

2019

Action Intentions and the Puzzle of Non-Willed Behaviour

Colton, James

<http://hdl.handle.net/10026.1/14644>

<http://dx.doi.org/10.24382/814>

University of Plymouth

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's prior consent.



**UNIVERSITY OF
PLYMOUTH**

**ACTION INTENTIONS AND THE PUZZLE OF
NON-WILLED BEHAVIOUR**

by

JAMES COLTON

A thesis submitted to the University of Plymouth
in partial fulfilment for the degree of

DOCTOR OF PHILOSOPHY

School of Psychology

June 2019

Acknowledgements

I would first like to thank Hannah for her love, support and unique ability to detect whether I was in need of a kick or a cuddle in the face of insurmountable odds and all flavours of personal strife. We travelled once around the world in going back and forth between Bournemouth and Plymouth and still found time for a life (and a wedding) outside of work. For that, and many other things, I am eternally grateful. My family also deserve special mention, for providing all kinds of help and support. Not least a bed to sleep in and more than my fair share of food from the fridge, but also for their part in providing spare cars, motorcycles, homemade diesel fuel and facilitating four house moves.

I would especially like to thank my entire supervisory team for their considerable expertise, tireless support (both technical and emotional) and extreme patience in the face of a never-ending torrent of emails. In particular, I would like to thank Prof. Christopher Mitchell (DoS) for encouraging me to apply for this PhD, for his wise guidance during the subsequent academic odyssey and the insightful commentary he consistently delivered throughout. I would also like to thank Dr Patric Bach for his positive attitude and endless enthusiasm, his ability to find the value in my half-baked ideas and for his ability to ask the right combination of leading questions to convince me that I'd found my own way back to the righteous (i.e. ideomotor) path. Finally, I also owe a debt of gratitude to Dr Ben Whalley for repeatedly digging me out of an enormous variety of conceptual, technical and statistical holes with his ability to find new ways of explaining phenomenally complex ideas in a way that a mere mortal could understand.

I would also like to thank Plymouth University's Action-Perception Lab – Patric, Katrina, Helen, Eddie, Ellie, Toby, Kim and Matt. You made my entire PhD experience

and created the perfect convivial atmosphere for me to get over my imposter syndrome and speak in front of colleagues about my work. Together we conquered our first posters, first talks, first papers and, even, first brain dissection. Having a group of friends like you to celebrate and commiserate the highs and lows of academic life made the last four years worthwhile.

Finally, I would like to offer an enormous thank you to Plymouth University's Tech Office for their incredible support. In particular, I would like to single out the contribution of Roger Creagh-Osborne who was able to translate my e-Prime experiment and create custom software and hardware which made the pressure pad version of the intention insertion task possible. He took an active and enthusiastic interest in the project, incorporated every feature I could possibly dream of and even some that I hadn't even realised I might need (until I did). Roger went above and beyond for me during my PhD and taught me that you never truly understand an idea until you've had to explain it to an engineer.

Author's Declaration

At no time during the registration for the degree of *Doctor of Philosophy* has the author been registered for any other University award without prior agreement of the Doctoral College Quality Sub-Committee.

Work submitted for this research degree at the University of Plymouth has not formed part of any other degree either at the University of Plymouth or at another establishment.

This study was financed with the aid of a studentship from Plymouth University School of Psychology. Relevant scientific seminars and conferences were attended at which work was often presented.

Publications:

Colton, J., Bach, P., Whalley, B., & Mitchell, C. (in prep.). Goal state or current state? Illusory perception of body ownership precludes the generation of non-willed behaviour.

Colton, J., Bach, P., Whalley, B., & Mitchell, C. (in prep.). Intention insertion: the facilitative effects of movement observation on the generation of non-willed behaviour.

Colton, J., Bach, P., Whalley, B., & Mitchell, C. (2018a). Intention insertion: activating an action's perceptual consequences is sufficient to induce non-willed motor behavior. *Journal of Experimental Psychology: General*, 147(8), 1256–1263.
<https://doi.org/10.1037/xge0000435>

Colton, J., Bach, P., Whalley, B., & Mitchell, C. (2018b). Intention insertion: activating an action's perceptual consequences is sufficient to induce non-willed motor behaviour (Analysis and Materials). Open Science Framework. Retrieved from <https://osf.io/s9qgu/>

Presentations at conferences:

Colton, J., Bach, P., Whalley, B., Mitchell, C. (2017, May). Acting in spite of intention: action-effect anticipation is sufficient to induce non-willed motor responses. A poster presentation delivered at the 2017 Concepts, Actions & Objects (CAOs) conference, Centre for Mind and Brain Sciences, University of Trento, Rovereto, Italy.

Colton, J., Bach, P., Whalley, B., Mitchell, C. (2016, July). From ideomotor suggestions to motor control. An oral presentation delivered at the Cognition Institute Conference. Plymouth University, UK.

Colton, J., Bach, P., Whalley, B., Mitchell, C. (2015, June). Action in spite of intention: Imagination and Ideomotor Suggestion. A poster presentation delivered at the annual School of Psychology Conference, Plymouth University, UK.

Word count of main body of thesis: 52,871

Signed:

Date:

Abstract

Action Intentions and the Puzzle of Non-Willed Behaviour

James Colton

It feels intuitive that our actions are intentional, but there is considerable debate about whether (and how) humans control their behaviour. Ideomotor models of action argue that action intentions have a fundamentally perceptual format. In these views, actions are not only *controlled* by anticipating – imagining – their intended perceptual consequences but are also *initiated* when this “action effect” activation is sufficiently strong. Yet, this latter initiating function of action effect activation has not yet been tested directly.

In Experiments 1-4, participants mentally rehearsed a movement sequence and were unexpectedly presented with salient visual cues that were either compatible or incompatible with their currently imagined action. As predicted, the combined activity from imagery and perception was sufficient to trigger non-willed action slips, even when participants were asked to withhold responses. The experiments provided the first direct evidence that forming an action intention may involve nothing more than evoking a strong enough mental image of the perceptual effect one wishes to achieve until a motor threshold is reached and the corresponding action is initiated. Experiments 5-6 showed that anticipated action effects attain this initiating role when they are treated as potential future goal states, rather than the current state of the sensorimotor system. Finally, Experiment 7 provided evidence that the tendency to produce action slips is positively associated with individual differences in ideomotor suggestibility.

These findings reveal how people are able to interact so effortlessly with the world by specifying lightly constraining action goals which mediate the translation of

sensory input into motor output - allowing the environment to trigger actions which correspond with the goal. In addition, these findings suggest that action slips are a product of the same processes that guide voluntary actions. That is, perceptual input can co-opt endogenously activated action effects (e.g. via imagery) and, sometimes, trigger inappropriate behaviour.

Table of Contents

ACKNOWLEDGEMENTS	I
AUTHOR'S DECLARATION	III
ABSTRACT	V
TABLE OF CONTENTS	VII
LIST OF FIGURES	X
1. CHAPTER 1: THE PUZZLE OF NON-WILLED ACTION	1
1.1. OVERVIEW	1
1.1.1 <i>Mundane Errors of Action Control</i>	4
1.1.2 <i>Esoteric Errors of Action Control</i>	6
1.2. THE PERCEPTUAL FORMAT OF ACTION INTENTIONS	9
1.3. THE INTERPLAY OF ENDOGENOUS AND EXOGENOUS PROCESSES IN ACTION	15
1.4. OVERVIEW OF THE THESIS	22
2. CHAPTER 2: THE RELATIONSHIP BETWEEN ACTION INTENTIONS AND INTERNAL REPRESENTATIONS IN VOLITIONAL ACTION CONTROL.....	26
2.1. OVERVIEW	26
2.1.1 <i>Introduction</i>	26
2.2. EXPERIMENT 1: ELICITING NON-WILLED ACTION SLIPS DURING IMAGERY WITH NEWLY LEARNED AND PREVIOUSLY ESTABLISHED ACTION EFFECT ASSOCIATIONS	29
2.2.1 <i>Method</i>	32
2.2.2 <i>Results</i>	36
2.2.3 <i>Discussion</i>	42
2.3. EXPERIMENT 2: ELICITING NON-WILLED ACTION SLIPS DURING IMAGERY WITHOUT TRAINED ASSOCIATIONS.....	46
2.3.1 <i>Method</i>	47
2.3.2 <i>Results</i>	48
2.3.3 <i>Discussion</i>	54
2.4. EXPERIMENT 3: ELICITING NON-WILLED ACTION SLIPS DURING IMAGERY WITH ACTION OBSERVATION	55

2.4.1.	<i>Method</i>	56
2.4.2.	<i>Results</i>	57
2.4.3.	<i>Discussion</i>	63
2.5.	GENERAL DISCUSSION	63
3.	CHAPTER 3: DO ACTION CUES PRIMARILY FACILITATE OR INTERFERE WITH THE GENERATION OF NON-WILLED ACTION SLIPS?.....	69
3.1.	INTRODUCTION.....	69
3.2.	EXPERIMENT 4: INTENTION INSERTION TASK 2.0 – FINGER PRESSURE AS AN IMPROVED BEHAVIOURAL MEASURE OF NON-WILLED ACTION SLIPS	73
3.2.1.	<i>Method</i>	76
3.2.2.	<i>Results</i>	80
3.2.3.	<i>Discussion</i>	92
3.3.	CONCLUSION.....	97
4.	CHAPTER 4: WHY DON'T WE SELF-IMITATE? CURRENT AND GOAL STATES IN SENSORIMOTOR PROCESSING	99
4.1.	INTRODUCTION.....	99
4.2.	EXPERIMENT 5: THE EFFECT OF A BODY OWNERSHIP ILLUSION ON THE GENERATION OF NON-WILLED ACTION SLIPS	102
4.2.1.	<i>Method</i>	106
4.2.2.	<i>Results</i>	108
4.2.3.	<i>Discussion</i>	114
4.3.	EXPERIMENT 6: THE EFFECT OF A BODY OWNERSHIP ILLUSION ON THE GENERATION OF NON-WILLED ACTION SLIPS WITHOUT VISUAL FEEDBACK	117
4.3.1.	<i>Method</i>	119
4.3.2.	<i>Results</i>	120
4.3.3.	<i>Discussion</i>	129
4.4.	GENERAL DISCUSSION	131
4.5.	CONCLUSION.....	137
5.	CHAPTER 5: THE RELATIONSHIP BETWEEN THE GENERATION OF ACTION SLIPS AND IDEOMOTOR (HYPNOTIC) SUGGESTIBILITY	139

5.1. OVERVIEW	139
5.2. INTRODUCTION.....	139
5.2.1. <i>Action Slips and Existing Perspectives on Hypnotic Responding.</i>	142
5.3. EXPERIMENT 7: A TEST OF THE RELATIONSHIP BETWEEN THE INTENTION INSERTION EFFECT AND IDEOMOTOR SUGGESTIBILITY	145
5.3.1. <i>Method</i>	146
5.3.2. <i>Results</i>	149
5.3.3. <i>Discussion</i>	154
5.4. CONCLUSION.....	159
6. CHAPTER 6: GENERAL DISCUSSION.....	161
6.1. INTRODUCTION.....	161
6.2. SUMMARY OF FINDINGS	163
6.3. THE PERCEPTUAL FORMAT OF ACTION INTENTIONS	166
6.4. VOLITION IN ACTION.....	173
6.4.1. <i>Volitional Acts are Endogenously Driven</i>	173
6.4.2. <i>Volitional Acts are Goal-directed</i>	175
6.4.3. <i>Volitional Acts Have a Distinctive Subjective Experience</i>	176
6.4.4. <i>The Broader Question of Volitional Control of Mental Events</i>	177
6.4.5. <i>Conclusion</i>	179
6.5. OPEN QUESTIONS AND FUTURE DIRECTIONS	180
6.5.1. <i>Subjective Experience of Non-Willed Action</i>	181
6.5.2. <i>Direct Assessment of Current/Goal State Proposal</i>	182
6.5.3. <i>Attending to Current or Anticipated Sensation in Predictive Processing</i>	183
6.5.4. <i>The Brain Basis of Volition in Action</i>	185
6.6. CONCLUSION.....	186
REFERENCES	189

List of Figures

FIGURE 1. A SCHEMATIC REPRESENTATION OF THE ASSOCIATION TRAINING TASK OF EXPERIMENT 1.	33
FIGURE 2. A SCHEMATIC REPRESENTATION OF A 5-BEAT COMPATIBLE CUED TRIAL FROM EXPERIMENT 1.	35
FIGURE 3. A BOX PLOT OF THE OBSERVED FREQUENCY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 1.....	39
FIGURE 4. A COMBINATION PLOT INDICATING THE PREDICTED PROBABILITY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM ACROSS EXPERIMENT 1.	40
FIGURE 5. A PLOT OF THE PREDICTED PROBABILITIES OF SPECIFIC ACTION SLIPS, DERIVED BY SIMULATION, ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 1.	41
FIGURE 6. A BOX PLOT OF THE OBSERVED FREQUENCY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 2.....	50
FIGURE 7. A COMBINATION PLOT INDICATING THE PREDICTED PROBABILITY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM ACROSS EXPERIMENT 2.	51
FIGURE 8. A PLOT OF THE PREDICTED PROBABILITIES OF SPECIFIC ACTION SLIPS, DERIVED BY SIMULATION, ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 2..	53
FIGURE 9. TYPICAL AND UNEXPECTED MOVEMENT CUES DISPLAYED DURING IMAGINATION PHASE OF EXPERIMENT 3.	57

FIGURE 10. A BOX PLOT OF THE OBSERVED FREQUENCY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 3.....	59
FIGURE 11. A COMBINATION PLOT INDICATING THE PREDICTED PROBABILITY OF GENERAL ACTION SLIPS ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM ACROSS EXPERIMENT 3.	61
FIGURE 12. A PLOT OF THE PREDICTED PROBABILITIES OF SPECIFIC ACTION SLIPS, DERIVED BY SIMULATION, ACCORDING TO THE CLASS OF ACTION CUE WHICH PRECEDED THEM IN EXPERIMENT 3.	62
FIGURE 13. THE CUSTOM HARDWARE DEVELOPED FOR THE PRESSURE-BASED VERSION OF THE INTENTION INSERTION TASK IN EXPERIMENT 4.	77
FIGURE 14. A MULTI-PANEL PLOT SHOWING FINGER PRESSURE IN THE CRITICAL 800MS INTERVAL FOLLOWING THE ONSET OF THE ONSET OF A TYPICAL METRONOME BEAT OR THE ONSET OF A SALIENT VISUAL CUE IN EXPERIMENT 4.	85
FIGURE 15. A MULTI-PANEL PLOT SHOWING FINGER PRESSURE IN THE CRITICAL 800MS INTERVAL FOLLOWING THE ONSET OF THE ONSET OF EACH TYPE OF VISUAL CUE IN EXPERIMENT 4.	88
FIGURE 16. A MULTI-PANEL PLOT SHOWING DIFFERENCES IN PRESSURE (AUC) DURING INTERVALS WHEN THE RECORDED FINGER WAS EITHER CONGRUENT OR INCONGRUENT WITH RESPECT TO THE OBSERVED MOVEMENT VS. A NEUTRAL CUE IN EXPERIMENT 4.	89
FIGURE 17. THREE EXAMPLES FROM THE UPDATED STIMULUS SET USED IN EXPERIMENT 5. THE LEFT IMAGE DEPICTS TWO GLOVED HANDS IN A NEUTRAL POSITION WITH A CENTRAL FIXATION CROSS.....	107

FIGURE 18. A MULTI-PANEL PLOT SHOWING FINGER PRESSURE IN THE CRITICAL 800MS INTERVAL FOLLOWING THE ONSET OF THE ONSET OF EACH TYPE OF VISUAL CUE IN EXPERIMENT 5.	109
FIGURE 19. A MULTI-PANEL PLOT SHOWING DIFFERENCES IN PRESSURE (AUC) DURING INTERVALS WHEN THE RECORDED FINGER WAS EITHER CONGRUENT OR INCONGRUENT WITH RESPECT TO THE OBSERVED MOVEMENT VS. A NEUTRAL CUE IN EXPERIMENT 5.	110
FIGURE 20. A MULTI-PANEL PLOT SHOWING FINGER PRESSURE IN THE CRITICAL 800MS INTERVAL FOLLOWING THE ONSET OF THE ONSET OF EACH TYPE OF VISUAL CUE IN EXPERIMENT 6.	122
FIGURE 21. A MULTI-PANEL PLOT SHOWING DIFFERENCES IN PRESSURE (AUC) DURING INTERVALS WHEN THE RECORDED FINGER WAS EITHER CONGRUENT OR INCONGRUENT WITH RESPECT TO THE OBSERVED MOVEMENT VS. A NEUTRAL CUE.	123
FIGURE 22. A MULTI-PANEL PLOT SHOWING DIFFERENCES IN FINGER PRESSURE DURING INTERVALS FOLLOWING NEUTRAL VS. OTHER TYPES OF ACTION CUE AGAINST IMAGERY AND ACROSS THE FULL RANGE OF IDEOMOTOR SUGGESTIBILITY ON THE ORIGINAL SCALE (I.E. FROM THE MINIMUM, ZERO, TO THE MAXIMUM, ONE, THEORETICAL SCORE).....	151

1. Chapter 1: The Puzzle of Non-Willed Action

1.1. Overview

Many people will have experienced the frustration of putting their car keys in the fridge while leaving a jug of milk out on the kitchen worktop. Or the embarrassment of answering “Yes” when asked by an aeroplane steward if one would like the chicken or the fish. As 19th-century Psychologist William James (1890) noted, “Who is there that has never wound up his watch on taking off his waistcoat in the daytime, or taken his latch-key out on arriving at the doorstep of a friend?” Freud (1914) studied these strange deviations from typical behaviour which he referred to as “Fehlleistungen” (or “faulty actions” in English) in “The Psychopathology of Everyday Life”. Freud’s writings gave rise to the phrase “Freudian slip” which is commonly used to describe errors of speech, memory or action which, according to Freud, reveal one’s unconscious desires and innermost feelings. In the modern age, the phrase “action slip” was deployed as a catch-all term to describe these kinds of instances of unusual behaviour (Reason, 1984). However, the precise origin of action slips and their broader relevance to the mechanisms which guide voluntary action have not since received a great deal of empirical attention.

There are three things of note with respect to this anomalous behaviour. First, action slips are, mercifully, rare in one’s everyday life, but are common in the sense that most people will have had such experiences at some point in their lives (Reason, 1984). Second, action slips are not *random* events but tend to reflect the impact of environmental cues or one’s inner train of thought (e.g. if one incorrectly anticipated a yes/no question from the aforementioned aeroplane steward). Finally, and perhaps most importantly, action slips are often referred to as being “unintentional” and yet, objectively, the erroneous behaviour must be self-generated. Accordingly, the incentive

for studying these kinds of slips of action control is that, far from being random events, they are in all likelihood borne of the same processes which guide all human behaviour. As such, the nature of our mistakes and the situations in which they occur can provide unique insights into the origins of our intentions and the underlying mechanisms of volitional action. Indeed, if action slips do represent a curious by-product of processes involved in generating voluntary actions, then it is important to explore how they can be accommodated by modern theories of action control. This line of thinking motivated the question at the centre of this thesis: how is it that people can produce behaviour if it is not, in some sense, willed?

Willed action is one of the biggest puzzles in psychology: do humans *will* their actions, or do they just follow previously established stimulus-response associations, or respond to affordances offered by the environment (e.g., Tucker & Ellis, 1998)? While a long tradition of psychological research, starting with behaviourism (Hull, 1943), has denied the possibility of voluntary action, or at least exempted it from rigorous psychological enquiry, a parallel tradition has made it a central topic. Ideomotor theories of action argue that people do have voluntary control over their actions, and that these action intentions have a fundamentally perceptual format (Hommel, 2009; Hommel, Müsseler, Aschersleben, & Prinz, 2001; James, 1890; Kunde, Elsner, & Kiesel, 2007; Pacherie & Haggard, 2010; Prinz, 1997; Shin, Proctor, & Capaldi, 2010). Ideomotor models claim that the processes which serve action, imagery and perception are functionally intertwined because they share a common language – that of perceptually-coded internal representations – or “action effects” (e.g. Common Coding theory, Prinz, 1990, 1997; van der Wel, Sebanz, & Knoblich, 2013). To initiate an action, these models argue, humans simply bring to mind (e.g. via imagery) the desired action effects. In other words, they anticipate how the action would look, feel and

sound, and how it would affect the environment. Automatic motor processes then make this imagined action reality.

Ideomotor models – and more recently those based on predictive processing (e.g. Active Inference, Adams, Shipp, & Friston, 2013; and Predictive Coding, Clark, 2013) – argue that perceptual control of action is possible because people, through a lifetime of interactions, form bi-directional associations between their motor behaviours and the sensory input that typically follows. They learn, for example, that tensing a specific muscle group will pull one's hand to the left, or that pressing a brake pedal will slow a car (Turella et al., 2016). As these sensorimotor associations are formed, willed behaviour becomes possible (James, 1890): one merely needs to think of (i.e. imagine) the sensory consequences one wants to achieve, and via the now established associations, the relevant motor behaviours are elicited. After sufficient practice, these acts of motor imagery become so effective at triggering the corresponding behaviours that the sophisticated motor processes, which produce them, operate largely outside of awareness. That is, we simply imagine our hand moving, and it does.

Importantly, ideomotor models assume that the internal perceptual representations which can be activated in a top-down manner during voluntary action (i.e. in order to achieve or advance towards some internally specified goal), can also be activated, bottom-up, by perceptual input. This could be achieved in two ways. On one hand, perceptual input could act as a retrieval cue for the action effects that one typically activates in response to such a stimulus. For example, when one encounters brake lights in a queue of traffic, this stimulus almost reflexively causes a press of the brake pedal to slow the car and avoid a collision. On the other hand, perceptual input could also drive action directly if, for example, the input corresponds with an action image that one could generate oneself. Indeed, the tendency for people to imitate other people's actions is a striking illustration of the potential for actions one observes to directly influence

one's ongoing behaviour when they strongly correspond with actions in one's motor repertoire (Chartrand & Bargh, 1999). Tentatively, this suggests that action slips might result from one's train of thought (i.e. mental imagery) or cues in the environment (i.e. perceptual input) inadvertently triggering an inappropriate action. In both cases, however, the resulting action slip is mediated by the activation of an internal representation of the action outcome. However, prior to considering models of action control in more detail, it is useful to first consider in more detail how closely this proposal maps onto the way in which action slips present in everyday life.

1.1.1 Mundane Errors of Action Control

Reason (1984) coined the term “action slips” to describe the variety of accidental errors of action control that people produce and conducted a study to explore their frequency and form. Action slips are not well suited to laboratory study, given that they are rare events, so participants were asked to keep a diary of their behaviour and record any aberrant behaviour over a seven day period. Reason (1984) developed a taxonomy with four distinct classes of action slip, distinguished by their behavioural characteristics, (a) Repetition, wherein actions in the intended sequence were repeated unnecessarily; (b) Wrong Object(s), where the intended actions were executed but directed at inappropriate objects; (c) Omission, where intended actions were left out of a particular action sequence; and (d) Intrusion, where novel unintended actions became incorporated into one's ongoing behaviour. Participants reported instances such as, “As I approached the turnstile on my way out of the library, I pulled my wallet out of my pocket as if to pay, although no money was required” (p. 544) and “I intended to close the window because it was cold, but closed the cupboard instead” (p.545).

Of course, the taxonomic labels (derived by factor analysis of questionnaire data) were not intended to represent an exhaustive or mutually exclusive description of action slips, but this study nevertheless provided many interesting cases which are useful to consider in the context of modern models of motor control. To illustrate, a typical tea making error provides such an example. A good cup of tea provides an ideal start to the day for many people and the task of preparing a brew involves a familiar ritual. Boil the kettle, place the teabag in the cup, pour boiling water onto the teabag and brew for 3-5 minutes, before removing the tea bag and adding the milk (“How to Make a Proper Brew,” 2019). However, on a particularly absent-minded morning, one might find oneself reaching for the milk jug, rather than the freshly boiled kettle, and pouring cold milk rather than hot water onto the teabag - ruining the day before it has even begun. But what might have caused this unusual error to occur?

For one, both the kettle and the milk jug are of a similar size, have a handle which affords a similar grasping action and, indeed, both are necessary for a proper cup of tea. For another, one might have just noticed that the milk was about to run out, inadvertently initiating a reaching and grasping towards the milk instead of adding milk to a mental shopping list. These observations highlight the role that cues in the environment (i.e. object affordances) and one’s inner train of thought (i.e. mental imagery) may play in the generation of action slips. That is, such action slips may be caused by, (1) poorly specified or monitored goals; (2) one’s goals simply being disrupted by salient information in the environment, or; (3) some combination of the two. In fact, many of the cases described by Reason (1984), indicated that both internal imagery and external sensory information contribute to the generation of action slips.

1.1.2 Esoteric Errors of Action Control

In the previous section, an unfortunate tea making error was used to illustrate some of the potential origins of the kinds of action slips many people will experience in everyday life. However, as noted earlier, these kinds of everyday action slips have not received a great deal of empirical attention, because they are rare events which appear to occur spontaneously. This presents a challenge for an empirical enquiry into the origin of action slips. As a result, it is useful to consider research concerning other instances of unusual behaviour which share the characteristic sense of involuntariness that accompanies everyday action slips. Classical ideomotor phenomena such as hypnosis, water dowsing, magic swinging pendula and the use of an Ouija board for contacting the dead have a long history in Psychological science (Stock & Stock, 2004). They stand as another test case of a similar kind of non-willed behaviour to the action slips described earlier, but that can be more readily reproduced in laboratory conditions.

For example, when subjected to hypnosis, some people appear in thrall to their hypnotist's suggestions, as if at that moment they exert no control over their own behaviour. The striking aspect of hypnotic behaviour is not *just* that the subject moves in accordance with the hypnotist's instructions; it is the subject's belief that they were not the causal agents of their actions and the apparent conviction with which this belief is held. Another example is provided by Chevreul's (1833) magic swinging pendulum, wherein superstitious individuals claim that the motion of a pendulum held over the belly of a pregnant woman can be used to determine the sex of her unborn child (Wegner, Ansfield, & Pilloff, 1998). Responsive individuals report the sense that supernatural forces, rather than their own actions, compel the pendulum to oscillate. Taken together, ideomotor phenomena serve as remarkable examples of how it is possible to blur the causal relationship between our intentions and our behaviour, even in the absence of any underlying clinical diagnosis. For the present purposes, ideomotor

phenomena are considered as particularly unusual manifestations of the same processes which give rise to action slips. One element common to each of these examples, that draws clear parallels with the action slips described by James (1890), Freud (1914) and Reason (1984), is that people commonly report that their behaviour is unintentional or somehow *non-willed*. Indeed, with both action slips and classical ideomotor phenomena, one's behaviour feels subjectively non-willed in spite of the fact that, objectively, it must be self-generated¹. Ideomotor principles were initially applied to debunking these supposedly paranormal phenomena which have now become interesting case studies in modern Psychology. As such, research concerning ideomotor phenomena can provide an important empirical insight into the interplay between internal and external processes in the generation of non-willed action slips.

Consider the example of Chevreul's pendulum. Chevreul (1833) was the first to demonstrate that the pendulum movements were eliminated if the subject's arm was supported, and the subject's movements were constrained. He argued that this finding indicated that the motion of the pendulum was, in fact, caused by tiny, imperceptible movements of the subject's arm and postural sway, rather than supernatural forces. Chevreul (1833) proposed that vivid mental imagery and one's expectations about the anticipated motion of the pendulum bob are the true causes of the movement. In other words, the individual holding the pendulum has an expectation of the sex of the mother's unborn baby, and the motion of the pendulum is merely a subconscious expression of that belief. In the modern age, this phenomenon was the subject of a formal empirical enquiry conducted by Easton and Shor (1975, 1976, 1977). In these experiments, individuals were specifically instructed to remain still, focus their attention

¹ As an interesting historical aside, the name "Ouija" is actually a trademark of Hasbro games who created the familiar planchette and board (McRobbie, 2013). Bizarrely, it was originally marketed as both a mysterious oracle capable of delivering prophesies from the netherworld and wholesome family entertainment. However, in order to be granted a patent, Hasbro had to prove that their Ouija board worked as described. The Chief Patent Officer insisted upon a demonstration and demanded the board be used to spell his name, apparently unknown to its creators. During the subsequent séance, the planchette accurately spelt out his name and, visibly shaken, he awarded Hasbro a US Patent.

on the pendulum bob and to vividly imagine it swinging from side to side while recording apparatus captured the magnitude and direction of the oscillation. This research demonstrated that the magnitude of swing was enhanced when external rhythmic oscillating audio-visual stimuli (e.g. a light which flashes in time with the natural swing of the pendulum) that corresponded with the motion were present, compared to when they were absent (Easton & Shor, 1975). Furthermore, the pendulum swing was reduced when participants view of the pendulum bob was obstructed, compared to when it was visible (Easton & Shor, 1977).

Notwithstanding an influence of the occult, the above-reviewed research on ideomotor phenomena suggests that merely imagining what the movement of a distal object (i.e. the pendulum bob) might look and feel like is sufficient to generate the behaviour through arm, wrist and postural sway like a self-fulfilling prophecy. Furthermore, the fact that eliminating visual feedback of participants' own movements reduced the effect highlights the importance of perceptual feedback to reinforcing the mental image which gives rise to the movement. Just as Chevreul (1833) proposed, people do not intend to swing the pendulum, merely to imagine it, but by summoning this perceptual mental image they also inadvertently, and automatically, activate the associated motor behaviour to a small degree.

While the paradigm used by Eason and Shor (1975, 1976, 1977) provides important insights, it is not well suited to determine the precise contribution of internal and external processes to the generation of action slips because the movements of the pendulum appear over a protracted period of time (30 seconds). Indeed, if one wanted to test whether imagery and perceptual input, in the right circumstances, could conspire to *cause* action slips, then the earliest stage of physical action production – the moment of movement onset - is most relevant to this question. Perhaps a better way to address this question would be to design a procedure that could reliably elicit discrete non-willed

action slips (e.g. sudden movements) and to manipulate the circumstances which immediately preceded the moment of action initiation. This way, one might more easily isolate the factors which contribute to this behaviour. For example, if one were to vividly imagine a specific movement and then suddenly perceived a salient cue in the environment that corresponded with the same action, would one execute the movement?

1.2. The Perceptual Format of Action Intentions

At this point, it is useful to discuss in more detail how modern theories of motor control conceive of the process of action generation. As part of a renewed interest in the mechanisms of volitional action control, ideomotor principles have been refined and elaborated in a number of related theories, but there is an apparent consensus with respect to the idea that action and perception are inextricably linked. This idea, first proposed by classical ideomotor theorists such as Carpenter (1852) and James (1890), has been formalised in modern ideomotor accounts (e.g. Theory of Event Coding, Hommel, 2009; Hommel, Müsseler, Aschersleben, & Prinz, 2001), extended by computational modelling approaches to action cognition (e.g. Forward/Inverse Models, Wolpert, 1997; Wolpert & Ghahramani, 2000) and, most recently, recast in terms of Bayesian predictive processing (e.g. Active Inference, Adams, Shipp, & Friston, 2013; and Predictive Coding, Clark, 2013). These modern perspectives on motor control (henceforth referred to, rather crudely, as “ideomotor models”) overlap in three important respects. First, action control is fundamentally anticipatory (that is, predictive) in nature. Second, actions are initiated by anticipating (i.e. imagining, predicting or otherwise “activating”) their associated perceptual consequences. Third, this anticipatory mental image (or prediction) has a role in generating action (Shin et al., 2010). If these core assumptions are taken seriously, then forming an intention to act is

nothing more than forming a sufficiently strong anticipatory mental image of the effects of possible actions. In this view, action slips could occur because cues in the environment or one's mental imagery activates the mental representation of an unsuitable perceptual outcome - inadvertently driving a strongly associated, but inappropriate, behaviour.

Common Coding theory (Greenwald, 1970; Knoblich & Prinz, 2005; Prinz, 1990, 1997) provides a useful framework for understanding how and why action generation may be vulnerable to corruption by inappropriate imagery and perceptual input. In this view, the mechanisms of action, perception and mental imagery are functionally intertwined because action control relies on a set of internal representations coded in fundamentally perceptual terms, commonly referred to as "action effects". Importantly, action effects can be activated either endogenously (e.g. through imagery) or exogenously (e.g. through visual input). As noted above, actions are initiated when the corresponding action effect representation becomes sufficiently activated by endogenous and exogenous processes (Shin et al., 2010). Rather than necessarily representing a specific behaviour, these internal representations correspond with the distal perceptual effects that these actions cause in the environment (Hommel, 2009; van der Wel et al., 2013). For example, when one observes an action (or otherwise perceives action-related information, e.g. the sound of footsteps is closely associated with walking), the sensory input induces activity in the same internal representation (i.e. action effect) that must be activated endogenously to produce the corresponding behaviour. To illustrate, imagine two experienced pianists taking turns to play a sequence of alternating notes with the thumb (e.g. an E note) and middle finger (e.g. a C note). In the Common Coding view, the same action effects that are activated by endogenous processes to produce the alternating pattern are also activated by exogenous perceptual input when one hears the characteristic sound of the notes produced by the

partner (i.e. the distal perceptual consequences of each movement). In this way, perceptual input can be used as a retrieval cue for (or indeed activate directly) the corresponding motor commands and can guide the production of appropriate actions.

One important consequence of the notion that perception, imagery and action rely on a common language of perceptually coded representations (i.e. action effects), is that these processes should interact with one another and, sometimes, in a counterproductive manner – producing action slips. For example, if a given perceptual input was strongly associated with a specific action (e.g. a piano tone and the finger movement required to produce it), then the presentation of such a stimulus may suffice to trigger an unintended action slip in a reflexive, stimulus-driven fashion. In addition, if one encountered a salient stimulus which strongly activated an action effect that happened to be incompatible with a planned action (and its corresponding action effects), then one might expect that this would interfere with the smooth production of the intended behaviour.

These ideas were directly confirmed in a study which recruited skilled pianists, due to their considerable prior experience of linking their finger movements on a piano with the desired perceptual effects (i.e. the movements required to generate specific notes). The participants were asked to play one of four two-note sequences as implied by the visual musical notation (Drost, Rieger, Brass, Gunter, & Prinz, 2005). Crucially, task-irrelevant piano notes were presented simultaneously with the onset of the imperative stimulus which was either congruent or incongruent with the action demanded by the musical notation. The data showed that the pianists' responses were slower and less accurate when the auditory information (e.g. an E note), and the associated motor response (e.g. thumb press), was incongruent with the planned action (e.g. a middle finger press to produce a C note), than when it was congruent.

Importantly, the authors noted that participants were prone to occasionally executing the action which corresponded with the auditory stimulus (i.e. the piano note) rather than the imperative visual stimulus (i.e. the musical notation), as required by the task. The findings of Drost et al. (2005) suggest that perceiving incompatible action effects does indeed interfere with the production of intentional actions and, sometimes, results in the production of erroneous behaviour, similar to the action slips described previously.

An important limitation of the above study (and others like it) is that it does not allow one to conclude whether the incongruent auditory tone truly *initiated* the action slip (or “induced” as the authors claimed) because participants certainly *intended* to produce an action at stimulus onset (i.e. to play one of four musical notes), but they were led by the environment to respond erroneously. Yet, if action effects activated by endogenous (e.g. via top-down imagery) and exogenous means (e.g. via bottom-up perceptual input) can influence *how* we act, then the activation of action effects may also determine *whether* we act at all if they are activated strongly enough. Indeed, in a recent review of contemporary ideomotor theory (Shin et al., 2010, p. 943) it was explicitly stated that “Any activation of the effect image, either endogenously or exogenously, will trigger the corresponding action.” In other words, these views assume that activating an action effect does not only specify the form of a forthcoming action but also provides the necessary impetus for its initiation – effectively representing the intention to act itself. In this way, action slips may simply be a product of salient perceptual input combining with broadly compatible imagery. If imagery and perceptual input combined to inadvertently (and strongly) activate an inappropriate action effect, the corresponding response may be initiated in defiance of one’s goals.

Strikingly, research has largely ignored the intentional or initiating role of action effect activation. Instead, research has focussed primarily on the contribution of action

effects to selecting or shaping the execution of planned actions. Accordingly, there is a wealth of evidence to demonstrate that people produce faster and more accurate responses when people are primed, just prior to execution, with the “proximal” effects of planned actions (e.g. a hand posture, or movement of a body part, see Bach, Peatfield, & Tipper, 2007; Hommel, 1995; Brass, Bekkering, & Prinz, 2001). For example, Brass and colleagues (2001) asked participants to hold their right index finger 2cm off the table and to execute either a lift or a tapping movement at the onset of the imperative stimulus. Crucially, the imperative stimulus consisted of a visual depiction of a lifting or tapping index finger movement which was either compatible or incompatible with the planned response. The data revealed evidence of faster responses when participants encountered compatible, rather than incompatible, movements (sometimes referred to as the “Brass effect”).

In addition, similar compatibility effects have also been reported for participants who were presented, prior to execution, with abstract stimuli which, in an earlier phase of the experiment, were contingent upon to-be-executed action (e.g. “distal” action effects in the environment, such as a light turning on or a tone playing, Elsner & Hommel, 2001). Moreover, similar compatibility effects are even evident when the action’s effects are not directly perceived, but merely anticipated and because they happen to share perceptual features with the required response (e.g. response effect compatibility, Keller & Koch, 2006; Kunde, 2001; Kunde, Koch, & Hoffmann, 2004; for review, see Badets, Koch, & Philipp, 2016). For example, Kunde et al., (2004) demonstrated that people initiate responses of varying intensity (e.g. soft key press) faster and more accurately when they are contingently followed by auditory tones of corresponding intensity (e.g. a soft tone) relative to a tone of a different intensity (e.g. a hard tone). Importantly, a follow-up experiment (Experiment 2; Kunde et al., 2004) specifically tested whether response-effect compatibility affects the initiation of

prepared responses. In this task, participants were informed about the response they would have to execute, prior to the onset of the imperative stimulus and task-irrelevant auditory stimulus. The data revealed a strong influence of response-effect compatibility on response times, in spite of the fact that participants had already decided which response to execute and were simply waiting for a go-cue. This finding, in particular, provides suggestive evidence that action effects can even influence the very early stages of action generation. However, although this research suggests that activated action effects may contribute to the initiation of a planned action, it is not clear whether the activation of action effects *alone* can suffice to initiate the corresponding response (i.e. in the absence of a prior intention to act).

Further support for the role of action effect representations in action initiation is provided by neuroimaging studies which show that action planning involves activation in lower-level perceptual regions that represent the specific outcome one wants to produce (Kühn, Keizer, Rombouts, & Hommel, 2011; van Steenbergen et al., 2017; Zimmermann et al., 2017; Zimmermann, Verhagen, de Lange, & Toni, 2016). In one such study, participants were asked to prepare either manual or facial actions while they underwent functional magnetic resonance imaging (fMRI; Kühn et al., 2011). This study aimed to assess the extent to which action preparation was mediated by activity in cortical areas which are specifically recruited when processing the perceptual effects of each class of action. The data revealed that preparation of manual actions was associated with both heightened activity in the hand-related areas of the motor cortex and the extrastriate body area (EBA; a region known to process visual information about body parts). In contrast, preparation of facial actions selectively recruited face-related areas of the motor cortex and the fusiform face area (FFA; a region implicated in facial perception). Together, these results provide important evidence that preparing an action involves the anticipation of the associated perceptual consequences, as assumed

by ideomotor models. Yet, while the aforementioned studies provide converging evidence that action effects contribute to the selection of an action against alternatives or shape its kinematics (e.g. Bach, Griffiths, Weigelt, & Tipper, 2010) they do not test the central assumption of ideomotor models: that the activation of an action's perceptual effects is tantamount to forming an intention to act. That is, the activation of an action's effects by endogenous imagery or exogenous perceptual input does not only determine *which* action is selected or *how* it is executed but *whether* it is executed at all. The experiments presented in the present thesis aim to test this idea empirically.

1.3. The Interplay of Endogenous and Exogenous Processes In Action

In the previous section, evidence was presented which characterised the intention to act as a product of the activation of perceptual action effect representations (i.e. “mental images”). This activation, which drives motor behaviour, can be brought about by endogenous (e.g. imagery) and exogenous means (e.g. action observation). However, if action effects can be activated by top-down imagery and bottom-up perceptual input and are also involved in action initiation, can these sources of activation really be considered functionally equivalent or are they implicated in generating different types of action?

Intuitively, the notion that endogenous and exogenous processes may drive different kinds of action has parallels with a proposal made by James (1890) that purposeful, goal-directed actions are distinct from simple reflexive responses to environmental stimuli. This conceptualisation still resonates with modern ideomotor theories of action control which define human behaviour according to two interdependent and often overlapping purposes; (a) intention-driven actions which serve

to influence the environment as a result of internal desires and motivations, and; (b) stimulus-driven actions which serve accommodate environmental demands in light of exogenous events (Herwig, Prinz, & Waszak, 2007; Waszak, Cardoso-Leite, & Hughes, 2012). On this account, stimulus-driven actions are largely driven by exogenous processes and are assumed to be the product of simple stimulus-response associations acquired through experience. These stimulus-response associations serve to bind the motor representation of a specific action with the conditions (i.e. specific exogenous sensory input) under which it must be initiated (e.g. “red light – brake” when driving a car). In contrast, intention-driven actions are assumed to be largely the product of endogenous processes (e.g. action imagery) and mediated by action effect representations which bind actions in one’s motor repertoire to their typical perceptual consequences. In this view, intention-driven actions are initiated by strongly activating the action effects which correspond with the desired behaviour which, in turn, automatically evokes the corresponding motor program and initiates the desired act.

Support for this categorical distinction between different sources of action comes from neuroimaging research which has identified several distinct cortical circuits which are specifically recruited during the execution of intention-driven (i.e. “voluntary”) and stimulus-driven (i.e. reflexive) behaviour (Dum & Strick, 2002). The circuits that contribute to the generation of action converge on the Primary Motor Cortex (M1) which, in turn, transmits motor commands, via the spinal cord, towards the relevant effector – initiating action (for reviews, see Haggard, 2008, 2019). Broadly speaking, M1 receives input from two major cortical circuits implicated in action generation. One circuit is preferentially recruited during voluntary action preparation and involves the supplementary motor area (SMA) and the pre-supplementary motor area (preSMA) which, in turn, receives projections from the basal ganglia and prefrontal cortex (PFC; Picard & Strick, 1996). A second important circuit is specifically implicated in the

sensory guidance of actions which are predominantly stimulus-driven. Information from the early sensory cortices is projected to the parietal cortex and, further, to the pre-motor cortex before engaging M1 (Rizzolatti, Luppino, & Matelli, 1998). This circuit is implicated in the use of online sensory input to guide object-oriented actions. Haggard (2008) proposed that although both of these pathways contribute to the generation of volitional actions, they may implement fundamentally different decisions. For example, when the environment demands an immediate response (e.g. the appearance of an oncoming car when crossing the road) the parietal circuit may be involved in deciding between action alternatives based on environmental constraints (e.g. jumping right, towards the pavement, or left, towards oncoming traffic). In contrast, less immediate and more deliberative actions (e.g. whether to visit a shop on the other side of the road to buy some milk) may be largely planned and initiated by the frontal circuit. However, these circuits do not operate independently and their apparent functional specialism should not be taken to imply that a given action will either be entirely endogenously formed (via the frontal route) or exogenously triggered (via the parietal route). In other words, although this research implies the existence of separate circuits that are responsible for producing different kinds of action, they may only appear distinct because they receive different types of input (i.e. signals arising from endogenous or exogenous sources). It is possible, therefore, that endogenous and exogenous information (processed in separate circuits) is integrated – via a common set of action effect representations – prior to engaging the motor system and initiating action.

Indeed, it is important to note that although the two categories of action (and the cortical circuits recruited in each case) are often defined in opposition to one another, this is partly due to convenience. A classical reflex, such as a highly stereotyped withdrawal response in reaction to a painful stimulus, stands as a convenient example of a stimulus-driven act in its most basic form. In contrast, it is harder to define what

constitutes a purely intention-driven action that emerges as the product of some internal deliberative process (Brass & Haggard, 2008). Instead, actions are said to be intention-driven if they exhibit a degree of stimulus-independence (i.e. they are not immediately triggered by an external stimulus). As Haggard (2008) noted, these categories of action should not necessarily be taken to represent a dichotomy, but rather a continuum with degrees of intentional oversight. Of course, the notion of a continuum in human behavioural control demands intermediate examples which share features of both intentional, internally generated actions and reflexive responses. The action slips described at the start of this chapter provide just such an example of behaviour which appears to be neither wholly intention- nor stimulus-driven. Indeed, the unusual subjective feeling associated with producing action slips may arise because one has internally generated an appropriate action goal which is then diverted towards an inappropriate outcome by cues in the environment - conflicting with the initial goal.

To illustrate, consider the tea making error described earlier in which one absentmindedly pours a mugful of cold milk, rather than hot water onto the teabag. Endogenously forming an action goal (e.g. make a cup of tea) may have led to the activation of a set of action effect codes which describe potential actions, necessary for making tea, in an intention-driven (i.e. internally generated) manner. This suggestion is similar to Hommel's (2009) concept of "event files" for voluntary action which comprise sets of broadly relevant action effects bound to specific contexts in which they are typically required in the Theory of Event Coding. In this way, the endogenous retrieval of the relevant action effects may allow the environment to trigger the appropriate reaching and grasping actions. However, in this case, attending to the milk jug (and the actions it affords) inadvertently drives the activation of an appropriate action effect (e.g. a reach and grasp towards a tea-relevant object with a handle), but at an inappropriate moment (e.g. pouring milk into the mug before the hot water). The

conceptualisation inherently, and deliberately, proposes that the moment of action initiation is simply a product of the activity of a corresponding action effect representation. In this way, action effects which become sufficiently activated by endogenous and exogenous means, do not only define the form of a forthcoming action but may be able to directly trigger its initiation. This proposal has not yet been tested directly and motivates all the experiments which follow.

It is important to note that this idea clashes with the received wisdom that action generation involves a number of distinct computational steps which might preclude the ideomotor interpretation of the role of action effects in action initiation. In classical stage theory, the process of voluntary action generation is conceived of as comprising a number of distinct stages or serial set of decisions, such as action selection, initiation and execution (Sanders, 1980). More recent research has retained these labels, principally for pragmatic reasons, but emphasises that each “stage” is assumed to reflect a different time point in a unitary process of action generation (Hommel, 1997). Typically, the action selection stage reflects the earliest stages of response production in which the intended act is chosen; action initiation refers to the phase immediately preceding the production of the overt response and action execution refers to the actual motor output (Kunde et al., 2004). Nevertheless, ideomotor models emphasise the contribution that the dynamic interaction of internal processing and external stimulation makes to activating action effects and initiating actions – an idea which appears to be at odds with the notion of discrete stages in action generation.

The dynamic nature of action control processes is, perhaps, more readily accommodated by more recent accounts of action generation which cast the process of action initiation in terms of an integration-to-bounds mechanism (Churchland, Kiani, & Shadlen, 2008; Murakami, Vicente, Costa, & Mainen, 2014; Schurger, Mylopoulos, & Rosenthal, 2016; Schurger, Sitt, & Dehaene, 2012; Schurger & Uithol, 2015). These

accounts describe a process in which evidence is accumulated in favour of action alternatives until a motor threshold is reached and the corresponding action is initiated. This research was inspired by the pioneering work of Benjamin Libet and colleagues, who asked participants to make spontaneous button presses while undergoing electroencephalography (EEG) and to report the moment at which they had intended to move (Libet, Gleason, Wright, & Pearl, 1983). This research indicated that a build-up of cortical signals (referred to as the readiness potential) reliably preceded participant's conscious awareness of their intention to act by some 300ms. This finding created incredible controversy because it was taken as evidence that the preparation of a supposedly voluntary action preceded one's conscious awareness of the intention to move – implying the feeling of intentionality is an effect, rather than a cause of voluntary action. However, more recent work has cast doubt on this interpretation and suggests that a significant proportion of the readiness potential may be attributable to the natural ebb and flow of background neuronal noise, rather than the outcome of some internal pre-conscious decision to initiate movement. For example, Schurger, Sitt, & Dehaene (2012) repeated Libet's task, with the additional instruction that the task would occasionally be interrupted by an auditory "click", at which point participants were required to press the button as soon as possible. The data revealed that the shortest response times followed instances in which background neural activity was arbitrarily high. This suggests that the moment of action initiation can be considered a threshold-crossing neural event following a period of "evidence accumulation" in favour of candidate action driven by endogenous and exogenous processes. Crucially, if the assumptions of ideomotor models are taken seriously, then the evidence accumulated by such mechanisms is nothing more than the activation of multimodal action effect codes: how vividly the desired action outcome is represented. In this view, an imagined action becomes overt behaviour if the perceptually-coded action effects become sufficiently

active and surpass the motor threshold. That is, the precise nature and timing of the actions people produce may be determined by nothing more than the moment a strongly activated action effect representation crosses the threshold for execution.

Supportive evidence in favour of the equivalence of action effect anticipation and action initiation comes from a body of research that demonstrates the considerable neural overlap between action imagery and action execution. Neuroimaging studies suggest that action imagery and action execution give rise to similar patterns of brain activity, even if the activity of these neuronal ensembles is significantly less intense during imagery (Kraeutner, Gionfriddo, Bardouille, & Boe, 2014). Similarly, it is well established that action imagery is associated with increased muscle activity (as recorded by electromyogram, EMG) in muscles specific to the imagined movement relative to baseline (Jacobson, 1930). Moreover, the magnitude and rhythm of the EMG activity have been shown to be proportional to the imagined effort and timing of the moment (Shaw, 1940; as cited in Jeannerod & Decety, 1995). Taken together, these findings suggest that the neural processes recruited when one engages in action imagery (e.g. by evoking action effect representations) are intimately associated with generating actions, just as one would expect if action effect anticipation was sufficient to drive and initiate actions. It is important to note however that none of these studies shows that fully formed actions can actually be *initiated* through imagery. Instead, action imagery typically results in a subliminal motoric activation of the corresponding muscles, which is only detectable if amplified by stimulation via TMS (e.g. Vargas et al., 2004).

One might argue that the classical ideomotor phenomena, such as Chevreul's (1833) magic swinging pendulum also provides support for the role of action effects in action initiation. In this case, participants experience movements that they did not consciously intend to produce (and therefore attribute to supernatural causes), but which nevertheless reflect their perceptual expectations of what will happen (Chevreul, 1833;

Easton & Shor, 1975, 1976, 1977). However, such effects are typically small and fall far below the threshold for formed responses, even if amplified by these “magic” devices. Indeed, other researchers have argued that motor imagery mainly affects the planning of the action, specifying its form prior to execution (Bach, Allami, Tucker, & Ellis, 2014; Caldara et al., 2004) and that different processes energise and release the action (Tecuapetla, Jin, Lima, & Costa, 2016). Others have argued that the effects of ideomotor cues (e.g. movement observation, Brass et al., 2001) on action execution reflect interference rather than facilitation effects (Ramsey, Cumming, Eastough, & Edwards, 2010) and would therefore – per definition – only be able to prevent, not cause, action. In light of this open question, the experiments in the present thesis aim to test the contribution of action effect representations, activated by imagery and perceptual input, to the *initiation* of non-willed action slips. In other words, the forthcoming experiments aim to test whether the combined activation of an action’s effects through imagery and external stimulation can become strong enough to elicit action when none was intended, effectively inserting an intention to act.

1.4. Overview of the Thesis

This thesis began by asking the question, how is it that people can produce apparently “non-willed” action? Ideomotor theories of action control claim that imagery and perception are tightly linked processes, mediated by activity in perceptually coded action effect representations. Importantly, there is converging evidence that implicates the activation of action effects in the initiation of actions. While there is a wealth of evidence to demonstrate a role for action effects in the selection and guidance of actions, no research has specifically tested whether action effect representations can

serve a truly generative purpose – creating spontaneous action *de novo* - regardless of one's broader goals. The experiments in this thesis specifically test this proposal. Specifically, these experiments aim to test the hypothesis that a mental image (i.e. action effect image) can initiate the corresponding action when it becomes sufficiently activated by a combination of top-down imagery and bottom-up perceptual input. Importantly, the mental image should be translated into overt movement, even if one had not previously intended to execute the action.

To that end, a novel experimental paradigm was designed in an effort to induce non-willed behaviour, or “action slips” (Chapter 2, Experiments 1-3). The predictions regarding the form that these action slips should take and the frequency with which they should occur are motivated by ideomotor models and related theories (Hommel, 2009; Hommel et al., 2001; James, 1890; Kunde et al., 2007; Pacherie & Haggard, 2010; Prinz, 1997; Shin et al., 2010). Participants rehearsed a finger tapping sequence using mental imagery in time with an auditory beat but were asked to withhold responding (Chapter 2-4). A pair of virtual hands was presented onscreen from an egocentric perspective which corresponded with the position and orientation of the participants' hands. Ostensibly, the virtual hands were used to indicate the correct finger tapping sequence at the start of each trial. However, crucially, the index fingers of the virtual hands occasionally made highly salient (but task-irrelevant) movements that were either congruent or incongruent with the participant's imagined movement (e.g. similar to the observed finger movements used by Brass et al., 2001). If (exogenous) visual stimulation and (endogenous) imagery combine to activate the same action effect representation, then congruent stimulation should be, in some trials, enough to induce non-willed finger tap (i.e. an action slip) in spite of a specific instruction to remain still.

After gaining proof of concept for this novel design (Chapter 2, Experiments 1-3) the procedure was redesigned to allow a more fine-grained analysis and to test more

detailed predictions about the nature of the mechanism which gives rise to action slips (Chapter 3, Experiment 4). Specifically, the design used in Experiments 3 was adapted to record continuous pressure data in Experiment 4. This made it possible to capture extremely small movements and to rule out an alternative explanation (simple motoric rehearsal) for the findings of Chapter 2 (Experiments 1-3). Further, Experiment 4 added a neutral baseline condition which made it possible to determine whether perceptual input (i.e. the salient movement cues), when combined with congruent mental imagery, primarily facilitates or inhibits the generation of action slips.

Chapter 4 (Experiments 5 and 6) applied a version of the task used in Chapter 3 (Experiment 4) which was modified to incorporate a body ownership illusion. Some participants were assigned to a heightened ownership group, wherein their left/right finger movements elicited corresponding movements from the virtual (onscreen) hands at the start of each trial. This design was intended to increase the extent to which participants in the heightened group mapped their own left/right index finger movements onto the movements of the onscreen fingers (i.e. to increase “ownership” of the virtual hands). Visual feedback of the participant’s own hands was also manipulated in Chapter 4. These experiments aimed to test whether ownership and the presence/absence of visual feedback for participant’s own movements would affect the tendency for observed movements to activate the motor system.

Chapter 5 (Experiment 7) tested the relationship between individual differences in responsiveness to a hypnotic screening procedure and the participant's prior performances in Chapter 4 (Experiments 5 and 6). That is, Experiment 7 aimed to test whether one’s propensity to produce action slips is related to individual differences in ideomotor suggestibility – the motor component of hypnotic suggestibility (Woody, Barnier, & McConkey, 2005).

Finally, Chapter 6 (General Discussion) summarises the results of Chapters 2-5 (Experiments 1-7) and integrates them into a model of intentional action. Together, the findings suggest that action slips are indeed a curious by-product of the mechanisms by which voluntary actions are generated. In addition, the results provide direct, behavioural evidence for some of the most fundamental features of ideomotor models of action control. These findings demonstrate the utility of using non-willed action slips as a device to test esoteric and often highly abstracted claims about the mechanisms of voluntary action.

2. Chapter 2: The Relationship Between Action Intentions and Internal Representations in Volitional Action Control

2.1. Overview

This chapter is based on Colton, Bach, Whalley and Mitchell (2018a), “ Intention insertion: activating an action’s perceptual consequences is sufficient to induce non-willed motor behaviour”, published in *Journal of Experimental Psychology: General*.

As part of this publication, all data, materials and an R script which reproduces the analyses and plots in Experiment 3 (i.e. the main experiment in the published article) were made available online with the Open Science Framework (Colton, Bach, Whalley, & Mitchell, 2018b). Each experiment in the forthcoming chapter uses this same analytical approach.

2.1.1. Introduction

The experiments in Chapter 2 were designed to test, for the first time, the assumption of ideomotor models that action intentions are essentially action effect representations that are activated strongly enough. As reviewed in the introductory chapter (Chapter 1), there is ample evidence that activated action effects shape the form and the timing of responses. For example, people produce faster and more accurate finger movements if they are primed, just prior to execution, with a visual depiction of a corresponding, than non-corresponding action (Brass et al., 2001; Brass, Bekkering, Wohlschläger, & Prinz, 2000). Such studies have provided direct evidence that actions are a manifestation of the same “common codes” (Knoblich & Prinz, 2005; Prinz, 1990, 1997) that also guide imagery and perception of stimuli of the environment. However, a

central assumption of ideomotor models – the activation of this action effect image is sufficient to initiate the corresponding action (i.e. in the absence of a prior intention to act) – has remained untested.

The studies described in Chapter 2 tested whether it was possible to elicit non-willed action slips by presenting participants with highly salient, unexpected action cues (i.e. a sudden observed movement) during a mental imagery task – even if they were instructed to actively withhold all movement – and essentially “insert” an intention to act into the participant’s motor system. This idea relies on the assumption of ideomotor models that perception, action, and imagery operate on the same or commensurable representations which have a fundamentally perceptual format (i.e. action effects) and can therefore summate (Berends, Wolkorte, Ijzerman, & Van Putten, 2013; Hommel et al., 2001; Kunde et al., 2004; Prinz, 1997; Vogt et al., 2013; Waszak et al., 2012). Consider a task in which participants are instructed to mentally rehearse – but not execute – a simple motor action, such as a left or right finger movement. Under the assumption of ideomotor models, presenting participants with the perceptual consequences of their imagined actions, both newly learned (e.g. task-irrelevant tones; Elsner & Hommel, 2001) and those acquired during previous action experience (e.g. an observed finger movement; Brass, Bekkering, & Prinz, 2001), should exert an additional bottom-up effect on activity already induced in a top-down manner by motor imagery. If the process of generating a strong mental image of a bodily movement can be considered equivalent with forming an action intention, then the combined influence of imagery and perception on the activity of shared action effect codes should occasionally pass the threshold required for execution, triggering responses even if none were intended (i.e. effectively inserting an intention to act). Importantly, the likelihood that presenting movement cues during imagery should elicit a non-willed act should emerge as a function of the dimensional overlap (i.e. shared perceptual features)

between the action effect codes activated by the observed and imagined action (Kunde et al., 2004). For example, perceiving a left index finger movement should be more likely to elicit action slips if the participants simultaneously imagined a left, rather than right finger movement. Therefore, non-willed action slips should be more likely to follow instances when the observed and imagined actions match than when they do not.

At this stage, it is important to emphasise how non-willed behaviour was operationalised in the following studies and how “action slips” in the current context can be distinguished from the kinds of errors of action control that have been widely reported in the literature. Classical stimulus-response-compatibility paradigms employed in ideomotor research (e.g. Brass et al., 2001) often require participants to execute a response at the onset of an imperative stimulus. Participants occasionally respond in error (i.e. they produce an alternative, incorrect response) when faced with a stimulus which is incompatible with their planned action. However, these error responses cannot be considered *truly* non-willed since participants certainly intended to execute a response from a small set of alternatives but were led by an experimental manipulation to select an incorrect action. This kind of manipulation, therefore, merely induces an error in action selection, not in action initiation per se. If participants had a broad goal of executing finger movements in response to, for example, any observed movement then such a prospective intention to act may be responsible, in some sense, for initiating the subsequent, if erroneous, action (Brass & Haggard, 2008). Instead, in Experiment 1, participants took part in an imagery task and were asked to withhold any physical responses until a pre-specified imperative stimulus appeared. As such, any responses that are collected prior to the true imperative stimulus can be considered antagonistic with the explicitly stated task goal of refraining from movement during imagery.

2.2. Experiment 1: Eliciting Non-Willed Action Slips During Imagery With Newly Learned and Previously Established Action Effect Associations

In his attempt to develop a taxonomy of accidental behaviour, Reason (1984) noted that slips of action control (hereafter, “action slips”) are not well suited to laboratory study, given that they are naturally, and necessarily, rare. In light of this, Experiment 1 was designed with the aim of providing the maximal conditions for an action slip to emerge in laboratory conditions. While the taxonomic labels derived by Reason (1984) were not intended to be exhaustive or mutually exclusive, Experiment 1 specifically intended to elicit so-called “intrusive” action slips, wherein novel unintended actions become incorporated into one’s ongoing behaviour. That is, Experiment 1 tested whether it was possible to effectively *insert* an intention to act (i.e. an action slip) by presenting intrusive movement cues during an imagery task.

Experiment 1 consisted of two parts. The first part followed the logic of Elsner and Hommel (2001), in which stimuli (e.g. tones) that reliably follow a motor response (e.g. a finger movement) become integrated with the action effect representation of the action, and can then – after training – prime the corresponding behaviour. In the first part of Experiment 1 (hereafter, the “association training task”) was similarly designed to forge associations between abstract stimuli (e.g. colours and tones) and specific finger movements, which could then be used to elicit non-willed action slips in the subsequent imagery task. In each trial, participants responded to specific coloured visual stimuli with either left or right index finger presses according to a fixed mapping (e.g. yellow – left index, blue – right index). When participants produced the correct response, specific auditory tones that corresponded with the participants’ keypresses

were played via headphones (e.g. a left index press produced a “buzz” sound and a right index produced a “chirp”). Hence, the association training task was designed to encourage the formation of novel action effect associations between the motor representations of the executed responses (e.g. left/right index finger movements) and the sensory representations of the imperative visual (e.g. colours) stimuli and the auditory stimuli (e.g. tones) that reliably followed specific actions: the action’s consequences. According to ideomotor models (e.g. Elsner & Hommel, 2001), when participants encounter the colours and tones in the subsequent imagery task, the corresponding motor representation should also be activated as a result of these newly formed associations.

In the second part of Experiment 1 (hereafter, “intention insertion task”), participants rehearsed sequences of left and right index finger taps by imagining what the movements would look and feel like if they were physically executing them in time with an auditory metronome beat. In one third of trials, participants encountered a surprising audio-visual action cue, which either matched the currently imagined action in the sequence (e.g. seeing the same finger depressed, along with the associated colour and tone) or mismatched the imagined action (e.g. a different finger being depressed, with the associated colour and tone). Participants were instructed to ignore these images and withhold any responses until a pre-specified imperative stimulus appeared. However, research suggests that activity induced by imagining a specific movement and suddenly encountering a visual depiction of a matching action (along with the strongly associated auditory and visual action effects) should nevertheless summate (see Kunde et al., 2004) and, sometimes, trigger involuntary action, when their activation becomes super-threshold. Instances in which participants respond accidentally following the onset of a sudden action cue (e.g. an observed movement and associated colour and tone) are referred to as action slips (mirroring the language used by Reason, 1984).

Crucially, manipulating the compatibility of imagined and perceived movements allows involuntary movements that may result from the premature release of pre-planned actions due to a startle response (P. Brown et al., 1991; Maslovat, Chua, & Hodges, 2013) to be distinguished from those that genuinely result from the summation of action effect codes. If the involuntary movements produced by participants are merely caused by the surprise of encountering an unexpected stimulus, then the likelihood of producing an action slip will not be affected by the compatibility between imagined and observed movements. Thus, when imagining a left finger key press, the likelihood of an incidental execution of this action should be just as high if participants are exposed the sensory cues associated with this left key press and a right key press.

In contrast, under ideomotor assumptions, action slips should be more likely when the imagined action and observed movements are congruent (such that the action effect codes overlap and can summate) than when they are incongruent (when such an addition is not possible), as predicted by ideomotor models. Furthermore, if this is the case, then action slips should primarily reflect the imagined action (e.g. a left index movement) and not the alternative response (e.g. a right finger movement) when the imagined and observed actions are congruent (e.g. simultaneously imagining and observing left index finger movements). In contrast, if participants imagine, for example, a left movement, then they should be more likely to produce the alternative, non-imagined action slip (i.e. a right response) when they observe a right (i.e. incongruent), rather than left (i.e. congruent), index finger movement on the screen.

2.2.1. Method

2.2.1.1. Participants

A total of 43 participants were tested, of whom 18 were male and 25 female (6 Left-handed; Age in years: $M=27.3$, $SD=12.5$). The sample consisted of both undergraduate psychology students from the University of Plymouth (course requirement) and members of the university's paid participation pool (remunerated £4).

2.2.1.2. Materials and Equipment

The experiment was conducted in a dark, soundproof room. Participants were asked to wear over-ear headphones for the presentation of auditory stimuli and then seated in front of a 19" LED computer monitor (Resolution: 1900x1200; Refresh rate: 60Hz). The experiment was designed using E-Prime Professional (SP2.0) and responses were collected with a standard computer keyboard. Participants were asked to place their left and right index fingers over the 'W' and 'O' keys respectively for the duration of the experiment. Auditory stimuli consisted of three tones; a 200ms 440Hz tone which served as our metronome beat, along with a 200ms "chirp" and 200ms "buzz" (i.e. trained action effects) created with the tone generator in Audacity (panned hard left/right according to the laterality of the associated response).

In the association training task, yellow and blue full-screen colours signalled the response participants needed to execute. In the intention insertion task, additional visual stimuli consisted of an image of two virtual egocentric hands with artificial 'shadows' as if viewed from above (each with a visual angle of approximately 16 ° vertically and 14° horizontally). The virtual hands were either presented with the fingers in a neutral position, left index finger depressed or right index finger depressed. In the latter two

cases, the shadow was made to converge with the tip of the index finger in each case to give an illusion of apparent downward and upwards motion (i.e. a keypress) when displayed in quick succession with the typical hand stimuli.

2.2.1.3. Procedure: Association Training Task

The association training task consisted of 150 trials in which participants were obliged to execute either a left or right index finger keypress according to the colour of the imperative stimulus (Figure 1). In each trial, a black fixation cross was presented on a white screen for 500ms followed by a blank screen for 100ms and then a visual imperative stimulus. The stimulus consisted of a full-screen colour change which appeared for 400ms followed by a blank screen of 600ms, creating an overall response window of 1000ms.




Figure 1. A schematic representation of the association training task of Experiment 1. Participants were presented with one of two colours (Blue/Yellow) and obliged to respond with a Left/Right finger keypress. When the participant registered their response, a lateralised (i.e. panned hard left or right) “buzz” or “chirp” tone was presented to the corresponding ear via headphones. After 150 trials, the participant will have learned to associate colour-tone pairs with specific movements.

Half of the participants were given instructions that presentation of a yellow stimulus required a right key press, whereas a blue screen required a left key press and vice versa for the second half of participants. In addition, a response to a yellow stimulus always produced a “buzz” tone and a blue stimulus always produced a “chirp”

tone. These colour-tone pairings were identical for all participants. Feedback about response accuracy and time was then displayed for 500ms after each trial. An inter-trial interval of 500ms separated the feedback from the subsequent trial.

2.2.1.4. Procedure: Intention Insertion Task

The intention insertion task consisted of four blocks of 24 trials. In each trial, participants were first instructed about the particular rhythm they had to imagine. An image of a pair of virtual hands was presented, from an egocentric perspective, and participants heard a slow, steady metronome beat (440Hz tone presented at 800ms interval). An arrow moved, in time with the metronome beat, between the left or right index fingers, indicating a simple four-beat finger tapping sequence (Figure 2). For any given trial, participants could encounter one of six different finger-tapping patterns which consisted of all possible permutations of a four-beat sequence involving two left and two right index finger movements. Participants were asked to imagine tapping along with this sequence and memorise it in this manner.

After four beats the arrow disappeared, but the metronome continued playing. Participants were instructed to continue to mentally rehearse the finger tapping sequence in time with the metronome while refraining from overt movement. The virtual hands remained onscreen in neutral position along with a central fixation cross that flashed in time with the metronome beat, providing an audio-visual representation of tempo, and drawing participants' attention to a point equidistant between the hands. After a variable duration of 5-8 beats, the word "Go!" appeared for 2000ms and participants were asked to perform the keypress they happened to be imagining at that point (accurate responses indicate task compliance). Participants were explicitly instructed to refrain from overt movement before this imperative stimulus appeared. It

was emphasized to even refrain from moving when surprising visual stimuli would appear during the imagination phase.




Figure 2. A schematic representation of a 5-beat compatible cued trial from Experiment 1. The depicted sequence is Left – Right – Left – Right – indicated by the arrows shown in the learning phase. In the imagination phase, a Right index finger cue (screen colour, deviant tone and observed finger movement) falls on beat 4. The cue was presented when the participant imagined a Right response and so this was compatible with the concurrent imagined response. The presence or absence of an action slip (keypress) in the 800ms inter-beat period following cue onset was the outcome variable. Participants were told to focus on their imagery and refrain from overt movement until presentation of the true imperative stimulus, “Go!”

The crucial manipulation was that in eight of the 24 trials in each block, unexpected action cues appeared at the onset of one of the metronome beats (on either beat 6, 7 or 8). The action cues consisted of images of the left or right index fingers moving downwards and back to a neutral position, mirroring what participants would see if they actually carried out a finger tap themselves. The finger movement was presented in tandem with the colour-tone pair that had been previously associated with that specific action (i.e. the screen colour that preceded the action, and the lateralised tone that followed the action in the association training task). The action cues were either compatible or incompatible with the finger movement that participants concurrently imagined using (if the participant correctly followed the tapping sequence). Within these 32 cued trials, all combinations of cue compatibility (compatible/incompatible), cued finger (left/right), trial length (5, 6, 7 and 8 beats) and cue placement in the sequence (1 or 2 beats prior to sequence end) occurred once and were therefore fully counterbalanced.

Prior to the main task, participants completed a short practice of eight trials in which feedback was provided and incorrect responses were repeated. Participants were given the option of repeating the practice if required. The experimenter monitored participants' performance in this practice session. He emphasized that participants' ability to imagine – not execute – the movements was being tested. He instructed them not to move their fingers in the imagery interval and intervened if this instruction was not followed. He also unobtrusively remained in the room during the experiment proper in order to intervene if active finger movements were detected.

2.2.2. Results

To assess whether participants engaged appropriately with the intention insertion task, the percentage of correct responses to the true imperative stimulus (i.e. when "Go!" appeared and participants executed the currently imagined action) during trials which did not contain visual action cues was scrutinised for each participant. The subsequent analysis depends on the assumption that participants imagined the correct movements during each trial in order to accurately determine whether the action cues were congruent or incongruent with the participants' imagery. This assumption is less likely to hold for those participants who scored poorly, especially if they were responding with near-chance accuracy (i.e. ~50%) which would indicate very low engagement with the task. As a result, conservative criteria were set and only those participants with accuracy rates within one standard deviation of the group mean ($M=83.9\%$; $SD=15.3\%$; $M-SD=68.6\%$) were considered, leading to the removal of 8 participants' data (leaving $N = 35$ participants). Exclusion of these participants did not affect the results but did make it safer to assume that the remaining participants engaged properly with the task.

First, to check whether participants followed the instruction to refrain from actively moving their fingers in the imagery intervals, the number of imagined and non-imagined action slips on non-cued trials were compared (i.e. trials in which no action effect cues were presented). If participants pressed the keys during imagery, then one would expect to see a higher proportion of keypresses which reflected the imagined movement, relative to the non-imagined movement. In the absence of visual stimuli, accidental button presses were exceedingly rare and occurred on only .3% of beat intervals. These non-cued slips were roughly equally distributed between button presses with the imagined and non-imagined finger, providing little evidence that participants were consistently using finger movements to rehearse the imagined pattern (66 imagined vs. 47 non-imagined slips across all participants, $\chi^2 = 3.19$, $p = .074$).

The data were analysed using generalised linear mixed-effects regression, GLMER, using the `glmer()` function of the `lme4` package (Bates, Mächler, Bolker, & Walker, 2015) in R Version 3.4.2 (R Core team, 2017). This approach was taken because action slips which followed action cues were rare, meaning that the resulting data were zero-inflated (i.e. there were a lot of observations in which no responses were registered) and were not normally distributed. Therefore, the data were not well suited to more traditional statistical methods (e.g. analysis of variance) and could be better accommodated by specifying a suitable statistical model using GLMER. When modelling binary responses with GLMER, null hypothesis significance testing is implemented by comparing a full model and a reduced model without the fixed effect of interest with a likelihood ratio test (hereafter, LRT). The resulting chi-square test statistic is analogous to the F-statistic in conventional models. The estimates given by the GLMER model are reported along with a bootstrapped 95% confidence interval based on 1000 simulations. All categorical predictors were effect coded and when referring to model parameters the first letter is capitalised for ease of comprehension.

The analysis proceeded in three steps. The first step of the analysis tested whether the surprising visual stimuli induced non-willed key presses (i.e. action slips) in the imagery phase. To do this, each 800ms interval between metronome beats (or, in the model, Beat) was treated as a discrete event in which an action slip could have occurred or not (coded *one* or *zero* respectively). The model included fixed effects of Cue Presence (2 levels: Yes/No) and Trial (1-96, logged and mean centred) and a random intercept for each Participant and each Beat. In addition, the effect of Cue Presence, Trial and their interaction were allowed to vary within Participant by including random slopes for these terms. Such an approach is roughly analogous to a repeated measures one-way ANOVA with the factor Cue Presence.

As predicted, the presence of unexpected cues significantly increased the probability of action slips, LRT: $\chi^2(1) = 31.17, p < .001$; GLMER estimate: $\beta = 2.98$, 95% CI = [2.22, 3.58]. Predicted probabilities derived from repeated simulation of the model reveal that there was a 9.23% [4.06, 19.2] chance of an action slip following the onset of a cue, but only .27% [.12, .59] when no cue was present. In addition, there was a significant main effect of Trial, indicating participants became less likely to accidentally press a key as the experiment progressed, LRT: $\chi^2(1) = 38.33, p < .001$; GLMER estimate: $\beta = -.31$, 95% CI = [-.55, .03]. Finally, Cue Presence and Trial interacted, LRT: $\chi^2(1) = 17.75, p < .001$; GLMER estimate: $\beta = -1.12$, 95% CI = [-1.69, -.67], indicating that action cues became less effective at eliciting action slips over the course of the experiment.




Figure 3. A box plot of the observed frequency of general action slips according to the class of action cue which preceded them in Experiment 1. The observed frequencies are broken down by Perception-Imagery Compatibility with a thick black bar to represent the median, a box showing the interquartile range and whiskers to indicate maxima and minima.

The second step tested whether action slips would generally be elicited by cues that match the currently imagined finger keypress, not by those that mismatched. Visual inspection of the observed frequencies of action slips recorded in response to each class of action cue (Figure 3) suggests that action slips were more likely when the observed finger press and the imagined finger press were identical (i.e. congruent, $M = 20.2\%$) compared to when they were different (i.e. incongruent, $M = 15.7\%$). A smaller dataset was constructed which only consisted of observations which followed the onset of surprising action cues (i.e. each participant contributed 32 observations), and dummy variables encoded whether the cue was compatible or incompatible with the currently imagined response. The model included fixed effects of Perception-Imagery Compatibility (2 levels: Compatible/Incompatible) and Cue Order within the trial series (ranked 1st to 32nd, log-transformed and mean centred) and a random intercept for each Participant. In addition, random slopes for Perception-Imagery Compatibility, Cue

Order and their interaction, were added which allowed each participant to vary in their general response bias.

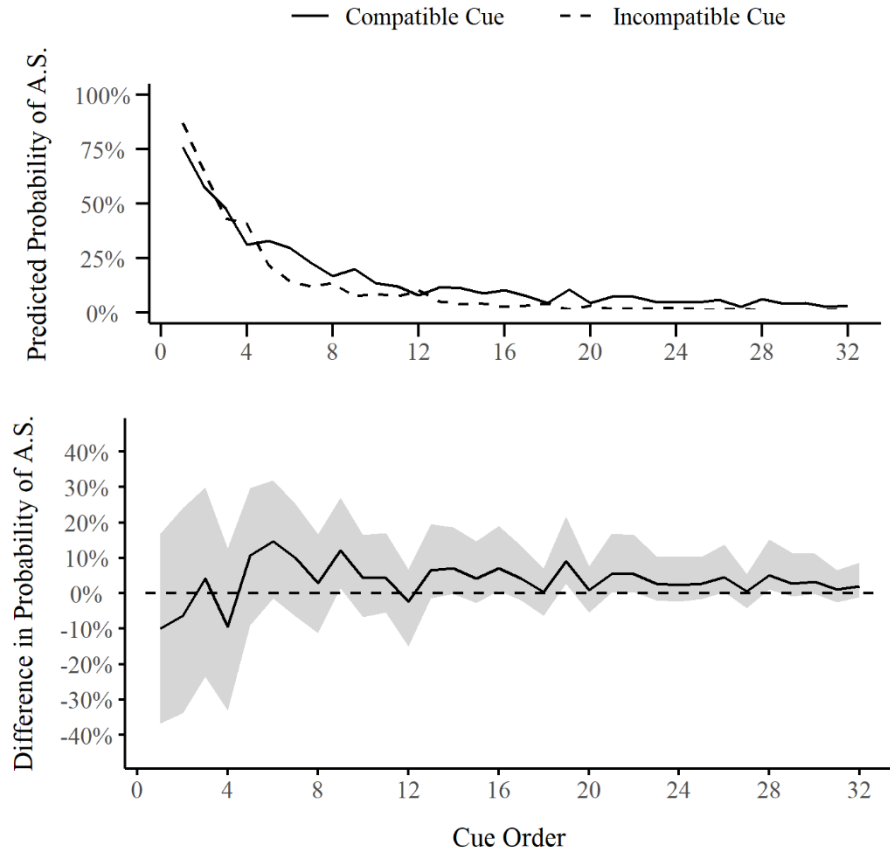


Figure 4. A combination plot indicating the predicted probability of general action slips according to the class of action cue which preceded them across Experiment 1. The upper panel shows a plot of the predicted probabilities, derived by simulation, that cue onset would elicit a general action slip (i.e. any response) according to Perception-Imagery Compatibility and Cue Order. The lower panel shows a difference plot (i.e. Compatible - Incompatible) in which the shaded area shows a 95% confidence interval derived by simulation.

The analysis revealed a significant main effect of Cue Order, $LRT: \chi^2(1) = 35.74$, $p < .001$; GLMER estimate: $\beta = -1.65$, 95% CI = [-2.19, -1.26], reflecting the fact that action slips were most likely in the early stages of the experiment and became less frequent as it progressed (Figure 4). Importantly, as predicted, there was a main effect of Perception-Imagery Compatibility, $LRT: \chi^2(1) = 4.49$, $p = .034$; GLMER estimate: $\beta = .4$, 95% CI = [.05, .86]. This indicates that action slips were significantly more likely

when participants encountered action cues that were compatible (e.g. right index), rather than incompatible (e.g. left index), with their currently imagined movement (e.g. right index). Finally, tests revealed no evidence of an interaction between Perception-Imagery Compatibility and Cue Order (LRT: $\chi^2(1) = 2.54, p = .11$; GLMER estimate: $\beta = .25$, 95% CI = [-.1, .64]).




Figure 5. A plot of the predicted probabilities of specific action slips, derived by simulation, according to the class of action cue which preceded them in Experiment 1. A stacked bar plot indicates the predicted proportion of specific action slips (i.e. either the imagined action or alternative action) made according to Perception-Imagery Compatibility.

Finally, the third step of the analysis tested the hypothesis that the specific action slips captured by Experiment 1 should primarily reflect the currently imagined key press. In addition, it was predicted that participants should execute the alternative action at a higher rate in response to incompatible cues than compatible cues. A new dataset was created which consisted of only observations in which an action slip occurred. In addition, a new outcome variable was added to indicate whether the participant had executed the imagined or alternative action (coded *one* and *zero* respectively). The model included fixed effects of Perception-Imagery Compatibility, Cue Order and their

interaction along with a random intercept for each Participant. In addition, a random slope for Perception-Imagery Compatibility was specified. The test revealed a significant main effect of Perception-Imagery Compatibility, LRT: $\chi^2(1) = 10.56$, $p = .001$; GLMER estimate: $\beta = 3.23$, 95% CI = [1.48, 14.1]. Predicted probabilities, derived by simulation, reveal that there was a 99.9% [62.5, 100] chance that participants would execute the imagined response when faced with a compatible cue, but only an 84.3% [74, 91] chance in response to an incompatible cue (Figure 5). Importantly, confidence intervals for compatible cues, in particular, did not overlap with chance (i.e. 50%), showing that participants were generally more likely to produce the currently imagined response compared to the alternative response. Cue Order (LRT: $\chi^2(1) = .09$, $p = .76$; GLMER estimate: $\beta = -.14$, 95% CI = [-4.74, 2.67]) and the interaction between Perception-Imagery Compatibility and Cue Order (LRT: $\chi^2(1) = .02$, $p = .9$; GLMER estimate: $\beta = -.06$, 95% CI = [-4.6, 2.78]) did not affect the type of action slip produced.

2.2.3. Discussion

The aim of Experiment 1 was to create optimal conditions to test whether it was possible to elicit a non-willed action slip by presenting unexpected perceptual action cues during an imagery task. In particular, Experiment 1 aimed to determine whether the compatibility between an imagined and observed movement would modulate the probability that a *general* action slip would emerge, in accordance with predictions derived from Common Coding theory (Prinz, 1990, 1997). Additionally, the experiment explored the nature of the *specific* action slips which were captured, under the assumption that participants should be more likely to execute the alternative response during incompatible trials (i.e. participants accidentally imitated the response primed by the observed movement cue). The results suggest that perceptual cues can indeed elicit

general action slips from participants in spite of an explicitly stated task goal to withhold responding. In addition, the data revealed that general action slips were more likely following compatible, rather than incompatible trials, as predicted by ideomotor models (e.g. Theory of Event Coding, Hommel, 2009). Furthermore, although participants tended to execute their currently imagined response, they were more likely to produce the non-imagined response if they encountered an incompatible cue (which, under our definition, corresponded with the non-imagined response).

The present findings are consistent with the idea that both general and specific action slips emerge as a function of the same ideomotor principles that have been shown to guide voluntary action. Within the ideomotor framework, motor imagery can be characterised as a sub-threshold activation of shared internal representations (i.e. action effects) implicated in the mechanisms of action and perception. In this way, it was predicted that action slips should be more likely in the presence of a movement cue than when no cues were present. Indeed, there was an order of magnitude difference in frequency of actions slips in trials in which such cues were present compared to trials without action cues. Such a low proportion of responses in no-cue trials is hardly surprising given that participants were specifically instructed to refrain from movement during imagery. It suggests that participants had no trouble adhering to this rule when no cues were present. In addition, the marked decline in action slips across Experiment 1 suggests that participants learned to inhibit their urge to respond to the onscreen movement.

Importantly, Experiment 1 revealed that action slips were most likely when the perceived action cues (i.e. colours, sounds and pictures of hands) were compatible with the participant's concurrent mental image. That is, action slips emerged as a function of the degree of dimensional overlap between the internal representations of the observed and imagined movements (Kunde et al., 2004). It was predicted that this should occur

because the process of imagining, for example, a left finger movement, involves activating the same action effect codes which are required to execute the overt act, only to a subthreshold level (i.e. activating internal representations referring to “left”, “down” and “index finger”). The specific action effect representation of a left finger movement would also be activated by the perceived movement cues on compatible trials (i.e. observed left finger movement along with newly associated screen colour and tone), but not incompatible trials. Therefore, when the observed movement cue matched the imagined response on compatible trials, activity in shared action effects was more likely to exceed the threshold for an overt response due to greater conceptual similarity, as predicted by ideomotor models.

However, it must be noted that even incompatible cues elicited action slips. This is, perhaps, not surprising since there remains a great deal of conceptual similarity between the imagined and perceived events on these incompatible trials. For example, an incompatible cue still strongly implies a downward movement which shares spatial features with the imagined downward keypress. Furthermore, it seems likely that there may also have been a very general effect of encountering a surprising visual stimulus on the initial action slips produced by each participant (e.g. a general release from inhibition, or “startle” effect, Maslovat et al., 2013). Recall that there was little difference in the probability of producing a general action slip to either kind of cue at the beginning of the experiment (Figure 4).

A related finding was that the *specific* action slips made by participants were also modulated by the compatibility between the observed and imagined movements. While participants typically executed the motor responses they were currently imagining, they sometimes responded with the alternative finger, specifically if the observed action cue corresponded with this alternative action. This was predicted on the basis that observation of an incompatible cue would induce activity in the action effect codes that

served the alternative response, competing with the endogenously activated action effects that serve the imagined response. If participants had simply been compelled to respond out of sheer surprise by an unexpected movement cue (e.g, Maslovat et al., 2013), then they should have *only* executed the imagined response in both conditions (i.e. there would be no compatibility effect). It is possible that participants may have made errors in tracking the sequence, in that they imagined an incorrect movement when the cue was presented. However, this cannot account for the fact that the vast majority of these non-imagined action slips were observed in the presence of action cues which corresponded with the imagery-incompatible action. If participants were prone to making errors in tracking the sequence, the ratio of imagined to non-imagined action slips would have been the same in both conditions.

It was predicted that action slips should be rare if participants engaged properly with the task and truly tried to inhibit their overt movement until the onset of the predefined imperative stimulus. Thus, the experiment sought to create the maximal conditions for such an action slip to occur by placing participants in an ambiguous situation in which movement was explicitly prohibited but strongly *implied*. To this end, the association training task taught participants an abstract association between colour-tone pairs and specific responses (i.e. action effect association). These trained associations were then exploited in conjunction with an unexpected visual depiction of the corresponding movement in the intention insertion task. However, this creates an inherent problem for a straightforward interpretation of these data, because participants could have been making a form of perseverative error. That is, participants initially learnt that the onscreen colours represented an imperative to act. For example, some participants were taught that yellow stimuli required a left index finger response and blue stimuli a right fingered response. Participants may, therefore, have continued to respond according to this previously learned stimulus-response contingency during the

intention insertion task. In order to address this issue, a new sample of participants was tested (Experiment 2) using the same procedure as Experiment 1, only without the association training task.

2.3. Experiment 2: Eliciting Non-Willed Action Slips During Imagery Without Trained Associations

Experiment 1 created the maximal conditions by which one might induce an action slips under laboratory conditions. However, now that the compatibility effect between action cues and imagery had been established, some of the unnecessary complexity in the design was stripped away to allow more specific claims to be made about the relationship between the activity of internal sensorimotor representations and action intentions. The data collected in Experiment 1 were consistent with the idea that it is possible to elicit a specific, non-willed action slip when the activity induced by endogenous imagery combines activity induced by exogenous action cues. Although the distribution of action slips was as predicted in Experiment 1, it was not possible to rule out that some of the responses could have been caused by a simple perseverative error. That is, a carry-over from the instructions given in the association training task in which participants learnt that coloured stimuli required responses. Under this view, it is the previously established (but erroneously applied) stimulus-response rule which initiated the action, rather than the observed movement. In order to exclude this possibility, Experiment 2 subjected a fresh sample of participants to the same intention insertion task as Experiment 1 but omitted the preceding association training task.

As well as ruling out the effects of perseveration, Experiment 2 aims to evaluate the contribution of the training phase and determine whether it was necessary to establish such an explicit link between abstract stimuli (i.e. colour-tone pairs) and specific responses in order to induce action slips. A lifetime of first-person experience

of visually guided hand movement means that observation of a first-person action – seeing one’s own fingers being depressed when making such a movement in the present case – should inherently represent a potent action cue and that the activity induced by action observation may be driving this effect alone. As in Experiment 1, the action cues still consisted of a visual depiction of congruent and incongruent finger movements presented with colour-tone pairs, but participants no longer learned an explicit association between them prior to the intention insertion task. The flashes of colour and distracting tone merely served to increase the salience of cue onset in this experiment (e.g. a movement of the left or right virtual finger). If Experiment 2 demonstrates the same compatibility effect between observed and imaged movements (i.e. replicating Experiment 1), then this would provide more direct evidence for the role of action effect representations in action initiation than Experiment 1.

2.3.1. Method

2.3.1.1. Participants

Thirty-two participants took part in the experiment (25 female; 5 Left-handed; Age in years: $M=23.2$, $SD=4.9$), recruited from the Plymouth University’s paid participation pool (remunerated £4). The study was approved by Plymouth University’s ethics committee.

2.3.1.2. Materials and Apparatus

The design employed the same set of audio and visual stimuli as were used in Experiment 1. While participants were not explicitly taught to associate abstract Colour-

Tone Pairs with specific movements, the mapping was consistent within participants (e.g. participants either received a yellow: left, blue: right mapping or the opposite, counterbalanced) because the intention insertion task was unaltered from Experiment 1. As in Experiment 1, all factors were fully counterbalanced.

2.3.1.3. Procedure

This experiment used the same intention insertion task that was presented to participants in Experiment 1 but omitted the prior association training task. The protocol for delivering the intention insertion task was identical in all respects.

2.3.2. Results

As in Experiment 1, the percentage of correct responses at the end of each trial (when the “Go!” stimulus was presented, and participants were required to respond) was used to assess whether participants correctly imagined the instructed movements at each point in the sequence. Only participants with accuracy rates within one standard deviation of the group mean ($M=87.9.1\%$; $SD=10.5\%$; $M-SD=77.5\%$) were considered, leading to the removal of 6 participants (leaving $N = 26$ participants).

As in Experiment 1, the first step was to check whether participants followed the instruction to refrain from actively moving their fingers in the imagery intervals by comparing the number of slips on non-cued trials. In the absence of visual stimuli, accidental button presses were rare and occurred on only .2% of beat intervals. As in Experiment 1, the non-cued action slips were roughly equally distributed between button presses with the imagined and non-imagined finger. Accordingly, the test of non-

cued action slips provided no evidence that participants were using finger movements to rehearse the imagined pattern (32 imagined vs. 26 non-imagined slips across all participants, $\chi^2 = .62$, $p = .43$).

The remaining data were analysed using generalised linear mixed-effects regression, GLMER, using the `glmer()` function of the `lme4` package (Bates et al., 2015) in R Version 3.4.2 environment (R Core team, 2017). As in Experiment 1, the analysis proceeded in three steps. The first step was to test whether the surprising visual stimuli would induce unintended key presses (i.e. action slips) in the imagination phase of the intention insertion task. The model included a fixed effect for Cue Presence (2 levels: Yes/No) and Trial (1-96, logged and mean centred) and a random intercept for each Participant and each Beat. In addition, the effect of Cue Presence, Trial and their interaction were allowed to vary within Participant by including a random slope for these terms. As in Experiment 1, the unexpected action cues increased the probability of action slips, LRT: $\chi^2(1) = 30.1$, $p < .001$; GLMER estimate: $\beta = 3.15$, 95% CI = [2.28, 3.87]. Predicted probabilities derived from repeated simulation of the model suggest that there was a 6.85% [3.02, 15.6] chance of an action slip when a cue was present, but only .18% [.08, .47] when no cue was present. Furthermore, the analysis indicated that there was a significant main effect of Cue Order, LRT: $\chi^2(1) = 20.5$, $p < .001$; GLMER estimate: $\beta = -.25$, 95% CI = [-.5, .27], and a significant interaction with Cue Presence, LRT: $\chi^2(1) = 12.5$, $p < .001$; GLMER estimate: $\beta = -1.03$, 95% CI = [-1.67, -.66]. This suggests that participants became less likely to make any responses (i.e. whether following cued or non-cued metronome beats) as the experiment progressed, but that there was a more pronounced decline with respect to trials in which an action cue appeared than non-cued trials.

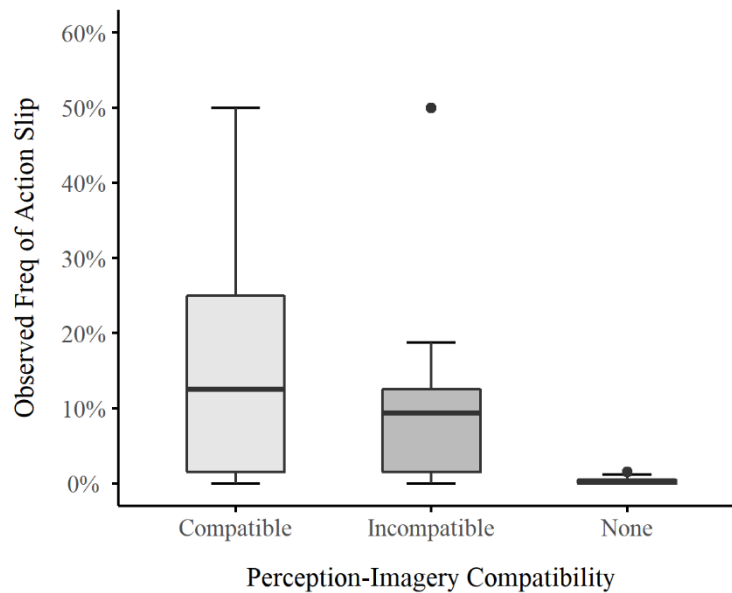


Figure 6. A box plot of the observed frequency of general action slips according to the class of action cue which preceded them in Experiment 2. The observed frequencies are broken down by Perception-Imagery Compatibility with a thick black bar to represent the median, a box showing the interquartile range and whiskers to indicate maxima and minima.

The second step of the analysis was to test the prediction that, in general, action slips were elicited by cues that matched the currently imagined finger keypress, not by those that mismatched. A smaller dataset was constructed which consisted only of intervals in which cues were present (i.e. each participant contributed 32 observations). The observed frequency of action slips recorded in response to each class of action cue (Figure 6) suggests that action slips were more likely when the observed finger press and the imagined finger press were identical (i.e. congruent, $M=16.6\%$) compared to when they were different (i.e. incongruent, $M=9.9\%$). The model included fixed effects for Perception-Imagery Compatibility (2 levels: Compatible/Incompatible), Cue Order (1st to 32nd, logged and mean centred), Colour-Tone Pair (2 levels: Blue/Yellow) and all interactions between these three variables. In addition, the model included a random intercept for each Participant with random slopes for Perception-Imagery Compatibility, Cue Order, Colour-Tone Pair and all interactions. This specification allowed each effect to vary within-participants.




Figure 7. A combination plot indicating the predicted probability of general action slips according to the class of action cue which preceded them across Experiment 2. The upper panel shows a plot of the predicted probabilities, derived by simulation, that cue onset would elicit a general action slip (i.e. any response) according to Perception-Imagery Compatibility and Cue Order. The lower panel shows a difference plot (i.e. Compatible - Incompatible) in which the shaded area shows a 95% confidence interval derived by simulation.

As in Experiment 1, there was a main effect of Perception-Imagery Compatibility, LRT: $\chi^2(1) = 4.91, p = .027$; GLMER estimate: $\beta = -.83$, 95% CI = [.21, 3.99], which suggests that action slips were most likely, in general, in the compatible condition where the observed finger press and the imagined finger press were identical (Figure 7). There was also a main effect of Cue Order, LRT: $\chi^2(1) = 7.6, p = .006$; GLMER estimate: $\beta = -1.53$, 95% CI = [-2.99, -.67], which indicated that action slips became less likely as the experiment progressed. However, the analysis provided no evidence of an interaction between Perception-Imagery Compatibility and Cue Order (LRT: $\chi^2(1) = .19, p = .663$; GLMER estimate: $\beta = -.17$, 95% CI = [-1.18, 1.06]). Tests of the

remaining parameters revealed no evidence of a main effect of Colour-Tone Pair (LRT: $\chi^2(1) = 1.53, p = .215$; GLMER estimate: $\beta = .51$, 95% CI = [-.39, 3.81]). Furthermore, there was no evidence of an interaction between Perception-Imagery Compatibility and Colour-Tone Pair (LRT: $\chi^2(1) = 2.04, p = .152$; GLMER estimate: $\beta = -.6$, 95% CI = [-3.35, .15]), Cue Order and Colour-Tone Pair (LRT: $\chi^2(1) = .24, p = .625$; GLMER estimate: $\beta = -.19$, 95% CI = [-1.44, .86]). or a three-way interaction between Perception-Imagery Compatibility, Colour-Tone Pair, Cue Order (LRT: $\chi^2(1) = .002, p = .965$; GLMER estimate: $\beta = -.01$, 95% CI = [-1.17, 1.13]).

Finally, the third step of the analysis tested the hypothesis that participants should produce the alternative (i.e. non-imagined) action slip at a higher rate when faced with an incompatible cue than a compatible cue. A new dataset was created which consisted of only observations in which an action slip occurred and a new outcome variable was created to indicate whether the participant had executed the imagined or alternative action (coded *one* and *zero* respectively). The model included a fixed effect of Perception-Imagery Compatibility and Cue Order. In addition, the model fitted a random intercept for each Participant and a random slope for Perception-Imagery Compatibility. The Colour-Tone Pair term was omitted from the model because of convergence issues with GLMER due to over parameterisation.




Figure 8. A plot of the predicted probabilities of specific action slips, derived by simulation, according to the class of action cue which preceded them in Experiment 2. A stacked bar plot indicates the predicted proportion of specific action slips (i.e. either the imagined action or alternative action) made according to Perception-Imagery Compatibility.

Figure 8 shows how the estimated proportion of action slips which reflected either the participant's imagined movement or the non-imagined movement was influenced by each kind of action cue. Visual inspection indicates that participants were more likely to produce the imagined response following the onset of both compatible and incompatible cues, as in Experiment 1. In addition, it appears that participants were more likely to produce the alternative (i.e. non-imagined) action slip if they observed a movement that corresponded with the currently non-imagined finger (i.e. an incongruent action cue), rather than the imagined finger (i.e. a congruent cue). Accordingly, the test revealed a significant main effect of Perception-Imagery Compatibility, LRT: $\chi^2(1) = 9.52, p = .002$; GLMER estimate: $\beta = 2.64, 95\% \text{ CI} = [1.62, 22.2]$, indicating that non-imagined action slips were more likely following incompatible cues than compatible cues, as predicted. Moreover, the analysis revealed a significant interaction between Perception-Imagery Compatibility and Cue Order, LRT: $\chi^2(1) = 5.14, p = .023$; GLMER estimate:

$\beta = 1.12$, 95% CI = [-.13, 11.3], indicating that the compatibility effect grew more pronounced as the experiment progressed.

2.3.3. Discussion

Experiment 2 provided further evidence that a combination of imagined and observed action effects is sufficient to induce non-willed behaviour, even though participants are actively trying to withhold it. These data replicate the effects that were observed in Experiment 1, despite the fact that participants were no longer trained to associate an abstract set of stimuli (i.e. colours and tone) with specific responses. This finding rules out that the action slips captured by the intention insertion task could have emerged as a form of perseverative error, caused by participants learning, in an earlier training phase, that the onset of coloured stimuli required them to press a key.

In addition, the results demonstrate that well-established and strongly associated perceptual consequences of our acts (i.e. visual feedback typically associated with first-person action observation) can summate with activity already induced endogenously via imagery and can, sometimes, produce non-willed behaviour. However, the stimuli in Experiment 2 retained a number of incidental features that preclude a straightforward account of the actions slips which were observed. In particular, the intention insertion task used the same set of stimuli as Experiment 1, which combined virtual finger movements with different colours and lateralised tones (a relic of the omitted association training task, Experiment 1) and were retained for continuity's sake. As a result, it is not possible to firmly conclude whether the action slips were caused by the onscreen movement or the laterality of the auditory tone because these elements are confounded in the current design (i.e. left virtual finger movements were always paired with left-panned tones). Indeed, it is possible that the compatibility effect could be

driven by either or both of these elements. In order to allow for a more specific conclusion about the origin of the action slips observed, a new sample of participants was tested (Experiment 3), which applied the same intention insertion task as Experiment 2, but with simpler stimuli. This is because the influence of action observation (i.e. the finger movements) on subsequent action slips, rather than lateralised tones, are of most interest to the present inquiry.

2.4. Experiment 3: Eliciting Non-Willed Action Slips During Imagery with Action Observation

Experiment 2 replicated the effects that were observed in Experiment 1, and in so doing, provided more direct evidence for the notion that activating an actions perceptual consequences, by endogenous (i.e. imagery) and exogenous means (i.e. action observation), can effectively generate the corresponding action. However, the action cues presented to participants in Experiment 2 contained a number of elements (e.g. finger movements presented with lateralised tones and blue/yellow colour flashes) which complicate this interpretation. Indeed, it is impossible to disentangle the influence of an observed movement from the auditory tones on participants' imagery in Experiments 1 and 2, because the action cues were always lateralised (i.e. panned hard left/right in the headphones) according to the observed action (i.e. observed left movement – tone played to left ear). In order to allow a more specific interpretation of the effect of an observed movement on mental imagery, a new sample of participants was tested with the same procedure as Experiment 2, but with simplified stimuli. In Experiment 3, visual depictions of finger movement appeared in tandem with a uniform

yellow background colour change and centrally-panned tone (i.e. presented to both ears by headphones) to increase the salience of the cues, but without influencing the compatibility effect. If Experiment 3 still provides evidence that action slips are modulated by the compatibility between the imagined and observed movements (i.e. replicating Experiments 1 and 2), then this would provide the most direct evidence yet for the relationship between activity in action effect representations and action initiation.

2.4.1. Method

2.4.1.1. Participants

Thirty-two participants took part in the experiment (21 female; 5 Left-handed; Age in years: $M=22.1$, $SD=3.5$), recruited from the Plymouth University's paid participation pool (remunerated £4). The study was approved by Plymouth University's ethics committee.

2.4.1.2. Procedure

This iteration of the intention insertion task was identical to the one used in Experiment 2 with the exception that the task-irrelevant perceptual features of the action cues were now more uniform. Action cues (i.e. left or right finger movements) were presented with uniform yellow backgrounds and tones panned centrally (Figure 9).




Figure 9. Typical and unexpected movement cues displayed during Imagination Phase of Experiment 3. The image “No Cue” is a pair of hands on a white background and in a neutral position with a central fixation cross. Participants saw this image for the bulk of their trials because action cues were rare. The left and right movement cues were accompanied with a yellow background and centrally-panned effect tone and when presented could be compatible or incompatible with the concurrent imagined movement.

2.4.2. Results

As in Experiments 1 and 2, conservative criteria were set and only participants with accuracy rates within one standard deviation of the group mean were considered ($M=84.3\%$; $SD=15.9\%$; $M-SD=68.5\%$), leading to the removal of 4 participants (leaving $N = 28$). Exclusion of these participants does not affect the results. All principal effects of interest remain significant even when these participants are included.

First, to check whether participants followed the instruction to refrain from actively moving their fingers in the imagery intervals, the number of imagined and non-imagined slips on non-cued trials were compared. In the absence of salient visual stimuli, accidental button presses were rare and occurred on only .1% of beat intervals. The test revealed that non-cued action slips were roughly equally distributed between button presses with the imagined and non-imagined finger and provided no evidence that participants used finger movements to rehearse the imagined pattern (14 imagined vs. 10 non-imagined slips across all participants, $\chi^2 = .667$, $p = .414$).

The data were analysed using the same approach as Experiments 1 and 2 and an R script which details this specific analysis was published with the Open Science Framework (Colton et al., 2018b). As with previous experiments, the first step of the analysis was to test whether the surprising visual stimuli would induce unintended key presses (i.e. action slips) in the imagination phase of the intention insertion task. The model included fixed effects of Cue Presence (2 levels: Yes/No) and Trial (1-96, logged and mean centred) and a random intercept for each Participant and each Beat. In addition, the effect of Cue Presence, Trial and their interaction were allowed to vary within Participant by including random slopes for these terms. As predicted, and as in Experiments 1 and 2, the presence of unexpected cues significantly increased the probability of action slips, LRT: $\chi^2(1) = 56.94, p < .001$; GLMER estimate: $\beta = 4.2$, 95% CI = [3.55, 5.26]. Predicted probabilities derived from repeated simulation of the model revealed that there was an 8.57% [3.89, 18.03] chance of an action slip following the onset of a cue, but only .09% [.03, .22] when no cue was present. In addition, Cue Presence and Trial interacted, LRT: $\chi^2(1) = 8.63, p = .003$; GLMER estimate: $\beta = -.79$, 95% CI = [-1.46, -.34], indicating that cues became less effective at eliciting action slips over the course of the experiment. The analysis provided no evidence of a main effect of Trial (LRT: $\chi^2(1) = 2.32, p = .128$; GLMER estimate: $\beta = -.46$, 95% CI = [-.88, .11]) on the probability of producing an action slip.




Figure 10. A box plot of the observed frequency of general action slips according to the class of action cue which preceded them in Experiment 3. The observed frequencies are broken down by Perception-Imagery Compatibility with a thick black bar to represent the median, a box showing the interquartile range and whiskers to indicate maxima and minima.

The second step of the analysis tested whether, in general, action slips would primarily be elicited by cues that match the currently imagined finger keypress (i.e. compatible cues), not by those that mismatched (i.e. incompatible cues). The model included fixed effects of Perception-Imagery Compatibility (2 levels: Compatible/Incompatible) and Cue Order within the trial series (ranked 1st to 32nd, log-transformed and mean centred) and a random intercept for each Participant. In addition, random slopes for Perception-Imagery Compatibility, Cue Order and their interaction were added, which allowed each participant to vary in their general response bias. The observed frequency of action slips recorded in response to each type of action cue (Figure 10) suggests that action slips were more likely when the observed finger press and the imagined finger press were identical (i.e. congruent, $M= 17.4\%$) compared to when they were different (i.e. incongruent, $M= 10.3\%$). The analysis revealed a significant main effect of Cue Order, $LRT: \chi^2(1) = 44.43, p < .001$; GLMER estimate: $\beta = -1.36$, 95% CI = [-1.79, -1.12], reflecting the fact that action slips were most likely in

the early stages of the experiment and became less frequent as it progressed.

Importantly, as in Experiments 1 and 2, there was a main effect of Perception-Imagery Compatibility, LRT: $\chi^2(1) = 7.59, p = .006$; GLMER estimate: $\beta = .54$, 95% CI = [.22, 1.07] confirming that action slips were more likely when observed finger press and the imagined finger press were identical compared to when they were different. In addition, Perception-Imagery Compatibility and Cue Order interacted, LRT: $\chi^2(1) = 5.59, p = .018$; GLMER estimate: $\beta = .34$, 95% CI = [.07, .76]. Figure 11 suggests that this is because action slip frequency was roughly equivalent in the compatible and incompatible conditions early in the session (when the number of action slips across both conditions was high). However, the consistent difference between the compatible and incompatible conditions appears to have emerged later in the session when participants had adjusted to the surprising cues and the total number of action slips had declined.

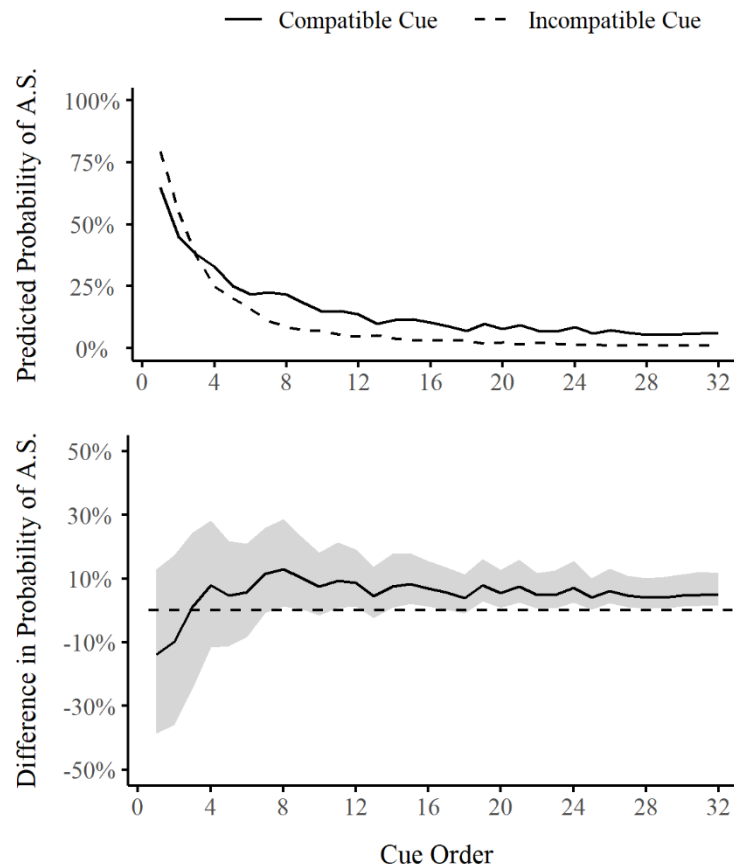


Figure 11. A combination plot indicating the predicted probability of general action slips according to the class of action cue which preceded them across Experiment 3. The upper panel shows a plot of the predicted probabilities, derived by simulation, that cue onset would elicit a general action slip (i.e. any response) according to Perception-Imagery Compatibility and Cue Order. The lower panel shows a difference plot (i.e. Compatible - Incompatible) in which the shaded area shows a 95% confidence interval derived by simulation.

The third step of the analysis tested the hypothesis that participants' specific action slips should generally reflect the currently imagined finger movement and that the alternative action would primarily be elicited by an incompatible cue. The model included fixed effects of Perception-Imagery Compatibility, Cue Order and their interaction along with a random intercept for each Participant. In addition, a random slope for Perception-Imagery Compatibility was specified. Figure 12 shows how the estimated proportion of action slips which reflected either participants' imagined movement or the non-imagined movement is influenced by each kind of action cue. Visual inspection indicated that participants were more likely to produce the imagined

response following the onset of both compatible and incompatible cues. As in Experiments 1 and 2, participants appeared to be more likely to produce the alternative (i.e. non-imagined) action slip if they observed a movement that corresponded with the currently non-imagined finger (i.e. an incongruent action cue), rather than the imagined finger (i.e. a congruent cue).




Figure 12. A plot of the predicted probabilities of specific action slips, derived by simulation, according to the class of action cue which preceded them in Experiment 3. A stacked bar plot indicates the predicted proportion of specific action slips (i.e. either the imagined action or alternative action) made according to Perception-Imagery Compatibility.

Accordingly, the analysis revealed a significant main effect of Perception-Imagery Compatibility, LRT: $\chi^2(1) = 19.01, p < .001$; GLMER estimate: $\beta = 5.73$, 95% CI = [3.81, 16.19] and fitted values have been plotted in Figure 12. Predicted probabilities, derived by simulation, reveal that there was a 100% [93.83, 100] chance that participants would execute the imagined response when faced with a compatible cue, but only a 65.57% [49, 79.2] chance in response to an incompatible cue. Importantly, confidence intervals for compatible cues, in particular, did not overlap with chance (i.e. 50%), showing that participants were much more likely to produce the currently

imagined response compared to the alternative response in the compatible condition.

Finally, the analysis revealed no evidence that either Cue Order (LRT: $\chi^2(1) = 1.38, p = .241$; GLMER estimate: $\beta = -1.64$, 95% CI = [-8.4, -.1]), or the interaction between Perception-Imagery Compatibility and Cue Order (LRT: $\chi^2(1) = .85, p = .356$; GLMER estimate: $\beta = -1.37$, 95% CI = [-8.1, .28]) affected the specific action slip that participants produced.

2.4.3. Discussion

Experiment 3 represents the simplest iteration of the intention insertion task so far and, despite removing several elements that could have contributed to the effect (e.g. association training task and lateralised tones), the outcomes of all principle comparisons were identical to those of Experiments 1 and 2. These data provide the most direct support yet (in this thesis) for the hypothesis that increasing activity in internal action effect representations is analogous with forming an intention to act. Importantly, Experiment 3 demonstrates that this effect is primarily driven by first-person observation of the associated proximal action effects (i.e. visual depiction of a downwards index finger movement) rather than the laterality of the tones (Experiment 2) or an explicitly trained association (Experiment 1).

2.5. General Discussion

In Experiments 1-3, participants were asked to imagine – but not execute – sequences of finger movements, and were then unexpectedly presented with visual images of those actions being executed at some point during the imagination phase. It was predicted that, if forming an action intention is equivalent to activating the relevant

action effect codes, as implied by ideomotor models, then the summed activation of these codes through imagery and perception may suffice to push them super-threshold, causing participants to sometimes execute even movements that they were asked to withhold. Such effects would go beyond prior work showing that observing (or imagining) actions modulates selection or executing planned actions – speeding up responses to congruent cues or inducing subtle changes in associated muscle activity – and confirm a long-held but untested assumption of ideomotor models: that activating an action's effects is equivalent to inducing a motor intention that leads to the initiation of overt action when activated strongly enough.

Consistent with these predictions, the combined influence of top-down imagery and bottom-up perceptual input was sufficient to effectively initiate non-willed button presses (i.e. action slips). This effect was strongest when the observed and imagined actions were identical and was weaker when they were different. This rules out the possibility that the action slips were merely "startle responses" or the release of prepared actions triggered by surprising stimuli (P. Brown et al., 1991; Maslovat et al., 2013) because this should have occurred equally for compatible and incompatible stimuli. Instead, it is consistent with the hypothesized summation of action effect codes, which can happen only if imagined and observed action effects are compatible, but not if they are incompatible. The data, therefore, show, for the first time, that activating an action's perceptual consequences does not only support action planning or selection (see Keller & Koch, 2006, for an application to similar finger sequences as used here), but – in line with ideomotor models – is equivalent to forming an action intention, which when strong enough, suffices to trigger the action itself.

Two further aspects of our results confirm this interpretation. First, the effect of perception-imagery summation became clearer as the experiment progressed. Experiments 1-3 indicated that, at the beginning of the intention insertion task,

participants responded at similar rates following the onset of both kinds of action cue (i.e. compatible and incompatible). However, across the entire task, action slips were consistently more likely when the observed and imagined action were identical than when they were different. This finding rules out that action slips reflect mere startle responses or a release of pre-planned responses. Second, in both experiments, the presentation of action effects primarily caused participants to execute the currently imagined action, compared to the alternative finger press. As predicted, the frequency with which the alternative (non-imagined) action was executed only increased when imagined and presented action effects were not identical, and the visual action cue, therefore, activated the currently non-imagined action. This suggests that two simultaneously activated actions intentions – one derived from imagery and one from observation – summate, such that both can trigger involuntary responses independently, albeit less frequently than if both plans are identical and do not compete (Hommel, 2009; Hommel et al., 2001).

In order to test whether it was possible to induce non-willed action slips in a laboratory setting, Chapter 2 began by creating the maximal conditions for an action slip to occur (Experiment 1). This was achieved by training abstract associations between the sensory representations activated by abstract stimuli (i.e. colours and tones) and the motor representation of the corresponding responses. The subsequent experiments gradually simplified the design by removing the association training task (Experiments 2-3) and then simplifying the action cues (i.e. virtual movements accompanied by a uniform yellow flash of colour and “buzz” tone; Experiment 3). Importantly, the principal effect of interest - the compatibility effect between an observed movement and participant’s imagery on the probability of producing an action slip – was reliably replicated in each version of the intention insertion task (Experiments 1-3). However, one interesting observation is that the compatibility effect exhibited in

the specific action slips analyses appeared to be stronger without the association training task (Experiment 2 and 3) than with it (Experiment 1). Tentatively, and perhaps unsurprisingly, this suggests that well established sensorimotor associations between an observed movement and its typical visual outcome (i.e. acquired through a lifetime of experience observing one's own actions) were more effective at triggering the corresponding action than the trained associations (i.e. colours and tones). In fact, the association training task, if anything, obscured the effect. Perhaps the colours and tones created a simpler response contingency (i.e. a flash of colour demands a response) and simply made participants more likely to respond when they encountered any coloured stimulus.

Previous work has characterised action initiation as an integration-to-bounds mechanism, which accumulates evidence for appropriate actions and executes the most suitable one if a threshold is reached, irrespective of whether this evidence results from the evaluation of imperative task cues or from internal decisions about whether to act or to wait (Churchland et al., 2008; Murakami et al., 2014). Once an action has surpassed the threshold, it triggers an “avalanche” of motor processes, supported by deep brain structures such as the basal ganglia, which facilitate fluent execution. The data from Experiments 1-3 are entirely consistent with this view but reveals that activated action effects – either through imagery or action observation – can push this activity super-threshold without additional processing. The data, therefore, provide direct evidence for the notion that forming an action intention may be identical to activating – imagining – the action's perceptual consequences, which is executed when the activation becomes strong enough (Prinz, 1997; Shin et al., 2010). These findings go beyond prior research that has shown only that motor imagery (Ramsey et al., 2010), action observation (e.g. Bach, Bayliss, & Tipper, 2011; Bach et al., 2007; Brass et al., 2001) and learned action effect associations (Elsner & Hommel, 2001) can modulate the selection, speed and

accuracy of already planned responses, but which has left open whether action effect activation suffices to initiate actions independently.

At this stage, the nature of the mechanism by which action observation induces motor responses from participants remains an open question. On the one hand, an observed action may trigger the imagined response by exerting an additional (i.e. additive) effect on activity in action effect codes already activated by imagery until the imagined response is activated so strongly that it is inadvertently executed (i.e. effectively inserting an intention to act). On the other hand, it is also possible that either the visual or the imagery-based action effects act by *removing inhibition* from one particular action (e.g. Rieger, Dahm, & Koch, 2017), in line with James' dictum (1890) that a mental idea of an action will lead to overt execution, as long as it is "not kept from so doing by an antagonistic representation present simultaneously to the mind". However, ideomotor models are agnostic about these two possibilities.

On a conceptual level, both are equivalent, as both summation of activation and removal of inhibition have a net positive effect in bringing activation closer to the threshold for execution and inducing overt behaviour. Indeed, an intention *not* to act has been shown to be similarly controlled by mental images of non-action (e.g. Kühn, Elsner, Prinz, & Brass, 2009; Ridderinkhof, van den Wildenberg, & Brass, 2014), and both inhibiting and excitatory roles have been attributed to complementary basal ganglia-prefrontal circuits for action execution (for a recent review, see Calabresi, Picconi, Tozzi, Ghiglieri, & Di Filippo, 2014). However, while Experiments 1-3 were not designed to resolve this issue, they may provide tentative support for an "energizing" role of action effect activation.

If the visual action cues primarily act through removing inhibition, then they should have induced action slips only for the currently imagined action; removing

inhibition should not cause a fully-fledged button press of the non-imagined finger. However, this was not the case. In Experiments 1-3, even the currently non-imagined action was executed - albeit less frequently than the imagined response - if it was so cued by the respective visual action cue, providing some evidence for an additive, “energizing” influence of action effect activation on action initiation. In order to more thoroughly evaluate the evidence for the proposed facilitating effect of action observation on the generation of non-willed action slips, a new version of the intention insertion task was developed (Experiment 4, Chapter 3) which incorporated a more sensitive behavioural measure – finger pressure. Importantly, Experiment 4 also presented participants with “neutral” visual cues that comprised of the same salient flash of colour and deviant tone as accompanied the action cues in Experiment 3 but contained no movement element (i.e. the virtual fingers did not move). Together, these changes allow a direct comparison of the manner in which action observation, in combination with mental imagery, initiates non-willed motor action slips relative to a perceptually similar, but motorically neutral, visual cue.

3. Chapter 3: Do Action Cues Primarily Facilitate or Interfere with the Generation of Non-willed Action Slips?

3.1. Introduction

The purpose of Chapter 3 is to present an updated intention insertion task which uses a similar design to Experiment 3 (Chapter 2) with a more sensitive behavioural measure – finger pressure – and analysis in order to reveal how action observation interacts with imagery to initiate actions. In addition, Experiment 4 (present chapter) incorporates a novel baseline condition which is designed to test whether movement cues primarily facilitate or interfere with the generation of non-willed action slips during imagery. Furthermore, several other changes were incorporated into Experiment 4 to address some of the limitations of the previous experiments, while taking steps to improve the quality and quantity of the data produced by each participant.

The findings of Chapter 2 (Experiments 1-3) supported the assumption that the mechanisms of action, perception and imagery rely on a common set of internal representations (i.e. action effect codes) which are coded in a fundamentally perceptual format. Importantly, Experiments 1-3 (Chapter 2) provided evidence in favour of the notion that heightened activity in action effects, whether induced by endogenous imagery or exogenous perceptual input, can initiate the corresponding action (i.e. effectively inducing an intention to act). Thus far, the data have suggested that action effect representations (or, “mental images”) do not only shape *how* and *when* actions are performed but serve a generative purpose - influencing *whether* actions are executed at all. Participants were consistently more likely to execute an action they were imagining, but trying to withhold, if they were presented with a surprising visual image of a compatible action, rather than incompatible action, occurring in front of them.

However, it was not possible to determine precisely how the movement cues generated action slips and gave rise to the compatibility effect described in Chapter 2 because Experiments 1-3 lacked an appropriate baseline condition for comparison. On one hand, congruent movement cues may have facilitated action slips, wherein perceptual input exerted an additive effect on the action effect codes already activated by imagery. On the other hand, incongruent cues may have simply interfered with the initiation of the imagined action more strongly than congruent cues. Chapter 3 describes a single experiment (Experiment 4) aimed at addressing this question with the addition of a “neutral” cue, which consists of the same salient flash of colour and surprising tone as the action cues used in Experiment 3 (Chapter 2), but without the movement element. By subtracting the responses made to neutral cues from those which involved movement (i.e. action cues), it is possible to control for the portion of the effect caused by encountering a surprising stimulus. That is, the remainder of the effect can only be attributed to the observed movement. This design will be used to evaluate the evidence for several related explanations relating to the manner in which movement observation and movement imagery interact to generate non-willed action slips.

One possibility is that the action cues in Chapter 2 largely *facilitated* the initiation of action slips, albeit to a varying extent in response to compatible and incompatible cues. On this account, the observation of a sudden virtual hand movement exerted an additional bottom-up effect on action effects which overlapped with those simultaneously activated in a top-down manner by imagery. On this view, the effect of movement observation was to push the summed activation higher and make action initiation generally more likely than a similarly salient cue with no movement element. In other words, when an observed movement aligned with an imagined movement, perception and imagery activate a common set of action effect codes, leading to a higher overall activation than would follow a similarly salient, but motorically neutral stimulus

– enhancing the execution of action slips. When the observed movement was incompatible with the imagined movement, the imagined and observed actions recruited differing sets of action effect codes. In this condition, the common perceptual elements of the imagined and observed act still leads to the activation of some overlapping action effects (e.g. all action cues thus far have presented a downward movement of an index finger). However, because incompatible cues shared fewer perceptual elements with the imagined movement (e.g. mismatching laterality of the finger movement) the additive effect of activation by the observed movement was weaker, and thus action initiation was less likely. Experiment 3 provided some evidence of facilitation, in that incompatible action observation occasionally led participants to execute the visually primed response, even though they weren't imagining that movement at the time. This demonstrates that the observed movement must facilitate action initiation to some extent, but competition between rival sets of action effect codes may well involve more than just additive effects.

An alternative possibility is that the sudden onset of a salient visual stimulus was largely responsible for triggering subsequent responses and that action observation only *interfered* with the execution of these incidental responses. On this view, when an observed movement matches an imagined movement, the resulting interference effect would be minimal because a similar set of action effect codes are activated. However, when an observed and imagined movement differed, the two simultaneously active sets of action effects may have inhibited and interfered with one another (Munakata et al., 2011). Thus, according to this account, observation of an incompatible action will exert a largely negative effect on the summed activation of action effects supporting the imagined movement and, as a result, the response crosses the threshold for execution less frequently. Indeed, it has been claimed in previous research that visuomotor priming effects on action initiation are largely the result of interference between the

perceived and planned actions (Ramsey et al., 2010). Ramsey et al. (2010) found that imagining an action that was different from the one participants were required to perform impaired subsequent action production. However, there was no evidence that imagery facilitated action production when imagery and the required response were identical. This provides suggestive evidence that interference may account for some of the reduction in the effect when the observed and imagined movements did not match in Experiments 1-3 (Chapter 2). However, given that in Chapter 2, non-imagined actions were consistently elicited at a marginally higher rate when visually primed (relative to when the imagined action was primed), it is unlikely that interference alone can account for all of the effects reported.

In fact, the final and most plausible possibility is that the action cues in Chapter 2 both facilitated and interfered with the generation of non-willed behaviour during imagery. Under this account, when participants' mental imagery was compatible with the observed action, the visual input exerted an additional, facilitating effect on activity induced by imagery. That is, the summed activation of shared action effects brought the corresponding action closer to the threshold for execution. However, when mental imagery and the observed action were incompatible, the action effect codes activated by the visual movement competed and interfered with the action effect activated by imagery which, in turn, reduced the likelihood that either set of action effects will cross the threshold for execution. Certainly, it has been claimed that basal ganglia-prefrontal circuits can exert both an inhibitory and excitatory effect on action production (for review, see Calabresi, Picconi, Tozzi, Ghiglieri, & Di Filippo, 2014). Accordingly, the purpose of Experiment 4 is to examine the manner in which a confluence of endogenous (e.g. mental imagery) and exogenous (e.g. action observation) activation generates non-willed action slips compared to a neutral baseline condition (which contains no movement element).

3.2. Experiment 4: Intention Insertion Task 2.0 – Finger Pressure as an Improved Behavioural Measure of Non-Willed Action Slips

The general design of the intention insertion task used in Experiment 1-3 (Chapter 2) provided an excellent starting point in the capture of action slips under laboratory conditions. However, the basic task was limited in at least two respects which are addressed in Experiment 4. The first limitation is that the presence/absence of a keypress on a given trial (i.e. a binary dependent variable) does not provide a great deal of detail about *how* the action slip occurred (e.g. the magnitude or timing of movement), merely that it did occur. This represented a convenient simplification because it operationalised what non-willed action slips represented in an intuitive way (i.e. whether a keypress did or did not occur, following a specific event). Nevertheless, as a behavioural measure in this context, not least given their relative scarcity, binary outcomes represent an information-poor choice of dependent variable.

The second limitation emerges from the observation that, in Experiments 1-3 (Chapter 2), participants were sometimes observed making small twitches at cue onset that were too minor to register as keypresses. The pressure required to fully depress a key on a computer keyboard represents an arbitrary threshold that was used to determine whether a movement had occurred. However, these subthreshold responses were essentially wasted data and are an important part of the picture. If the proposed summation of action effects was responsible for producing the action slips observed in Chapter 2, then all movement produced by participants in response to cues is relevant to this enquiry and should follow ideomotor principles. As a result, the intention insertion task was redesigned to function with a new response interface, abandoning the computer keyboard in favour of recording the pressure exerted by the participant's fingers using

pressure sensors. The hardware served to precisely capture any and all movements made by each of the participant's fingers, rather than simply record whether or not their movements have exceeded an arbitrary threshold (i.e. a keypress). As a consequence, in Experiment 4, action slips were operationalised as the magnitude of induced pressure following cue onset (i.e. a continuous variable). In addition, a number of significant improvements were made to allow for more nuanced predictions in this experiment and those which follow.

In the new version of the intention insertion task, participants were again instructed to imagine performing imagined sequences of finger taps, one movement at a time. However, the pressure was recorded from both index fingers simultaneously. As a result, a new independent variable, Imagery-Finger Congruence, was created to identify the pressure values which corresponded with the finger that was congruent with the currently imagined movement, and which values corresponded with the alternative, incongruent finger. Contrasting the induced pressure recorded at the imagery-congruent and imagery-incongruent fingers in response to each kind of cue makes it possible to quantify the contribution that imagery makes to the size of the movement.

By continuously recording finger pressure, rather than discrete keypresses, Experiment 4 addressed another limitation of the previous experiments. The keyboard-based intention insertion task from Chapter 2 was unable to definitively rule out that participants were engaging in a covert movement strategy in which they lightly tapped their fingers on the keys as a non-imagery based means of rehearsing the movement sequence. Although the experiment recorded a very low proportion of key presses during non-cued beats, it was not possible to completely exclude this interpretation because, by definition, only responses that exceeded the keypress threshold were captured. Experiment 4, which relied on continuous pressure data, is much more effective at capturing the full range of responses produced by participants. If the action

slips captured by Experiment 4 arises from an interaction of imagery and movement observation, rather than simply a physical rehearsal strategy, then participants should produce much stronger presses with their imagery-congruent finger following cues than on non-cued beats.

Further improvements were made to the trial structure by incorporating a phase in which participants were required to execute the briefly presented movement sequence. In Experiment 4, following the initial presentation of the finger tapping sequence, participants were required to *perform* the sequence for four beats (e.g. one complete repetition of the sequence), prior to the imagery phase of each trial. Instructing participants to physically rehearse the movement sequence, prior to imagery, was intended to enhance their ability to conjure a strong mental image of how these actions would look and feel in the subsequent imagery phase. In Chapter 2, some participants reported that they had struggled to build a strong mental image of the perceptual consequences of the required movements, so this feature provided a useful sensory reference for the feelings they were asked to simulate. In addition, asking participants to perform the sequence prior to imagery served to further emphasise task instructions, namely, when physical movements were permitted and when they were not (e.g. during imagery). Finally, unambiguous labels which explicitly highlighted the current task demands (e.g. “Watch”, “Perform” and “Imagine”) were added to the top of the screen during each phase of the experiment. Together, these changes should both enhance participants’ ability to engage in effective imagery and limit the extent to which they rely on covert movement during imagery as a means of rehearsing the sequences.

Experiment 4 was principally concerned with dissociating facilitative and interference effects arising from the sudden onset of visual movement cues during imagery. It is possible to tease apart facilitative effects from interference effects on induced pressure by taking the responses that followed congruent and incongruent

action cues and comparing them to responses on neutral cue trials in which no movement occurred. Subtracting the pressure induced by a neutral cue from the pressure induced by an observed movement at the cue-congruent and cue-incongruent finger (i.e. the finger which matched/mismatched the observed movement) effectively controls for the effect of being confronted with a surprising visual cue. The remaining portion of the effect can only be attributed to the influence of the observed action. If action observation exerts an additional, facilitative effect on activity in action effect codes (i.e. beyond the effect of encountering an unexpected stimulus), then it should generate larger pressure spikes at the cue-congruent finger than neutral cues. In addition, if action observation can also interfere with the generation of action slips then smaller pressure spikes should be induced at the cue-incongruent finger than neutral cues. Finally, if endogenous activation of action effects via imagery is driving the initiation of non-willed action slips, then the effects of action observation should be much stronger at the imagery-congruent finger, than the imagery-incongruent finger.

3.2.1. Method

3.2.1.1. Participants

Forty participants took part in the experiment (31 female; 8 Left-handed; Age in years: $M=20.5$, $SD=3.9$), recruited from the Plymouth University's paid participation pool (remunerated £4). The study was approved by Plymouth University's ethics committee.

3.2.1.2. Materials and Apparatus

The experiment was conducted in a dark, soundproof room and was administered on a Windows 10 computer using custom hardware (see Figure 13), an Arduino Leonardo interface and displayed on a 19" LED computer monitor (Resolution: 1900x1200; Refresh rate: 60Hz). Custom experimental software required to interface with the new hardware was adapted from the original E-Prime script used in the previous chapter by Plymouth University's Tech Office.



Figure 13. The custom hardware developed for the pressure-based version of the intention insertion task in Experiment 4. This interface consisted of a standard musical keyboard and two force-sensitive resistors, or “pressure pads” to record the pressure exerted by each of the participant’s fingers for the duration of the experiment. The keys facilitated a good range of movement and the spacing between the two active keys was similar to the keyboard-based intention insertion task used in Chapter 2.

The new hardware consisted of two force-sensitive resistors, hereafter “pressure pads”, attached to two of the keys of a musical keyboard. The resistance of the pressure pads varied according to the force applied to the surface of the sensor and they were used to record the pressure applied by the participant’s fingers on the keys. The experimental software was programmed to record the pressure from each pressure pad at 20ms intervals (± 3 ms due to sampling variability) for the duration of the

experiment. The keyboard itself was not functional, but merely facilitated a good range of motion and smooth action. Masking tape was placed over the unused keys to ensure participants were interacting with the keyboard correctly and consistently. Participants wore over-ear headphones and placed the tips of their left and right index fingers onto the pressure pads. The stimuli for this experiment corresponded exactly with those used in Experiment 3 (Chapter 2).

3.2.1.3. Procedure

Experiment 4 consisted of four blocks of 24 trials in which participants are presented with a series of finger tapping sequences which they are first asked to watch, then perform and finally imagine performing prior to the onset of the true imperative stimulus (“Go!”). Each trial had the same general structure with a slow, steady metronome beat (440Hz tone presented at 800ms interval) helping to guide the underlying pace of each trial. For the first four beats of every trial, participants were presented with an image of a pair of hands (egocentric perspective) with the word “Watch” written above them. At this stage, an arrow indicated a four-beat tapping sequence by pointing at the left and right index fingers of the virtual hands. During the watch phase, participants were merely required to observe and retain the implied sequence. After four beats, the instruction at the top of the screen changed to “Perform” and participants were instructed to press keys according to the sequence implied by the arrow and in time with the metronome beat. After one complete repetition of the sequence (i.e. four beats), the instruction changed to “Imagine”, and participants were required to rehearse the sequence in their minds by imagining that they were performing it in time with the beat, but without physically moving. This final imagery phase had a variable duration of 5-8 beats, after which the imperative stimulus – the word “Go!” –

was presented and the participant was required to execute the response they currently imagined.

The crucial manipulation was that in 12 of the 24 trials in each block, unexpected visual cues appeared at the onset of one of the metronome beats (on either beat, 5-8). The visual cues consisted of a sudden yellow background change and either no movement (i.e. neutral cues) or an implied movement of the left or right index fingers moving downwards and back to a resting position (i.e. action cues). In other words, the action cues presented participants with what they would see if they actually carried out a finger tap themselves (i.e. the visual perceptual consequences). The visual cues were presented with a background colour change to yellow and a sharp “buzz” tone in place of the typical metronome beat to increase the salience of the event.

Two sets of pressure values were recorded throughout the whole experiment which corresponded with the left and right index finger. The outcome variable was the finger pressure induced in the 800ms interval between metronome beats in each trial. For each observation, the two sets of pressure values were dummy coded to identify whether the value corresponds with the imagery-congruent and imagery-incongruent finger. Participants encountered 48 cues in total which were counterbalanced according to Cue-Finger Congruence (congruent/incongruent/neutral), Imagery-Finger Congruence (congruent/incongruent), trial length (5, 6, 7 and 8 beats) and cue placement in the sequence (1 or 2 beats prior to sequence end).

Prior to the main task, participants were exposed to a short practice of up to 24 control trials (e.g. no surprising visual cues) in which accuracy feedback was provided. Participants had to maintain at least 75% accuracy for a minimum of 8 trials in order to progress. Participants were given the opportunity to repeat the practice if required until they had become sufficiently proficient at tracking the sequences. In addition,

participants completed an acclimatisation task in which they had to maintain three different reference pressures for five seconds at a time. The pressure training task served to emphasise the sensitivity of the apparatus, participants' awareness of their finger position and the instruction to avoid lightly tapping the keys.

3.2.2. Results

3.2.2.1. Exclusion Criteria

As with Experiments 1-3 in Chapter 2, participant accuracy during control trials was used as an index of task engagement. All participants successfully completed the practice task and understood the task requirements. In the main task, participants with accuracy rates below one standard deviation of the group mean (chance=50%; $M=81.6\%$; $SD=17.5\%$) were excluded from further analysis, leading to the removal of 7 individuals (leaving $N = 33$).

3.2.2.2. Data Preparation

The resulting data were subjected to some pre-processing and organising steps which were intended to, (a) improve the signal-to-noise ratio; and (b) distil a dependent variable consisting of a single value for each 800ms interval which captured the magnitude of pressure exerted by each finger during this period. During each trial of the experiment, the pressure was sampled (i.e. recorded) at approximately 20ms intervals during each 800ms interval (equates to 40 samples per interval, ± 3 due to sampling variability). Participants were instructed to rest each of their index fingers on the pressure pads for the duration of the experiment. However, the resting pressure of each

finger naturally varied slightly at the start of each 800ms interval. As a result, the pressure data was baselined by subtracting the mean pressure recorded at each finger in a 200ms window prior to beat onset (i.e. from the final 200ms of the preceding interval) from each sample following beat onset in the subsequent interval.

In each of the forthcoming analyses, a pair of values were derived which represented the magnitude of induced pressure at each finger between the metronome beats of the imagery phase of every trial. Calculating the area under the curve (AUC) is an efficient means of capturing the magnitude of any pressure spikes that were captured. It was implemented using the trapezoid rule; a common numerical integration method which involves summing the area of a series of trapezoids of known number and width which fit inside the curve.

One important consequence of simultaneously recording continuous pressure data from both fingers is that the coding scheme used to classify the visual cues in Chapter 2 (Experiments 1-3) was no longer suitable. In Chapter 2, action cues were categorised according to their compatibility with the concurrently imagined movement, but Perception-Imagery Compatibility is very confusing when broken down by Imagery-Finger Congruence. For example, compatible cues are expected to facilitate large movements of the imagery-congruent finger, whereas incompatible cues are expected to interfere and induce smaller pressure spikes. In contrast, the opposite prediction is made for the imagery-incongruent finger, such that incompatible cues (which visually prime the imagery-incongruent movement) should elicit the largest spikes in pressure. Furthermore, compatible cues (which visually prime the imagery-congruent movement) should interfere with the initiation of an imagery-incongruent response. In order to make inference and interpretation straightforward, participants' fingers were classified as cue-congruent, cue-incongruent or neutral to the observed response following cue onset. This new variable is referred to as Cue-Finger Congruence.

3.2.2.3. Data Analysis

In contrast to the action slips described in Chapter 2, the dependent variable represents a continuous measure of pressure induced at each finger, rather than a discrete binary outcome. These data were over-dispersed (i.e. the values occupied a wide numerical range) and zero-inflated (i.e. no pressure was recorded in a large number of intervals). As a result, these data were subjected to Bayesian Mixed Effects Regression using the BRMS package (Bürkner, 2017) which allowed a model to be specified which accurately accounted for the structure of the data. All subsequent models relied upon four chains of 3,500 iterations (1000 iteration burn-in) in length to converge on stable parameter estimates. Bayesian modelling requires the specification of prior knowledge about the distribution of likely values of subsequent parameter estimates. Weakly informative priors were specified based on recommendations of the package authors (Bürkner, 2017; Gelman, 2006). An additional term was added after inspecting model diagnostic plots to allow residual variance in the model to vary (in this case, increase) as the experiment progressed; this substantially improved model fit.

Model fit was checked for convergence visually using trace plots and via the \hat{n} -effective and \hat{R} statistics; in all cases, convergence appeared to be satisfactory. To visualize and draw inferences from the results (e.g., to examine the effect of action cues relative on induced pressure), it is possible to make linear predictions from the posterior distribution of the model for each category, and across the full range of observed values of all other parameters. This approach provides samples from the posterior distribution of responses for each type of action cue and, by subtraction, the differences between cue types which are of interest.

To evaluate each prediction in the forthcoming analyses, data were simulated from the relevant model and used to calculate summary statistics (i.e. to indicate the effect size) and Bayes-p values (i.e. to indicate the strength of evidence in support of the predicted effect; Wagenmakers et al., 2018). The posterior mean and 95% posterior credible interval (PCI) for the differences of interest represents an estimate of the size of the effect under scrutiny (given the model and the data). Credible intervals in Bayesian statistics have a much more intuitive definition than confidence intervals in the frequentist tradition. For example, a 95% credible interval refers to the numerical range with a 95% probability of including the true value of a particular parameter of interest (e.g. the true effect size).

The posterior predictive Bayes-p values associated with the forthcoming contrasts represent the probability (given the model and the data) that the effect was in the predicted direction (i.e. H_1), as opposed to the opposite direction (i.e. H_0). To aid interpretation, Bayes-p values are recapitulated as evidence ratios (ER; analogous with Bayes factors) to describe the strength of evidence in favour of the directed hypothesis, weighted against the alternative hypothesis. For consistency, the weight of evidence in support of each directed hypothesis was described using terminology regularly applied to Bayes factors (Wagenmakers et al., 2018); “extremely strong” (ER: >100), “very strong” (ER: 30-100), “strong” (ER: 10 - 30), “moderate” (ER: 3 - 10), “anecdotal” (ER: $>1 - 3$) and “no evidence” (ER: 1). Importantly, evidence ratios less than one indicate that the evidence supports the alternative hypothesis; even more so for values which are a tiny fraction of one. For example, an $ER < .01$ represents *extremely strong* evidence in favour of the *alternative hypothesis*.

3.2.2.4. Analysis of Cued and Non-Cued Beats

The purpose of the first analysis step was to evaluate the evidence that the pressure induced by the visual cues was largely driven by participants engaging in a physical rehearsal strategy (i.e. rather than purely imagery) in order to track the finger tapping sequences. The model outcome was finger pressure (AUC) in the 800ms interval between each metronome beat of the imagination phase of every trial. Finger pressure was recorded from both the left and right index fingers during both cued (i.e. 48 cued observations) and non-cued beats (i.e. 264 non-cued observations). One predictor, Imagery-Finger Congruence, coded each observation according to whether the finger being recorded at that moment matched the participants' imagined movement (i.e. imagery-congruent), according to the learned sequence, or mismatched the imagined movement (e.g. imagery-incongruent). A second predictor, Cue Presence, coded each observation according to whether a surprising visual cue was presented at the start of the critical interval or whether no cue was present. A third predictor, Experimental Block, coded each observation according to the block of trials in which it was made (i.e. numbered one to four). For cued and non-cued intervals finger pressure was analysed with a model which specified fixed effects of Imagery-Finger Congruence, Cue Presence, Experimental Block, and all interactions between these terms. In addition, Cue Presence, Imagery-Finger Congruence and the interaction between these variables were allowed to vary within and between participants. For clarity, the predictors of interest are capitalised when they are referred to in the text.




Figure 14. A multi-panel plot showing finger pressure in the critical 800ms interval following the onset of the onset of a typical metronome beat or the onset of a salient visual cue in Experiment 4. The top and bottom panels separate responses from fingers which matched (i.e. imagery-congruent) or did not match (i.e. imagery-incongruent) the imagined response. The coloured lines indicate whether the cue was present (i.e. red) or not (i.e. blue) at the start of the 800ms interval.

Figure 14 shows the induced pressure that was recorded in the 800ms interval between metronome beats according to whether or not a cue was presented at beat onset and whether the finger being recorded matched (i.e. imagery-congruent) or mismatched (i.e. imagery-incongruent) the participants' currently imagined movement. Importantly, the size of the finger presses captured by Experiment 4 appears to vary considerably with respect to Cue Presence and Imagery-Finger Congruence. When the recorded finger was congruent with participants' mental imagery (Figure 14, top panel) it appears that finger pressure was similar for both cued and non-cued beats in the 250ms period post-beat-onset. However, visual inspection suggests that the onset of visual cues

(indicated by the red line) induced a substantial spike in finger pressure, which began at approximately 250ms and peaked at 400ms. In contrast, when no cues were present (indicated by the blue line), finger pressure showed the opposite pattern and steadily decreased in the same time period. When the recorded finger was incongruent with participants' mental imagery (Figure 14, bottom panel), finger pressure following cued and non-cued beats showed a similar pattern to the imagery-congruent responses, although the absolute size of the press appeared to be much smaller.

In order to gauge the strength of evidence for these observations, data were simulated from the model and summary statistics, Bayes-p values and evidence ratios regarding the differences of interest were calculated. It was predicted that if the visual cues were responsible for driving the action slips captured by Experiment 4, rather than a more trivial covert movement account, then participants should have produced much larger responses when cues were present than when they were absent. It was anticipated that participants should press much harder with their imagery-congruent, than the imagery-incongruent finger and, importantly, this difference should be greater following cues than on non-cued beats (i.e. Cue Presence and Imagery-Finger Congruence should interact). Importantly, for brevity and to safeguard against false positives due to multiple testing (Cramer et al., 2016), this analysis focused only on the comparisons that capture this latter prediction, namely the two-way interaction of Cue Presence and Imagery-Finger Congruence. A test of this interaction revealed that there is extremely strong evidence ($ER < 999$) in favour of the prediction that imagery-congruent responses were much stronger than imagery-incongruent responses following cued beats relative to non-cued beats (Posterior Mean = 6.5, 95% PCI [4.26, 8.79], Bayes-p > .999). In other words, this finding indicates that the presence of visual cues was associated with much more substantial presses at the participants imagined finger, than their non-imagined finger, relative to when no cues were present. This finding argues

against covert movement as a viable explanation for the action slips captured by Experiment 4, replicating the findings of Experiments 1-3 (Chapter 2)

3.2.2.5. Analysis of Cued Beats

This analysis modelled the pressure induced with each finger when participants were presented with different types of action cues during imagery. In addition, pressure induced following action cues was contrasted with responses to neutral cues (i.e. no movement element) to control for the effects of encountering salient visual stimuli during imagery. Following the findings of Chapter 2 (Experiments 1-3), it was anticipated that induced pressure would be greater when the recorded finger (e.g. left index) matched the observed movement (e.g. virtual left movement) than when it mismatched (e.g. virtual right movement). This is the *intention insertion effect*. Further, if pressure is caused by the summation of action effect codes activated by perception and imagery, the difference in induced pressure between by congruent and incongruent observation should be greater on the finger (e.g. left index) that is currently imagined moving (e.g. imagined left index press) than on the finger that is imagined not moving (e.g. imagined right index press).

The model outcome was finger pressure (AUC) in the 800ms critical intervals after visual cues were presented. Finger pressure was recorded from both the left and right index fingers in each of the 48 cued intervals. One predictor encoded each observation according to whether the finger being recorded at that moment was congruent or incongruent with the finger the participant was imagining moving (Imagery-Finger Congruence). A second predictor encoded whether the visual cue was congruent with the currently recorded finger, whether it was incongruent with the currently recorded finger, or whether a neutral, no-movement cue was presented (Cue-

Finger Congruence). The model also included Experimental Block (1 to 4) and Cue Order (1 to 48), along with a full set of interactions between these predictors. In addition, Imagery-Finger Congruence, Cue-Finger Congruence, Cue Order and all interactions were allowed to vary within-participants.




Figure 15. A multi-panel plot showing finger pressure in the critical 800ms interval following the onset of the onset of each type of visual cue in Experiment 4. The top and bottom panels separate responses from fingers which did or did not match the imagined response. The coloured lines indicate whether the cue contained no movement (i.e. neutral – blue), or if it included an animation of the same (i.e. congruent – red)/other finger (i.e. incongruent – green) from that being recorded.

Figure 15 shows the pressure induced in the critical 800ms period which followed the onset of visual cues with respect to Cue-Finger Congruence and Imagery-Finger Congruence. Visual inspection suggests that the pressure induced following visual cues was much higher if participants already imagined a movement of the recorded finger. Furthermore, the pressure induced by the visual stimuli appears to vary substantially

according to whether the cue was neutral (i.e. no movement, blue line) or contained a movement which was congruent (red line) or incongruent (green line) with the finger producing the response.




Figure 16. A multi-panel plot showing differences in pressure (AUC) during intervals when the recorded finger was either congruent or incongruent with respect to the observed movement vs. a neutral cue in Experiment 4. Positive scores indicate participants pressed harder than in equivalent trials where a neutral cue was presented (i.e. implying facilitation). Negative scores indicate that participants pressed less hard than following neutral cues (i.e. implying interference). The top and bottom panels distinguish responses where the participant was/was not also required to imagine a movement in the recorded finger. Error bars are the 95% PCI.

In order to derive a summary measure of induced pressure, the area under the curve (AUC) was calculated for individual observations within each condition and fitted values from the model were plotted in Figure 16. Figure 16 shows how much pressure was induced when the recorded finger was either congruent or incongruent with an observed movement, as compared with pressure after a neutral cue (i.e. no movement,

baseline condition). In this model, the magnitude of the intention insertion effect was indexed by contrasting the baseline pressure (i.e. after subtracting neutral responses) induced when the participant's finger matched the observed movement (i.e. cue-congruent) with the pressure induced when the finger mismatched the observed movement (i.e. cue-incongruent). Furthermore, the intention insertion effect was derived separately for instances in which the recorded finger corresponded with the participant's currently imagined movement (i.e. imagery-congruent) and when it corresponded with the alternative, non-imagined movement (i.e. imagery-incongruent). Importantly, the magnitude of the difference between cue-congruent and cue-incongruent responses corresponds with the magnitude of the intention insertion effect. For example, Figure 16 (top panel) indicates that action cues induced much stronger presses at the cue-congruent finger than the cue-incongruent finger when participants also imagined a movement of the currently recorded finger. In contrast, Figure 16 (bottom panel) indicates that when the recorded finger was incongruent with participants' mental imagery, the action cues did not strongly influence the magnitude of cue-congruent or cue-incongruent responses. That is, there does not appear to be a strong intention insertion effect with incongruent mental imagery.

As in the previous analysis, to gauge the strength of evidence for these observations, data were simulated from the model and summary statistics, Bayes-p values and evidence ratios regarding the differences of interest were calculated. Following the findings of Chapter 2 (Experiments 1-3), it was predicted that the intention insertion effect should be greater when the recorded finger matched the participant's imagined movement (i.e. imagery-congruent) than when it mismatched (i.e. imagery-incongruent). As before, this analysis focused only on the comparisons that capture these predictions, namely the two-way interaction of Cue-Finger Congruence, Imagery-Finger Congruence and corresponding pairwise comparisons. A

test of the overall interaction revealed very strong evidence ($ER = 65.7$) that the overall intention insertion effect was strongest when the finger being recorded was congruent, rather than incongruent, with the participants currently imagined movement, as predicted (Posterior Mean = 4.49, 95% PCI [.48, 8.55], Bayes- $p = .985$).

The next step was to evaluate the two-way interaction between Cue-Finger Congruence and Imagery-Finger Congruence with a series of step down tests. As expected, the analysis of the finger which participants currently imagined moving (Figure 16, top panel) revealed extremely strong evidence ($ER = 499$) of a substantial intention insertion effect (Posterior Mean = 4.1, 95% PCI [1.29, 6.97], Bayes- $p = .998$). It was anticipated that, with congruent imagery, presentation of action cues should induce stronger presses at the finger which matched the observed movement (i.e. cue-congruent) than when a similarly salient, but motorically neutral cue was presented. Accordingly, there was strong evidence ($ER = 19.8$) that action cues facilitated, rather than interfered with, finger presses at the participant's corresponding finger (i.e. cue-congruent) than when neutral (i.e. no movement) cues were presented (Posterior Mean = 2.38, 95% PCI [.38, 5.18], Bayes- $p = .952$). It was also predicted that action cues should induce weaker presses (i.e. interfere) with the finger which mismatched the observed movement (i.e. cue-incongruent), relative to neutral. Consistent with that hypothesis, there was moderate evidence ($ER = 7.85$) that action cues interfered with, rather than facilitated, participants responses at the non-corresponding finger (i.e. cue-incongruent) relative to neutral stimuli (Posterior Mean = -1.73, 95% PCI [-4.53, 1.09], Bayes- $p = .887$). Taken together, these results suggest that, on the finger that participants currently imagined moving, action cues both facilitated and interfered with participant's responses depending upon the compatibility between the observed movement and the recorded finger.

The analysis of the finger which participants were not currently imagining moving (Figure 16, bottom panel) provided only anecdotal evidence ($ER = .66$) that action cues modulated finger pressure induced at the participants corresponding and non-corresponding finger (Posterior Mean = .39, 95% PCI [-3.31, 2.48], Bayes-p = .399). Interestingly, the evidence ratio (i.e. less than one) implies that there is slightly more evidence for the alternative hypothesis, that cue-incompatible responses were stronger than cue-compatible responses. However, it must be emphasised that both the strength of evidence (i.e. evidence ratio) and effect size (i.e. posterior mean) of this apparent “reversal” of the intention insertion effect, are very small. Indeed, tests of the individual components of the intention insertion effect revealed anecdotal evidence ($ER = 3$) that action cues facilitated larger presses with the participants corresponding (i.e. cue-compatible) finger than neutral (Posterior Mean = .97, 95% PCI [-1.81, 3.82], Bayes-p = .75). Furthermore, there was moderate evidence ($ER = .21$) in favour of the alternative hypothesis, that action cues facilitated stronger presses with the participants non-corresponding (i.e. cue-compatible) finger than neutral (Posterior Mean = 1.36, 95% PCI [-1.49, 4.18], Bayes-p = .172). In other words, these results suggest that finger pressure did not vary substantially according to the compatibility of the observed movement and the finger of response when it was also incompatible with the participant’s mental imagery. Instead, action cues had a small and broadly facilitating effect on responses which were incompatible with participants’ imagined movement.

3.2.3. Discussion

The purpose of Chapter 3 was to develop a more sensitive behavioural measure of participants’ movements which captured *how* participants responded, not simply whether or not they did, as in the keyboard-based intention insertion task (Chapter 2,

Experiments 1-3). Chapter 2 provided evidence that the presentation of action cues during a mental imagery task caused participants to produce action slips at a rate determined by the similarity of the observed and imagined acts. However, the keyboard-based intention insertion task (Chapter 2, Experiments 1-3) relied on a rather coarse behavioural measure (i.e. binary keypress) and lacked an appropriate baseline condition. Therefore, the previous design was insufficient to determine whether action observation primarily supports or detracts from the generation of non-willed action slips following the onset of an unexpected, salient stimulus.

Experiment 4 aimed to both replicate the findings of the earlier experiments (Chapter 2) using a continuous behavioural measure (i.e. finger pressure) and to incorporate a neutral, no-movement condition. More specifically, the neutral condition aimed to distinguish between facilitative and interference effects in the interaction of movement observation and mental imagery on participants behaviour (i.e. the intention insertion effect). Importantly, Experiment 4 successfully captured action slips and replicated the compatibility effect reported in Chapter 2 (Experiments 1-3), despite using a very different and more sensitive procedure. At a basic level, the presentation of action cues during imagery induced larger finger presses than when no cues were present. In addition, the magnitude of the pressure spikes reflected the similarity between the observed and imagined movements, replicating the findings of Experiments 1-3 (Chapter 2). Extending the findings of Chapter 2, Experiment 4 also provided evidence that, with compatible imagery, the action cues induced additional activity in the action effect codes supporting the corresponding behaviour. Action cues enhanced motor activation when compatible with imagery (i.e. summation of the two sources of activation) and detracted from motor activation when incompatible with imagery (i.e. competition between the two sources of activation).

Experiment 4 was principally concerned with testing whether the action cues in the intention insertion task primarily facilitate, or interfere with, the action slips generated by the intention insertion task. The data revealed that the overall difference between the induced pressure at the cue-congruent and cue-incongruent fingers (i.e. the intention insertion effect) was strongest when the finger being recorded matched the participants' imagined movement, as predicted. Importantly, Experiment 4 demonstrated that, with respect to the finger that participants currently imagined moving, action cues (e.g. a left movement) induced larger presses at the corresponding finger (e.g. left index finger) than responses to neutral cues. This suggests that the movement component of the action cues elicited additional activity in the action effects shared with the corresponding behaviour, over and above activity induced by merely encountering a surprising stimulus when combined with congruent mental imagery. This finding strongly contrasts with research which claimed that all visuomotor priming effects are the result of interference (e.g. Ramsey et al., 2010), and clearly demonstrates that, with congruent imagery, action cues exerted an additive (i.e. facilitating) effect on the participants' corresponding behaviour, as predicted.

Furthermore, Experiment 4 showed that observation of a downward finger movement induced smaller presses with the cue-incongruent finger, relative to a neutral, on the finger that participants currently imagined moving. This suggests that action cues also *interfered* with the generation of non-willed action slips. This apparent interference effect could be taken as evidence of competition between rival sets of action effect codes induced in a top-down manner by imagery and a bottom-up manner by the action cues. On this view, the endogenous and exogenous sources of activation interfered with one another and caused a reduction in the summed activity of the action effects which corresponded with the imagery-congruent finger, resulting in a weaker overt movement than elicited by neutral cues. Taken together, these findings suggest that the spatial

component of the action cues (i.e. the laterality of the observed movement) is an important driver of the intention insertion effect. Moreover, these findings are consistent with the prediction, motivated by ideomotor theory, that an observed movement can elicit additional activity in the specific action effect codes which serve the corresponding movement (i.e. in addition to activity induced by endogenous imagery).

Following the findings of Experiments 1-3 (Chapter 2), it was predicted that the action cues would have minimal effect on pressure at the finger that participants were currently not imagining moving. This was predicted on the basis that any activation of imagery-incongruent response would be strongly inhibited by imagined movement of the alternative finger. Consistent with this prediction, the data provided no evidence of an intention insertion effect when the recorded finger was also imagery-incongruent and, instead, showed that action cues exerted only a general and weakly facilitating effect induced pressure. This finding implies that the tendency for visual input (i.e. an observed movement) to induce the corresponding behaviour depends upon prior activation of the relevant action effect via endogenous (i.e. top-down) imagery.

Two lines of evidence support this interpretation. First, when endogenous imagery was focussed on the alternative action (i.e. not the currently recorded finger), the pressure induced by action cues was only slightly stronger than the pressure induced by neutral cues. Second, when the recorded finger corresponded with participants imagined movement (i.e. imagery-congruent) but was incongruent with the observed movement (i.e. cue-incongruent), finger presses were only attenuated, rather than eliminated. This finding indicates that activity induced by imagery was able to overcome a rival set of action effect codes (e.g. activated via action observation) which served the alternative finger movement.

Taken together, these results suggest that imagery and perceptual input (i.e. action observation) do not simply have an additive effect on the activation of the corresponding behaviour. Instead, the effect appears to be *super-additive* (i.e. imagery and perceptual input interact). This is because finger pressure was only induced if both endogenous imagery and exogenous perceptual input drove the activation of shared action effects close to (or, indeed, over) the motor threshold, as predicted by threshold models based on an integration-to-bounds mechanism (e.g. Churchland, Kiani, & Shadlen, 2008; Schurger, Mylopoulos, & Rosenthal, 2016). Importantly, the results also demonstrate that action observation can both facilitate and, to a slightly lesser extent, interfere with the generation of action slips in combination with mental imagery. They are therefore in line with the assumption (see Chapter 1 and 2) that the evidence accumulated by such a mechanism may reflect precisely the activation of the action's intended consequence (i.e. action effects), as suggested by ideomotor models.

In the introduction, it was suggested that participants may be prone to lightly tapping their fingers on the keys instead of using mental imagery as a means of rehearsing the movement sequence. With the previous design (Chapter 2, Experiments 1-3), the binary keypress variable gave the impression that when no response was registered the participant had not moved. In fact, at the onset of a surprising cue, participants may have already been exerting significant pressure on the key under the imagery-congruent finger. As a result, it was predicted that if the action cues, rather than a physical rehearsal strategy, were driving the effect, the participants should press much harder with their imagined finger than their non-imagined finger following cued than non-cued beats. The data were entirely consistent with this prediction and provided extremely strong evidence that participants imagined responses were more than six times larger than presses captured following cued than non-cued beats. This finding

argues against participants' use of a physical rehearsal strategy as a viable explanation for the intention insertion effect observed in Experiment 4.

3.3. Conclusion

In conclusion, Experiment 4 replicated and extended the findings of Chapter 2 (Experiments 1-3) by demonstrating a robust compatibility effect with a more sensitive continuous behavioural measure. Importantly, Experiment 4 demonstrated that action cues presented during imagery can elicit additional activity in action effect representations, facilitating the activation of the corresponding behaviour (i.e. when imagery and perception converge on a common set of action effect codes) and interfering with the activation of the non-corresponding behaviour (i.e. when imagery and perception activate different action effect codes). In other words, the findings provide direct evidence that perceptual input can induce additional activity in action effect codes and drive the corresponding overt behaviour, rather than simply detracting from action effects evoked by endogenous processes. In addition, they provide evidence for a super-additive effect of imagery and observation on motor output: actions are *only* initiated when both endogenous imagery and exogenous perceptual input activate the corresponding action effect representation.

More broadly, the findings provide a useful insight into how people manage to interact with the world around them with considerable automaticity. Bottom-up perceptual input can trigger specific, strongly associated responses by driving activity in action effects, but only if these are simultaneously activated by endogenous processes. This also explains why sometimes, salient cues in the environment are able to elicit behaviour which might be contextually inappropriate, but that is broadly compatible

with one's goals. For example, a loud whistle of a boiling kettle might compel a person to pour its contents into a waiting mug, even if they had planned to wash it first.

At this stage, one interesting observation is that Experiments 1-4 have relied on the same basic stimuli; a set of white, male hands presented from an egocentric point of view. The stimuli were designed to closely correspond with how participants would visualise their fingers moving. The rationale was that when a participant imagined and observed an action from the same visual perspective, the confluence of activity would be focussed on a very similar set of action effects. However, the implicit assumption that increasing the correspondence between the participant's own hands and the virtual hands should enhance the intention insertion effect presents a problem for a straightforward interpretation of these data. If action observation and imagery rely on the same set of action effects, why doesn't the observation of one's own movements automatically trigger the corresponding behaviour, as implied by the intention insertion task? One possibility is that visual input which corresponds with our own actions is treated very differently to visual input which corresponds with another's actions and does not have the same capacity to engage the motor system. In order to test address this question, Chapter 4 (Experiments 5 and 6) combined a virtual hand illusion (related to the classical rubber hand illusion; Ma & Hommel, 2015) with the same pressure-based intention insertion task used in Experiment 4.

4. Chapter 4: Why don't we self-imitate? Current and Goal States in Sensorimotor Processing

4.1. Introduction

Chapter 3 replicated and extended the findings of Chapter 2 by demonstrating that observation of a finger movement elicits stronger presses with one's corresponding, rather than the non-corresponding finger. Importantly, this compatibility effect (or "intention insertion effect") was strongest – or was only observed – when the visually depicted movement also matched the participant's imagined response. Experiments 1-4 therefore provided strong evidence that the activation of an action's perceptual consequences, by endogenous (e.g. mental imagery) and exogenous means (e.g. an observed action), induces activity in the motor system in accordance with ideomotor principles.

It has been argued that observed and imagined actions can drive behaviour because direct associations between visual (i.e. sensory) representations and the corresponding motor representations are formed through the repeated, correlated experience of observing and executing one's own actions (Heyes, Bird, Johnson, & Haggard, 2005; Hommel, Müsseler, Aschersleben, & Prinz, 2001; Prinz, 1990; for review, see Shin, Proctor, & Capaldi, 2010). In this way, both observed and imagined actions can effectively act as goal states for the motor system, triggering behaviour directly given a sufficiently strong association. However, if action observation and action initiation involve the same sensorimotor representation then why, when we observe our own behaviour, doesn't this process continue to drive the motor system, resulting in an endless positive feedback loop of repetitive imitative behaviour? The purpose of Chapter 4 is to combine the intention insertion task with a body ownership illusion (Experiments 5 and 6) in order to address this question.

In order to account for the distinct lack of self-imitation in everyday life, it is useful to consider how this observation could be accommodated by modern extensions of ideomotor theory with a specific focus on the computational principles that support voluntary action control (e.g. Theory of Event Coding, Hommel, 2009; Inverse/Forward Models, Wolpert, 1997; Active Inference, Adams et al., 2013). These models (hereafter, collectively referred to as “ideomotor models”) make two related, but often implicit, assumptions which could explain why observation of one’s own behaviour does not appear to engage the motor system in the same way as observation of another person’s actions. First, these models assume that sensorimotor processes make use of two, relatively independent, types of representation. One type of representation captures the *current state* of the sensorimotor system, and another describes a predicted – or intended – future *goal state* of the sensorimotor system. Second, goals are translated into action when the comparison between one’s current state and an activated goal state representation reveals a strong discrepancy which can be resolved by engaging the motor system (Hommel, 2009). In this conceptualisation, the role of the motor system is to resolve discrepancies between sensorimotor representations by affecting bodily transitions from the estimated current state towards target goal states.

Intuitively, the notion that actions are triggered when sensorimotor processes signal a discrepancy between the current state and an activated goal state corresponds nicely with the findings of the intention insertion task (Chapter 3, Experiment 4). Participants were instructed to imagine a sequence of finger movements while refraining from physical movement. In ideomotor terms, the imagery phase of each trial required participants to activate the relevant sets of action effect codes in order to build a vivid mental image (i.e. sensorimotor goal state) of the simulated sensory effects one would anticipate if each action in the sequence was actually executed. At the same time, participants were required to monitor their current state representation (informed by

visual and proprioceptive feedback) in order to ensure that the goal (i.e. the imagined finger movement) was not inadvertently triggered. At the moment of cue onset, the visual stimulus (i.e. observed finger movement) induced additional activity in the action effect codes that were shared with the goal state representation (already activated via imagery). As a result, the heightened activation of the goal state representation signalled a strong discrepancy with the current state, which was resolved by executing a downward finger movement in accordance with the observed movement. In other words, the sensorimotor system resolved the discrepancy between the two sensorimotor representations by initiating a transition from the current state towards the target goal state.

Ideomotor models suggest that people are not prone to self-imitation because the visual input that corresponds with the observation of one's own behaviour, by definition, describes the current state of the sensorimotor system (e.g. limb position). Visual information which corresponds with one's own body will, therefore, be integrated with the current state representation, rather than treated as a potential goal state. As a result, the comparator mechanism will no longer need to resolve any discrepancy between sensorimotor representations (i.e. there will be no sufficiently activated goal state) and the motor system will not be engaged. In contrast, when one voluntarily evokes (i.e. imagines) a goal state or when one observes the actions of another person, the same visual representation will not provide information relevant to the current state representation and instead will be treated as a candidate goal state with the capacity to engage the motor system because it differs from the current state.

Taken together, these assumptions predict that observing one's own actions will not engage the motor system in the same way as observing another's actions, even if the visual input is identical. Experiments 5 and 6 of the present chapter aim to test this prediction by incorporating a body ownership illusion into a new version of the

intention insertion task (Chapter 3, Experiment 4) and examining the performance of participants relative to a control group without the illusion.

Experiment 5 directly tests the hypothesis that, if an unexpected virtual movement in the intention insertion task is perceived (due to an illusion of ownership) to have been caused by oneself then the intention insertion effect should be reduced, relative to a control group without the illusion. Suggestive evidence in support of this prediction comes from a study by Schütz-Bosbach, Mancini, Aglioti, and Haggard (2006) which combined a classical rubber hand illusion with transcranial magnetic stimulation (TMS) to measure corticospinal excitability (an indicator of motor activation) during action observation. The study found that observing the actions of a rubber hand that was *not* attributed to one's body increased corticospinal excitability, suggesting an increase in motor activity and mirroring of the observed action. In contrast, this effect disappeared when the movements of the artificial limb were illusorily attributed to oneself. While the authors linked this effect to the motor system being able to distinguish self from other, it is also entirely consistent with the current state/goal state comparison implicitly assumed by ideomotor models. When the visual representation of the observed rubber hand movement was integrated with the current state, rather than activating the corresponding goal state, sensorimotor processes no longer detected conflict and the motor system was not engaged.

4.2. Experiment 5: The Effect of a Body Ownership Illusion on the Generation of Non-willed Action Slips

In order to test the predictions motivated by the ideomotor model, the intention insertion task developed in Chapter 3 (Experiment 4) was modified to incorporate a

body ownership illusion with respect to the virtual hands. The classical rubber hand illusion (Botvinick & Cohen, 1998) demonstrates that one's current state representation shows a considerable degree of plasticity and can be readily disrupted. In a typical experimental setup, participants observe a brush stroke a rubber hand while receiving identical tactile stimulation to their own unseen hand. When the experimenter strokes both the rubber hand and the participant's hand at the same time (i.e. in synchrony) it can result in the illusory perception that the artificial limb is part of the body. In contrast, asynchronous tactile stimulation does not cause an illusory feeling of ownership over the rubber hand. It was suggested that the illusion occurs because the sensory feedback driven by visual stimulation is tightly correlated with tactile and proprioceptive signals. This process is commonly referred to as *intermodal matching* and it has been proposed to reflect the mechanism by which people are able to generate a coherent representation of the current state of their sensorimotor system (Armel & Ramachandran, 2003; Botvinick & Cohen, 1998). Recently, the classical paradigm has been extended to demonstrate that people develop similar illusory feelings of ownership with respect to a virtual hand that moves in synchrony with their own (Ma & Hommel, 2015b) and even controllable non-corporeal objects such as a balloon which changes in size or a geometric shape which changes in colour (Ma & Hommel, 2015a). For the present purposes, a virtual hand illusion provided a convenient means to enhance participant's feelings of ownership with respect to the virtual hands in Experiment 5.

To illustrate how the illusion was incorporated, recall two important elements of Experiment 4 (Chapter 3). First, participants completed a brief "pressure training task" prior to beginning the intention insertion task. The pressure training task required participants to press the keys and maintain a series of reference pressures for five seconds at a time in order to become acquainted with the equipment. Second, following the initial presentation of the finger tapping sequence (e.g. "watch phase"), participants

were required to *perform* the sequence for four beats (e.g. one complete repetition of the sequence), prior to the imagery phase of each trial. In Experiment 5, the pressure training task and performance phase (during which movement was permitted) were altered so that participant's keypresses either triggered movements of the corresponding virtual finger or had no effect on the virtual fingers. According to Ma and Hommel (2015b), repeated exposure to a set of virtual hands which move in synchrony with one's own is sufficient to induce a feeling of illusory body ownership. In Experiment 5, participants were assigned to one of two groups; a heightened ownership group, who completed the intention insertion task with an ownership illusion, or a control group, without the illusion. Importantly, the virtual fingers only moved in response to participants' own presses during an acclimatisation task at the beginning of each block (i.e. the pressure training exercise, as in Experiment 4, Chapter 3) and at the beginning of each trial when participants were asked to physically execute the sequence (i.e. the "Perform" phase, as in Experiment 4, Chapter 3). That is, the virtual fingers did not respond to participants' movements during the imagery phase of each trial. This ensured that stimulus presentation (e.g. the surprising visual cues) and the conditions which immediately preceded it, were identical in both groups.

There is some debate in the literature concerning body ownership illusions about the extent to which the effect is influenced by the perceptual similarity of the artificial effector and the participants' own limbs (e.g. Tsakiris, Carpenter, James, & Fotopoulou, 2010, or for an opposing view, Ma & Hommel, 2015b). In order to sidestep this problem, a new stimulus set was created in which the virtual hands wore blue gloves and the arms were covered with long white sleeves. These stimuli were designed to be neutral with respect to participants own hands by removing skin tone and gender cues. This change aims to minimise any unintended influence of perceptual similarity on the strength of the body ownership illusion. In addition, in Experiment 5, the pressure

training task was repeated prior to each of the four experimental blocks for two reasons. First, the task served to emphasise task instructions and ensured participants returned both fingers to a neutral resting position prior to starting each new block. Second, it presented more opportunities for participants in the heightened ownership condition to form a strong association between their keypresses and the movements of the virtual hands.

If Experiment 5 demonstrates that the ownership illusion modulates whether combined imagery and action observation drives the motor system, then this finding would also help distinguish the ideomotor explanation of the effects in Chapter 3 from competing explanations. These competing predictions are motivated by research which specifically aims to account for the influence of movement observation on overt behaviour, Associative Sequence Learning (ASL, Heyes & Bird, 2008; Heyes & Ray, 2000). ASL assumes that visual (i.e. sensory) representations of one's own actions form direct links with the motor commands required to generate them, via Hebbian learning. As soon as an action is observed, it can therefore – via the established links – directly elicit the associated motor commands and give rise to involuntary behaviour. Importantly, unlike ideomotor models, the links between sensory and motor representations are not mediated by higher level representations, of current and goal states. As a result, if the translation of visual input into motor output simply relies on general associative mechanisms, the proposed ownership illusion should not reduce the intention insertion effect.

In fact, it is conceivable that the intention insertion effect may even be enhanced by the ownership illusion. In ASL, associations are formed whenever our actions reliably cause perceptual effects (see “response-effect compatibility”; Keller & Koch, 2006; Kunde, Koch, & Hoffmann, 2004). The ownership illusion itself creates a situation in which one's own actions (a finger press) creates such a reliable consequence

(an observed keypress on the screen) that it should therefore strengthen the association between the participant's keypress (i.e. response) and the corresponding movement of the virtual hand (i.e. effect). If participants' responses reflect such visuomotor association, then the body ownership training should, if anything, heighten the extent to which the visual representation activates the motor representation. As a consequence, the body ownership illusion should then actually *enhance* – not reduce - the intention insertion effect, relative to a control group.

4.2.1. Method

4.2.1.1. Participants

85 participants were tested, all of whom were female (72 Right-handed; Age in years: $M=23.6$, $SD=5.3$). The sample was comprised of both undergraduate psychology students from the University of Plymouth (course requirement) and members of the university's paid participation pool (remunerated £8).

4.2.1.2. Materials

The experimental hardware was identical to Experiment 4 (Chapter 3) and the software was lightly modified to allow participants in the heightened ownership condition to control the virtual fingers. A new stimulus set (Figure 17) was created in which the virtual hands appeared in blue gloves with long white sleeves in order to mitigate the influence of gender or skin tone cues. The same stimuli were used for both the heightened ownership and control condition.




Figure 17. Three examples from the updated stimulus set used in Experiment 5. The left image depicts two gloved hands in a neutral position with a central fixation cross. The middle image depicts a downward movement of the left index finger of the virtual hand which was presented to participants in the heightened ownership condition when they pressed the corresponding key during the performance phase of each trial. The right image depicts a salient right index finger movement cue presented during imagery to participants in both the control and heightened ownership condition.

4.2.1.3. Procedure

As in Experiment 4 (Chapter 3), the trial structure was such that participants first watched the four beat finger tapping sequence indicated by an arrow, then “performed” the four beat tapping sequence once, then “imagined” performing the sequence in time with a metronome until they encountered the predefined “Go!” stimulus. Experiment 5 was the same as Experiment 4 in most respects but differed in three key areas. First, participants were assigned to one of two ownership conditions. In the heightened ownership condition, keypresses triggered a downward movement of the corresponding virtual finger during the “perform” phase of each trial. In the control condition, the virtual fingers did not move in response to participants’ keypresses. Second, participants completed a brief pressure training task prior to the beginning of each experimental block. Participants were required to maintain a series of reference pressures for five seconds at a time before they were allowed to progress to the next block. This served to both encourage participants to reset their finger position prior to each block of trials and to highlight the sensitivity of the apparatus. Third, in the heightened ownership condition, participants’ keypresses also triggered movements of the virtual hands during the pressure training task in order to further strengthen this association.

4.2.2. Results

4.2.2.1. Exclusion Criteria

As with Chapter 2 and 3, participant accuracy during control trials was used as an index of task engagement. Eighteen participants with accuracy rates below one standard deviation of the group mean (chance=50%; $M=85.3\%$; $SD=12.5\%$) were excluded (11 Control, 7 Heightened Ownership), leaving $N = 67$ to be analysed. All data were pre-processed as described in Experiment 4 (Chapter 3), and induced finger pressure was calculated by deriving the area under the curve (AUC) in the 800ms interval between cue onset and the subsequent metronome beat.

4.2.2.2. Analysis of Finger Pressure During Cued Intervals

This analysis modelled the magnitude of pressure induced at each finger when participants were presented with different types of visual cues during imagery. Following Experiment 4 (Chapter 3), it was anticipated that induced pressure would be greater when the recorded finger (e.g. left index) matched the observed movement (e.g. virtual left movement) than when it mismatched (e.g. virtual right movement) - the *intention insertion effect*. Further, this effect should be greater when the recorded finger (e.g. left index) corresponds with the participant's imagery (e.g. imagined left index press) than when participants are imagining a movement of the other finger (e.g. imagined right index press). Extending Experiment 4 (Chapter 3), this analysis evaluated whether the intention insertion effect would vary as a result of the ownership illusion. The model outcome was the magnitude of induced finger pressure in the 800ms critical intervals after an action cue was presented. Finger pressure was recorded from

both the left and right index fingers in each of the 48 cued intervals. The model was an identical specification to the one used in Experiment 4 (Chapter 3), with the addition of another predictor (Ownership) which encoded the presence/absence of the ownership illusion. For clarity, the predictors of interest are capitalised when they are referred to in the text.

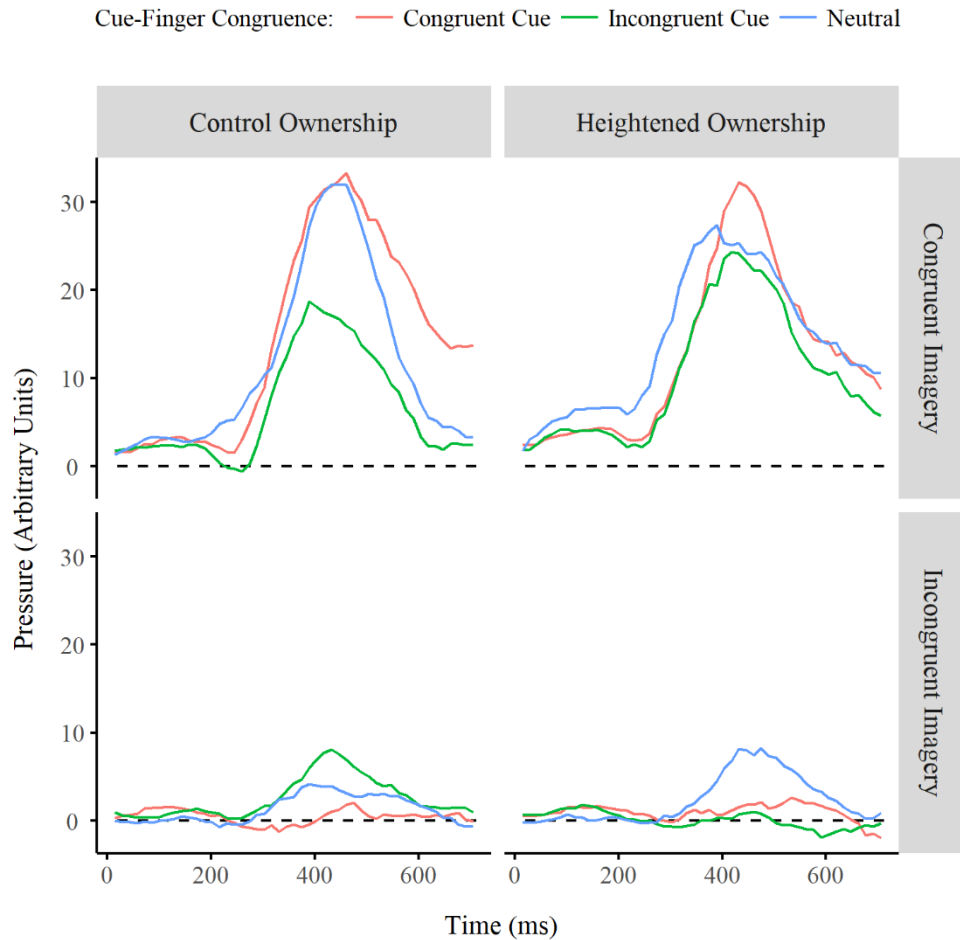


Figure 18. A multi-panel plot showing finger pressure in the critical 800ms interval following the onset of the onset of each type of visual cue in Experiment 5. The top and bottom panels separate responses from fingers which did or did not match the imagined response. The left and right portions of each panel separate responses from the control and heightened ownership groups. The coloured lines indicate whether the cue contained no movement (i.e. neutral – blue), or if it included an animation of the same (i.e. congruent – red)/other finger (i.e. incongruent – green) from that being recorded.

Figure 18 shows the magnitude of pressure that was induced in the critical 800ms period which followed the onset of a visual cue with respect to Ownership, Cue-Finger

Congruence and Imagery-Finger Congruence. Visual inspection suggests that pressure after a cue was much higher if participants also imagined a movement of the recorded finger (i.e. imagery-congruent; Figure 18, top panel) than if they imagined a movement of the other finger (i.e. imagery-incongruent; Figure 18, bottom panel). In addition, this was the case for both the heightened and control ownership groups. Importantly, the pressure induced by the visual stimuli appears to vary substantially according to whether the cue was neutral (i.e. no movement), or contained a movement which was congruent or incongruent with the finger producing the response.

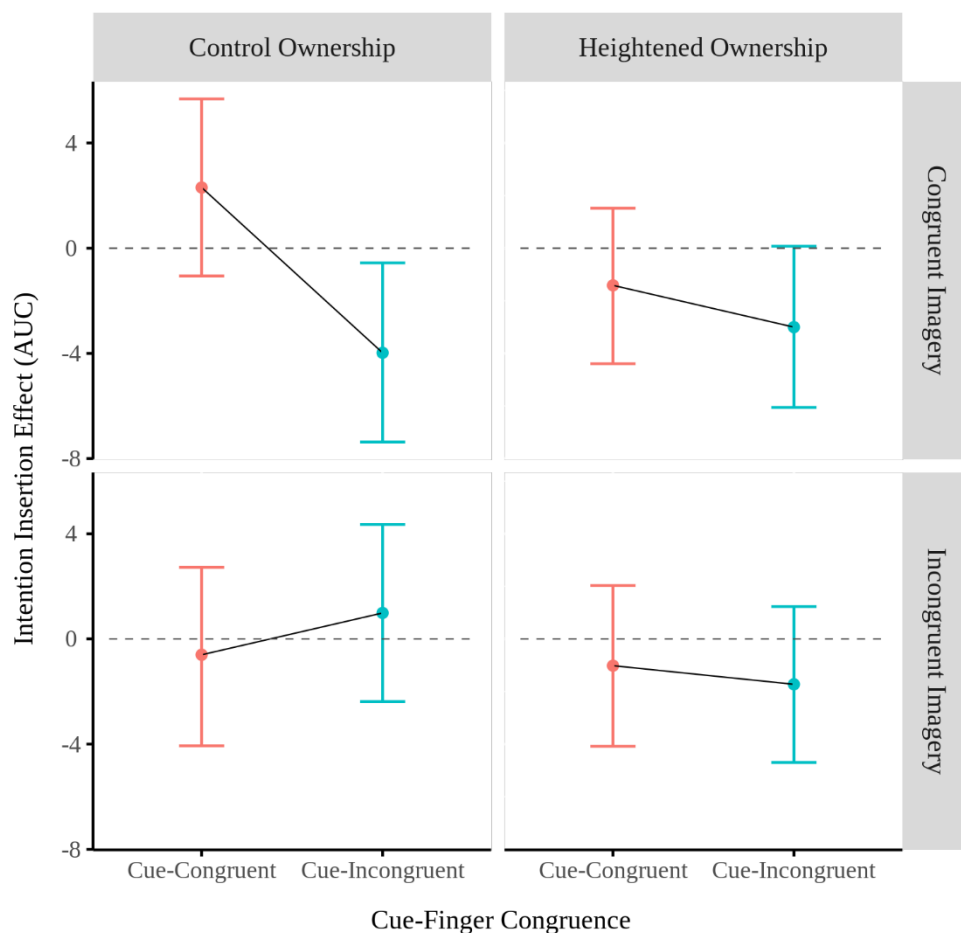


Figure 19. A multi-panel plot showing differences in pressure (AUC) during intervals when the recorded finger was either congruent or incongruent with respect to the observed movement vs. a neutral cue in Experiment 5. Positive scores indicate participants pressed harder than in equivalent trials where a neutral cue was presented (i.e. implying facilitation). Negative scores indicate that participants pressed less hard than following neutral cues (i.e. implying interference). The top and bottom panels distinguish responses where the participant was/was not also required to imagine a

movement in the recorded finger. The left and right portions of each panel separate responses from the control and heightened ownership groups. Error bars are the 95% PCI.

As in Experiment 4 (Chapter 3), in order to derive a summary measure of induced pressure, the area under the curve (AUC) was calculated for individual observations within each condition and fitted values from the model were plotted in Figure 19. Figure 19 shows the magnitude of induced pressure when the recorded finger was either congruent or incongruent with an observed movement, as compared with pressure after a neutral cue. In this model, the magnitude of the intention insertion effect was indexed by contrasting the pressure induced when the participant's finger matched the observed movement (i.e. cue-congruent) with the pressure induced when the finger mismatched the observed movement (i.e. cue-incongruent). Furthermore, the intention insertion effect was derived separately for instances in which the recorded finger corresponded with the participant's currently imagined movement (i.e. imagery-congruent) and instances in which the participants imagined a non-corresponding movement (i.e. imagery-incongruent). For example, the top-left panel of Figure 19 shows the intention insertion effect for the finger that participants in the control group were currently imagining moving. For this group, participants appear to press harder when the recorded finger is cue-congruent than cue-incongruent. Bigger differences between the two scores indicate a larger intention insertion effect. Visual inspection suggests that the intention insertion effect was greatest when participants also imagined a movement with the finger being recorded. It also appears the effect was larger in the control ownership group than the heightened group, as predicted by the ideomotor hypothesis. Although there appears to be a strong intention insertion effect for the control group when the recorded finger was congruent with participants' imagery, this effect is only small or even slightly reversed (e.g. Figure 19, bottom left panel) in the remaining three conditions. That is, there does not seem to be a strong intention insertion effect on the

imagery-incongruent finger in the control group, or on either finger in the heightened ownership group.

As in Experiment 4 (Chapter 3), in order to gauge the strength of evidence for these observations, data were simulated from the model and summary statistics, Bayes-p values and evidence ratios regarding the differences of interest were calculated. To briefly recap, the posterior predictive Bayes-p values associated with the forthcoming contrasts represent the probability (given the model and the data) that the effect was in the predicted direction (i.e. H_1), as opposed to the opposite direction (i.e. H_0). To aid interpretation, Bayes-p values are recapitulated as evidence ratios (ER; analogous with Bayes factors) to describe the strength of evidence in favour of the directed hypothesis, weighted against the alternative hypothesis. For consistency, the weight of evidence in support of each directed hypothesis was described using terminology regularly applied to Bayes factors (Wagenmakers et al., 2018); “extremely strong” (ER: >100), “very strong” (ER: 30-100), “strong” (ER: 10 - 30), “moderate” (ER: 3 - 10), “anecdotal” (ER: $>1 - 3$) and “no evidence” (ER: 1). Importantly, evidence ratios less than one indicate that the evidence supports the alternative hypothesis; even more so for values which are a tiny fraction of one. For example, an $ER < .01$ represents *extremely strong* evidence in favour of the *alternative hypothesis*.

It was predicted that the intention insertion effect would be reduced for participants in the heightened ownership group, relative to the control group, and that this difference would be strongest when the recorded finger matched the participant’s imagined movement (i.e. imagery-congruent), relative to the mismatching finger (i.e. imagery-incongruent). Importantly, for brevity and to safeguard against false positives due to multiple testing (Cramer et al., 2016), this analysis focused only on the comparisons that capture this prediction, namely the three-way interaction of Cue-Finger Congruence, Imagery-Finger Congruence and Ownership. Indeed, a test of this

interaction revealed very strong evidence ($ER = 61.5$) that the magnitude of the intention insertion effect in the presence of congruent, relative to incongruent imagery (see Chapter 3) was weaker (i.e. H_1), rather than stronger (i.e. H_0) in the heightened ownership group (Posterior Mean = 6.99, 95% PCI [.61, 13.31], Bayes-p = .984). The next step tested whether the ownership illusion had a general effect on the pressure induced by an observed movement, reducing the influence of the visual cues, independently of the participant's imagery (i.e. Imagery-Finger Congruence). This is captured by the contrast of the intention insertion effect for participants in the control group, relative to the heightened group, which is analogous with a two-way interaction between Ownership and Cue-Finger Congruence. This test revealed only anecdotal evidence ($ER = 3.31$) that the body ownership illusion reduced, rather than increased, the intention insertion effect in general (i.e. independent of the participants' imagined responses; Posterior Mean = 1.2, 95% PCI [-2.03, 4.42], Bayes-p = .768).

The next step was to evaluate the three-way interaction between Cue-Finger Congruence, Ownership and Imagery-Finger Congruence with a series of step-down tests. As expected, the analysis of pressure at the finger that participants currently imagined moving (Figure 19, upper panels) revealed a strong effect of body ownership on the intention insertion effect. The test revealed very strong evidence ($ER = 44.4$) for a larger difference between cue-congruent and cue-incongruent responses (i.e. intention insertion effect) in the control group than the heightened ownership group (Posterior Mean = 4.7, 95% PCI [.12, 9.27], Bayes-p = .978). Accordingly, in the control group (Figure 19, top left panel), there was extremely strong evidence ($ER = 999$) that visual cues that matched this (imagined) finger induced stronger presses than cues that mismatched the participants' imagined movements (Posterior Mean = 6.29, 95% PCI [2.95, 9.61], Bayes-p > .999). In contrast, in the heightened ownership group (Figure 19,

top right panel), there was only moderate evidence ($ER = 5.94$) for such an effect (Posterior Mean = 1.59, 95% PCI [-1.38, 4.59], Bayes- $p = .856$).

Figure 19 (lower panels) indicates that the intention insertion effect was much weaker with respect to pressure on the finger that did not correspond with participant's imagery (e.g. recording from the left finger when the participant imagined a right finger movement). A test of the difference between ownership groups provides moderate evidence in favour of the alternative hypothesis ($ER = .19$) that the difference between induced pressure at the cue-congruent and cue-incongruent finger (i.e. intention insertion effect) was stronger in the heightened than control ownership group (Posterior Mean = -2.3, 95% PCI [-6.88, 2.18], Bayes- $p = .157$). More specifically, there was moderate evidence ($ER = .21$) in favour of the alternative hypothesis that participants pressed slightly harder with the cue-incongruent finger than the cue-congruent finger in the control ownership group (Posterior Mean = -1.59, 95% PCI [-4.98, 1.79], Bayes- $p = .174$). Furthermore, there is only anecdotal evidence ($ER = 2.1$) to suggest that pressure induced at the cue-congruent finger was stronger than the cue-incongruent finger (i.e. the predicted intention insertion effect) in the heightened ownership group (Posterior Mean = .71, 95% PCI [-2.35, 3.76], Bayes- $p = .678$). Taken together, these tests suggest that the intention insertion effect was very small, or even slightly reversed when the response was incompatible with participants' imagery. However, it should be noted that effect sizes of these differences (i.e. the estimates of the posterior mean for the contrast) were much smaller than those observed with congruent mental imagery.

4.2.3. Discussion

Experiment 5 tested whether inducing an illusory sense of body ownership for a pair of virtual hands would influence its ability to engage the motor system with respect

to the finger that one currently imagines moving. Two potential outcomes were possible. First, ideomotor models claim that observation of one's own behaviour does not engage the motor system because the observed movement (and corresponding sensory representation) reflects the current state of the sensorimotor system, rather than a potential goal state. Participants in the heightened ownership group, who experienced a body ownership illusion should, therefore, exhibit a weaker intention insertion effect than those in the control group, who did not experience the illusion. In contrast, if the visual representation is linked directly to the motor representation, as assumed by models based on general associative processes, (e.g. ASL, Heyes & Bird, 2008; Heyes & Ray, 2000), then the intention insertion effect in the heightened ownership group should be the same as control, or even enhanced (e.g. due to response-effect compatibility, see Keller & Koch, 2006; Kunde et al., 2004). The results favoured the ideomotor view. The intention insertion effect was reduced in the heightened ownership group compared to the control group.

The predictions were informed by two related assumptions which are common, but not always explicit, in ideomotor models of action control (e.g. TEC, Hommel, 2009) and extensions of ideomotor principles (e.g. Inverse/Forward Models, Wolpert, 1997; Active Inference, Adams et al., 2013). The first assumption is that sensorimotor processes make use of two kinds of representations, one which captures the (estimated) current state of the sensorimotor system and another which describes predicted future goal states. The second assumption is that a comparator mechanism identifies discrepancies between the current state and goal state representations and resolves them by engaging the motor system and initiating a bodily transition (i.e. movement) towards target goal state. On this basis, it was predicted that heightening participants' sense of ownership with respect to the virtual hands would mean that the observed movements would be taken to reflect the current state, rather than a potential goal state. If the visual

representation of the observed movement was integrated with the current state, rather than the goal state representation, the intention insertion effect should be decreased, relative to a control group with no ownership illusion. As predicted, Experiment 5 revealed a much weaker intention insertion effect in the heightened ownership group, despite both groups being presented with identical visual stimuli, in accordance with predictions motivated by ideomotor models.

In addition, the data again confirmed that imagery is an important driver of the intention insertion effect because the differences in induced pressure were strongest with compatible imagery, replicating the findings of Experiment 4 (Chapter 3). In the ideomotor view, imagery may serve to pre-activate a goal state representation, to a sub-threshold level, which is then amplified by observing a movement of the corresponding finger on the virtual hand. This finding suggests that the sub-threshold activation of a goal state representation (via imagery) is a prerequisite for an externally activated visual representation (i.e. action observation) to trigger non-willed responses.

One interesting observation from Experiment 5 was that, in the control ownership group, the intention insertion effect appeared to be slightly reversed (e.g. Figure 19, lower left panel). Recall that in the intention insertion task, when responses are considered both cue- and imagery-incongruent, by definition, the participant's other finger must have been both cue- and imagery-congruent (i.e. the condition which induced the strongest presses). With the caveat that the effect size and evidence in support of this apparent reversal were small, one might speculate that this apparent reversal in the effect may reflect participants responding so strongly that they inadvertently respond with both fingers in this condition.

Although it is tempting to conclude that the elimination of the intention insertion effect in the heightened ownership group was caused by the observed action becoming

integrated with the current state, rather than the goal state representation, there is an alternative possibility. At cue onset, participants in the heightened ownership group may simply have been more inclined to visually check their own finger position when a cue appeared. This is because the heightened ownership group were led to believe they had control over the virtual hands and therefore would be more likely to believe they had caused the movement to occur. Given that overt movement was expressly prohibited, participants in the heightened ownership group may have checked to see if they had inadvertently responded and caused the movement they perceived on screen. This tendency for participants to attend to the fingers by visually monitoring their finger position (and thus, directing attention away from the virtual hands) may have been responsible for the reduction in the intention insertion effect in this group, rather than the predicted interaction of current and goal state representations. As a result, it was necessary to test a new sample of participants (Experiment 6) which used the same procedure as Experiment 5, but with their own hands occluded from view. If the same pattern of results is present – or if this pattern is enhanced – when participants are unable to visually monitor their own actions then it would rule out this alternative explanation for the reduction in the intention insertion effect in the heightened ownership group, relative to the control group.

4.3. Experiment 6: The Effect of a Body Ownership Illusion on the Generation of Non-willed Action Slips without Visual Feedback

In order to rule out that participants in the heightened ownership condition were simply more inclined to observe their own actions following cue onset, it was necessary

to test a new sample of participants with their hands occluded from view. By eliminating accurate visual feedback of participants' own limbs, participants in both groups could only rely on proprioceptive information, rather than vision, to monitor and control their behaviour. If the reduction in intention insertion effect observed in Experiment 5 was caused by participants in the heightened ownership group checking their own fingers following cue onset, then forcing both groups to rely on proprioception should eliminate this difference in Experiment 6. Hence, if Experiment 6 produces strong evidence of a difference between ownership groups when visual feedback is eliminated, then this will rule out the alternative explanation.

If the between-group difference in the intention insertion effect in Experiment 5 really did reflect a change in the contribution of visual input to current and goal state representations then, according to accounts based on Bayesian estimation (e.g. Active Inference, Adams et al., 2013), eliminating visual feedback from participants' own fingers may, in fact, enhance the effect. Accounts of motor control based on Bayesian estimation emphasise that sensory input is inherently noisy and, at any given moment, one can only estimate (i.e. predict) the current position of the body in space. Furthermore, the precision of the estimated current state representation is strongly influenced by the modality of the sensory input and environmental constraints (Körding & Wolpert, 2004, 2006). For example, waking up on a dark morning, one faces considerably more uncertainty in the estimation of the current state because of the lack of visual feedback and the corresponding reliance upon proprioception to navigate to the light switch. In Experiment 5, participants were free to observe their hands at any time, although they were specifically instructed to direct their visual attention towards a central fixation cross for the duration of the experiment. Therefore, for participants in either ownership group, there was ready access to visual information which could be used to accurately estimate the current position of their hands and control their non-

willed action slips. According to this view, visual input should retain some basic degree of relevance to the current state representation in both ownership groups, which might affect the contribution of action observation to the activation of the corresponding goal state representation.

If visual input makes a dynamic contribution to both the current and goal states, weighted by environmental constraints, then eliminating visual feedback for one's own movements may make visual information, in general, less relevant to the current state representation. If visual input can no longer provide information about the current state of one's body, then action effects activated by visual stimulation may make a correspondingly larger contribution to the goal state representation. This would result in a general increase in the intention insertion effect in both ownership groups (Experiment 6), relative to when participants' hands were not occluded (Experiment 5). In order to rule out the alternative explanation for the reduction in the intention insertion effect and to test assumptions about the dynamic weighting of visual information to current and goal states, Experiment 6 tested a new sample of participants whose hands were occluded from view.

4.3.1. Method

4.3.1.1. Participants

A total of 80 participants were tested, all of whom were female (70 Right-handed; Age in years: $M=21.5$, $SD=5.2$). The sample was comprised of both undergraduate psychology students from the University of Plymouth (course requirement) and members of the university's paid participation pool (remunerated £8).

4.3.1.2. Materials and Stimuli

The stimulus set, hardware and software were identical in all respects to those used in Experiment 5, with the exception that a large box was used to cover both the apparatus and participants hands in order to prevent visual monitoring of their physical movements. In addition, a fabric shroud was attached to the top of the box and covered participants' forearms to completely eliminate visual feedback of their movements.

4.3.1.3. Procedure

Participants were assigned to one of two ownership groups, as in Experiment 5, and were presented with an identical experiment. The only difference was that with this new cohort, a large box blocked the view of the participants' own hands.

4.3.2. Results

4.3.2.1. Exclusion Criteria

As previously, participant accuracy during control trials was used as an index of task engagement. Twelve participants with accuracy rates below one standard deviation of the group mean (chance=50%; $M=82.3\%$; $SD=17.9\%$) were excluded (7 Control, 5 Heightened Ownership), leaving $N = 68$ to be analysed. All data were pre-processed as described in Experiment 5, and finger pressure was calculated by deriving the area under the curve (AUC) in the 800ms interval between cue onset and the subsequent metronome beat.

4.3.2.2. Analysis of Finger Pressure During Cued Intervals

The first analysis was concerned with modelling how hard participants pressed with each finger when presented with action cues. Following Experiment 4 (Chapter 3) and Experiment 5 (present chapter), it was anticipated that induced pressure would be greater when the recorded finger matched the observed movement than when it mismatched (i.e. the intention insertion effect). In addition, the intention insertion effect should be greater when the recorded finger corresponds with the participant's imagery than when participants are imagining a movement of the other finger. Furthermore, it was predicted that the intention insertion effect would vary as a result of the body ownership illusion (i.e. replicating the results of Experiment 5). Importantly, it was predicted that the ownership illusion should have the same effect in spite of the fact that participants' hands were occluded from view. This finding would rule out that the effect of the ownership illusion (Experiment 5) was to draw participants' visual attention towards their own hands and away from the virtual hands. As in Experiment 5, the model outcome was finger pressure in the 800ms critical intervals after an action cue was presented and the model specification was identical.




Figure 20. A multi-panel plot showing finger pressure in the critical 800ms interval following the onset of the onset of each type of visual cue in Experiment 6. The top and bottom panels separate responses from fingers which did or did not match the imagined response. The left and right portions of each panel separate responses from the control and heightened ownership groups. The coloured lines indicate whether the cue contained no movement (i.e. neutral – blue), or if it included an animation of the same (i.e. congruent – red)/other finger (i.e. incongruent – green) from that being recorded.

Figure 20 shows the induced pressure in the critical 800ms period which followed the onset of a visual cue according to Ownership, Cue-Finger Congruence and Imagery-Finger Congruence. It appears that induced pressure following all types of visual cue was much higher if participants were also imagining a movement of that finger and this was true for both the heightened and control ownership groups. In addition, induced pressure appears to vary substantially according to whether the visual stimulus was neutral, or contained a virtual movement which was congruent or incongruent with the finger which produced the response. This replicates the performance pattern observed in

Experiment 5. Interestingly, Figure 20 indicates that, in absolute terms, participants whose hands were occluded exerted a higher peak pressure in all conditions, relative to Experiment 5 (Figure 19).




Figure 21. A multi-panel plot showing differences in pressure (AUC) during intervals when the recorded finger was either congruent or incongruent with respect to the observed movement vs. a neutral cue. Positive scores indicate participants pressed harder than in equivalent trials where a neutral cue was presented. The top and bottom panels distinguish responses where the participant was/was not also required to imagine a movement in the recorded finger. The left and right portions of each panel separate responses from the control and heightened ownership groups. Error bars are the 95% PCI.

Figure 21 shows the pressure induced when the recorded finger was either congruent or incongruent with an observed movement, as compared with pressure after a neutral cue. As in Experiment 5, the magnitude of the intention insertion effect was indexed by contrasting the pressure induced when the participant's finger matched the

observed movement (i.e. cue-congruent) with the pressure induced when the finger mismatched the observed movement (i.e. cue-incongruent). In addition, the intention insertion effect was derived separately for instances in which the recorded finger corresponded with the participant's currently imagined movement (i.e. imagery-congruent) and instances in which the participants imagined a non-corresponding movement (i.e. imagery-incongruent). Larger differences between the cue-congruent and cue-incongruent finger pressure indicate a larger intention insertion effect. Visual inspection suggests that the intention insertion effect was greatest when the participant was also imagining a movement with the finger being recorded (Figure 21, top panels). Importantly, there appears to be a strong intention insertion effect for both the control group and the heightened ownership group on the imagery-congruent finger, but the effect is larger in the control group. In addition, when the finger of response was incongruent with participants' imagery (Figure 21, bottom panels), there appears to be strong intention insertion effect in the control group, but not in the heightened ownership group.

As in Experiment 4 (Chapter 3) and 5, in order to gauge the strength of evidence for these observations, data were simulated from the model and summary statistics, Bayes-p values and evidence ratios regarding the differences of interest were calculated. Following Experiment 5, it was predicted that the intention insertion effect should be reduced in the heightened ownership group, relative to the control group, and that this difference should be greater when the recorded finger matched the participants imagined movement (i.e. imagery-congruent), relative to the mismatching finger (i.e. imagery-incongruent). As before, this analysis focused only on those aspects of the data which are relevant to this prediction, namely the interaction of Cue-Finger Congruence, Imagery-Finger Congruence and Ownership. However, a test of the overall interaction provided only moderate evidence ($ER = 4$) of a slightly larger effect of Ownership on

the intention insertion effect with congruent, relative to incongruent imagery (Posterior Mean = 2.62, 95% PCI [-3.66, 8.8], Bayes-p = .799). This suggests that, in contrast to Experiment 5, the effect of the body ownership illusion on the intention insertion effect did not greatly differ when the finger being recorded matched the participants imagined movement (i.e. imagery-congruent), relative to when it mismatched (i.e. imagery-incongruent). The next step was to evaluate the general effect of the body ownership illusion on pressure induced by visual cues, independently of Imagery-Finger Congruence. This test involves a contrast of the intention insertion effect for participants in the control group, relative to the heightened group, which corresponds with the two-way interaction of Ownership and Cue-Finger Congruence. In contrast to Experiment 5, this test revealed extremely strong evidence (ER= 999) that the body ownership illusion caused a reduction (rather than an increase) in the intention insertion effect, independently of the participants' imagined movement (Posterior Mean = 4.88, 95% PCI [1.79, 8.06], Bayes-p = .999). This suggests that, when participants' hands were occluded, the intention insertion effect was generally stronger in the control group, relative to the heightened ownership group, regardless of the participant's imagined movement.

Following the same procedure as Experiment 5, participant performance was evaluated with planned step-down comparisons of the interaction between Ownership, Cue-Finger Congruence, and Imagery-Finger Congruence. As predicted, and replicating Experiment 5, the analysis of pressure on the finger that participants currently imagined moving (Figure 21, upper panels) revealed a very strong effect of body ownership on the intention insertion effect. A pairwise comparison between the intention insertion effects in the two groups provided extremely strong evidence (ER = 332) that the effect was much larger in the control group than the heightened ownership group (Posterior Mean = 6.19, 95% PCI [1.73, 10.48], Bayes-p = .997). In other words, the evidence

ratio suggests that there was 332 times more evidence that the intention insertion effect was reduced, rather than enhanced, by the ownership illusion when participant imagery was congruent with the recorded finger (i.e. imagery-congruent). Accordingly, in the control condition (Figure 21, upper panels), there was extremely strong evidence ($ER > 999$) of a substantial intention insertion effect when the recorded finger matched the observed movement (i.e. cue-congruent) compared to the finger which mismatched (i.e. cue-incongruent; Posterior Mean = 11.09, 95% PCI [7.93, 14.21], Bayes- $p > .999$). In contrast to Experiment 5, there was also extremely strong evidence ($ER = 999$) of a large intention insertion effect in the heightened ownership group (Figure 21, top right panel; Posterior Mean = 4.9, 95% PCI [1.87, 7.98], Bayes- $p = .999$). This provides suggestive evidence that the intention insertion effect was generally enhanced for both groups when the participant's hands were occluded, relative to when they were visible (Experiment 5).

In support of that observation, the analysis of pressure at the finger that participants did not currently imagine moving (Figure 21, lower panels) also revealed an effect of the body ownership illusion on the intention insertion effect. The pairwise comparison of the intention insertion effect size in each ownership group revealed strong evidence ($ER = 16.9$) of a larger intention insertion effect for the control than heightened condition (Posterior Mean = 3.57, 95% PCI [.79, 8.05], Bayes- $p = .944$). In other words, the evidence ratio suggests that there was 17 times more evidence in favour of the hypothesis that the intention insertion was reduced, not increased, by the ownership illusion when the finger being recorded was incongruent with participants' imagery (i.e. imagery-incongruent). Within the ownership groups, tests revealed extremely strong evidence ($ER = 124$) of a substantial intention insertion effect in the control condition (Figure 21, lower left panel; Posterior Mean = 3.81, 95% PCI [.65, 6.96], Bayes- $p = .992$), but no evidence ($ER = 1.24$) of an intention insertion effect in

the heightened ownership condition (Figure 21, lower right panel; Posterior Mean = .24, 95% PCI [-2.81, 3.35], Bayes-p = .564).

4.3.2.3. Analysis of Experiments 5 and 6: The Effect of Visual Feedback

The next step was to conduct a combined analysis of the data from Experiments 5 and 6 in order to perform a between-experiments contrast to test predictions about the influence of eliminating visual feedback of participants' own movements on the intention insertion effect. The model has an identical specification to the one used in the cued trials analysis, with the addition of another predictor, Visual Feedback, which described whether participants hands were visible (Experiment 5) or occluded (Experiment 6). This allows us to test whether the strength of the intention insertion effect was enhanced when participants' hands were occluded from view, as predicted.

A test of the interaction between Visual Feedback and Cue-Finger Congruence revealed strong evidence (ER = 332) that occluding participants' hands, in general (i.e. independently of all other conditions), enhanced the intention insertion effect relative to when their own hands were visible (Posterior Mean = 3.2, 95% PCI [.98, 5.4], Bayes-p = .997). In addition, this model allows us to test whether eliminating visual feedback (Experiment 6) generally enhanced the intention insertion effect produced by each ownership group, in the presence of congruent or incongruent mental imagery (i.e. Imagery-Finger Congruence). A test of the interaction between Visual Feedback, Cue-Finger Congruence and Imagery-Finger Congruence revealed moderate evidence (ER = .19) in favour of the alternative hypothesis. This suggests that, counter to expectations, when participants' hands were occluded, the effect of body ownership on the intention insertion effect was stronger for the finger which mismatched the currently imagined movement (i.e. imagery-incongruent), relative to when it matched (i.e. imagery-

incongruent; Posterior Mean = -4.46, 95% PCI [-13.38, 4.42], Bayes-p = .161). This finding corroborates the notion that the occluding participants' hands generally enhanced the tendency for visual cues to engage the motor system, regardless of the participants imagined movement.

The next step was to compare participant performance in Experiments 5 (hands visible) and 6 (hands occluded) with a series of planned step-down comparisons of the interaction between Visual Feedback, Cue-Finger Congruence, Imagery-Finger Congruence and Ownership group. It was predicted that, if occluding participants' hands increased the tendency for an observed movement to activate the corresponding action, then the intention insertion effect should be stronger in both the heightened and control ownership group when the finger of response matched the participant's imagery (i.e. imagery-congruent). A pairwise comparison revealed only anecdotal evidence (ER = 2.03) that the effect of Visual Feedback was stronger for the control than heightened ownership group (Posterior Mean = 1.43, 95% PCI [-4.94, 7.87], Bayes-p = .67). More simply, the size of the intention insertion effect at the imagery-congruent finger was generally increased in both the control and heightened ownership groups when participants' hands were occluded, relative to when they were visible. Accordingly, the analysis of the responses which were congruent with participants' imagery revealed very strong evidence (ER = 40.7) that the intention insertion effect was increased, rather than reduced, when participants hands were occluded in the control ownership group (Posterior Mean = 4.69, 95% PCI [.08, 9.33], Bayes-p = .976) and strong evidence (ER = 13.3) of an increase in the heightened ownership group (Posterior Mean = 3.26, 95% PCI [-1.08, 7.45], Bayes-p = .93).

Next, the analysis of the responses which were incongruent with participants' mental imagery revealed that occluding participants hands modulated the strength of the intention insertion effect in each ownership group. If occluding participants hands

generally enhanced the tendency for action observation to activate the corresponding response, then it was predicted that the intention insertion effect should also be increased at the finger which was incongruent with participants' mental imagery. A pairwise comparison revealed very strong evidence ($ER = 32.3$) that the elimination of visual feedback caused a stronger increase in the intention insertion effect for the control than the heightened ownership group (Posterior Mean = 5.89, 95% PCI [.2, 12.05], Bayes- $p = .97$). In other words, the evidence ratio suggests that there was 32 times more evidence supporting the hypothesis that the change in the intention insertion effect (caused by occluding participants hands in Experiment 6, relative to Experiment 5) at the imagery-incongruent finger was stronger (as opposed to weaker) in the control group than the heightened ownership group. As expected, there was extremely strong evidence ($ER = 99$) that occluding participants hands substantially increased, rather than decreased, the size of the intention insertion effect in the control ownership group (Posterior Mean = 5.39, 95% PCI [.8, 9.86], Bayes- $p = .99$), but no such evidence ($ER = .69$) of an increase in the heightened ownership group (Posterior Mean = .49, 95% PCI [-4.69, 3.7], Bayes- $p = .41$).

4.3.3. Discussion

The main purpose of Experiment 6 was to rule out an alternative explanation for the reduction in the intention insertion effect in the heightened ownership group, relative to the control group, observed in Experiment 5. In the alternative view, because participants in the heightened ownership group were led to believe that they were responsible for the virtual movements, they may have been more inclined to visually inspect their own hands following cue onset than participants in the control group. As a result, the effect of the body ownership illusion may have simply been to draw participants' visual attention away from the virtual hand movement, limiting the extent

to which visual input was able to drive the motor system. In order to exclude this possibility, participants' hands were occluded from view in Experiment 6.

If the alternative explanation of the findings of Experiment 5 were valid, then preventing participants from visually checking their own finger position should reduce the difference in the magnitude of the intention insertion effect between ownership groups. In contrast, if the reduction in the intention insertion effect in the heightened ownership group was caused by the observed action being integrated with the current, rather than the goal state, as predicted by ideomotor models, then the pattern of results would be the same as Experiment 5. As expected, the data suggested that there was a much stronger intention insertion effect in the control than heightened ownership condition in Experiment 6, just as there was in Experiment 5. This finding not only rules out a viable alternative explanation but provides a valuable replication of the effect of an ownership illusion on the intention insertion effect.

In addition, occluding participants' view of their own hands in Experiment 6 was predicted to cause a general increase in the magnitude of the intention insertion effect relative to Experiment 5, when participants' hands were visible, in a between-experiment comparison. According to ideomotor models, the visual information must be appropriately integrated (i.e. bound) with either the current or goal state representation. This effect was predicted on the basis that occluding participants' hands would eliminate an important source of current state-relevant information (i.e. visual feedback of one's own body). Therefore, if visual input contains less current state-relevant information, the activated action effects may be more effectively integrated (i.e. bound) with goal state representation. The data indicated that, as predicted, the intention insertion effect was indeed substantially stronger without visual feedback of participants' own movements. Interestingly, the data also suggested that there was a strong intention insertion effect for responses which matched the participant's imagined

movement in the heightened ownership group of Experiment 6 which was absent in Experiment 5. That is, the occlusion of the hands overrode the impact of the ownership illusion, and revealed an intention insertion effect in the ownership condition in Experiment 6 where there wasn't one in Experiment 5. In fact, by occluding participants' hands from view, the virtual movement activated the corresponding response so strongly that, for the first time, there was evidence of an intention insertion effect at the finger which *did not* match the participants' imagery in the control condition (but not in the heightened condition). These findings can be readily accounted for by ideomotor models if one assumes that by reducing the current state-relevance of visual input *in general* (i.e. by eliminating visual feedback), the visual input should activate the corresponding goal state more strongly. In this view, action observation triggered the corresponding response so substantially in the control group that the exogenously activated goal state representation was even able to overcome the endogenously activated goal (i.e. via imagery) when participants imagined the alternative response (i.e. the two goal states differed).

4.4. General Discussion

This chapter used the pressure-pad intention insertion task (Experiment 4, Chapter 3) to test a commonly held, but largely theoretical, assumption about the role of the “self” in the translation of sensory input into motor output. Experiments 1-4 provided strong evidence that observing a downward finger press can effectively trigger the corresponding movement, when combined with congruent mental imagery, in spite of specific instructions to remain still. According to the ideomotor account of the intention insertion effect, action observation activates the same perceptually coded sensorimotor representations (i.e. action effects) that are responsible for executing the corresponding

action. The question that motivated this chapter was if this is the case, then why aren't we driven to endlessly self-imitate when observing our own behaviour? Experiments 5 and 6 of the present chapter suggest that we do not self-imitate because visual input which corresponds with one's own actions inherently reflects the current state of the sensorimotor system (i.e. it corresponds with the current location of the body in space), *not* a potential goal state and, therefore, does not engage the motor system.

Research has demonstrated that while action observation can elicit activity in the motor system, the effect disappears when the same observed action is illusorily attributed to oneself (Schütz-Bosbach et al., 2006). This suggests that body ownership illusions can cause people to process identical visual and proprioceptive inputs quite differently. Ideomotor models account for this phenomenon by assuming that, by definition, the visual input associated with observation of our movements reflects the current position of the body in space (i.e. the current state of the sensorimotor system), not a potential future state with the capacity to engage the motor system. Accordingly, Experiments 5 and 6 aimed to test whether an illusory perception of body ownership would reduce the extent to which observed movements drive the motor system during imagery (i.e. indexed by the intention insertion effect), as predicted by ideomotor theories of action control. To that end, Experiment 5 and 6 replicated the basic intention insertion effect observed in Experiment 4 (Chapter 3). Further, these experiments extended the effect by demonstrating that it was strongly modulated by factors such as a body ownership illusion induced by correlated sensory feedback (e.g. intermodal matching, Botvinick & Cohen, 1998) and the availability of sensory information (e.g. presence of visual feedback). More specifically, the intention insertion effect was substantially weaker for participants who received the body ownership illusion than participants in the control group in both Experiments 5 and 6, as predicted by ideomotor models.

The data reported in this chapter provided empirical support for the notion that sensorimotor processes make use of two, relatively separate, representations which capture either the current state of the sensorimotor system or predicted future goal states. It was predicted that if an observed action was taken to reflect one's current state (e.g. due to a body ownership illusion) then the corresponding sensory representation would not be processed as a goal state and would no longer engage the motor system. Consistent with this prediction, the data revealed that participants in the heightened ownership group showed a significantly weaker intention insertion effect than those in the control group. Importantly, this substantial difference was evident in spite of the fact that there was no perceptual difference between the virtual hand movements encountered by participants in either group and that the conditions immediately preceding stimulus presentation were identical. Importantly, this finding cannot be accommodated by accounts of action control which assume that visual representations trigger motor responses directly via general associative processes (e.g. ASL, Heyes & Ray, 2000). If the visual representation which corresponded with the observed action simply activated the associated motor representation directly then the intention insertion effect should have, if anything, been increased by the body ownership illusion that would strengthen any association between observed and own finger movements. The strong difference in the intention insertion effect between ownership groups suggests that visual input (i.e. an observed action) only activates the corresponding behaviour (i.e. goal state) if it is *not* considered self-generated. This argues against a direct coupling of the visual (i.e. sensory) representation and the motor representation and is, instead, more consistent with the current state/goal state distinction described by ideomotor models.

The finding that eliminating visual feedback increased the intention insertion effect supports the assumption that action initiation results from the interaction of

current and goal state representations. It was argued that occluding participants' hands would mean that that visual input could no longer provide useful information about the current state of the participants' fingers during the task. On this basis, it was predicted that this would cause the contribution of visual information to the current state representation to be downregulated (Körding & Wolpert, 2004, 2006). Correspondingly, it was predicted that the visual input that was available – the virtual hands - would make a larger contribution to the activation of the goal state representation which would heighten sensorimotor conflict and produce a larger intention insertion effect in general. The comparison of Experiment 5 and 6 showed that participant performance was entirely consistent with this conceptualisation of action control processes. The intention insertion effect was increased by occluding the participants' hands. Interestingly, one consequence of occluding participants' hands was that the observed movement activated the corresponding response so substantially that there was, for the first time, clear evidence of an intention insertion effect even at the finger which mismatched the participant's mental imagery. This is especially important because in Experiments 4 (Chapter 3) and 5 the virtual movements only triggered corresponding responses in the presence of matching mental imagery. Previously, this was taken to suggest that the pre-activation of a given sensory goal state via imagery (i.e. the simulated perceptual effects corresponding with the imagined sequences) might be a prerequisite for the intention insertion effect to occur. Under this view, observed actions could only amplify or interfere with existing goal states activated by imagery but are insufficient to trigger behaviour alone. Instead, the findings of Experiment 6 suggest that in the right circumstances, action observation can insert the corresponding response even when participants imagine a different movement. In other words, an exogenously activated goal state (i.e. observed action) would appear to be able to initiate the corresponding

response when it is sufficiently salient, even when participant's imagery was focussed on a different action.

At this stage, one might be tempted to suggest that some of the differences in the intention insertion effect observed between body ownership groups could have been due to differences in internal monitoring of behaviour. In other words, once participants realised that they were producing accidental actions, they should have tried to inhibit these movements. Further, participants in the heightened ownership group may have been more likely to engage in such monitoring because they had been led to believe their own actions *caused* the virtual finger movements. As a result, the onset of a sudden movement cue could have been taken as evidence that a key had been inadvertently pressed. Indeed, Experiments 1-3 (Chapter 2) provided clear evidence that participants became progressively less likely to produce action slips, which clearly highlights the role of action monitoring processes in this task. However, it must be emphasised that an increase in a more specific monitoring explanation is not necessarily antagonistic with the proposed notion of separate current state and goal state sensorimotor representations. This is because internal monitoring, by definition, involves scrutinising any detected changes to the current state of one's body. For example, in the ideomotor view, if the heightened ownership group monitored their actions more closely than the control group, then this enhanced scrutiny may have arisen because participants assumed that movements of the virtual hands represented a change in their current state, rather than a potential goal state. On this logic, the control group, who were less likely to assume that the virtual movements reflected a sudden and undesirable change in their current state, did not engage monitoring (i.e. scrutiny of current state) and the observed action was able to trigger the corresponding goal state representation.

An alternative possibility is that the experience of causing virtual finger movements in the heightened ownership group caused participants to limit their motor output in a general, non-specific manner (e.g. by raising the threshold required to initiate movement). Certainly, the pressure which was induced by the onset of unexpected visual stimuli was reduced by the ownership illusion, in both Experiments 5 and 6. However, a general increase in monitoring would predict a consistent downregulation of motor output following all kinds of the cue (i.e. neutral as well as congruent/incongruent movement cues) in the heightened ownership group, relative to the control group. Therefore, a general tendency to monitor motor output does not readily account for the specific decrease in the pressure induced by an observed movement at the cue-congruent finger, relative to cue-incongruent finger (i.e. the intention insertion effect).

One final point to consider the extent to which the implementation of the virtual hand illusion in order to manipulate ownership might complicate the interpretation of these data. It is possible that, by exposing participants in the heightened ownership group to a significantly higher number of visual finger movement stimuli, this group may have, to some extent, habituated to the visual stimuli. This would have the effect of reducing the effectiveness of the visual stimuli when they were presented during the subsequent imagery phase of each trial, relative to the control group who *only* encountered the finger movement cues during imagery. However, there are several points worth noting in this regard.

First, participants in the heightened ownership condition had *only* encountered such visual stimuli (i.e. finger movements) as a *consequence* of their actions (at least prior to the first cue) and, if anything, they should be especially surprised by (and subsequently attend to) the sudden onset of a finger movement during imagery. Indeed, the purpose of the experiment was to heighten the correspondence between the

participant's behaviour and those specific visual stimuli (i.e. virtual hand movements), before placing them in a situation where those same physical responses were explicitly prohibited in order to ensure participants avoided intentionally responding. On that basis, one would not expect that the finger movement stimuli should be *less* disruptive (i.e. less able to facilitate or interfere with action) simply because participants saw them before (albeit without the obnoxious yellow flash). In fact, the visual stimuli served as a potent signal that the very behaviour they were trying to avoid might have just occurred (i.e. a sudden and undesirable change in the current state), rather than a task-irrelevant stimulus made inconspicuous by prior experience.

Second, habituation as an explanation for the effects described in Chapter 4 cannot readily account for the *increase* in the intention insertion effect in Experiment 6 (when participants' hands were occluded) relative to Experiment 5 (when participants' hands were visible). In both experiments, the visual stimuli were identical, so one would have expected participants in the heightened ownership to have habituated to the presence of finger movements to the same extent – whether their hands were occluded or not. However, to more explicitly rule out the effects of simple habituation it would be useful to use a non-motor ownership induction to allow for a simpler interpretation. In this way, a future study might consider repeating the task but, instead, applying an ownership induction that relies on synchronous/asynchronous tactile stimulation (e.g. Botvinick & Cohen, 1998), or synchronous/asynchronous motor movement (Ma & Hommel, 2015b, 2015a), rather than merely comparing absence or presence of seen movement in response to one's own movements.

4.5. Conclusion

In conclusion, the data gathered in Chapter 4 have provided empirical evidence in support of some of the assumed features of modern accounts of action control.

Experiments 5 and 6 suggest that the tendency for an observed action to activate the corresponding response can be modulated by a body ownership illusion induced by correlated visual and proprioceptive feedback (i.e. reducing the effect) and the availability of accurate visual feedback of one's own movements (i.e. enhancing the effect). While there is a variety of research probing the mechanisms which govern human behaviour at various levels of abstraction, from the anatomical to the conceptual, there is an increasing acceptance of the notion that action control is fundamentally perceptual in nature (Gentsch et al., 2016). The brain is seen largely as a prediction engine, combining prior experience with information about the current state of the sensorimotor system to anticipate upcoming stimulation. However, a common criticism of these accounts is that they are largely theoretical and are difficult to test empirically (Huang, 2008). The experiments described in this chapter provide a novel approach to driving the motor system without relying on participants to intentionally initiate the response. As a result, the intention insertion task can be used to reveal important details about the mechanisms which translate sensory input into motor output.

5. Chapter 5: The Relationship Between the Generation of Action Slips and Ideomotor (Hypnotic) Suggestibility

5.1. Overview

The results from Chapter 2 and 3 have supported the assumption that perceptually-coded action effect representations engage the motor system when activated strongly enough by endogenous (e.g. mental imagery) and exogenous means (e.g. action observation) and can induce overt action. Moreover, Chapter 4 revealed that a body ownership illusion decreased the intention insertion effect; a finding consistent with the ideomotor view that movement occurs in case of a strong discrepancy, or prediction error, between one's current self-representation and a future predicted goal state. In Chapter 5, these same concepts are applied to ideomotor suggestibility, a phenomenon that is related to intention insertion. Following testing, participants who took part in Experiments 5 and 6 (Chapter 4) were screened to assess individual differences in ideomotor suggestibility, the motor component of the multifarious trait of hypnotic suggestibility (Woody et al., 2005). Chapter 5 aims to evaluate how the tendency for observed actions to engage the motor system during imagery (i.e. the intention insertion effect) relates to participants' responsiveness to simple ideomotor suggestions.

5.2. Introduction

The research presented in this thesis was initially motivated by the bizarre phenomenology of hypnosis and how modern ideomotor models of action control might be able to explain hypnotic motor behaviour. Indeed, hypnosis and related parapsychological phenomena, such as the magic swinging pendulum, table turning and

the divining rod, have a long history in classical ideomotor research (for review, see Stock & Stock, 2004). Early proponents of ideomotor concepts, such as Carpenter (1852), Chevreul (1833) and James (1890) developed a plausible, scientific account of these reportedly supernatural experiences. They claimed that vivid imagery is responsible for generating the seemingly non-willed behaviour which commonly occurs following hypnotic suggestions (Carpenter, 1852), just as it is responsible for every day volitional action, just in an unusual setting (e.g. a “hypnotic induction”). Inspired by this idea, the intention insertion task was designed to test the assumption that super-threshold activation of action effect representations (i.e. mental images) brings about overt action, just as with classical ideomotor phenomena.

When subjected to hypnosis some people appear in thrall to their hypnotist’s suggestions as if, at that moment, they exert no control over their own behaviour. During a typical hand lowering suggestion, the hypnotist delivers carefully worded instructions (i.e. ideomotor suggestions) for their subject to hold their hand out at shoulder height and to imagine holding a heavy ball (Bowers, 1998). Directions are given to take note of the ball’s shape and to imagine the feeling of a heavy mass pulling on the fingers and the outstretched arm. The hypnotic subject imagines the ball getting heavier and heavier while an invisible and overwhelming force pulls their arm downwards. As this scenario unfolds, responsive subjects will begin to lower their arm, but will often report afterwards that the movements they made were involuntary. Importantly, these self-reported feelings of non-volition occur despite the fact that, objectively, people must produce the movement themselves. To account for this seemingly non-willed behaviour, recall the notion that action reflects the resolution of the conflict between one’s current state and anticipated sensory goal state representations (Chapter 4). Under this view, if one’s current state representation and a sensory goal state (e.g. activated via imagery) are sufficiently divergent, then the motor

system will be engaged and will affect a bodily transition towards the target state to eliminate the discrepancy. Intuitively, such a mechanism could readily account for how vivid mental imagery could become translated into unintentional motor output in a hypnotic context.

Now that the pressure-based intention insertion paradigm has been thoroughly validated (Chapter 3) and the effects replicated and extended (Chapter 4), the aim is to come full circle and evaluate the assumed relationship between hypnotic behaviour and the ideomotor processes manifested in the intention insertion effect. While there is a vast spectrum of research related to hypnosis (for review, see Halligan & Oakley, 2014), there is a growing consensus that hypnotic suggestibility should not be treated as a unitary concept. Instead, Woody, Barnier and McConkey (2005) argue that hypnotic suggestibility is actually a manifestation of a constellation of sub-components which influence participant's responsiveness to specific hypnotic items (e.g. ideomotor suggestions vs. post-hypnotic amnesia). In light of this, the present chapter is specifically concerned with the insights that the intention insertion effect can bring to the ideomotor component of hypnosis. If the intention insertion task reflects the same mechanisms as the involuntary behaviour exhibited during hypnosis, then individual differences in susceptibility to the intention insertion manipulation should be related to ideomotor (hypnotic) suggestibility. In order to test this prediction, participants who completed Experiments 5 and 6 (Chapter 4) were screened in order to determine their responsiveness to simple ideomotor suggestions (e.g. arm levitation) and the scores were incorporated into an overall analysis of the pressure data.

Following the ideomotor account of the intention insertion effect, in order for a movement to occur following a specific suggestion, the goal representation (i.e. the simulated effects of the suggested movement) must become super activated (e.g. via elaborate mental imagery) which inadvertently engages the corresponding motor

processes. If this is the case, it suggests that individual differences in the ideomotor component of the general trait of hypnotic suggestibility (for review, see Lynn, Laurence, & Kirsch, 2015) may reflect a tendency to integrate activated perceptual representations with goal-related processes. Conceivably, this might arise because of a predisposition to focus on anticipated sensation (i.e. future-oriented), rather than current sensation (i.e. present-oriented). As a consequence, in the ideomotor view, a heightened focus on activated goal states would naturally drive conflict with the current state representation during vivid imagery and more readily engage the motor system. If this conceptualisation is valid, then more suggestible participants should exhibit a stronger intention insertion effect than those who are less suggestible. That is, there should be a positive relationship between the intention insertion effect and ideomotor suggestibility. First, however, it is useful to explore how such a prediction relates to existing research concerning the origins of hypnotic behaviour.

5.2.1. Action Slips and Existing Perspectives on Hypnotic Responding

The above proposal is similar to active inference-based accounts of sensory illusions (H. Brown, Adams, Parees, Edwards, & Friston, 2013) and hypnotic responding (Jamieson, 2016) which cast hypnosis as a form of sensory hypothesis testing which relies upon hierarchically organised generative models of the sources of sensory input. These accounts characterise the brain as a prediction engine with the sole function of anticipating future sensory input in order to minimise prediction errors (or “free energy”, Friston, 2010) generated by a comparison of bottom-up perceptual input and top-down predictions at multiple hierarchical levels. With respect to exogenous sensory inputs (e.g. vision), prediction errors are typically minimised by updating one’s

hypothesis about the causes of those inputs, based on prior knowledge, in accordance with the incoming sensory information. For example, when looking at an object in the distance (i.e. visual sensory input), one might form an initial hypothesis about the object's identity based on past experience with similar looking items, which is then updated upon closer inspection (e.g. Is it a bird? Is it a plane?). However, when proprioception is involved, prediction errors are driven by proprioceptive feedback which corresponds with muscle activity (e.g. signals from muscle fibres) and has the capacity to be resolved by movement (Adams, Shipp, & Friston, 2013). That is, the sensory inputs can be made to conform to the predictions of the generative model by initiating action to resolve proprioceptive prediction errors. For example, when trying to avoid scratching an itch (i.e. tactile sensory input), the imagined relief (i.e. tactile goal state) almost reflexively drives the scratching movement rather than a change in the hypothesis about the causes of that input (e.g. an insect bite). This is because imagining scratching one's itch naturally involves forming a proprioceptive prediction about how one might achieve it (i.e. proprioceptive goal state) which, in turn, drives prediction errors (i.e. the urge to scratch). This process will continue to drive prediction errors until a scratching action is initiated or the itching sensation subsides. In other words, movement reflects a process of bringing one's current state into alignment with one's expectations about an anticipated proprioceptive goal state.

Jamieson (2016) proposed a similar mechanism. In this view, the hypnotist's words activate high-level representations in the mind of the hypnotic subject that reflect the meaning of the hypnotic suggestion (i.e. the perceptual consequences of the body movement implied by it) which, in turn, drives activity in lower-level perceptual units. To illustrate, consider the arm lowering suggestion described earlier. Entertaining the hypothesis that a heavy ball is pulling one's outstretched arm downwards involves generating a set of predictions about how such an event might look and feel. Given that

this sensory hypothesis, by definition, does not correspond with objective reality, it will naturally drive prediction errors, rather than prediction matches, in lower-level representational units (Jamieson, 2016). This surge of prediction errors, driven by proprioceptive feedback from the target limb, will either cause the generative model to shift back to reality (i.e. revising the prior hypothesis) or initiate a movement to bring sensory input in line with the prediction (i.e. validating the prior hypothesis). Following this logic, less suggestible individuals may be more inclined to correct or reject a given sensory hypothesis, perhaps due to a heightened focus on their current state representation (i.e. present-oriented). In contrast, highly suggestible individuals may focus more strongly on the anticipated goal state (i.e. future-oriented) and will, therefore, be more inclined to change the sensory input to conform to the expectations of the generative model by engaging the motor system.

In support on this conceptualisation, an interesting line of research has cast hypnotic responding as a form of strategic self-deception with respect to one's current mental state (i.e. a lack of awareness of, or focus on, one's current state representation; Dienes et al., 2016). In this account, hypnosis is contrasted with Buddhist meditative practices and secular mindfulness techniques which are described as a form of strategic self-insight. Consistent with this idea, research has demonstrated a negative correlation between hypnotic suggestibility and mindfulness, such that individuals with low scores on measures of hypnotic suggestibility tend to score highly on mindfulness scales (Grover, Jensen, Patterson, Gertz, & Day, 2018). Intuitively, this conceptualisation of hypnotic suggestibility corresponds nicely with the ideomotor prospective. If people are less responsive to hypnotic suggestion because they are more mindful (i.e. stronger self-insight) then this is consistent with the notion of a heightened focus on the current state, which might preclude the activation of anticipated goal states. Conversely, if people are highly responsive to hypnosis because they are less mindful (i.e. weaker self-insight),

then this could reflect a reduced focus on the current state and a correspondingly stronger focus on anticipated goal states, inadvertently driving action (i.e. stronger self-deception). If individual differences in hypnotic suggestibility are related to one's predisposition to focus on current states (i.e. self-insight) or future goal states (i.e. self-deception), then hypnotic suggestibility should exhibit a generally positive relationship with the intention insertion effect. To evaluate this prediction, participants' suggestibility scores were combined with the data from Experiments 5 and 6.

5.3. Experiment 7: A Test of the Relationship Between the Intention Insertion Effect and Ideomotor Suggestibility

In order to test the proposed relationship between the intention insertion effect and ideomotor suggestibility, participants who had completed Experiments 5 and 6 were subjected to a suggestibility screening procedure immediately following completion of the intention insertion task. Suggestibility screening was carried out with the use of a novel protocol which combined the ideomotor items from established scales such as the Harvard Group Scale of Hypnotic Susceptibility (Shor & Orne, 1962) and the Stanford Scale of Hypnotic Suggestibility: Form C (Weitzenhoffer, Hilgard, & Kihlstrom, 1996). For the present study, the simple pass/fail scoring system used in the aforementioned scales was rejected in favour of operationalising suggestibility in terms of the extent (e.g. on a 0 to 3 scale) of overt movement following each item. This suits the purpose of assessing the relationship between the size of the intention insertion effect and the extent of motor engagement caused by simple ideomotor suggestions.

5.3.1. Method

5.3.1.1. Participants

All participants who completed the task in Experiments 5 and 6 and met the criteria for inclusion in the respective analyses (N=67 and N=68 respectively) provided written consent to take part in the suggestibility screening procedure.

5.3.1.2. Materials

The suggestibility protocol consisted of a four-item scale that indexes the magnitude of the physical response to each suggestion. Typically, hypnotic suggestibility is assessed using one of a range of well-established protocols, such as the Harvard Group Scale of Hypnotic Susceptibility (HGSHS: Form A, Shor & Orne, 1962) or the Stanford Scale of Hypnotic Suggestibility (SSHS: Form C, Weitzenhoffer & Hilgard, 1962). These scales include different classes of suggestion such as ideomotor (e.g. arm levitation), challenge (e.g. arm rigidity) and executive function (e.g. amnesia) items. However, with the present study, the scope was deliberately narrowed to focus on the correspondence between motor output following ideomotor suggestions and the intention insertion effect. Therefore, the protocol used in Experiment 7 takes the established scales but uses only ideomotor items.

The screening protocol was delivered by a trained hypnotherapist via headphones with a pre-recorded tape. A video camera was used to capture the participant's responses to each item. In addition, a splitter was used to send the audio signal to both the headphones and the camera in order to appropriately timestamp each video for subsequent coding. The procedure lasted around 15 minutes and consisted of a brief introduction and four ideomotor suggestions, comprising (a) an eye closing suggestion;

(b) a head falling suggestion; (c) a hand levitation suggestion; and (d) a combination hand lowering and levitation suggestion. The recording began with an introduction to the procedure and explained the manner in which participants should engage with the task and then each item was delivered in succession.

The first item was an eye closing suggestion in which participants were asked to relax and imagine how heavy their eyes were beginning to feel. The suggestions continued until 1m 55s into the recording, at which point the phrase, “Close your eyes...” was repeated. Participants were scored on a scale of 0 (e.g. eyes fully open) to 2 (e.g. eyes fully closed) at this key moment in the tape.

The second item was a head falling suggestion in which participants were asked to imagine that their head was feeling heavy and starting to fall forwards. The item began at 3m 10s with an instruction to sit upright in a comfortable position in the chair, still with the eyes closed. The suggestion ended with an instruction to, “stop there” at 5m 30s and was scored according to the magnitude of the difference between starting and final posture. The scale ranged from 0 (e.g. no movement, sat upright) to 3 (e.g. both head and body fully slumped forwards, having started in an upright posture).

The third item began at 6m 10s and consisted of a hand levitation suggestion. Participants were asked to place their non-dominant hand on their thigh and raise it slightly, such that only their fingertips are in contact with their leg. The recording gave instructions to imagine how light the hand feels and to picture it floating upwards. The item ended at 9m 10s and was scored according to the vertical displacement of the imagined hand. The scale ranged from 0 (e.g. the hand is still in contact with the leg) to 3 (e.g. hand raised above the shoulder).

The final item began at 10m 05s and consisted of a suggestion to imagine holding a pile of books in one hand and a bunch of balloons in the other. Participants began with

the arms outstretched, palms facing upwards. In one hand they were asked to picture a stack of books and to attend to feelings of heaviness as well as the look, shape and smell of the books. With the other hand, they were simultaneously asked to imagine their hand being pulled upwards by some helium balloons. The item formally finished at 13:50, but participants varied in their ability to keep their arms outstretched for the full duration of the suggestion, so this item was coded according to the maximum difference between the vertical displacements of the two hands. The scale ranged from 0 (e.g. hands level) to 3 (e.g. one hand well above the shoulder and one well below).

5.3.1.3. Procedure

After completion of the intention insertion task (Experiments 5 and 6, Chapter 4), participants gave written consent to complete a brief screening procedure in order to assess individual differences in ideomotor suggestibility. Participants sat in a comfortable chair in a small, dimly lit room with large screens which divided the room in half. These privacy screens were intended to minimise feelings of self-consciousness and allowed the experimenter to stay out of view during the entire procedure. Participants were given a verbal briefing prior to screening in order to emphasise some important elements of the procedure. In particular, participants were specifically instructed that they must actively engage with each suggestion by constructing a vivid mental image of the scenario described by the recording, rather than listening passively. In addition, it was highlighted that the researchers were not interested in the participant's final score and that less suggestible participants were just as important as those who were highly suggestible. Following this verbal briefing, the experimenter set the camera to record and began the tape. Participants were then left in relative privacy

for the duration of the screening procedure. When the recording finished, the camera was paused and the video file downloaded for later coding.

5.3.1.4. Data Analysis

The video files recorded were coded according to the criteria reported earlier. The scores on each item were scaled to fit the range zero to one and an average movement score was derived for each participant, which was used as an index of motor engagement. These scores were then mean-centred (e.g. with range $-.5$ to $.5$) and added to the combined datasets generated by Experiments 5 and 6. The data were then reanalysed using a model with an identical specification as those used in Chapter 4 with the addition of the ideomotor suggestibility score as a continuous covariate. As a consequence of the scaling of the ideomotor suggestibility score, a one unit change in this variable corresponds to the difference between the minimum and the maximum theoretical score under this protocol.

5.3.2. Results

5.3.2.1. Exclusion Criteria

In total, 4 participants were excluded for making no attempt to engage with the screening procedure and 5 participants were excluded due to a recording malfunction. As a result of these exclusions, the final sample of $N=126$ consisted of participants who had satisfactorily completed both the intention insertion task and ideomotor suggestibility screening. Screening revealed that, on the original scale (i.e. not mean centred), the average ideomotor suggestibility score across all participants was $.42$ ($SD = .24$).

5.3.2.2. The Intention Insertion Effect and Ideomotor Suggestibility

As in Chapter 4 (Experiments 5 and 6), this analysis models the pressure that was induced when participants were presented with action cues at the finger which matched and the finger that mismatched the observed movement. The intention insertion effect reflects a contrast of the presses produced by the cue-compatible and cue-incompatible finger in the critical 800ms period following cue onset. As in Experiments 4-6, cue-compatible and cue-incompatible responses were baselined by subtracting the pressure exerted by participants in response to neutral cues (i.e. no movement). Importantly, the analysis focused only on those aspects of the data that are relevant to this prediction, namely the interaction of the intention insertion effect with the participant's current mental imagery (i.e. imagining the same or a different movement to the currently recorded finger) and with their responsiveness to ideomotor suggestions.

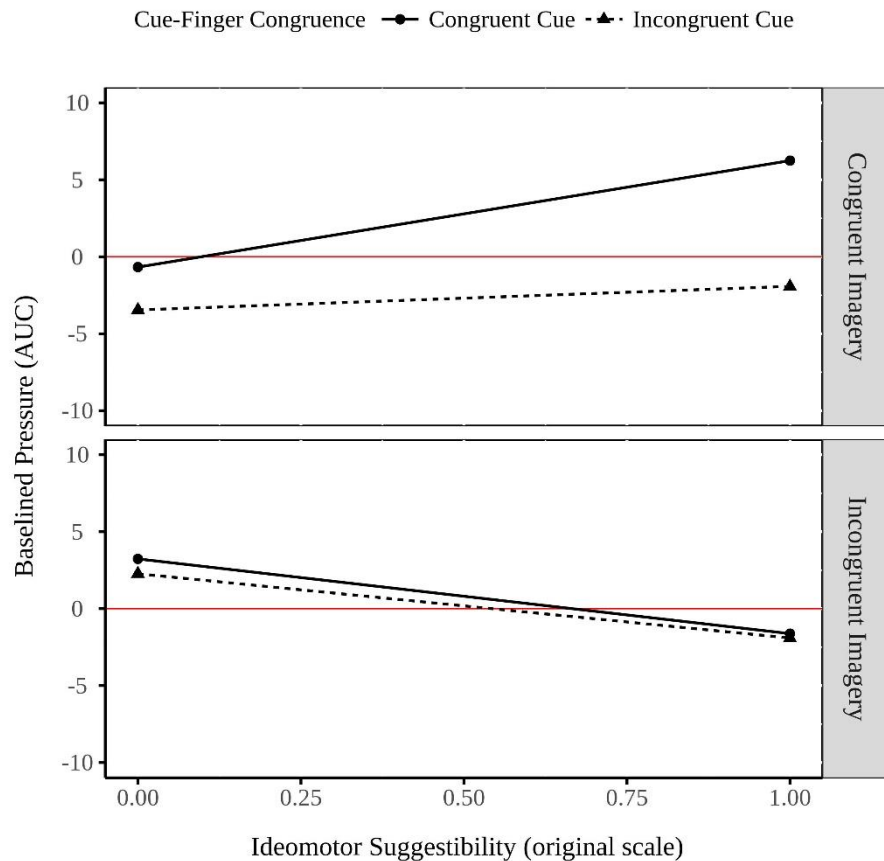


Figure 22. A multi-panel plot showing differences in finger pressure during intervals following neutral vs. other types of action cue against imagery and across the full range of ideomotor suggestibility on the original scale (i.e. from the minimum, zero, to the maximum, one, theoretical score). Positive scores indicate participants pressed harder than in equivalent trials where a neutral cue was presented. Negative scores indicate that participants pressed less hard than following neutral cues (i.e. implying interference). The upper and lower panels distinguish responses where the participant was/was not required to imagine a movement in the recorded finger. The lines indicate whether the cue contained an animation of the same (i.e. congruent – solid)/other finger (i.e. incongruent – dotted) from that being recorded.

For example, the top panel of Figure 22 shows that the pressure exerted by the cue-congruent and cue-incongruent finger generally increased with ideomotor suggestibility (relative to baseline) when participants also imagined a movement with that finger. Furthermore, visual inspection suggests that the intention insertion effect, captured by the difference in pressure at the cue-congruent and cue-incongruent finger, appears to increase with ideomotor suggestibility (i.e. the two slopes diverge). In contrast, the bottom panel of Figure 22 suggests that the magnitude of both cue-

congruent and cue-incongruent responses, relative to baseline, decreases with ideomotor suggestibility. In addition, it appears that the intention insertion effect is very small in general (i.e. the difference between the two lines is small), and does not vary according to ideomotor suggestibility (i.e. the two lines are almost parallel) with incongruent imagery.

As in Chapter 3 and 4 (Experiments 4-6), in order to gauge the strength of evidence for these observations, data were simulated from the model and summary statistics, Bayes-p values and evidence ratios regarding the differences of interest were calculated. To briefly recap, the posterior predictive Bayes-p values associated with the forthcoming contrasts represent the probability (given the model and the data) that the effect was in the predicted direction (i.e. H_1), as opposed to the opposite direction (i.e. H_0). To aid interpretation, Bayes-p values are recapitulated as evidence ratios (ER; analogous with Bayes factors) to describe the strength of evidence in favour of the directed hypothesis, weighted against the alternative hypothesis. For consistency, the weight of evidence in support of each directed hypothesis was described using terminology regularly applied to Bayes factors (Wagenmakers et al., 2018); “extremely strong” (ER: >100), “very strong” (ER: 30-100), “strong” (ER: 10 - 30), “moderate” (ER: 3 - 10), “anecdotal” (ER: >1 – 3) and “no evidence” (ER: 1). Importantly, evidence ratios less than one indicate that the evidence supports the alternative hypothesis; even more so for values which are a tiny fraction of one. For example, an $ER < .01$ represents *extremely strong* evidence in favour of the *alternative hypothesis*.

It was predicted that if highly suggestible participants are more prone to focusing on anticipated goal states, then the intention insertion effect should increase with ideomotor suggestibility and that the relationship should be strongest with congruent imagery. A test of the overall interaction provides moderate evidence ($ER = 6.94$) for a positive relationship between the intention insertion effect and hypnotic suggestibility

with congruent, relative to incongruent imagery (Posterior Mean = 6.07, 95% PCI [-4.3, 16.6], Bayes-p = .874). In other words, the test suggests that it was 6.94 times more likely that participants who were more suggestible also tended to respond more strongly when viewing visual cues that match the currently imagined finger key press, as predicted.

Step-down analyses confirm this interpretation. The analysis of the finger that the participants imaged moving (Figure 22, top panel) provides moderate evidence (ER = 5.85) of a positive relationship between the intention insertion effect and ideomotor suggestibility (Figure 22, upper panel; Posterior Mean = 5.38, 95% PCI [-4.66, 15.39], Bayes-p = .854). Closer inspection reveals that there is strong evidence (ER = 14.4) that an increase in ideomotor suggestibility was associated with stronger presses at the cue-congruent finger than baseline (Posterior Mean = 6.92, 95% PCI [-2.03, 16.13], Bayes-p = .935). That is, there is 14.4 times more evidence in favour of a positive relationship between the size of cue-compatible responses and ideomotor suggestibility than a negative relationship. However, counter to expectations, there is anecdotal evidence (ER = .57) that the interference effect caused by action cues on the cue-incongruent finger was weaker (i.e. the effect becomes less negative, implying less interference; Figure 22, upper panel, dotted line) with increasing suggestibility (Posterior Mean = 1.53, 95% PCI [-6.93, 10.34], Bayes-p = .365).

In contrast, the analysis of the finger that the participants did not imagine moving (Figure 22, bottom panel) provides anecdotal evidence (ER = .78) of a negative relationship between the intention insertion effect and ideomotor suggestibility (Posterior Mean = .68, 95% PCI [-9.28, 8.14], Bayes-p = .438). Accordingly, the tests revealed moderate evidence (ER = .16) that, with incongruent mental imagery, there was a negative relationship between ideomotor suggestibility and presses made with the cue-congruent finger (Figure 22, lower panel, solid line; Posterior Mean = -

4.86, 95% PCI [-13.52, 3.97], Bayes-p = .136). In addition, there was moderate evidence (ER = 6.14) that when the recorded finger was incongruent with the imagined movement, action cues interfered with cue-incongruent responses more strongly, relative to baseline, with increasing ideomotor suggestibility (Figure 22, lower panel, dotted line, Posterior Mean = -4.17, 95% PCI [-11.77, 3.53], Bayes-p = .86).

5.3.3. Discussion

The purpose of Chapter 5 was to test whether the non-willed action slips captured by the intention insertion task (Chapter 3 and 4) and the similar non-willed behaviour produced in simple hypnotic suggestions are linked. More specifically, the analysis presented here evaluated whether individual differences in ideomotor suggestibility (i.e. the motor component of hypnotic suggestibility) could be used to predict the magnitude of the intention insertion effect. In the introduction, it was proposed that highly suggestible individuals might be more prone to focusing on anticipated goal states (e.g. the simulated feelings associated with the suggestion) rather than their current state representation (e.g. the current position of the body in space). The findings of Chapter 4 (Experiment 5 and 6) indicate that a heightened focus on anticipated events could inherently drive conflict with the current state and, by extension, predispose some people to produce the suggested behaviour. On this basis, it was predicted that if individual differences in ideomotor suggestibility reflect a predisposition to focus on goal states, then it should demonstrate a positive relationship with the intention insertion effect. As predicted, Experiment 7 suggested that, with compatible imagery, there was moderate evidence to suggest a positive relationship between the intention insertion effect and ideomotor suggestibility, particularly on the finger that participants' concurrently imagined moving.

The prediction that the intention insertion effect should increase with ideomotor suggestibility implicitly assumed that both components of the effect would increase at the same rate (i.e. facilitation of cue-congruent responses and interference of cue-incongruent responses would be in proportion). Indeed, the data revealed that there was evidence that higher suggestibility scores were associated with substantially stronger cue-congruent responses (i.e. which matched the observed movement) than baseline (i.e. the neutral, no movement condition). In contrast, there was only anecdotal evidence to suggest that the interference effect of action observation on the cue-incongruent response, relative to baseline, interacted with suggestibility. These findings appear to be quite consistent with the notion that those participants who score highly in ideomotor suggestibility may tend to focus more strongly on anticipated goal states if one assumes that the endogenously-driven goal state (e.g. via imagery) has primacy over an exogenously evoked goal state (e.g. via action observation). On this view, exogenous stimuli are able to evoke activity in action effect codes shared with the imagined movement (i.e. anticipated goal state) and enhance the generation of the corresponding (i.e. cue-congruent) response. In addition, one might speculate how this might have contributed to the absence of a strong correlation between ideomotor suggestibility and the interference effect typically observed at the cue-incongruent finger could occur. That is, a stronger focus on the endogenously driven goal state may also reduce the extent to which exogenous stimuli are able to activate competing (i.e. non-overlapping) action effect codes.

The analysis of the individual components of the intention insertion effect at the finger which did not correspond with participants' imagery provides further support for this interpretation. Recall that, when the recorded finger did not correspond with the imagined movement, both cue-congruent and cue-incongruent responses became steadily weaker than baseline with increasing ideomotor suggestibility. That is, visual

cues interfered more strongly, in general, with the initiation of the non-imagined response with increasing ideomotor suggestibility (i.e. a general negative relationship, but no overall interaction with the intention insertion effect). These findings are consistent with the idea that the action cues were only able to further magnify goal states already activated by imagery in highly suggestible participants, rather than activate novel, non-imagined goal states. Such an account corresponds nicely with the notion that highly suggestible participants focus more strongly on their endogenously driven active goal states, which might preclude the activation of competing goals driven by exogenous stimuli. Importantly, this view accounts for both the positive relationship between ideomotor suggestibility and the intention insertion effect with congruent imagery (i.e. amplifying an endogenous goal) and the lack of interaction with incongruent imagery (i.e. the exogenously activated goal is inhibited more strongly by the endogenous goal).

5.3.3.1. A Brief History of Underwhelming Attempts at Predicting Hypnotic Suggestibility

At this point, it is worth highlighting that hypnosis research has a long history of searching, often fruitlessly, for robust personality traits, cognitive constructs and behavioural patterns which can predict individual differences in hypnotic suggestibility (for review, see Lynn et al., 2015). For example, personality traits such as openness to experience demonstrate only weak correlations with hypnotic suggestibility (Cardeña & Terhune, 2014). Similarly, individual differences in cognitive constructs such as imagery vividness and absorption only show weak to moderate correlations with hypnotic suggestibility (e.g. Glisky, Tataryn, & Kihlstrom, 1995), which become weaker and statistically non-significant when researchers control for the hypnotic

context (Amanda J Barnier & McConkey, 1999). These findings have typically been taken as evidence that hypnosis is not a unitary concept and, instead, can be better characterised as a constellation of sub-components with relevance to specific classes of hypnotic suggestion (e.g. ideomotor, ideomotor challenge, perceptual-cognitive and posthypnotic amnesia, see Woody, Barnier, & McConkey, 2005). For example, highly suggestible participants who also scored highly on an imagery vividness scale were more prone to incorporating false memory suggestions (Labelle, Dixon, & Laurence, 1989) and those high in imagery absorption were more convinced that their false memories were, in fact, real (Kandyba & Laurence, 1996). Indeed, the multifaceted nature of hypnosis research motivated the specific focus on the ideomotor aspect of the phenomenon under our definition of “ideomotor suggestibility”. However, given the apparent diversity of overlapping sub-components, it is perhaps not surprising that Experiment 7 only demonstrated a moderate relationship between the intention insertion effect and ideomotor suggestibility. It is entirely possible that there were individual differences in other components of hypnotic suggestibility (e.g. attention) which are not strictly ideomotor – and thus not well indexed by the screening procedure – but are still relevant to this form of hypnotic suggestion.

5.3.3.2. Study Limitations

While Experiment 7 did provide an important test of the assumption that the intention insertion effect arises from similar processes to those that produce hypnotic motor responses, the study was limited in two important respects which may account for the weakness of the observed relationship. First, ideomotor suggestibility was indexed by scoring the extent of overt movement in response to direct ideomotor suggestions borrowed from a range of established scales of hypnotic suggestibility (e.g. Harvard

Group Scale of Hypnotic Susceptibility, Shor & Orne, 1962). This was a deliberate choice, motivated by research which suggested ideomotor suggestibility is a sub-component of a more general trait of hypnotic suggestibility. However, ideomotor suggestions are considered the “easiest” to pass in the hierarchy of hypnotic suggestions (Amanda J Barnier, Dienes, & Mitchell, 2008) and therefore may not be sufficiently challenging to discriminate between participants who are moderately and highly suggestible. Anecdotally, a number of participants who scored in the upper range of ideomotor suggestibility behaved differently with respect to some of the more implicit suggestions in the screening procedure. For example, in between following the head falling and hand levitation, there was no explicit instruction to return to an upright seating position before the next suggestion. While most participants reset their posture, as instructed in previous items, a handful remained in a fully slumped position for the duration of the following suggestion, as if trapped in place. This suggests that ceiling effects may have limited the ability of the screening procedure to discriminate between moderately and highly suggestible participants. In future, a broader range of more difficult suggestions which are still relevant to the motor aspect of the phenomenon (e.g. ideomotor challenge suggestions, such as arm rigidity) would allow for greater precision for evaluating individual differences in ideomotor suggestibility. Importantly, this observation does not undermine the conclusions drawn from Experiment 7 and merely highlights that a more sensitive screening procedure would provide a better index of the phenomenon.

Second, Experiment 7 was further limited in the respect that screening occurred after participants had completed the intention insertion task (Chapter 4, Experiments 5 and 6). Again, this was a deliberate choice which was motivated by the need to complete test and screening in a single sitting in order to ensure efficient data collection and avoid data loss through participant non-attendance of follow-up appointments. In

addition, placing the screening procedure after the test was intended to limit the influence of any unintended demand characteristics of a hypnotic induction on the intention insertion task. However, it is possible that the completion of the intention insertion task (30-minute duration) may have fatigued some participants more than others with a potential to influence their performance in the subsequent screening procedure. In future, screening participants prior to the test in a setting that was not obviously connected with the subsequent experiment would ensure that equal numbers of low and highly suggestible participants could be tested and evaluated. This amendment, combined with more difficult ideomotor suggestions, would improve the ability of the model to estimate an accurate intention insertion effect for more highly suggestible participants.

5.4. Conclusion

In conclusion, Experiment 7 did indeed provide evidence in favour of a positive relationship between the intention insertion effect (with compatible imagery) and the behavioural responses to simple ideomotor suggestions, as predicted. However, while these data suggest that we can be fairly confident about the existence of the proposed conceptual relationship, several aspects of the design (e.g. the screening procedure) may mean that the extent of the relationship is actually underestimated. As a result, Chapter 5 represents a starting point for future exploration of the extent of the relationship between these two phenomena. With an improved screening procedure which makes use of more difficult “challenge” suggestions, it may be possible to more thoroughly validate the intention insertion effect as a measure of ideomotor suggestibility, one of the key sub-components of the broader trait of hypnotic suggestibility (Woody et al., 2005).

6. Chapter 6: General Discussion

6.1. Introduction

The experiments presented in this thesis have examined non-willed behaviour (i.e. action slips) produced by participants during an imagery task. The studies tested fundamental assumptions about voluntary action control that are common to classical ideomotor theories and modern extensions of these principles (collectively, and rather crudely, referred to as “ideomotor models”; Adams, Shipp, & Friston, 2013; Clark, 2013; Hommel, 2009; Prinz, 1997; Wolpert, 1997). Among these theories, there is broad agreement that perception and action are inextricably linked. Remarkably, however, there remains an enormous variety of different models which make subtly different claims about the manner in which perception and action interact (for a review, see Gentsch et al., 2016). One reason for this apparent divergence of theoretical opinion is the highly abstracted plain on which the discussion takes place and this obscures some of the natural conceptual overlap between different models. In addition, these accounts are largely reliant on theoretical computational models and some of their most fundamental claims have yet to be tested directly.

Ideomotor models provide an intuitive framework for understanding how voluntary action control may be achieved simply by anticipating (i.e. imagining, predicting or otherwise “activating”) the perceptual consequences of desirable action goals. One only has to imagine a particular sensory effect one wishes to achieve (e.g. sipping from a glass of water) and automatic motor processes make this imagined act a reality (e.g. the reaching and grasping action required to reach this goal). When selecting an action, people voluntarily evoke (e.g. via imagery) the perceptual consequences of candidate actions (i.e. action effects) in order to generate the

appropriate behaviour. Importantly, action effects are also assumed to be activated by perceptual input.

A great deal of research has demonstrated that people produce faster and more accurate responses if they are presented, just prior to execution, with the consequences of their intended action, relative to an alternative action (Shin et al., 2010). For example, people respond faster and more accurately when presented with a visual depiction of a corresponding movement (that is, “proximal” action effects; e.g. Brass, Bekkering, & Prinz, 2001; Brass, Bekkering, Wohlschläger, & Prinz, 2000), and when people encounter a tone which previously followed the to-be-executed action (“distal” action effects; e.g. Elsner & Hommel, 2001). Indeed, in the opening lines of their review, Shin et al. (2010, p. 943) explicitly stated that “Any activation of the effect image, either endogenously or exogenously, will trigger the corresponding action.” However, despite being such a fundamental assumption, the *initiating* function of action effect activation on the corresponding behaviour has only rarely (if at all) been addressed directly. Therefore, it is not clear whether imagery and perceptual input can really be considered functionally equivalent, such that perceptual inputs could accidentally trigger an imagined response if the shared action effect codes were activated strongly enough. That is, if one’s mental imagery happens to align with some strongly associated stimulus in the environment, would the corresponding action be automatically initiated? This line of thinking motivated the experiments presented in this thesis in which participants were presented with unexpected finger movements while they mentally rehearsed a finger tapping sequence in an effort to elicit action slips. If the assumptions of ideomotor models are taken seriously, then imagining and observing the same action should produce especially strong activation of shared action effects and, as a consequence, trigger the corresponding action – in spite of instructions to remain still.

6.2. Summary of Findings

Across a number of chapters, the experiments presented in this thesis have provided robust evidence in favour of the idea that imagery, perception and action rely on a common neural infrastructure (e.g. Common Coding theory, Prinz, 1990, 1997; van der Wel, Sebanz, & Knoblich, 2013). More specifically, these data support a long-held but, until now, an untested assumption of ideomotor models that action effect activation is translated into action when this “mental image” becomes sufficiently vivid; a core ideomotor concept that can be traced back to Carpenter (1852) and James (1890).

Experiments 1-3 (Chapter 2) consistently found that observed finger movements triggered larger responses with the corresponding finger than the non-corresponding finger, particularly when the imagined and observed action were congruent. Thus, these experiments demonstrated, for the first time, that evoking the perceptual consequences of a given action (i.e. action effects), by endogenous (e.g. action imagery) and exogenous means (e.g. action observation), is sufficient to trigger the corresponding behaviour, in accordance with ideomotor principles (Shin et al., 2010).

Experiment 4 (Chapter 3) further specified these effects by demonstrating that, with congruent mental imagery, observed finger movements both facilitated the execution of the corresponding response and interfered with the non-corresponding response, relative to a similarly salient, but motorically neutral stimulus. This finding suggests two things. First, that observing an action can selectively facilitate the corresponding response by further activating action effects which are shared with the imagined movement. Second, when the observed and imagined movements differ (e.g. an observed left finger tap, while imagining a right finger tap) the action effects evoked via perception (e.g. left) and imagery (e.g. right) interfere with one another – impairing

action initiation and reducing the magnitude of the subsequent action slip relative to baseline.

Experiment 5 and 6 (Chapter 4) demonstrated that the tendency for an observed action to engage the motor system (i.e. the intention insertion effect) can be modulated with a body ownership illusion and by restricting visual feedback of participants' limbs. That is, the effect was consistently weaker for those who were subjected to a body ownership illusion (i.e. virtual hand illusion) than for the control group without the illusion (Experiments 5 and 6). Furthermore, the intention insertion effect was enhanced when participants' hands were occluded (Experiment 6) relative to when they were visible (Experiment 5). These findings argue against a direct coupling of visual (i.e. sensory) and motor representations serving the corresponding action, as assumed by models of motor control which seek to account for automatic imitation effects with general associative mechanisms (Associative Sequence Learning; Heyes & Bird, 2008; Heyes & Ray, 2000). For example, recall that in Experiments 5 and 6, the observed movements had a much weaker effect on the size of participants' action slips in the heightened ownership group than the control group, even though the same set of visual stimuli were presented to both groups. If the visual input which corresponded with the observed movement activated the associated motor representation directly, then one would expect that heightening the correspondence between participants' movements and the onscreen hands (and thus, the action effect binding) would enhance, not reduce the effect. Instead, this finding can be readily explained if one assumes that sensory information is integrated with higher-order multimodal representations, prior to engaging the motor system; an assumption which is common, but not necessarily explicit, in models based on ideomotor principles (Clark, 2013; Hommel, 2009; Wolpert, 1997). In this view, incoming sensory information is weighted towards either an internal representation which captures (i.e. estimates) the current state of one's body

or, conversely, a representation which describes anticipated future goal states. Importantly, only goal states have the capacity to drive the motor system. To illustrate, consider the clapping game “Pat-a-cake”, in which people are able to coordinate complex sequences of individual, two-handed and crossed claps with a partner. If visual attention is directed toward one’s own movements, then these inputs correspond with the current state of the body and do not engage the motor system. Alternatively, observing a partner’s movement means that the visual inputs can be effectively used as goal states for our own motor system, allowing one to efficiently imitate the movements and learn new sequences. In other words, these findings support modern ideomotor accounts which assume that people maintain separate internal representations of the estimated current state of their body and potential future goal states.

Furthermore, when participants’ hands were occluded (Experiment 6: Control Ownership group), there was evidence of a strong difference in the pressure induced by action cues at participants’ corresponding and non-corresponding fingers (i.e. a robust intention insertion effect) even when they were not imagining using this finger. This indicates that exogenous stimuli do not merely enhance (or interfere with) existing action plans (e.g. pushing an imagined action past the threshold for execution), but can also trigger even *non*-imagined behaviour when the corresponding action effects are activated strongly enough by the visual sensory input. Crucially, these data contrast with research which has only reported interference effects of imagery (Ramsey et al., 2010) and action observation (Kilner, Hamilton, & Blakemore, 2007; Kilner, Paulignan, & Blakemore, 2003) on action execution.

Finally, in Chapter 5, promising evidence of the assumed positive relationship between the magnitude of the intention insertion effect and ideomotor suggestibility was presented (i.e. the motor component of the broader trait of hypnotic suggestibility; Woody, Barnier, & McConkey, 2005). As predicted, participants who were more

responsive to simple ideomotor suggestions during a standardized hypnotic induction exhibited a stronger intention insertion effect (with compatible mental imagery) than those who were less responsive to the suggestions. This prediction was motivated by theoretical accounts of hypnotic responding (Jamieson, 2016) and sensory illusions (H. Brown et al., 2013) based on Active Inference (Adams et al., 2013). The findings of Experiment 7 were consistent with the notion that ideomotor suggestibility (and perhaps, hypnotic suggestibility more broadly) may reflect a tendency to focus on anticipated future goal states, rather than the current state of one's body and the wider world. In this view, if highly suggestible individuals have a natural tendency to focus on anticipated (i.e. predicted) sensation, then they may naturally weight incoming sensory information towards integration with active goal states and are, therefore, more prone to inadvertently engaging the motor system (i.e. producing action slips). This conceptualisation draws parallels with modern accounts which cast the highly hypnotic subject as one who is prone to self-delusion (i.e. a tendency to focus on anticipated sensation), rather than self-insight (i.e. a tendency to focus on current sensation, or, one's current state; Dienes et al., 2016). Taken together, these experiments provide novel insights into the process by which we come to engage with the world. The findings have implications, not just for ideomotor models and the information processing steps involved in motor control, but our understanding of volition more generally.

6.3. The Perceptual Format of Action Intentions

The introductory paragraphs of this thesis (Chapter 1) raised the prospect that instances of accidental behaviour in everyday life (i.e. action slips; Freud, 1914; Reason, 1984) and responses to classical ideomotor phenomena (e.g. Chevreul's

Pendulum; Easton & Shor, 1975, 1976, 1977) are not random but, instead, are borne of the same processes which govern willed action. For example, one might find oneself absentmindedly reaching for a jug of milk rather than a freshly boiled kettle and pouring a full cup of cold milk rather than hot water onto the waiting teabag. In this case, both the milk jug and kettle are of similar size, have a handle and afford a similar grasping action and, indeed, both are required to make tea – just not in that order or in those quantities! Crucially, such action slips are often explained away after the fact by suggesting that these obviously erroneous actions were somehow “unintentional”. The question that motivated this thesis was, how is it that people can produce behaviour if it is not, in some sense, willed? To that end, Experiments 1-7 of the present thesis produced several findings which have provided novel insights into the role of imagery and perception in the generation of non-willed action slips.

Action in the real world has been conceptualised as an interaction of higher-level goals with lower-level cues from the environment: actions are executed if they match both the individual’s goals and are afforded by the environment (Cisek, 2007; Guérard & Brodeur, 2015; Kunde et al., 2007; Ondobaka & Bekkering, 2012). The keyboard-based (Experiment 1-3) and pressure-based (Experiments 4-6) intention insertion tasks can be considered as a model for these interactions. In ideomotor frameworks, higher-level goals are nothing more than the imagined action effects one wants to achieve (Hommel et al., 2001), similar to the imagined actions in the intention insertion task. The findings of Experiments 1-6 suggest that bottom-up activity induced by perceptual cues in the environment – in this case, the observed actions – activate the same action effect representations as top-down mental imagery and, when congruent, can combine to trigger overt action. In essence, this view implies that forming an intention to act is merely a case of strongly activating the action effects that are shared with the desired movement to the extent that they engage the motor system. Importantly, provided that

the resulting activation surpasses the response threshold, action can be driven by endogenous (e.g. imagery) and/or exogenous (e.g. action observation) sources of activation.

It has been argued that there are two classes of action which rely on separate neural substrates; “stimulus-based” actions that function as prepared reflexes to environmental stimuli (i.e. simple sensorimotor contingencies) and “intention-based” actions which aim to produce some endogenously specified effect (Herwig et al., 2007). In this view, stimulus-based actions are driven by exogenous processes and are assumed to be mediated by simple stimulus-response associations which bind a motor representation of the desired action with the conditions (i.e. specific exogenous sensory input) under which it must be executed (e.g. “red light – brake” when driving a car). In contrast, intention-based actions are assumed to be driven by endogenous processes and mediated by action effect associations which bind a motor representation and its typical perceptual effects. That is, the mere anticipation of the perceptual consequences of the action one wishes to achieve will automatically evoke the corresponding motor program. Inherent in this conceptualisation is the assumption that a given action will be either stimulus-driven or intentional (goal/effect-driven). However, action slips captured by the intention insertion tasks (Experiments 1-6) appear to provide an intermediate example of a behaviour which is neither wholly stimulus-driven nor intention-driven. Indeed, a natural consequence of the way ideomotor models conceive of voluntary action control is that it is precisely this combination of endogenous imagery and exogenous perceptual input which allows people to interact so effortlessly with the world. In other words, in such views, the vast majority of actions (excluding classical reflexes, e.g. the knee jerk response) are, to some extent, the product of endogenous and exogenous processes.

In Chapter 1, it was proposed that the purpose of an action goal (i.e. the intention-driven component) is to broadly specify a range of behaviours that might be appropriate in a given context. In the ideomotor view, this is achieved by anticipating (i.e. predicting, imagining or otherwise “activating”) relevant action effects, to a subthreshold level, in order to allow for the efficient translation of sensory input into motor output. Importantly, this “pre-activation” of candidate actions (or potential goal states) allows stimuli in the environment to trigger responses automatically when action effects activated by exogenous perceptual input correspond with those activated by one’s endogenously formed goals. The findings of Experiments 1-6 were entirely consistent with this claim. To illustrate, all experiments consistently showed that participants were more likely to produce action slips with the finger they currently imagined using than the non-imagined finger, independently of the observed action. This finding implies that action slips may be more similar to intention-driven than stimulus-driven actions because endogenously driven activation of action effect codes (i.e. via imagery) appeared to make a stronger contribution to the initiation of the corresponding action than perceptual input. Indeed, taken in isolation, this finding could be interpreted as evidence that a strong mental image is some kind of prerequisite for action slips to occur. This proposal is similar to the notion that a broad prospective intention to produce the imagined movement when appropriately cued (i.e. as per task instructions) provides the preconditions for the environment to subsequently trigger the immediate intention to act - initiating the imagined response (Pacherie, 2008; Pacherie & Haggard, 2010). However, the findings of Experiment 6 demonstrated that external stimuli (e.g. observed actions) can elicit the corresponding action, even in the absence of any specific endogenous activation of the goal. Recall that in Experiment 6, when participants’ hand were occluded (in the control ownership condition), visual action cues enhanced the execution of the corresponding action, even when it did *not* match

the participants imagined movement. These enhanced movements of the non-imagined finger are more evocative of stimulus-driven actions, in which an observed movement reflexively triggered the corresponding response so substantially that it was able to overcome competition with the imagined (i.e. intention-driven) movement.

Considering the role of both stimulus- and intention-driven activation of action effects, two aspects of these findings are of note. First, the interactions observed in all experiments showed that action slips were most likely when the visual stimulus matched the imagined (i.e. intended) movement, to an extent that did not simply reflect a simple addition of the main effects of imagery- and cue-congruence (i.e. independently of one another). In fact, the effect was super-additive. That is, the same visual cue made action slips both much more likely (Experiments 1-3) and of significantly greater magnitude (Experiments 4-5, but not Experiment 6) when it corresponded with the currently imagined finger than the non-imagined finger. This finding can readily be explained if one assumes that the endogenous (i.e. intention-driven) activation of a set of action effect codes brings the corresponding response closer to the threshold for execution, such that the additional influence of observing a compatible movement readily pushes the action super-threshold. In contrast, the finger which is not currently imagined is much further from the threshold for execution and the same visual cue does nothing (unless participants hands were occluded, as in Experiment 6). This conceptualisation parallels with threshold models of voluntary action which describe the moment of action initiation in terms of an integration-to-bounds mechanism (Churchland et al., 2008; Murakami et al., 2014; Schurger et al., 2016). In such models, the evidence in favour of candidate actions accumulates until the motor threshold is reached and the most suitable action is initiated. Experiments 1-6 indicate that this proposed process of “evidence accumulation” may correspond with the

activity in multimodal action effect codes. Actions are initiated when an internal representation of an action's perceptual consequences becomes sufficiently active.

Second, it is important to emphasise that both Experiments 5 and 6 (Chapter 4) revealed the effect of observing movements during imagery was reduced when participants experienced a body ownership illusion, relative to a control group without the illusion. This evidence runs counter to the Associative Sequence Learning theory of imitation, in which visual stimulation activates the corresponding action plan directly via associations acquired through repeated correlated experience of initiating actions and then perceiving movement contingent effects (e.g. Heyes, Bird, Johnson, & Haggard, 2005; Heyes & Ray, 2000). Indeed, according to this approach, one might have expected that the ownership illusion would have had no effect or, if anything, enhance the effect as a result of strengthening the correspondence between participants' movements and the movements of the virtual hand. Accordingly, the opposite results of the manipulation of ownership (Experiment 6) require a quite different explanation. In line with ideomotor views, it was suggested that this finding provides evidence that sensorimotor processes make use of two, relatively independent, types of representation; one which captures the estimated current state of the body and one which refers to anticipated future goal states. Under this view, visual inputs (e.g. an observed action) can only engage the motor system if they are integrated with one's goal state and, importantly, not if they are integrated with one's current state (i.e. the observed movements are *not* considered self-generated). That is, the ownership illusion promotes the integration of visual sensory input with the current state (correspondingly decreasing the contribution to the goal state) and, as a result, reduces the magnitude of the intention insertion effect.

The notion of separate representations of current and goal states becomes especially important in light of the results of Experiment 7 (Chapter 5). Experiment 7

revealed a positive relationship between individual differences in ideomotor suggestibility and the magnitude of the intention insertion effect (i.e. the extent to which perceptual cues modulated the imagined response). This positive relationship could suggest that some people are more prone to producing action slips because of a natural tendency to focus on anticipated future goal states, rather than their current state. In this view, a tendency to focus on anticipated, rather than actual sensations may be the driving force behind responsiveness to all kinds of hypnotic suggestion, not just ideomotor suggestions. For example, visual hallucinations are considered to be some of the most difficult hypnotic suggestions to experience and are, typically, the preserve of so-called “virtuoso” hypnotic subjects (Woody & Sadler, 2005). This may reflect the fact that visual input, as the dominant sense, is very precise and, therefore, requires an especially strong focus on the anticipated (i.e. imagined) visual stimulation to experience. These “virtuoso” subjects may so inclined to attend to anticipated visual stimulation meaning that it more readily obscures the true visual input – creating the visual illusion.

Taken together, these findings suggest that non-willed behaviour really is a curious by-product of the processes which are implicated in volitional action control. Actions are initiated when an internal representation of the associated perceptual consequences becomes sufficiently active. This may, therefore, explain why many everyday actions are executed seemingly automatically, simply because they match the goals one has previously committed to and the structure of the environment (Bargh & Chartrand, 1999; Kunde et al., 2007). Further, these findings may account for why one’s longer-term action plans (e.g. healthy eating) are followed more easily when one has visualised one’s behaviour in the relevant situation (Brandstatter, Lengfelder, & Gollwitzer, 2001). They may also explain the wide range of action slips that people make, where they sometimes perform unintended actions (e.g., picking one’s nose in

public, taking a wrong turn in the car), find themselves copying the postures and mannerisms of others (Chartrand & Bargh, 1999), or respond, in a seemingly unintentional way, to ideomotor hypnotic suggestions. This view suggests that the purpose of a broader action goal (e.g. commute to work) is to specify light constraints (e.g. maintaining speed) and sensorimotor contingencies (e.g. green light means go) which allow for the smooth translation of sensory input into motor output. Furthermore, these findings potentially link the generation of non-willed action slips, not just to mechanisms involved in volitional action control, but to illusions of predictive processing found in other, non-motor, domains (e.g. visual hallucination).

6.4. Volition in Action

At this stage, it is important to consider the implications that this research has for the broader concept of volition in action – a diffuse concept without a clear definition (Haggard, 2019). The discussion thus far has focused on the mechanisms which gave rise to the action slips captured by the intention insertion task. However, the findings clearly have relevance to the more general question of the extent to which we control our behaviour in everyday life – the notion of willed action. To this end, the responses observed in Experiments 1-6 are discussed in the context of several of the characteristic features of volitional actions described in a recent review by Haggard (2019).

6.4.1. Volitional Acts are Endogenously Driven

One feature which is commonly attributed to volitional actions is that of being endogenously generated (i.e. intention-driven), as opposed to a reaction to some external stimulus (i.e. stimulus-driven; Herwig, Prinz, & Waszak, 2007). On one

account, volitional acts are those in which endogenous processes, rather than exogenous perceptual input, determine *whether* an action should be made, *what* form it should take and *when* it should be executed (i.e. the WWW model of Brass & Haggard, 2008). As noted in the previous section, action slips appear to provide an intermediate example of an action that is partially driven (and specified) by internal processes (i.e. imagery) and partly driven by external stimulation. For example, Experiments 1-5 demonstrated that observed movements had a substantially larger effect if the participants also imagined a similar action. This suggests that endogenous processes (i.e. imagery) were largely responsible for determining the form (i.e. the “what” component) of the accidentally produced action, but the visual cues were able to override the “whether” and “when” attributes of the imagined action, in effect pushing the pre-activated action goal over the response threshold. Then again, as mentioned earlier in this chapter, the data could also be conceptualised in light of an integration-to-bounds mechanism (Churchland et al., 2008; Murakami et al., 2014; Schurger et al., 2016) if one assumes that activated action effects constitute the evidence accumulated by such a mechanism. In such a model, the activated action effects both specify the form of the action and - via the strength of their activation relative to the response threshold – whether and when it is executed.

It is interesting to note that, in the normal condition (i.e. no ownership illusion) of Experiment 6, in which participants hands were occluded, visual cues primed the corresponding response so substantially that this action was enhanced even when participants did not concurrently imagine moving this finger. Similarly, even in the original experiments (Chapter 2), external stimuli could sometimes cause participants to press the alternative finger that they did not currently imagine moving. This suggests that external stimuli, when sufficiently salient, are able to override endogenously specified actions and dictate a different form (i.e. the “what”) of initiated action, in addition to the “whether” and “when” components. This conceptualisation is entirely

consistent with the assumed common medium of perceptual action effect activation that dictates all three of these components (e.g. Common Coding Theory, Prinz, 1990, 1997; van der Wel, Sebanz, & Knoblich, 2013). For pragmatic reasons, it may be useful to conceptualise volition in action in terms of a continuum between those which are more strongly determined by either endogenous, intentional process or exogenous stimuli in the environment. In this context, the action slips observed at the participant's non-imagined finger, which were more strongly determined by exogenous stimulation, could be considered *less* volitional than those which corresponded with the participants' imagined movements.

6.4.2. Volitional Acts are Goal-directed

Another feature typically attributed to volitional actions is that they are generated for a reason. That is, volitional actions can be considered a means of achieving or advancing towards desirable goal states (i.e. they are goal-directed). In this context, the responses collected by Experiments 1-6 could be considered non-willed because they did not advance the explicitly stated task goal of remaining still during imagery. However, this is not to suggest that the action slips generated by the intention insertion task were not goal-directed in any way. Indeed, Experiments 1-6 demonstrated how one's intended goal to engage in movement imagery and to simulate the perceptual effects of finger movements created fertile preconditions for the environment to trigger responses which violated the higher-level (and explicitly stated) goal of remaining still. In this way, one could argue that imagery and perceptual input were each able to activate lower-level and more immediate goal states (with the potential to activate the motor system) which corresponded with the imagined and observed movements. From this perspective, the effect of movement observation on participants' non-imagined

responses (Experiment 6, control ownership group) still advanced their behaviour towards an outcome that was, in a broad sense, similar to their imagined goal. For example, the observed and imagined actions both corresponded with a downward movement of an index finger, even though the movements mismatched in terms of laterality (e.g. observing a left finger tap while imagining a right tap). However, in Experiment 6, the visual input made a larger contribution to the initiation and form of the executed action and was sufficiently salient to overcome competition with the participant's imagined goal (i.e. due to mismatching laterality of observed and imagined movements). That is, when visual cues elicited action slips with participants' non-imagined fingers, the intention to produce these movements was strongly driven by the environment, rather than largely driven by an endogenously evoked goal. In this context, one might consider the action slips produced by participants' non-imagined fingers to be *less* volitional than those produced by the finger the participant imagined using because the endogenously activated goal had less of an influence over the initiated action.

6.4.3. Volitional Acts Have a Distinctive Subjective Experience

Another feature which is claimed to distinguish volitional actions from more habitual, automatic actions is the distinctive subjective experience of having caused a particular movement. Experiments 1-6 did not directly assess participants' sense of awareness and agency, so it is beyond the scope of these data to draw conclusions about participants' subjective experiences of having produced action slips. This was a deliberate decision made on the basis that any attempt to question participants about their responses (e.g. asking "did you produce a response?" following each visual cue) during the task itself would draw unwanted attention to the purpose of the manipulation

and may have influenced the subsequent responding. However, it is possible to indirectly infer that participants were aware of having responded due to the characteristic decline in responsiveness during the course of the keyboard-based intention insertion task reported in Experiments 1-3 (Chapter 2). It suggests that participants actively either inhibited their responses after they noticed responding in error following the first visual cue, or they increased their response threshold and, in effect, required greater activation (or, “evidence” in threshold models, e.g. Churchland et al., 2008) to execute a given movement. Anecdotally, during the experimental debrief of all experiments, participants commonly reported that they were aware they were responding in error and that it was “effortful” to avoid doing so. However, the precise nature of the subjective experience of producing action slips, which appear to have characteristics of both willed and non-willed behaviour, remains an open question for future research.

6.4.4. The Broader Question of Volitional Control of Mental Events

Haggard’s (2019) review deliberately focused on the notion of volition as it relates to action initiation and motor control, and explicitly excluded the question of volitional control of purely mental contents (e.g. thoughts, images, memories, and emotions). However, constraining the discussion in this way naturally obscures the dynamic nature and context sensitivity of the processes which generate overt behaviour (Schurger & Uithol, 2015; Uithol, Burnston, & Haselager, 2014). Indeed, while it may seem obvious that some thoughts (or mental images) are more strongly “motor” than others (e.g. imagining a ball flying through the air vs. imagining throwing it); in ideomotor models (e.g. TEC; Hommel, 2009), both can drive action equally. What matters is only whether the mental images can serve as potential action outcomes,

irrespective of whether they are proximally related to the motor behaviours (e.g. imagining the actual hand movements) or distally (e.g. the specific ball trajectory one intends to achieve), as long as they are reliably associated with the motor command that produced them.

Note also that the extent to which an external stimulus may trigger a thought that closely corresponds with a specific action outcome will vary according to context. For example, one might encounter a pop-up advert for beer which simply consists of the name of the brewery. If coincidentally, there was a glass of this brand of beer within reach, one might feel compelled by the advert to drink from it. In this case, the initial visual stimulus had no obvious motor content, but the chain of events led to the evocation of a specific action goal (i.e. sip from the glass) which produced the corresponding behaviour. The specific action goal that produced the movement appears to be endogenously formed and, yet, the preceding stimulus (i.e. pop-up advert) can only be said to have caused the action in a rather indirect way. Consider then, whether it is useful or appropriate to attempt to trace the causal chain of events backwards in time and, indeed, whether such a search is likely to yield a single, focal origin for any given action. The same is true for the action slips investigated here.

The intention insertion task, by design, considers only the specific conditions that were associated with the 800ms critical interval following the onset of a surprising visual cue (i.e. the identity of the imagined and observed movement). However, as mentioned earlier, participants may well have formed a broadly specified *prospective* intention (Pacherie, 2008; Pacherie & Haggard, 2010) to respond (in any way) to visual stimuli. Such a prospective intention may then have created the preconditions for the subsequent observed movements to elicit responses (i.e. to insert an *immediate* intention to act) and, in some sense, “caused” the subsequent action slips. That said, it must be emphasised that this seems antagonistic with the instruction to ignore “distracting

stimuli” that was explicitly and repeatedly stated in Experiments 1-6. Indeed, if a prospective intention to “respond to visual stimuli” exerted a greater influence on the generation of action slips than the interaction of the participants’ imagery and the observed movement, then one would have expected participants to produce highly stereotyped responses (e.g. only producing the imagined response to all kinds of visual cue). Instead, the robust compatibility effect observed in Experiments 1-6 demonstrates that the interaction of imagery and perceptual input had a substantial effect on both the form and timing of the subsequent action slips. Crucially, when participants’ hands were occluded (Experiment 6, control ownership group), the observed movements activated the corresponding response so substantially that there was clear evidence of an intention insertion effect, even for responses which mismatched participants’ imagery. This finding, and the fact it was *not* evident in Experiments 4-5, implies that perceptual inputs really can *cause* movements (i.e. actions that would, otherwise, not be executed). Together, these findings suggest that it is the moment-by-moment dynamic interaction of endogenous mental imagery and exogenous perceptual input makes a strong *causal* contribution to the generation of action slips. Furthermore, this dynamic interaction has a stronger influence on the form and timing of the action slips produced in any given moment than a vague prospective intention (e.g. an “early whether” decision in the WWW model of Brass & Haggard, 2008) to respond to visual stimuli.

6.4.5. Conclusion

According to the criteria for volitional actions discussed by Haggard (2019), the responses captured by Experiments 1-6 appear to share features of both willed and non- (or, at least, *less-*) willed acts. Together, therefore, the data from Experiments 1-7 suggest that while typical motor behaviour is always goal-directed – and therefore, in

some sense “volitional” – people might not have full control of their goals. The data suggest that visual stimuli can, in effect, insert lower level goals into participants’ action control mechanisms. The characteristic feeling of involuntariness which tends to accompany accidental slips of action control might be better described as the feeling of transitioning towards an unexpected (or inappropriate) goal, rather than noticing that one has executed an incorrect action which may not have been consciously specified (e.g. the specific trajectory for a reaching and grasping movement).

On the basis of these findings, one might suggest that a more useful description of "volition" in action may be the extent to which the internal representations that produce a given movement are driven by top-down imagery, rather than bottom-up perceptual input. However, if the internal representation responsible for the initiated action is evoked endogenously, but as the consequence of a longer-term train of thought evoked by an external stimulus (e.g. the beer advert mentioned previously), then this distinction obscures some of the complexity in volitional action control. Indeed, as noted by Schüür and Haggard (2011), the role of volition in action may be to integrate sensory inputs and prior knowledge over an extended time period or to direct attention towards stimuli which support one’s long-term goals in order to guide behaviour. As a result, when considering the extent to which a given action is “intentional”, it is important to consider all the antecedent events which might have given rise to the activation of the specific action effect representation which triggered the executed action.

6.5. Open Questions and Future Directions

In addition to addressing questions about the relationship between imagery and perception on the initiation of non-willed actions, the current thesis also motivates

several lines of inquiry for future research. In the following section, some interesting open questions are discussed, along with some potential avenues for future research designed to address them.

6.5.1. Subjective Experience of Non-Willed Action

Volitional actions are characteristically self-referential and are typically distinguished from reflex actions by both a characteristic feeling of awareness that an action has taken place and a higher-order feeling of one's own role in having caused the behaviour to occur (i.e. a sense of agency). Indeed, one might define the subjective sense of volition as all that separates the feeling that "I" pressed a key with my finger, from the fact that my index finger moved downwards (Patrick Haggard, 2019). In the previous section, it was noted that Experiments 1-6 were not designed to assess participants' subjective experience of producing non-willed behaviour and, as a result, this is an obvious line of enquiry. Specifically, it was considered important to avoid drawing too much attention to the purpose of the visual cues (i.e. to initiate actions) by asking participants about any action slips that they might have produced. It was assumed that by asking participants to scrutinise their behaviour, the visual cues would have been less effective at inducing movements. However, now that the pressure-based intention insertion task (e.g. Chapter 3, Experiment 4) has been validated and the intention insertion effect has been replicated and extended (Chapter 4, Experiments 5-6), it is important to question this assumption.

One might begin by creating a new version of the intention insertion task in which, after each visual cue, participants reported, (a) whether they produced a response following the onset of each cue, and; (b) which response they executed, and; (c) a rating of their subjective sense of agency or control with respect to the reported movement

(e.g. on a Likert scale). Such a design would motivate four interesting questions. First, to what extent are people aware they are making non-willed responses? Second, are those individuals who are more susceptible to producing the effect (and therefore produce larger overt responses) more or less aware of their behaviour? Third, would asking people to make self-referential judgements about their own behaviour drastically reduce the size of the overall size of the intention insertion effect? Fourth, and finally, do people feel a general lack of agency with respect to the induced movements or do judgements of agency interact with the compatibility between participants' imagery and the observed movement. This latter question is particularly interesting because Experiments 4-6 consistently demonstrated that participants produced the largest action slips when the observed movement was compatible with their imagined movement. However, research has revealed that priming participants, just prior to execution, with compatible action effect information can actually enhance explicit judgements of agency, compared to when those effects are incompatible with the executed action (Aarts, Custers, & Wegner, 2005). This motivates the counterintuitive prediction that people might experience a greater sense of agency following trials in which they produced the strongest action slips (i.e. when the imagined and observed movements were compatible) than following trials in which participants produced the weakest action slips (i.e. when the imagined and observed movements were incompatible).

6.5.2. Direct Assessment of Current/Goal State Proposal

Another important avenue for future research is the proposal that sensorimotor processes make use of two, relatively separate internal representations which capture either; (1) one's estimated current state, or; (2) anticipated goal states (Chapter 4). While Experiments 5 and 6 provided evidence that was consistent with this

conceptualisation, the use of an ownership illusion complicates the interpretation of these data. In order to provide a more direct test of whether action emerges from the conflict between one's current state and an activated goal state, the intention insertion task could be slightly modified. To illustrate, participants could be presented with a number in place of the central fixation cross to indicate their current finger position (i.e. the current state) which ranged from one (i.e. resting position) to five (i.e. key fully depressed). This would provide the opportunity to present participants with false feedback about their current state immediately prior to the presentation of visual cues (i.e. inserted goal state) and, by extension, would allow the magnitude of the difference between the current state and the goal state to be manipulated. If, as proposed in Chapter 4, action is a consequence of a discrepancy between one's current state and an active goal state, then artificially reducing this discrepancy by implying that participants' fingers are already partially or completely depressed should limit or eliminate the motor output.

6.5.3. Attending to Current or Anticipated Sensation in Predictive Processing

Experiment 7 (Chapter 5) provided evidence in support of a positive relationship between individual differences in the tendency to produce action slips and individual differences in ideomotor suggestibility (a sub-component of the broader trait of hypnotic suggestibility, Woody, Barnier, & McConkey, 2005). One explanation for this finding is that people were more prone to producing non-willed behaviour in both the main task and screening procedure due to a natural tendency to focus on (and integrate sensory information with) anticipated goal states than the estimated current state. The notion that there may be strong individual differences in the manner in which perceptual input is processed (i.e. whether sensory information weighted towards the current or

goal state) could, therefore, be of relevance to research regarding other forms of predictive processing, such as with visual perception and the representational momentum effect. For example, it has been demonstrated that people make a small but consistent error when asked to estimate the position of moving objects. The representational momentum effect describes how people will reliably judge the object to be further along its predicted trajectory than its true location – as if their percept is modulated by their expectations (Freyd & Finke, 1984). More recently, these findings have been extended to the domain of predictive social processing – how we predict and interpret the behaviour of others. Research suggests that one’s perception of an inefficient reaching motion (i.e. one that takes an indirect route from A to B) shows a consistent bias towards the efficient trajectory (i.e. a direct route from A to B) that one would expect based on prior experience of reaching and grasping movements (Hudson, McDonough, Edwards, & Bach, 2018). These kinds of representational momentum effects could be accounted for if one assumes that perceptual judgements are a product of both current and anticipated visual input. For example, one would expect highly suggestible people to show a stronger representational momentum effect because of a tendency to focus on anticipated goals (i.e. the efficient route), rather than the current state (e.g. the observed inefficient route). This is because, on this logic, the perceptual judgements of highly suggestible people would be more strongly informed by anticipated sensation than actual sensation in comparison to less suggestible people. Accordingly, if individual differences in the tendency to produce action slips reflect a predisposition to attend to current or anticipated internal representations (both within and without the motor domain) then the intention insertion effect (and, indeed, ideomotor suggestibility) should demonstrate a positive relationship with the representational momentum effect.

6.5.4. The Brain Basis of Volition in Action

Research concerning the neural origins of voluntary action defines “volition” as “a state or process that leads to action” (Patrick Haggard, 2019). Under this view, it is argued that voluntary actions can be distinguished from more reflexive, stimulus-driven actions in the brain by using neuroimaging techniques to identify the unique neural networks recruited during the execution of each kind of action. Decades of neuroimaging research has implicated several cortical circuits which predominantly serve either internally generated “volitional” actions or stimulus-driven responses to environmental stimuli (for review, see Haggard, 2008). These circuits, it is claimed, converge on the primary motor cortex (M1) which provides a “final common path” for action as the signals are transmitted via the spinal cord to the relevant effector. If the action slips generated by the intention insertion task represent an “intermediate” example of behaviour, then it would be interesting to investigate whether their neural activation corresponded more closely with volitional actions or reflexive responses to external stimuli.

Two important frontal regions which project afferent signals to M1 are the Supplementary Motor Area (SMA) and the Premotor Cortex (PMC; Haggard, 2019). While both regions are involved in the coordination of complex movements, there is evidence of a double dissociation in their role in generating action. For example, a study which applied repetitive transcranial magnetic stimulation (rTMS) during different types of action suggested that the SMA is preferentially involved in the preparation of voluntary, rather than externally cued actions, whereas the PMC shows the opposite pattern (Lu, Arai, Tsai, & Ziemann, 2012). If afferent signals from the SMA reflect the contribution of endogenous imagery and signals from the PMC reflect the contribution of exogenous perceptual input, then this could suggest that action slips occur when the combined activity of the SMA and the PMC crosses the threshold to engage M1. This

proposal provides an indication of how the behaviour generated by the intention insertion task could be used to gain novel insights into the neural basis of volition in action.

6.6. Conclusion

In conclusion, the current thesis provides direct, behavioural evidence in favour of some of the fundamental assumptions about the mechanisms of voluntary action first made by 19th century ideomotor theorists (Carpenter, 1852; James, 1890), which were then formalised by modern theories of ideomotor action control (e.g. Theory of Event Coding, Hommel, 2009; Hommel, Müsseler, Aschersleben, & Prinz, 2001), extended by computational models (Wolpert, 1997; Wolpert & Ghahramani, 2000) and, most recently, recast in terms of Bayesian predictive processing (Adams et al., 2013; Clark, 2013). When reviewing the diverse assortment of models of motor control, one is struck by the extent to which they are in broad agreement about a central feature of voluntary action – that action, perception and mental imagery are inextricably linked because they rely on a common neural infrastructure and are coded in a fundamentally perceptual format (e.g. Common Coding theory, Prinz, 1990, 1997; van der Wel, Sebanz, & Knoblich, 2013).

The experiments presented in this thesis support the assumption that action intentions are fundamentally anticipatory (i.e. intentions correspond with future goal states). Furthermore, the data were consistent with the notion that forming an intention to act involves nothing more than evoking a mental image (i.e. activating action effects) associated with the perceptual effect one wishes to achieve (i.e. an action goal) until a motor threshold is reached and the corresponding action is initiated. Crucially,

Experiments 1-6 demonstrated that both endogenous (i.e. imagery) and exogenous processes (i.e. perception) can modulate the activation of action effect representations and, by extension, engage the motor system and cause action slips – even when physical responses were explicitly prohibited. Taken together, these findings provide an important insight into how people are able to interact so effortlessly with the world by specifying lightly constraining action goals (e.g. make a cup of tea) to mediate the translation of sensory input into motor output - allowing the environment to directly trigger actions which correspond with the goal (i.e. the precise set of sensorimotor transformations required to make tea). In addition, these findings are consistent with the notion that accidental slips of action control are a product of the same processes that guide voluntary actions – wherein perceptual input can co-opt internally generated action plans to trigger inappropriate actions – as proposed at the start of this thesis. These experiments also reveal why errors of action control can often be ironic (Wegner, 1994; Wegner et al., 1998) such that fixating on avoiding a given behaviour can sometimes produce the very action one was trying to prevent (e.g. stifling a laugh at an inappropriate joke at a funeral).

Extending these findings, Experiment 7 provided evidence that the tendency to produce action slips exhibits a positive relationship with ideomotor suggestibility. This finding is significant, because it may suggest that some people are more prone to producing action slips due to a natural tendency to focus on anticipated future goal states, rather than their current state. In fact, individual differences in the tendency to attend to anticipated, rather than actual sensations could account for curious errors of predictive processing found in other, non-motor domains, such as hypnotic visual hallucination (Woody & Sadler, 2005) and the representational momentum effect (Freyd & Finke, 1984; Hudson et al., 2018).

In part, the design of the intention insertion task was motivated by a desire to provide tangible, behavioural evidence for some of the common features of modern models of motor control, which are often discussed (and indeed, researched) in highly abstracted terms. For example, Friston's (2010) concept of "free energy", which forms the basis for Active Inference (Adams et al., 2013), has the potential to provide a unified brain theory with the power to account for all cognitive functions and, further, all self-organising biological systems. Yet, despite the potentially groundbreaking implications of this theoretical framework, it is routinely derided for being unintelligible to all but the most devoted of minds (see, "How to Read Karl Friston (In the Original Greek)", Maren, 2017). The findings produced by the intention insertion task demonstrate the utility of using accidental behaviour, induced by imagery and perception, as a device to directly test these esoteric claims about how volitional actions are generated. However, a natural consequence of developing a novel experimental method is that the findings it generated have motivated further questions about, among other things, participants' subjective experience of producing non-willed action slips and the brain basis for volition in action. As a consequence, the experiments presented in this thesis represent a starting point for further exploration with many, potentially fruitful, lines of enquiry.

References

- Aarts, H., Custers, R., & Wegner, D. M. (2005). On the inference of personal authorship: Enhancing experienced agency by priming effect information. *Consciousness and Cognition*, 14(3), 439–458.
<https://doi.org/10.1016/j.concog.2004.11.001>
- Adams, R. A., Shipp, S., & Friston, K. J. (2013). Predictions not commands: Active inference in the motor system. *Brain Structure and Function*, 218(3), 611–643.
<https://doi.org/10.1007/s00429-012-0475-5>
- Armel, K. C., & Ramachandran, V. S. (2003). Projecting sensations to external objects: Evidence from skin conductance response. *Proceedings of the Royal Society B: Biological Sciences*, 270(1523), 1499–1506.
<https://doi.org/10.1098/rspb.2003.2364>
- Bach, P., Allami, B. K., Tucker, M., & Ellis, R. (2014). Planning-related motor processes underlie mental practice and imitation learning. *Journal of Experimental Psychology: General*, 143(3), 1277–1294. <https://doi.org/10.1037/a0035604>
- Bach, P., Bayliss, A., & Tipper, S. P. (2011). The predictive mirror: interactions of mirror and affordance processes during action observation. *Psychonomic Bulletin & Review*, 18(1), 171–176. <https://doi.org/10.3758/s13423-010-0029-x>
- Bach, P., Griffiths, D., Weigelt, M., & Tipper, S. P. (2010). Gesturing Meaning: Non-action Words Activate the Motor System. *Frontiers in Human Neuroscience*, 4(November), 214. <https://doi.org/10.3389/fnhum.2010.00214>
- Bach, P., Peatfield, N. A., & Tipper, S. P. (2007). Focusing on body sites: The role of spatial attention in action perception. *Experimental Brain Research*, 178(4), 509–

517. <https://doi.org/10.1007/s00221-006-0756-4>

Badets, A., Koch, I., & Philipp, A. M. (2016). A review of ideomotor approaches to perception, cognition, action, and language: advancing a cultural recycling hypothesis. *Psychological Research*, 80(1), 1–15. <https://doi.org/10.1007/s00426-014-0643-8>

Bargh, J. A., & Chartrand, T. L. (1999). The Unbearable Automaticity of Being. *American Psychologist*, 54(7), 462–479. <https://doi.org/10.1037/0003-066X.54.7.462>

Barnier, Amanda J, Dienes, Z., & Mitchell, C. J. (2008). How hypnosis happens: New cognitive theories of hypnotic responding. In A. J. Barnier & M. R. Nash (Eds.), *The Oxford handbook of hypnosis: Theory, research, and practice* (pp. 141–177). New York, NY: Oxford University Press.

Barnier, Amanda J, & McConkey, K. M. (1999). Absorption, hypnotizability and context: non-hypnotic contexts are not all the same. *Contemporary Hypnosis*, 16(1), 1–8. <https://doi.org/10.1002/ch.144>

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1–48. <https://doi.org/10.18637/jss.v067.i01>

Berends, H. I., Wolkorte, R., Ijzerman, M. J., & Van Putten, M. J. A. M. (2013). Differential cortical activation during observation and observation-and-imagination. *Experimental Brain Research*, 229(3), 337–345. <https://doi.org/10.1007/s00221-013-3571-8>

Botvinick, M. M., & Cohen, J. (1998). Rubber hand feels touch that eyes see. *Nature*,

391(February), 756. <https://doi.org/10.1038/35784>

Bowers, K. S. (1998). Waterloo-stanford group scale of hypnotic susceptibility, form C: Manual and response booklet. *International Journal of Clinical and Experimental Hypnosis*, 46(3), 250–268. <https://doi.org/10.1080/00207149808410006>

Brandstatter, V., Lengfelder, A., & Gollwitzer, P. M. (2001). Implementation Intentions and Efficient Action Initiation. *Journal of Personality and Social Psychology*, 81(5), 946–960. <https://doi.org/10.1037//0022-3514.81.5.946>

Brass, M., Bekkering, H., & Prinz, W. (2001). Movement observation affects movement execution in a simple response task. *Acta Psychologica*, 106(1–2), 3–22. [https://doi.org/10.1016/S0001-6918\(00\)00024-X](https://doi.org/10.1016/S0001-6918(00)00024-X)

Brass, M., Bekkering, H., Wohlschläger, A., & Prinz, W. (2000). Compatibility between observed and executed finger movements: Comparing symbolic, spatial, and imitative cues. *Brain and Cognition*, 44(2), 124–143. <https://doi.org/10.1006/brcg.2000.1225>

Brass, M., & Haggard, P. (2008). The what, when, whether model of intentional action. *Neuroscientist*, 14(4), 319–325. <https://doi.org/10.1177/1073858408317417>

Brown, H., Adams, R. A., Parees, I., Edwards, M., & Friston, K. (2013). Active inference, sensory attenuation and illusions. *Cognitive Processing*, 14(4), 411–427. <https://doi.org/10.1007/s10339-013-0571-3>

Brown, P., Rothwell, J. C., Thompson, P. D., Britton, T. C., Day, B. L., & Marsden, C. D. (1991). New observations on the normal auditory startle reflex in man. *Brain*, 114(4), 1891–1902. <https://doi.org/10.1093/brain/114.4.1891>

Bürkner, P.-C. (2017). brms: An R Package for Bayesian Multilevel Models Using Stan.

- Journal of Statistical Software*, 80(1), 1–28. Retrieved from <https://www.jstatsoft.org/v080/i01>
- Calabresi, P., Picconi, B., Tozzi, A., Ghiglieri, V., & Di Filippo, M. (2014). Direct and indirect pathways of basal ganglia: a critical reappraisal. *Nature Neuroscience*, 17(8), 1022. <https://doi.org/10.1038/nn.3743>
- Cardeña, E., & Terhune, D. B. (2014). Hypnotizability, personality traits, and the propensity to experience alterations of consciousness. *Psychology of Consciousness*, 1(3), 292–307. <https://doi.org/10.1037/cns0000026>
- Carpenter, W. B. (1852). *On the Influence of Suggestion in Modifying and Directing Muscular Movement, Independently of Volition*. Proceedings of the Royal Institution. Retrieved from <https://books.google.co.uk/books?id=ajDTnAEACAAJ>
- Chartrand, T. L., & Bargh, J. A. (1999). The Chameleon Effect: The Perception-Behaviour Link and Social Interaction. *Journal of Personality and Social Psychology*, 76(6), 893–910. <https://doi.org/10.1037/0022-3514.76.6.893>
- Chevreul, M. E. (1833). Lettre à M. Ampère sur une classe particulière de mouvements musculaires. *Review Des Deux Mondes*, (2), 258–266.
- Churchland, A. K., Kiani, R., & Shadlen, M. N. (2008). Decision-making with multiple alternatives. *Nature Neuroscience*, 11(6), 693–702. <https://doi.org/10.1038/nn0708-851c>
- Cisek, P. (2007). Cortical mechanisms of action selection: the affordance competition hypothesis. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 362(1485), 1585–1599. <https://doi.org/10.1098/rstb.2007.2054>

- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *The Behavioral and Brain Sciences*, 36(3), 181–253.
<https://doi.org/10.1017/S0140525X12000477>
- Colton, J., Bach, P., Whalley, B., & Mitchell, C. (2018a). Intention insertion: Activating an action's perceptual consequences is sufficient to induce non-willed motor behavior. *Journal of Experimental Psychology: General*, 147(8), 1256–1263.
<https://doi.org/10.1037/xge0000435>
- Colton, J., Bach, P., Whalley, B., & Mitchell, C. (2018b). Intention insertion: activating an action's perceptual consequences is sufficient to induce non-willed motor behaviour (Analysis and Materials). Retrieved from <https://osf.io/s9qgu/>
- Cramer, A. O. J., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P. P. P., ... Wagenmakers, E. J. (2016). Hidden multiplicity in exploratory multiway ANOVA: Prevalence and remedies. *Psychonomic Bulletin and Review*, 23(2), 640–647. <https://doi.org/10.3758/s13423-015-0913-5>
- Dienes, Z., Lush, P., Semmens-Wheeler, R., Parkinson, J., Scott, R., & Naish, P. (2016). Hypnosis as self-deception; meditation as self-insight. In M. Raz, Amir and Lifshitz (Ed.), *Hypnosis and meditation: Towards an integrative science of conscious planes*. (pp. 107–125). Oxford: Oxford University Press.
- Drost, U. C., Rieger, M., Brass, M., Gunter, T. C., & Prinz, W. (2005). When hearing turns into playing: Movement induction by auditory stimuli in pianists. *The Quarterly Journal of Experimental Psychology*, (8), 58–1376.
<https://doi.org/10.1080/02724980443000610>
- Dum, R. P., & Strick, P. L. (2002). Motor areas in the frontal lobe of the primate. *Physiology & Behavior*, 77(4), 677–682. <https://doi.org/10.1016/S0031->

- Easton, R. D., & Shor, R. E. (1975). Information processing analysis of the Chevreul pendulum illusion. *Journal of Experimental Psychology: Human Perception and Performance*, 1(3), 231–236. <https://doi.org/10.1037/0096-1523.1.3.231>
- Easton, R. D., & Shor, R. E. (1976). An Experimental Analysis of the Chevreul Pendulum Illusion. *The Journal of General Psychology*, 95(1), 111–125. <https://doi.org/10.1080/00221309.1976.9710871>
- Easton, R. D., & Shor, R. E. (1977). Augmented and Delayed Feedback in the Chevreul Pendulum Illusion. *The Journal of General Psychology*, 97(2), 167–177. <https://doi.org/10.1080/00221309.1977.9920835>
- Elsner, B., & Hommel, B. (2001). Effect anticipation and action control. *Journal of Experimental Psychology: Human Perception and Performance*, 27(1), 229–240. <https://doi.org/10.1037/0096-1523.27.1.229>
- Freud, S. (1914). *The Psychopathology of Everyday Life*. New York: The Macmillan Company.
- Freyd, J. J., & Finke, R. A. (1984). Representational momentum. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 10(1), 126. <https://doi.org/10.1037/0278-7393.10.1.126>
- Friston, K. (2010). The free-energy principle: a unified brain theory? *Nature Reviews. Neuroscience*, 11(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (Comment on Article by Browne and Draper). *Bayesian Analysis*, 1(3), 515–534. <https://doi.org/10.1214/06-BA117A>

- Gentsch, A., Weber, A., Synofzik, M., Vosgerau, G., Schutz-Bosbach, S., & Schütz-Bosbach, S. (2016). Towards a common framework of grounded action cognition: Relating motor control, perception and cognition. *Cognition*, 146(April), 81–89. <https://doi.org/10.1016/j.cognition.2015.09.010>
- Glisky, M. L., Tataryn, D. J., & Kihlstrom, J. F. (1995). Hypnotizability and mental imagery. *International Journal of Clinical and Experimental Hypnosis*, 43(1), 34–54. <https://doi.org/10.1080/00207149508409374>
- Greenwald, A. G. (1970). Sensory feedback mechanisms in performance control: With special reference to the ideo-motor mechanism. *Psychological Review*, 77(2), 73–99. <https://doi.org/10.1037/h0028689>
- Grover, M. P., Jensen, M. P., Patterson, D. R., Gertz, K. J., & Day, M. A. (2018). The Association Between Mindfulness and Hypnotizability: Clinical and Theoretical Implications. *American Journal of Clinical Hypnosis*, 61(1), 4–17. <https://doi.org/10.1080/00029157.2017.1419458>
- Guérard, K., & Brodeur, M. (2015). Manipulability agreement as a predictor of action initiation latency. *Behavior Research Methods*, 47(2), 572–579. <https://doi.org/10.3758/s13428-014-0495-6>
- Haggard, Patrick. (2008). Human volition: Towards a neuroscience of will. *Nature Reviews Neuroscience*, 9(12), 934–946. <https://doi.org/10.1038/nrn2497>
- Haggard, Patrick. (2019). The Neurocognitive Bases of Human Volition. *Annual Review of Psychology*, 70(1), 9–28. <https://doi.org/10.1146/annurev-psych-010418-103348>
- Halligan, P. W., & Oakley, D. A. (2014). Hypnosis and beyond: Exploring the broader

domain of suggestion. *Psychology of Consciousness: Theory, Research, and Practice*, 1(2), 105–122. <https://doi.org/10.1037/cns0000019>

Herwig, A., Prinz, W., & Waszak, F. (2007). Two modes of sensorimotor integration in intention-based and stimulus-based actions. *The Quarterly Journal of Experimental Psychology*, 60(11), 1540–1554. <https://doi.org/10.1080/17470210601119134>

Heyes, C., & Bird, G. (2008). Mirroring, association, and the correspondence problem. In P. Haggard, Y. Rossetti, & M. Kawato (Eds.), *Sensorimotor Foundations of Higher Cognition* (Vol. 22, pp. 460–480). Oxford, UK: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199231447.003.0021>

Heyes, C., Bird, G., Johnson, H., & Haggard, P. (2005). Experience modulates automatic imitation, 22, 233–240. <https://doi.org/10.1016/j.cogbrainres.2004.09.009>

Heyes, C., & Ray, E. (2000). What Is the Significance of Imitation in Animals? In P. J. B. Slater, J. S. Rosenblatt, C. T. Snowdon, & T. J. Roper (Eds.), *Advances in the Study of Behavior* (Vol. 29, pp. 215–245). Academic Press. [https://doi.org/10.1016/S0065-3454\(08\)60106-0](https://doi.org/10.1016/S0065-3454(08)60106-0)

Hommel, B. (1995). Stimulus-response compatibility and the Simon effect: Toward an empirical clarification. *Journal of Experimental Psychology: Human Perception and Performance*, 21(4), 764–775. <https://doi.org/10.1037/0096-1523.21.4.764>

Hommel, B. (1997). Toward an action-concept model of stimulus-response compatibility. In B. Hommel & W. B. Prinz (Eds.), *Theoretical Issues in Stimulus-Response Compatibility* (Vol. 118, pp. 281–320). Amsterdam: North-Holland. [https://doi.org/10.1016/S0166-4115\(97\)80041-6](https://doi.org/10.1016/S0166-4115(97)80041-6)

- Hommel, B. (2009). Action control according to TEC (theory of event coding). *Psychological Research*, 73(4), 512–526. <https://doi.org/10.1007/s00426-009-0234-2>
- Hommel, B., Müsseler, J., Aschersleben, G., & Prinz, W. (2001). The Theory of Event Coding (TEC): A framework for perception and action planning. *Behavioral and Brain Sciences*, 24(05), 849–878. <https://doi.org/10.1017/S0140525X01000103>
- How to Make a Proper Brew. (2019). Retrieved January 4, 2019, from <https://www.yorkshiretea.co.uk/our-teas/how-to-make-a-proper-brew>
- Hudson, M., McDonough, K. L., Edwards, R., & Bach, P. (2018). Perceptual teleology: expectations of action efficiency bias social perception. *Proceedings. Biological Sciences*, 285(1884), 20180638. <https://doi.org/10.1098/rspb.2018.0638>
- Hull, C. L. (1943). *Principles of behavior: an introduction to behavior theory*. Oxford, UK: Appleton-Century.
- Jacobson, E. (1930). Electrical measurements of neuromuscular states during mental activities: 1. Imagination of movement involving skeletal muscle. *American Journal of Physiology*, 91, 567–608.
- James, W. (1890). *The principles of psychology (Vols. 1 & 2)*. New York, NY: Henry Holt and Company. <https://doi.org/10.1037/10538-000>
- Jamieson, G. A. (2016). A unified theory of hypnosis and meditation states: The interoceptive predictive coding approach. In A. Raz & M. Lifshitz (Eds.), *Hypnosis and meditation: Towards an integrative science of conscious planes*. (pp. 313–342). New York, NY: Oxford University Press.
- Jeannerod, M., & Decety, J. (1995). Mental motor imagery: a window into the

representational stages of action. *Current Opinion in Neurobiology*, 5(6), 727–732.

[https://doi.org/10.1016/0959-4388\(95\)80099-9](https://doi.org/10.1016/0959-4388(95)80099-9)

Kandyba, K., & Laurence, J. R. (1996). Hypnotic age-regression, cuing and memory creation. In *International Journal of Psychology* (Vol. 31, p. 4475). Sussex, UK: Psychology Press.

Keller, P. E., & Koch, I. (2006). The planning and execution of short auditory sequences. *Psychonomic Bulletin & Review*, 13(4), 711–716.

<https://doi.org/10.3758/BF03193985>

Kilner, J. M., Hamilton, A. F. de C., & Blakemore, S. J. (2007). Interference effect of observed human movement on action is due to velocity profile of biological motion. *Social Neuroscience*, 2(3–4), 158–166.

<https://doi.org/10.1080/17470910701428190>

Kilner, J. M., Paulignan, Y., & Blakemore, S. J. (2003). An interference effect of observed biological movement on action. *Current Biology*, 13(6), 522–525.

[https://doi.org/10.1016/S0960-9822\(03\)00165-9](https://doi.org/10.1016/S0960-9822(03)00165-9)

Knoblich, G., & Prinz, W. (2005). Linking perception and action: An ideomotor approach. In H.-J. Freund, J. Marc, H. Mark, & R. Leiguarda (Eds.), *Higher-order motor disorders: From Neuroanatomy and neurobiology to clinical neurology* (pp. 79–104). Oxford, UK: Oxford University Press.

Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247. <https://doi.org/10.1038/nature02169>

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, 10(7), 319–326.

<https://doi.org/10.1016/j.tics.2006.05.003>

Kraeutner, S., Gionfriddo, A., Bardouille, T., & Boe, S. (2014). Motor imagery-based brain activity parallels that of motor execution: Evidence from magnetic source imaging of cortical oscillations. *Brain Research*, 1588, 81–91.

<https://doi.org/10.1016/j.brainres.2014.09.001>

Kühn, S., Elsner, B., Prinz, W., & Brass, M. (2009). Busy doing nothing: Evidence for nonaction-effect binding. *Psychonomic Bulletin & Review*, 16(3), 542–549.

<https://doi.org/10.3758/PBR.16.3.542>

Kühn, S., Keizer, A., Rombouts, S. A. R. B., & Hommel, B. (2011). The functional and neural mechanism of action preparation: roles of EBA and FFA in voluntary action control. *Journal of Cognitive Neuroscience*, 23(1), 214.

<https://doi.org/10.1162/jocn.2010.21418>

Kunde, W. (2001). Response-Effect Compatibility in Manual Choice Reaction Tasks.

Journal of Experimental Psychology: Human Perception and Performance, 27(2),

387–394. <https://doi.org/10.1037/0096-1523.27.2.387>

Kunde, W., Elsner, K., & Kiesel, A. (2007). No anticipation-no action: The role of anticipation in action and perception. *Cognitive Processing*, 8(2), 71–78.

<https://doi.org/10.1007/s10339-007-0162-2>

Kunde, W., Koch, I., & Hoffmann, J. (2004). Anticipated action effects affect the

selection, initiation, and execution of actions. *Quarterly Journal of Experimental Psychology: Part A: Human Experimental Psychology*, 57(1), 87.

<https://doi.org/10.1080/02724980343000143>

Labelle, L., Dixon, M., & Laurence, J.-R. (1989). Pseudomemory creation and

confidence ratings: A theory still in search of supportive data. *British Journal of Experimental & Clinical Hypnosis*, 6(3), 163–165.

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): The unconscious initiation of a freely voluntary act. *Brain*, 106(3), 623–642.
<https://doi.org/10.1093/brain/106.3.623>

Lu, M. K., Arai, N., Tsai, C. H., & Ziemann, U. (2012). Movement related cortical potentials of cued versus self-initiated movements: Double dissociated modulation by dorsal premotor cortex versus supplementary motor area rTMS. *Human Brain Mapping*, 33(4), 824–839. <https://doi.org/10.1002/hbm.21248>

Lynn, S. J., Laurence, J.-R. R., & Kirsch, I. (2015). Hypnosis, Suggestion, and Suggestibility: An Integrative Model. *American Journal of Clinical Hypnosis*, 57(3), 314–329. <https://doi.org/10.1080/00029157.2014.976783>

Ma, K., & Hommel, B. (2015a). Body-ownership for actively operated non-corporeal objects. *Consciousness and Cognition*, 36, 75–86.
<https://doi.org/10.1016/j.concog.2015.06.003>

Ma, K., & Hommel, B. (2015b). The role of agency for perceived ownership in the virtual hand illusion. *Consciousness and Cognition*, 36, 277–288.
<https://doi.org/10.1016/j.concog.2015.07.008>

Maren, A. (2017). How to Read Karl Friston (In the Original Greek). Retrieved December 6, 2018, from <http://www.aliannajmaren.com/2017/07/27/how-to-read-karl-friston-in-the-original-greek/>

Maslovat, D., Chua, R., & Hodges, N. J. (2013). When unintended movements “leak”

out: A startling acoustic stimulus can elicit a prepared response during motor imagery and action observation. *Neuropsychologia*, 51(5), 838–844.
<https://doi.org/10.1016/j.neuropsychologia.2013.01.016>

McRobbie, L. R. (2013). The Strange and Mysterious History of the Ouija Board.

Retrieved February 13, 2018, from <https://www.smithsonianmag.com/history/the-strange-and-mysterious-history-of-the-ouija-board-5860627/>

Munakata, Y., Herd, S. A., Chatham, C. H., Depue, B. E., Banich, M. T., & O'Reilly, R.

C. (2011). A unified framework for inhibitory control. *Trends in Cognitive Sciences*, 15(10), 453–459. <https://doi.org/10.1016/j.tics.2011.07.011>

Murakami, M., Vicente, M. I., Costa, G. M., & Mainen, Z. F. (2014). Neural

antecedents of self-initiated actions in secondary motor cortex. *Nature Neuroscience*, 17(11), 1574–1582. <https://doi.org/10.1038/nn.3826>

Ondobaka, S., & Bekkering, H. (2012). Hierarchy of idea-guided action and perception-

guided movement. *Frontiers in Psychology*, 3(December), 1–5.

<https://doi.org/10.3389/fpsyg.2012.00579>

Pacherie, E. (2008). The phenomenology of action: A conceptual framework.

Cognition, 107(1), 179–217. <https://doi.org/10.1016/j.cognition.2007.09.003>

Pacherie, E., & Haggard, P. (2010). What are intentions? In W. Sinnott-Armstrong & N.

Lynn (Eds.), *Conscious Will and Responsibility. A tribute to Benjamin Libet*. (pp. 70–84). Oxford University Press, USA.

Picard, N., & Strick, P. L. (1996). Motor Areas of the Medial Wall: A Review of Their

Location and Functional Activation. *Cerebral Cortex*, 6(3), 342–353. Retrieved from <http://dx.doi.org/10.1093/cercor/6.3.342>

- Prinz, W. (1990). A Common Coding Approach to Perception and Action. In O. Neumann & W. Prinz (Eds.), *Relationships Between Perception and Action: Current Approaches* (pp. 167–201). Berlin, Heidelberg: Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-75348-0_7
- Prinz, W. (1997). Perception and Action Planning. *European Journal of Cognitive Psychology*, 9(2), 129–154. <https://doi.org/10.1080/713752551>
- R Core team. (2017). R: A Language and Environment for Statistical Computing. Vienna, Austria.: R Foundation for Statistical Computing. Retrieved from <https://www.r-project.org/>
- Ramsey, R., Cumming, J., Eastough, D., & Edwards, M. G. (2010). Incongruent imagery interferes with action initiation. *Brain and Cognition*, 74(3), 249–254. <https://doi.org/10.1016/j.bandc.2010.08.005>
- Reason, J. (1984). Lapses of Attention in Everyday Life. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of Attention* (pp. 515–549). New York, NY: Academic Press.
- Ridderinkhof, K. R., van den Wildenberg, W. P. M., & Brass, M. (2014). “Don’t” versus “Won’t”: Principles, mechanisms, and intention in action inhibition. *Neuropsychologia*, 65(September), 255–262. <https://doi.org/10.1016/j.neuropsychologia.2014.09.005>
- Rieger, M., Dahm, S., & Koch, I. (2017). Inhibition in motor imagery: a novel action mode switching paradigm. *Psychonomic Bulletin & Review*, 24(2), 459–466. <https://doi.org/10.3758/s13423-016-1095-5>
- Rizzolatti, G., Luppino, G., & Matelli, M. (1998). The organization of the cortical

motor system: new concepts. *Electroencephalography and Clinical Neurophysiology*, 106(4), 283–296. [https://doi.org/10.1016/S0013-4694\(98\)00022-4](https://doi.org/10.1016/S0013-4694(98)00022-4)

Sanders, A. F. (1980). Stage Analysis of Reaction Processes. In G. E. Stelmach & J. B. T.-A. in P. Requin (Eds.), *Tutorials in Motor Behavior* (Vol. 1, pp. 331–354). Amsterdam: North-Holland. [https://doi.org/10.1016/S0166-4115\(08\)61955-X](https://doi.org/10.1016/S0166-4115(08)61955-X)

Schurger, A., Mylopoulos, M., & Rosenthal, D. (2016). Neural Antecedents of Spontaneous Voluntary Movement: A New Perspective. *Trends in Cognitive Sciences*, 20(2), 77–79. <https://doi.org/10.1016/j.tics.2015.11.003>

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(42), E2904–E2913. <https://doi.org/10.1073/pnas.1210467109>

Schurger, A., & Uithol, S. (2015). Nowhere and Everywhere: The Causal Origin of Voluntary Action. *Review of Philosophy and Psychology*, 6(4), 761–778. <https://doi.org/10.1007/s13164-014-0223-2>

Schütz-Bosbach, S., Mancini, B., Aglioti, S. M., & Haggard, P. (2006). Self and Other in the Human Motor System. *Current Biology*, 16(18), 1830–1834. <https://doi.org/10.1016/j.cub.2006.07.048>

Schüür, F., & Haggard, P. (2011). What are self-generated actions? *Consciousness and Cognition*, 20(4), 1697–1704. <https://doi.org/10.1016/j.concog.2011.09.006>

Shin, Y. K., Proctor, R. W., & Capaldi, E. J. (2010). A review of contemporary ideomotor theory. *Psychological Bulletin*, 136(6), 943–974.

<https://doi.org/10.1037/a0020541>

Shor, R. E., & Orne, E. C. (1962). *Harvard group scale of hypnotic susceptibility*. Consulting Psychologists Press.

Stock, A., & Stock, C. (2004). A short history of ideo-motor action. *Psychological Research*, 68(2–3), 176–188. <https://doi.org/10.1007/s00426-003-0154-5>

Tecuapetla, F., Jin, X., Lima, S. Q., & Costa, R. M. (2016). Complementary Contributions of Striatal Projection Pathways to Action Initiation and Execution. *Cell*, 166(3), 703–715. <https://doi.org/10.1016/j.cell.2016.06.032>

Tsakiris, M., Carpenter, L., James, D., & Fotopoulou, A. (2010). Hands only illusion: Multisensory integration elicits sense of ownership for body parts but not for non-corporeal objects. *Experimental Brain Research*, 204(3), 343–352. <https://doi.org/10.1007/s00221-009-2039-3>

Tucker, M., & Ellis, R. (1998). On the relations between seen objects and components of potential actions. *Journal of Experimental Psychology: Human Perception and Performance*, 24(3), 830–846. <https://doi.org/10.1037/0096-1523.24.3.830>

Turella, L., Tucciarelli, R., Oosterhof, N. N., Weisz, N., Rumiati, R., & Lingnau, A. (2016). Beta band modulations underlie action representations for movement planning. *NeuroImage*, 136, 197–207. <https://doi.org/10.1016/j.neuroimage.2016.05.027>

Uithol, S., Burnston, D. C., & Haselager, P. (2014). Why we may not find intentions in the brain. *Neuropsychologia*, 56(1), 129–139. <https://doi.org/10.1016/j.neuropsychologia.2014.01.010>

van der Wel, R. P. R. D., Sebanz, N., & Knoblich, G. (2013). Action perception from a

- common coding perspective. In K. Johnson & M. Shiffrar (Eds.), *People Watching: Social, Perceptual, and Neurophysiological Studies of Body Perception* (pp. 101–120). New York, NY: Oxford University Press.
- van Steenbergen, H., Warren, C. M., Kühn, S., de Wit, S., Wiers, R. W., & Hommel, B. (2017). Representational precision in visual cortex reveals outcome encoding and reward modulation during action preparation. *NeuroImage*, *157*, 415–428.
<https://doi.org/10.1016/j.neuroimage.2017.06.012>
- Vargas, C. D., Olivier, E., Craighero, L., Fadiga, L., Duhamel, J. R., & Sirigu, A. (2004). The Influence of Hand Posture on Corticospinal Excitability during Motor Imagery: A Transcranial Magnetic Stimulation Study. *Cerebral Cortex*, *14*(11), 1200–1206. <https://doi.org/10.1093/cercor/bhh080>
- Vogt, S., Di Rienzo, F., Collet, C., Collins, A., Guillot, A., Rienzo, F. Di, ... Guillot, A. (2013). Multiple roles of motor imagery during action observation. *Frontiers in Human Neuroscience*, *7*(November), 807.
<https://doi.org/10.3389/fnhum.2013.00807>
- Wagenmakers, E.-J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., ... Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin & Review*, *25*(1), 58–76.
<https://doi.org/10.3758/s13423-017-1323-7>
- Waszak, F., Cardoso-Leite, P., & Hughes, G. (2012). Action effect anticipation: Neurophysiological basis and functional consequences. *Neuroscience and Biobehavioral Reviews*, *36*(2), 943–959.
<https://doi.org/10.1016/j.neubiorev.2011.11.004>
- Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review*.

<https://doi.org/10.1037/0033-295X.101.1.34>

Wegner, D. M., Ansfield, M., & Pilloff, D. (1998). The Putt and the Pendulum: Ironic Effects of the Mental Control of Action. *Psychological Science*, 9(3), 196–199.

<https://doi.org/10.1111/1467-9280.00037>

Weitzenhoffer, A. M., & Hilgard, E. R. (1962). *Stanford hypnotic susceptibility scale, form C* (Vol. 27). Palo Alto, CA: Consulting Psychologists Press.

Weitzenhoffer, A. M., Hilgard, E. R., & Kihlstrom, J. F. (1996). Stanford hypnotic susceptibility scale, form C. *Philosophical Psychology*, 16(2), 325–340.

<https://doi.org/10.1017/CBO9781107415324.004>

Wolpert, D. M. (1997). Computational approaches to motor control. *Trends in Cognitive Science*, 1(6), 209–216. [https://doi.org/10.1016/S1364-6613\(97\)01070-X](https://doi.org/10.1016/S1364-6613(97)01070-X)

Wolpert, D. M., & Ghahramani, Z. (2000). Computational principles of movement neuroscience. *Nature Neuroscience*, 3(11s), 1212–1217.

<https://doi.org/10.1038/81497>

Woody, E. Z., Barnier, A. J., & McConkey, K. M. (2005). Multiple hypnotizabilities: differentiating the building blocks of hypnotic response. *Psychological Assessment*, 17(2), 200. <https://doi.org/10.1037/1040-3590.17.2.200>

Woody, E. Z., & Sadler, P. (2005). On the virtues of virtuosos. *Contemporary Hypnosis*, 22(1), 8–13. <https://doi.org/10.1002/ch.15>

Zimmermann, M., Mars, R. B., de Lange, F. P., Toni, I., Verhagen, L., Lange, F., ... Verhagen, L. (2017). Is the Extrastriate Body Area part of the dorsal visuomotor stream? A study of structural and functional connectivity. *Brain Structure and Function*, 223(July), 31–46. <https://doi.org/10.1007/s00429-017-1469-0>

Zimmermann, M., Verhagen, L., de Lange, F. P., & Toni, I. (2016). The extrastriate body area computes desired goal states during action planning. *ENeuro*, 3(April), 1–13. <https://doi.org/10.1523/ENEURO.0020-16.2016>