervals in rule-based and voluntary conditions

There was no difference between the rule-based and voluntary conditions in the number of rewarded, that is, accurately timed trials, $t(17) = 0.41$, $p > .05$. The temporal intervals between successive actions tended to be shorter in the rule-based than the voluntary condition, although this effect did not reach significance $t(17) = -1.93$, $p = .07$. There was no difference in the duration of omission intervals, $t(17) = -1.18$, $p > .05$. This suggests that the conditions did not differ substantially in difficulty and participants maintained broadly similar timing across both. In the voluntary condition, participants made 5.54 actions on average between omissions with a standard deviation of 1.1. In the rule-based condition, by contrast, the number of actions between omissions was instructed to be 3 or 5 (for “omit every fourth action” and “omit every sixth action,” respectively; see Design).

EEG data

Figure 5.1 shows ERPs for the different types of action at electrode C3.

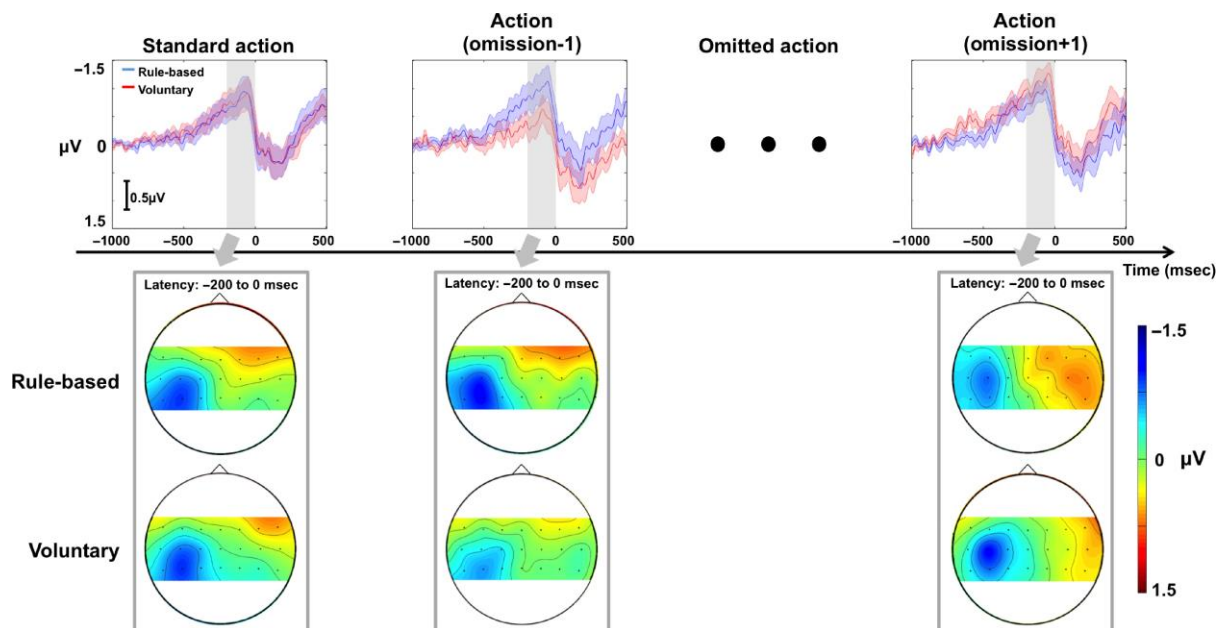


Figure 5.1. ERP data ($n = 18$) for standard, omission $- 1$, and omission $+ 1$ actions at electrode C3—the position of the omitted action in the sequence is indicated by the dots. Shaded colours around ERPs show standard error. Data are time-locked to action onset (keypress). Note difference between conditions in RPs for actions immediately prior to omission.

RP amplitudes were calculated prior to actions in the late RP period (mean of the signal during the period -500 to 0 m). We separately averaged actions immediately before and after the omission and all other “standard” actions. A 2×3 ANOVA with factors of Omission Type (rule-based/voluntary) and Action Position (standard actions/omission $- 1$ actions/omission $+ 1$ actions) showed no main effects of either Omission Type, $F(1, 17) = 1.34$, $p > .05$, or Action Position, $F(2, 34) = 1.15$, $p > .05$. However, a significant interaction was found between these factors, $F(2, 34) = 6.68$, $p < .01$.

To explore this interaction, simple effects tests were used to compare voluntary and rule-based RPs for each action. There were no differences in RP amplitudes between rule-based and voluntary conditions for standard actions across the conditions, $t(17) = 0.45$, $p > .05$, or omission $+ 1$ actions, $t(17) = 1.61$, $p > .05$, but a significant difference was found for omission $- 1$ actions, $t(17) = -4.01$, $p = .001$. The crucial difference between conditions

therefore lies in preparatory activity for actions that immediately precede an omission. Simple effects were also tested by comparing each RP within the rule-based and voluntary conditions. RP amplitude for omission - 1 actions was significantly reduced compared with standard actions in the voluntary condition, $t(17) = 2.56, p < .05$, but no difference was found in the rule-based condition, $t(17) = -1.36, p > .05$. Omission - 1 RPs were reduced relative to omission + 1 RPs in the voluntary condition, $t(17) = 2.89, p = .01$, but not in the rule-based condition, $t(17) = -0.32, p > .05$. Omission + 1 potentials did not differ from standard actions in either the voluntary condition, $t(17) = 1.58, p > .05$, or the rule-based condition, $t(17) = 0.37, p > .05$.

We also explored the EEG activity related to action omission by time-locking to the expected time of action. We had no strong prior hypothesis about the scalp location or form of omission potentials. However, we note that response inhibition has been frequently localized to medial frontal areas (Simmonds, Pekar, & Mostofsky, 2008; Picton et al., 2007; Mostofsky et al., 2003) and that departures and omissions from regular sequences of events are typically measured by vertex potentials (e.g., Nordby, Hammerborg, Roth, & Hugdahl, 1994). We therefore compared the omission potential at Cz using the same time window as for RPs preceding action. The data are shown in figure 5.2. In the period before rule-based omissions, the trace showed a negative-going deflection. This deflection was maximal over the contralateral sensorimotor cortex and had a similar form to an RP, though a somewhat smaller amplitude. Furthermore, the abrupt shift to positivity just before movement onset that marks the end of the classical RP was not present for rule-based omission trials. Because the classical RP is characterized by a negative-going ramp-like form, we used linear fits to the averaged EEG in the -500 to 0 msec time window corresponding to the late RP to investigate whether a component similar to RP might be present on omission trials. We found a trend toward a negative slope for rule-based omissions, $t(17) = -1.89, p = .08$. In contrast, we found a positive-going slope before voluntary omissions. This slope was significantly greater than zero: $t(17) = 3.4, p < .01$. Furthermore, the slopes significantly differed between the two conditions: $t(17) = -5.02, p < .001$.

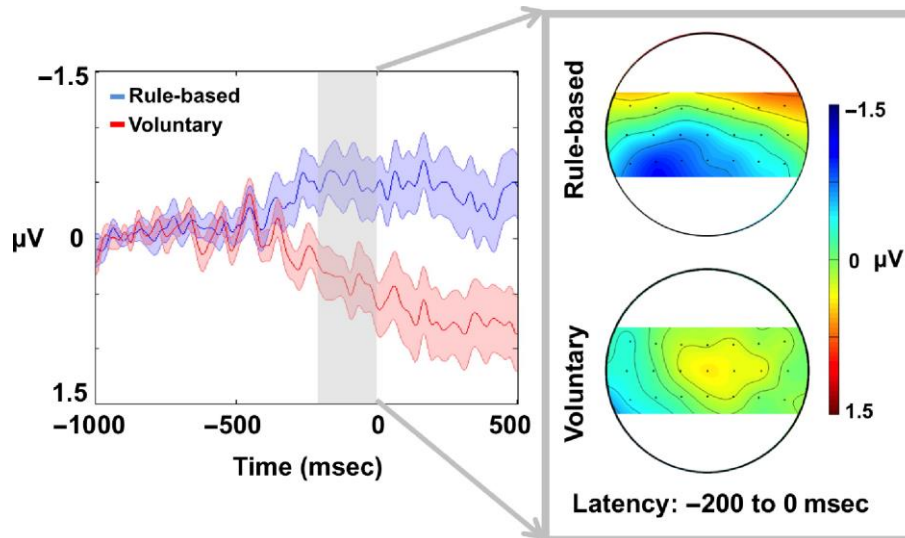


Figure 5.2. ERP data ($n = 18$) for omissions at electrode Cz (left) and across the scalp (right). Shaded colours around ERPs show standard error. Zero milliseconds indicates the inferred time of action omission.

Finally, we compared mean amplitudes during the time window of the late RP for omissions in the same way as we previously did for actions. This showed a significantly greater amplitude for voluntary than for rule-based omission potentials, $t(17) = 2.99$, $p < .01$.

5.2.4. Discussion

We measured event-related EEG associated with actions and with decisions to omit actions in a regular sequence. Our results indicated that the “voluntary” decision to omit a particular action was prefigured by a reduced RP for the immediately preceding action in the sequence. The motor system seems primed in advance to omit an action voluntarily. This was not the case in a rule-based condition, where the particular action to be omitted was specified beforehand by a rule. Therefore, the decrease in action-related processing prior to voluntary action is unlikely to reflect simple advance planning for the forthcoming omission, because such advance planning should be even more evident when using a rule than when choosing “freely.” Any difference in task difficulty across the rule-based and voluntary tasks also seems unable to account for the findings. Although a difference in difficulty could explain a main effect of task, it cannot readily explain the interaction between task and action type (i.e., standard/preomission/postomission actions) found in our RPs. One might suggest that the

effect of task difficulty could be temporally focused on the period just prior to omission, but this seems implausible, because scalp topographies show the main modulation during these periods to be over motor areas, rather than the distributed network associated with task difficulty (e.g., Sunaert, Van Hecke, Marchal, & Orban, 2000). Instead, these results are consistent with a loop model for action generation, in which the level of activation in an iterating motor loop influences decisions to act or not. In particular, background fluctuations in the activation level of such a loop could affect a high-level choice process. A purely hierarchical, feedforward model, in which decisions originate at high levels and cascade unidirectionally to lower levels for execution (e.g., Kounieher et al., 2009), cannot easily explain the association between the activation associated with execution for one action and the decision to omit the next. Our result does not rule out hierarchical models per se, but it does imply that the higher levels in the hierarchy are influenced by feedback from lower-level execution processes.

Because our sequential action task had regular timing, we were also able to calculate ERPs associated with actions that are inhibited by time-locking to the moment when the action would have been expected to occur. The latencies of ERP components for these averages should be interpreted with caution, because the time-locking point is only inferred, rather than measured. Any inaccuracy in this inference will produce a temporal smearing of ERP components and a reduction in ERP amplitude. Therefore, comparing amplitudes between omission-related potentials and action-related potentials may be problematic. However, we can compare omission-related potentials between rule-based and voluntary conditions. Both conditions should be equally affected by smearing, and indeed temporal intervals for actions and omissions did not differ between the conditions (see Results). We found a trend toward a negative-going, ramp-like component, similar in some ways to the form of an RP, prior to instructed, rule-based omissions. Voluntary omissions, however, were preceded by a positive-going potential. We cannot exclude the possibility that this prolonged positive shift may include other components. For example, studies of stimulus-locked no-go potentials found a no-go-P3 component (Bokura, Yamaguchi, & Kobayashi, 2001; Falkenstein et al., 1999), directly related to response inhibition. A temporally smeared version of these components could contribute to our omission-related potentials, but the difference between conditions nevertheless remains. Alternatively, the negative-going omission-related potential in the rule-based condition could conceivably represent a “simulated” but not executed action. For example, participants might represent the rule-based

action sequence as a rhythmic pattern. They might then simulate the “missed beat” of this rhythm to follow the omission rule.¹ In any case, this component was not present prior to voluntary omissions, suggesting an important difference between voluntary and rule-based omissions.

Previous studies of intentional inhibition have been hampered by the difficulty of assessing whether a failure to act was because of a predecision to not initiate any action processing or because of a specific last moment process of inhibition (Filevich et al., 2012). Because there is no clear behavioral marker to indicate the time of inhibition, it is difficult to distinguish whether inhibition occurs early or late relative to action preparation. This distinction is important, because an early inhibition effectively reduces the decision to inhibit to action selection (Mostofsky & Simmonds, 2008), whereas late inhibition implies a specialized cognitive process reminiscent of self-control. The design of our task controls this aspect. In particular, rule-based omission implies predecision about inhibition. In contrast, our voluntary condition invited participants to decide spontaneously to omit an action, as a result of their own real-time choice. If participants had in fact predecided which action to omit in the voluntary condition, then no difference would be expected between voluntary and rule-based conditions. The results showed that voluntary omission was associated with a reduced RP for preceding actions, whereas rule-based omission was not. Reduction of preceding RPs may reflect a mechanism underlying voluntary decisions to inhibit which is independent of conscious prior decision and which could potentially explain the spontaneous nature of some action choices (Libet et al., 1983). In essence, free decisions may capitalize on the momentary state of activation in motor circuits.

Competition between Action and Omission

Recent models of action selection have proposed that plans for multiple response alternatives compete simultaneously in the motor system (Klaes, Westendorff, Chakrabarti, & Gail, 2011; Cisek & Kalaska, 2010; Cisek, 2007). Accumulating information and biasing influences from the BG and pFC contribute to resolving the competition and to selecting a particular response.

Ongoing fluctuations in the level of activity of the cortico- BG-thalamocortical circuit (Alexander & Crutcher, 1990; Alexander et al., 1986) might provide the basis for binary

decisions about action and inhibition behavior. Low activity would make omission more likely, whereas higher activity would favor action. Progressively reducing motor activity would lead to the “voluntary” decision to omit, as shown by our RPs for actions immediately preceding omission.

The lack of preparatory motor activity or excitability during voluntary omissions is consistent with the idea that low spontaneous motor activity biases the choice to omit an action. At electrode Cz, voluntary omissions showed an increasing positivity replacing the ramp-like negativity of the normal RP. EEG provides only a weak indication about the location of underlying generators, which may not correspond to the maximal location of the component on the scalp. However, the central distribution of this positive-going potential could reflect activity in pre- SMA or other medial frontal areas involved in inhibition (Simmonds et al., 2008; Brass & Haggard, 2007; Picton et al., 2007). In any case, our results suggest a specific inhibitory process associated with voluntary inhibition and marked by a positive-going potential at the scalp.

Although both voluntary and rule-based action selections employ a prefrontal control component in the planning of the temporal structures of events and the control of goal-directed action behaviors (Tanji & Hoshi, 2008), a clear external rule appears to override the biasing effects of spontaneous activity in the motor systems. Thus, we found no reduction in RP for actions prior to omission in the rule-based condition. In fact, we found a trend toward a negative-going shift even before rule-based omissions themselves. We speculate that this may correspond to an internal preparation or simulation process, which occurs even when action execution is omitted (Osman, Albert, Ridderinkhof, Band, & van der Molen, 2006) — indeed several fMRI studies confirm that primary motor cortex is often activated during simulated or imagined action (Lacourse, Orr, Cramer, & Cohen, 2005; Gerardin et al., 2000; Lotze et al., 1999). Interestingly, this component was not present for voluntary decisions to omit. This finding is consistent with our hypothesis that voluntary omissions are associated with a reducing level of motor activation.

Another possible account for the present findings is based on changing levels of uncertainty about action in the voluntary condition. Greater probability (certainty) of an impending action has been associated with larger preparatory motor potentials (Scheibe, Schubert, Sommer, & Heekeren, 2009; Low & Miller, 1999). In the voluntary task of this study, there is ongoing uncertainty about whether one should act or omit action in each

moment. Given that actions should be omitted at some point in the sequence, voluntary decisions to omit action may involve a hazard function. Each successive action in the sequence up until the omission involves greater uncertainty than the one before, because it is increasingly likely that the omission will shortly occur. This growing uncertainty might progressively reduce RP amplitudes and could therefore account for the reduced RP we found for omission – 1 actions. In contrast, there is never uncertainty about action in the rule-based task. An account based on uncertainty therefore makes two clear predictions. First, the voluntary task should have globally smaller RPs compared with the rule-based task. The near-identical RPs found for standard actions across the voluntary and rule-based conditions goes against this prediction. Second, in the voluntary task only, RPs should progressively decrease in amplitude as the omission approaches, following a hazard function. In principle, this second prediction could be tested by comparing RPs for omission – 1, omission – 2, ..., omission – n actions. However, our design gave too few trials at omission – 2 and earlier actions to calculate reliable ERPs for testing this prediction. A study with more trials and longer sequences of actions prior to omission would be required to test the uncertainty hypothesis in detail.

The Influential Role of Spontaneous Motor Activity in Voluntary Decisions

There is strong evidence indicating that ongoing spontaneous activity in neural systems influences motor behaviors (Mazaheri, Nieuwenhuis, van Dijk, & Jensen, 2009; Fox, Snyder, Vincent, & Raichle, 2007; Churchland, Afshar, & Shenoy, 2006; Connolly, Goodale, Goltz, & Munoz, 2005). In fact, the RP itself may reflect spontaneously fluctuating neural activity. If the fluctuation is sufficient to cross a threshold, a movement may be triggered, whereas subthreshold fluctuations spontaneously decay without causing a movement (Schurger, Sitt, & Dehaene, 2012). According to this model, the ramp-like shape of the RP may simply reflect averaging those fluctuations, which eventually succeed in crossing the threshold level for action. Several researchers emphasize the functional relevance of spontaneous brain activity and the risk of dismissing such activity as irrelevant noise (Sadaghiani, Hesselmann, Friston, & Kleinschmidt, 2010; Fox & Raichle, 2007). Such fluctuating activity is typically shown to affect behavior when it is examined immediately prestimulus (Sadaghiani et al., 2010). Our results extend this view to the case of action inhibition.

The relation between the EEG signal and these hypothesized fluctuating activation levels is an important issue. The EEG time series contains low-frequency components, including “slow drifts,” whose physiological significance is unclear. As a result, there will be trivial dependence between EEG amplitudes at successive time points. Baseline correction removes these low-frequency components. Our analysis is based on ERPs during brief, nonoverlapping epochs, baseline-corrected at the start of each epoch. This approach reduces the probability of detecting a spurious association between the activity levels for successive actions/omissions because of both events riding upon the same slow drift. Thus, our findings imply fluctuation in signals related to processing successive action events rather than long-lasting shifts in global brain state. For example, figure 5.1 shows a greater postmovement positivity for the action just prior to omission in the voluntary condition relative to the rule-based condition. However, this cannot be assumed to continue into the positive-going shift seen during subsequent voluntary omissions, because a baseline correction intervenes between the two components. That is, our results demonstrate a relation between event-related processing for successive events rather than a trivial continuity of the EEG time series. The focus on event-related processing rather than EEG level across interevent intervals is methodologically appropriate for removing artifactual correlations because of global slow drifts. It is also scientifically appropriate, given that our task was constructed around action and omission events, with no instruction regarding the relatively long 2-sec interval in between.

Inhibition, Cognitive Control, and Consciousness

Many theories of inhibition accord it special status, as a specific cognitive control function with a privileged link to consciousness (Hughes, Velmans, & De Fockert, 2009; Dehaene et al., 2003). Our data suggest a rather different view. Instead of intervention by a high-level process, intentional inhibition may also reflect fluctuating activity in low-level motor circuits. In the real world, when the external motivations for acting or withholding an action are clear-cut and an easy decision between competing options can be made—as is the case in rule-based action selection—spontaneous fluctuations in the motor system may have little influence on voluntary decisions. However, in situations where there are competing motivations to act or not to act, the action selection system must resolve the conflict by choosing between equally attractive alternatives. Our results suggest that such “decisions”

could simply capitalize on the preceding state of the motor system. Many scientists have suggested that inhibition is a necessarily conscious top-down override (Dehaene & Naccache, 2001; Jack & Shallice, 2001; Merikle, Joordens, & Stolz, 1995) or “free won’t” veto mechanism (Libet et al., 1983). Indeed, legal systems assume that the capacity to refrain from inappropriate action is the basis of human moral responsibility. Our results suggest a less homuncular view, namely that “decisions” to inhibit may be consequences of the ongoing state of the cortical and subcortical motor systems. Some theories of volition have involved the very strong contention that “free will” simply reduces to random fluctuations in neural processes (Carpenter, 1999; Eccles, 1985). However, these theories could not convincingly identify the locus of randomness. Our findings suggest that intentional inhibition, like other cognitive processes, indeed interacts with background fluctuations of neural activity. However, we add two important caveats: Our study cannot provide direct evidence that these fluctuations are random, and our study does not preclude other processes contributing to intentional inhibition. However, the design of our task does offer some hints about the mechanism where this fluctuating influence acts. Specifically, intentional inhibition may be determined by the interaction between cognitive, prefrontal decision processes and levels of activation in an iterative motor execution loop that links each action to the next (Marsden, 1984).

In conclusion, endogenous decisions to spontaneously inhibit an action are influenced by lower-level motor activity. Neural activity related to actions preceding and leading up to intentional inhibition can be predictive of its upcoming occurrence. In situations where we are making a difficult decision about whether to act or not, spontaneous levels of activity in our motor system may feed into our decision, biasing us to go one way or the other.

5.3. *Experiment 5.2. Omissions in action sequences – fMRI*

5.3.1. *Introduction*

The brain areas relevant to instructed inhibition are relatively well known. Evidence from stop signal and go-nogo tasks suggest that the right IFG, pre-SMA, and STN areas play important roles in inhibiting actions in response to an external signal (Aron & Poldrack, 2006; Duann et al., 2009). The level of overlap between instructed and intentional inhibition in the involvement of these areas is less well known. Some evidence suggests significant

overlap in activated regions for the two types of inhibition, in particular within pre-SMA and IFG areas (Schel et al., 2014). The division of labour between these areas in inhibition is debated. However, some argue that the pre-SMA has a preparatory task-configuration role for inhibition, while the IFG is more involved in the process of stopping or ‘braking’ action (Aron et al., 2007; Swann et al., 2012). Both of these functions should be common to intentional inhibition and instructed inhibition.

The dorsal fronto-median cortex (dFMC) has frequently been reported as a brain area specifically related to intentional inhibition (Brass & Haggard, 2007; Kuhn et al., 2009; Schel et al., 2014). These studies have focused on the immediate decision between whether to act or inhibit. It is not known whether the dFMC is relevant to intentional inhibition in other contexts, such as during *what* or *when* decisions instead of *whether* decisions (Brass & Haggard, 2008). We often need to decide *when* we want to apply inhibition, rather than whether we need to apply it. For example, when we are walking up to a pedestrian road crossing and notice a ‘red man’ light, we know we must stop at some point near the edge of the road, but when exactly we stop is not well defined. Instead, we must make a free intentional decision. We could halt our walking early on the pavement behind other waiting pedestrians or we could push forward further and wait near the very edge of the road.

The *when* aspect of intentional action decisions has been studied with fMRI. Activity in the rostral cingulate zone (RCZ) has been associated with selecting *what* action to execute, while activity in a more superior medial frontal region near the pre-SMA has been associated with the timing of actions (Krieghoff et al., 2009). Similarly, activity in the SMA and pre-SMA reflects self-generated timing of actions more than external timing (Cunnington et al., 2002; Jahanshahi et al., 1995; Jenkins et al., 2000). In another study conducted by Soon et al. (2008), participants viewed a stream of changing letters, and freely chose when to execute a left or right handed action. The frontopolar cortex and precuneus/posterior cingulate areas produced activity predictive of which action participants would choose, while the pre-SMA and SMA areas showed predictive information about *when* participants would perform the action. The *what* and *when* components of a free decision may therefore depend on separate neural systems.

In this experiment, we applied a very similar experimental design to that used by Soon et al. (2008), in order to study the *when* component in intentional inhibition. Free choices to inhibit were embedded in a continuous stream of action. We implemented this

using a similar task to experiment 5.1. We asked participants to entrain a rhythmic sequence of key presses to a visual stream of letters, and inhibit a single action in two different conditions: 1) when they freely chose; 2) when they saw a letter pre-specified by the experimenter at the start of each trial. We then compared brain activity during these intentional and rule-based (instructed) types of inhibition.

5.3.2. *Methods*

5.3.2.1. *Participants*

Twenty right-handed participants (16 female) were tested. Eight participants were excluded (four had >3mm head motion across the experiment; two performed no useable action omission events in one or more runs of trials; two could not maintain rhythmic actions in order to perform the task and no data could be collected), leaving 12 participants with usable data. All had normal or corrected-to-normal vision, and none had a history of neurological or psychiatric disorders. Scanning was performed at the University of Ghent in Belgium.

5.3.2.2. *Design*

Participants viewed a single letter presented in the centre of the screen which changed to a different letter every 500ms (Soon et al., 2008). Participants were asked to repeatedly press a key with their right index finger, in time with the changing letter stream. There were three blocked conditions. Two of these blocks were voluntary conditions. In these conditions, the letters in the stream were randomised in order with the restriction that the same letter could not be presented within 8 consecutive letters. Participants were asked to keep pressing the key along with the changing letters, and to inhibit making a single key press on one letter whenever they freely chose. They were asked to decide at the very last moment before their inhibition response and to avoid planning ahead. Participants were asked to report either the time of their inhibition or the time of their intention to inhibit, in a similar way to the M and W judgements in previous studies (Libet et al., 1983). Given that the judgements relate to inhibition rather than action in this experiment, we will refer to them as M' and W' judgements. One of the voluntary blocks was the M'-judgement block in which participants reported which letter they actually inhibited their action on at the end of the trial. The other

was the W'-judgement block in which participants reported which letter they saw when they first had the intention to inhibit action.

In the rule-based condition, the letter stream changed in the normal ABC...XYZ order, starting at a random letter. Participants were given a single random letter at the start of each trial which they had to inhibit their action on. This letter was 20-25 letters away from the starting letter in the ABC sequence (returning to A after Z). Therefore, the timing of inhibition was highly predictable in this condition.

Voluntary blocks were always presented consecutively and counterbalanced across participants. The rule-based condition was presented either before or after both voluntary blocks; this order was also counterbalanced across participants.

5.3.2.3. *Procedure*

Participants were told before each block which type of task they would be performing: either the rule-based, voluntary M'-judgement, or voluntary W'-judgement task. They were explicitly told during the voluntary tasks that they should always make the decision to inhibit action spontaneously and at the very last moment before inhibiting. During each trial of the rule-based task, participants were informed of a specific letter to inhibit their action on before the trial started.

The computer initiated each trial and all trials started with a countdown from 6 to 1 before the letter stream. These numbers changed at the same rate as the letters with each number presented on-screen for 500ms, and served to entrain participants' key presses before the letters appeared. Participants were told to initiate rhythmic key pressing in time with the changing numbers as soon as they appeared. Immediately after the offset of number 1, the letter stream started and participants continued their time-locked key presses along with the changing letters. Participants then inhibited their action according to the instructions of the task block. If an inhibition was made within 3 seconds of the start of the letter stream (i.e. 6 letters), a 'TOO EARLY' signal appeared, and this trial was excluded from analysis.

After a successful inhibition, participants were asked to immediately continue actions again as the letter stream continued for another 4-8 letters. In the rule-based condition, the next trial then started after a random interval between 2-5 seconds. In the voluntary-M' condition, participants were instead asked which letter they saw when they *actually* inhibited

their action. In the voluntary-W' condition, participants were instructed to judge which letter they saw when they first *intended* to inhibit their action. These instructions remained on screen for 1 second. Next, to obtain the judgement, 4 items were presented in quadrangular form. These were: 1) the letter presented on-screen during the inhibited action; 2) the letter presented on-screen 1 letter before the inhibited action; 3) the letter presented on-screen 2 letters before the inhibited action; 4) a '#' symbol. Participants selected one of these items corresponding to the answer to the immediately preceding question. If the letter corresponding to their response was not present, they were asked to select the '#' symbol. After their response, the next trial started after a random interval between 2-5 seconds. The basic stimulus sequence can be seen below in figure 5.3.

There were 6 separate sessions for each of the 3 task conditions, making 18 experimental sessions in total. Each session was approximately 3 minutes long. Trials continued in each session until the 3 minute mark was reached, at which point the session finished after the end of the current trial. Participants then had a brief break before starting the next session.

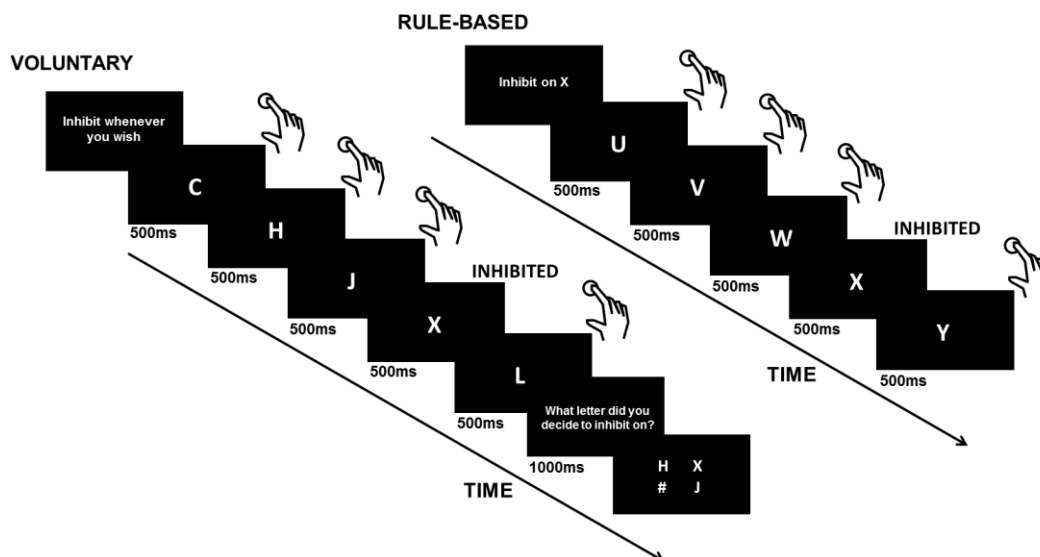


Figure 5.3. Basic stimulus sequence for voluntary-W' and rule-based conditions. In the voluntary-M' condition, the question asking 'What letter did you decide to inhibit on?' was replaced with 'What letter did you actually inhibit on?'

5.3.2.4. *Scanning procedure*

Participants were scanned in a 3 T Magnetom Trio MRI scanner system (Siemens Medical Systems, Erlangen, Germany). A standard eight-channel radio-frequency head coil was used. Participants were asked to minimise head motion to avoid motion artifacts. Any participants who produced over 3mm of head movement were excluded from analysis. Before the experiment started, 176 high-resolution anatomical images were acquired using a T1-weighted 3D MPRAGE sequence (TR = 1550 ms, TE = 2.39 ms, TI = 900 ms, acquisition matrix = $256 \times 256 \times 176$, sagittal FOV = 220 mm, flip angle = 9° , voxel size = $0.9 \times 0.86 \times 0.86$ mm³ (resized to $1 \times 1 \times 1$ mm³)). Whole brain functional images were collected using a T2*-weighted EPI sequence, sensitive to BOLD contrast (TR = 2000 ms, TE = 35 ms, image matrix = 64×64 , FOV = 224 mm, flip angle = 80° , slice thickness = 3 mm, distance factor = 17%, voxel size $3.5 \times 3.5 \times 3.5$ mm³, 30 axial slices).

5.3.2.5. *Data preprocessing and analysis*

Data were preprocessed and analysed using SPM8 (Wellcome Trust Centre for Neuroimaging, London, UK). The first four scans of each individual time series were removed to account for T1 relaxation artifacts. Then the data was slice time corrected, realigned to the mean image by rigid body transformation, coregistered with the structural image, normalised to the Montreal Neurological Institute (MNI) template, and smoothed with an 8mm full-width at half maximum filter.

For first-level analysis, all inhibition events in each session (one per trial) for each participant were convolved with the haemodynamic response function (HRF). Due to the highly frequent action responses (multiple actions in every scan), modelling every action was likely to be problematic for general linear model (GLM) estimation. Therefore, only the middle key press action response in each action sequence for a trial was modelled as an additional regressor. This gave the same number of action and inhibition events (1 per trial). Both action and inhibition events were modelled along with their first temporal derivative creating 4 regressors in the first-level general linear model (GLM). Six additional head motion regressors were also included in the model to explain any remaining head motion variance in the data. The six separate experimental sessions were also modelled as separate sessions in the GLM. This model was built separately for each of the three task conditions.

Contrast images from the first-level model were taken through to second-level analysis to create a 3 x 2 design (voluntary-M'/voluntary-W'/rule-based x action/inhibition). Group contrast images were built for whole-brain analysis. These images were thresholded at $p < 0.05$ FWE corrected or $p < 0.001$ uncorrected for inhibition contrasts, as is typical in inhibition studies (Horn et al., 2003; Jaffard et al., 2008). Note that due to the very high frequency of action (approximately 2Hz), and low temporal precision of fMRI (TR = 1550ms in this experiment), all scans during an inhibition were also certain to have an action performed within them. This makes action contrasts uninterpretable, because scans with actions cannot be compared to scans without actions. However, comparisons of rule-based and voluntary inhibitions are still interesting. Any effects of action should cancel out in a contrast, leaving only the critical differences between instructed inhibition and intentional inhibition.

5.3.3. Results

Behavioural data

Figure 5.4 below displays overall M' and W' judgements relative to the time of inhibition. When participants reported the same letter as the letter they inhibited on this had a value of 0ms. When the letter immediately before was reported, this had a value of -500ms, and when the letter two letters before was reported, this had a value of -1000ms. The values displayed in figure 5.4 shows the means across all of these judgements.




Figure 5.4. *M' and W' judgements relative to time of actual inhibition. Error bars show standard error across participants.*

As expected, W' judgements were significantly earlier than M' judgements, $t(11) = -8.08$, $p < 0.001$. In addition to this, M' judgements were significantly different to 0, $t(11) = 4.32$, $p = 0.001$, and W' judgements were also significantly different to 0, $t(11) = 10.10$, $p < 0.001$.

fMRI data

Actions

Figure 5.5 below shows activation for action responses relative to the implicit baseline of all unmodelled events.




Figure 5.5. Brain activations for action > baseline contrast. The table shows the peak voxel for all clusters at $p < 0.05$ (FWE correction).

Interestingly, no clear activation could be identified on the precentral gyrus to reflect M1 activity in the action > baseline contrast. M1 activity may have been cancelled out in the contrast due to the high frequency of action in all scans. However, there were other significant areas of activity, which may have reflected fluctuations in action-related activity over a trial. During the central action response in each trial, the cerebellum produced significant activation, which may have reflected the developing rhythmic nature of the action task (Ivry & Keele, 1989). Similarly, activations in the superior temporal gyrus could also be explained by perceptions of the simple metric rhythm (Grahn & Brett, 2007). The postcentral gyrus also produced activation, perhaps reflecting the somatosensory qualities of finger movements which may vary in intensity over the trial.

Rule-based and voluntary inhibition

Figure 5.6 below shows activations for the inhibition contrasts. Here, we used a more relaxed significance threshold of $p < 0.001$ uncorrected. FWE correction may have been too conservative in these contrasts as it left no activations to interpret. The low sample size in this study is an important limitation, and should be considered when examining the data.

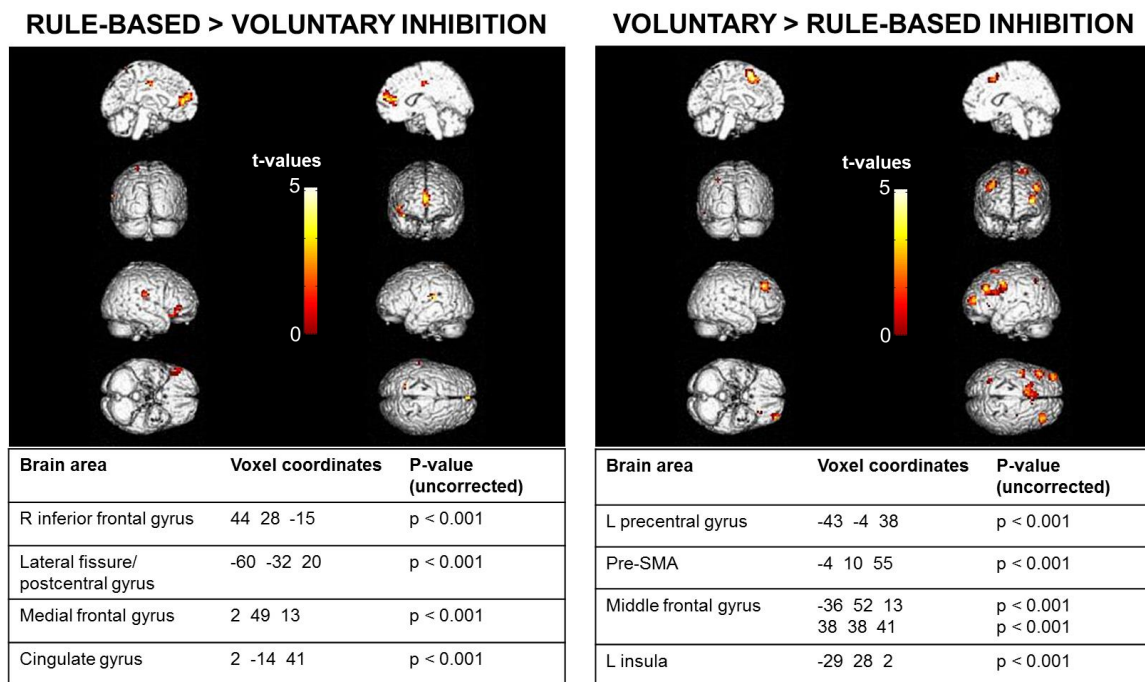


Figure 5.6. Brain activations for rule-based and voluntary inhibition contrasts. The tables show the peak voxel for all clusters at $p < 0.001$ (uncorrected).

Rule-based inhibition activated the right IFG significantly more than voluntary inhibition. In contrast, voluntary inhibition activated more posterior medial frontal areas including the pre-SMA. In addition, voluntary inhibition showed more activation in dorsolateral prefrontal cortex compared to rule-based inhibition. Significant insular activation was also found in the voluntary inhibition condition, consistent with previous evidence (Brass & Haggard, 2007).

5.3.4. Discussion

At least in temporal terms, the decision to inhibit appears to be very similar to the decision to act. Here, we show for the first time that the M' and W' judgements during intentional inhibition are broadly consistent with the typical M and W judgements that participants report during intentional actions (Haggard & Eimer, 1999; Soon et al., 2008). This suggests that chronometric methods used for measuring the moment of intention in action (e.g. Libet et al., 1983), could also be used for intentional inhibition.

The fMRI data in this experiment contributes to what we know about the differences between instructed and intentional inhibition. Consistent with previous work, instructed inhibition produced significant activation in the right IFG (Aron & Poldrack, 2006). The right IFG is thought to play a critical role in stopping or 'braking' action (Aron et al., 2014), with activity beginning around 100-250ms after instructed stop signals (Swann et al., 2009). Inhibitory activity in the IFG is typically characterised by increased beta band oscillations at the moment of action stopping (Swann et al., 2012). Lesions to the IFG also impair performance in stop-signal inhibition tasks (Aron et al., 2003), supporting the view that the IFG has a critical role in action inhibition. Therefore, the rule-based task in this experiment may have activated a stronger stopping process than the intentional inhibition task.

Rule-based inhibition also displayed a medial frontal activation, which included the anterior cingulate cortex (ACC). The ACC has been associated with response conflict and monitoring (Botvinick et al., 1999; Botvinick et al., 2004; Braver et al., 2001). Inhibition responses in the current task were very infrequent relative to actions, making them highly conflicting when executed. This could account for the ACC activation but it is perhaps surprising that such activation was significantly higher during rule-based than voluntary inhibitions. One possible explanation is that during rule-based inhibition, participants were more focused on ensuring inhibition during one specific letter (i.e. that specified by the rule). In the voluntary condition on the other hand, the letter during which inhibition occurred was less important. The more careful monitoring and letter monitoring/prediction required by the rule-based inhibition condition may have driven greater ACC activity. Similarly, ACC activity is also found during the monitoring and detection of errors (Carter et al., 1998). Inhibition 'errors' in the current task may be more relevant to the rule-based condition than the intentional condition, particularly because participants are given a specific letter that they must inhibit action on. This clear instruction means that participants may be carefully monitoring for errors as they attempt to correctly inhibit action based on the rule. In the

voluntary condition on the other hand, participants can inhibit action whenever they wish. Intentional inhibitions are therefore more ambiguous, and errors are less objectively defined.

The ACC activation for rule-based inhibition also extended more anteriorly into medial frontopolar cortex (BA10) (Ridderinkhof et al., 2004). BA10 function may be related to prospective memory (Burgess et al., 2003). The rule-based task in this experiment did involve a prospective memory component, in that participants were instructed to inhibit an action when they saw a particular letter in a sequence. The voluntary condition on the other hand involved no such component because participants inhibited an action whenever they freely chose. This may explain the asymmetric medial frontopolar activation across the two tasks. Interestingly, voluntary inhibition showed some lateral frontopolar activation. There may be a lateral versus medial distinction in frontopolar function, with medial areas associated with attention to external stimuli and lateral areas associated with a focus on internal thoughts (Burgess et al., 2003). The data from experiment 5.2 support this distinction: rule-based inhibition encourages attentional focus on the external letter stimuli as target letters are monitored, and voluntary inhibition encourages attentional focus on internal thoughts, motivations, and intentions.

Dorsolateral prefrontal cortex (DLPFC) activity was associated with voluntary inhibition more than rule-based inhibition. Activity in this area has previously been linked to willed voluntary actions (Frith et al, 1991). Similarly, DLPFC activity has also been reported during voluntary suppression of sexual arousal to erotic stimuli (Beauregard et al., 2001), and the voluntary suppression of emotional sadness (Levesque et al., 2003). The data presented here in experiment 5.2 suggests that DLPFC activity is also relevant to voluntary inhibition of manual actions in rhythmic sequences. During voluntary inhibition, there was also some contralateral precentral gyrus activity. This may reflect greater motor preparation for action immediately before voluntary inhibition events than rule-based inhibition events. This motor activity importantly suggests that voluntary inhibition was spontaneous and unplanned, as participants prepared an action before choosing to inhibit it.

Voluntary inhibition activated the pre-SMA significantly more than rule-based inhibition. The pre-SMA is frequently linked to inhibitory function, but this includes stimulus-driven inhibition (Aron & Poldrack, 2006). Interestingly, in the context of actions, pre-SMA activity is associated with 'free' rather than instructed actions (Lau et al. 2004). Lesions to the pre-SMA in monkeys results in an impaired ability to produce self-generated

actions to receive food, but an intact ability to respond to external stimuli (Thaler et al., 1995). The direct comparison between voluntary and rule-based inhibition presented here, suggests that the pre-SMA may also have a stronger involvement in ‘free’ intentional inhibition compared to instructed inhibition.

The division of labour between pre-SMA vs IFG in action inhibition is not entirely clear. Recent evidence has attempted to clarify this issue, by using intracranial EEG to show that pre-SMA activity precedes IFG activity during the action stopping process by around 750ms (Swann et al., 2012). The pre-SMA may therefore be involved in preparing neural systems for action control, while the IFG is more involved in the action stopping mechanism itself. This account would suggest that the rule-based inhibition task in the present experiment involves stronger action stopping while the intentional inhibition task involves a greater need for action monitoring and internal planning for inhibition.

This experiment did not identify any dFMC activity previously found in other work (Brass & Haggard, 2007; Kuhn et al., 2009). Experiments reporting dFMC activation usually explore brain activity around the time at which the decision about *whether* to inhibit action is being made (Brass & Haggard, 2007; Kuhn et al., 2009). In contrast, this experiment explored brain activity during the decision of *when* to inhibit. There may therefore be an interesting division of labour between the more anterior dFMC area and the posterior pre-SMA area. During decision-making, activity in more anterior regions may first code for whether actions should be executed or inhibited. The pre-SMA may then be responsible for the timing of such action execution or inhibition. A similar division of labour has previously been applied to intentional action, with the rostral cingulate zone (RCZ) responsible for action selection, and the pre-SMA responsible for action timing (Kriehoff et al., 2009; Mueller et al., 2007). Anterior to posterior organisational principles are also common for frontal function (Koechlin & Summerfield, 2007). The data presented here along with previous fMRI work on intentional inhibition, suggests that a hierarchical structure may apply to intentional inhibition function along the medial frontal cortex.

Finally, the insula also showed significantly greater activation during intentional compared to instructed inhibition. This is consistent with the idea that the insula may have a role in affective monitoring of intentional action and intentional inhibition outcomes (Brass & Haggard, 2010). Evaluating outcomes is particularly important during self-generated decisions compared to instructed decisions, because self-generated decisions can be adjusted

in the future. If we choose to inhibit an action and evaluate it as a bad decision, in the future we can choose to act instead in a similar situation. On the other hand, when responses are instructed by a stimulus, behaviour is less flexible, and therefore there are fewer possibilities for outcome evaluation to guide future selection. Another important insula function is in interoceptive awareness (Critchley et al., 2004; Craig, 2009). The greater insula activation during intentional inhibition than instructed inhibition may reflect the greater demand to monitor internal signals in making a 'free' decision to inhibit action. Participants may monitor internal motor states as they consider when to apply intentional inhibition, in contrast to rule-based inhibition which is defined clearly by an exogenous rule. Intentional inhibition in the current task may also produce a much stronger internal experience of action omission in self-generated sequences, compared to instructed inhibition. The omission event itself is less predictable and is characterised by a stronger internal motivation. This greater salience may also contribute to the enhanced insula activity.

5.4. *General discussion*

This chapter presented two experiments in which instructed inhibition was based on a 'rule'. Most previous work on inhibition has used an unexpected external sensory signal as a stop signal. The choice of which instructed inhibition paradigm to use is clearly important when comparing instructed inhibition to intentional inhibition. Expectancy and predictability of inhibition clearly differ when comparing unexpected stop signal inhibition to freely decided inhibition is expectancy. When freely deciding to inhibit action, there is an element of expectancy as the moment of inhibition approaches after the initial decision is made. In contrast, during stop signal and go/nogo paradigms, actions must be inhibited rapidly after seeing or hearing a signal which was not entirely predictable (Verbruggen & Logan, 2008). This confound between free choice and predictability has been shown to be important for action studies. In particular, medial frontal activation can be associated with predictability, rather than intentional selection per se (Jenkins et al., 2000). The current experiments control for this confound by using an instructed inhibition condition in which the moment of inhibition is predictable. Instead of unpredictable signals, participants maintain and apply an internalised rule to inhibit action, allowing them to predict and prepare for the moment of inhibition.

Many of the previous experiments in this thesis have highlighted critical differences between intentional action and intentional inhibition. Importantly, the fMRI results from experiment 2 in this chapter have provided evidence of *similarities* between the two functions: first, in the temporal qualities of subjective intention judgements, and second, in the greater pre-SMA involvement in free compared to forced responses. The fMRI data also provides an interesting complement to the experiments in chapter 2. Those experiments, in particular experiment 2.3, showed that perceptual outcomes are evaluated differently for intentional actions and intentional inhibitions. The fMRI experiment in this chapter suggests that outcomes are also evaluated differently for intentional inhibitions and instructed inhibitions, and that insula activity may underlie these differences.

Overall, this chapter has added significantly to what we know about the neural correlates of intentional inhibition. Previous work focused mainly on the dFMC as a key relevant area. The experiments here expanded on this by highlighting a potential role for the pre-SMA, and demonstrating how inhibition decisions may depend on spontaneous neural activity fluctuations. The chapter proposes a new paradigm for studying intentional inhibition, which controls for predictability. It further identifies potential neural causes of intentional inhibition, and neuroanatomical bases of this function. These areas presumably correspond to the intentional decision to inhibit, and the consequences of this decision, though fMRI lacks the time resolution to distinguish conclusively between these two aspects.

Chapter 6. Tourette syndrome and tic inhibition

Gilles de la Tourette syndrome is characterised by involuntary motor movements (tics). Some patients have the capacity to inhibit their tics voluntarily. This final experimental chapter reports two experiments with Tourette syndrome patients. Experiment 6.1 studied the effects of attentional manipulation on tic frequency, and found that focusing attention away from tics significantly reduced tic occurrence. Experiment 6.2 examined the neural correlates of tic inhibition, and the interactions between tic inhibition and voluntary action behaviour. It showed that tic inhibition directly affects oscillatory brain activity around the time of voluntary action. This chapter highlights the relevance of intentional inhibition in disorders such as Tourette syndrome, and shows how these inhibitory processes affect involuntary actions.

6.1. Introduction

The experiments presented so far have primarily focused on intentional inhibition in healthy populations. Patient groups often present unusual instances of inhibition, and studying these can be informative in developing new treatments and in improving our understanding of inhibitory mechanisms.

Tourette syndrome (TS) is a neuropsychiatric disorder in which patients present multiple motor tics and at least one phonic tic for more than one year with onset before the age of 18 (4th ed., text rev.; DSM-IV-TR; American Psychiatric Association, 2000). Tics typically start around 6-7 years of age (Freeman et al., 2000; Robertson, 2000). The status of these tics in the context of volition is unclear: although they are often preceded by an involuntary and unwanted urge to move (Jackson et al., 2011), they can also be quite well suppressed by voluntary inhibition (Ganos et al., 2012; Peterson et al., 1998; Serrien et al., 2005). For this reason, there is disagreement about whether tics should be described as ‘voluntary’ or ‘involuntary’.

Tic inhibition is a common strategy used to suppress tics (Ganos et al., 2012, Specht et al., 2013), and TS patients may in fact maintain tonic inhibition prior to movement, in order to prevent disruption by tics (Heise et al., 2010). The real-world choice to suppress tics can be considered as intentional inhibition. Patients make an endogenous self-driven decision

to prevent involuntary motor urges from emerging as overt movements. The following experiments investigate this capacity to intentionally inhibit tics. Experiment 1 compares the tic reductions achieved by such conscious control with reductions achieved by attentional manipulations. Experiment 2 examines the neural correlates of tic inhibition.

6.2. Experiment 6.1: The role of attention in tics

6.2.1. Introduction

Tics are repetitive context-independent actions. Although they occur persistently in Tourette syndrome, their overall frequency can vary depending on external and internal conditions. Stress and anxiety are known to exacerbate tics (Conelea & Woods, 2008; Robertson, 2000), but the reasons for this interaction between state of arousal and tic generation remain unknown. Cognitive behavioural theories postulate ‘vicious cycle’ processes for anxiety disorders, in which excessive attention and sensitivity to anxiety-related body signals leads to increasing symptoms of anxiety (Clark, 1986; Taylor et al., 1992). We propose that a similar mechanism may underlie tic occurrence: excessive attention to tics, leads to a cycle of increasing tic frequency. This hypothesis has yet to be tested experimentally, to our knowledge. Here we report an experiment in which patients’ focus of attention was manipulated either towards their tics or to other stimuli while they performed a voluntary motor task. We investigated whether tic frequency varied with attentional focus, predicting that diverting attention away from tics would prevent enhancement of neuromotor noise and therefore reduce tics.

6.2.2. Methods

6.2.2.1. Participants

16 adult Tourette syndrome patients (mean age = 31 +/- 10.2 SD; 15 male) participated in the experiment with ethical approval. All procedures were in accordance with the Declaration of Helsinki. Patients were recruited in the Department of Paediatric and Adult Movement Disorders and Neuropsychiatry, Institute of Neurogenetics, University of Lübeck, Germany. All patients were diagnosed with Tourette syndrome according to DSM-IV-TR criteria (American Psychiatric Association, 2000). No patient fulfilled DSM-IV criteria for a

diagnosis of obsessive-compulsive disorder (OCD).

GTS symptom severity for the last week before testing was assessed by a clinician using the Yale Global Tic Severity Scale (YGTSS) (Leckman et al., 1989). Premonitory urges were measured using the validated German version of the “Premonitory Urge for Tics Scale” (PUTS) (Rössner et al., 2010; Woods et al., 2005), which assesses the quality, as well as the severity of bodily sensations preceding tics. Symptoms of Attention-deficit hyperactivity disorder (ADHD) were rated on the German short version of the “Wender Utah Rating Scale” (WURS-K) (Ward et al., 1993). Symptoms of obsessive-compulsive disorder (OCD) were rated on the “Yale-Brown Obsessive Compulsive Disorder Scale” (Y-BOCS) (Goodman et al., 1989).

At the time of the study, all patients reported having motor tics and an additional 8 reported having vocal tics. The mean YGTSS total tic severity was 16.8 +/- 7.7 SD, the mean YGTSS motor tic severity was 12.4 +/- 3.5 SD, and the mean YGTSS vocal tic severity was 4.4 +/- 5.5 SD. All patients reported premonitory urges. The mean PUTS score was 24 +/- 5.7 SD. Three patients were taking medication for their tics. Values of the Y-BOCS ranged from 0-11 (overall cut-off for OCD = 16), the mean total Y-BOCS score was 2.3 +/- 3.9 SD. Values of the WURS-K ranged from 0-44 with a mean of 16.8 +/- 12.1 SD. According to the WURS-K cut-off value of 30, 3 patients scored in the clinical range but only one of those patients fulfilled DSM-IV criteria for ADHD.

6.2.2.2. *Apparatus and materials*

The experiment was run on a computer. A video camera recorded the head and shoulders of each patient. The recordings were used for counting tics. Finger pressure sensors were attached to each of the four fingers of their dominant hand. These provided a digital output each time one of the four fingers was opposed to the thumb.

6.2.2.3. *Design and procedure*

The key features of the design are shown in figure 6.1. Patients were asked to oppose a finger that they freely chose against the thumb of their dominant hand. They chose anew, and repeated the finger opposition approximately once every 2 seconds. They were asked to avoid

using the same finger for consecutive actions, and to avoid adopting specific strategies or patterns of responses. Instead, patients were asked to decide spontaneously on a new finger for each action. Each trial lasted 1 minute, giving approximately 30 actions per trial. An initial practice trial used an auditory metronome at a rate of 0.5Hz to allow patients to synchronise their pinching actions and become familiar with a 2 second rhythm. Practice data was not recorded or analysed.

Each time patients performed an action, a large coloured circle immediately appeared in the centre of the laptop screen for 750ms. There were four possible colours (blue, green, red, and orange). The sequence of colours was random, except that a colour could not be repeated twice in a row. The finger used for each action, and the colour seen on-screen were completely independent.

In each 1 minute trial, a 100ms auditory tone was presented 250ms after the onset of some colours. This tone instructed participants to remember a pre-specified item. Each trial contained 3, 4, or 5 such tones, randomly and equiprobably. The timing of tones within trials was random, so the to-be-remembered items occurred at random times within each trial. The to-be-remembered items depended on the condition. Patients were asked to focus on, and remember, either the finger they used when they heard the beep (finger condition), the colour they saw when they heard the beep (colour condition), or whether they had produced a tic between the previous colour and the current tone during the approximate 2 second interval (tic condition).

In summary, a voluntary motor finger-opposition task was performed in all conditions. Conditions only differed in requiring patients to focus their attention on perception of their fingers, external colours, or their own tics. A memory task was used to verify that participants complied with the task instructions. Attention conditions were blocked so patients always knew which events to attend to on each trial. They were asked to remember the correct item for every “remember this” tone, in each trial, in the correct order.

At the end of the trial, patients were first asked how many tone memory cues had occurred. The next question varied according to condition. In the finger attention condition, they were asked to report the finger used when each beep occurred during the trial. They responded motorically, by opposing the appropriate fingers against the thumb in the correct order, to reconstruct the memory sequence. In the colour attention condition, they were

instead asked about the colours and entered their responses by typing ‘r’ for red, ‘g’ for green, ‘b’ for blue, or ‘o’ for orange on the computer keyboard, again in the order of their occurrence. In the tic attention condition, they typed a sequence of ‘j’ for yes, and ‘n’ for no, to indicate whether or not they produced a tic before each beep presented. After entering their responses, they continued with the next trial after pressing return on the keyboard.




Figure 6.1. Key features of the experimental design. Patients opposed one finger of their choice, against the thumb, every 2s. Each opposition produced a colour on-screen. When random auditory “remember this” cues occurred, they had to remember a corresponding item. This was the finger moved, the colour displayed, or the occurrence of a tic, according to condition.

In addition, we manipulated tic control by instruction. Patients performed each of the 3 conditions (a) while being instructed to voluntarily inhibit their tics to the best of their ability, and (b) without such instruction, so they could tic freely. This created a 3x2 repeated measures design, in which the first factor was the attentional manipulation and the second factor was instructed tic control. Nine 1-minute trials were performed in each condition of this design, giving a total of 54 trials across the experiment. The main outcome measure was tic frequency in each condition. Our main research questions were: whether attentional focus influenced tic frequency, and whether this mechanism was related to voluntary tic inhibition, or was rather independent of it.

6.2.2.4. *Tic counting*

Before the main tasks in the experiment, each patient was video recorded at rest for 1 minute, without instruction regarding tics. In addition, a further 1 minute recording was made after giving the instruction to voluntarily inhibit tics. These baseline conditions were always performed in this order. This was to avoid any potential ‘rebound’ effect of increased tics following tic inhibition which patients often report (Verdellen et al., 2007). The number of tics in these trials was used as a baseline measure of tics, to compare with tic frequency during trials in the experiment.

Two experimenters independently counted tics in all videos for all conditions. The mean of their independent counts was used in all analyses. These experimenters were blind to all experimental conditions during tic counting but not to experimental hypotheses or design. A third rater naïve to all aspects of the experimental design including hypotheses and conditions independently counted a subset of tic videos (the first five trials from one condition drawn at random from each patient’s data). Their count was used to test inter-rater reliability with the counts used for analysis.

6.2.3. *Results*

Behavioural data

Patients generally had no difficulties in maintaining the 2 second rhythm with finger pinching actions. The mean interval between their finger movements was 1.79s (average standard deviation over movements of 0.30s). Patients also avoided using the same finger consecutively in 96.21% of their actions, as instructed.

Memory performance was high across all finger and colour conditions (see table 6.1). A 2x2 repeated measures ANOVA on memory performance showed no main effect of attentional focus, $F(1,15) = 2.9$, $p = 0.11$, no main effect of tic inhibition, $F(1,15) = 0.02$, $p = 0.88$, and no significant interaction, $F(1,15) = 0.56$, $p = 0.46$. This suggests that finger and colour tasks did not differ significantly in difficulty. Measures of objective memory performance were not possible for the tic attention task for several reasons. First, patients’ perceptions of their own tics may be very different to observers’ perceptions (Müller-Vahl et al., 2014; Pappert et al., 2003). Second, patients may perceive tics in body regions not recorded by the video camera. Third, any false positives and misses by the experimenter

during tic counting would greatly influence measures of memory performance for tics. For these reasons, a reliable estimate of memory for tics based on matching tic rater and patient reports is impossible in the context of the current experiment.

	Finger	Colour
Freely tic	79%	89%
Tic inhibition	82%	86%

Table 6.1. Percentage of correctly recalled items in the finger and colour tasks during freely tic and tic inhibition conditions.

Tic frequency

A strongly significant correlation between independent raters' tic counts was found ($r = 0.94$, $p < 0.001$, see figure 6.2), suggesting that they produced highly similar data following tic counting.

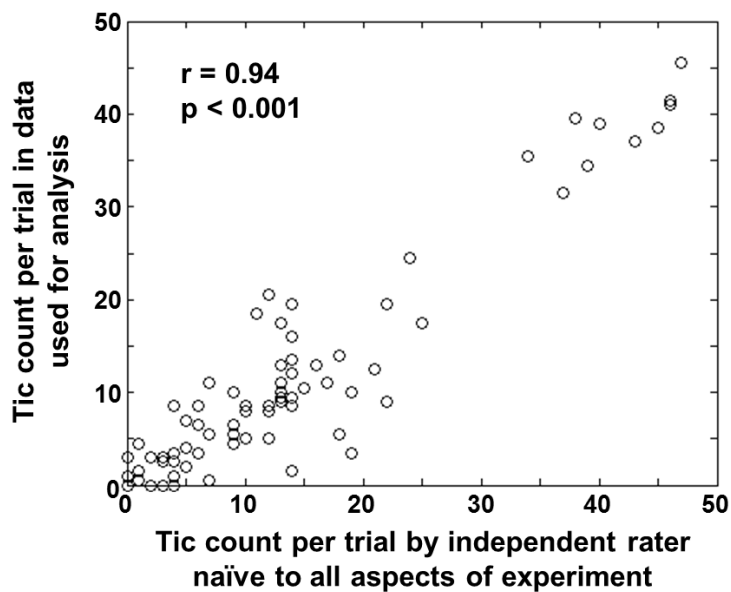


Figure 6.2. Inter-rater reliability between rater counts used for analysis and independent rater naïve to all aspects of the experiment. Each datapoint is the tic count for a single trial across a subset of 80 trials. Ratings on the y-axis show the mean of independent counts performed by authors EM and VB. Ratings on the x-axis represent counts performed by an independent rater not involved in the experiment.

The mean number of tics per trial in each task condition is shown in figure 6.3. A 2x3 repeated measures ANOVA showed a significant main effect of attention, $F(2,30) = 6.54$, $p = 0.004$, and a significant main effect of tic inhibition, $F(1,15) = 9.27$, $p = 0.008$. A Mauchly's test of sphericity indicated a significant violation for the attention x inhibition interaction ($p < 0.05$). After applying a Greenhouse-Geisser correction, a significant interaction was found, $F(1.41, 21.09) = 5.31$, $p = 0.02$. We explored this interaction with simple effects t-tests. These results are shown in figure 6.3.




Figure 6.3. Mean number of tics per trial for each task condition and baseline. Black connecting lines indicate significant differences between freely tic conditions, and grey connecting lines indicate significant tic reductions when inhibiting tics within an attention condition (* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$). The freely tic baseline had significantly greater tics than all within-task freely tic conditions (dashed lines, $p < 0.05$, Bonferroni corrected). There were no differences between baseline tic inhibition and within-task tic inhibition conditions. Note that no difference was found between freely tic and tic inhibition states in the finger attention condition.

When allowed to tic freely, patients produced the highest number of tics in the tic attention condition during task performance. Tics were significantly reduced when patients focused their attention on colours compared to tics, $t(15) = -2.16$, $p = 0.047$ (effect size: Cohen's $d = 0.34$). Tics were reduced even further when focusing attention on fingers relative to colours, $t(15) = -2.15$, $p = 0.048$ (Cohen's $d = 0.24$). Interestingly, tic frequency was at an identical level for all attention conditions when patients were asked to inhibit their tics. When comparing to the passive baseline in which patients performed no task, it can be seen that tic inhibition during tasks reduced tics to an 'inhibition baseline' level. Individual t-tests between the inhibition baseline and each of the three tic inhibition task conditions, showed no significant differences even without correcting for multiple comparisons ($p > 0.05$). In contrast, Bonferroni corrected t-tests showed that the baseline tic frequency during free tics was significantly higher than all three of the free tic task conditions ($p < 0.05$). This suggests that the attentional demands of any general task performance (e.g. maintaining a steady rhythm in actions; remembering items) reduced tics to some extent. The important point is that the level of tic reduction was significantly affected by the focus of attention.

6.2.4. Discussion

Participants were asked to attend to, and count, specific events. The nature of the events attended strongly influenced tic frequency. The results highlight the important role that attention plays in the presentation of tics. Simply engaging attention in a task reduced the frequency of tics from a baseline level with no task. This is consistent with a general distracting effect of any cognitive task. However, the specific content of attention strongly influenced tic frequency. Tics were most frequent when attending to tics, least frequent when attending to voluntary finger movements, and showed an intermediate level when attending to an external stimulus caused by the participant's voluntary finger movements.

Attention and motor noise

Attention enhances neural signals in the pathways and areas associated with the relevant signals (Kastner et al., 1998; Moran & Desimone, 1985; Murray & Wojciulik, 2004; Saalman et al., 2007; Somers et al., 1999; Treue & Maunsell, 1996). Neural noise is naturally present in the motor system and can affect motor planning and movement (Hamilton et al.,

2004; Harris & Wolpert, 1998). However, normal context-embedded actions depend on voluntary planning and intention, rather than motor noise. We propose that motor noise in TS patients is enhanced, and that such noise processes contribute to tic generation. Excessive attention to these involuntary motor processes and tics could therefore enhance tic generation.

Tics are very common in young children with prevalence estimates up to 18% in the early school years (Ludolph et al., 2012; Robertson, 2000). Tics (i.e. superfluous and repetitive context-independent actions) might be a reflection of increased motor noise and perhaps the product of an immature motor system going through a ‘tuning’ process. They are often not noticed and not troublesome, and typically subside within a few months. However, in some children and adolescents, tics persist and may cause distress. When motor and phonic tics are present for more than a year, this is referred to as Tourette syndrome. Why tics persist in some patients is unclear. Some possible factors have been proposed, for example streptococcal infection (Allen et al., 1995; Mell et al., 2005; Toufexis et al., 2014), but none have been conclusive or widely accepted. The role of attention in tic persistence has not been studied systematically. The results of the present study show that attention enhances tic generation. It may further be conceivable that increased attention to tics by relatives, teachers and peers, could also promote tic persistence, by directly influencing how much the patient pays attention to his/her tics.

The results of the present study show that distraction of attention away from tics reduces tic frequency. Tic reductions were found in all attention conditions relative to baseline. Task demands such as retaining information in memory or maintaining a steady action rhythm, are sufficient to divert attention from tics to some extent. These were common to all tasks, including the tic attention task. However, the extent of distraction could be improved by manipulating the specific object of attention. Selective attention to the colour of the stimuli presented on the screen meant that attention was focussed on external stimuli, and allocated away from tics. We propose that this prevented attentional facilitation of tic-related neural signals (Woodman & Luck, 2003). The distraction provided by attending specifically to finger actions during the task appeared to exhibit even greater tic reduction benefits. This may be because attention is focused on voluntary as opposed to involuntary actions. Some have argued that the boundary between voluntary and involuntary actions becomes blurred in TS (Cavanna & Nani, 2013). Our data suggest that attention to voluntary action generation may lead to a stronger separation of the two systems. Attention to voluntary actions may draw attentional resources away from involuntary movement. This focus on the voluntary

stream of action may further inhibit any contribution to overt action from the involuntary stream.

Previous models have suggested multiple, dissociable cortical streams for action control. For example, internally-generated and externally-triggered actions are thought to involve different pathways (Jenkins et al., 2000). However, the pathway for tic generation remains controversial (Ganos et al., 2013). Our results argue against the view that tic generation originates in the voluntary motor pathway. Had that been the case, attention to voluntary action and attention to tics should have comparable effects on tic frequency. In fact, a clear dissociation was found. Thus, our results support a dissociation, and even an inhibitory link between the voluntary and tic pathways, which can be modulated by attention. Recent neuroimaging findings propose a similar inhibitory interaction between tic inhibition and voluntary action pathways (Ganos et al., 2014; Thomalla et al., 2014)

Tic inhibition

The instruction to voluntarily inhibit tics was generally effective consistent with previous studies (Ganos et al., 2012; Peterson et al., 1998; Serrien et al., 2005). This instruction brought tics down to the same low baseline level in every attentional condition, despite large differences between attentional conditions in the number of tics without inhibition. This suggests that tic inhibition is applied at the final motor output stage of tic generation. Any prior influences on the tic generation process might be simply cancelled out by a ‘cut-off’ mechanism which acts to block the output stream and reduce or prevent overt tic expression.

Although tic suppression is successful in reducing tics, it may not remove the urge to tic. Overall, there is no correlation between trait level of urge to tic as measured by PUTS and tic inhibition (Ganos et al., 2012; Müller-Vahl et al., 2014). However, tic inhibition in some patients may in fact produce a continuing intensification of the immediate urge to tic (Himle et al., 2007). The mechanism of tic reduction by attentional focus is likely to be quite different to tic inhibition. We have suggested that attention may operate by modulating noise at the tic generation stage. In contrast, we suggest that inhibition is applied at a later stage, gating the output of the generator. Therefore, reductions in tic frequency through attentional manipulation seem to occur without the need for effortful inhibition of tics. Note that the number of tics during attention to voluntary finger movements in the current task was reduced to the same extent as active tic inhibition (see figure 6.3). We show that patients may

be able to obtain tic reduction benefits as strong as those with tic inhibition, without the need for continuous internal monitoring and active tic prevention.

Attention in other disorders

Several other psychiatric disorders involve an attentional component. Anxiety and panic disorders involve excessive attention to internal body signals and misinterpretation of normal bodily processes and sensations (Clark, 1986; Taylor et al., 1992). Psychogenic disorders have also been characterised by maladaptive attentional processes (Edwards & Bhatia, 2012; Edwards et al., 2012). For example, attention to abnormal prior beliefs about bodily sensations and movements excessively weights their contribution to perception and action, creating the functional symptoms of movement disorders (Edwards et al., 2012). Psychogenic movement disorders are of course very different to TS, both in aetiology and in phenomenology. However, our data suggest that maladaptive attention in TS may exacerbate tic symptoms. For example, it might augment random fluctuations in motor noise, generating urges to tic that might not occur in the absence of attention.

Implications for behavioural treatment

The data we present here has several implications for behavioural treatments of TS. One important current behavioural treatment is habit reversal therapy (HRT) (Bate et al., 2011; Deckersbach et al., 2006; Himle et al., 2006; Wilhelm et al., 2003). This therapy aims to reduce the association between urge and tic by encouraging patients to become more aware of their tics as they occur, and teaching them to produce a movement which is incompatible with the urge to tic. It appears to work through improving the ability to control and inhibit tics (Bate et al., 2011). In contrast to HRT, our results might suggest a treatment based on attentional distraction. This putative therapy would minimise attention to tics, with the aim of reducing the urge to tic in the first place, rather than improving the control of overt tics. Instead of teaching patients to become highly aware of the warning signs of a tic, as HRT does, patients would be encouraged to focus their attention on external events and fully voluntary actions and intentions. By employing these distraction techniques under conditions in which they are most likely to tic, such as anxious and stressful social situations, patients may be able to reach a level of tic frequency similar to effortful tic inhibition without having to actively control and monitor involuntary urges. This putative therapy remains speculative, but it is clearly suggested by the evidence presented here.

Limitations

We lacked any subjective self-report measures of online urge during the task from patients. We have suggested that attentional distraction reduces urges to tic, but we have no direct measure to show this. However, the way in which tic frequency changed across attentional conditions was very different from the way it changed with tic inhibition. Tic inhibition consistently reduced tics to the same level, whereas attentional distraction had a more graded effect depending on the quality of the distraction. We believe there are two possible ways to reduce tics: either actively inhibiting them, or having a reduced urge to tic in the first place. While tic inhibition involves the first mechanism, attentional distraction should involve the second.

A second potential limitation is the possible confound of task difficulty. We could not directly measure task difficulty in the tic attention condition due to the inability to accurately categorise a patient's tic judgements as correct or incorrect (see Results). However we can certainly rule out task difficulty as an explanation of the difference in tic frequency between finger and colour conditions. Objective measures of task performance did not differ between these conditions. Therefore attention alone can affect tic production, irrespective of the difficulty of the attended task. Although the tic attention task may be easier overall than finger and colour tasks, task difficulty alone cannot explain our data on changing tic frequency with attention.

Conclusions

In conclusion, we provide evidence that tic generation is strongly influenced by attention. When patients attend directly to tics, they exhibit a higher number of tics than when attending to other events or objects. An active desire or instruction to stop tics further facilitates attention to tics, potentially creating a cycle of increasing symptoms. This process may also explain the reported increase in tics during stress outside the laboratory (Conelea & Woods, 2008), since tics may become the focus of attention in such situations, particularly when there is some social relevance. We propose that directing attention away from tics may improve the signal-to-noise ratio in the motor system. This ultimately means that patients tic less than they normally would, without the need to use effortful tic suppression.

6.3. *Experiment 6.2: Intentional action and tic inhibition: an EEG and psychophysical study in Tourette syndrome*

6.3.1. *Introduction*

The Libet task (Libet et al., 1983) is frequently used to measure the experience of voluntary intention before action. A recent experiment found that the experience of intention is delayed in Tourette syndrome (Moretto et al., 2011). Instead of the typical report of intention approximately 200 ms before the judged time of action in healthy controls, Tourette syndrome patients report the time of intention as only around 50 ms before the time of their action judgement (Moretto et al., 2011). This suggests that the experience of intention is less distinguishable than normal, and it often coincides with the experience of actual action. The following experiment will use EEG to investigate the neural correlates of this delayed experience of volition in Tourette syndrome.

Another important question for Tourette syndrome is how volition and intentional action interact with tic inhibition. Some evidence suggests that during task performance, top-down inhibitory processes suppress tics so that task-related voluntary behaviour is effective (Heise et al., 2010). The neural mechanisms underlying tic inhibition and its interaction with the voluntary action system remain relatively unknown. This study will directly examine the neural correlates of tic inhibition during a voluntary action task, using both event-related potential (ERP) and spectral power (ERSP) approaches.

6.3.2. *Methods*

6.3.2.1. *Participants*

Ten TS patients were recruited from the TS specialty out-patient Clinic in the Department of Neurology, University Medical Center Hamburg-Eppendorf. In addition, ten healthy age and gender-matched volunteers were recruited as a control group. TS patients and healthy volunteers were tested with ethical permission, and in accordance with the principles of the Declaration of Helsinki.

6.3.2.2. *Design and procedure*

The experimental tasks were based on those used previously (Libet et al., 1983, Moretto et al., 2011). Participants made self-paced voluntary keypress actions, while watching a clock hand rotating every 2560 ms. The clock was small enough to be seen without eye movement:

the hand was 12 mm in length. The experimenter initiated the rotation of the clock on each trial, and invited the participant to make a voluntary action at some point after the first rotation of the clock hand. The right hand was used for the keypress action. Participants were particularly instructed that they should make the actions at a time of their own choice, and not in response to any specific time indicated on the clock. After each keypress, the clock hand stopped after a random interval. Participants verbally reported the clock time at which they pressed the button (M judgement), or at which they first experienced the conscious intention to move (W judgement), in separate blocks of trials. In the W condition, participants were invited to report the moment of “the urge to perform the movement” (“Als Sie den Drang, die Bewegung auszuführen gespürt haben”). M and W judgements were tested in two separate blocks. Blocks were presented in a randomised order between subjects. Ten training trials for each type of judgment were run before the experimental session. Each block consisted of 40 voluntary action trials. The duration of trials varied as a function of when participants chose to act, but was typically between 3 and 10 s.

The patients, but not the volunteers, performed two further blocks in which they were asked to voluntarily suppress their tics during the period of each trial (i.e., while the clock was rotating).

6.3.2.3. *Measurements and EEG processing*

The judged times of M and W relative to the actual keypress were stored on a computer. In addition, continuous video recording of each patient was used to count the number of tics occurring during each trial. One patient showed numerous hand tics which, due to technical error, often fell outside the video frame: therefore tic count data was available for only 9 patients.

EEG was recorded continuously with a 64-channel passive electrode set-up using a BrainAmp system. Data was sampled at 1000Hz and downsampled to 100Hz during preprocessing. Online filtering was applied with 0.1Hz high pass and 30Hz low pass cut-offs. Data was referenced offline to an electrode attached to the nose. EOG electrodes were placed on the supra-orbital area and external canthus of the right eye. EMG was recorded from the right hand first dorsal interosseus (FDI) muscle with two electrodes in a belly-tendon montage. Only the core subset of electrodes most relevant for movement-related potentials were analysed here, since our primary interest was in component strength rather topography

or source localisation. EEG processing was performed in EEGLAB (Delorme & Makeig, 2004). The onset of voluntary EMGs in the right FDI muscle associated with each keypress action was marked by a cursor. Epochs from 1500 ms before to 1000 ms after the voluntary EMG onset were extracted. The epoch length chosen had a shorter premovement period than many previous RP studies (e.g. Haggard & Eimer, 1999), because longer epochs are more likely to contain tic-related artefacts, which could make EEG data less comparable between patients and healthy controls. Thus, we chose the shortest epoch compatible with established RP measures (Shibasaki et al., 1980). However, we cannot exclude that we have missed effects earlier in the sequence of events that precedes voluntary action. Each epoch was baseline corrected using the mean of the data from 1500 to 1400 ms prior to the EMG onset. Each epoch was inspected to check for obvious artefacts and EEG quality. Next, independent component analysis (ICA) was used to detect blink artefacts (Delorme & Makeig, 2004; Delorme et al., 2007), by identifying components with a strong frontal topography and removing them from the data. In addition to ERP values, spectral properties of the EEG were measured by plotting the EEG signal power as a function of frequency and time. Because the spectral power computation requires a minimum window width to compute oscillatory power, the epoch for EEG spectral analysis could not be calculated for the first and last 500 ms of the selected epoch. Therefore, EEG results are reported for the period -1000 to +500 ms relative to EMG onset. The baseline was taken from -1000 to -800 ms relative to the movement onset.

6.3.3. Results

Voluntary action in TS and volunteers

Our first analyses used a factorial design to compare M and W judgement blocks between patients and controls.




Figure 6.4. Judgements of intention and action in TS patients and volunteers (* $p < 0.05$). Error bars show standard deviation across participants.

The behavioural results are shown in figure 6.4. ANOVA on the psychophysical judgement data revealed an unsurprising significant main effect of judgement type, with W judgements being significantly earlier than M judgements ($F(1,18)=31.717, p<0.001$). There was also a significant main effect of group ($F(1,18)=9.574, p=0.006$), with TS patients making later judgements than volunteers. Most interestingly, there was a significant interaction between these factors ($F(1,18)=4.593, p=0.046$). Post hoc simple effects showed that this arose because the experience of intention differed between the two groups ($t(18)=2.718, p=0.014$), whereas the experience of action itself did not ($t(18)=1.471, p=0.159$).

The grand average readiness potentials recorded at key electrodes are shown in figure 6.5. Following previous research (Shibasaki & Hallett, 2006), we measured the mean evoked amplitude at electrode Cz, for early (-1000 to -500 ms) and late (-500 to 0 ms) RP components. Analysis of the early RP amplitude revealed no between-subjects effect of group ($F(1,18) = 2.45, p = 0.14$), no effect of judgement condition ($F(1,18) = 0.37, p = 0.55$), and no significant interaction between judgement condition and group ($F(1,18) = 2.34, p =$

0.14). Analysis of the late RP also revealed no significant effect of group ($F(1,18) = 2.14$, $p = 0.16$), judgement ($F(1,18) = 0.11$, $p = 0.75$) and no interaction between judgement and group ($F(1,18) = 2.26$, $p = 0.15$).

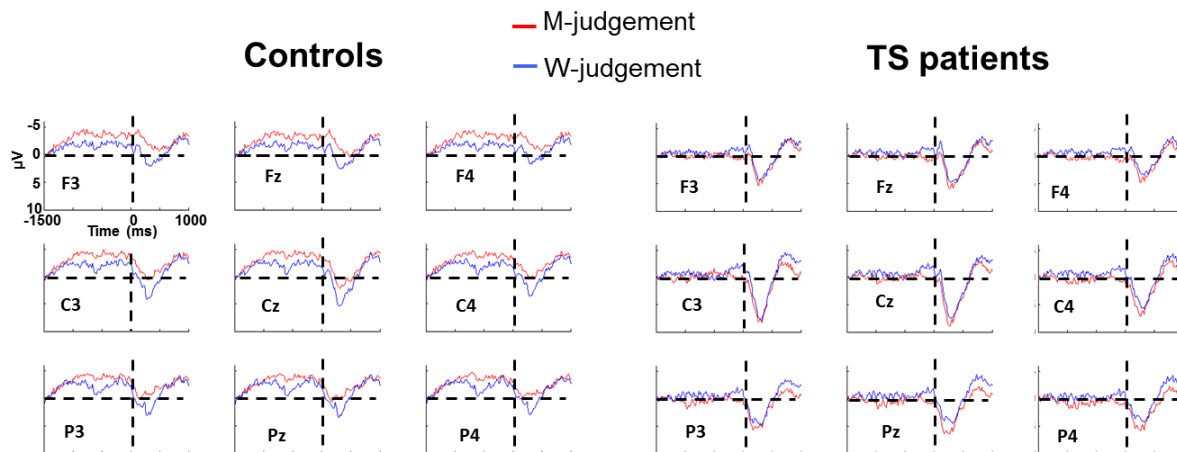


Figure 6.5. Readiness potentials for volunteers and patients in M and W judgement conditions. Vertical line at 0 ms corresponds to voluntary EMG onset.

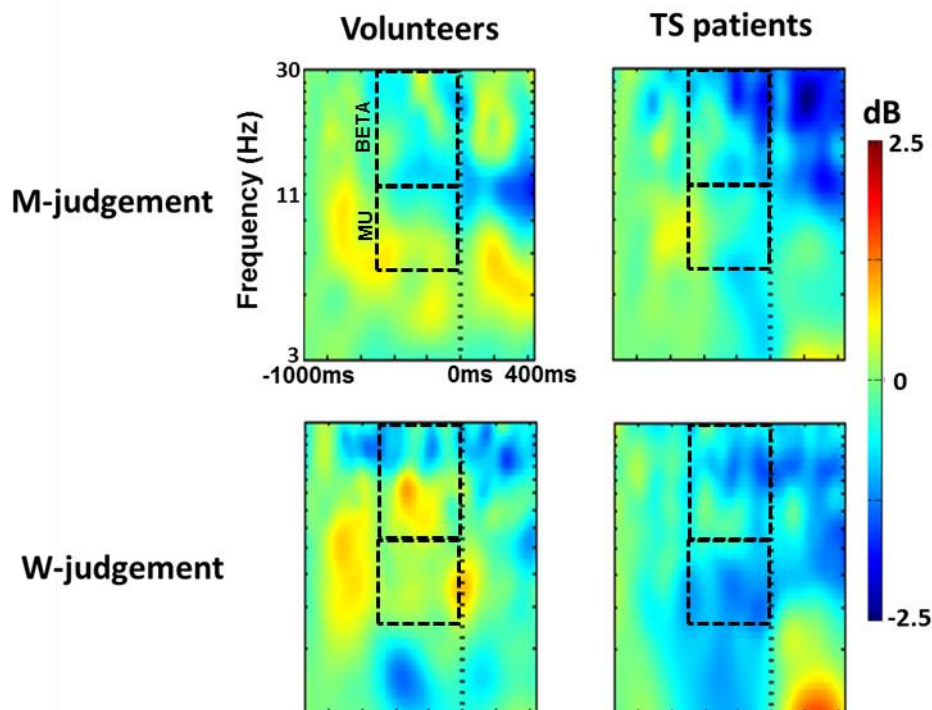


Figure 6.6. Event-related spectral power (ERSP) for volunteers and patients in M and W judgement conditions, shown as frequency x time plots. Vertical line at 0 ms is voluntary EMG onset. Blue represents event-related desynchronisation (ERD) while red represents event-related synchronisation (ERS).

Figure 6.6 also shows the characteristic event-related desynchronisation (ERD) around the time of voluntary action. The periods selected for EEG analysis could not precisely match those used for readiness potential analysis, because of the minimum window-width required for calculating the EEG spectrum (see methods). Based on previous reports, we focussed on event-related desynchronisation in the Mu (8-12 Hz) and beta (12-30 Hz) bands at contralateral motor electrode C3, for the epoch from -500 to 0 relative to EMG onset, corresponding to the late RP (Pfurtscheller, & Lopes da Silva, 1999). Results in the mu band showed no main effect of group ($F(1,18) = 3.40, p = 0.08$), or judgement type ($F(1,18) = 0.02, p = 0.90$). The interaction between these factors was significant however ($F(1,18) = 4.43, p = 0.046$). Simple effects analysis showed that this interaction occurred because the mu rhythm power was lower (i.e., stronger ERD) for the patients than for the controls in the W-

judgement condition, $t(18) = 2.58$, $p = 0.02$, while there was no difference in the M-judgement conditions, $t(18) = 0.13$, $p = 0.90$. Results in the beta band showed no significant main effect of group ($F(1,18) = 2.24$, $p = 0.15$), judgement condition ($F(1,18) = 1.55$, $p = 0.23$) and no interaction between judgement and group ($F(1,18) = 0.23$, $p = 0.64$).

Effects of voluntary tic suppression on voluntary action in TS

For the patient group only, we compared voluntary actions made with vs. without the additional instruction to inhibit tic behaviours during each trial. First, we performed a quantitative tic count of each patient's video record. The number of tics occurring during each trial, defined as the period between the experimenter initiating clock rotation, and the clock rotation stopping after the participant's voluntary action, was noted. The average number of tics per trial was then analysed using ANOVA with factors of judgement condition (M/W) and suppression instruction ("just let yourself tic"/"please try to suppress your tics"). The data are shown in table 6.2. The results showed a predicted decrease in tic frequency when patients were instructed to suppress their tics compared to when they were not (mean (SE) 0.196 (0.074) and 0.418 (0.089) tic events per trial respectively; $F(1,8)=6.440$, $p=0.035$). There was no significant effect of judgement condition (W or M), and no interaction (both $F<1$, NS).

Temporal judgement data are shown in table 6.2. Judgement data were compared between instructed suppression and no instructed suppression conditions. ANOVA again showed a predicted main effect of judgement type, with W judgements preceding M judgements ($F(1,9)=20.090$, $p=0.002$). There was no significant effect of the instruction to suppress tics on temporal judgements, and no interaction (both $F<1$).

Judgement	Instruction	Mean (SD) tic count (tics/minute)	Mean (SD) temporal judgement (ms)
M	No tic suppression	4.1 (2.6)	-64.1(41.3)
M	Tic suppression	2.2 (2.2)	-68.9 (45.5)
W	No tic suppression	4.5 (3.1)	-176.0 (79.7)
W	Tic suppression	1.8 (2.7)	-178.3 (97.0)

Table 6.2. Temporal judgements (ms) and tic counts for TS patients in M and W judgement conditions during suppression and no suppression of tics.

The effects of instructions to suppress tics are shown in figure 6.7. RPs at Cz were almost indistinguishable for the suppression and no suppression conditions. RP amplitudes at Cz were calculated and analysed as before. ANOVA showed no effects of tic suppression ($F(1,9) = 0.35$, $p = 0.57$) or judgement condition ($F(1,9) = 2.46$, $p = 0.15$), nor any interaction ($F(1,9) = 0.01$, $p = 0.93$), at the early RP time window. Similarly no main effect of suppression ($F(1,9) = 0.41$, $p = 0.54$), judgement ($F(1,9) = 1.68$, $p = 0.23$) or interaction between them ($F(1,9) = 0.03$, $p = 0.86$) could be found in the late RP. Thus, we found no evidence for change in RP depending on whether tics were voluntarily suppressed or not.

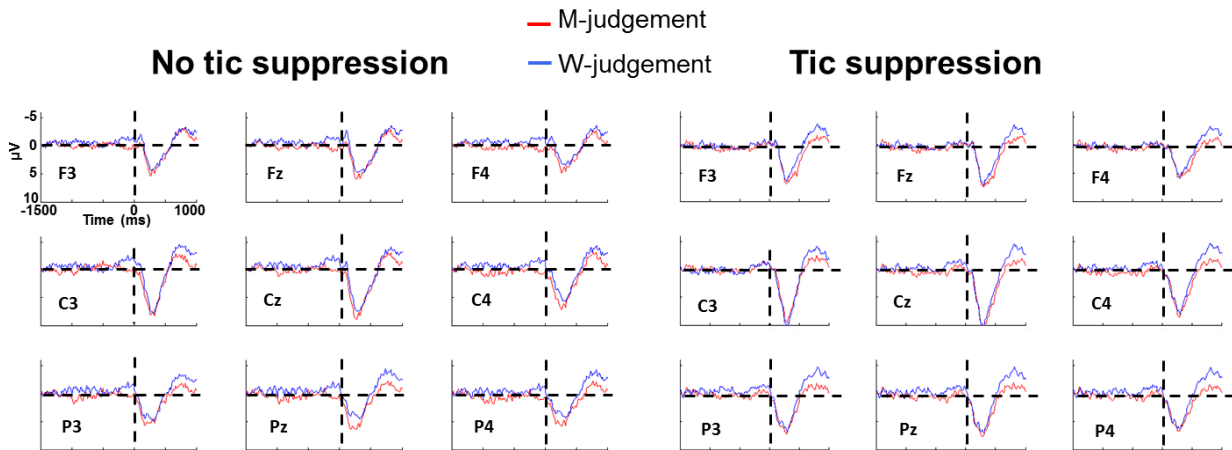


Figure 6.7. Effects of voluntary tic suppression on readiness potentials in M and W judgement conditions. Vertical line at 0 ms corresponds to voluntary EMG onset.

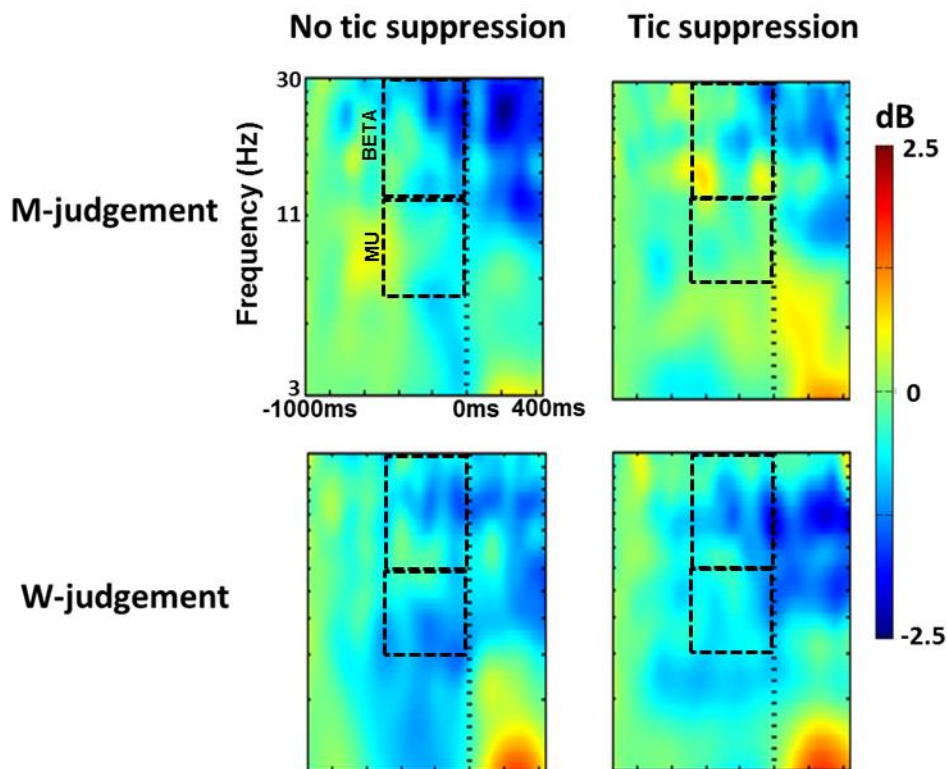


Figure 6.8. Effects of voluntary tic suppression on event-related spectral power (ERSP) in M and W judgement conditions, shown as frequency x time plots. Vertical line at 0 ms is voluntary EMG onset. Blue represents event-related desynchronisation (ERD) while red represents event-related synchronisation (ERS).

The spectral EEG is shown in figure 6.8. The characteristic movement-related ERD was observed only in the M judgement condition. ERD occurred later, and to a lesser extent, when patients were instructed to suppress their tics. ANOVA of the C3 electrode data in the mu band showed no significant effect of tic suppression ($F(1,9) = 0.31, p = 0.59$). However, there was a significant effect of judgement ($F(1,9) = 8.24, p = 0.02$), with lower mu power (i.e., more ERD) for W than for M judgement. There was no significant interaction ($F(1,9) = 0.01, p = 0.94$). In contrast, in the beta band, there was no significant effect of tic suppression ($F(1,9) = 2.98, p = 0.12$), no effect of judgement ($F(1,9) = 1.04, p = 0.33$), but a significant interaction between these factors ($F(1,9) = 6.37, p = 0.03$). Post-hoc simple effects testing showed that this arose because the instruction to suppress tics significantly reduced ERD in the M judgement condition, compared to no instruction ($t(9) = -2.93, p = 0.02$). There was no significant difference for W judgement trials ($t(9) = 0.48, p = 0.64$). Thus, the instruction to suppress tics prevented the change in beta-band oscillatory power that normally precedes voluntary action.

6.3.4. Discussion

We examined EEG readiness potentials and oscillatory power associated with voluntary actions and tic inhibition in Tourette syndrome. We replicated a previously found effect in which Tourette syndrome patients report a delayed experience of volition (Moretto et al., 2011). In addition to this, we found that readiness potential amplitudes did not reflect delayed intention in patients compared to controls, but that mu power displayed an interesting difference between the groups. During judgements of intention, mu power was significantly lower for patients than for controls. During judgements of action on the other hand, no significant difference was present.

The abnormal experience of intention that patients show may be related to abnormal mu oscillations. Event-related desynchronisation (ERD) in mu power (i.e. reduced mu power) is typically related to action processing and action observation (Lepage & Theoret, 2006; Muthukumaraswamy & Johnson, 2004). The stronger mu ERD during intention judgements for patients compared to controls may reflect excessive motor activity during voluntary action, consistent with previous evidence (Franzkowiak et al., 2010). This may drive the abnormally delayed sense of intention in patients, as intention signals are less effectively detected amongst the increased background motor noise.

Although readiness potentials did not vary with tic inhibition, beta oscillations showed weaker desynchronisation when inhibiting tics compared to ticing freely during M judgements. Increased mu/beta power has previously been linked with instructed and intentional inhibition of actions (Klimesch, Sauseng, & Hanslmayr, 2007; Swann et al., 2009, 2012; Walsh et al., 2010; Yamanaka & Yamamoto, 2009). Our data suggests that inhibition of involuntary actions in Tourette syndrome leaks into the voluntary action system, increasing pre-movement beta power. This effect was not found during W judgements, suggesting that attentional factors play a strong role in the neural activity related to tic inhibition (see also experiment 1 in this chapter for the role of attention in tic generation). When patients focus on their voluntary actions, the brain activity preceding the actions is strongly affected by the attempt to suppress tics. This is consistent with findings showing a functionally beneficial interaction between tic inhibition and voluntary task performance in Tourette syndrome (Heise et al., 2010). When the focus is instead on preceding intentions rather than actions, such differences in inhibition-related activity are less pronounced.

In conclusion, TS patients' delayed experience of volition may be driven by increased motor noise during action execution. This makes detection of intentional signals more difficult. The inhibition of involuntary tics also increases beta activity preceding voluntary actions, suggesting that inhibition is rather global in its functioning (Poo & Isaacson, 2009), rather than specifically tuned to involuntary tics in TS. This is in line with previous evidence demonstrating significant changes in general inhibitory control in TS (Raz et al., 2009).

6.4. *General discussion*

This chapter explored inhibition in the context of TS. Patients with this disorder often try to suppress tics during their everyday life. Their decision to inhibit involuntary movements is frequently due to stress and anxiety, in particular within social situations (Conelea & Woods, 2008), but is almost always an intentional self-generated decision with no immediate external instruction. Such inhibition is likely to be challenging for patients due to the continued effort it requires.

Experiment 1 in this chapter demonstrated that tic inhibition can be very successful in reducing the overt expression of tics, consistent with previous evidence (Ganos et al., 2012; Specht et al., 2013). However, it also showed that simply directing attention away from tics can reduce tics to levels comparable with active tic inhibition. This has important

implications for treatment and suggests that attentional control strategies may be very useful in helping everyday functioning for patients. Experiment 2 showed that the neural signature of tic inhibition may be increased beta activity around the time of voluntary action.

This chapter has highlighted the important role that intentional inhibition plays in neuropsychiatric disorders such as TS. It has shown that research into inhibitory function can contribute significantly to our knowledge of these disorders, but also that such research can help in developing new practical treatments. An important message from this chapter is that intentional inhibition is an active and effortful process, which can be used to counteract urges to act. Maintaining constant inhibition can be difficult, as shown by the work on 'ego depletion' (Baumeister, Vohs, & Tice, 2007; Muraven & Baumeister, 2000). Therefore, understanding how inhibitory mechanisms operate in disorders can be useful in two important ways: 1) it can help to build stronger models of how inhibition relates to motor symptoms; 2) it can lead the way in comparing how alternative symptom relief approaches compare with inhibitory strategies.

Chapter 7. General discussion

This thesis has explored intentional inhibition and its broad range of associated processes and constructs. It has demonstrated that perceptual decision-making and perceptual processing depend strongly on whether intentional action or intentional inhibition systems are in operation. It has shown that our previous history of behaviour influences our future decisions to inhibit action, and that asymmetries in action and inhibition capacities can lead to dramatic differences in behavioural decision-making. Furthermore, the thesis has provided evidence that there are distinct neural processes underlying intentional inhibition and instructed inhibition. It has also looked at the role of intentional inhibition in Tourette syndrome with reference to inhibitory strategies for tic control.

This final chapter will bring together the evidence presented in each of the preceding experimental chapters, in order to present a coherent and complete account of intentional inhibition in humans. It will first identify the critical points that make intentional inhibition unique in comparison to instructed inhibition and intentional action. It will then describe the remaining problems for the concept of intentional inhibition, before outlining a potential model for resolving these problems and developing a stronger descriptive framework.

7.1. Intentional inhibition as a distinct construct

7.1.1. Intentional inhibition vs intentional action

The majority of this thesis has highlighted an important asymmetry between intentional action and inhibition. Intentional actions can be habitual, agentic, and easy or satisfying to execute. Intentional inhibition on the other hand is relatively inert with little effect on the external world. At times, it can also be very difficult to execute, or maintain successfully, over extended periods. As chapter 2 showed, problems framed as action vs inhibition can contain inherent biases in favour of an inhibition response. Symmetrical or ‘fair’ decision-making requires balanced alternatives, usually in the form of one action vs another (Gold & Shadlen, 2007; Shadlen & Newsome, 2001). Chapter 4 provided a demonstration of a real-world competition in which action vs inhibition imbalances in decision-making can lead to dramatic differences in decision outcomes. Chapter 3 illustrated that prior action and inhibition events in immediate behavioural history also have asymmetrical effects on future

decisions. Repeated actions produce a habitual tendency to continue acting, while inhibitions lack this self-reinforcing cumulative effect on future behaviour. Chapter 6 presented the neural correlates and ultimate outcomes of direct conflict between over-active action tendencies and inhibitory control. In Tourette syndrome patients, involuntary movement tends to overpower inhibition, and attentional allocation strategies may perform better in reducing action impulses.

7.1.2. Intentional inhibition vs instructed inhibition

Another critical distinction must be drawn between instructed and intentional inhibition. Chapter 5 explored in detail the different neural mechanisms underlying these functions. It was clear in showing that spontaneous neural activity can drive intentional inhibition decisions, but not instructed inhibition decisions. It also singled out the pre-SMA as a potential area recruited more strongly for intentional than instructed inhibition. Pre-SMA and dFMC (Brass & Haggard, 2007; Kuhn et al., 2009) areas both seem relevant to intentional inhibition, while right IFG may be more specific to instructed and rule-based inhibition. When making a free decision between action and inhibition, activity in the dFMC may code for which of them is chosen, while more downstream pre-SMA activity may code for when the decision should be implemented. Similar models of intentional action also posit a role for the pre-SMA in response timing (Mueller et al., 2007; Kriehoff et al., 2009; Soon et al., 2008). This suggests that decision-making related to action and inhibition may converge on the pre-SMA when it comes to response timing. When moving further upstream, action and inhibition decision systems diverge into the rostral cingulate zone (Mueller et al., 2007; Kriehoff et al., 2009) and dFMC (Brass & Haggard, 2007; Kuhn et al., 2009) respectively. Each of these areas is responsible for freely deciding between response alternatives. The medial anterior-posterior organisation for intentional inhibition decision-making may be a critical mechanism necessary for everyday voluntary self-control.

Chapter 3 examined the interactions between intentional and instructed inhibition rather than the differences. It showed that repeated instructed inhibition encourages an intentional decision to act when presented with a free choice. This is consistent with work proposing a limited inhibitory resource (Baumeister et al., 2007; Muraven & Baumeister, 2000), but additionally suggests that instructed inhibition processes inform voluntary inhibition decisions. There is therefore a strong relation between intentional and instructed

inhibition, as the prior history of one biases the current output of the other. It seems plausible that the two decisional systems accumulate and use a common source of evidence. Although chapter 5 highlighted some distinct neural processes underlying each type of inhibition, there are also some important commonalities that should be considered when developing models of inhibition.

7.2. Remaining problems for intentional inhibition

Although this thesis has contributed significantly to what is known about intentional inhibition, several problems still remain. Intentional inhibition is a very broad concept covering many different types of behaviour. Most experiments have looked at how we suppress manual actions. However, there are many other types of real-world behaviour that also involve decisions to inhibit, including speech (Severens et al., 2011a; Xue et al., 2008), emotion (Knyazev, 2007; Kuhn et al., 2011), and thoughts (Dillon & Pizzagalli, 2007; Wenzlaff & Wegner, 2000). It seems unlikely that an identical inhibitory mechanism is used across all of these domains, although some commonalities may exist such as increased alpha oscillations (Klimesch et al., 2007). It also seems unlikely that an entirely independent system is required for each domain, given the large number of possible behavioural targets for inhibition. In fact, inhibitory processes generally have rather broad and global effects in contrast to more selective excitatory processes (Majid et al., 2012; Poo & Isaacson, 2009).

Everyday decisions vary greatly in the amount of information available to inform choices. Often, we make quick decisions based purely on an unambiguous instruction signal in the environment such as a traffic light. At other times however, we can be more unsure about whether to act or inhibit. We could be considering whether we should reveal a secret to a friend, or whether we should take another chocolate from the box. These decisions involve less instructional information in the environment to determine our behaviour. With less external information, our inhibition decisions therefore become more ‘free’ or intentional. Accordingly, the next section proposes a continuum model for human behavioural inhibition. It aims to present a perspective that captures the wide range of possible contexts for intentional decisions.

7.3. *A continuum of instructed to intentional inhibition*

A suitably parsimonious model is needed to account for the wide areas of influence of intentional inhibition. Many areas of psychology and neuroscience have benefited from positing continuums of function rather than categorical constructs. For example, characteristics of psychiatric disorders such as ADHD (Levy et al., 1997), psychosis (van Os et al., 2000, 2009), and many others, are better framed as personality states or behavioural states, which vary along a continuum in the entire population. Diagnosable disorders are simply extremes on that continuum. The concept of human consciousness has also been considered as varying along a continuous scale. This perspective has been helpful in understanding patients with impaired consciousness (Kübler & Kotchoubey, 2007), and has contributed to improving our knowledge of the notoriously complicated ‘consciousness’ concept (Mandler, 2005).

It may be valuable to consider inhibition as varying along a continuum from instructed to intentional. Two key variables may define where inhibition lies on this continuum: 1) the amount of external evidence available in the environment to encourage an inhibition response; 2) the endogeneity/exogeneity of the current task demands. Figure 7.1 below presents these two variables along separate axes and plots the locations of the various types of inhibition from the thesis. The top right corner of the plot is for more instructed inhibition and the bottom left is for more intentional inhibition. Distinct signals that direct an immediate inhibition response, such as those in stop signal and go/nogo paradigms (Verbruggen & Logan, 2008), would be high in external evidence and highly exogenous in the task demands. This makes them highly instructed inhibition responses. Simple ‘free decision’ tasks on the other hand, in which participants make an entirely ‘free’ decision to inhibit action (Brass & Haggard, 2007), would be low in external evidence and highly endogenous in task demands – this places them at the bottom left ‘intentional’ end of figure 7.1.

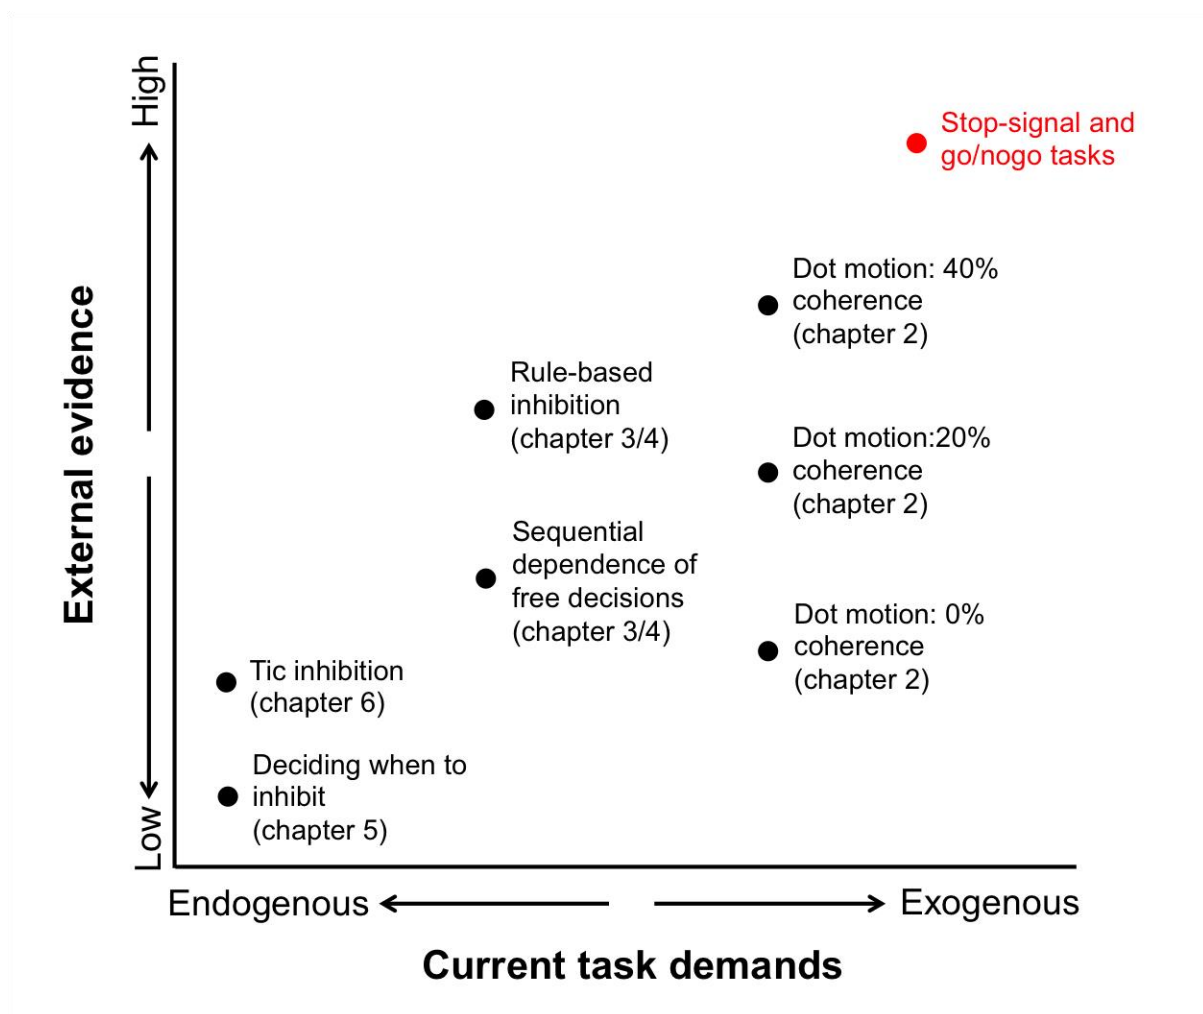


Figure 7.1. Two dimensions of an instructed-intentional continuum for behavioural inhibition: available evidence in the external world to support decisions, and exogeneity/endogeneity of current task demands. The top right of the plot indicates more instructed inhibition, while the bottom left indicates more intentional inhibition.

Inhibitory decision-making in experiments could in principle be placed anywhere along the continuum of instructed to intentional inhibition. Figure 7.1 illustrates the general positive correlation between how exogenous a task tends to be and how much external evidence is available during decision making. With more external evidence available, there is a more ideal environment for exogenous decision-making. On the other hand, with less external evidence, there must be a greater reliance on endogenous decision-making. As the experiments in chapter 2 demonstrated however, there are some cases in which a highly exogenous task may be low in the external evidence that it provides (0% dot motion coherence). Defining whether inhibition is instructed or intentional depends on both

dimensions in figure 7.1. The two dimensions could also be thought of as affordances (y-axis), and processes/tasks on which people are currently engaged (x-axis). Mind-wandering for example is a process characterised by endogenous cognition independent of external stimuli (Mason et al., 2007). This would position it on the extreme left of the x-axis in figure 7.1. Typical ‘free choice’ intentional inhibition experiments may therefore be associated with greater mind-wandering and less task-focused attention. This should be considered when designing intentional inhibition paradigms and when interpreting brain imaging and behavioural data.

In future experimental tasks, careful consideration of where task-related inhibition lies on the instructed-intentional continuum, could provide a basis for unifying data across different inhibition experiments. Furthermore, it could promote a new perspective in the study of inhibition, in which continuous changes in brain activity/behavioural function are identified, rather than simply discrete differences between extremes.

The initial experiments from chapter 2 of this thesis present a particularly good example of inhibition decisions varying along the instructed-intentional scale. Participants’ decisions depended on random dot motion stimuli of varying stimulus strength and so the task had highly exogenous demands. When stimuli had high stimulus strength (i.e. 40% motion coherence), this provided high external evidence for decisions. Therefore, decisions were relatively easy and based almost entirely on the actual direction of the dots, making inhibition decisions largely instructed. 20% motion coherence stimuli presented weaker and less obvious external evidence, making inhibitions a little more difficult. This left more room for internal variability in decision-making and therefore inhibition became more intentional. Similarly, 0% motion coherence presented no consistent external evidence at all and therefore became even more intentional. However, due to the exogenous nature of the task, participants were still presumably looking for external evidence and *attempting* to base their decisions on it, even though no evidence was available. Therefore decisions were unlikely to be *purely* intentional.

The experiments in chapter 3 examined the effects of immediate behavioural history and stimulus history on decision-making. The decision to inhibit in any single trial was ‘free’, because participants were asked to choose to act or inhibit depending on their prediction for an unseen external event. However, during each experiment, participants were also monitoring prior external evidence and accumulating information about runs of repeated

events. This added some degree of exogeneity to the task. Participants clearly used this external information in their decision-making, and therefore were making self-generated decisions that depended on prior external information. As chapter 4 showed, professional goalkeepers also employed this mix of intentional and instructed decision-making. However, in addition to this, goalkeepers demonstrated that self-generated intentional inhibition is difficult to produce in the context of frequent action-related external evidence.

Chapter 5 investigated perhaps the most intentional of inhibition decisions from the thesis. Participants were entirely unconstrained by external factors in when they executed their voluntary inhibition responses. They were simply told to ‘inhibit actions whenever they felt like it’. The only instructed component to such voluntary inhibition was that participants knew they had to inhibit actions *at some point* in each trial. Although inhibition decisions were driven by internal neural activity, there were no predictive external stimuli in the experiment to guide inhibition. The rule-based inhibition decisions in each task on the other hand were driven by an exogenous rule. This made them primarily instructed rather than intentional inhibition. However, participants were required to internalise such rules, and maintain them over extended periods in their continuous behaviour. This may make ‘rule-based’ inhibition more endogenous and more intentional than basic stop-signal inhibition, in which inhibition responses must be generated *immediately* as soon as a clear and unambiguous external signal is detected.

Chapter 6 finally presented the case of tic inhibition with Tourette syndrome patients. This is a rather unusual type of inhibition in which relatively *involuntary* actions must be suppressed – all other experiments throughout the thesis explored inhibition of voluntary actions. Tic inhibition can be considered intentional because it is generally a conscious endogenous strategy used by patients to prevent the occurrence of overt motor tics (Peterson et al., 1998; Verdellen et al., 2007). However, in the context of the chapter 6 experiments, participants were explicitly instructed to do their best to inhibit all tics throughout certain task conditions. Therefore, in this case, tic inhibition could be considered more instructed than intentional. Despite this, the self-generated endogenous nature of everyday tic inhibition, and the self-driven maintenance of tonic inhibitory processes throughout the experimental tasks, adds significant relevance to intentional inhibition.

7.4. Conclusions

The 7 chapters of this thesis have defined, explored, tested, and developed the concept of intentional inhibition. The thesis has demonstrated the wide diversity of behavioural inhibition and attempted to develop a suitable continuum model to represent this diversity. Behavioural inhibition may be more or less ‘free’ or ‘intentional’, depending on the external evidence available in the immediate environment to guide that decision, and the endogeneity of the task demands. The proposition that intentional inhibition belongs on a behavioural continuum is supported by the strong neural overlap it shares with other types of instructed inhibition (Schel et al., 2014). Future work may illuminate a continuum of neural activity patterns that reflect a continuous scale ranging from more instructed to more intentional inhibition.

Converging evidence throughout the thesis from behavioural, neuroimaging, and patient group methods, have shown that intentional inhibition is a cognitive function worth studying. There are four key points that justify the importance of research into intentional inhibition: 1) intentional inhibition is influenced by unique effects of sequential dependence that do not apply to other action or inhibition functions; 2) decision-making parameters for intentional inhibition are distinct from those associated with intentional action; 3) intentional inhibition has distinct neural correlates that are not identical to instructed inhibition; 4) intentional inhibition is a function with direct relevance to real-world behaviour. The behavioural influences and neural mechanisms involved in deciding to inhibit action cannot be identified by work on instructed inhibition or intentional action, yet they are critical to our understanding of everyday self-control. Without a strong account of how we choose to inhibit action in ambiguous conditions, we cannot fully understand a function which is central to human behaviour. A critical faculty of the human condition is the ability to produce flexible behaviour, which is not constrained to immediate responses driven by external stimuli.

Achieving a comprehensive understanding of intentional inhibition is important for many areas of knowledge. It informs us about the mechanisms involved when we stop ourselves acting in inappropriate ways in wider society. Without this ability for self-control, everyday life would likely be more unpredictable, more dangerous, and more violent (Pinker, 2011). The more we understand about inhibition, and the factors involved in failures of inhibition, the more effectively we can address important self-control issues in societal and legal settings. The recent 2014 legal trial of Oscar Pistorius in South Africa highlighted the ‘involuntary action’ defence (Marszal, 2014). Pistorius was accused of shooting and killing

his partner. One of the defences raised was that Pistorius “did not think about pulling the trigger”. Pistorius stated that he pulled the trigger before he could think about what he was doing. If this statement were indeed true, then the remark clearly refers to a failure to consciously decide to exert intentional inhibition. Before we can fairly assess how a failure to intentionally inhibit action is relevant in judgements of legal responsibility, we need to fully understand the inhibition mechanism itself. A stronger understanding may facilitate the development of more effective methods for improving self-control and deterring automatic action in provocative situations.

References

1. Adolphs, R. (2003). Cognitive neuroscience of human social behaviour. *Nature Reviews Neuroscience*, 4(3), 165–178. doi:10.1038/nrn1056
2. Alexander, G. E., & Crutcher, M. D. (1990). Functional architecture of basal ganglia circuits: Neural substrates of parallel processing. *Trends in Neurosciences*, 13, 266–271.
3. Alexander, G. E., DeLong, M. R., & Strick, P. L. (1986). Parallel organization of functionally segregated circuits linking basal ganglia and cortex. *Annual Review of Neuroscience*, 9, 357–381.
4. Allen, A. J., Leonard, H. L., & Swedo, S. E. (1995). Case Study: A New Infection-Triggered, Autoimmune Subtype of Pediatric OCD and Tourette's Syndrome. *Journal of the American Academy of Child & Adolescent Psychiatry*, 34(3), 307–311. doi:10.1097/00004583-199503000-00015
5. American Psychiatric Association. (2000). *Diagnostic and statistical manual of mental disorders* (4th ed., text rev.). Washington, DC: Author.
6. Apesteguia, J., & Palacios-Huerta, I. (2010). Psychological Pressure in Competitive Environments: Evidence from a Randomized Natural Experiment. *The American Economic Review*, 100(5), 2548–2564.
7. Arellano, M., & Honoré, B. (2001). Chapter 53 Panel data models: some recent developments. In J.J. Heckman and E. Leamer (Ed.), *Handbook of Econometrics* (Vol. Volume 5, pp. 3229–3296). Elsevier. Retrieved from <http://www.sciencedirect.com/science/article/pii/S1573441201050061>
8. Aron, A. R., Behrens, T. E., Smith, S., Frank, M. J., & Poldrack, R. A. (2007). Triangulating a Cognitive Control Network Using Diffusion-Weighted Magnetic

Resonance Imaging (MRI) and Functional MRI. *The Journal of Neuroscience*, 27(14), 3743–3752. doi:10.1523/JNEUROSCI.0519-07.2007

9. Aron, A. R., Fletcher, P. C., Bullmore, E. T., Sahakian, B. J., & Robbins, T. W. (2003). Stop-signal inhibition disrupted by damage to right inferior frontal gyrus in humans. *Nature Neuroscience*, 6, 115–116. doi:10.1038/nn1003
10. Aron, A. R., & Poldrack, R. A. (2005). The cognitive neuroscience of response inhibition: relevance for genetic research in attention-deficit/hyperactivity disorder. *Biological Psychiatry*, 57(11), 1285–1292. doi:10.1016/j.biopsych.2004.10.026
11. Aron, A. R., & Poldrack, R. A. (2006). Cortical and Subcortical Contributions to Stop Signal Response Inhibition: Role of the Subthalamic Nucleus. *The Journal of Neuroscience*, 26(9), 2424–2433. doi:10.1523/JNEUROSCI.4682-05.2006
12. Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: one decade on. *Trends in Cognitive Sciences*, 18(4), 177–185. doi:10.1016/j.tics.2013.12.003
13. Awh, E., Barton, B., & Vogel, E. K. (2007). Visual Working Memory Represents a Fixed Number of Items Regardless of Complexity. *Psychological Science*, 18(7), 622–628. doi:10.1111/j.1467-9280.2007.01949.x
14. Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: Two faces of subjective randomness? *Memory & Cognition*, 32(8), 1369–1378. doi:10.3758/BF03206327
15. Azar, O. H., & Bar-Eli, M. (2011). Do soccer players play the mixed-strategy Nash equilibrium? *Applied Economics*, 43(25), 3591–3601. doi:10.1080/00036841003670747

16. Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. D. (2010). Optimally interacting minds. *Science (New York, N.Y.)*, *329*(5995), 1081–1085. doi:10.1126/science.1185718
17. Bannon, S., Gonsalvez, C. J., Croft, R. J., & Boyce, P. M. (2002). Response inhibition deficits in obsessive–compulsive disorder. *Psychiatry Research*, *110*(2), 165–174. doi:10.1016/S0165-1781(02)00104-X
18. Barkley, R. A. (1997). Behavioral inhibition, sustained attention, and executive functions: Constructing a unifying theory of ADHD. *Psychological Bulletin*, *121*(1), 65–94. doi:10.1037/0033-2909.121.1.65
19. Bate, K. S., Malouff, J. M., Thorsteinsson, E. T., & Bhullar, N. (2011). The efficacy of habit reversal therapy for tics, habit disorders, and stuttering: A meta-analytic review. *Clinical Psychology Review*, *31*(5), 865–871. doi:10.1016/j.cpr.2011.03.013
20. Baumeister, R. F., Vohs, K. D., & Tice, D. M. (2007). The Strength Model of Self-Control. *Current Directions in Psychological Science*, *16*(6), 351–355. doi:10.1111/j.1467-8721.2007.00534.x
21. Baym, C. L., Corbett, B. A., Wright, S. B., & Bunge, S. A. (2008). Neural correlates of tic severity and cognitive control in children with Tourette syndrome. *Brain*, *131*(1), 165–179. doi:10.1093/brain/awm278
22. Bays, P. M., & Husain, M. (2008). Dynamic Shifts of Limited Working Memory Resources in Human Vision. *Science*, *321*(5890), 851–854. doi:10.1126/science.1158023
23. BBC News (2004, June 25). Record audience for England match. Retrieved Feb 21, 2014, from <http://news.bbc.co.uk/1/hi/entertainment/3839067.stm>
24. Beauregard, M., Lévesque, J., & Bourgouin, P. (2001). Neural correlates of conscious self-regulation of emotion. *The Journal of Neuroscience*, *21*(18), 6993–7000.

25. Beaver, K. M., Ratchford, M., & Ferguson, C. J. (2009). Evidence of Genetic and Environmental Effects On the Development of Low Self-Control. *Criminal Justice and Behavior*, *36*(11), 1158–1172. doi:10.1177/0093854809342859
26. Beaver, K. M., Schutt, J. E., Boutwell, B. B., Ratchford, M., Roberts, K., & Barnes, J. C. (2009). Genetic and Environmental Influences on Levels of Self-Control and Delinquent Peer Affiliation Results from a Longitudinal Sample of Adolescent Twins. *Criminal Justice and Behavior*, *36*(1), 41–60. doi:10.1177/0093854808326992
27. Beaver, K. M., Wright, J. P., DeLisi, M., & Vaughn, M. G. (2008). Genetic influences on the stability of low self-control: Results from a longitudinal sample of twins. *Journal of Criminal Justice*, *36*(6), 478–485. doi:10.1016/j.jcrimjus.2008.09.006
28. Bednekoff, P. A., & Lima, S. L. (1998). Randomness, chaos and confusion in the study of antipredator vigilance. *Trends in Ecology & Evolution*, *13*(7), 284–287. doi:10.1016/S0169-5347(98)01327-5
29. Bickel, W. K., & Marsch, L. A. (2001). Toward a behavioral economic understanding of drug dependence: delay discounting processes. *Addiction*, *96*(1), 73–86. doi:10.1046/j.1360-0443.2001.961736.x
30. Bickel, W. K., Miller, M. L., Yi, R., Kowal, B. P., Lindquist, D. M., & Pitcock, J. A. (2007). Behavioral and Neuroeconomics of Drug Addiction: Competing Neural Systems and Temporal Discounting Processes. *Drug and Alcohol Dependence*, *90*(Suppl 1), S85–S91. doi:10.1016/j.drugalcdep.2006.09.016
31. Boecker, H., Dagher, A., Ceballos-Baumann, A. O., Passingham, R. E., Samuel, M., Friston, K. J., et al. (1998). Role of the human rostral supplementary motor area and the basal ganglia in motor sequence control: Investigations with H215O PET. *Journal of Neurophysiology*, *79*, 1070–1080.
32. Bohlhalter, S., Goldfine, A., Matteson, S., Garraux, G., Hanakawa, T., Kansaku, K., ... Hallett, M. (2006). Neural correlates of tic generation in Tourette syndrome: an

event-related functional MRI study. *Brain*, 129(8), 2029–2037.

doi:10.1093/brain/awl050

33. Bokura, H., Yamaguchi, S., & Kobayashi, S. (2001). Electrophysiological correlates for response inhibition in a Go/NoGo task. *Clinical Neurophysiology*, 112(12), 2224–2232. doi:10.1016/S1388-2457(01)00691-5
34. Botvinick, M., Nystrom, L. E., Fissell, K., Carter, C. S., & Cohen, J. D. (1999). Conflict monitoring versus selection-for-action in anterior cingulate cortex. *Nature*, 402(6758), 179–181. doi:10.1038/46035
35. Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: an update. *Trends in Cognitive Sciences*, 8(12), 539–546. doi:10.1016/j.tics.2004.10.003
36. Brass, M., & Haggard, P. (2007). To Do or Not to Do: The Neural Signature of Self-Control. *The Journal of Neuroscience*, 27(34), 9141–9145. doi:10.1523/JNEUROSCI.0924-07.2007
37. Brass, M., & Haggard, P. (2008). The What, When, Whether Model of Intentional Action. *The Neuroscientist*, 14(4), 319–325. doi:10.1177/1073858408317417
38. Brass, M., & Haggard, P. (2010). The hidden side of intentional action: the role of the anterior insular cortex. *Brain Structure and Function*, 214, 603–610. doi:10.1007/s00429-010-0269-6
39. Brass, M., & Von Cramon, D. Y. (2002). The role of the frontal cortex in task preparation. *Cerebral Cortex*, 12, 908–914.
40. Braver, T. S., Barch, D. M., Gray, J. R., Molfese, D. L., & Snyder, A. (2001). Anterior cingulate cortex and response conflict: effects of frequency, inhibition and errors. *Cerebral Cortex (New York, N.Y.: 1991)*, 11(9), 825–836.

41. Brembs, B. (2011). Towards a scientific concept of free will as a biological trait: spontaneous actions and decision-making in invertebrates. *Proceedings of the Royal Society B: Biological Sciences*, 278(1707), 930–939. doi:10.1098/rspb.2010.2325
42. Brotchie, P., Iansak, R., & Horne, M. K. (1991). Motor function of the monkey globus pallidus. 2. Cognitive aspects of movement and phasic neuronal activity. *Brain*, 114, 1685–1702.
43. Bullen, C., McRobbie, H., Thornley, S., Glover, M., Lin, R., & Laugesen, M. (2010). Effect of an electronic nicotine delivery device (e cigarette) on desire to smoke and withdrawal, user preferences and nicotine delivery: randomised cross-over trial. *Tobacco Control*, 19(2), 98–103. doi:10.1136/tc.2009.031567
44. Burgess, P. W., Scott, S. K., & Frith, C. D. (2003). The role of the rostral frontal cortex (area 10) in prospective memory: a lateral versus medial dissociation. *Neuropsychologia*, 41(8), 906–918. doi:10.1016/S0028-3932(02)00327-5
45. Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: “Gambler’s fallacy” versus “hot hand“. *Psychonomic Bulletin & Review*, 11(1), 179–184. doi:10.3758/BF03206480
46. Camerer, C. F., Ho, T.-H., & Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games. *The Quarterly Journal of Economics*, 119(3), 861–898. doi:10.1162/0033553041502225
47. Carpenter, R. H. S. (1999). A neural mechanism that randomises behaviour. *Journal of Consciousness Studies*, 6, 13–22.
48. Carter, C. S., Braver, T. S., Barch, D. M., Botvinick, M. M., Noll, D., & Cohen, J. D. (1998). Anterior Cingulate Cortex, Error Detection, and the Online Monitoring of Performance. *Science*, 280(5364), 747–749. doi:10.1126/science.280.5364.747

49. Cavanna, A. E., & Nani, A. (2013). Tourette Syndrome and Consciousness of Action. *Tremor and Other Hyperkinetic Movements*, 3. Retrieved from <http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3782755/>
50. Chambers, C. D., Bellgrove, M. A., Stokes, M. G., Henderson, T. R., Garavan, H., Robertson, I. H., ... Mattingley, J. B. (2006). Executive “brake failure” following deactivation of human frontal lobe. *Journal of Cognitive Neuroscience*, 18(3), 444–455.
51. Chamberlain, S., Fineberg, N., Blackwell, A., Robbins, T., & Sahakian, B. (2006). Motor Inhibition and Cognitive Flexibility in Obsessive-Compulsive Disorder and Trichotillomania. *American Journal of Psychiatry*, 163(7), 1282–1284.
52. Chambon, V., Wenke, D., Fleming, S. M., Prinz, W., & Haggard, P. (2013). An Online Neural Substrate for Sense of Agency. *Cerebral Cortex*, 23(5), 1031–1037. doi:10.1093/cercor/bhs059
53. Chiappori, P. –A., Levitt, S., & Groseclose, T. (2002). Testing Mixed-Strategy Equilibria When Players Are Heterogeneous: The Case of Penalty Kicks in Soccer. *The American Economic Review*, 92(4), 1138-1151.
54. Churchland, M. M., Afshar, A., & Shenoy, K. V. (2006). A central source of movement variability. *Neuron*, 52, 1085–1096.
55. Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693–702. doi:10.1038/nn.2123
56. Cisek, P. (2007). Cortical mechanisms of action selection: The affordance competition hypothesis. *Philosophical Transactions of the Royal Society, Series B, Biological Sciences*, 362, 1585–1599.
57. Cisek, P., & Kalaska, J. F. (2010). Neural mechanisms for interacting with a world full of action choices. *Annual Review of Neuroscience*, 33, 269–298.

58. Clark, D. M. (1986). A cognitive approach to panic. *Behaviour Research and Therapy*, 24(4), 461–470. doi:10.1016/0005-7967(86)90011-2
59. Cohen, J. D., McClure, S. M., & Yu, A. J. (2007). Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 933–942. doi:10.1098/rstb.2007.2098
60. Conelea, C. A., & Woods, D. W. (2008). The influence of contextual factors on tic expression in Tourette's syndrome: A review. *Journal of Psychosomatic Research*, 65(5), 487–496. doi:10.1016/j.jpsychores.2008.04.010
61. Congdon, E., Lesch, K. P., & Canli, T. (2008). Analysis of DRD4 and DAT polymorphisms and behavioral inhibition in healthy adults: implications for impulsivity. *American Journal of Medical Genetics. Part B, Neuropsychiatric Genetics: The Official Publication of the International Society of Psychiatric Genetics*, 147B(1), 27–32. doi:10.1002/ajmg.b.30557
62. Connolly, J. D., Goodale, M. A., Goltz, H. C., & Munoz, D. P. (2005). fMRI activation in the human frontal eye field is correlated with saccadic reaction time. *Journal of Neurophysiology*, 94, 605–611.
63. Craig, A. D. B. (2009). How do you feel--now? The anterior insula and human awareness. *Nature Reviews. Neuroscience*, 10(1), 59–70. doi:10.1038/nrn2555
64. Cressman, E. K., Lam, M. Y., Franks, I. M., Enns, J. T., & Chua, R. (2013). Unconscious and out of control: subliminal priming is insensitive to observer expectations. *Consciousness and Cognition*, 22(3), 716–728. doi:10.1016/j.concog.2013.04.011
65. Critchley, H. D., Wiens, S., Rotshtein, P., Ohman, A., & Dolan, R. J. (2004). Neural systems supporting interoceptive awareness. *Nature Neuroscience*, 7(2), 189–195. doi:10.1038/nn1176

66. Cunnington, R., Windischberger, C., Deecke, L., & Moser, E. (2002). The preparation and execution of self-initiated and externally-triggered movement: a study of event-related fMRI. *NeuroImage*, *15*(2), 373–385. doi:10.1006/nimg.2001.0976
67. David, N., Newen, A., & Vogeley, K. (2008). The “sense of agency” and its underlying cognitive and neural mechanisms. *Consciousness and Cognition*, *17*(2), 523–534. doi:10.1016/j.concog.2008.03.004
68. Davis, C., Katzman, D., & Kirsh, C. (1999). Compulsive physical activity in adolescents with anorexia nervosa: A psychobehavioral spiral of pathology. *The Journal of Nervous and Mental Disease*, *187*(6), 336-342.
69. Deckersbach, T., Rauch, S., Buhlmann, U., & Wilhelm, S. (2006). Habit reversal versus supportive psychotherapy in Tourette’s disorder: A randomized controlled trial and predictors of treatment response. *Behaviour Research and Therapy*, *44*(8), 1079–1090. doi:10.1016/j.brat.2005.08.007
70. Dehaene, S., Artiges, E., Naccache, L., Martelli, C., Viard, A., Schürhoff, F., et al. (2003). Conscious and subliminal conflicts in normal subjects and patients with schizophrenia: The role of the anterior cingulate. *Proceedings of the National Academy of Sciences*, *100*, 13722–13727.
71. Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition*, *79*, 1–37.
72. Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*(1), 9–21. doi:10.1016/j.jneumeth.2003.10.009
73. Delorme, A., Sejnowski, T., & Makeig, S. (2007). Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *Neuroimage*, *34*, 1443–1449.

74. Dickstein, S. G., Bannon, K., Xavier Castellanos, F., & Milham, M. P. (2006). The neural correlates of attention deficit hyperactivity disorder: an ALE meta-analysis. *Journal of Child Psychology and Psychiatry*, *47*(10), 1051–1062. doi:10.1111/j.1469-7610.2006.01671.x
75. Dillon, D. G., & Pizzagalli, D. A. (2007). Inhibition of action, thought, and emotion: A selective neurobiological review. *Applied and Preventive Psychology*, *12*(3), 99–114. doi:10.1016/j.appsy.2007.09.004
76. Dimoska, A., Johnstone, S. J., Barry, R. J., & Clarke, A. R. (2003). Inhibitory motor control in children with attention-deficit/hyperactivity disorder: event-related potentials in the stop-signal paradigm. *Biological Psychiatry*, *54*(12), 1345–1354. doi:10.1016/S0006-3223(03)00703-0
77. Dror, I. E., Basola, B., & Busemeyer, J. R. (1999). Decision making under time pressure: An independent test of sequential sampling models. *Memory & Cognition*, *27*(4), 713–725. doi:10.3758/BF03211564
78. Duann, J.-R., Ide, J. S., Luo, X., & Li, C. R. (2009). Functional Connectivity Delineates Distinct Roles of the Inferior Frontal Cortex and Presupplementary Motor Area in Stop Signal Inhibition. *The Journal of Neuroscience*, *29*(32), 10171–10179. doi:10.1523/JNEUROSCI.1300-09.2009
79. Duckworth, A. L., & Seligman, M. E. P. (2005). Self-Discipline Outdoes IQ in Predicting Academic Performance of Adolescents. *Psychological Science*, *16*(12), 939–944. doi:10.1111/j.1467-9280.2005.01641.x
80. Eagle, D. M., Baunez, C., Hutcherson, D. M., Lehmann, O., Shah, A. P., & Robbins, T. W. (2008). Stop-Signal Reaction-Time Task Performance: Role of Prefrontal Cortex and Subthalamic Nucleus. *Cerebral Cortex*, *18*(1), 178–188. doi:10.1093/cercor/bhm044
81. Eccles, J. C. (1985). Mental summation: The timing of voluntary intentions by cortical activity. *Behavioral and Brain Sciences*, *8*, 542–543.

82. Edwards, M. J., Adams, R. A., Brown, H., Pareés, I., & Friston, K. J. (2012). A Bayesian account of “hysteria.” *Brain*, *135*(11), 3495–3512.
doi:10.1093/brain/aws129
83. Edwards, M. J., & Bhatia, K. P. (2012). Functional (psychogenic) movement disorders: merging mind and brain. *The Lancet Neurology*, *11*(3), 250–260.
doi:10.1016/S1474-4422(11)70310-6
84. Efron, B., & Tibshirani, R. J. (1994). *An Introduction to the Bootstrap*. CRC Press.
85. Eimer, M. (1993). Effects of attention and stimulus probability on ERPs in a Go/Nogo task. *Biological Psychology*, *35*(2), 123–138. doi:10.1016/0301-0511(93)90009-W
86. Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: Behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(6), 1737–1747. doi:10.1037/0096-1523.24.6.1737
87. Eisner, M. (2001). Modernization, Self-Control and Lethal Violence. The Long-term Dynamics of European Homicide Rates in Theoretical Perspective. *British Journal of Criminology*, *41*(4), 618–638. doi:10.1093/bjc/41.4.618
88. Everitt, B. J., & Robbins, T. W. (2005). Neural systems of reinforcement for drug addiction: from actions to habits to compulsion. *Nature Neuroscience*, *8*(11), 1481–1489. doi:10.1038/nn1579
89. Fagerström, K. O. (2005). Can reduced smoking be a way for smokers not interested in quitting to actually quit? *Respiration; International Review of Thoracic Diseases*, *72*(2), 216–220. doi:10.1159/000084057

90. Falba, T., Jofre-Bonet, M., Busch, S., Duchovny, N., & Sindelar, J. (2004). Reduction of quantity smoked predicts future cessation among older smokers. *Addiction*, *99*(1), 93–102. doi:10.1111/j.1360-0443.2004.00574.x
91. Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in Go/Nogo tasks and their relation to inhibition. *Acta Psychologica*, *101*(2-3), 267–291. doi:10.1016/S0001-6918(99)00008-6
92. Faraone, S. V., Perlis, R. H., Doyle, A. E., Smoller, J. W., Goralnick, J. J., Holmgren, M. A., & Sklar, P. (2005). Molecular Genetics of Attention-Deficit/Hyperactivity Disorder. *Biological Psychiatry*, *57*(11), 1313–1323. doi:10.1016/j.biopsych.2004.11.024
93. Farrer, C., Franck, N., Georgieff, N., Frith, C. D., Decety, J., & Jeannerod, M. (2003). Modulating the experience of agency: a positron emission tomography study. *NeuroImage*, *18*(2), 324–333.
94. Farrer, C., & Frith, C. D. (2002). Experiencing Oneself vs Another Person as Being the Cause of an Action: The Neural Correlates of the Experience of Agency. *NeuroImage*, *15*(3), 596–603. doi:10.1006/nimg.2001.1009
95. Favaro, A., & Santonastaso, P. (1998). Impulsive and compulsive self-injurious behavior in bulimia nervosa: Prevalence and psychological correlates. *The Journal of Nervous and Mental Disease*, *186*(3), 157-165.
96. Feinberg, T. E., Schindler, R. J., Flanagan, N. G., & Haber, L. D. (1992). Two alien hand syndromes. *Neurology*, *42*(1), 19–19. doi:10.1212/WNL.42.1.19
97. Ferrão, Y. A., Miguel, E., & Stein, D. J. (2009). Tourette’s syndrome, trichotillomania, and obsessive–compulsive disorder: How closely are they related? *Psychiatry Research*, *170*(1), 32–42. doi:10.1016/j.psychres.2008.06.008

98. Filevich, E., Kühn, S., & Haggard, P. (2012). Intentional inhibition in human action: The power of “no.” *Neuroscience & Biobehavioral Reviews*, 36(4), 1107–1118. doi:10.1016/j.neubiorev.2012.01.006
99. Filevich, E., Kühn, S., & Haggard, P. (2013). There Is No Free Won’t: Antecedent Brain Activity Predicts Decisions to Inhibit. *PLoS ONE*, 8(2), e53053. doi:10.1371/journal.pone.0053053
100. Fillmore, M. T., & Rush, C. R. (2002). Impaired inhibitory control of behavior in chronic cocaine users. *Drug and Alcohol Dependence*, 66(3), 265–273.
101. Fincham, F. D., & Jaspars, J. M. (1980). Attribution of responsibility: From man the scientist to man as lawyer. In L. Berkowitz (Ed.), *Advances in Experimental Social Psychology*, (Vol. 13, pp 81-138). New York: Academic Press.
102. Fleming, S. M., Mars, R. B., Gladwin, T. E., & Haggard, P. (2009). When the Brain Changes Its Mind: Flexibility of Action Selection in Instructed and Free Choices, *19*(10), 2352–2360. doi:10.1093/cercor/bhn252
103. Fox, M. D., & Raichle, M. E. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nat Rev Neurosci*, 8(9), 700–711. doi:10.1038/nrn2201
104. Fox, M. D., Snyder, A. Z., Vincent, J. L., & Raichle, M. E. (2007). Intrinsic fluctuations within cortical systems account for intertrial variability in human behavior. *Neuron*, 56, 171–184.
105. Franzkowiak, S., Pollok, B., Biermann-Ruben, K., Südmeyer, M., Paszek, J., Jonas, M., ... Schnitzler, A. (2010). Altered pattern of motor cortical activation-inhibition during voluntary movements in Tourette syndrome. *Movement Disorders: Official Journal of the Movement Disorder Society*, 25(12), 1960–1966. doi:10.1002/mds.23186

106. Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature*, 40(2), 351–401. doi:10.1257/002205102320161311
107. Freeman, R. D., Fast, D. K., Burd, L., Kerbeshian, J., Robertson, M. M., & Sandor, P. (2000). An international perspective on Tourette syndrome: selected findings from 3500 individuals in 22 countries. *Developmental Medicine & Child Neurology*, 42(7), 436–447. doi:10.1111/j.1469-8749.2000.tb00346.x
108. Frith, C. D., & Done, D. J. (1986). Routes to action in reaction time tasks. *Psychological Research*, 48(3), 169–177. doi:10.1007/BF00309165
109. Frith, C. D., Friston, K., Liddle, P. F., & Frackowiak, R. S. J. (1991). Willed Action and the Prefrontal Cortex in Man: A Study with PET. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 244(1311), 241–246. doi:10.1098/rspb.1991.0077
110. Gallagher, H. L., & Frith, C. D. (2003). Functional imaging of “theory of mind.” *Trends in Cognitive Sciences*, 7(2), 77–83. doi:10.1016/S1364-6613(02)00025-6
111. Ganos, C., Kahl, U., Schunke, O., Kühn, S., Haggard, P., Gerloff, C., ... Münchau, A. (2012). Are premonitory urges a prerequisite of tic inhibition in Gilles de la Tourette syndrome? *Journal of Neurology, Neurosurgery & Psychiatry*, 83(10), 975–978. doi:10.1136/jnnp-2012-303033
112. Ganos, C., Kühn, S., Kahl, U., Schunke, O., Feldheim, J., Gerloff, C., et al. (2014). Action inhibition in Tourette syndrome. *Movement Disorders*. doi:10.1002/mds.25944
113. Ganos, C., Roessner, V., & Münchau, A. (2013). The functional anatomy of Gilles de la Tourette syndrome. *Neuroscience & Biobehavioral Reviews*, 37(6), 1050–1062. doi:10.1016/j.neubiorev.2012.11.004

114. Garavan, H., Ross, T. J., & Stein, E. A. (1999). Right hemispheric dominance of inhibitory control: An event-related functional MRI study. *Proceedings of the National Academy of Sciences*, *96*(14), 8301–8306. doi:10.1073/pnas.96.14.8301
115. Gerardin, E., Sirigu, A., Lehericy, S., Poline, J.-B., Gaymard, B., Marsault, C., et al. (2000). Partially overlapping neural networks for real and imagined hand movements. *Cerebral Cortex*, *10*, 1093–1104.
116. Gizer, I. R., Ficks, C., & Waldman, I. D. (2009). Candidate gene studies of ADHD: a meta-analytic review. *Human Genetics*, *126*(1), 51–90. doi:10.1007/s00439-009-0694-x
117. Gold, J. I., & Shadlen, M. N. (2007). The Neural Basis of Decision Making. *Annual Review of Neuroscience*, *30*(1), 535–574. doi:10.1146/annurev.neuro.29.051605.113038
118. Goodman, W. K., Price, L. H., Rasmussen, S. A., Mazure, C., Fleischmann, R. L., Hill, C. L., ... Charney, D. S. (1989). The Yale-Brown Obsessive Compulsive Scale. I. Development, use, and reliability. *Archives of General Psychiatry*, *46*(11), 1006–1011.
119. Grahn, J. A., & Brett, M. (2007). Rhythm and beat perception in motor areas of the brain. *Journal of Cognitive Neuroscience*, *19*(5), 893–906. doi:10.1162/jocn.2007.19.5.893
120. Graybiel, A. M., & Rauch, S. L. (2000). Toward a neurobiology of obsessive-compulsive disorder. *Neuron*, *28*(2), 343–347. doi:10.1016/S0896-6273(00)00113-6
121. Green, L., & Myerson, J. (2004). A Discounting Framework for Choice With Delayed and Probabilistic Rewards. *Psychological Bulletin*, *130*(5), 769–792. doi:10.1037/0033-2909.130.5.769

122. Green, D. M., & Swets, J. A. (1974). *Signal detection theory and psychophysics* (Vol. xiii). Oxford, England: Robert E. Krieger.
123. Grèzes, J., & Decety, J. (2002). Does visual perception of object afford action? Evidence from a neuroimaging study. *Neuropsychologia*, *40*(2), 212–222. doi:10.1016/S0028-3932(01)00089-6
124. Grèzes, J., Tucker, M., Armony, J., Ellis, R., & Passingham, R. E. (2003). Objects automatically potentiate action: an fMRI study of implicit processing. *European Journal of Neuroscience*, *17*(12), 2735–2740. doi:10.1046/j.1460-9568.2003.02695.x
125. Haggard, P. (2008). Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, *9*(12), 934–946. doi:10.1038/nrn2497
126. Haggard, P., Clark, S., & Kalogeras, J. (2002). Voluntary action and conscious awareness. *Nat Neurosci*, *5*(4), 382–385. doi:10.1038/nn827
127. Haggard, P., & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, *126*, 128–133. doi:10.1007/s002210050722
128. Haggard, P., Poonian, S., & Walsh, E. (2009). Representing the consequences of intentionally inhibited actions. *Brain Research*, *1286*, 106–113. doi:10.1016/j.brainres.2009.06.020
129. Haggard, P., & Tsakiris, M. (2009). The Experience of Agency Feelings, Judgments, and Responsibility. *Current Directions in Psychological Science*, *18*(4), 242–246. doi:10.1111/j.1467-8721.2009.01644.x
130. Hamilton, A. F. de C., Jones, K. E., & Wolpert, D. M. (2004). The scaling of motor noise with muscle strength and motor unit number in humans. *Experimental Brain Research*, *157*(4), 417–430. doi:10.1007/s00221-004-1856-7

131. Hanks, T. D., Ditterich, J., & Shadlen, M. N. (2006). Microstimulation of macaque area LIP affects decision-making in a motion discrimination task. *Nature Neuroscience*, *9*(5), 682–689. doi:10.1038/nn1683
132. Harris, C. M., & Wolpert, D. M. (1998). Signal-dependent noise determines motor planning. *Nature*, *394*(6695), 780–784. doi:10.1038/29528
133. Heise, K.-F., Steven, B., Liuzzi, G., Thomalla, G., Jonas, M., Müller-Vahl, K., ... Hummel, F. C. (2010). Altered modulation of intracortical excitability during movement preparation in Gilles de la Tourette syndrome. *Brain*, *133*(2), 580–590. doi:10.1093/brain/awp299
134. Hesselmann, G., Kell, C. A., Eger, E., & Kleinschmidt, A. (2008a). Spontaneous local variations in ongoing neural activity bias perceptual decisions. *Proceedings of the National Academy of Sciences*, *105*(31), 10984–10989. doi:10.1073/pnas.0712043105
135. Himle, M. B., Woods, D. W., Conelea, C. A., Bauer, C. C., & Rice, K. A. (2007). Investigating the effects of tic suppression on premonitory urge ratings in children and adolescents with Tourette's syndrome. *Behaviour Research and Therapy*, *45*(12), 2964–2976. doi:10.1016/j.brat.2007.08.007
136. Himle, M. B., Woods, D. W., Piacentini, J. C., & Walkup, J. T. (2006). Brief review of habit reversal training for Tourette syndrome. *Journal of Child Neurology*, *21*(8), 719–725.
137. Horn, N. R., Dolan, M., Elliott, R., Deakin, J. F. W., & Woodruff, P. W. R. (2003). Response inhibition and impulsivity: an fMRI study. *Neuropsychologia*, *41*(14), 1959–1966. doi:10.1016/S0028-3932(03)00077-0
138. Hsieh, J. C., Hagermark, O., Stahle-Backdahl, M., Ericson, K., Eriksson, L., Stone-Elander, S., & Ingvar, M. (1994). Urge to scratch represented in the human cerebral cortex during itch. *Journal of Neurophysiology*, *72*(6), 3004–3008.

139. Hsu, T.-Y., Tseng, L.-Y., Yu, J.-X., Kuo, W.-J., Hung, D. L., Tzeng, O. J. L., ... Juan, C.-H. (2011). Modulating inhibitory control with direct current stimulation of the superior medial frontal cortex. *NeuroImage*, *56*(4), 2249–2257.
doi:10.1016/j.neuroimage.2011.03.059
140. Hughes, G., Velmans, M., & De Fockert, J. (2009). Unconscious priming of a no-go response. *Psychophysiology*, *46*, 1258–1269.
141. Huk, A. C., & Shadlen, M. N. (2005). Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Motion Signals during Perceptual Decision Making. *The Journal of Neuroscience*, *25*(45), 10420–10436.
doi:10.1523/JNEUROSCI.4684-04.2005
142. Humphreys, G. R., & Buehner, M. J. (2009). Magnitude estimation reveals temporal binding at super-second intervals. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1542–1549. doi:10.1037/a0014492
143. Ivry, R. B., & Keele, S. W. (1989). Timing functions of the cerebellum. *Journal of Cognitive Neuroscience*, *1*(2), 136–152. doi:10.1162/jocn.1989.1.2.136
144. Jack, A. I., & Shallice, T. (2001). Introspective physicalism as an approach to the science of consciousness. *Cognition*, *79*, 161–196.
145. Jackson, S. R., Parkinson, A., Kim, S. Y., Schüermann, M., & Eickhoff, S. B. (2011). On the functional anatomy of the urge-for-action. *Cognitive Neuroscience*, *2*(3-4), 227–243. doi:10.1080/17588928.2011.604717
146. Jacobson, L., Javitt, D. C., & Lavidor, M. (2011). Activation of Inhibition: Diminishing Impulsive Behavior by Direct Current Stimulation over the Inferior Frontal Gyrus. *Journal of Cognitive Neuroscience*, *23*(11), 3380–3387.
doi:10.1162/jocn_a_00020

147. Jaffard, M., Longcamp, M., Velay, J.-L., Anton, J.-L., Roth, M., Nazarian, B., & Boulinguez, P. (2008). Proactive inhibitory control of movement assessed by event-related fMRI. *NeuroImage*, *42*(3), 1196–1206. doi:10.1016/j.neuroimage.2008.05.041
148. Jahanshahi, M., Jenkins, I. H., Brown, R. G., Marsden, C. D., Passingham, R. E., & Brooks, D. J. (1995). Self-initiated versus externally triggered movements I. An investigation using measurement of regional cerebral blood flow with PET and movement-related potentials in normal and Parkinson's disease subjects. *Brain*, *118*(4), 913–933. doi:10.1093/brain/118.4.913
149. Jenkins, I. H., Jahanshahi, M., Jueptner, M., Passingham, R. E., & Brooks, D. J. (2000). Self-initiated versus externally triggered movements II. The effect of movement predictability on regional cerebral blood flow. *Brain*, *123*(6), 1216–1228. doi:10.1093/brain/123.6.1216
150. Johnson-Laird, P. N. (1994). Mental models and probabilistic thinking. *Cognition*, *50*(1–3), 189–209. doi:10.1016/0010-0277(94)90028-0
151. Kaladjian, A., Jeanningros, R., Azorin, J.-M., Grimault, S., Anton, J.-L., & Mazzola-Pomietto, P. (2007). Blunted activation in right ventrolateral prefrontal cortex during motor response inhibition in schizophrenia. *Schizophrenia Research*, *97*(1-3), 184–193. doi:10.1016/j.schres.2007.07.033
152. Karch, S., Mulert, C., Thalmeier, T., Lutz, J., Leicht, G., Meindl, T., ... Pogarell, O. (2009). The free choice whether or not to respond after stimulus presentation. *Human Brain Mapping*, *30*(9), 2971–2985. doi:10.1002/hbm.20722
153. Kastner, S., Weerd, P. D., Desimone, R., & Ungerleider, L. G. (1998). Mechanisms of Directed Attention in the Human Extrastriate Cortex as Revealed by Functional MRI. *Science*, *282*(5386), 108–111. doi:10.1126/science.282.5386.108
154. Kaufman, J. N., Ross, T. J., Stein, E. A., & Garavan, H. (2003). Cingulate Hypoactivity in Cocaine Users During a GO-NOGO Task as Revealed by Event-

Related Functional Magnetic Resonance Imaging. *The Journal of Neuroscience*, 23(21), 7839–7843.

155. Kirby, K. N., & Herrnstein, R. J. (1995). Preference Reversals Due to Myopic Discounting of Delayed Reward. *Psychological Science*, 6(2), 83–89.
doi:10.1111/j.1467-9280.1995.tb00311.x
156. Kirby, K. N., & Petry, N. M. (2004). Heroin and cocaine abusers have higher discount rates for delayed rewards than alcoholics or non-drug-using controls. *Addiction*, 99(4), 461–471. doi:10.1111/j.1360-0443.2003.00669.x
157. Klaes, C., Westendorff, S., Chakrabarti, S., & Gail, A. (2011). Choosing goals, not rules: Deciding among rule-based action plans. *Neuron*, 70, 536–548.
158. Klimesch, W., Sauseng, P., & Hanslmayr, S. (2007). EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, 53(1), 63–88.
doi:10.1016/j.brainresrev.2006.06.003
159. Knyazev, G. G. (2007). Motivation, emotion, and their inhibitory control mirrored in brain oscillations. *Neuroscience & Biobehavioral Reviews*, 31(3), 377–395. doi:10.1016/j.neubiorev.2006.10.004
160. Koechlin, E., Ody, C., & Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science*, 302, 1181–1185.
161. Koechlin, E., & Summerfield, C. (2007). An information theoretical approach to prefrontal executive function. *Trends in Cognitive Sciences*, 11(6), 229–235.
doi:10.1016/j.tics.2007.04.005
162. Kok, A., Ramautar, J. R., De Ruiter, M. B., Band, G. P. H., & Ridderinkhof, K. R. (2004). ERP components associated with successful and unsuccessful stopping in a stop-signal task. *Psychophysiology*, 41(1), 9–20. doi:10.1046/j.1469-8986.2003.00127.x

163. Kornhuber, H. H., & Deecke, L. (1965). [CHANGES IN THE BRAIN POTENTIAL IN VOLUNTARY MOVEMENTS AND PASSIVE MOVEMENTS IN MAN: READINESS POTENTIAL AND REAFFERENT POTENTIALS]. *Pflügers Archiv Für Die Gesamte Physiologie Des Menschen Und Der Tiere*, 284, 1–17.
164. Kounieher, F., Charron, S., & Koechlin, E. (2009). Motivation and cognitive control in the human prefrontal cortex. *Nature Neuroscience*, 12, 939–945
165. Kovash, K., & Levitt, S. (2009). *Professionals Do Not Play Minimax: Evidence from Major League Baseball and the National Football League* (NBER Working Paper No. 15347). National Bureau of Economic Research, Inc. Retrieved from <http://econpapers.repec.org/paper/nbrnberwo/15347.htm>
166. Krieghoff, V., Brass, M., Prinz, W., & Waszak, F. (2009). Dissociating what and when of intentional actions. *Frontiers in Human Neuroscience*, 3, 3. doi:10.3389/neuro.09.003.2009
167. Kritikos, A., Breen, N., & Mattingley, J. B. (2005). Anarchic hand syndrome: Bimanual coordination and sensitivity to irrelevant information in unimanual reaches. *Cognitive Brain Research*, 24(3), 634–647. doi:10.1016/j.cogbrainres.2005.03.015
168. Kübler, A., & Kotchoubey, B. (2007). Brain–computer interfaces in the continuum of consciousness: *Current Opinion in Neurology*, 20(6), 643–649. doi:10.1097/WCO.0b013e3282f14782
169. Kuhnen, C. M., & Knutson, B. (2005). The Neural Basis of Financial Risk Taking. *Neuron*, 47(5), 763–770. doi:10.1016/j.neuron.2005.08.008
170. Kühn, S., Gallinat, J., & Brass, M. (2011). “Keep Calm and Carry On”: Structural Correlates of Expressive Suppression of Emotions. *PLoS ONE*, 6(1), e16569. doi:10.1371/journal.pone.0016569

171. Kühn, S., Haggard, P., & Brass, M. (2009). Intentional inhibition: How the “veto-area” exerts control. *Human Brain Mapping, 30*(9), 2834–2843. doi:10.1002/hbm.20711
172. Lacourse, M. G., Orr, E. L. R., Cramer, S. C., & Cohen, M. J. (2005). Brain activation during execution and motor imagery of novel and skilled sequential hand movements. *Neuroimage, 27*, 505–519.
173. Land, M. F., & McLeod, P. (2000). From eye movements to actions: how batsmen hit the ball. *Nature Neuroscience, 3*(12), 1340–1345. doi:10.1038/81887
174. Lau, H. C. (2007). A higher order Bayesian decision theory of consciousness. In Rahul Banerjee and Bikas K. Chakrabarti (Ed.), *Progress in Brain Research* (Vol. Volume 168, pp. 35–48). Elsevier. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0079612307680042>
175. Lau, H. C., & Passingham, R. E. (2006). Relative blindsight in normal observers and the neural correlate of visual consciousness. *Proceedings of the National Academy of Sciences, 103*(49), 18763–18768. doi:10.1073/pnas.0607716103
176. Lau, H. ., Rogers, R. ., Ramnani, N., & Passingham, R. (2004). Willed action and attention to the selection of action. *NeuroImage, 21*(4), 1407–1415. doi:10.1016/j.neuroimage.2003.10.034
177. Leckman, J. F., Bloch, M. H., Scahill, L., & King, R. A. (2006). Tourette Syndrome: The Self Under Siege. *Journal of Child Neurology, 21*(8), 642–649. doi:10.1177/08830738060210081001
178. Leckman, J. F., Riddle, M. A., Hardin, M. T., Ort, S. I., Swartz, K. L., Stevenson, J., & Cohen, D. J. (1989). The Yale Global Tic Severity Scale: initial testing of a clinician-rated scale of tic severity. *Journal of the American Academy of Child and Adolescent Psychiatry, 28*(4), 566–573. doi:10.1097/00004583-198907000-00015

179. Lee, D., Conroy, M. L., McGreevy, B. P., & Barraclough, D. J. (2004). Reinforcement learning and decision making in monkeys during a competitive game. *Cognitive Brain Research*, 22(1), 45–58. doi:10.1016/j.cogbrainres.2004.07.007
180. Lepage, J.-F., & Théoret, H. (2006). EEG evidence for the presence of an action observation–execution matching system in children. *European Journal of Neuroscience*, 23(9), 2505–2510. doi:10.1111/j.1460-9568.2006.04769.x
181. Levy, F., Hay, D. A., McStephen, M., Wood, C., & Waldman, I. (1997). Attention-deficit hyperactivity disorder: a category or a continuum? Genetic analysis of a large-scale twin study. *Journal of the American Academy of Child and Adolescent Psychiatry*, 36(6), 737–744. doi:10.1097/00004583-199706000-00009
182. Lévesque, J., Eugène, F., Joannette, Y., Paquette, V., Mensour, B., Beaudoin, G., ... Beauregard, M. (2003). Neural circuitry underlying voluntary suppression of sadness. *Biological Psychiatry*, 53(6), 502–510. doi:10.1016/S0006-3223(02)01817-6
183. Li, C. R., Huang, C., Constable, R. T., & Sinha, R. (2006). Imaging Response Inhibition in a Stop-Signal Task: Neural Correlates Independent of Signal Monitoring and Post-Response Processing. *The Journal of Neuroscience*, 26(1), 186–192. doi:10.1523/JNEUROSCI.3741-05.2006
184. Li, C. R., Huang, C., Yan, P., Bhagwagar, Z., Milivojevic, V., & Sinha, R. (2008). Neural Correlates of Impulse Control During Stop Signal Inhibition in Cocaine-Dependent Men. *Neuropsychopharmacology*, 33(8), 1798–1806. doi:10.1038/sj.npp.1301568
185. Li, C. R., Luo, X., Yan, P., Bergquist, K., & Sinha, R. (2009). Altered Impulse Control in Alcohol Dependence: Neural Measures of Stop Signal Performance. *Alcoholism: Clinical and Experimental Research*, 33(4), 740–750. doi:10.1111/j.1530-0277.2008.00891.x

186. Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness potential). *Brain*, *106*(3), 623–642. doi:10.1093/brain/106.3.623
187. Livesey, E. J., & Costa, D. S. J. (2014). Automaticity and conscious control in single and choice response time versions of the Perruchet effect. *Quarterly Journal of Experimental Psychology (2006)*, *67*(4), 646–664. doi:10.1080/17470218.2013.824014
188. Lochner, C., Seedat, S., du Toit, P. L., Nel, D. G., Niehaus, D. J., Sandler, R., & Stein, D. J. (2005). Obsessive-compulsive disorder and trichotillomania: a phenomenological comparison. *BMC Psychiatry*, *5*, 2. doi:10.1186/1471-244X-5-2
189. Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.), *Inhibitory processes in attention, memory, and language* (pp. 189–239). San Diego, CA, US: Academic Press.
190. Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295–327. doi:10.1037/0033-295X.91.3.295
191. Lotze, M., Montoya, P., Erb, M., Hülsmann, E., Flor, H., Klose, U., et al. (1999). Activation of cortical and cerebellar motor areas during executed and imagined hand movements: An fMRI study. *Journal of Cognitive Neuroscience*, *11*, 491–501.
192. Low, K. A., & Miller, J. (1999). The usefulness of partial information: Effects of go probability in the choice/nogo task. *Psychophysiology*, *36*, 288–297.
193. Ludolph, A. G., Roessner, V., Munchau, A., & Muller-Vahl, K. (2012). Tourette Syndrome and Other Tic Disorders in Childhood, Adolescence and Adulthood. *Deutsches Ärzteblatt International*, *109*(48), 821–288. doi:10.3238/arztebl.2012.0821

194. Macrae, C. N., Moran, J. M., Heatherton, T. F., Banfield, J. F., & Kelley, W. M. (2004). Medial Prefrontal Activity Predicts Memory for Self. *Cerebral Cortex*, *14*(6), 647–654. doi:10.1093/cercor/bhh025
195. Majid, D. S. A., Cai, W., George, J. S., Verbruggen, F., & Aron, A. R. (2012). Transcranial magnetic stimulation reveals dissociable mechanisms for global versus selective corticomotor suppression underlying the stopping of action. *Cerebral Cortex*, *22*(2), 363–371. doi:10.1093/cercor/bhr112
196. Makris, S., Hadar, A. A., & Yarrow, K. (2011). Viewing objects and planning actions: On the potentiation of grasping behaviours by visual objects. *Brain and Cognition*, *77*(2), 257–264. doi:10.1016/j.bandc.2011.08.002
197. Maltby, N., Tolin, D. F., Worhunsky, P., O’Keefe, T. M., & Kiehl, K. A. (2005). Dysfunctional action monitoring hyperactivates frontal–striatal circuits in obsessive–compulsive disorder: an event-related fMRI study. *NeuroImage*, *24*(2), 495–503. doi:10.1016/j.neuroimage.2004.08.041
198. Mandler, G. (2005). The consciousness continuum: from “qualia” to “free will.” *Psychological Research*, *69*(5-6), 330–337. doi:10.1007/s00426-004-0206-5
199. Marchetti, C., & Della Sala, S. (1998). Disentangling the alien and anarchic hand. *Cognitive Neuropsychiatry*, *3*(3), 191–207. doi:10.1080/135468098396143
200. Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, *9*(6), 296–305. doi:10.1016/j.tics.2005.04.010
201. Marsden, C. D. (1984). Which motor disorder in Parkinson’s disease indicates the true motor function of the basal ganglia? *Ciba Foundation Symposium*, *107*, 225–241.

202. Marszal, A. (2014, April 15). Oscar Pistorius murder trial: April 15 as it happened. Retrieved from <http://www.telegraph.co.uk/news/worldnews/oscar-pistorius/10769304/Oscar-Pistorius-murder-trial-April-15-as-it-happened.html>
203. Mason, M. F., Norton, M. I., Van Horn, J. D., Wegner, D. M., Grafton, S. T., & Macrae, C. N. (2007). Wandering Minds: The Default Network and Stimulus-Independent Thought. *Science*, *315*(5810), 393–395. doi:10.1126/science.1131295
204. Matsushashi, M., & Hallett, M. (2008). The timing of the conscious intention to move. *European Journal of Neuroscience*, *28*, 2344–2351.
205. Mazaheri, A., Nieuwenhuis, I. L. C., van Dijk, H., & Jensen, O. (2009). Prestimulus alpha and mu activity predicts failure to inhibit motor responses. *Human Brain Mapping*, *30*, 1791–1800.
206. McClure, S. M., Ericson, K. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2007). Time Discounting for Primary Rewards. *The Journal of Neuroscience*, *27*(21), 5796–5804. doi:10.1523/JNEUROSCI.4246-06.2007
207. McClure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science*, *306*(5695), 503–507. doi:10.1126/science.1100907
208. McMorris, T. & Hauxwell, B. (1997). In *Science and Football III*, T. Reilly, J. Bangsbo, M. Hughes, Eds. (London: Taylor & Francis Group), pp. 290-294.
209. Mell, L. K., Davis, R. L., & Owens, D. (2005). Association Between Streptococcal Infection and Obsessive-Compulsive Disorder, Tourette's Syndrome, and Tic Disorder. *Pediatrics*, *116*(1), 56–60. doi:10.1542/peds.2004-2058
210. Menon, V., Adleman, N. E., White, C. d., Glover, G. h., & Reiss, A. I. (2001). Error-related brain activation during a Go/NoGo response inhibition task. *Human Brain Mapping*, *12*(3), 131–143. doi:10.1002/1097-0193(200103)12:3<131::AID-HBM1010>3.0.CO;2-C

211. Merikle, P. M., Joordens, S., & Stolz, J. A. (1995). Measuring the relative magnitude of unconscious influences. *Consciousness and Cognition*, 4, 422–439.
212. Metcalfe, J., & Mischel, W. (1999). A hot/cool-system analysis of delay of gratification: Dynamics of willpower. *Psychological Review*, 106(1), 3–19.
doi:10.1037/0033-295X.106.1.3
213. Mischel, W., & Ebbsen, E. B. (1970). Attention in delay of gratification. *Journal of Personality and Social Psychology*, 16(2), 329–337. doi:10.1037/h0029815
214. Mischel, W., Ebbsen, E. B., & Zeiss, A. R. (1972). Cognitive and attentional mechanisms in delay of gratification. *Journal of Personality and Social Psychology*, 21(2), 204–218.
215. Mischel, W., Shoda, Y., & Rodriguez, M. I. (1989). Delay of gratification in children. *Science*, 244(4907), 933–938. doi:10.1126/science.2658056
216. Misirlisoy, E., & Haggard, P. (2014). Veto and Vacillation: A Neural Precursor of the Decision to Withhold Action. *Journal of Cognitive Neuroscience*, 26(2), 296–304. doi:10.1162/jocn_a_00479
217. Monterosso, J. R., Aron, A. R., Cordova, X., Xu, J., & London, E. D. (2005). Deficits in response inhibition associated with chronic methamphetamine abuse. *Drug and Alcohol Dependence*, 79(2), 273–277. doi:10.1016/j.drugalcdep.2005.02.002
218. Moore, J. W., Ruge, D., Wenke, D., Rothwell, J., & Haggard, P. (2010). Disrupting the experience of control in the human brain: pre-supplementary motor area contributes to the sense of agency, *Proc. R. Soc. B*, 277(1693), 2503–2509.
doi:10.1098/rspb.2010.0404
219. Moran, J., & Desimone, R. (1985). Selective attention gates visual processing in the extrastriate cortex. *Science*, 229(4715), 782–784. doi:10.1126/science.4023713

220. Moretto, G., Schwingenschuh, P., Katschnig, P., Bhatia, K. P., & Haggard, P. (2011). Delayed experience of volition in Gilles de la Tourette syndrome. *Journal of Neurology, Neurosurgery & Psychiatry*, 82(12), 1324–1327.
doi:10.1136/jnnp.2010.221143
221. Mostofsky, S. H., Schafer, J. G. B., Abrams, M. T., Goldberg, M. C., Flower, A. A., Boyce, A., et al. (2003). fMRI evidence that the neural basis of response inhibition is task-dependent. *Cognitive Brain Research*, 17, 419–430.
222. Mostofsky, S. H., & Simmonds, D. J. (2008). Response Inhibition and Response Selection: Two Sides of the Same Coin. *Journal of Cognitive Neuroscience*, 20(5), 751–761. doi:10.1162/jocn.2008.20500
223. Mueller, V. A., Brass, M., Waszak, F., & Prinz, W. (2007). The role of the preSMA and the rostral cingulate zone in internally selected actions. *NeuroImage*, 37(4), 1354–1361. doi:10.1016/j.neuroimage.2007.06.018
224. Muraven, M., & Baumeister, R. F. (2000). Self-regulation and depletion of limited resources: Does self-control resemble a muscle? *Psychological Bulletin*, 126(2), 247–259. doi:10.1037/0033-2909.126.2.247
225. Muraven, M., Tice, D. M., & Baumeister, R. F. (1998). Self-control as a limited resource: Regulatory depletion patterns. *Journal of Personality and Social Psychology*, 74(3), 774–789. doi:10.1037/0022-3514.74.3.774
226. Murray, S. O., & Wojciulik, E. (2004). Attention increases neural selectivity in the human lateral occipital complex. *Nature Neuroscience*, 7(1), 70–74.
doi:10.1038/nn1161
227. Muthukumaraswamy, S. D., & Johnson, B. W. (2004). Changes in rolandic mu rhythm during observation of a precision grip. *Psychophysiology*, 41(1), 152–156.
doi:10.1046/j.1469-8986.2003.00129.x

228. Müller-Vahl, K. R., Riemann, L., & Bokemeyer, S. (2014). Tourette patients' misbelief of a tic rebound is due to overall difficulties in reliable tic rating. *Journal of Psychosomatic Research*, *76*(6), 472–476. doi:10.1016/j.jpsychores.2014.03.003
229. Nordby, H., Hammerborg, D., Roth, W. T., & Hugdahl, K. (1994). ERPs for infrequent omissions and inclusions of stimulus elements. *Psychophysiology*, *31*, 544–552.
230. Oosterlaan, J., Logan, G. D., & Sergeant, J. A. (1998). Response Inhibition in AD/HD, CD, Comorbid AD/HD+CD, Anxious, and Control Children: A Meta-analysis of Studies with the Stop Task. *Journal of Child Psychology and Psychiatry*, *39*(3), 411–425. doi:10.1111/1469-7610.00336
231. Osman, A., Albert, R., Ridderinkhof, K. R., Band, G., & van der Molen, M. (2006). The beat goes on: Rhythmic modulation of cortical potentials by imagined tapping. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 986–1005.
232. Palacios-Huerta, I. (2003). Professionals Play Minimax. *The Review of Economic Studies*, *70*(2), 395–415. doi:10.1111/1467-937X.00249
233. Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376–404. doi:10.1167/5.5.1
234. Pappert, E. J., Goetz, C. G., Louis, E. D., Blasucci, L., & Leurgans, S. (2003). Objective assessments of longitudinal outcome in Gilles de la Tourette's syndrome. *Neurology*, *61*(7), 936–940. doi:10.1212/01.WNL.0000086370.10186.7C
235. Parkinson, J., & Haggard, P. (2014). Subliminal priming of intentional inhibition. *Cognition*, *130*(2), 255–265. doi:10.1016/j.cognition.2013.11.005

236. Penadés, R., Catalán, R., Rubia, K., Andrés, S., Salamero, M., & Gastó, C. (2007). Impaired response inhibition in obsessive compulsive disorder. *European Psychiatry*, 22(6), 404–410. doi:10.1016/j.eurpsy.2006.05.001
237. Perony, N., Tessone, C. J., König, B., & Schweitzer, F. (2012). How Random Is Social Behaviour? Disentangling Social Complexity through the Study of a Wild House Mouse Population. *PLoS Comput Biol*, 8(11), e1002786. doi:10.1371/journal.pcbi.1002786
238. Perruchet, P., Cleeremans, A., & Destrebecqz, A. (2006). Dissociating the effects of automatic activation and explicit expectancy on reaction times in a simple associative learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(5), 955–965. doi:10.1037/0278-7393.32.5.955
239. Perry, J. L., & Carroll, M. E. (2008). The role of impulsive behavior in drug abuse. *Psychopharmacology*, 200(1), 1–26. doi:10.1007/s00213-008-1173-0
240. Peterson, B. S., Skudlarski, P., Anderson, A. W., Zhang, H., Gatenby, J. C., Lacadie, C. M., ... Gore, J. C. (1998). A functional magnetic resonance imaging study of tic suppression in Tourette syndrome. *Archives of General Psychiatry*, 55(4), 326–333.
241. Pfurtscheller, G., & Lopes da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical Neurophysiology*, 110(11), 1842–1857. doi:10.1016/S1388-2457(99)00141-8
242. Picton, T. W., Stuss, D. T., Alexander, M. P., Shallice, T., Binns, M. A., & Gillingham, S. (2007). Effects of Focal Frontal Lesions on Response Inhibition. *Cerebral Cortex*, 17(4), 826–838. doi:10.1093/cercor/bhk031
243. Pinker, S. (2011). *The Better Angels of Our Nature: The Decline of Violence In History And Its Causes*. Penguin UK.

244. Pfurtscheller, G., & Lopes da Silva, F. H. (1999). Event-related EEG/MEG synchronization and desynchronization: basic principles. *Clinical Neurophysiology*, *110*(11), 1842–1857. doi:10.1016/S1388-2457(99)00141-8
245. Polosa, R., Caponnetto, P., Morjaria, J. B., Papale, G., Campagna, D., & Russo, C. (2011). Effect of an electronic nicotine delivery device (e-Cigarette) on smoking reduction and cessation: a prospective 6-month pilot study. *BMC Public Health*, *11*, 786. doi:10.1186/1471-2458-11-786
246. Poo, C., & Isaacson, J. S. (2009). Odor Representations in Olfactory Cortex: “Sparse” Coding, Global Inhibition, and Oscillations. *Neuron*, *62*(6), 850–861. doi:10.1016/j.neuron.2009.05.022
247. Ramautar, J. R., Kok, A., & Ridderinkhof, K. R. (2004). Effects of stop-signal probability in the stop-signal paradigm: The N2/P3 complex further validated. *Brain and Cognition*, *56*(2), 234–252. doi:10.1016/j.bandc.2004.07.002
248. Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, *20*(4), 873–922. doi:10.1162/neco.2008.12-06-420
249. Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, *9*(5), 347–356. doi:10.1111/1467-9280.00067
250. Raz, A., Zhu, H., Yu, S., Bansal, R., Wang, Z., Alexander, G. M., ... Peterson, B. S. (2009). Neural substrates of self-regulatory control in children and adults with Tourette syndrome. *Canadian Journal of Psychiatry. Revue Canadienne De Psychiatrie*, *54*(9), 579–588.
251. Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M. C., Lehericy, S., Bergman, H., ... Obeso, J. A. (2010). Goal-directed and habitual control in the basal ganglia: implications for Parkinson’s disease. *Nature Reviews Neuroscience*, *11*(11), 760–772. doi:10.1038/nrn2915

252. Ridderinkhof, K. R., van den Wildenberg, W. P. M., Segalowitz, S. J., & Carter, C. S. (2004). Neurocognitive mechanisms of cognitive control: The role of prefrontal cortex in action selection, response inhibition, performance monitoring, and reward-based learning. *Brain and Cognition*, *56*(2), 129–140. doi:10.1016/j.bandc.2004.09.016
253. Rigoni, D., Kühn, S., Gaudino, G., Sartori, G., & Brass, M. (2012). Reducing self-control by weakening belief in free will. *Consciousness and Cognition*, *21*(3), 1482–1490. doi:10.1016/j.concog.2012.04.004
254. Robertson, M. M. (2000). Tourette syndrome, associated conditions and the complexities of treatment. *Brain*, *123*(3), 425–462. doi:10.1093/brain/123.3.425
255. Robinson, T. E., & Berridge, K. C. (1993). The neural basis of drug craving: An incentive-sensitization theory of addiction. *Brain Research Reviews*, *18*(3), 247–291. doi:10.1016/0165-0173(93)90013-P
256. Roessner, V., Albrecht, B., Dechent, P., Baudewig, J., & Rothenberger, A. (2008). Normal response inhibition in boys with Tourette syndrome. *Behavioral and Brain Functions : BBF*, *4*, 29. doi:10.1186/1744-9081-4-29
257. Roitman, J. D., & Shadlen, M. N. (2002). Response of Neurons in the Lateral Intraparietal Area during a Combined Visual Discrimination Reaction Time Task. *The Journal of Neuroscience*, *22*(21), 9475–9489.
258. Rosenberg, D. R., Dick, E. L., O’Hearn, K. M., & Sweeney, J. A. (1997). Response-inhibition deficits in obsessive-compulsive disorder: an indicator of dysfunction in frontostriatal circuits. *Journal of Psychiatry and Neuroscience*, *22*(1), 29–38.
259. Rössner V, Müller-Vahl K, Neuner I (2010). PUTS - premonitory urge tics scale: Fragebogen für Kinder. In K. Müller-Vahl (Ed.), *Tourette-Syndrom und*

andere Tic-Erkrankungen im Kindes- und Erwachsenenalter. Berlin: MWV Medizinische Wissenschaftliche Verlagsgesellschaft.

260. Rubia, K., Russell, T., Bullmore, E. T., Soni, W., Brammer, M. J., Simmons, A., ... Sharma, T. (2001). An fMRI study of reduced left prefrontal activation in schizophrenia during normal inhibitory function. *Schizophrenia Research*, *52*(1–2), 47–55. doi:10.1016/S0920-9964(00)00173-0
261. Rubia, K., Smith, A. B., Brammer, M. J., & Taylor, E. (2003). Right inferior prefrontal cortex mediates response inhibition while mesial prefrontal cortex is responsible for error detection. *NeuroImage*, *20*(1), 351–358. doi:10.1016/S1053-8119(03)00275-1
262. Rubia, K., Smith, A. B., Brammer, M. J., Toone, B., & Taylor, E. (2005). Abnormal Brain Activation During Inhibition and Error Detection in Medication-Naive Adolescents With ADHD. *American Journal of Psychiatry*, *162*(6), 1067–1075. doi:10.1176/appi.ajp.162.6.1067
263. Saalmann, Y. B., Pigarev, I. N., & Vidyasagar, T. R. (2007). Neural Mechanisms of Visual Attention: How Top-Down Feedback Highlights Relevant Locations. *Science*, *316*(5831), 1612–1615. doi:10.1126/science.1139140
264. Sadaghiani, S., Hesselmann, G., Friston, K. J., & Kleinschmidt, A. (2010). The Relation of Ongoing Brain Activity, Evoked Neural Responses, and Cognition. *Frontiers in Systems Neuroscience*, *4*. doi:10.3389/fnsys.2010.00020
265. Sato, Y., Akiyama, E., & Farmer, J. D. (2002). Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, *99*(7), 4748–4751. doi:10.1073/pnas.032086299
266. Savelsbergh, G. J. P., Van der Kamp, J., Williams, A. M., & Ward, P. (2005). Anticipation and visual search behaviour in expert soccer goalkeepers. *Ergonomics*, *48*(11-14), 1686–1697. doi:10.1080/00140130500101346

267. Savelsbergh, G. J. P., Williams, A. M., Kamp, J. V. D., & Ward, P. (2002). Visual search, anticipation and expertise in soccer goalkeepers. *Journal of Sports Sciences*, *20*(3), 279–287. doi:10.1080/026404102317284826
268. Scheibe, C., Schubert, R., Sommer, W., & Heekeren, H. R. (2009). Electrophysiological evidence for the effect of prior probability on response preparation. *Psychophysiology*, *46*, 758–770
269. Schel, M. A., Kuhn, S., Brass, M., Haggard, P., Ridderinkhof, K. R., & Crone, E. A. (2014). Neural correlates of intentional and stimulus-driven inhibition: a comparison. *Frontiers in Human Neuroscience*, *8*. doi:10.3389/fnhum.2014.00027
270. Scheres, A., Oosterlaan, J., Geurts, H., Morein-Zamir, S., Meiran, N., Schut, H., ... Sergeant, J. A. (2004). Executive functioning in boys with ADHD: primarily an inhibition deficit? *Archives of Clinical Neuropsychology*, *19*(4), 569–594. doi:10.1016/j.acn.2003.08.005
271. Schneiderman, N., & Gormezano, I. (1964). Conditioning of the nictitating membrane of the rabbit as a function of CS-US interval. *Journal of Comparative and Physiological Psychology*, *57*, 188–195.
272. Schurger, A., Sitt, J. D., & Dehaene, S. (2012). An accumulator model for spontaneous neural activity prior to self-initiated movement. *Proceedings of the National Academy of Sciences*, *109*, E2904–E2913.
273. Schütz-Bosbach, S., & Prinz, W. (2007). Perceptual resonance: action-induced modulation of perception. *Trends in Cognitive Sciences*, *11*(8), 349–355. doi:10.1016/j.tics.2007.06.005
274. Serrien, D. J., Orth, M., Evans, A. H., Lees, A. J., & Brown, P. (2005). Motor inhibition in patients with Gilles de la Tourette syndrome: functional activation patterns as revealed by EEG coherence. *Brain*, *128*(1), 116–125. doi:10.1093/brain/awh318

275. Severens, E., Janssens, I., Kühn, S., Brass, M., & Hartsuiker, R. J. (2011a). When the brain tames the tongue: Covert editing of inappropriate language. *Psychophysiology*, *48*(9), 1252–1257. doi:10.1111/j.1469-8986.2011.01190.x
276. Severens, E., Kühn, S., Hartsuiker, R. J., & Brass, M. (2011b). Functional mechanisms involved in the internal inhibition of taboo words. *Social Cognitive and Affective Neuroscience*. doi:10.1093/scan/nsr030
277. Shadlen, M. N., & Newsome, W. T. (2001). Neural Basis of a Perceptual Decision in the Parietal Cortex (Area LIP) of the Rhesus Monkey. *Journal of Neurophysiology*, *86*(4), 1916–1936.
278. Shibasaki, H., Barrett, G., Halliday, E., & Halliday, A. M. (1980). Components of the movement-related cortical potential and their scalp topography. *Electroencephalography and Clinical Neurophysiology*, *49*(3–4), 213–226. doi:10.1016/0013-4694(80)90216-3
279. Shibasaki, H., & Hallett, M. (2006). What is the Bereitschaftspotential? *Clinical Neurophysiology*, *117*, 2341–2356.
280. Shoda, Y., Mischel, W., & Peake, P. K. (1990). Predicting adolescent cognitive and self-regulatory competencies from preschool delay of gratification: Identifying diagnostic conditions. *Developmental Psychology*, *26*(6), 978–986. doi:10.1037/0012-1649.26.6.978
281. Siegal, M., & Varley, R. (2002). Neural systems involved in “theory of mind.” *Nature Reviews Neuroscience*, *3*(6), 463–471. doi:10.1038/nrn844
282. Simmonds, D. J., Pekar, J. J., & Mostofsky, S. H. (2008). Meta-analysis of Go/No-go tasks demonstrating that fMRI activation associated with response inhibition is task-dependent. *Neuropsychologia*, *46*(1), 224–232. doi:10.1016/j.neuropsychologia.2007.07.015

283. Simpson, A., Riggs, K. J., Beck, S. R., Gorniak, S. L., Wu, Y., Abbott, D., & Diamond, A. (2012). Refining the understanding of inhibitory processes: how response prepotency is created and overcome. *Developmental Science*, *15*(1), 62–73. doi:10.1111/j.1467-7687.2011.01105.x
284. Sirigu, A., Daprati, E., Ciancia, S., Giraux, P., Nighoghossian, N., Posada, A., et al. (2004). Altered awareness of voluntary action after damage to the parietal cortex. *Nature Neuroscience*, *7*, 80–84.
285. Skinner, B. F. (1971). *Beyond Freedom & Dignity*. Indianapolis, IN: Hackett Publishing.
286. Smith, M. C. (1968). CS-US interval and US intensity in classical conditioning of the rabbit's nictitating membrane response. *Journal of Comparative and Physiological Psychology*, *66*(3), 679–687.
287. Smith, J. L., Johnstone, S. J., & Barry, R. J. (2008). Movement related potentials in the go/nogo task: The P3 reflects both cognitive and motor inhibition. *Clinical Neurophysiology*, *119*, 704–714.
288. Somers, D. C., Dale, A. M., Seiffert, A. E., & Tootell, R. B. H. (1999). Functional MRI reveals spatially specific attentional modulation in human primary visual cortex. *Proceedings of the National Academy of Sciences*, *96*(4), 1663–1668. doi:10.1073/pnas.96.4.1663
289. Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nat Neurosci*, *11*(5), 543–545. doi:10.1038/nn.2112
290. Spatt, J., & Goldenberg, G. (1993). Components of random generation by normal subjects and patients with dysexecutive syndrome. *Brain and Cognition*, *23*(2), 231–242. doi:10.1006/brcg.1993.1057

291. Specht, M. W., Woods, D. W., Nicotra, C. M., Kelly, L. M., Ricketts, E. J., Conelea, C. A., ... Walkup, J. T. (2013). Effects of tic suppression: Ability to suppress, rebound, negative reinforcement, and habituation to the premonitory urge. *Behaviour Research and Therapy*, *51*(1), 24–30. doi:10.1016/j.brat.2012.09.009
292. Stivers, T., Enfield, N. J., Brown, P., Englert, C., Hayashi, M., Heinemann, T., ... Levinson, S. C. (2009). Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, *106*(26), 10587–10592. doi:10.1073/pnas.0903616106
293. Sunaert, S., Van Hecke, P., Marchal, G., & Orban, G. A. (2000). Attention to speed of motion, speed discrimination, and task difficulty: An fMRI study. *Neuroimage*, *11*, 612–623.
294. Swann, N. C., Cai, W., Conner, C. R., Pieters, T. A., Claffey, M. P., George, J. S., ... Tandon, N. (2012). Roles for the pre-supplementary motor area and the right inferior frontal gyrus in stopping action: electrophysiological responses and functional and structural connectivity. *NeuroImage*, *59*(3), 2860–2870. doi:10.1016/j.neuroimage.2011.09.049
295. Swann, N., Tandon, N., Canolty, R., Ellmore, T. M., McEvoy, L. K., Dreyer, S., ... Aron, A. R. (2009). Intracranial EEG Reveals a Time- and Frequency-Specific Role for the Right Inferior Frontal Gyrus and Primary Motor Cortex in Stopping Initiated Responses. *The Journal of Neuroscience*, *29*(40), 12675–12685. doi:10.1523/JNEUROSCI.3359-09.2009
296. Swedo, S. E., & Leonard, H. L. (1992). Trichotillomania: An obsessive compulsive spectrum disorder? *Psychiatric Clinics of North America*, *15*(4), 777–790.
297. Symons, C. S., & Johnson, B. T. (1997). The self-reference effect in memory: a meta-analysis. *Psychological Bulletin*, *121*(3), 371–394.
298. Tamm, L., Menon, V., Ringel, J., & Reiss, A. L. (2004). Event-related FMRI evidence of frontotemporal involvement in aberrant response inhibition and task

switching in attention-deficit/hyperactivity disorder. *Journal of the American Academy of Child and Adolescent Psychiatry*, 43(11), 1430–1440.

doi:10.1097/01.chi.0000140452.51205.8d

299. Tangney, J. P., Baumeister, R. F., & Boone, A. L. (2004). High Self-Control Predicts Good Adjustment, Less Pathology, Better Grades, and Interpersonal Success. *Journal of Personality*, 72(2), 271–324. doi:10.1111/j.0022-3506.2004.00263.x
300. Tanji, J., & Hoshi, E. (2008). Role of the lateral prefrontal cortex in executive behavioral control. *Physiological Reviews*, 88, 37–57.
301. Tanji, J., & Shima, K. (1994). Role for supplementary motor area cells in planning several movements ahead. *Nature*, 371, 413–416.
302. Taylor, S., Koch, W. J., & McNally, R. J. (1992). How does anxiety sensitivity vary across the anxiety disorders? *Journal of Anxiety Disorders*, 6(3), 249–259. doi:10.1016/0887-6185(92)90037-8
303. Thaler, D., Chen, Y. C., Nixon, P. D., Stern, C. E., & Passingham, R. E. (1995). The functions of the medial premotor cortex. I. Simple learned movements. *Experimental Brain Research*, 102(3), 445–460.
304. Thomalla, G., Jonas, M., Bäumer, T., Siebner, H. R., Biermann-Rubén, K., Ganos, C., ... Münchau, A. (2014). Costs of control: decreased motor cortex engagement during a Go/NoGo task in Tourette's syndrome. *Brain*, 137(1), 122–136. doi:10.1093/brain/awt288
305. Toufexis, M., Deoleo, C., Elia, J., & Murphy, T. K. (2014). A Link Between Perianal Strep and Pediatric Autoimmune Neuropsychiatric Disorder Associated With Streptococcal Infection (PANDAS). *The Journal of Neuropsychiatry and Clinical Neurosciences*, 26(2), 164–168. doi:10.1176/appi.neuropsych.12050126

306. Treue, S., & Maunsell, J. H. R. (1996). Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*, 382(6591), 539–541. doi:10.1038/382539a0
307. Tune, G. S. (1964). Response preferences: A review of some relevant literature. *Psychological Bulletin*, 61(4), 286-302.
308. Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105-110.
309. Van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *The Journal of Neuroscience*, 28(32), 8053–8062. doi:10.1523/JNEUROSCI.1278-08.2008
310. Van Gaal, S., Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. A. F. (2010). Unconscious activation of the prefrontal no-go network. *The Journal of Neuroscience*, 30(11), 4143–4150. doi:10.1523/JNEUROSCI.2992-09.2010
311. Van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P. M., & Lamme, V. A. F. (2009). Dissociating consciousness from inhibitory control: evidence for unconsciously triggered response inhibition in the stop-signal task. *Journal of Experimental Psychology. Human Perception and Performance*, 35(4), 1129–1139. doi:10.1037/a0013551
312. Van Os, J., Hanssen, M., Bijl, R. V., & Ravelli, A. (2000). Strauss (1969) revisited: a psychosis continuum in the general population? *Schizophrenia Research*, 45(1-2), 11–20.
313. Van Os, J., Linscott, R. J., Myin-Germeys, I., Delespaul, P., & Krabbendam, L. (2009). A systematic review and meta-analysis of the psychosis continuum: evidence for a psychosis proneness–persistence–impairment model of psychotic disorder. *Psychological Medicine*, 39(02), 179–195. doi:10.1017/S0033291708003814

314. Vandekerckhove, J., & Tuerlinckx, F. (2008). Diffusion model analysis with MATLAB: a DMAT primer. *Behavior Research Methods*, *40*(1), 61–72.
315. Verbruggen, F., & Logan, G. D. (2008). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, *12*(11), 418–424.
doi:10.1016/j.tics.2008.07.005
316. Verdellen, C. W. J., Hoogduin, C. A. L., & Keijsers, G. P. J. (2007). Tic suppression in the treatment of Tourette's syndrome with exposure therapy: The rebound phenomenon reconsidered. *Movement Disorders*, *22*(11), 1601–1606.
doi:10.1002/mds.21577
317. Vickery, T. J., Chun, M. M., & Lee, D. (2011). Ubiquity and Specificity of Reinforcement Signals throughout the Human Brain. *Neuron*, *72*(1), 166–177.
doi:10.1016/j.neuron.2011.08.011
318. Wagenaar, W. A. (1972). Generation of random sequences by human subjects: A critical survey of literature. *Psychological Bulletin*, *77*(1), 65–72.
doi:10.1037/h0032060
319. Walsh, E., Kuhn, S., Brass, M., Wenke, D., & Haggard, P. (2010). EEG activations during intentional inhibition of voluntary action: An electrophysiological correlate of self-control? *Neuropsychologia*, *48*(2), 619–626.
doi:10.1016/j.neuropsychologia.2009.10.026
320. Ward, M. F., Wender, P. H., & Reimherr, F. W. (1993). The Wender Utah Rating Scale: an aid in the retrospective diagnosis of childhood attention deficit hyperactivity disorder. *The American Journal of Psychiatry*, *150*(6), 885–890
321. Wasserman, S., & Bockenholt, U. (1989). Bootstrapping: Applications to Psychophysiology. *Psychophysiology*, *26*(2), 208–221. doi:10.1111/j.1469-8986.1989.tb03159.x

322. Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition, 115*(1), 26–38. doi:10.1016/j.cognition.2009.10.016
323. Wenzlaff, H., Bauer, M., Maess, B., & Heekeren, H. R. (2011). Neural Characterization of the Speed–Accuracy Tradeoff in a Perceptual Decision-Making Task. *The Journal of Neuroscience, 31*(4), 1254–1266. doi:10.1523/JNEUROSCI.4000-10.2011
324. Wenzlaff, R. M., & Wegner, D. M. (2000). Thought Suppression. *Annual Review of Psychology, 51*(1), 59–91. doi:10.1146/annurev.psych.51.1.59
325. Wessel, J. R., Conner, C. R., Aron, A. R., & Tandon, N. (2013). Chronometric Electrical Stimulation of Right Inferior Frontal Cortex Increases Motor Braking. *The Journal of Neuroscience, 33*(50), 19611–19619. doi:10.1523/JNEUROSCI.3468-13.2013
326. Wilhelm, S., Deckersbach, T., Coffey, B. J., Bohne, A., Peterson, A. L., & Baer, L. (2003). Habit Reversal Versus Supportive Psychotherapy for Tourette’s Disorder: A Randomized Controlled Trial. *American Journal of Psychiatry, 160*(6), 1175–1177. doi:10.1176/appi.ajp.160.6.1175
327. Woodman, G. F., & Luck, S. J. (2003). Serial deployment of attention during visual search. *Journal of Experimental Psychology: Human Perception and Performance, 29*(1), 121–138. doi:10.1037/0096-1523.29.1.121
328. Woods, D. W., Piacentini, J., Himle, M. B., & Chang, S. (2005). Premonitory Urge for Tics Scale (PUTS): initial psychometric results and examination of the premonitory urge phenomenon in youths with Tic disorders. *Journal of Developmental and Behavioral Pediatrics: JDBP, 26*(6), 397–403.
329. Wooldridge, J.M. (2002). *Econometric analysis of cross section and panel data* (Cambridge, MA: MIT Press).

330. Woolley, J., Heyman, I., Brammer, M., Frampton, I., McGuire, P. K., & Rubia, K. (2008). Brain activation in paediatric obsessive–compulsive disorder during tasks of inhibitory control. *The British Journal of Psychiatry*, *192*(1), 25–31. doi:10.1192/bjp.bp.107.036558
331. Xue, G., Aron, A. R., & Poldrack, R. A. (2008). Common Neural Substrates for Inhibition of Spoken and Manual Responses. *Cerebral Cortex*, *18*(8), 1923–1932. doi:10.1093/cercor/bhm220
332. Yamanaka, K., & Yamamoto, Y. (2009). Single-trial EEG Power and Phase Dynamics Associated with Voluntary Response Inhibition. *Journal of Cognitive Neuroscience*, *22*(4), 714–727. doi:10.1162/jocn.2009.21258
333. Yarrow, K., Brown, P., & Krakauer, J. W. (2009). Inside the brain of an elite athlete: the neural processes that support high achievement in sports. *Nature Reviews Neuroscience*, *10*(8), 585–596. doi:10.1038/nrn2672