

# **The Role of Associatively-Mediated Processes in Shaping Driving Behaviour**

How Experience of Contingencies Interacts with Response  
Inhibition

Submitted by William George Nicholson to the University of Exeter as a thesis  
for the degree of *Doctor of Philosophy* in Psychology, November 2019.

This thesis is available for Library use on the understanding that it is copyright  
material and that no quotation from the thesis may be published without proper  
acknowledgement.

I certify that all material in this thesis which is not my own work has been  
identified and that any material that has previously been submitted and  
approved for the award of a degree by this or any other University has been  
acknowledged.

Signed: .....

November 2019



Driving is a necessary, but inherently risky, daily activity. One behaviour exacerbating these risks occurs when drivers illegally cross amber traffic lights, and an improved ability to inhibit this behaviour would promote safer driving. This type of inhibitory control has previously been conceptualised as being itself under conscious control, and therefore requiring deliberate thought and intention. However, driving is cognitively demanding, and this is likely to reduce the ability to maintain the intention to inhibit the amber-crossing response. Recent research has demonstrated that response inhibition can become associatively-mediated with the right type of training and is thus not exclusively reliant on control processes. This finding has led to the development of inhibition training techniques to develop associatively-mediated inhibitory responses to cues that might lead to an incorrect behaviour. However, it is unclear to what extent this work could be generalised to driving. The first question addressed in this thesis centres on what kind of behaviour at traffic lights might be primed as a result of experiencing the contingencies produced at traffic light-controlled junctions. The second focuses on how training could be developed to change the products of this learning so that it primes safer behaviours.

Chapter One introduces the theoretical background to the thesis and includes a discussion of dual-process models of associative learning and associatively-mediated inhibition. Chapters Two and Three ask what is learnt at an associative level at traffic lights. Chapter Two begins the development of a laboratory paradigm that aims to capture the contingencies linked to traffic lights, and Chapter Three continues this by introducing sequences into the paradigm. Chapter Four investigates the importance of task set for associative learning and begins the development of a training task to change the learnt associative behaviour towards amber traffic lights. This work is continued in Chapter Five where the task is taken out of a pure associative learning context and applied in a real-world intervention. Finally, Chapter Six summarises the empirical work and links it to the theories and issues introduced in Chapter One.



## ACKNOWLEDGEMENTS

---

Firstly, I would like to thank my wife, my family, my friends, and my fellow (and past) PhD students at the University of Exeter and further afield, without whom I would not have completed this thesis.

I am grateful to my supervisors, Professors Ian McLaren and Frederick Verbruggen, and Doctors Natalia Lawrence and Cris Burgess for their advice and guidance. My thanks particularly to my first supervisor Ian, who lent me his expertise in associative learning and guided me through the PhD.

I would like to thank the Economic and Social Research Council and the South West Doctoral Training Centre for funding this thesis. I am particularly appreciative of the Overseas Institutional Visit grant that enabled me to undertake a study visit to the University of Sydney. Thank you to Doctors Evan Livesey and Dominic Tran for hosting me.

Finally, I would like to thank my undergraduate interns who assisted me in data collection for my thesis, and all the students and staff who participated in my experiments.



# CONTENTS

	<b>Page</b>
<b>List of Tables</b>	<b>11</b>
<b>List of Figures</b>	<b>13</b>
<b>Author's Declaration</b>	<b>15</b>
<b>1 Introduction</b>	<b>17</b>
1.1 Executive control.....	18
1.1.1 A brief history of inhibition .....	18
1.1.2 Fractioning the executive.....	19
1.2 Response inhibition.....	20
1.2.2 Models of inhibition.....	27
1.2.3 Proactive or reactive inhibition.....	28
1.3 Associatively-mediated response inhibition.....	30
1.3.1 Primed inhibition .....	30
1.3.2 Learnt inhibition: bottom-up inhibition .....	32
1.3.3 Applied response inhibition.....	34
1.4 Dual process theories of human behaviour.....	37
1.4.1 Propositional learning in humans.....	38
1.4.2 Associative learning in humans .....	41
1.4.3 Association and Cognition .....	44
1.5 Behaviour change in driving.....	46
1.5.1 Associatively-mediated behaviour in driving.....	48
1.5.2 Controlling associatively-mediated behaviours.....	50
1.6 Conclusions and next steps .....	53
<b>2 Associative Processes I: Basic Contingencies Experienced At Traffic Lights</b>	<b>55</b>
2.1 Contingencies at traffic lights .....	55
2.2 Learning about an outcome .....	56
2.3 Present experiments .....	57
2.4 Experiment 1.....	57
2.4.1 Method .....	59

2.4.2	Analysis .....	63
2.4.3	Summary .....	69
2.5	Experiment 2 .....	70
2.5.1	Method .....	73
2.5.2	Analysis .....	74
2.5.3	Summary .....	79
2.6	Experiment 3 .....	81
2.6.1	Method .....	81
2.6.2	Analysis .....	82
2.6.3	Summary .....	85
2.7	Joint analysis of Experiment 2 and Experiment 3 .....	86
2.7.1	Joint feature-positive analysis .....	87
2.7.2	Joint traffic light cues analysis .....	87
2.7.3	Summary .....	90
2.8	General discussion.....	91
2.8.1	Implications for the associatively-mediated stopping hypothesis....	92
2.8.2	The Feature-positive effect.....	94
2.8.3	The effect of task set .....	95
2.8.4	Modelling the contingencies .....	95
2.8.5	Conclusion.....	96
<b>3</b>	<b>Associative Processes II: The Effect Of Sequences upon Contingency Learning</b>	<b>99</b>
3.1	The importance of sequences.....	99
3.2	Associative learning of sequences.....	100
3.3	Present experiments .....	103
3.4	Experiment 4 .....	103
3.4.1	Method .....	106
3.4.2	Analysis and results.....	109
3.4.3	Summary .....	124
3.5	Experiment 5.....	126
3.5.1	Method .....	126
3.5.2	Analysis and results.....	128



3.5.3	Summary .....	138
3.6	Joint analysis of Experiment 4 and Experiment 5 .....	139
3.6.1	Joint feature-positive analysis .....	139
3.7	General discussion.....	140
3.7.1	Sequence learning.....	141
3.7.2	Feature-positive effect .....	141
3.7.3	Learning and awareness .....	142
3.7.4	Conclusion.....	143
<b>4</b>	<b>Associative Processes III: The Role Of Task Set In Contingency Learning</b>	<b>145</b>
4.1	Task set .....	145
4.2	The task-cueing paradigm.....	146
4.3	Present experiments .....	147
4.4	Experiment 6.....	147
4.4.1	Method .....	151
4.4.2	Analysis and results.....	154
4.4.3	Summary .....	167
4.5	Experiment 7 .....	169
4.5.1	Method .....	172
4.5.2	Analysis and results.....	174
4.5.3	Summary .....	185
4.6	General discussion.....	187
4.6.1	Learning of traffic light cues.....	187
4.6.2	Inhibition training .....	188
4.6.3	Conclusion.....	188
<b>5</b>	<b>Inhibition Training I: The STOP-CHANGE Paradigm</b>	<b>191</b>
5.1	Cognitive control in driving .....	191
5.2	Training control .....	192
5.3	Learnt control .....	192
5.4	Present experiments .....	193
5.5	Experiment 8.....	194

5.5.1	Method .....	195
5.5.2	Analysis and results.....	200
5.5.3	Summary .....	206
5.6	Experiment 9.....	208
5.6.1	Method .....	208
5.6.2	Analysis and results.....	213
5.6.3	Summary .....	221
5.7	Experiment 10.....	223
5.7.1	Method .....	223
5.7.2	Analysis and results.....	225
5.7.3	Summary .....	230
5.7.4	Bayesian meta-analysis.....	231
5.8	General discussion.....	232
5.8.1	Inhibition training in driving .....	232
5.8.2	Inhibition training across domains .....	234
5.8.3	Future directions.....	234
5.8.4	Conclusions.....	236
<b>6</b>	<b>Conclusions</b>	<b>237</b>
6.1	The Experimental Work.....	237
6.1.1	Chapter 2: Associative learning at traffic lights .....	237
6.1.2	Chapter 3: The role of sequences .....	238
6.1.3	Chapter 4: The effect of task set.....	238
6.1.4	Chapter 5: Effecting behaviour change .....	239
6.1.5	Summary .....	239
6.2	Associatively-Mediated Learning .....	240
6.2.1	Dual process theories of human behaviour .....	240
6.2.2	Applying associatively-mediated learning paradigms .....	241
6.3	Implications For Road Safety .....	242
6.4	Inhibition Training.....	243
6.5	Conclusions .....	245
<b>7</b>	<b>Appendices</b>	<b>247</b>
<b>8</b>	<b>References</b>	<b>263</b>

## LIST OF TABLES

<b>Table</b>	<b>Page</b>
Table 2.1. Summary of Experiment 1 design.....	60
Table 2.2. Descriptive statistics for Experiment 1.....	67
Table 2.3. AICc scores for models run for Experiment 1 .....	69
Table 2.4. Summary of Experiment 2 design.....	74
Table 2.5. Descriptive statistics for Experiment 2.....	76
Table 2.6. AICc scores for models run for Experiment 2 .....	78
Table 2.7. Summary of Experiment 3 design.....	82
Table 2.8. Descriptive statistics for Experiment 3.....	83
Table 2.9. AICc scores for models run for Experiment 3 .....	85
Table 2.10. AICc scores for models run for joint analysis.....	89
Table 2.11. Descriptive statistics for the joint analysis. ....	90
Table 3.1. Sequence runs for Experiment 4 .....	104
Table 3.2. Filler sequences for Experiment 4. ....	105
Table 3.3. Summary of Experiment 4 design.....	108
Table 3.4. Training descriptive statistics for Experiment 4. ....	112
Table 3.5. AICc scores for models for Experiment 4 .....	114
Table 3.6. Test descriptive statistics for Experiment 4. ....	118
Table 3.7. Descriptive statistics for expectancy ratings for Experiment 4.....	120
Table 3.8. Sequence runs for Experiment 5 .....	127
Table 3.9. Filler sequences for Experiment 5. ....	127
Table 3.10. Training descriptive statistics for Experiment 5. ....	130
Table 3.11. AICc scores for models run for Experiment 4 .....	131
Table 3.12. Test descriptive statistics for Experiment 5. ....	134
Table 3.13. Descriptive statistics for expectancy ratings for Experiment 5.....	136
Table 4.1. Sequence runs for Experiment 6 .....	149
Table 4.2. Filler sequences runs for Experiment 6. ....	151
Table 4.3. Summary of Experiment 6 design.....	152
Table 4.4. Training descriptive statistics for Experiment 6. ....	155

Table 4.5. AICc scores for models for Experiment 6 .....	158
Table 4.6. Test descriptive statistics for Experiment 6. ....	161
Table 4.7. Descriptive statistics for expectancy ratings for Experiment 6.....	163
Table 4.8. Summary of Experiment 7 design.....	174
Table 4.9. Descriptive statistics for the training phase of Experiment 7. ....	176
Table 4.10. AICc scores for training models for Experiment 7.....	178
Table 4.11. Descriptive statistics for the inhibition phase of Experiment 7. ....	179
Table 4.12. AICc scores for inhibition models for Experiment 7 .....	180
Table 4.13. Descriptive statistics for test by condition for Experiment 7.....	183
Table 4.14. AICc scores for test models for Experiment 7 .....	185
Table 5.1. Descriptive statistics for GO trials for Experiment 8.....	203
Table 5.2. Descriptive statistics for CHANGE trials for Experiment 8.....	205
Table 5.3. Descriptive statistics for stops at amber lights for Experiment 8....	206
Table 5.4. Coding of crosses at traffic lights experienced at amber trials for Experiment 9. ....	215
Table 5.5. Descriptive statistics for GO trials for Experiment 9.....	217
Table 5.6. Descriptive statistics for CHANGE trials Experiment 9. ....	218
Table 5.7. Descriptive statistics for stops at amber lights for Experiment 9....	221
Table 5.8. Descriptive statistics for GO trials for Experiment 10.....	227
Table 5.9. Descriptive statistics for CHANGE trials for Experiment 10.....	228
Table 5.10. Descriptive statistics for stops at amber lights for Experiment 9..	230

## LIST OF FIGURES

Figure	Page
Figure 1.1. Schematic of a typical go/no-go task.....	23
Figure 1.2. Schematic of a typical stop signal task.....	24
Figure 1.3. Schematic of a typical stop change task. ....	27
Figure 1.4. Schematic of the Theory of Planned Behaviour. ....	47
Figure 2.1. Traffic lights to cues in Experiment 1.....	61
Figure 2.2. Schematic of a single cue stop trial for Experiment 1.....	62
Figure 3.1. Experiment 4 training phase correlations .....	122
Figure 3.2. Experiment 4 test phase correlations .....	123
Figure 3.3. Experiment 5 correlations.....	137
Figure 4.1. Conditional cues used in Experiment 6. ....	148
Figure 4.2. Schematic of a trial for Experiment 6.....	153
Figure 4.3. Experiment 6 training phase correlations .....	165
Figure 4.4. Experiment 6 test phase correlations .....	166
Figure 4.5. Schematic of a go trial for the inhibition phase of Experiment 7... ..	171
Figure 4.6. Schematic of a stop trial for the inhibition phase of Experiment 7. ....	172
Figure 5.1. Schematic of the inhibition training task for Experiment 8. ....	197
Figure 5.2. Schematic of the inhibition training task for Experiments 9. ....	210
Figure 5.3. Photo of the driving simulator used for Experiment 9. ....	212
Figure 5.4. Photograph of the driving simulator used in Experiment 10. ....	224



## AUTHOR'S DECLARATION

---

Some of the work contained in this thesis has been prepared, submitted, or published in journals. As such, the chapters contain sections of manuscripts that were written by me but edited by fellow authors. Specifically, Chapter 2 was submitted to JEP:AL&C, and Experiment 1,2 and 4 were submitted to Cognitive Science, with the former two being published in the Proceedings of the 40th Annual Conference of the Cognitive Science Society.

In undertaking the research contained in this thesis I supervised several student interns. These students, Alice Rostant, Vanessa Evans, Paula Kastrati, Chloe Doyle, Emma Harman, Nicole Russell Pascual, Anna McCutcheon, Emma Parker, and Daisy Harris undertook recruitment for Experiments 7-10. Experiment 6 was conducted as a third-year undergraduate project, co-supervised by Ian McLaren and myself, with Eda Bulca undertaking recruitment. In all cases, I was solely responsible for the design, analysis, and interpretation of the data.

Unless stated in text, all Figures and Tables are authors own.

Finally, for reasons of power, Experiment 8 contains data from my Masters project.

William G Nicholson

November 2019





# CHAPTER 1

---

## Introduction

In 2017 there were 1,793 deaths on roads in the United Kingdom (UK) despite years of highway safety improvements and campaigns (Department of Transport, 2018). Road traffic collisions are estimated to cost the UK economy £35 billion a year (Department of Transport, 2018). Given that some estimates put driver error as a critical factor in 94% of road accidents (U.S. Department of Transportation, 2015), interventions to address the cognitive characteristics of drivers would seemingly be of considerable utility (Cheng, Ng, & Lee, 2012).

One cause of road accidents is people contravening traffic light signals. Specifically, with 22% of urban road incidents caused by drivers ignoring a stop signal at traffic lights (Retting, Williams, Preusser, & Weinstein, 1995), and 38% of drivers rarely stopping at amber traffic lights (Thrifty, 2011), there is a need to reduce the crossing of amber and red traffic lights on the part of drivers (Polders et al., 2015). A potential solution that has been adopted in some cases is the use of cameras to enforce red traffic lights. While these cameras can lead to safer driving and increased compliance (Baratian-Ghorghi, Zhou, & Franco-Watkins, 2017) such a reactive approach does not address the root causes of the behaviour.

So, what causes this behaviour? How does one decide to stop (or not) at traffic lights? Such a process relies on the human ability to adjust behaviour in response to environmental cues, and an important component of this ability is an activity termed response inhibition. While this process has been studied from many different perspectives for a long time (Verbruggen, McLaren, & Chambers, 2014) this thesis explores the control of human behaviour by appealing to current cognitive psychology theories, that is through the use of executive control processes. Crucially, while these processes have historically been ascribed to top-down conscious thought, recent research has highlighted how this is not necessarily the case, demonstrating that control can become mediated through bottom-up associatively-mediated processes. The thesis goes on to explore how techniques to change cognitive behaviour might help ameliorate dangerous driving at traffic lights. Specifically, the work explores how behaviour at traffic lights might be, in part, associatively-mediated, and

investigates the implications of this possibility for the development of behaviour change techniques.

To begin, the role of executive control in human behaviour (with a focus on motor control) is outlined, with an emphasis on how this can be voluntary or automatic. I will then highlight the current role of behaviour change techniques in the domain of driving, addressing the shortcomings of these approaches, and how more appropriate techniques could be developed by using frameworks already established in cognitive psychology.

## **1.1 EXECUTIVE CONTROL**

The term executive control is an umbrella term covering a variety of systems that allow people to modify their behaviour in response to environmental changes. It includes planning and monitoring as well as response inhibition (Baumeister & Heatherton, 1996; Miyake et al., 2000; Hofmann, Schmeichel, & Baddeley, 2012; Diamond, 2013) and is important to many aspects of human life, such as school success (Blair & Razza, 2007). A review of the whole executive control literature is beyond the scope of this thesis (for a review see Monsell and Driver (2000)); the focus here is on response inhibition and how this is important in motor control.

### **1.1.1 A brief history of inhibition**

Scientists and philosophers have long studied the nature of human control. As Bari and Robbins (2013) note, Plato's allegory of the chariot, where the human soul is a charioteer being driven by two horses having opposite characters, is symbolic of the operation of inhibition, whereby to reach the intended destination (in this case heaven) the horses have to be controlled with the opposing forces successfully balanced. Similar ideas come from Descartes (trans. 1989) and Diamond, Balvin, and Diamond (1963).

In the 19<sup>th</sup> century, the study of inhibition shifted from a philosophical to a scientific perspective (Smith, 1992). These approaches were rooted in the field of psychiatry. As early as 1843 the German psychiatrist Wilhelm Griesinger was arguing that, in the terminology of the day, insanity was due to impaired inhibition (as reported by Macmillan, 1996). Sechenov (1863) and Ferrier (1886) began the search for a neural basis of inhibition. Such work often involved

patients with lesions to the frontal cortex. These patients tended to be unimpaired on most mental functions (e.g., vision and hearing), but had deficits in goal-directed behaviour and were distracted by salient but irrelevant stimuli (Milner, 1963; Perret, 1974; Milner & Petrides, 1984).

A famous demonstration of the effects of lesions to the frontal cortex comes from Shallice and Burgess (1991). In this study three individuals with prefrontal brain damage were set a number of simple but open-ended tasks to complete in a shopping centre whilst also obeying a number of rules. For example, one of the tasks was to “buy a packet of throat pastilles” (p. 733), and a rule was “no shop can be entered other than to buy something” (p. 734). The patients were told to complete the task quickly. All three patients made more errors in the task than controls matched for age and IQ. The types of mistakes demonstrate the deficits in goal-direction. For example, one participant made an error when she entered a chemist shop to buy soap (one of the task items) but did not buy anything (and thus broke a rule) as she did not like the soap on sale, even though personal preference was irrelevant to the task at hand. In fact, all three patients displayed inefficiencies in the task and broke at least five rules each. This study demonstrates that patients with prefrontal lobe damage can complete basic tasks (such as remembering instructions), but that the efficient coordination of behaviour to do so proves difficult. These findings led to the development of the idea that the prefrontal cortex is not involved in a particular faculty (such as language) but rather is central to control and goal-directed behaviour (Miller & Cohen, 2001). Such work was key in the development of influential theories such as that of the supervisory attentional system by Norman and Shallice (1986) and the concept of the ‘central executive’ in Baddeley and Hitch's (1996, 2003) model of working memory.

### **1.1.2 Fractioning the executive**

With the development of the concept of an ‘executive controller’, control was conceived as being located in a unitary homunculus that was responsible for pulling the levers to regulate low-level systems (Baddeley, 1996). However, it has become clear that this concept is untenable as it relies on circular reasoning, and begs the questions *Quis custodiet ipsos custodes?* (Juvenal, trans. 2014 from Watson & Watson, 2014). If the homunculus is in control, what

(or 'who', see the doctrine of divine universal causality, Grant, 2010) is controlling the homunculus?

In an attempt to move beyond explaining what is controlled to how control is exercised, Monsell and Driver (2000) proposed the slogan "Dissolve, deconstruct, or fractionate, the executive! Let a hundred idiots flourish!" (p. 7). They argued that to fully explain executive control, the processes ("idiots") that underlie the homunculus need to be defined. There is evidence against the concept of a unitary homunculus. While Shallice and Burgess' (1991) patients all showed deficits in the task, each patient had differing impairments. For example, Patient 1 showed no task failures, while Patient 2 did. This suggests that different parts of executive control are impaired following damage to specific parts of the prefrontal cortex. Further examples of the existence of multiple control processes comes from work by Godefroy, Cabaret, Petit-Chenal, Pruvo, and Rousseaux (1999) who showed that some individuals with frontal brain damage were impaired on one cognitive task but not another; while the opposing pattern of results were found in other patients with posterior brain damage, a double-dissociation implying separable processes.

These results led to the fractioning of the executive controller into subcomponents. Focusing on individual differences, Miyake et al. (2000) developed three distinct functions of executive control: updating and monitoring information, switching between responses and task sets, and inhibiting irrelevant actions. Thus, the current prevailing view is that 'executive control' contains separate specialised components, though see Verbruggen, McLaren, et al. (2014) for a commentary on how the homunculus has not yet been fully banished. The focus of this thesis now turns to the third component of executive control from the Miyake et al. (2000) taxonomy: inhibition.

## **1.2 RESPONSE INHIBITION**

Inhibition is often considered to be a key facet of executive control (Baddeley, 1996; Aron, 2007; Aron, Robbins, & Poldrack, 2014; Nigg, 2017). While its role in some processes such as memory and attention is debated (MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003; Aron, 2007), there is a large body of evidence suggesting that inhibition is used in motor control, specifically the ability to cancel an already actioned motor response (Nigg, 2000; Coxon, Stinear, &

Byblow, 2006; Verbruggen & Logan, 2008a; Majid, Cai, George, Verbruggen, & Aron, 2011; Bari & Robbins, 2013).

One issue in assessing response inhibition is task purity. By definition, response inhibition is a component of a wider network of executive control, and as such these other processes (e.g., attention) are likely to impact task performance. A solution around this is to focus on the inhibition of observable motor behaviours. Such an approach enables an objective measure of inhibition, either via a button press (Verbruggen & Logan, 2009a) or through the direct investigation of motor cortex via Transcranial Magnetic Stimulation (Pascual-Leone, Valls-Solé, Wassermann, & Hallett, 1994; Cirillo, Cowie, MacDonald, & Byblow, 2017). Thus, for present purposes, response inhibition, or just inhibition, will be defined as the “ability to suppress a motor response that is no longer appropriate or required” (Bowditch, 2016, p. 5; adapted from Chambers, Garavan, & Bellgrove, 2009). Accordingly, response inhibition can be studied in the laboratory by encouraging the development of prepotent, dominant, responses to a cue, and then introducing the need to suddenly inhibit this set response in reaction to a rarely seen stop signal (Logan, 1994).

It is worth noting that motor response inhibition (the stopping of an action, e.g., not pressing a key in a laboratory experiment) and cognitive inhibition, which can be seen as the stopping of a mental process (MacLeod, 2007), e.g., the suppression of a task irrelevant memory (Anderson & Green, 2001), might not be distinct processes but rather represent similar constructs. For example, brain imaging studies have found overlapping areas of neural activity in both cognitive and response inhibition (Cohen & Lieberman, 2010). Further support comes from studies which show that deficits in response inhibition and cognition inhibition can be comorbid, e.g., in those suffering from Obsessive Compulsive Disorder (Chamberlain, Fineberg, Blackwell, Robbins, & Sahakian, 2006; for review see Bari & Robbins, 2013). Indeed, response inhibition has been found to be important for normal and healthy functioning, which suggests a wider, more complex network of processes is at play than just the inhibition of motor responses. For example, poor response inhibition (and poor executive control more generally) has been linked to attention deficit/hyperactivity disorder (Nigg, 2001; Berryessa, 2017), the development of anti-social and criminal behaviour (Tremblay, Pihl, Vitaro, & Dobkin, 1994; Moffitt, 2017), and to drug use (Moffitt

et al., 2011). Of interest for the present thesis, poor response inhibition has also been linked to risky driving (Bachoo, Bhagwanjee, & Govender, 2013; O'Brien & Gormley, 2013; Bıçaksız & Özkan, 2016; Sani, Tabibi, Fadardi, & Stavrinos, 2017), and those with attention deficit/hyperactivity disorder have been shown to have an increased tendency to commit traffic violations (Groom, van Loon, Daley, Chapman, & Hollis, 2015). Furthermore, performing motor inhibition tasks has been shown to lead to a change in actual behaviour, with this type of training being used to reduce alcohol consumption (Houben, Nederkoorn, Wiers, & Jansen, 2011) and unhealthy food intake (Veling, van Koningsbruggen, Aarts, & Stroebe, 2014). These applications of inhibition are explored further in section 1.3.3. As an aside, it is worth noting that while a lack of inhibition certainly has negative consequences, it is not entirely without benefits. Dickman's (1990) construct of *functional impulsivity* implies that in some situations lack of control is beneficial. For example, Dickman and Meyer (1988) found that when under time pressure individuals with high impulsivity were more accurate in a visual-comparison task than those with low impulsivity. Overall, inhibition is key for healthy functioning; yet depending on the context low inhibition is not necessarily bad.

### **1.2.1.1 Paradigms**

As described above, motor inhibition fundamentally reflects a person's ability to stop a prepotent motor response. Attention now turns to the paradigms used to investigate inhibition of motor responses. The three main paradigms used in the literature, and that are the primary methodologies employed in subsequent empirical chapters of this thesis, are the go/no-go, stop-signal, and stop-change paradigms.

#### **1.2.1.1.1 GO/NO-GO**

In this paradigm, first developed by Donders (1868), participants are presented with stimuli and told to respond (e.g., press a key) when a go stimulus is presented, but to withhold their response when a no-go stimulus is presented. In Figure 1.1 participants would have to respond to the letter Y and withhold a response to letter B. Typically, trials are presented rapidly and there is a low probability of no-go. This design leads to a go response becoming the default, 'prepotent', response and ensures that participants are withholding a response rather than merely deciding not to make a response (Aron, 2011). This issue is

discussed further in section 1.2.3. These paradigms often use mapping to visual cues, such as responding to different coloured circles (Nicholson, Verbruggen, & McLaren, 2018), or to categories of stimuli, such as responding to words describing living rather than non-living objects (Verbruggen & Logan, 2008b). Experiments have also used two different auditory tones as go and no-go cues (Steinmann et al., 2011).

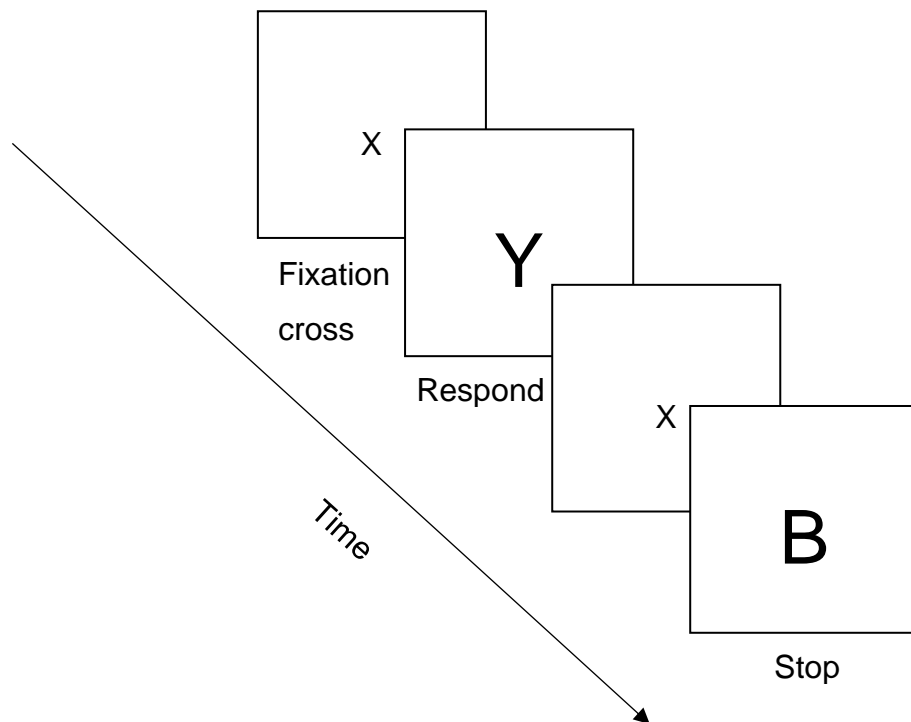


Figure 1.1. Schematic of a typical go/no-go task. In this task participants have to respond to the letter 'Y' and withhold a response to the letter 'B'.

#### 1.2.1.1.2 Stop-signal task

The stop-signal task was developed by Lappin and Eriksen (1966) and Logan and Cowan (1984). In the task, participants have to make a choice discrimination (e.g., in Figure 1.2, left key response for letter A and right key response for letter W). If a stop signal is presented (which can be another visual stimulus or auditory cue) then participants must withhold their response.

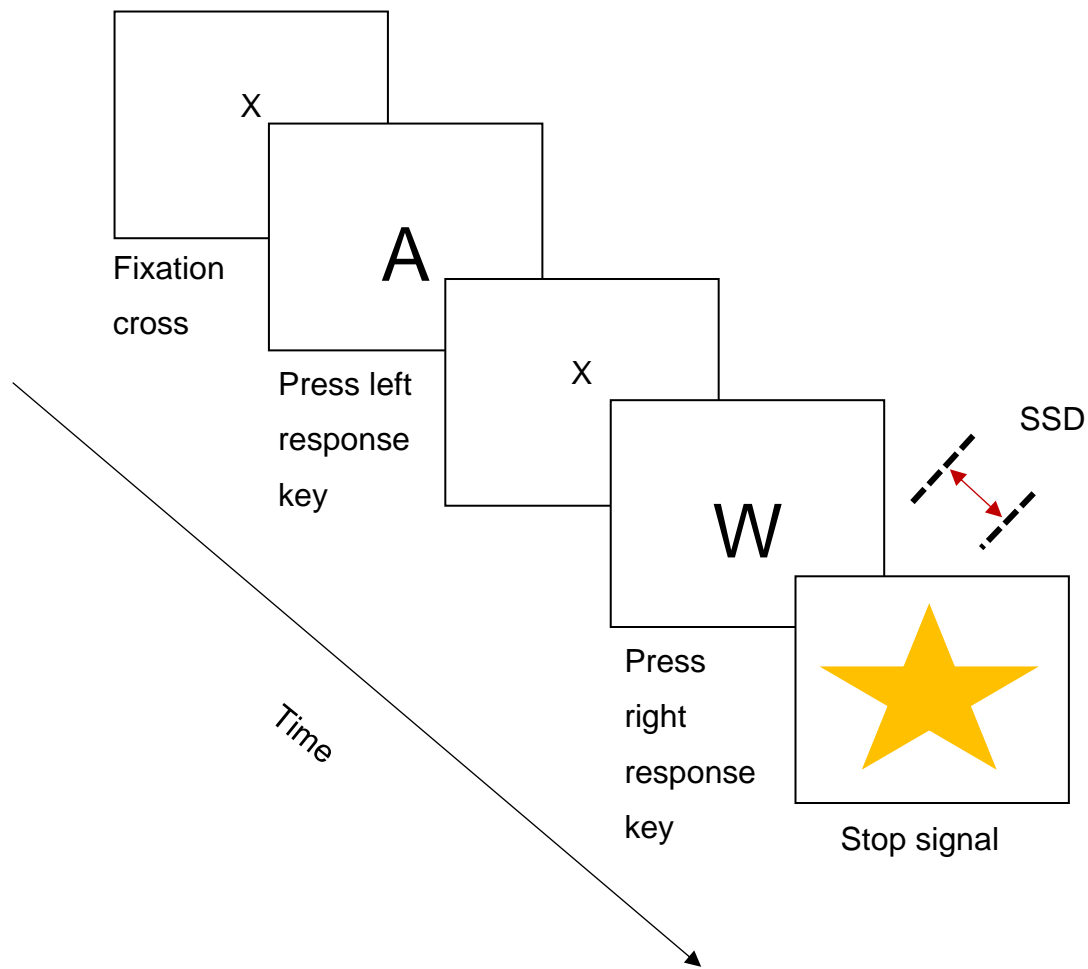


Figure 1.2. Schematic of a typical stop signal task. In this task participants have to respond to letter ‘A’ with a left key press and respond to letter ‘W’ with a right key press. However, when the stop signal is presented (the star) participants must withhold their response. The SSD is the Stop Signal Delay, the time between stimulus presentation and the stop signal occurring.

The delay between the presentation of a stimulus and the stop signal is varied and called the stop signal delay. The duration of this delay directly affects the probability of correct inhibition of response. If the delay is short it is likely that participants will be able to withhold their response, but if the delay is long then participants tend to be unable to cancel their response and make a commission error (Matzke, Verbruggen, & Logan, 2018). By linking the length of the delay to participants’ performance (called a staircase design) the differences between participants and within-participants across the experiment can be controlled for (Verbruggen & Logan, 2009a; Coulacoglou & Saklofske, 2017). For example, the delay could increase by 50ms if correct inhibition occurred and decrease by 50ms if a commission error occurred. Such a one-up/one-down tracking



procedure would result in a 50% chance of inhibiting a response. An additional benefit of varying the delay in this way is that it becomes hard for participants to develop a strategy to increase the chance of correctly stopping (Logan, 1994).

Regarding the required response in the task, while key presses are the most frequently used responses that participants are required to inhibit, researchers have also explored other modalities. These include requiring participants to interrupt their own speech (Ladefoged, Silverstein, & Papçun, 1973) or eye movements (Logan & Irwin, 2000). As Bari and Robbins (2013) note, each modularity has its advantages and disadvantages, with a key feature being experience in the primary task, for example an actor might have more difficulty interrupting their own speech than a normal person due to the automatic nature of the task and practiced behavioural sequences. Similarly to the go/no-go task, discrimination can be between a variety of visual (Senderecka, Grabowska, Szewczyk, Gerc, & Chmylak, 2012) and auditory stimuli (Manuel, Bernasconi, & Spierer, 2013). Naturally, stimuli that are easier to detect will result in better performance than stimuli that are harder to detect (Palmer, Huk, & Shadlen, 2005), and as noted above there are likely to be individual differences in task difficulty. Interference from the stimulus can also affect performance, with emotional stimuli leading to worse performance compared to neutral stimuli (Verbruggen & De Houwer, 2007). Both visual (Verbruggen & Logan, 2008a) and auditory stop signals (Van Der Schoot, Licht, Horsley, & Sergeant, 2005) are commonly used. The ease with which the stop signal is detected has a direct effect on performance in the task. For example, louder stop signal tones produce better inhibition than quieter tones (Van Der Schoot et al., 2005). To help minimise such issues, the experiments in this thesis will build upon the work of Bowditch, Verbruggen, and McLaren (2016) and Verbruggen and Logan (2008a) to use simple visual discriminations involving novel stimuli which require key presses as a response, a task that it is reasonable to assume that all participants will be equally (un)familiar with.

#### **1.2.1.1.3 Stop-change paradigm**

The final paradigm used in this thesis is the stop-change paradigm. Here, participants are told to stop their response to a go cue and make a 'change' response when an appropriate cue appears (Verbruggen, Schneider, & Logan, 2008; Verbruggen & Logan, 2009a). In Figure 1.3, participants must respond to

letter A with a left-hand response and to the letter W with a right-hand response, but if this changes to a red W then it requires a left-hand response. The delay between the letter W appearing and changing red is called the stop-change signal delay and is equivalent to the stop signal delay in stop signal paradigms. The most typical change response is to press a different key from that of the first response (Logan & Burkell, 1986) or to make an opposite response, such as in Nachev, Wydell, O'Neill, Husain, and Kennard (2007) where participants pressed the key opposite to that primed by the first cue. However, studies have also required participants to respond depending on the identity of the stop-change signal, e.g., word discrimination based on whether the signal was a high or low tone (Verbruggen & Logan, 2008c). The paradigm has also been used in animal experiments, including with rats (Beuk, Beninger, & Paré, 2014) and pheasants (Meier et al., 2017).

One issue regarding the stop-change paradigm is whether participants inhibit their first response before changing to the second. In other words, is inhibition required to complete the task? To explore this question, Verbruggen, Schneider, and Logan (2008) introduced a variable delay between the stop signal and the second go response to distinguish between two models of performance. The first, the GO1-GO2 model, assumes that inhibition is not required for successful completion of stop-change tasks, with participants replacing their goal to respond to GO1 with the goal to respond to GO2 when instructed. In a sense STOP is achieved by the replacement of GO1 with the GO2 goal. The second, the GO1-STOP-GO2 model, assumes that when instructed to change responses, participants first stop GO1 responses and then make GO2 responses, with a STOP goal being required to inhibit GO1 responses before GO2 responses can be made. Interestingly, the two models would predict different effects of increasing the delay between the stop signal and the second go response. Model 1 would predict that increasing delays would not impact upon performance, with reaction times to the second GO cue not varying with delay. This is because the GO2 processing happens immediately when the change signal is presented. However, model 2 would predict that reaction times to the GO2 cue are linked to the delay, with a longer delay period leading to faster GO2 reaction times. This is because the STOP process must be finished before a GO2 response can be made, and so a longer delay will increase the chance that this STOP process is completed. In their

experiments, Verbruggen et al. (2008) found evidence to support the model 2 account of stop-change performance, with reaction times to the GO2 cue decreasing when the delay between the change signal and GO2 increased. This demonstrates that successful performance within stop-change tasks involves the use of inhibition. Such conclusions are supported by computational (Camalier et al., 2007) and neuroimaging (Boecker, Gauggel, & Druke, 2013; Jha et al., 2015) experiments.

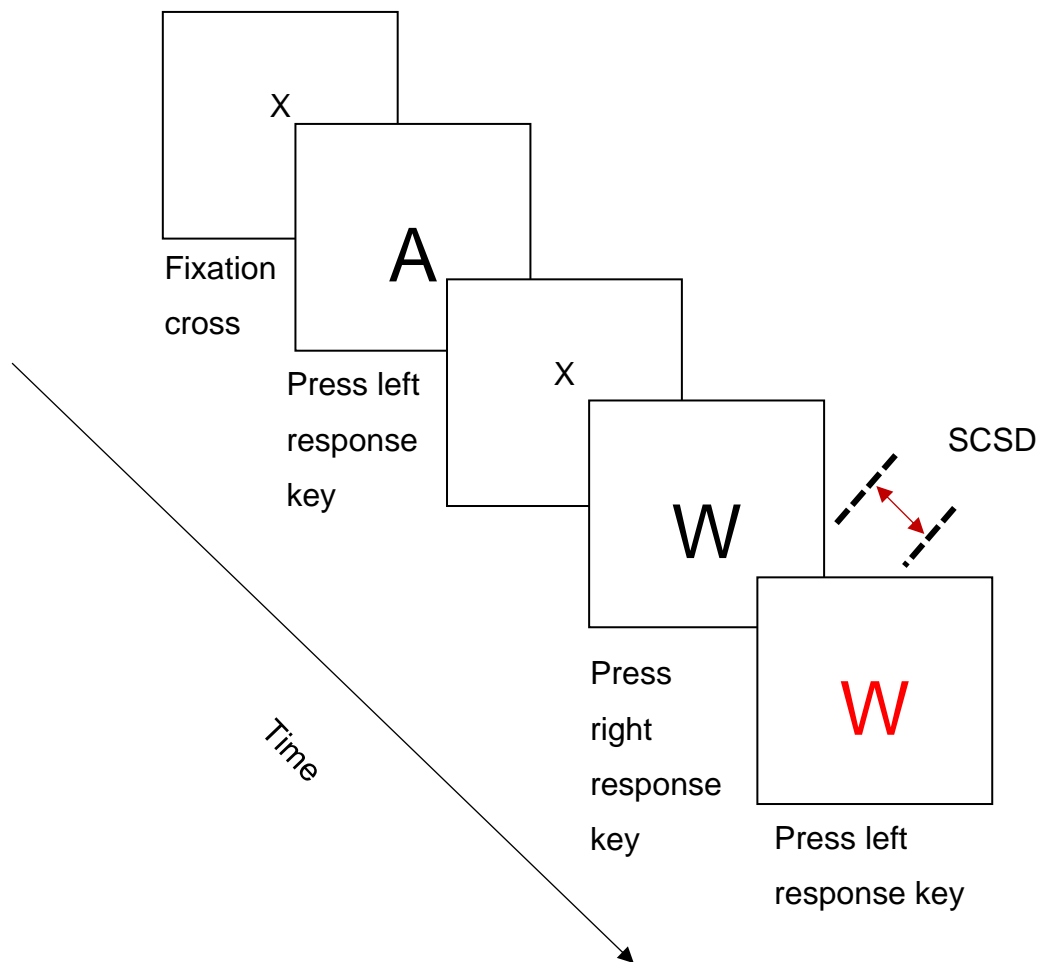


Figure 1.3. Schematic of a typical stop change task. In this task participants must respond to letter 'A' with a left key press and respond to letter 'W' with a right key press. However, when the change signal is presented (W turns red), participants must change their response to a left key press. The SCSD refers to the Stop Change Signal Delay, the time between the stimulus presentation and the stop change signal occurring.

### 1.2.2 Models of inhibition

Performance on the above response inhibition tasks is often modelled as a 'horse race' between two independent processes; a go process triggered by the

presentation of a go stimulus, and a stop process triggered by the presentation of a stop stimulus (Logan & Cowan, 1984; Logan, Van Zandt, Verbruggen, & Wagenmakers, 2014). On any trial, performance (that is, responding or not responding) depends on the relative finishing times of the two processes. If the stop process finishes before the go process then participants will withhold a response, but if the go process finishes first then participants will respond. By definition, the independent race model assumes that the two processes operate completely separately, with behavioural (Lea, Chow, Meier, McLaren, & Verbruggen, 2019) and neurological evidence (Schmidt, Leventhal, Mallet, Chen, & Berke, 2013) supporting this view. However this assumption has not been accepted by all, with some studies suggesting that the two processes are more interrelated (Özyurt, Colonius, & Arndt, 2003; Gulberti, Arndt, & Colonius, 2014).

### **1.2.3 Proactive or reactive inhibition**

A key distinction in response inhibition research is between proactive inhibition (also referred to as action restraint) and reactive inhibition (also termed action cancellation). Proactive inhibition (Whitely & Blankfort, 1933) refers to the inhibition of motor control prior to a response being made, while reactive inhibition describes the inhibition of a motor response during its execution (Hull, 1943). Braver (2012) characterised reactive inhibition as being triggered by contextual cues in the environment, with proactive inhibition entailing processing of goal-relevant information to bias behavioural responses.

Recently this distinction has been applied to inhibition training paradigms (Aron, 2011). It has been suggested that the go/no-go paradigm measures proactive inhibition, whereas the stop signal task measures the ability to cancel an ongoing motor response, that is reactive inhibition (Schachar et al., 2007; Littman & Takács, 2017). The argument is that for go/no-tasks participants must make response decisions based on the trial: that is, they must first select a response before initiating it. In comparison, in stop-signal tasks the go signal is presented first, and as such activates the go process, with any subsequent stop signal requiring the inhibition of the already active go pathway. Such a characterisation has found support in partially dissociable neural networks (Eagle, Bari, & Robbins, 2008; Swick, Ashley, & Turken, 2011).

However, as discussed earlier, go/no-go experiments are often designed so that go is the prepotent response, either by having more go than stop trials, for example in Adams, Lawrence, Verbruggen, and Chambers (2017) 75% of trials were go, or by having short response windows, as seen in Leiva, Parmentier, Elchlepp, and Verbruggen (2015). Given this, the argument is that as go is the default response, the go process is activated immediately on the presentation of a stimulus. Evidence from Leocani, Cohen, Wassermann, Ikoma, and Hallett (2000) supports this case. The study used a go/no-go task to investigate corticospinal excitability (by using the proxy of motor-evoked potentials induced by Transcranial Magnetic Stimulation). Results showed that there was inhibition of excitability around 200-300ms following presentation of no-go trials. This suggests that the prepotent response was go, and that participants inhibited a response rather than merely not making a response. Crucially, this effect was found even when there was an even ratio of stop to go trials (50:50) and the response window was long for such a task (6-8 seconds). Participants were instructed to respond “as quickly as possible” (p. 1163) to stimuli, indicating that this simple instruction alone can be enough to enable the formation of prepotent responses in go/no-go paradigms.

While proactive control is an important facet of inhibition, it does seem clear that reactive inhibition is often involved in not responding during these tasks, and that certain design features can lead to a greater role for reactive inhibition. The use of instructions emphasising going, high go ratios, and quick response windows are likely to push participants to use reactive inhibition in both stop-signal and go/no-go tasks. Saying this, successful performance on inhibition tasks requires a balance between responding to go trials but also withholding responses on stop trials (Verbruggen & Logan, 2009b). It is likely that both processes are running concurrently (Duque, Lew, Mazzocchio, Olivier, & Ivry, 2010). For example, reaction times are typically longer in blocks where there are stop trials, compared to blocks where there are no stop trials (Ramautar, Kok, & Ridderinkhof, 2004; Verbruggen, Best, Bowditch, Stevens, & McLaren, 2014). Moreover, this effect can happen on a trial-by-trial basis when participants are informed about the likelihood of a stop signal occurring (Chikazoe et al., 2009; Jahfari, Stinear, Claffey, Verbruggen, & Aron, 2010; Jahfari et al., 2012; Zandbelt, Bloemendaal, Neggers, Kahn, & Vink, 2013).

To summarise, when participants receive instructions promoting going, and are responding at speed, in tasks designed with relatively high go ratios, then one can be confident that action cancellation is occurring, with the involvement of proactive inhibition being limited. With this in mind, the empirical work in this thesis will use the above strategies to minimise the role of proactive inhibition. In the next section I highlight another aspect of inhibition: how, over time, inhibition can become associatively-mediated rather than relying on top-down conscious control.

### **1.3 ASSOCIATIVELY-MEDIATED RESPONSE INHIBITION**

As discussed earlier, inhibition is often assumed to be part of a suite of executive control processes that are under conscious control, so that behaviour in stop-signal and go/no-go tasks results from goal-directed and deliberate actions (Diamond, 2013). Indeed, it has been argued that inhibition by its very nature requires conscious processing. As Parkinson and Haggard (2014) note, to curb a prepotent action one needs to consciously resist it and to know that this is the case. Furthermore, throughout history there are accounts of the intensity of feelings when trying to overcome prepotent urges (Augustine, trans. 2006; Dostoevsky, trans. 2018). This can be contrasted with the position taken by learning theorists, who argue that with sufficient practice, responses can be automated (Dickinson, 1985; McLaren, Green, & Mackintosh, 1994; McLaren et al., 2014). Indeed, there is now converging evidence that inhibition is influenced by bottom-up processing and that, with appropriate training, inhibition itself can become driven by bottom-up rather than top-down processes.

#### **1.3.1 Primed inhibition**

Priming effects have a long history in psychology and occur when exposure to one stimulus affects the response to another. For example the word 'Mug' is recognised more quickly if preceded by the word 'Hot' than by the word 'Car' (for review see Neely, 1991). Subliminal priming is when the priming stimuli are presented too quickly to be noticed so that they fall below the threshold of perception and so are not consciously processed (Elgendi et al., 2018).

Work conducted by van Gaal et al. (van Gaal, Ridderinkhof, Fahrenfort, Scholte, & Lamme, 2008; van Gaal, Ridderinkhof, van den Wildenberg, & Lamme, 2009; van Gaal, Ridderinkhof, Scholte, & Lamme, 2010; van Gaal, Lamme,

Fahrenfort, & Ridderinkhof, 2011) has been key in showing how inhibition can be triggered by subliminal primes. In a series of experiments participants were presented with subliminal primes in both go/no-go and stop signal tasks (van Gaal et al., 2008; van Gaal et al., 2011). Behavioural results showed that the presentation of masked no-go stimulus or stop-signals before a go or no-signal trial led to slower reaction times and increased the likelihood of omission errors. These findings were supported by functional magnetic resonance imaging recordings which showed that the presentation of the primes led to neural activity similar to that linked to top-down inhibition, indicating that inhibitory control functions in the brain can be triggered subconsciously (van Gaal et al., 2010). Furthermore, the magnitude of electroencephalographic components N2 and P3 (typically associated with response inhibition in standard go/no-go (Lavric, Pizzagalli, & Forstmeier, 2004) and stop signal task (Ramautar et al., 2004)) correlated with the slowing of behavioural responses to subliminal no-go and stop signals (van Gaal et al., 2008; van Gaal et al., 2011). However, though the authors (van Gaal, De Lange, & Cohen, 2012) attribute the pattern of results to subconscious activation of the inhibition-related neural networks, this remains controversial as one cannot be sure that participants were entirely unaware of the subliminal stimuli, with such priming methods being shown to underestimate conscious perception (Newell & Shanks, 2014).

Further evidence for the role of priming, and thus automation, in response inhibition comes from Verbruggen and Logan (2009c). In a series of studies, the role of irrelevant (but visible) primes upon stopping in a stop-signal task was investigated. In the experiments the words GO or STOP were superimposed over simple shapes (circles or squares) with participants being told to ignore the words and respond to the shapes unless an auditory tone was presented. Despite the words being irrelevant to the task they had a clear effect on performance, with slower responding for go trials when STOP appeared over the shapes compared to GO. The results were replicated in a go/no-go variant with the finding that the effect was dependent upon task context. That is, the STOP word only impacted performance when stopping was an outcome, having no effect when presented in a go-only condition.

Considering these experiments in combination, it seems that inhibitory control, typically considered to be exclusively delivered through top-down processes,

can be influenced by stimulus driven processes. Indeed, it seems possible that awareness is not required. However, the context effects seen in Verbruggen and Logan (2009c) demonstrate that such priming does not necessarily override top-down processes and suggests that bottom-up and top-down processes interact (an issue I will return to in section 1.4.3).

### **1.3.2 Learnt inhibition: bottom-up inhibition**

In the previous section it was demonstrated how priming effects can increase stopping effects in response inhibition tasks. Given the evidence that shows how practice at responding to a cue can lead to the learnt response becoming automatised over time, through a process of forming stimulus–response mappings (Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977; Dickinson, 1985; Logan, 1988), it is natural to ask if the effect works the other way. Does repeatedly stopping (via response inhibition) to a cue lead to an associatively-mediated stop response?

The first experiment to explore this idea was by Verbruggen and Logan (2008b). In one of their experiments, participants were presented with living or non-living words where the category determined the required response, e.g., GO for the word ‘apple’ (living) but STOP for the word ‘glass’ (non-living). Following training, these contingencies were reversed for the test phase. If pairing stimuli (in this case category of words) with specific responses leads to associatively-mediated stopping, then one would expect words that were previously paired with stopping to have slower reaction times than novel stimuli. Consistent with this hypothesis, it was found that responses were slower in the test phase to stimuli previously associated with stopping, compared to stimuli associated with going or to novel stimuli. The authors concluded that the slowing witnessed in the test phase was caused by the retrieval of stimulus–stop associations that automatically inhibited responding. Similar results were found by Noel et al. (2016) who paired words with consistently going or consistently stopping, and then at test reversed the mappings. At test, in line with Verbruggen and Logan (2008b), go performance was impaired for old stop stimuli. Further work by Verbruggen, Best, et al. (2014) confirmed the findings were not driven by sequential learning after-effects (repetition priming effects).

Further support for the notion of associatively-mediated stopping comes from Best, Lawrence, Logan, McLaren, and Verbruggen (2016). In the experiment



participants had to respond to vowels or consonants superimposed on top of images (e.g., a bucket) for some blocks and respond to numbers bigger or smaller than five superimposed on top of the same images in other blocks. The design meant that though the go/stop signals changed over the experiment, participants always responded to some images (e.g., a bucket) and always stopped to others (e.g., a hat). At test the contingences were reversed. Analysis at test showed a significant difference between responding to old-go and old-stop images such that responding was slower to old-stop compared to old-go. These results suggest that participants can acquire direct stimulus-stop associations.

Such work led to the development of the associatively-mediated inhibition hypothesis (Verbruggen & Logan, 2008b; Verbruggen, Best, et al., 2014; Verbruggen, McLaren, et al., 2014), namely that inhibitory control can be triggered automatically via the retrieval of stimulus-stop associations. This process relies on the consistent presentation of stimulus-stop trials (Logan, 1988), or more accurately with practice of these consistent mappings, with A. Jones et al. (2016) showing that it is the proportion of successful stimulus-inhibition responses that is important rather than the total number of stimulus-stop trials. In a relevant study for the current thesis, Hochman, Henik, and Kalanthroff (2018) explored the effect of images of traffic lights upon going and stopping. The experiment used a stop-signal task, in which participants had to respond to a picture of either a red or a green traffic light and withhold a response if they heard an auditory tone. Given the ubiquitous nature of red and green traffic lights to mean 'stop' and 'go' respectively (which extends outside traffic management, for example, the use of the traffic light rating system for food labels (Department of Health, 2016)), the associatively-mediated inhibition hypothesis would predict that reaction times to going at green traffic light images would be faster than going to red traffic light pictures. As predicted, reaction times in go trials were significantly faster when paired with a green traffic light than a red traffic light.

The work on priming and learnt inhibition shows how inhibition can become associatively-mediated through the pairing of a stimulus and a response. Such findings can be seen in the wider context of the associative learning literature,

which is mentioned in section 1.4, but now the focus of this introduction turns to the applications of the associatively-mediated inhibition hypothesis.

### **1.3.3 Applied response inhibition**

The evidence that associatively-mediated response inhibition can result from simple computerised tasks offers a tantalising practical application: that tasks could be developed to pair images (e.g., red traffic lights) to a certain response (e.g., stop) and so result in increased real-world braking to red traffic lights. As discussed previously there is some debate in the literature as to whether motor inhibition can bring about more general inhibition of thoughts and other mental processes (MacLeod et al., 2003; Berkman, Burklund, & Lieberman, 2009). However, the promise of the associatively-mediated inhibition hypothesis has led to a number of studies demonstrating that pairing images of alcohol (Houben et al., 2011; Houben, Havermans, Nederkoorn, & Jansen, 2012; A. Jones & Field, 2013) or food (Veling, Aarts, & Papies, 2011; Houben & Jansen, 2011, 2015; Veling, Aarts, & Stroebe, 2013a, 2013b; Veling et al., 2014; N. S. Lawrence, O'Sullivan, et al., 2015; N. S. Lawrence, Verbruggen, Morrison, Adams, & Chambers, 2015; Poppelaars et al., 2018; Camp & Lawrence, 2019; Forman et al., 2019) with stopping (through go/no-go or stop-signal tasks) can lead to a subsequent reduction in consumption. Trials have also investigated using such training to reduce cocaine use (Alcorn III, Pike, Stoops, Lile, & Rush, 2017) and to see if the training can be used in military contexts, such as using go/no-go tasks to reduce civilian casualties caused by a failure to inhibit shooter performance (Biggs, Cain, & Mitroff, 2015).

A striking example of the beneficial impact such training could have is provided by N. S. Lawrence, O'Sullivan, et al. (2015). The study examined the ability of online inhibition training, in this case a go/no-go task, to lead to a reduction in food intake. In the study, adult participants were randomly assigned to either a control or active condition. The active condition had to inhibit motor responses to images of high-calorie food, while the control condition had to withhold responses to non-food items (such as socks). After four 10-minute sessions over a one-week period, those in the active group showed significant weight loss and reduced calorific intake, compared to the control group at a two-week follow-up. Furthermore, at the six months follow-up the self-reported weight loss was maintained for the active group. This study demonstrates the substantial

positive impact such training could have upon lifestyles. Yet, Carbine and Larson (2019) performed a *p*-curve analysis (where the distributions of significant *p*-values are plotted against different expected distributions, see Simonsohn, Nelson, and Simmons (2014)) to assess the overall evidential value of research using inhibition training to reduce food consumption. The authors performed four curve analyses and consistently found a 'U'-shape distribution, arguing that this was evidence for a true underlying effect of the training but also evidence of selective reporting in the literature. Analysis also showed that the effect of training was not as robust as tests initially revealed and that the effect was dependent on a single small *p*-value. However, since the analysis was published, the first pre-registered inhibition training experiment on food choice has been published. Z. Chen, Holland, Quandt, Dijksterhuis, and Veling (2019) found clear evidence of training leading to increased preference for go foods compared to stop foods. An update to the *p*-curve analysis by Veling, Chen, Huaiyu, Quandt, and Holland (2019) showed that, with the addition of the experiments in the Z. Chen et al. (2019) paper, the *p*-curve now suggests that response inhibition training is effective in bringing about behaviour change.

However, in some domains research has been less successful. For example, in a randomised controlled trial using online go/no-go training, where images of smoking were paired with stopping, Bos et al. (2019) found no evidence of the effectiveness of the training. Indeed, smoking cessation rates reduced in line with a control group. Additionally, research investigating the use of inhibition training to reduce alcohol consumption has met with mixed success, with some empirical work not finding any difference in alcohol intake between training and control groups (Smith, Dash, Johnstone, Houben, & Field, 2017; A. Jones et al., 2018). Conversely, one recent trial by Strickland, Hill, Stoops, and Rush (2019) did find real-world reductions in drinking following go/no-go inhibition training. There is also a variable picture in the use of inhibition training to target problem gambling. Initial work showed how a short task that promoted cautious motor responses led to a reduction in betting scores in a gambling task conducted at least two hours after training compared to controls (Verbruggen, Adams, & Chambers, 2012). However, a follow-up study where the delay between training task and gambling was 24 hours found strong support for the null hypothesis, that is, there was no difference between controls and those who received the training (Verbruggen et al., 2013). Subsequent studies again found an effect,

with stop-signal training leading to people placing lower monetary bets in a subsequent gambling task (Stevens et al., 2015), although the effect was small.

These mixed effects have been confirmed in several meta-analyses and reviews. For example, a meta-analysis by Allom, Mullan, and Hagger (2016) found that whilst go/no-go training was effective its effects did not seem to persist over time (see also Turton, Bruidegom, Cardi, Hirsch, and Treasure, 2016). A. Jones et al. (2016) found that inhibition training led to a significant decrease in food and alcohol consumption compared to controls, although the overall effect size was small. Further analysis showed that the effect size was dependent on the training paradigm used, with the effect size being medium if only go/no-go training was included. It is worth reflecting on these results further. Research has found stop-signal training to be less effective than go/no-go training (Adams et al., 2017). However, the reasons behind this difference are not clear. It could be that due to the inherently lower successful inhibition in stop-signal tasks (resulting from failures to stop), such tasks provide a lower amount of inhibition training and thus produce lower effects; or perhaps stop-signal tasks, with less consistent stimulus-signal mappings (due to the delay in presentation of the stimuli and stop signal), encourage learning towards the stop cue, rather than the stimuli. Thus, without the stop cue being present in real-life inhibition is less successful (Veling, Lawrence, Chen, van Koningsbruggen, & Holland, 2017). One important factor of the go/no-go paradigms used in such work is that they tend to be incidental versions. By this I mean, that while participants categorise the images (for example as appearing on the left- or right-hand of the screen), stopping is signalled by another cue. For example in N. S. Lawrence, O'Sullivan, et al. (2015), participants had to stop responding when a rectangle surrounding the image turned bold. This differs from traditional go/no-go tasks where the images themselves determine the appropriate response (see Figure 1.1).

One outstanding issue surrounds the mechanisms that enable inhibition training to affect behaviour change. Though a full review of this question is outside the scope of this thesis, broadly there are two distinct pathways. One pathway suggests that training strengthens top-down inhibitory control towards no-go foods (Guerrieri, Nederkoorn, & Jansen, 2012), while another argues that the training creates associations between stopping and no-go foods (Verbruggen,

Best, et al., 2014; Stice, Lawrence, Kemps, & Veling, 2016). In addition, research also supports the notion that pairing stimuli with stopping responses leads to devaluation of strongly reward-associated cues that drive go responses (N. S. Lawrence, O'Sullivan, et al., 2015; Camp & Lawrence, 2019). Indeed, it may be that devaluation and bottom-up processes work together at different stages of the training to effect behaviour change, e.g., initial training creates response conflict (resulting from stopping to go foods) followed by the development of associatively-mediated inhibition (for a discussion see Veling et al., 2017). As a separate issue, it is important to distinguish between the effects of stimulus-specific response inhibition training (i.e. such as is the case in food and alcohol studies) and non-cue specific inhibition training effects which are more applicable to gambling research. A. Jones, Hardman, Lawrence, and Field (2017) in their review clearly distinguish between tasks that train top-down general inhibition (where participants must withhold responses to arbitrary cues) and tasks that train bottom-up stimulus-response associations (such as the food literature reviewed above). Overall, non-cue specific inhibition training has shown minimal effects in effecting behaviour change (Verbruggen et al., 2013; Bartsch, Kothe, Allom, Mullan, & Houben, 2016; yet see A. Jones et al., 2018).

To summarise, research supports the idea that pairing a stimulus with a stopping response can lead to slower reaction times (Verbruggen & Logan, 2008b) when these images are subsequently presented. These effects have been used to improve inhibition in the real-world and have been shown to be effective (though with small effect sizes) in a range of behaviours. The idea of applied response inhibition will be explored further in Chapter 5, but I shall now set the associatively-mediated inhibition hypothesis into a wider research context that can be used to further investigate human learning and which is key for the early empirical chapters of this thesis.

## **1.4 DUAL PROCESS THEORIES OF HUMAN BEHAVIOUR**

A popular form of theorising about human cognition is the notion of duality, that cognition can be divided into two distinct systems (Deutsch, 2016). One system is characterised as slow, effortful and deliberate, while the other is fast, automatic and effortless (Evans, 2008; Stacy & Wiers, 2010; Kahneman, 2011; McLaren et al., 2014; Sherman, Gawronski, & Trope, 2014; Strack & Deutsch,

2015). While some suggest that learning arises as result of a single, effortful, propositional process (Lovibond & Shanks, 2002; Mitchell, De Houwer, & Lovibond, 2009; Shanks, 2009), others (Stanovich & West, 2000; McLaren et al., 2014) argue that an additional, associative, system exists. This system is mechanical in nature and develops through detecting the frequency of events and the contingencies between them (McLaren et al., 2014). The general view in the field of learning is that humans (and other animals) are capable of learning using rules and reasoning, with the existence of a purely associative learning route being more contentious. In this section I review the evidence for the role of propositional learning in human behaviour, and then the research that demonstrates that associative learning can operate in humans independently of propositional knowledge.

### **1.4.1 Propositional learning in humans**

Mitchell et al. (2009; see also De Houwer, 2009) argue that associative learning results from the operation of a single controlled reasoning process. This system captures the nature of the relationships between events and as such contains both 'truth value' (Strack & Deutsch, 2004) and causality. Therefore, the resulting hypotheses can be proved true or false, and as such this leads to further learning. Support for the propositional approach rests on four strands of evidence: 1) that verbal instructions should be sufficient to produce learning even in the absence of an event (and thus the absence of the formation of associative links); 2) awareness is key to learning, and learning should only occur with awareness; 3) learning will be impaired with cognitive load; and 4) learning should be rational and rule-based. Of course, the point here is that all these things can be true at various times. So, simply demonstrating that these principles hold, at least some of the time, is not enough to rule out a dual-process theory.

In support of the first claim, studies have found that verbal instructions on their own have been sufficient to produce learning similar to that gained through experience of the actual contingency. For example, in Cook and Harris (1937) when participants were informed that a shock would always follow a tone, subsequent presentations of a tone lead to increased skin conductance even though a shock and tone were never presented together (see Smyth, Barnes-

Holmes, & Barnes-Holmes, 2008). The effect of verbal instructions holds for complex learning such as retrospective revaluation (Lovibond, 2003).

The second point, on the importance of awareness, finds support in work by Lovibond and Shanks (2002) showing that while there are many examples of work claiming unaware conditioning, in most cases the measures used to check for awareness were much more sensitive tests of learning than those used to assay awareness, and as such there is little evidence for learning without awareness.

The third point is that learning is affected by cognitive load. The propositional account argues that cognitive resources are required to attend, learn, and deploy appropriate rules. Therefore, under conditions where resources are low, either through secondary tasks or diverted attention, learning should be impaired. For example, the work by Dawson and colleagues (Dawson, 1970; Dawson & Biferno, 1973) showed that the presence of a masking task could impair conditional learning. In the experiments a classical conditioning design, where tones were or were not paired with shocks, was embedded within an auditory perception task, whereby participants were asked questions about the pitch of preceding tones. As is the case in conditioning designs, one tone predicted a shock, and another was never paired with a shock. Awareness was manipulated, with one group being told that tones predicted the probability of a shock and another receiving no such instructions. Galvanic skin response measures were taken as the dependent measure and knowledge of the contingencies was assessed by questionnaires and online expectancy ratings. If learning of the contingencies arises due to a resource intensive propositional system those in the unaware group should have impaired learning, with the masking task interfering with their ability to learn the contingencies. Of course, those in the aware group should be able to learn the contingencies as they have been informed about them. As predicted by the propositional account the imposition of the masking task impaired learning of the tone-shock contingencies. Participants classed as unaware of the conditioning contingencies did not show any difference in galvanic skin response between the tone paired with a shock and the tone never paired with a shock. This suggests that load is important to learning about contingencies, with the fact that an increased cognitive load impaired contingency learning suggesting that

learning within this task relied on propositional, rather than associative, processes (also see De Houwer & Beckers, 2003).

Yet, recent work paints a more complex picture. Seabrooke, Wills, Hogarth, and Mitchell (2019) found that while cognitive load affected performance in a complex outcome-response priming task, performance was unaffected in a simple task. This dissociation points to two separate routes. The fact that performance in a simple task was unaffected by cognitive load suggests that control processes were not key for this task, whilst impairment of performance in the complex task under cognitive load indicates that here controlled reasoning was required. However, the results from the simple task could occur through, as the authors' note, the deployment of a very simple rule. If the rule was very simple then it would be easily deployed and, if requiring only relatively little cognitive capacity, would be unaffected by the imposed cognitive load.

The final strand of support for the propositional account of learning comes from evidence that learning is always rational. For example, Shanks and Darby (1998) found that participants learnt rules in an allergy prediction paradigm rationally and acted in a manner inconsistent with associative learning accounts. In the experiment, participants were presented with various cues and outcomes. Crucially the design followed a rule whereby a compound cue (e.g., AB-) was the reverse of the outcome of its constituent parts (A+ and B+). To assay learning, in training, participants were presented with cues I+, J+, M-, and N- and at test shown their compounds, IJ and MN, having to predict the likelihood of cues leading to an allergic reaction. Crucially, the propositional and associative accounts would predict differing learning of these unseen compounds. If participants had learnt the rule, then one would predict that participants would rate MN over IJ as more likely to induce an allergic reaction (as the compounds' constituent parts did not result in a reaction). However, if associative learning was guiding learning, then participants should rate IJ, not MN, as more likely to induce a reaction, as the compounds constituent parts also led to a reaction. That is, seeing the cue IJ should activate representations of I+ and J+ summing to IJ++. Thus, the associative account would predict that unseen compounds would lead to the same outcome as the constituent parts, while propositional learning would predict the opposite outcome. In support of the propositional account, participants judged that the allergic reaction was



more likely to occur for cue MN than IJ. Therefore, there is evidence that participants had seemingly learnt and deployed the appropriate rule. Yet, Wills, Graham, Koh, McLaren, and Rolland (2011) found that learning depended on cognitive load; specifically when participants were under heavy cognitive load in a set-up similar to Shanks and Darby (1998), participants displayed learning consistent with surface similarity rather than the application of the rule (also see continued discussion in section 1.4.2).

### **1.4.2 Associative learning in humans**

The dual-process route as posited by McLaren et al. (2014) does not seek to dismiss the existence of a propositional account for human learning. It accepts that people can solve problems by rational hypothesis testing and using rules, but argues that a simpler, associative based, system can also influence behaviour. Furthermore, as McLaren et al. (2014) note, in some circumstances one would expect the two systems to operate in parallel and both to contribute to behaviour.

Before discussing examples of human learning that cannot be easily explained by propositional knowledge, and thus support a dual-process account, it is worth noting the occurrence of associative learning in other animals. For example, there is wide ranging evidence that the sea slug *Aplysia californica* learns by the development of increasing or decreasing strength in synaptic connections between neurons (Kandel, 1976; Hawkins, Clark, & Kandel, 2006) and can produce behaviour predicted by learning theories, such as classical conditioning and conditional discrimination (Jami, Wright, & Glanzman, 2007). Whilst it is clear that making strong claims about human learning on the evidence from work with sea slugs would be dangerous, if we move closer to our own genetic heritage there is evidence that the activity of dopamine neurons in primates codes for prediction-error (Schultz, 1998), and that this signal is similar to a teaching signal as predicted by reinforcement learning theories such as the Rescorla–Wagner model (Rescorla & Wagner, 1972). Indeed, it has been argued that the development of associative learning marked the beginning of a new stage in the development of life on earth approximately 541 million years ago, the Cambrian explosion (Ginsburg & Jablonka, 2010). Therefore, to argue that human learning is solely propositional is to suggest that, despite common evolutionary ancestors, humans developed a separate system of

learning – which in evolutionary biology terms is highly improbable. It is more parsimonious to assume that both associate and propositional systems could have developed in humans.

As well as theoretical arguments there are also examples of learning in humans that cannot be explained by propositional learning accounts. One such example is the Perruchet effect (Perruchet, 1985). In the original experiment, participants were presented with a partial reinforcement schedule in which a tone (the conditioned stimulus, CS) was played on every trial, and on 50% of trials was followed by a puff of air (the unconditioned stimulus, US). This schedule resulted in a conditioned response (CR) of an eye-blink on presentation of the tone. The order of trials was pseudo-randomised which led to the creation of varying lengths of runs. These consisted of either CS-US pairings or just CS trials. Thus, trials were either followed by the same type of trials (CS-US, CS-US) or followed by the other trial type (CS-US, CS). After each presentation of the trials participants were asked to rate their expectancy of receiving the US (the air-puff) in the following trial. Whilst the chance of receiving an air puff was constant throughout the experiment at 50%, participants' prediction of the chance of receiving the US decreased as the number of consecutive CS-US trials increased. Thus, participants displayed the gamblers fallacy (Burns & Corpus, 2004), i.e. the erroneous belief that an outcome is less likely to occur in the future if it has already occurred. However, their conditioned responding (in effect another measure of learning) displayed the opposite pattern, with more consecutive presentations of the CS-US leading to greater predicted probability of a CR. That is, participants showed an increase in CR and a decrease in reported predictions of the US following a run of CS-US presentations, with the opposite pattern being seen on non-reinforced trials, i.e. CR reducing but expectancies for the US increasing. Thus, there are two directly opposing results; one measure shows a decrease following reinforcement, while another shows an increase. The effect has been found using other measures of learning such as reaction times (Perruchet, Cleeremans, & Destrebecqz, 2006; yet see critique by Mitchell, Wardle, Lovibond, Weidemann, & Chang, 2010; with a response by Barrett & Livesey, 2010) and galvanic skin response (McAndrew, Jones, McLaren, & McLaren, 2012). As Mitchell et al. (2009) concede this double dissociation between awareness and automatic conditioning is hard to explain within a single learning system but is exactly the pattern predicted by

associative learning theory (McLaren et al., 2014). Such a model would argue that autonomic measures (such as the eye-blink) are governed by the strength of the associations between the CS and the US, while the expectancy predictions are a product of conscious processes falling to the gamblers fallacy (see Tversky & Kahneman, 1974). Despite close scrutiny (Weidemann, Tangen, Lovibond, & Mitchell, 2009; Weidemann, Broderick, Lovibond, & Mitchell, 2012) and greater understanding of the contribution of non-associative processes (Livesey & Costa, 2014; Weidemann, McAndrew, Livesey, & McLaren, 2016) the effect still provides robust evidence in support of a dual-process account of behaviour (Perruchet, 2015), suggesting a separation between awareness and conditioning that cannot be easily explained via a single propositional system.

Research reviewed earlier also provides support for the notion that associative learning processes can influence behaviour. Previously I discussed work by Shanks and Darby (1998) who found evidence of rule-based learning in humans that cannot be readily explained by dual-process accounts. However, this was not the whole picture. Within the study there were a subset of participants whose performance was comparatively poor suggesting that they had not learnt the rule. In fact, the participants in question displayed performance consistent with an associative learning account, and at complete odds to behaviour predicted if one was using the task rule. Therefore, there seems to be a dissociation between participants who use rule-based learning and those whose behaviour supports an associative learning account. These results support the idea that once they have been learnt, rules are easy to apply to related problems (e.g., the novel MN compound seen in the test phase), but that in their absence participants fall back on associative learning processes. This is of course not to say that all participants would not have learnt the rules with enough training, nor does it dismiss rule-based learning. Indeed, the fact that most participants (by a 2:1 ratio) were able to learn the rules is a strong indication that conscious processes dominate behaviour. I return to this idea of a unitary account below but to summarise the current section it is my view that a single, propositional model cannot account for all the effects I have described so far, particularly the research on response inhibition. Furthermore, a dual-process framework offers the flexibility to explain not only the research outlined but offers a clear theoretical anchor from which to investigate how associative learning underpins human learning.

### 1.4.3 Association and Cognition

In this section so far, I have presented the evidence for each pathway as if they worked independently. However, there is good reason to believe they operate in an interrelated fashion, with the associative system providing the basis for propositional behaviour (McLaren et al., 2014; Verbruggen, McLaren, et al., 2014; later developed in McLaren et al., 2019; similarly see Abrahamse, Braem, Notebaert, & Verguts, 2016, and Evans & Stanovich, 2013). This view argues that rather than there being two separate and competing theories of human learning, both processes feed into one system which mostly acts through propositional processes but can act via associative processes. To put this theory into context, as McLaren et al. (2019) p.16 state:

“We are convinced that cognitive processes can prevent the expression of any associative learning. They don’t have to, but they can do so, and this is the default. Otherwise, we would be at the mercy of events and our environment. As an example of what might happen if this were not the case, if you saw a chair, you would inevitably sit in it because of the long-standing association between stimulus and response. If this is not to be the outcome, then the expression of associative learning has to be inhibited by cognitive control in most circumstances. However, associative processes do support learning in the background. This learning might not inevitably be expressed, but it does automatically take place.”

Therefore, this account argues that behaviour is mostly consciously driven, but that this cognition is built on top of associative learning, which provides the basic building blocks for propositional learning to occur. Importantly, the associative learning takes place automatically in the background and can influence behaviour when explicit processes are weak (McLaren et al., 2019), a point I will return to later with regards to driving.

A piece of evidence in support of this theory comes from a task switching experiment in McLaren et al. (2019). The authors wished to investigate if participants could change from one mode of learning to another and what learning would transfer. The experiment was a typical bi-conditioning design in which two tasks were cued by certain shapes: either categorising digits as odd or even, or categorising digits as higher or lower than five. Typical with these experiments’ participants were informed about the rules of the task. However,

without the rules, the design becomes more akin to associative learning experiments where participants must learn the cue+stimuli – response (CSR) mappings by trial and error. Thus, in the experiment, participants either started with rule based (TASK group) or associative learning (CSR group) instructions and then, half-way through, switched to the other type of instructions. Results showed that for participants in the CSR-TASK group, performance after switch worsened and was similar to performance in the first block for a control TASK-TASK group. However, for those in the TASK-CSR group, performance was somewhat protected, and at switch their learning matched that of the second block of a CSR-CSR control group. Thus, those in the CSR-TASK had to ‘start again’ in terms of performance, while those in the TASK-CSR group did not. The fact that performance in the TASK-CSR group did not collapse suggests that whilst under task instructions participants were able to learn about the CSR mappings, with such learning taking place ‘in the background’, and that once cognition had been surrendered by the switch of instructions, such learning came to the fore. The results of the CSR-TASK group suggest that once propositional learning is engaged then the behaviour arising from associative learning is suppressed. It should be noted that for this group the results are unexpected in traditional dual-processing accounts as one would predict to see some effect of the learning of the CSR mappings in the CSR-TASK manifest itself in performance in the second block. Overall, these results support the notion of a ‘two processes, one system’ model.

One consequence of the argument that associative learning is at least partly the basis of all behaviour is that it would be expected that the basic ability to inhibit a response could be associatively driven. Evidence in support of this notion comes from task-switching experiments. Here, participants are exposed to stimuli and required to categorise them based on some feature of the stimulus. Humans often show switch costs, longer reaction times to a stimulus when the feature or dimension they are asked to use is switched rather than repeated across trials (Vandierendonck, Liefoghe, & Verbruggen, 2010). This is often ascribed to the use of rules to complete the task, with the switch costs representing the time required to recall a new set of rules (task set) into working memory (Monsell, 2003). However, there is evidence that humans can solve these tasks without the use of rules (Dreisbach, Goschke, & Haider, 2007; Dreisbach, 2012), particularly with relatively small stimulus sets (Forrest,

Monsell, & McLaren, 2014), with participants who complete the tasks without rules showing no, or relatively little, switch costs. Furthermore, even pigeons, who have generally been found to lack executive control functions (Lea & Wills, 2008, but see Rose & Colombo, 2005 and Castro & Wasserman, 2016), have been found to complete task-switching paradigms (Meier, Lea, & McLaren, 2016). Though, like the human participants, they show no switch costs. Similar findings have been found in monkeys (Smith & Beran, 2018). Thus, performance for pigeons, monkeys, and humans on these tasks indicate that response inhibition could be to be, in part, a product of associative learning.

## **1.5 BEHAVIOUR CHANGE IN DRIVING**

Having reviewed the wider literature relating to executive control, inhibition, and the idea of dual-processes in learning, I now turn to the applied aspect of this thesis, driving behaviour. Here the state of behaviour change is reviewed with a focus on the role that associatively-mediated processes play in driving behaviour.

One of the most well-known models in road safety is Ajzen's (1991) Theory of Planned Behaviour (itself an extension of the Theory of Reasoned Action; Ajzen & Fishbein, 1980). A core component of the theory is intention; the more intent an individual has to commit an action, the more likely it is that an individual will engage in that behaviour. The theory argues that three factors determine the strength of an intention (see Figure 1.4); belief about the likely consequences of the behaviour (attitude), perceived expectation of others (subjective norms), and perceived ability to perform the behaviour (perceived behaviour control). These three factors form an overall behavioural intention which is the immediate antecedent to behaviour. Taking them in reverse order, perceived behaviour control refers to an individual's perception of their ability to undertake a behaviour. It can be seen as a superordinate construct (Ajzen, 2002a) that contains two separate and individual components which both impact the strength of the overall construct: self-efficacy (an individual's view of the difficulty or not of a behaviour) and controllability (the extent to which the behaviour is up to an individual). However, this construct is not universally accepted (Kiriakidis, 2017), with some researchers adding self-efficacy as a separate construct to the Theory of Planned Behaviour model (Terry & O'Leary,

1995; Tolma, Reininger, Evans, & Ureda, 2006). The next element of the model is subjective norms, which refers to the societal norms around an action; the pressure to conform, or not, to a behaviour. The more favourable the norms to completing an action, the stronger the intention to carry it out will be. The final factor in the model is attitude toward the behaviour. This refers to a person's appraisal of a behaviour. This is driven by the consequences of an action, with positive consequences leading to more positive appraisals, and thus a higher chance that the behaviour will be performed. The model also suggests that perceived behaviour control can be used to directly predict real behaviour.

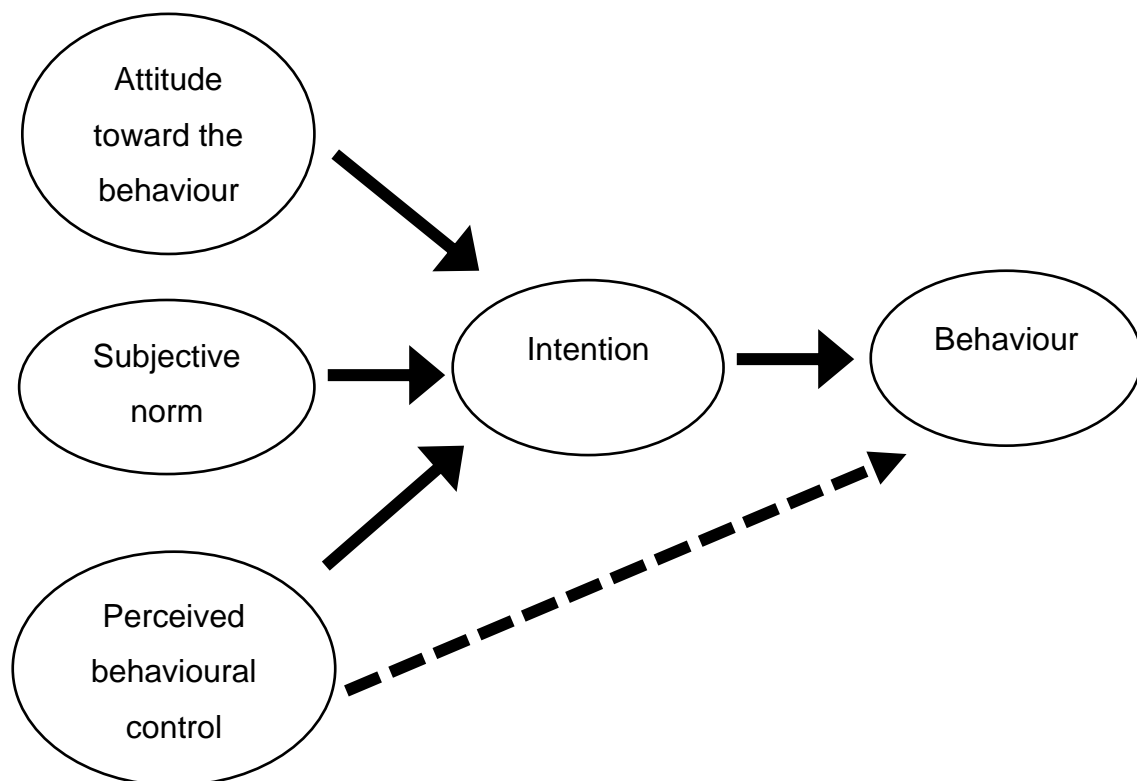


Figure 1.4. Schematic of the Theory of Planned Behaviour. Original drawing.

The model has been found to successfully predict a range of real-world driving violations. For example, Nemme and White (2010) found that the model predicted 35% of the total variance in sending texts while driving, and Elliott et al (Elliott, Armitage, & Baughan, 2003; Elliott, Armitage, & Baughan, 2007) reported that the model explained 31-39% of the variance in observed speeding in a driving simulator. The model has also been used to account for 33% of the variance in driving over the alcohol limit (Castanier, Deroche, & Woodman, 2013) and can predict red light jumping in motorcyclists (Satiennam, Satiennam, Triyabutra, & Rujopakarn, 2018).

However, despite the success of the Theory of Planned Behaviour in predicting behavioural outcomes, there has been less success in using the theory to change behaviour. Poulter and McKenna (2010) evaluated a road safety course used in UK schools aimed at pre-drivers based on the Theory of Planned Behaviour. They found that the course only produced a short-term increase in safer driving attitudes. Glendon, McNally, Jarvis, Chalmers, and Salisbury (2014) found that compared to a matched control group, novice drivers reported riskier driving attitudes over the course of a driver education course and at a six-week follow-up. The authors argued that such unexpected results might be a result of participants using defence mechanisms (e.g., optimism bias, Weinstein & Klein, 1996) to help allay fears of mortality caused by the use of crash statistics in the driving course. Furthermore, many reviews and meta-analyses have found limited evidence that driver improvement courses benefit road safety. For example a recent meta-analysis by Steinmetz, Knappstein, Ajzen, Schmidt, and Kabst (2016) found an effect size of 0.26 for interventions using the Theory of Planned Behaviour to change traffic behaviour. For contrast, the effect size for interventions aimed at physical activity was 0.54 (also see Roberts & Kwan, 2001; Ker et al., 2005; Peck, 2011). However, some research has found benefits of using Theory of Planned Behaviour-based interventions, with Quine, Rutter, and Arnold (2001) finding that such an intervention increased wearing of cycle helmets and Stead, Tagg, MacKintosh, and Eadie (2004) finding sizeable effects resulting from an intervention aimed at reducing speeding. Others have argued that the effects of such education programmes have been underestimated in meta-analyses (af Wählberg, 2018). There is general agreement that more well-controlled studies are required to fully understand the benefits of such educational interventions (Beanland, Goode, Salmon, & Lenné, 2013).

### **1.5.1 Associatively-mediated behaviour in driving**

It seems that focusing on the constructs in the Theory of Planned Behaviour to effect behaviour change does not capture all aspects of human cognition that influence risky driving behaviour (Conner & Sparks, 2005). It is worth highlighting that the Theory of Planned Behaviour relies on an actor's behaviour arising from conscious, rational decisions (the 'economist's perspective' on decision-making) rather than behaviour that stems from automatic or routinised



processes (for further see Kahneman, 2011). As reviewed earlier there is evidence that human behaviour can be associatively-mediated, and, in part, cued by environmental stimuli. This next section reviews the evidence for such learning in a driving context.

Before proceeding further, it should be noted that the research communities of road safety and associative learning are distinct from one another with relatively little theoretical overlap. This has led to differing language to describe similar constructs. In the road safety literature, behaviour which the associative learning literature might describe as associatively-mediated is termed habitual. Habits are argued to be formed through associative learning resulting from the regular pairing of actions and events (Wood, Quinn, & Kashy, 2002), and thus the activation of a habit leads to the enactment of a specific, well-defined, response (Wood & R nger, 2016). However, of course, not all frequently enacted behaviour is habitual (Verplanken & Wood, 2006). Additionally, once developed, habits are held to be activated in response to an environmental stimuli without mediation by conscious goals (Wood & Neal, 2007; Gardner, 2012) and are in some sense the 'default' setting for behaviour (Evans & Stanovich, 2013), though this viewpoint is not without critique (De Houwer, Tanaka, Moors, & Tibboel, 2018). Akin to the associative learning literature, the prevention of habitual behaviour is dependent on cognition control (Wood & R nger, 2016). A wider review of habits is outside the scope of this review, and rather the focus here is to highlight the conceptual overlaps between the two research fields relevant for this thesis.

In support of the role of automatic behaviours in driving, Verplanken, Aarts, Knippenberg, and Moonen (1998) found that car use was predicted by both habits and the Theory of Planned Behaviour. However, the relationship between the Theory of Planned Behaviour and actual behaviour was moderated by habit; intention only significantly predicted behaviour when habit was weak. Other research by Lheureux, Auzoult, Charlois, Hardy-Massard, and Minary (2016) explored the separate influences of habits and planned behaviour on drink driving. Controlling for the constructs in the Theory of Planned Behaviour, habit was still found to significantly predict behaviour. However, in all model's intention was found to be the greatest predictor of behaviour. This led the authors to conclude that behaviour is a result of both intentional and habitual

processes. A similar conclusion was reached by Elliott and Thomson (2010) who found that though intent was the largest predictor of speeding behaviour (explaining 47% of the behaviour), habit was a significant predictor, explaining an additional 4% of the variance. One issue regarding the above three studies is that they measure real-world behaviour via self-report measures; these have been found to overestimate the variance explained by the Theory of Planned Behaviour (McEachan, Conner, Taylor, & Lawton, 2011). This is likely to explain the differing conclusions, as issues around social desirability are more likely to affect sensitive topics such as drink driving and speeding than car use. Overall, the research demonstrates the importance of both rational decision making and automatic processes in governing behaviour related to driving (see reviews by Ouellette & Wood, 1998 and Wood & R nger, 2016).

#### **1.5.1.1 Measuring automaticity in driving**

Characterising habitual behaviours as a distinct construct from planned behaviours has not been without its critics (Ajzen, 2002b; de Wit et al., 2018). One of the main measures of habit is the Self-Report Habit Index, a 12-item index developed by Verplanken and Orbell (2003). Though popular, various methodological concerns have been raised. For instance, there have been questions surrounding the operationalisation of habit in the scale (Sniehotta & Pesseau, 2011; Gardner & Tang, 2014) and the difficulty in asking individuals to assess awareness of supposedly automatic processes (Hagger, Rebar, Mullan, Lipp, & Chatzisarantis, 2015). Labrecque and Wood (2015) suggest future research should use cue–response association tests (that address the relationship between a context, goals, and a behaviour) rather than self-report measures to provide more valid measures of automaticity. As Gardner (2015a) notes, more work is required to develop reliable measures of automatic behaviours.

#### **1.5.2 Controlling associatively-mediated behaviours**

The research reviewed thus far supports the argument that ‘habits’ are a key predictor in driving behaviour, and that they can promote unsafe practices. Given the notion that habits are the default behaviour option it is likely that the behaviours primed by driving habits will dominate driving, unless cognitive control can be brought to bear. Unfortunately, research has shown that driving is a cognitively demanding activity. For example, Wadley et al. (2009) found that

those with cognitive impairments, that effected executive control, showed lower driving performance, such as poor lane control, compared to healthy controls. Briggs, Hole, and Turner (2017) compared behaviour on a hazard perception test between one group who had to complete the task (low load condition) and another who simultaneously had a conversation with the experimenter on a hands-free mobile phone (high load condition). Findings showed those in the high load condition detected fewer unexpected events and took longer to react to events than those in the low load condition.

Given the nature of driving interventions that rely on cognition control are unlikely to be successful. For example, Elliott and Armitage (2006) explored the use of implementation intentions (an if-then plan, Gollwitzer, Sheeran, Trotschel, & Webb, 2011), to increase compliance with speed limits. Comparing self-reported compliance at a one-month follow-up, participants in the experimental condition, who had formed implementation intentions, showed significantly increased compliance with speed limits compared to a control condition who did not form implementation intentions. However, further analysis found that the intervention was only effective for participants who had a goal not to speed. Additionally, there is evidence that these implementation intentions are not completely automatised and that the deployment of such intentions is impaired under heavy cognitive load (McDaniel & Scullin, 2010).

There have been growing calls for a greater focus on 'habits' when designing health interventions (Marteau, Hollands, & Fletcher, 2012; Sheeran, Gollwitzer, & Bargh, 2013). In his review Gardner (2015b) found only 38% of interventions directly addressed habitual behaviour, and in a recent review of behaviour change techniques in road safety and Fylan (2017) emphasised the need to establish 'good' habits. Changing associatively-mediated processes is arguably the Holy Grail in designing an intervention that will deliver a new behavioural 'default'. Earlier evidence highlighted the fact that there had been some success in using associatively-mediated inhibition training with regards to food and drinking consumption – could such approaches benefit driving behaviours?

### **1.5.2.1 Inhibition in driving**

Given that driving frequently involves cancelling an already actioned motor response, e.g., the traffic light suddenly turns red, this suggests that inhibition training that targets the prepotent response to a cue could lead to safer driving.

Supporting the importance of inhibition in driving, studies have found that poor impulse control leads to more risky driving (Jongen, Brijs, Komlos, Brijs, & Wets, 2011; Bachoo et al., 2013; O'Brien & Gormley, 2013; Bıçaksız & Özkan, 2016; Sani et al., 2017), and that those with attention deficit/hyperactivity disorder (Barkley & Cox, 2007; Groom et al., 2015) are more likely to commit traffic violations. Results inconsistent with these findings have also been found. For example, Mäntylä, Karlsson, and Marklund (2009) explored the role of cognitive control in driving behaviour (also see Renner & Anderle, 2000). They measured participant's performance on tasks requiring mental shifting, working memory updating, and response inhibition, as well as behaviour in a simulated driving scenario. Only the correlation between driving performance and working memory updating was found to be significant. However, compared to other studies, the participants in Mäntylä et al. (2009) had minimal driving experience, with only 4% of the sample holding a driving licence compared to 100% in Sani et al. (2017).

To date there has only been one experiment investigating whether inhibition training could improve risky driving. Hatfield et al. (2018) compared performance (e.g., average speed) in a driving simulator scenario pre- and post-training between a control group and a group that received inhibition training. The experimental group received five consecutive days of go/no-go training against a control group who received a filler task for the same amount of time. The inhibition task required participants to respond to images of computer-generated junctions. In the experimental group, participants had to respond if it would be safe to turn right, and withhold a response if not safe to do so, while those in the control group had to press SAFE if they felt it would be safe to turn right, and to press 'UNSAFE' if they did not. Compared to the pre-training driving, the results showed there was little evidence of transfer, with no significant difference in driving behaviour between groups. However, there was a tendency for increase stopping at red traffic lights in the experimental group following training, relative to a control group. I shall return to this research in Chapter 5.

## 1.6 CONCLUSIONS AND NEXT STEPS

This introduction began by highlighting the dangers inherent in driving, specifically behaviour at traffic lights. We have seen how associatively-mediated behaviours play an important part in driving, and how traditional education-based interventions have not addressed such associative processes. In terms of developing road safety-related interventions it is essential to explore the role of associative processes in the target behaviour. If associative learning supports the development of safe driving, then efforts should be made to enhance existing interventions. If, on the other hand, associative processes promote risky driving then a new type of intervention will be needed. Therefore, the initial focus of this thesis will be to develop a laboratory paradigm that allows for the exploration of associative learning based on the contingencies at UK traffic lights. Chapter 2 will build on the work of Bowditch et al. (2016) in doing this, with Chapters 3 and 4 developing the paradigm further to take into account the role of sequences and task sets experienced at traffic lights.

The last empirical chapter of this thesis will consider the application of inhibition training to driving. The evidence presented here suggests that by changing the response associated with a cue (e.g., from go to stop) one can change behaviour (e.g., a reduction in eating of chocolate). In a similar vein as Hatfield et al. (2018), I wish to test if the associatively-mediated inhibition hypothesis can be applied to traffic lights: could response inhibition training reduce crossing of amber traffic lights? The development of this training will be informed by the results from the early chapters. Thus Chapter 5 will apply the associatively-mediated inhibition hypothesis to a novel domain of human behaviour (driving) and so help establish further evidence of the effectiveness of such training.



### ASSOCIATIVE PROCESSES I: BASIC CONTINGENCIES EXPERIENCED AT TRAFFIC LIGHTS

As outlined in Chapter 1 there is a need to develop interventions that address what might be termed associatively learnt behaviours at traffic lights. However, exactly *what* is learnt at an associative level is yet to be established. This chapter explores the learning of contingencies experienced at traffic lights and what effect they might have on behaviour. Additionally, the final two experiments in the chapter investigate the influence of the effective outcome on the learning, both as an issue in its' own right, and as evidence that associative processes are producing the effects.

#### 2.1 CONTINGENCIES AT TRAFFIC LIGHTS

The UK traffic light signal changes from green to amber to indicate drivers should prepare to stop; then to red meaning stop; then to a conjunction of red and amber to tell drivers to get ready to start, and finally back to green (which, of course, means go). The rules governing the responses that should follow at each stage of this sequence are clearly laid out in the Highway Code, yet they are not always observed in practice. One possible reason for this not a deliberate lack of compliance on the part of the driver, but rather the effect that the experience of the contingencies which occur at traffic lights whilst driving will, in time, have on the individual. Could it be that the experience leads to learning that captures these contingencies via association, and this then leads to behaviour that does not respect the explicit rules that apply to these situations? In other words, does the combination of rules, signals and typical driving behaviour lead people to learn stimulus-response reactions that then quite naturally predispose them to break those same rules? Previous research exploring the contingencies between traffic lights and behaviour has focused on engineering solutions to bring about safer driving, from altering the timings of the light pattern (Jason, Neal, & Marinakis, 1985) to adding countdown timers to the traffic light sequence (Felicio, Grepo, Reyes, & Yupingkun, 2015). Other research has looked from the perspective of understanding how personal factors (e.g., time and social context) are predictive of behaviour around traffic

lights (Palat & Delhomme, 2016), rather than how experience of the contingencies between lights and responses may come to cue a certain behaviour.

## **2.2 LEARNING ABOUT AN OUTCOME**

Bowditch et al. (2016) provides a framework to explore the learning of contingencies. They developed an incidental go/no-go task, whereby shapes appeared in the middle of the screen followed by a circle on the left or right. If this circle was white participants had to respond, and if it was coloured participants had to withhold a response. While the shapes in the middle were predictive of the response required this was not revealed to participants, and as such, any learning about the shapes was hypothesised to be associatively driven (see Yeates, Jones, Wills, Aitken, & McLaren, 2013). In the training phase one shape was paired with stopping (75% stop) and another with going (25% stop), while at test all shapes were 50% stop. If participants were developing stimulus–response associations during training it would be expected that performance at test would differ for the two shapes, with the 75% stop shape having longer reaction times and less commission errors compared to the 25% go shape. Results supported this hypothesis, with significantly longer reaction times and marginally significantly less commission errors for the 75% stop shape than the 25% go shape. Thus, the study demonstrates how participants can learn the contingencies between shapes and stopping in a relatively quick and practicable experiment.

Another feature of the experiment was the use of multiple stop signals. Analyses comparing performance with shapes paired with just one stop signal (that is one colour) against shapes paired with more than one coloured stop signal found enhanced learning in the multiple-signal group. That is, reducing the contingences between cues and specific signals resulted in more robust slowing of reaction times between stop and go cues. The authors argue that this enhancement was due to the development of cue–stop associations, rather than cue–signal associations, with the use of multiple-signals causing learning to shift to the consistent cue, rather than the inconsistent stop signal. The exact nature of the learning and the balance between a direct cue–stop association or an indirect cue-signal pathway is not the focus of this thesis (see Verbruggen,



McLaren, et al., 2014; Best et al., 2016). It is sufficient to note that the experiments in this, and subsequent, chapters will use multiple stop signals simply as a way of creating more effective learning, and thus clearer results (in terms of reaction times) between shapes with different contingencies.

## **2.3 PRESENT EXPERIMENTS**

The work by Bowditch et al. (2016) provides the departure point for this chapter. Their results showed that the incidental go/no-go task enabled the development of associative learning of the contingencies between cues and responses and, as such, is ideal for the exploration of contingency learning at traffic lights. As seen in Chapter 1, the framework by McLaren et al. (2019) hypothesises that when control is weak, underlying associative processes come to affect behaviour. Whilst, other work has shown that driving is cognitively demanding (Wadley et al., 2009). Therefore, it seems sensible to begin to understand the nature of associative learning at traffic lights. In the three experiments in this chapter I present an adapted version of the incidental go/no-go task, where arbitrary shapes are paired with various degrees of stopping to reflect the contingencies experienced at traffic lights. The experiments aimed to capture the contingencies between traffic light signals and the typical responses made to them in a task that superficially was quite unlike driving so as to be able to study them in a pure form, without the deployment of the rules used in driving. In summary, in this chapter I will investigate the effect that the contingencies experienced between traffic lights and permitted responses (stopping/starting) have upon driving behaviour.

## **2.4 EXPERIMENT 1**

To begin, it is necessary to decide what the contingencies experienced at UK traffic lights are. While a singleton green light (G) always means go and red (R) always means stop, the contingencies around red and amber (RA) and amber (A) are less clear. Taking amber first, while the Highway Code specifies that drivers should stop when they see this light, it is not always the case that a stop is required 100% of the time. The Highway Code allows drivers to go at amber lights if the stop line has already been crossed or if the driver is too near the stop line to stop safely. In practice, these provisions have afforded drivers some leeway and a solo amber light is not seen as a strong stop cue, with eight out of

ten drivers admitting to crossing amber lights, and nearly four out of ten drivers saying they rarely stop at amber lights (Thrifty, 2011). Accordingly, it could be argued that the Highway Code affords amber a fairly neutral contingency (and it is certainly experienced as such on the road), thus A will be treated as a 50% stop cue. In terms of red and amber lights in combination, while the Highway Code states that drivers should not cross the stop line while this light is showing, it is acceptable to get ready to move away, e.g., drivers might release their handbrake in order to move away when the green light shows. Therefore, to encompass this, RA will be a go cue rather than a stop cue. These contingencies mean that whilst red on its own signals stop, in conjunction with amber it cues readiness to start, while amber on its own is a neutral cue. Overall, ignoring the sequence information inherent in the typical experience of UK traffic lights, this leads to the following set of contingencies, where '+' denotes 'stop', '±' indicates '50/50 stop', and '-' denotes 'go', and R, A, and G stand for red, amber and green respectively: R+, G-, A±, RA-.

As well as deciding the appropriate responses to each of the cues, the experiment needs to capture the task set drivers are experiencing at traffic lights. When approaching lights, drivers could be in either a 'go' task set - that is, looking for signals that indicate permission to continue - or in a 'stop' task set, looking for signals to stop. In line with the work by Bowditch (2016), Experiment 1 used a procedure that should cause stopping to be the effective outcome from what is learnt. It was felt that making the outcome stop rather than go better reflected drivers' decision-making when approaching traffic lights, where one is looking for a stop signal in case the brake needs to be applied. Therefore, given that stop is the outcome and '+' is typically used to denote the outcome of a task this means that '+' will indicate a stop response and '-' a go response.

What would one expect to be learnt? Based on standard associative theories (for review see Pearce & Bouton, 2001) clearly cue R will become associated with stopping to some extent, and cue G with not stopping. The RA-contingency may tend to cause A± to become a go cue, while the A± contingency itself might promote a weak stop or go association to amber. Therefore, one question for Experiment 1 is: does experience of the

contingencies experienced at traffic lights while under a stop task set lead to amber becoming a stop or a go cue?

## **2.4.1 Method**

### **2.4.1.1 Participants**

Fifty participants participated in exchange for payment of £5 or one course credit (see Results section for details on the outlier removal process). Of these 50 participants, 41 were female with an overall mean age (with one missing data point) of 21.02 (SD = 5.43). The inclusion criteria were that participants had to be 18-65 years old, have normal or corrected vision, and not be colour blind. Given the use of mixed-effects models analysis, traditional power techniques were inappropriate (Johnson, Barry, Ferguson, & Müller, 2015). Therefore, the decision was made to use a sample size of 50 in line with similar past research (Bowditch et al., 2016). A post hoc power analysis using the R package SIMR (Green & MacLeod, 2016) found the study had a power of 91.40% to detect a difference of 18ms between a go vs. stop cue, where a priori one would expect to find an effect. Specifically, this test was run on G vs. B, see section 2.4.1.2 for details.

### **2.4.1.2 Design**

The experiment used a within-participants design to compare performance to cues over time (see Table 2.1 for design<sup>1</sup>). Overall, there was 1 calibration block, 8 training blocks, and 2 test blocks with a 10 second break between each block.

---

<sup>1</sup> A comment on the notation used in Table 2.1 and throughout. The tables use + and - to denote what type of response (go or stop) participants are most likely to make to a cue. Other notations were considered, such as subscript 'g' and 's' or 'go' and 'stop'. However it was decided that using + and - to denote the required response best suited this work, given that + and - not only help to convey the appropriate response but also the task set of the experiment (as + is used to denote that the response to the cue is the desired outcome). As will be clear later, the task set and changes to it form a key part of the experimental design. Whilst the use of + and - means that in some experiments - represents going counter to normal expectations, this usefully serves to highlight the change in task set between experiments.

The task was designed to mimic the contingencies around traffic lights with cue G being a go cue (analogous to a green traffic light) and cue R being a stop cue (corresponding to a red traffic light). However, to reduce the likelihood of participants explicitly realising the relationship between cues and required responses (as I was interested in the associative learning arising from experience rather than rule-based performance), and to obtain commission data, cues G and R were not 100% go and stop respectively but rather set at 75%. Cue A was designed to have contingencies equivalent to an amber traffic light and cue RA equalled a red/amber traffic light (i.e. a 75% go cue). Of course, this partial reinforcement, while helping to keep learning incidental, does affect the mapping of the task onto traffic lights as there are no instances where green traffic lights signal stop or red signals go (this issue is addressed in later chapters). Cues B, I, P, IP were control cues for G, R, A, RA respectively and were set at 50% stop. These cues were designed to provide a baseline to compare learning to the traffic light cues against. It is important to note that while the cue G was capturing the contingences experienced at a green traffic light, it did not in not in any way physically resemble a green traffic light (see Figure 2.1 for examples of this).

Phase	Blocks	Trials per block	N per type	Design	
Calibration	1	48	48		J±
Training	8	144	16	G-, R+, A±, RA-, B±, I±, P±, IP±	J±
Test	2	144	16	G, R, A, RA, B, I, P, IP	J±

Table 2.1. Summary of Experiment 1 design. Letters represent coloured cues. Trials were go 75% of the time (-), stop 75% of the time (+), or stop 50% of the time (±). At test all trials were 50% stop and so the cues were now non-predictive.

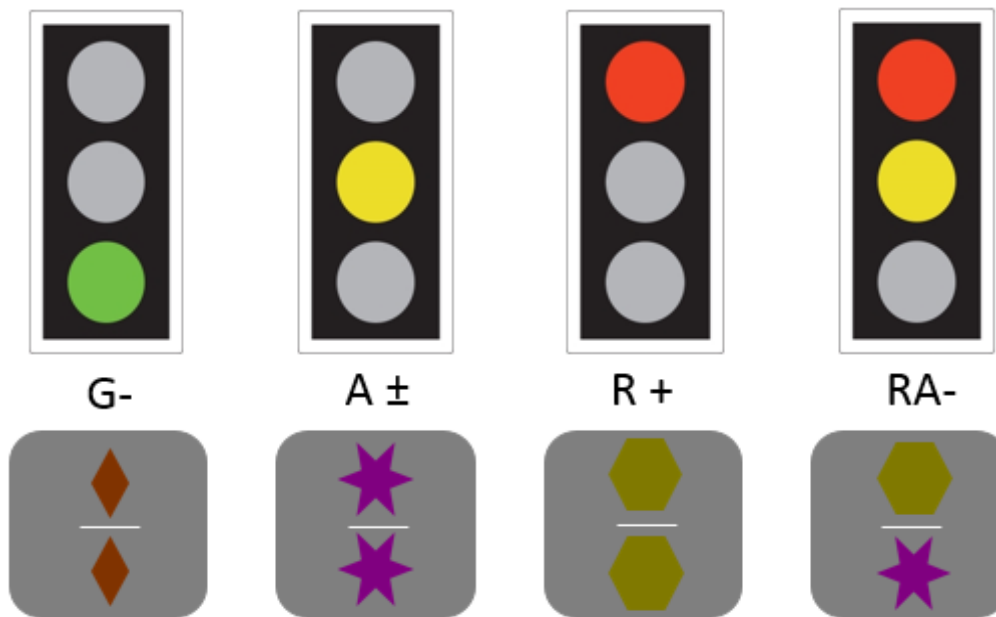


Figure 2.1. Traffic lights to cues in Experiment 1. This figure illustrates how each cue (and thus shape) represented a different traffic light but were not related visually to the traffic light in question. It also shows how RA combined both the cues R and A to form a compound cue. Traffic light images taken from The Highway Code. Contains public sector information licensed under the Open Government Licence v3.0.

#### 2.4.1.3 Procedure

The task required participants to press a key or to withhold a response depending on the colour of a presented circle stimulus (see Figure 2.2 for a trial schematic). Each trial started with two coloured shape cues being presented, one above the other, for 250ms, on a 50% grey background. Participants were informed that these cues indicated that the trial was beginning but, in fact, some of them were stochastically predictive of whether or not a response was required. Throughout the task a white horizontal bar measuring 19mm by 4mm was displayed in the centre of the screen. Coloured shape cues measuring 19mm were presented in vertical alignment above and below and equidistant from this horizontal bar. The same seven coloured shape cues were used throughout the experiment and were randomly assigned to a cue for each participant. On single cue trials (e.g., G-) the cue appeared in both the top and bottom positions, while on compound trials (e.g., RA-) each cue was randomised to appear in either the top or bottom position. Following presentation of the cues on go trials, a 19mm diameter white circle appeared to the left or right of the central bar (separated by 22mm edge-to-edge). This

indicated to participants that they needed to make a spatially congruent response, e.g., a left side response ('x' on a standard QWERTY keyboard) if a left-side circle was displayed (right-hand circles will require a '>' key press). On no-go trials the circle displayed was coloured, informing participants that they needed to withhold their response. For both singleton and compound cues the white circle appeared equally often on either side of the screen as did the four coloured circles used as stop signals. The colour of each cue and stop signal was randomised for each participant and sampled from the HSB colour-space (Joblove & Greenberg, 1978) by selecting equally spaced hues whilst constraining saturation (75-100%) and brightness (50-100%). The colour of the stop signals differed from those used in the cues.




Figure 2.2. Schematic of a single cue stop trial for Experiment 1.

Cue J was used for tracking purposes only, i.e. performance to this cue was used to set the response window duration. This tracking procedure for the task applied to both go and no-go trials involving the cue J and was a 3-down/1-up procedure (for similar procedures see Leiva et al., 2015; Elchlepp & Verbruggen, 2017), so that for every three correct trials the maximum response window (from signal onset, i.e. when a circle appeared) shortened by 50ms, whilst an error resulted in the window being increased by 50ms. The window started at 750ms and the calibration phase helped determine a reasonable starting window for each participant, with the maximum response window

capped at 2000ms, and the minimum 100ms. The tracking was applied to all blocks of the experiment. The idea was to ensure speeded responses throughout the experiment, and prevent people being too cautious because of the possibility of a no-go signal appearing.

Participants received on-screen feedback to errors. For commission errors, regardless of congruency of response or incorrect keypress, the feedback was 'No response required!'. For omission errors participants received feedback of 'You should have responded'. On go trials participants received feedback on incorrect key presses ("Incorrect key pressed, use X or .>') and wrong direction key presses ("Press the key that matches the side the white circle appears on'). All feedback was displayed for 500ms and accompanied with a 400Hz tone for 150ms delivered through closed headphones. Participants received no feedback at the end of each block. There was a variable inter-trial interval of 250ms to 500ms, throughout which the white bar remained on screen. As in Bowditch (2016) the experiment was designed so that participants' focus was on stopping. This was achieved in two ways. Firstly, coloured circles were stop trials, with white (in some way the default colour) being go. Secondly, the instructions promoted going as the default by mentioning it first. These manipulations were designed to encourage people to be looking for coloured circles as a signal to stop, that is, to employ a stop task set (this is further discussed in Chapter 4).

## **2.4.2 Analysis**

Data was processed and analysed using R v.3.5.1 (R Core Team, 2018). Due to the need for participants to respond at least once per cue to obtain reaction time measurements, data was averaged for each cue by each block (data from the calibration block was not analysed) with reaction times on error trials being excluded. To prevent excessive data loss, trials immediately following an error trial were retained. As the focus of the experiment was on performance to each cue, rather than the development of learning over training, training data was further summarised into grand means per cue per participant. For test data, only the first test block was analysed following Bowditch et al.'s (2016) observation that extinction can be a problem if testing is overly prolonged.

To take into account individual differences in reaction times and error rates, linear mixed-effects models were developed using lme4 (Bates, Mächler,

Bolker, & Walker, 2015). I used the Akaike Information Criterion Corrected (AICc; Akaike, 1974; Burnham & Anderson, 2002) to compare the models. The corrected criterion was used due to small sample size and so prevents overfitting of models. Model selection was undertaken by comparing the difference between AIC score of model *i* against the best model (the model with the lowest absolute AIC score), with a difference score greater than 2 indicating that difference was meaningful, and that the best model was the most appropriate (Burnham & Anderson, 2004).

Homoscedasticity and normality of the residuals were explored using a graphical approach. Contrasts were corrected using a Bonferroni procedure (where the standard alpha level of .05 was divided by the number of test run; see Frane, 2015). Contrasts were undertaken by changing the baseline category of the chosen models. Using mixed-effects models, as opposed to standard pairwise comparisons, allows for data from all groups to be pooled and used to estimate the global variance, this can lead to be more powerful tests (reflected in the higher degrees of freedom; Harrison et al., 2018). To confirm that variances between cues were similar, and thus pooling variances was justified, I used the (conservative) criterion outlined in Fox (2008), where a ratio of less than 1:4 between the largest and smallest variance indicates comparable variance. In instances where the variance ratio for a measure was greater than 1:4 standard *t*-tests were conducted. It should be noted that due to the use of inverse gamma models for the analysis of commission errors the variances for these models are in fact the negative reciprocal of the original values ( $-1/\mu$ ) rather than the data presented in the descriptive statistics tables.

Confidence intervals were calculated using the Wald method (for details see Pek & Wu, 2015). Conditional  $R^2$  values were estimated using the Nakagawa approach (Nakagawa & Schielzeth, 2013; Nakagawa, Johnson, & Schielzeth, 2017). The significance of effects was assessed using through the lmerTest package (Kuznetsova, Brockhoff, & Christensen, 2016).

Outlier replacement was undertaken with a view to removing atypical responses while maximising power and efficient use of resources. To this end, the process for identifying outliers for reaction times and commission and omission errors was different. For commission and omission errors (which represent failures to perform the task) participants with errors greater than two whiskers (i.e.



1.5IQR\*2 = an IQR of 3) from the upper and lower quartiles were replaced. This process was undertaken until there were no outliers for these two measures. In Experiment 1 six participants were replaced, five for having high omission errors, and one for having high commission errors. For reaction times a more nuanced approach was undertaken, using a combination of a priori screening followed by model criticism of the mixed-effect models (see Baayen & Milin, 2010). This approach was used because long reaction times do not necessarily indicate lack of engagement with the task, or failure to understand the task, but could be due other factors such as age (Der & Deary, 2006). Furthermore, the experiments here are concerned with differences between cues rather than overall reaction times. Initially, reaction time data was screened for participants with response times greater than two whiskers (i.e. 1.5IQR\*2 = an IQR of 3) from the upper and lower quartiles with these participants being replaced. Box-and-whisker plots were then run on this new sample. Any participants that would have also been outliers in the original sample were replaced, while participants that were outliers in the new sample but would have not been in the original sample were retained. Next, I used the sigtest function in the influence.ME package (Nieuwenhuis, te Grotenhuis, & Pelzer, 2012) to assess if the presence of these 'sub' outliers influenced the results of the mixed effects models. In this case the single 'sub' outlier present in reaction time data did not significantly influence the results and so full models are reported for Experiment 1.

Due to non-normal data, commission errors were analysed using a Gamma model through the glmmTMB package (Brooks et al., 2017). These models were performed using the standard inverse links and model the negative reciprocal of the mean, i.e.  $-1/\mu$  (where  $\mu$  is the expected mean). As the data included 0, prior to analysis the data transformation  $(y * (n - 1) + 0.5)/n$  was applied, where  $y$  was commission errors and  $n$  the sample size (Smithson & Verkuilen, 2006; Cribari-Neto & Zeileis, 2010). Transformed data is reported throughout. For reaction times, Cohen's  $d$  was calculated using the lme.dscore function in the EMAtools package (Kleiman, 2017). Due to the models used it was not possible to calculate confidence intervals or effect sizes for commission data.

I first present a model comparing G vs. B performance to show that the task is experienced as expected (i.e. as a manipulation check), followed by the key 'traffic light' contrasts: A vs. an average of I and P (hence-forth referred to as I/P), R vs. I/P, and RA vs IP. Omission data is presented but not analysed as due to the low error rate any conclusions drawn are likely to be spurious. As the G vs. B contrast was the manipulation check, the standard alpha level was used. For the other contrasts, to control for multiple comparisons, the alpha level was corrected to 0.017 (for means see Table 2.2). As I averaged cue's I and P this meant that the models I ran included the variance of cue I, cue P and cue I/P. To correct for this, cues I and P were removed from the models and only I/P left in.

Phase	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Training						
A-/+	371.75	54.22	0.02	0.01	0.01	0.01
B-/+	372.24	51.66	0.02	0.01	0.01	0.01
G-	367.85	53.20	0.02	0.02	0.01	0.01
I-/+	372.75	51.17	0.02	0.01	0.01	0.00
IP-/+	376.50	54.66	0.02	0.01	0.01	0.01
J±	370.82	51.39	0.02	0.01	0.01	0.01
P-/+	375.23	55.96	0.01	0.01	0.01	0.01
R+	370.10	46.34	0.02	0.01	0.01	0.01
RA-	373.29	50.87	0.02	0.01	0.01	0.01
Test						
A	375.99	68.61	0.01	0.02	0.01	0.03
B	383.94	68.26	0.02	0.04	0.01	0.02
G	366.38	62.28	0.01	0.02	0.01	0.02
I	384.59	71.30	0.02	0.03	0.01	0.00
IP	390.36	79.90	0.02	0.03	0.01	0.02
J	374.19	54.71	0.02	0.03	0.01	0.00
P	383.25	88.06	0.01	0.02	0.01	0.02
R	375.36	57.20	0.02	0.05	0.01	0.02
RA	378.58	64.97	0.01	0.02	0.01	0.02

Table 2.2. Descriptive statistics for Experiment 1. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

#### 2.4.2.1 Reaction times

For the training data, the best fitting model (see Table 2.3) was a model that included the main effects of cue with random intercepts (conditional  $R^2 = 0.95$ ). The G vs. B contrast approached significance,  $t(343) = 1.91$ ,  $p = .057$ , 95% CI [-0.12, 8.90],  $d = 0.21$ . The difference between R vs. I/P was marginally significant at a standard alpha level,  $t(343) = 1.69$ ,  $p = .092$ , 95% CI [-0.62, 8.41],  $d = 0.18$ , hinting at a trend for R to have faster responses to than I/P, suggesting R was a go cue. The A vs. I/P contrast was not significant,  $t(343) = 0.97$ ,  $p = .330$ , 95% CI [-2.27, 6.76],  $d = 0.11$ . The contrast between RA and IP

was non-significant,  $t(343) = 1.39$ ,  $p = .164$ , 95% CI [-1.30, 7.72],  $d = 0.15$ . The results suggest that participants were not particularly successful at learning the contingencies related to the cues, but there is some evidence of learning as the G vs. B contrast would be significant using a one-tailed test.

At test, the best fitting model was the main effects of cue with random intercepts (conditional  $R^2 = 0.78$ ). The G vs. B contrast was now significant,  $t(343) = 2.77$ ,  $p = .006$ , 95% CI [5.14, 29.98],  $d = 0.30$ , with faster reaction times to G suggesting that by the end of training some learning of the contingencies had taken place. The A vs. I/P contrast was not significant,  $t(343) = 1.25$ ,  $p = .212$ , 95% CI [-4.49, 20.35],  $d = 0.14$ . The R vs. I/P contrast was also not significant,  $t(343) = 1.35$ ,  $p = .178$ , 95% CI [-3.86, 20.98],  $d = 0.15$ . The contrast between RA and IP was marginally significant at the standard alpha level,  $t(343) = 1.86$ ,  $p = .064$ , 95% CI [-0.64, 24.20],  $d = 0.20$ , with faster reaction times to RA than IP, suggesting that RA was a go cue at test.

#### **2.4.2.2 p(respond)**

Training data models were run with a Gamma family to better fit the shape of the data. The final model (see Table 2.3) was a model with a Gamma family and inverse link and included the main effects of cue with random intercepts. Currently, it is not possible to calculate  $R^2$  for such models. The G vs. B contrast was marginally significant at the standard alpha level,  $z = 1.69$ ,  $p = .091$ , suggesting that this difference was not learnt well during training. The A vs. I/P contrast was not significant,  $z = 0.37$ ,  $p = .714$ . The R vs. I/P contrast was also not significant,  $z = 0.12$ ,  $p = .901$ . The RA vs. IP contrast was not significant,  $z = 0.87$ ,  $p = .383$ . Overall, it seems that learning was weak for this measure at training.

At test, the G vs. B contrast was significant,  $z = -3.52$ ,  $p = < .001$ . However, the results were not in the expected direction with significantly more errors for cue B (a 50% go cue) than cue G (a 75% go cue). The A vs. I/P contrast was not significant at the reduced alpha,  $z = -2.21$ ,  $p = .027$ , though there was a trend for more errors for I/P than A, tentatively suggesting that A was something of a stop cue. The R vs. I/P contrast was also not significant at the corrected alpha level,  $z = 2.05$ ,  $p = .040$ , though there was a trend for more commission errors to R than I/P, suggesting R was a go cue. The RA vs. IP contrast was not significant,  $z = -1.02$ ,  $p = .309$ .

Model	Training	Test
Reaction time models		
Main effects of cue	4305.68	4505.84
Main effects of cue with random intercept	<b>3327.35</b>	<b>4037.82</b>
p(respond)		
Main effects of cue with Gamma family and inverse link	-2737.61	-2512.04
Main effects of cue with Gamma family and inverse link and random intercept	<b>-2769.68</b>	<b>-2631.96</b>

Table 2.3. AICc scores for models run for Experiment 1 on reaction time and p(respond) data at training and test. Bold are the models chosen. Note that for both DVs the models with main effects of cue with random intercept and random slope failed to converge for both test and training and so are not reported.

### 2.4.3 Summary

In summary, the reaction time data for Experiment 1 supports past work in demonstrating how pairing a stimulus to a go response (cue G) can lead to faster reaction times than to a stimulus associated with stopping (cue B), with significant differences at test for reaction times. However, for p(respond) at test the direction of the effect is unexpectedly reversed, with significantly more errors for B than G, suggesting that cue B primed a go response more than cue G for this measure. Given this, I cannot unequivocally say that I have demonstrated learning of the contingencies implemented in the design. This is an interesting finding in terms of the wider literature, yet, as discussed below, this finding is likely explained by reference to the experimental design as opposed to providing evidence against associatively-mediated learning in such tasks.

The learning to the cues A and R is not overly clear. The fact that there were no significant reaction time contrasts does indicate that learning was weak with

regards to response times, yet the marginally significant (at a standard alpha level) contrast for R vs. I/P does present a slight trend for participants expressing cue R was a go cue in training. For commission errors, only results at test suggested any (weak) learning to the cues, and at face value the marginally significant results suggest a weak trend for cue A to prime a stop response (having more errors than cue I/P), and for cue R to promote a go response (with R having more errors than I/P). On their own these results would be troubling, with people learning at an associative level that R is a go cue!

However, there are strong reasons to assume that there are issues with the experimental design which question the internal validity of the experiment. There is a large body of evidence suggesting that participants can learn incidentally about the contingencies linked to cues (Best et al., 2016; Bowditch et al., 2016). However, this does not seem to be the case here, with cue B (a 50% stop cue) having more errors at test than cue G (a 75% go cue). Appealing to Occam's razor it seems the experiment is not being experienced in the manner expected. One suggested reason behind these findings is that the design is flawed. While the design mimicked the contingencies of traffic lights, the control cues (B, I, IP, P) did not balance the contingencies implemented by the traffic light cues, and so the overall experimental design means that participants were more likely to 'go' than to 'stop'. While there were two go cues, there was only one outright stop cue with the rest being 50/50. This might have led people to feel the task was pushing them to respond, while the task outcome was in fact geared to stopping – this could explain the unusual results.

## **2.5 EXPERIMENT 2**

Experiment 1 was the first attempt, as far as I am aware, to investigate the learning that occurs incidentally at UK traffic lights. However, there are two issues to note. Firstly, as discussed above, there was the unbalanced design. In Experiment 2 a slight change to the contingencies associated with the control cues rectified this issue (see Table 2.4) so that the overall incidence of stopping was now 50%. The second issue to address is the assertion that participants are learning to stop with this being the outcome. This issue arises from the fact that when one describes a design as R+ and RA-, the designation of what is + and what is - is to some extent arbitrary. One may have a particular outcome in

mind that is labelled +, but is this the psychologically real outcome experienced by people doing the task? To investigate this, it is possible to make use of the feature-positive effect (Lotz, Uengoer, Koenig, Pearce, & Lachnit, 2012) to confirm that stopping was the outcome. This is the effect that learning is generally faster to excitatory than inhibitory cues. Thus, according to Rescorla and Wagner (1972), a discrimination between X and XY is easier to solve if XY denotes the presence of an outcome, i.e. X-, XY+; rather than if the compound denotes the absence of an outcome, X+, XY-. This is because the first, feature-positive discrimination, requires the simultaneous excitatory learning to Y and the extinction of X, whereas the second, feature-negative discrimination requires a participant to first learn that X is an excitatory cue before learning that Y is an inhibitor, which takes longer. The feature-positive effect is robust, being found in pigeons (Jenkins & Sainsbury, 1970), rats (Reberg & Leclerc, 1977), honey bees (Abramson et al., 2013) and humans (Newman, Wolff, & Hearst, 1980; Richardson & Massel, 1982). The effect can also be found when comparing a difference in magnitude of an outcome rather than its presence or absence (Todd, Winterbauer, & Bouton, 2010).

A key piece of evidence for the feature-positive effect in humans comes from Lotz et al. (2012). This was the first experimental evidence that the effect could be found in humans using simple discriminations (i.e. similar to the current procedure) rather than complex discriminations used previously (see Newman et al., 1980). The Lotz experiments entailed participants completing a predictive learning task in which they were shown a letter or pairs of letters and asked to respond if they thought a green circle would follow, or not respond if they thought the circle would not be presented. In support of the feature-positive effect, participants were able to learn that the presence of another letter (AB) compared to a single letter (A) predicted the occurrence of a green circle better than when the presence of another letter predicted the absence of the outcome.

Crucially, it is possible to use this effect to learn what the outcome of a task is. If the outcome of a task is stop (+), then the discrimination between cues that signal an outcome X- vs. XY+ (a feature-positive discrimination) will be acquired more readily than a discrimination which signals absence of an outcome C+ vs. CD- (a feature-negative discrimination). If, instead, the effective outcome is going, then the C vs. CD discrimination should be learnt faster than X vs. XY,

because the former is now the feature-positive discrimination, i.e. C-, CD+, where + now denotes go. Using this logic, Bowditch (2016) was able to provide good evidence that using similar procedures to the ones presented here, participants were in fact looking for occasions when they needed to stop, with this being the outcome. The technique used by Bowditch was to compare the rate of acquisition of the two discriminations, and then use this to diagnose which was the feature positive discrimination, and hence deduce what was the effective outcome. Bowditch based his claim that stop was the effective outcome on the fact that a discrimination, X-, XY+, where + denotes stopping, was learnt significantly faster than another discrimination, C+, CD-, which, with this notation, is the standard feature-positive effect.

Noting that in Experiment 1 + represented stop, the pair R+, RA- is a feature-negative pair (if one ignores A±), thus in order to replicate the analysis implemented by Bowditch to assay the task outcome a feature-positive discrimination pair is needed, therefore I±, IP± becomes I-, IP+, a feature-positive discrimination if the outcome is indeed stop. If stop is the outcome then this new pair should be learnt more readily than R+, RA-. To equate conditions more precisely, I also include P± to allow for A±. This helps the design to be more balanced and allows me to assay the effective outcome using these procedures.

Of course, the discriminations present in the current design vary somewhat from the feature-positive discrimination described in Lotz et al. (2012). In the Lotz experiments (as is typical in the feature-positive literature) the compound cue is never seen separately, that is, participants never experience each part of the compound cue independently. In contrast, in the design employed here, the A in RA is experienced as a separate cue (as is cue R). The discrimination therefore is not R- RA+ but rather R- A± RA+, and thus is not a true feature-positive discrimination. Of course, it may be that cue A± (which is a 50/50 cue) does not convey any response tendency and is therefore irrelevant to the discrimination. In which case the discriminations become more like the traditional feature-positive discrimination.

However, given that the experiments (through the instructions) assume a certain task set it could be that cues come to be learnt as holding percentage outcomes. By this I mean that, if the outcome is stop, cue A will not be seen as



a 50/50 cue, but rather as a 50% stop cue. In these circumstances, the discriminations can be solved without reference to the feature-positive effect. The expected difficulty of learning a 'feature-negative discrimination' could arise from the fact that it requires configuration. Focusing on the R- RA+ discrimination, it involves cue R (a 75% stop cue), A (a 50% stop cue), and RA (a 25% stop cue). Therefore, combining R and A (both cues with moderate stop tendencies) leads to a cue (RA) that has weaker stop tendencies. In this instance the logic of Rescorla and Wagner (1972) does not hold, as R and A cannot easily summate to a weaker contingency. Therefore, to solve this discrimination participants would need to form a configural unit of RA. Such configuration is resource intensive and thus necessarily slow (Sutherland & Rudy, 1989). Conversely, the feature-positive discrimination does not require configuration. This is because the two stop cues summate to a greater stop cue, with cue I (a 25% stop cue) and cue P (a 50% stop cue) summing to IP (a 75% stop cue). Therefore, the feature-positive discriminations can be learnt through the logic of Rescorla and Wagner (1972), a quicker process when compared to the 'feature-negative discrimination'.

Ultimately, the experiments were not designed to provide evidence for or against competing learning accounts. However, it is assumed that participants come to learn both excitatory and inhibitory responses in the experiments (see section 2.8.1 for further). In this case, 50% cues (e.g., cue A or cue P) will not be seen as cuing 50% of a specific outcome but will be seen as 50/50 neutral cues - being equally likely to cue either outcome. Such learning could result in the 50/50 cues being somewhat irrelevant to the discriminations and so the discriminations come to be experienced in the tradition of a feature-positive discrimination design. As such, language appropriate to the feature-positive model will be used throughout. Yet the discussion above should serve to highlight that other explanations could be generated to explain the witnessed learning.

## **2.5.1 Method**

### **2.5.1.1 Participants**

Fifty-five participants participated in exchange for payment of £5 or one course credit (see Results section for details on the outlier removal process). Of these

55 participants, 41 were female with an overall mean age of 19.84 (SD = 3.47). The inclusion criteria and outlier removal process were identical to those used in Experiment 1. A power analysis using SIMR (Green & MacLeod, 2016) indicated that a sample size of 55 would give sufficient power (96.80%) to detect a 18ms difference in reaction times between a go vs. stop cue in test block 1, where a priori one would expect to find an effect. Specifically, this test was run on G vs. B, see section 2.5.1.2 for details.

### 2.5.1.2 Design

A number of changes were made to the design (see Table 2.4 for summary of the experimental design), 1) B $\pm$  became B+ to reduce the tendency to go to this cue, and 2) I $\pm$ , IP $\pm$  became I-, IP+, and thus a feature-positive discrimination.

Phase	Blocks	Trials per block	N per type	Design	
Calibration	1	48	48		J $\pm$
Training	8	144	16	G-, R+, A $\pm$ , RA-, B+, I-, P $\pm$ , IP+	J $\pm$
Test	2	144	16	G, R, A, RA, B, I, P, IP	J $\pm$

Table 2.4. Summary of Experiment 2 design. Letters represent coloured cues. Trials were go 75% of the time (-), stop 75% of the time (+), or stop 50% of the time ( $\pm$ ). At test all trials were 50% stop and so the cues were now non-predictive.

### 2.5.1.3 Procedure

The procedure was identical to Experiment 1.

## 2.5.2 Analysis

The data was processed and analysed as described in Experiment 1. Six participants were replaced for having high mean omission errors. There were no outliers for commission errors. As with Experiment 1 there was one 'sub' outlier in the reaction time data, but this did not significantly affect the results and so the participant was retained. Given that control cues were designed to balance

out the overall design rather than act as baseline cues the previous analyses are no longer appropriate. Therefore, a different set of contrasts were conducted. I still conducted the G vs. B contrast (which is still the manipulation check, but I then performed contrasts aimed at understanding to what extent cues A and R primed going or stopping. Thus, the key 'traffic light' contrasts are: A vs. B, A vs. G, A vs. R, R vs. B and R vs. G. These will be presented in text if significant or informative. Once again, omission data is presented but not analysed due to the low error rate. This time the standard alpha level was used for the G vs. B contrast and a corrected alpha level of .010 was used for all other 'traffic light' contrasts (for means see Table 2.5). To analyse the feature positive effect, paired *t*-tests comparing the differences between the feature-positive and feature-negative contrasts were also conducted. In this instance due to the assumed outcome being stop, R+ RA- is the feature negative contrast, with I- IP+ being the feature positive contrast. As these tests were also, in effect, a manipulation check, a standard alpha level was used. It should also be noted that *t*-tests rather than multilevel modelling were used for these analyses to enable more direct comparison to the earlier work by Bowditch (2016).

Phase	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Training						
A-/+	385.95	42.79	0.01	0.01	0.01	0.01
B+	392.57	43.26	0.01	0.01	0.01	0.01
G-	378.19	38.38	0.02	0.02	0.01	0.00
I-	383.47	41.52	0.02	0.02	0.01	0.01
IP+	394.27	48.1	0.01	0.01	0.01	0.01
J±	385.01	40.00	0.01	0.01	0.01	0.01
P-/+	389.99	40.28	0.01	0.01	0.01	0.01
R+	389.77	42.94	0.01	0.01	0.01	0.01
RA-	390.91	39.9	0.02	0.02	0.01	0.00
Test						
A	392.47	70.06	0.02	0.03	0.01	0.02
B	412.14	65.79	0.01	0.02	0.01	0.00
G	394.40	61.35	0.01	0.02	0.01	0.02
I	382.25	55.88	0.01	0.02	0.01	0.00
IP	401.15	54.7	0.02	0.04	0.01	0.00
J	400.04	59.46	0.01	0.02	0.01	0.02
P	404.69	73.9	0.02	0.03	0.01	0.00
R	401.45	62.91	0.02	0.04	0.02	0.03
RA	400.45	54.4	0.01	0.02	0.01	0.02

Table 2.5. Descriptive statistics for Experiment 2. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 2.5.2.1 Reaction times

For training data, the best fitting model (see Table 2.6) was one that included the main effects of cue with random intercepts (conditional  $R^2 = 0.91$ ). The G vs. B contrast was highly significant,  $t(432) = 5.86$ ,  $p = < .001$ , 95% CI [9.57, 19.19],  $d = 0.56$ , with G being faster than B, thus confirming that participants were learning the contingencies present in the design. In terms of the experimental contrasts, A vs. B was significant,  $t(432) = 2.70$ ,  $p = .007$ , 95% CI [1.81, 11.43],  $d = 0.26$ , with faster responses to cue A. The contrast A vs. G was also significant,  $t(432) = -3.16$ ,  $p = .002$ , 95% CI [-12.57, -2.95],  $d = -0.30$ ,

with slower responses to cue A. These results suggest that while cue A was more of a go cue than B, it was not such a strong go cue as G and thus neutral overall. The contrast for A vs. R was non-significant,  $t(432) = 1.56$ ,  $p = .120$ , 95% CI [-0.99, 8.64],  $d = 0.15$ . In terms of the R contrasts, the R vs. B contrast was not significant,  $t(432) = 1.14$ ,  $p = .255$ , 95% CI [-2.01, 7.61],  $d = 0.11$ , though the R vs. G contrast was significant,  $t(432) = -4.72$ ,  $p = < .001$ , 95% CI [-16.39, -6.77],  $d = -0.45$ , with responding in the presence of R being slower. This pattern of results suggests that cue R was overall a stop cue, or at least certainly not a go cue. There was a significant positive difference between the differences of IP+ vs. I- ( $M = 10.80$ ,  $SD = 19.65$ ) and R+ vs. RA- ( $M = -1.14$ ,  $SD = 18.20$ ),  $t(54) = 3.23$ ,  $p = .002$ , 95% CI [4.54, 19.35],  $d = 0.44$ , confirming that the feature-positive discrimination was easier to acquire than the feature-negative discrimination, if the effective outcome of the task is taken to be stop. To put this another way, the significant feature-positive effect in the data confirms that the effective outcome was stop.

At test, the best model was one that included the main effects of cue with random intercepts (conditional  $R^2 = 0.66$ ). The G vs. B contrast was again significant,  $t(432) = 2.54$ ,  $p = .011$ , 95% CI [4.05, 31.45],  $d = 0.24$ , confirming that the participants had learnt about the contingencies in the experiment. In terms of the experimental contrasts, A vs. B was significant,  $t(432) = 2.81$ ,  $p = .005$ , 95% CI [5.98, 33.38],  $d = 0.27$ , with faster responses in the presence of cue A indicating that cue A primed a go response. The A vs. G contrast was not significant,  $t(432) = 0.28$ ,  $p = .783$ , 95% CI [-11.77, 15.63],  $d = 0.03$ . The A vs. R contrast was non-significant,  $t(432) = 1.29$ ,  $p = .200$ , 95% CI [-4.72, 22.69],  $d = 0.12$ . For the contrasts against R, R vs. B was not significant,  $t(432) = 1.53$ ,  $p = .127$ , 95% CI [-3.01, 24.39],  $d = 0.15$ . The R vs. G contrast was also non-significant,  $t(432) = -1.01$ ,  $p = .313$ , 95% CI [-20.76, 6.64],  $d = -0.10$ . There was a marginally significant difference (at a standard alpha level) between the differences of IP+ vs. I- ( $M = 18.90$ ,  $SD = 41.13$ ) and R+ vs. RA- ( $M = 1.01$ ,  $SD = 57.59$ ),  $t(54) = 1.78$ ,  $p = .081$ , 95% CI [-2.28, 38.08],  $d = 0.24$ , reinforcing the view that the feature-positive discrimination was easier to acquire than the feature-negative discrimination.

Model	Training	Test
Reaction time models		
Main effects of cue	5116.26	5508.14
Main effects of cue with random intercept	<b>4163.31</b>	<b>5094.88</b>
p(respond)		
Main effects of cue with Gamma family and inverse link	-3476.56	-3293.62
Main effects of cue with Gamma family and inverse link and random intercept	<b>-3545.54</b>	<b>-3447.99</b>

Table 2.6. AICc scores for models run for Experiment 2 on reaction time and p(respond) data at training and test. Bold are the models chosen. Note that for both DVs the models with main effects of cue with random intercept and random slope failed to converge for both test and training and so are not reported.

### 2.5.2.2 p(respond)

For commission data I used Gamma models to analyse the results (see Table 2.6 for best fitting model). The G vs. B contrast was significant,  $z = 3.88$ ,  $p < .001$ , with more errors for cue G, thus confirming that participants were learning the contingencies present in the design (as they were more likely to go if G was presented). The A vs. B contrast was not significant,  $z = 1.37$ ,  $p = .166$ , nor was the A vs. R contrast,  $z = 1.24$ ,  $p = .215$ . However, the A vs. G contrast was significant,  $z = -2.60$ ,  $p = .009$ , with more errors for cue G, suggesting that A was not a strong go cue. Focusing on the R contrast, R vs. B was not significant,  $z = 0.15$ ,  $p = .881$ , yet the R vs. G contrast was,  $z = -3.75$ ,  $p < .001$ , with more errors for G than R, suggesting that R was not a go cue and overall seems to be rather like B on this measure. There was no significant difference between the differences of IP+ vs. I- ( $M = -0.003$ ,  $SD = 0.02$ ) and R+ vs. RA- ( $-0.004$ ,  $SD = 0.02$ ),  $t(54) = 0.29$ ,  $p = .773$ , 95% CI [-0.01, 0.01],  $d = 0.04$ , and the direction of the effect was not in line with the hypothesis.

However, the difference was not significant, and the effect size was small. For clarity it should be noted that for  $p(\text{respond})$  the contrast IP+ vs. I- is expected to be negative (i.e. more errors for I than IP) as is R+ vs. RA- (i.e. more errors for RA than R). Thus, if the difference, in terms of its absolute size, is bigger for IP+ vs. I-, then the  $t$  statistic will be negative. However, it was positive which means that the difference between R+ vs. RA was the bigger.

For test, the G vs. B contrast was not significant,  $z = 1.34$ ,  $p = .180$ . The A vs. B contrast was significant,  $z = 3.18$ ,  $p = .002$ , with more errors for A than B, indicating that compared to B, cue A was a go cue. The A vs. G contrast was not significant at the reduced alpha,  $z = 1.97$ ,  $p = .049$ , although there is a trend for more errors to cue A than G, suggesting that cue A was more of a go cue than G. The A vs. R contrast was not significant,  $z = -0.76$ ,  $p = .446$ . Focusing on cue R, the R vs. B contrast was significant,  $z = 3.78$ ,  $p < .001$ , with more errors for R than B, suggesting that cue R was experienced as more of a go cue than B. The R vs. G contrast was also significant,  $z = 2.65$ ,  $p = .008$ , with, unexpectedly, more errors for R than G. This suggests that cue R was a strong go cue, more so than the 75% go cue. There was no significant difference between the differences of IP+ vs. I- ( $M = 0.004$ ,  $SD = 0.04$ ) and R+ vs. RA- ( $0.007$ ,  $SD = 0.04$ ),  $t(54) = -0.28$ ,  $p = .784$ , 95% CI  $[-0.02, 0.01]$ ,  $d = -0.04$ , with I- and R+ showing greater learning in their pairs.

### 2.5.3 Summary

Experiment 2 found compelling evidence that participants were learning that G was a go cue and cue B a stop cue, giving confidence that participants were experiencing and learning from the incidental go/no-go task as expected. This contrasts with the mixed results in Experiment 1 and suggests that those findings were due to that particular experimental design.

In terms of the learning of the traffic light contingencies, there is some evidence that cue A was priming a weak go response. During training for reaction times, it was significantly slower than G, but significantly faster than B. Regarding commission errors during training, there was a significant difference between A and G (with more errors for G) and a non-significant difference between cue A and B (though A had numerically more errors). At test, cue A had significantly faster responses than cue B, and there were also significantly more errors for A than for B, which is consistent with a tendency to want to respond to A rather

than withhold a response (indeed, the contrast for A vs. G was not significant at the correct alpha for each measure at test, with cue A having numerically faster reaction times and more errors than G). So, in training cue A seems to be less go than cue G, but certainly not stop, while at test cue A is numerically more go than G. Looking at cue R, at training it seemed rather like cue B. The cue was significantly slower than G (like B), and not significantly faster to B, though numerically cue B seemed to prime more of a stop response. There were significantly fewer errors for R than for G, and the mean errors were identical for R and B. At test, it did not differ from B (or G this time) in terms of response times, with a mean reaction time midway between the two, but there were significantly more errors made to R than to both B and G. There is evidence that cue R was promoting a going response to some extent in the  $p(\text{respond})$  data for test, but otherwise it is best described as a weak stop cue. The contrasts between cues A and R were never significant, and that though cue A seemed to prime a go response and cue R stopping, it should be noted that going and stopping are relative terms.

The results regarding cue R do suggest that in a stop task set red traffic lights only become associated with a weak stop response. Considering that the task has stopping as the effective outcome, this result is rather surprising. One might think that when the default is to go, and one is looking out for a stop signal this is when learning about red will be optimal, but the evidence suggests that this is not necessarily the case. This is certainly an avenue worth further exploration as it indicates that the contingences of UK traffic lights prevent strong learning of stop cues, at least in a stop task set.

The significant feature-positive effect found in the reaction time training data (and the marginally significant result at test for reaction times) supports the idea that the effective outcome is stop as a result of the manipulation of task set. Following Bowditch (2016), it is believed that the task instructions and the use of a number of differing coloured stop signals are what promote this task set, and result in the feature-positive discrimination based on the outcome being stop being learnt more easily than the feature negative. The obvious test of this proposition is to change task set by changing these parameters, and that is one purpose of the next experiment. In both Experiments 1 and 2 stopping was the designated task outcome as this was felt to be the more plausible scenario in



modelling behaviour approaching traffic lights. However, there is a case to be made that the other task set which has going as the effective outcome also plays a role in driving behaviour. If the driver is stationary at the lights, then they will be looking for a go signal. Accordingly, this situation was explored in Experiment 3.

## **2.6 EXPERIMENT 3**

The purpose of this experiment was firstly to demonstrate that it was possible to manipulate the effective outcome during training, (i.e. go vs. stop), thus confirming that the technique used to achieve this was valid. Secondly, to see what effect this change of task set, to one where the effective outcome was going, had on learning based on experience. In order to change the focus of participants from a stop to a go task set two changes were made following Bowditch (2016):

- 1) the go/no-go signals were reversed, with coloured circles now being go signals and white circles being stop signals, and
- 2) the order of the instructions was changed so that stopping was mentioned first (making it the 'default' behaviour).

The net effect of these changes should be to change the participant's task set from learning when to stop, to learning when to respond.

### **2.6.1 Method**

#### **2.6.1.1 Participants**

The sample size, payment, and inclusion criteria and outlier removal process were the same as Experiment 2 (see Results section for details of those removed). Of the final sample, 38 were female, with an overall mean age of 21.29 (SD = 6.38).

#### **2.6.1.2 Design**

The design was similar to Experiment 2. However, as going was expected to be the effective outcome, + was now go rather than stop, and this is reflected in the summary of the design in Table 2.7.

Phase	Blocks	Trials per block	N per type	Design	
Calibration	1	48	48		J±
Training	8	144	16	G+, R-, A±, RA+, B-, I+, P±, IP-	J±
Test	2	144	16	G, R, A, RA, B, I, P, IP	J±

Table 2.7. Summary of Experiment 3 design. Letters represent coloured cues. Trials were go 75% of the time (+), stop 75% of the time (-), or stop 50% of the time (±). At test all trials were 50% stop and so the cues were now non-predictive.

### 2.6.1.3 Procedure

The procedure was identical to that of Experiment 2, except for the two changes made to shift the demands of the task from looking to stop to looking to respond. Coloured circles now required a go response and white circles required a stop response – so that the singleton colour was now a stop response. The order of instructions was changed so that stopping was mentioned first (making it the default) and responding second.

### 2.6.2 Analysis

The analysis followed the same approach as in Experiment 2. For Experiment 3, four participants were replaced: three for having high mean omission errors and one for withdrawing from the experiment. There were no outliers for commission errors or reaction times. As in Experiment 2 paired *t*-tests were run to analyse the feature positive effect. In Experiment 3 due to the assumed outcome being go, R- RA+ is the feature positive contrast, with I+ IP- being the feature negative contrast (for means see Table 2.8). The standard alpha level was used for the G vs. B contrasts and for the feature-positive contrasts, with a corrected alpha level of .010 being applied to all other contrasts.

Phase	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Training						
A-/+	405.61	39.87	0.02	0.01	0.01	0.01
B-	411.47	39.56	0.02	0.01	0.01	0.01
G+	397.06	36.40	0.02	0.02	0.01	0.01
I+	403.00	39.22	0.02	0.02	0.01	0.01
IP-	408.04	42.49	0.02	0.02	0.01	0.01
J±	400.53	37.29	0.02	0.02	0.01	0.01
P-/+	407.73	38.44	0.02	0.02	0.01	0.01
R-	408.64	42.68	0.02	0.01	0.01	0.01
RA+	401.95	40.11	0.02	0.02	0.01	0.01
Test						
A	408.79	58.66	0.02	0.04	0.01	0.02
B	414.91	74.26	0.02	0.03	0.01	0.00
G	409.33	70.58	0.02	0.04	0.01	0.02
I	399.43	55.56	0.01	0.02	0.02	0.04
IP	412.37	69.64	0.02	0.03	0.02	0.03
J	404.22	54.51	0.01	0.02	0.01	0.03
P	401.03	58.81	0.02	0.04	0.02	0.03
R	407.09	60.84	0.02	0.04	0.01	0.00
RA	407.55	63.62	0.03	0.05	0.01	0.02

Table 2.8. Descriptive statistics for Experiment 3. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 2.6.2.1 Reaction times

For the training data, the best fitting model (see Table 2.9) had a conditional  $R^2$  of 0.88. The G vs. B contrast was highly significant,  $t(432) = 5.44$ ,  $p < .001$ , 95% CI [9.22, 19.59],  $d = 0.52$ , with faster responses for cue G, thus confirming that participants were learning the contingencies present in the design. In terms of the experimental contrasts, A vs. B was not significant at the reduced alpha,  $t(432) = 2.21$ ,  $p = .027$ , 95% CI [0.67, 11.05],  $d = 0.21$ , but hints at participants experiencing cue A as distinct from cue B, i.e. not a strong stop cue. The contrast A vs. G was significant,  $t(432) = -3.23$ ,  $p = .001$ , 95% CI [-13.73, -3.36],

$d = -0.31$ , with faster responses for cue G, indicating that cue A was not experienced as a strong go cue. The A vs. R contrast was not significant,  $t(432) = 1.15$ ,  $p = .218$ , 95% CI [-2.15, 8.22],  $d = 0.11$ . In terms of the R contrasts, the R vs. B contrast was not significant,  $t(432) = 1.07$ ,  $p = .286$ , 95% CI [-2.36, 8.01],  $d = 0.10$ , and the R vs. G contrast was significant,  $t(432) = -4.38$ ,  $p < .001$ , 95% CI [-16.77, -6.40],  $d = -0.42$ . This indicates that cue R was not seen as a go cue, with slower responses than cue G (a 75% cue go). The difference between R- vs. RA+ ( $M = 6.70$ ,  $SD = 24.37$ ) and IP- vs. I+ ( $M = 5.04$ ,  $SD = 21.39$ ) was non-significant (but numerically in the right direction for a feature-positive advantage),  $t(54) = 0.38$ ,  $p = .707$ , 95% CI [-7.16, 10.47],  $d = 0.05$ .

At test, the best model was one that included the main effects of cue with random intercepts (conditional  $R^2 = 0.64$ ). The G vs. B contrast was not significant,  $t(432) = 0.77$ ,  $p = .439$ , 95% CI [-8.53, 19.68],  $d = 0.07$ . None of the other contrasts were significant either (see Appendix A for full results).

### 2.6.2.2 p(respond)

In the best fitting Gamma model (see Table 2.9), the G vs. B contrast was on the threshold of significance,  $z = 2.56$ ,  $p = .010$ , suggesting that learning was occurring as expected, with more errors for cue G than cue B. The A vs. G contrast was marginally significant at the standard alpha level,  $z = -1.85$ ,  $p = .064$ , and hints at a trend for more errors to G than A, suggesting that A was not a strong go cue and supports the findings in response times. The A vs. B contrast was not significant,  $z = 0.75$ ,  $p = .453$ , nor was the A vs. R contrast,  $z = 1.05$ ,  $p = .293$ . Focusing on the R cues, the R vs. B contrast was not significant,  $z = -0.30$ ,  $p = .763$ . Yet, the R vs. G contrast was significant,  $z = -2.84$ ,  $p = .004$ , with more errors for cue G than R indicating that cue R was not experienced as a go cue. The difference between R- vs. RA+ ( $M = -0.01$ ,  $SD = 0.02$ ) and IP- vs. I+ ( $M = 0.0004$ ,  $SD = 0.02$ ) was significant,  $t(54) = -2.08$ ,  $p = .043$ , 95% CI [-0.01, -0.0002],  $d = -0.28$ , demonstrating that participants were able to learn the feature-positive contrast more readily than the feature-negative, consistent with the effective outcome being go. For clarity, it should be noted that the negative mean score for the R- vs. RA contrast indicates more commission errors for RA+ than R, which is in the right direction to learn this discrimination. Thus, the negative  $t$  statistics means that the feature-positive discrimination was better learnt than the feature-negative.

At test the G vs. B contrast was not significant,  $z = 0.00$ ,  $p = 1.00$ . The difference between R- vs. RA+ and IP- vs. I+ was also non-significant,  $t(54) = -1.15$ ,  $p = .255$ , 95% CI [-0.03, -0.01],  $d = -0.16$ , supporting the notion that learning at test was weak. The rest of the contrasts were also not significant (see Appendix B for full results).

Model	Training	Test
Reaction time models		
Main effects of cue	5058.42	5522.40
Main effects of cue with random intercept	<b>4219.96</b>	<b>5120.98</b>
p(respond)		
Main effects of cue with Gamma family and inverse link	-3056.96	-3002.98
Main effects of cue with Gamma family and inverse link and random intercept	<b>-3201.52</b>	<b>-3187.23</b>

Table 2.9. AICc scores for models run for Experiment 3 on reaction time and p(respond) data at training and test. Bold are the models chosen. Note that for both DVs the models with main effects of cue with random intercept and random slope failed to converge.

### 2.6.3 Summary

Experiment 3 aimed to investigate the learning of traffic light contingences when go was the effective task set. The changes made between Experiments 2 and 3 appear to have affected the task set, with the feature-positive effect for commission errors during training providing some evidence that ‘go’ was now the task outcome. There was also good evidence that participants learnt that cue G was a go cue and cue B a stop cue during training, indicating that the participants were still learning about the task as expected, despite the changes made. There can be little doubt though that learning was somewhat weaker in this experiment, and hence the generally non-significant results at test.

The change in task set has affected the learning of traffic light contingencies. While Experiment 2 suggested that A primed a weak go cue, the results for Experiment 3 are somewhat different. Here A has significantly slower reaction times against G in the training phase, and is numerically closer to B, but the contrast does indicate it is not the same as cue B (being significantly faster at a standard alpha level). On test, none of the contrasts are significant, though numerically cue G and A are similar. Looking at errors, during training numerically there is a tendency to commit more errors for G than A (with the difference being marginally significant at a standard alpha level), whilst the difference between A and B is small. At test there is no significant difference, yet numerically cue A does have more errors than either B or G. The overall impression is that A has become a bit less like G and a bit more like B and was experienced as a neutral to weak stop cue in Experiment 3. For cue R, as in Experiment 2, during training it seems to be a stop cue, being significantly slower to G and numerically like B on both measures, in fact error rates were lowest in R than B. However, at test performance seems to have collapsed for all cues, with reaction times and commission error rates being similar for cues G, B and R. Indeed, it is notable that in contrast to Experiment 2 there is not the large difference in  $p(\text{respond})$  between cue R and cue G or B. Overall, the results suggest that R seems to be a stop cue.

## **2.7 JOINT ANALYSIS OF EXPERIMENT 2 AND EXPERIMENT 3**

As Experiments 2 and 3 are opposites (in the sense they use the same cue contingencies but with opposite outcomes) it is possible to combine the two experiments and undertake a between-participants analysis to investigate changes across the studies. First, I present an analysis that demonstrates that the changes between Experiment 2 and 3 were successful in swapping round the feature-positive effect between the two experiments. This analysis was conducted on training data where this effect should be most obvious and is where Bowditch (2016) observed the effect.

The second analysis is an attempt to address an issue in the design of Experiments 2 and 3 in that while they investigated a single task set, traffic lights are likely to be experienced in different task sets. For example, a red traffic light is clearly not going to lead to the same effective outcome as a green

traffic light. Combining the two experiments allowed for the effect of both task sets together to be explored. Though it should be noted that the effects in Experiment 2 are bigger and this task set is likely to influence results to a greater extent than the go task set in Experiment 3. The analysis was undertaken on the test data as this is where learning about the cues is likely to be clearest, i.e. following acquisition, and therefore gives an indication of what learning has developed. Of course, this analysis is imperfect as the task set changes between, rather than within, traffic lights. By this I mean that while the joint analysis allows me to see what sort of behaviour cue G (the cue representative of a green traffic light) might promote in a combined task set situation, green lights are likely to be mostly experienced in a go task set (where the effective outcome is stop). This line of reasoning is continued and developed further in Experiment 6 in Chapter 4. The standard alpha level was used for the G vs. B contrasts (the go vs. stop manipulation check) and for the feature-positive contrasts, with a corrected alpha level of .010 being applied to all other contrasts.

### **2.7.1 Joint feature-positive analysis**

An analysis comparing Experiments 2 and 3 found that the outcome of the tasks had been successfully manipulated as far as the training reaction times were concerned, with the feature-positive effect swapping round from Experiment 2 to Experiment 3. There was a significant difference for the 'difference between the differences', that is, taking the difference for (IP-I) and (R-RA) for each experiment and comparing these differences,  $t(108) = 2.37$ ,  $p = .019$ , 95% CI [2.22, 24.99],  $d = 0.45$ . This demonstrates that the changes made between the two experiments successfully changed the nature of the discriminations experienced by participants, such that the effective outcome changed from stopping to going. Though the effect was stronger in Experiment 2 than 3: Experiment 2 difference is  $M = 11.95$ ,  $SD = 27.41$ , Experiment 3 difference is  $M = -1.66$ ,  $SD = 32.61$ . The same analysis for training commission errors was non-significant,  $t(108) = -1.06$ ,  $p = .291$ , 95% CI [-0.02, 0.01],  $d = -0.20$ .

### **2.7.2 Joint traffic light cues analysis**

Adding 'experiment' as a between participant factor to the models allowed me to investigate overall learning for the key contrasts in the test phase. Otherwise analysis was conducted as described in Experiment 2 (see Table 2.10 for AICs

and Table 2.11 for descriptive statistics). For reaction time data, while the interaction with random intercept model had the lowest AICc (see Table 2.10), these interaction effects merely reflect the results found individually for the two experiments and do not represent overall effects. Therefore, the analysis reported is the model with the main effects of cue with random intercepts (conditional  $R^2 = 0.65$ ). In this model there was a main effect of G vs. B across the experiments,  $t(872) = 2.32$ ,  $p = .020$ , 95% CI [1.83, 21.49],  $d = 0.15$ , with faster response times for G compared to B. This indicates that participants were able to learn the contingencies present in the designs. There was a just significant (right on the threshold) difference between A and B,  $t(872) = 2.57$ ,  $p = .010$ , 95% CI [3.07, 22.73],  $d = 0.17$ , with faster responses for cue A, indicating that cue A was not as much of a stop cue overall as B. The A vs G contrast was not significant,  $t(872) = 0.25$ ,  $p = .805$ , 95% CI [-8.60, 11.07],  $d = 0.02$ , nor was the A vs R contrast,  $t(872) = 0.73$ ,  $p = .468$ , 95% CI [-6.19, 13.48],  $d = 0.05$ . In terms of R, the contrast against B was marginally significant at the standard alpha level,  $t(872) = 1.84$ ,  $p = .065$ , 95% CI [-0.58, 19.09],  $d = 0.12$ , suggesting a weak trend for faster responses to cue R than B. The contrast against G was non-significant,  $t(872) = -0.48$ ,  $p = .631$ , 95% CI [-12.24, 7.42],  $d = -0.03$ .

For p(respond), as with reaction times, the interaction with random intercept model had the lowest AICc (see Table 2.10), yet for the reasons noted above I used a main effect of cue with random intercepts model. The G vs. B contrast was not significant,  $z = 0.68$ ,  $p = .499$ . The A vs. B contrast was significant,  $z = 2.74$ ,  $p = .006$ , with more errors for A than for B, indicating that A was not a strong stop cue. The A vs. G contrast was not significant at the reduced alpha level,  $z = 2.11$ ,  $p = .035$ , with more errors for A than G tentatively suggesting A was a go cue. The A vs. R contrast was not significant,  $z = 0.00$ ,  $p = 1.00$ . The contrast for R vs. B was significant,  $z = 2.74$ ,  $p = .006$ , with more errors for R than B, demonstrating that R was not a strong stop cue. The R vs. G contrast was significant at a standard alpha level,  $z = 2.11$ ,  $p = .035$ , with more errors for R than G. The results indicate that cues B and G primed similar responses, while cues R and A primed responses that were more error prone.



Model	Training
Reaction time models	
Main effects of cue and experiment	11014.80
Interaction effects of cue and experiment	11028.52
Main effects of cue and experiment with random intercept	<b>10251.55</b>
Interaction of cue and experiment with random intercept	10212.01
p(respond)	
Main effects of cue with Gamma family and inverse link	- 6229.28
Interaction effects of cue with Gamma family and inverse link	-6281.57
Main effects of cue with Gamma family and inverse link and random intercept	<b>-6581.25</b>
Interaction effects of cue with Gamma family and inverse link and random intercept	-6628.49

Table 2.10. AICc scores for models run for joint analysis on test reaction time and p(respond) data. Bold indicates the chosen model.

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
A	400.63	64.84	0.02	0.04	0.01	0.02
B	413.53	69.84	0.01	0.03	0.01	0.00
G	401.86	66.25	0.02	0.03	0.01	0.02
I	390.84	56.13	0.01	0.02	0.01	0.03
IP	406.76	62.59	0.02	0.03	0.01	0.02
J	402.13	56.82	0.01	0.02	0.01	0.03
P	402.86	66.50	0.02	0.04	0.01	0.02
R	404.27	61.66	0.02	0.04	0.01	0.02
RA	404.00	59.03	0.02	0.04	0.01	0.02

Table 2.11. Descriptive statistics for the joint analysis. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 2.7.3 Summary

The joint analysis confirms that participants were learning about the task in the expected manner with G having significantly faster response times than B. The analysis also clearly showed the change in relative acquisition of the R vs. RA and I vs. IP discriminations across the two experiments at training for reaction times, indicating that the task set manipulations were successful. This demonstrates that the changes in the experimental design changed the effective outcome (an issue I will return to in Chapter 4).

In terms of the traffic light contingencies, the just significant difference for A vs. B for response time indicates that cue A (a 50% cue) primed a significantly faster go response than a 75% stop cue (cue B), suggesting that this was not a stop cue. The reaction time data suggests that cue A primed a similar response as cue G. Indeed, the p(respond) data indicates that it was more of a go cue than cue G, though the difference was only significant at a standard alpha level. This leads to the conclusion that cue A is a go cue. For cue R, the evidence suggests overall it was a weak go cue. The cue was marginally significant at a standard alpha level to cue B (with faster response to R), and while not significantly slower than cues G and A did have numerically slightly slower response times. For error rates, Cue R had significantly more errors than cue B, and significantly more errors at a standard alpha level than cue G. The fact that

R is not a stop cue is surprising. It may be the case that RA is such a strong go cue that some of the go association transfers to the singleton R cue. However, looking at the descriptive statistics, RA does not seem to be a strong go cue. This suggests a more complex interaction between cues A, R, and RA, one that results in both cue A and cue R becoming go.

## 2.8 GENERAL DISCUSSION

Experiment 1 aimed to provide initial evidence that the contingencies learnt under a stop task set at traffic lights might not be in line with the rules of the Highway Code. However, the experimental design meant it was not possible to reach conclusions about learning towards key cues. Fortunately, Experiment 2 was able to provide evidence of learning, finding that cue A was experienced as a weak go cue and cue R as a weak stop cue. Experiment 2 made use of the feature-positive effect to confirm that the design was indeed inducing a stop task set with the feature-positive contrast (I- vs. IP+) being more readily learnt than the feature-negative discrimination (R+ vs. RA-) during training. This suggests that participants were looking to successfully withhold rather than respond. Experiment 3 demonstrated how swapping the response signal colour and changing the instructions can affect task outcome, in this case from stop to go. There was support for a feature-positive effect in commission errors during training, providing limited evidence of the feature-positive contrast being better learnt than the feature-negative contrast, suggesting that participants were now learning to respond. Under a go task set it seems that cue A primed a neutral/weak stop cue, and cue R a stronger stop response. Joint feature-positive analysis of Experiment 2 and Experiment 3 found that the outcome of the tasks had been successfully manipulated as far as training reaction time measures were concerned. The analysis showed that the feature-positive effect swapped around from Experiment 2 to Experiment 3, with a significant difference between the differences capturing this effect in each experiment (i.e. (R-RA)-(IP-I)), though the magnitude of the effect was higher in Experiment 2 than in Experiment 3.

In terms of the traffic light contrasts, the joint analysis suggests that, in a situation when both go and stops task sets are in play, cue A promotes going (faster and more errors than B), while cue R is best described as a weak go

cue, though it should be noted that there is relatively little difference between these cues. The findings will now be considered in terms of implications for the associatively-mediated stopping theory, the feature-positive effect, as well in terms of future experimental designs.

### **2.8.1 Implications for the associatively-mediated stopping hypothesis**

Firstly, it should be noted that across Experiments 2 and 3 I found good evidence that participants learnt to discriminate between a go and a stop cue, with the G vs. B contrast being significant at test for reaction times in Experiment 2 and the joint analysis. This indicates that participants were able to learn about the contingencies for the cues at an associative level. The results also suggest that, in line with the associatively-mediated inhibition hypothesis (Verbruggen, Best, et al., 2014; McLaren & Verbruggen, 2016), pairing an arbitrary stimulus with withholding a response leads to the stimuli having slower reaction times when a response to it is required (the reasons why I believe cue B to be a stop cue are explored below).

As discussed in Chapter 1, there is good reason to ask if go/no-go paradigms require action cancellation. If go/no-go tasks are measuring strength of action initiation this undermines the claim that, using such tasks, inhibition can be developed through pairing stimuli to stop responses. To increase the likelihood of action cancellation occurring within go/no-go tasks they are typically designed so that go is the prepotent response. The assumption is that with go being the default response the go process is activated immediately on the presentation of a stimulus. However, this approach is not suitable for studies investigating human learning which, like Experiment 2 and 3, require an overall 50:50 go likelihood ratio so as not to bias overall learning. Does this in fact mean that the tasks here are simply measuring action initiation? One aspect of the design suggests this is not the case: the adaptive staircase ensures that participants are always responding as quickly as they can and keeps response times under 2000ms. The reaction time means for the experiments reported here were mostly around 400ms. This is a short latency, especially when considering that reaction times are unlikely to be quicker than 150ms, therefore I am reasonably confident that action cancellation is being measured by the go/no-go task in these experiments. This is, because the sheer speed of

responding suggests that the go response is initiated very rapidly, implying that responding would then have to be modulated by inhibition. This argument is further bolstered by the low omission errors seen across the experiments. This indicates that go responses are prepotent, as one would expect high omission errors if cues caused the absence of an action rather than inhibition of action.

One feature of the experiments presented here is the mix of 25%, 50%, and 75% cues. Throughout that chapter it has been held that cue B (75% stop) represents stopping and cue G (75% go) represents going. But can a 75% stop cue be used as evidence for associatively-mediated inhibition? Might it rather be called a weak go cue? If the terms of the experiments are reframed from inhibition to excitation, then the results could be phrased as the 25% cue becoming an excitatory cue, leading to more commission errors and faster reaction times. A similar concern was raised by Bowditch et al. (2016), who made the case that it is possible to use the 50% cues as a baseline as these cues are neither associated with going or stopping. This logic does not fully hold for the experiments presented here, as the 50/50 cues (ignoring Experiment 1) are either involved in other compounds (e.g., cue A is also seen in cue RA) or involved in tracking (cue J). While the contingency for J is 50/50 the cue receives additional training on its own at the start of the experiments and so is treated differently. Furthermore, due to the tracking procedure there is a positive feedback loop on cue J, with responses to J directly affecting the outcome of the next response. Yet, analysis by Bowditch et al. (2016) found evidence of differences between 25% stop and 50% stop cues and 50% stop and 75% stop cues, indicating that 75% stop cues can lead to inhibition and that 25% stop cues can become excitatory, suggesting that both excitatory and inhibitory effects can occur. Other research also supports the notion that cues requiring stop responses can become associated with inhibition. For example, imaging studies have found increased activation in areas linked to inhibition on presentation of cues previously linked to stopping even when a go response was required (Lenartowicz, Verbruggen, Logan, & Poldrack, 2011). Furthermore, Leocani et al. (2000) found that on presentation of no/go cues motor output fell below resting levels, suggesting that performance on no/go trials cannot just be inaction but rather involve active action cancellation.

## **2.8.2 The Feature-positive effect**

The research presented here further supports the notion that not all discriminations between singletons and compounds are equal, and that those that signal the absence of an outcome (R+ RA-) are, in some instances, harder to learn than compounds that signal the presence of an outcome (R- RA+). Building on the work of Bowditch (2016) the research reaffirmed that this can be learnt in an incidental procedure and, that depending on the task at hand, what would be the traditional feature-negative discrimination can have an advantage – quite at odds with past work in this field (Lotz et al., 2012) – but easily explained in terms of the effective outcome interpretation.

### **2.8.2.1 Variability of the effect**

The findings suggest that the feature-positive effect is not consistent. This could be because the effect itself is relatively weak or that overall learning is weak. In Experiment 2, the results from training reaction times indicated that the feature-positive discrimination was easier to learn than the feature-negative. However, the effect was only marginally significant at test and the difference was not found in the commission error data. While in Experiment 3, none of the critical differences in reaction time data were significant (although the training contrast was in the right direction), and only the discrimination for commission errors during training supported the feature-positive effect.

It is worth noting the commission error results might reflect the induced changing task priorities for participants. When participants are looking to stop (Experiment 2) then the measures that relate to 'going' seem to be less sensitive than when participants are looking to go (Experiment 3). Similar results were found by Bowditch (2016), and past work (Verbruggen, Stevens, & Chambers, 2014) suggests that stimulus detection is a limited resource with a balance needed to be struck between ignoring irrelevant information and monitoring for occasional but highly relevant signals. Thus, when participants are in a stop task set, more attention is directed to focusing on stop cues and therefore they are less likely to respond incorrectly, whilst when in a go task set attention is shifted to focusing on responding and so more commission errors occur.

### **2.8.3 The effect of task set**

The results do suggest that a participant's task set is key to the learning that takes place at traffic lights. When participants are in a 'go' task set (perhaps when waiting at traffic lights) then they are in a better position to learn that red traffic lights means stop, but when participants are in a 'stop' task set (perhaps when approaching traffic lights) learning is modified and red traffic lights are seen as a more neutral cue and now amber lights may even come to signal go.

This has a number of implications. First, and counter-intuitively, it appears that drivers may learn that red and amber cues mean 'stop' best when the effective outcome is going. I can re-phrase this to say that it is when they are looking for signals that allow them to make progress, but the default is to stop, or at least proceed with caution, that red and amber lights will act as stop cues. This is rather surprising, as one might have thought that when the default is to go, and one is looking out for a stop signal, then this is when learning about stop will be optimal, but the evidence suggests that this is not the case. Instead, it may be that establishing a task-set in drivers that has the default as stopping coupled with an active search for signals that denote permission to proceed will be most effective in curbing running red lights and jumping amber ones. Secondly, the current model of contingency learning at traffic lights used in the three experiments reported here is inadequate. Task set seems to be a key factor in participant's learning. However, when driving it is likely that different task-sets, and hence outcomes, are effectively in play for different contingencies. By this I mean that some lights are experienced by their very nature in a stop task set (when approaching lights; Green, Amber, and Red lights) and some lights will be experienced in a go task set (when stationary at lights; Red, Red and Amber, and Green), thus the approach taken thus far is, at best, incomplete. I will return to this idea in Chapter 4 where I design a task that more accurately captures the contingencies in play in these scenarios.

### **2.8.4 Modelling the contingencies**

It has been argued that the incidental go/no-go task provides a model for exploring the learning of contingencies experienced at traffic lights. However, the model needs refinement. Firstly, the model does not contain the sequential information that traffic lights provide. That is, in the real-world a red traffic light is always preceded by an amber, yet in the experiments here it could be preceded

by a 'green' light. Secondly, the contingencies used in the designs did not match exactly those of traffic lights in the real-world, for example cue G was predicting of going 75% of the time, whereas in the real-world green always signals go. Thirdly, traffic light changes are timed, e.g., red and amber is only displayed for a short while before being followed by green, whereas in the current design the 'lights' are not participant to any intentional timing effects. These design issues are all addressed in further chapters, but it is worth highlighting here how substantial the jump from the current design to actual experience of the contingencies at traffic lights is.

Additionally, this chapter started with the aim of exploring what learning is primed by associatively-mediated processes in order to inform development of later interventions. However, while the data suggests *what* is being learnt, the expression of this learning is less clear. It might be the case that the contingences highlighted in the task are not reflective of the decision-making processes drivers make in the real-world. A driver's response to traffic lights is likely heavily influenced by context, prior experience, and initially explicit instruction. Therefore, while it seems that at an associative level learning in this task does not support the rules in the Highway Code, the expression of this learning, and thus its effect on behaviour, will be tempered by a variety of factors.

### **2.8.5 Conclusion**

To conclude, the experiments presented here mark the initial development of a laboratory-based paradigm that tries to capture the experience of contingencies at UK traffic lights and its impact on associative learning of those contingencies. Caveats about the design limitations aside, the experiments demonstrated that cue A was experienced as weak go cue in a stop task set and a neutral or weak stop cue in a go task set. The joint analysis which was intended to capture more realistic experience of task set at traffic lights indicated that amber was experienced as a go cue. Thus, in situations when control processes are weak, the evidence suggests that associative processes could lead people to commit traffic violations. These findings also suggest that the development of interventions to target these stimulus-response associations could be useful in addressing dangerous behaviour at traffic lights. Finally, the results also demonstrate how small changes in the experimental design can lead to shifts in



task demands, and how task set could be an important factor in what is learnt at traffic lights.

So far, the task has presented each cue independently (i.e. cue G could be followed by cue R), but within UK traffic lights there is a set sequence and it possible that this sequence impacts upon learning. Therefore, the following chapter explores the role of sequences by adding them to the design already employed.



### ASSOCIATIVE PROCESSES II: THE EFFECT OF SEQUENCES UPON CONTINGENCY LEARNING

This chapter discusses the effects of sequences upon the learning of cues. Chapter 2 highlighted how the learning that takes place in response to cues at traffic lights does not necessarily mirror the rules of the road. However, the experiments did not incorporate the sequences experienced at traffic lights. This chapter rectifies this and shows the effect sequences have upon the associative learning resulting from the contingencies experienced at traffic lights.

#### 3.1 THE IMPORTANCE OF SEQUENCES

The importance of sequences in human cognition cannot be overstated. As Ashe, Lungu, Basford, and Lu (2006) note, sequences are key to language, episodic memory, and motor movements. A sequence is a list of cues, events, or digits that follow an order, e.g., an area code for a telephone number such as 01392. Sequences may be rule based, but these rules can operate at a specific level (such as '1' must be the second digit within the sequence for the above example) or the rule can be more abstract, with a particular set of items preceding a different class of items. For example, in the telephone number 01392 661000 the first set of digits would be the area code, with the succeeding digits being the specific number for the person you wish to call (in this case the University of Exeter main switchboard). Crucially, the area code needs to be entered first for the call to connect.

Research on sequence learning in humans has focused heavily on the study of motor sequences due to the ease in which learning effects can be found in a short period of time (Ashe et al., 2006). Many of these experiments use the serial reaction time task by Nissen and Bullemer (1987). In this task a single visual cue can appear in any one of four locations within a trial. When the cue appears, participants must press the appropriate key for the location of the cue; the reaction time between the appearance of the cue and the participant's response to it is the primary measure of performance. Unbeknown to the participant, the position of the cue is controlled by a sequence, e.g., positions 1-

2-4-2-1-3-3-2-1 occur in that order. These sequence trials are then followed by random trials where the cue is presented on the screen but without the sequence. Typically it is found that learning to the cue becomes enhanced, that is faster reaction times and fewer errors, when it is presented within sequence, compared to learning in the random order blocks (Willingham, Nissen, & Bullemer, 1989; for a review see Robertson, 2007).

While there are several types of sequences (for a review see Conway and Christiansen, 2001) this chapter will concern itself with fixed sequences. These are sequences where the pattern is pre-set. For example, the phrase 'It was a dark and stormy night' is a fixed sequence. The operation of traffic lights is another example. These sequences are generally easy to learn, with young human children (aged between two and 4 years old), capuchin monkeys, and chimpanzees all being shown to be able to learn them (Custance, Whiten, & Fredman, 1999; Whiten, Custance, Gomez, Teixidor, & Bard, 1996). Indeed, the chimpanzee Ai, who received training in numbers and symbols has been found to be able to learn a sequence of up to five numbers in line with young human children (Kawai & Matsuzawa, 2000). In general, the main point from this brief discussion of sequences is that they are important to human daily life and easy to learn (Clegg, DiGirolamo, & Keele, 1998), with cues in sequences showing enhanced learning compared to randomly ordered stimuli.

### **3.2 ASSOCIATIVE LEARNING OF SEQUENCES**

One area of contention in the field of sequence learning, especially relevant to this thesis, is the question of whether sequences can be learned through associative systems. Early evidence in support of this argument came from work by Nissen and Bullemer (1987) who found patients with amnesia were still able to learn the sequence contained within a serial reaction time task. Using a similar approach Cohen, Ivry, and Keele (1990) asked participants to complete a serial reaction time task while also reporting on the pitch of a presented tone. The idea here is that if sequence learning is not reliant on propositional learning then learning in the experimental group should be better than in the control group who received random sequences, despite the cognitive load induced by the tone task. Results demonstrated that those in the sequence group showed an improvement in performance (as measured by a reduction in reaction times)

compared to the control group. In order to assess awareness, participants completed a generation task at the end of the experiment. In this task participants were required to press the key indicating where the position of the next cue would be, rather than responding to the position of the current key. The results showed a non-significant difference in accuracy of predictions between the two groups. Thus, there was a dissociation between response times and awareness which was claimed to indicate that participants were learning the sequence associatively (Willingham et al., 1989; Reed & Johnson, 1994)

Work in the mid-1990's cast doubt on the appropriateness of the generation task to assess awareness. The seminal work of Shanks and St. John (1994) argued that such tests failed to establish that learning is occurring outside propositional routes. The authors argued that such awareness tests often do not provide feedback thus running the risks that participants might be forgetting the sequence and therefore be unable to report it. The authors also argued that lack of awareness could be due to participants not transferring knowledge from one task to another, that is, participants could have seen the two tasks as distinctly separate and so not apply knowledge learnt in the reaction time task to the awareness task. In support of this transference argument Perruchet and Amorim (1992) used an adapted generation task whereby participants were directly instructed to form sequences similar to those they had seen in the training blocks. The authors found evidence that successful reporting of the sequences matched the enhanced response times to these sequences. That is, participants were able to report explicit knowledge of the sequences and this corresponded to shorter response times, indicating a shared process for performance and conscious knowledge (but see Cohen & Curran, 1993; Willingham, Greeley, & Bardone, 1993).

However, there are two strands of evidence that suggest learning of sequences can occur separately from explicit processes. The first strand comes from the field of clinical psychology, while the second derives from laboratory-based experiments. It has been argued that reading deficits associated with dyslexia are, in part, caused by impairment in incidental learning of sequences (Lum, Ullman, & Conti-Ramsden, 2013). Therefore, in tasks where sequences are not made explicit it would be expected that dyslexic individuals would learn the

sequence to a lesser extent than healthy controls. This is what Jiménez-Fernández, Vaquero, Jiménez, and Defior (2011) found. The authors compared performance between dyslexic children and good readers in both intentional and incidental versions of a sequence learning task. Results showed that in the incidental version, normal controls exhibited typical learning patterns, with reaction times shortening over time, while dyslexic participants showed no such pattern. In contrast, when the task was intentional, in that participants were informed about the sequences, both controls and dyslexic participants showed the typical pattern of learning.

The second strand of evidence supporting the notion that sequences can be learnt through associative process comes from work by McLaren and colleagues over several years. For example, F. Jones and McLaren (2009) conducted two experiments looking at the differences between incidental and intentional learning within a serial reaction time task. Experiment 1 explored learning under incidental conditions. The experimental group was exposed to four sub-sequences (XXX, XYY, YYX, and YXX) while the control group only saw pseudo-random sequences. In a subsequent test phase, both groups received the same trial order of pseudo-random sequences. Performance was measured on occurrences of the experimental sub-sequences within the pseudo-random test block. Overall, there was evidence that sequence learning had occurred, with better performance (less errors, faster response times) in the experimental group at test compared to the control. Further analysis revealed the effects were driven by differences in sub-sequences YYX and XYY. Strikingly, a structured interview conducted after the test phase found that the experimental group were not able to verbalise knowledge of the sequences to explain their performance on the task. In a separate experiment where participants were informed of the sequences, the XXX sequence showed the greatest evidence of learning. Therefore, it seems there is a dissociation of learning, with enhanced learning of sequences XYY and YYX under incidental conditions, while under intentional conditions learning is best to the XXX sequence. These results lend weight to the view that propositional knowledge is not necessarily key to sequence learning, and that learning can occur at an associative level (Spiegel & McLaren, 2006; McLaren, Jones, McLaren, & Yeates, 2013; Yeates, Jones, Wills, McLaren, & McLaren, 2013). However, it should be noted that associative learning and propositional learning are

perhaps more closely interlinked than the discussion here suggests (Spiegel & McLaren, 2001), and it is likely that sequence learning occurs under something like the framework proposed by McLaren et al. (2019).

### **3.3 PRESENT EXPERIMENTS**

So far in my research on associative learning at traffic lights no account has been taken of the role of sequences. That is, in the previous experiments, a green traffic light could be followed by a red traffic light whereas, in the UK at least, it is followed by an amber light. Research has shown the advantage that sequences confer to learning of cues (Willingham et al., 1989) and how sequences can be learnt at an associative level (F. Jones & McLaren, 2009). Therefore, it is fair to assume that the current experimental design is lacking a key source of learning that is not only in play in the real-world but is also likely to have a large effect on the learning of the cues. Thus, to ensure that the laboratory experiments reflect real-world learning, the design will need to employ the relevant sequences.

### **3.4 EXPERIMENT 4**

Experiment 4 is a replication and extension of Experiment 2, in that it uses a stop task set, but with the addition of (some of) the sequences experienced at traffic lights. I have already talked about the main sequence at UK traffic lights, which is green-amber-red-red and amber-green, and so it would be easy to imagine that simply having Experiment 4 follow this pattern would be sufficient. However, while this is the main sequence, a driver can enter this at any point and thus there are several sub-sequences that might be experienced. For example, red and amber together followed by green – with this sequence not containing a singleton red or amber. Therefore, it was important that the sequences employed in the experimental design captured the range of potential sequences experienced at traffic lights in the real-world. Table 3.1 outlines the sequences that participants experienced in the experiment.

The sequences were chosen to not only reflect the various start points possible but also to capture some of the contextual experience of traffic lights. For example, in Green 1, G is go in both cases, mimicking the situation when the lights have recently changed to green, and thus it is likely a driver would have

time to cross the junction. Whereas in Green 2 G is stop, modelling the situation where a driver has not been able to cross a junction on a green light, or has entered the junction when the lights have been green for some time and so could be expecting them to change and hence might be more cautious. There is an argument for having a preceding go G before the stop G to mimic the shift in driver's perception of G being a go cue to a prepare to stop cue, but there was a limit to the number of sequences that could be included in order to keep the experiment relatively quick. More generally, it is worth highlighting that it is evident that the sequences employed here are not going to capture all experiences in the real-world. The use of double G and double R served to model the fact that traffic lights often display red or green, with amber and red/amber being transitional signals.

Light type	Sequence
Green 1	G- -> G-
Green 2	G+ -> A+ ->R+ ->R+ ->RA- ->G-
Amber 1	A- ->A-
Amber 2	A+ ->R+ ->R+ -> RA- ->G-
Red 1	R+ ->R+ ->RA- ->G-
Red 2	R- ->RA- ->G-
Red/Amber	RA- ->G-

Table 3.1. Sequence runs for Experiment 4 '+' is 100% stop, '-' is 100% go.

To balance the number of singleton '-' and '+' cues, the number of 50/50 cues, and compound '-' and '+' cues, a more complex array of filler cues were used than in previous experiments (see Table 3.2). Whereas the designs in Chapter 2 used 75% going and stopping, in this chapter due to the wish to focus more on the experience at traffic lights the overall going or stopping to a cue is determined by the sequences, for example cue R has a go ratio of 1:6, and cue G a go ratio of 8:1 (for comparison in Experiment 2 the ratio for R was 1:3, and cue go was 3:1). Therefore, cue B is no longer a 75% cue but now directly



balances out cue G (i.e. has a 8:1 stop ratio). I also introduced a new cue, YZ± which was a compound cue to balance the experience of other compound cues, RA and IP. Of course, these and other filler cues were not experienced identically to the experimental cues, as they are not themselves in a sequence. Given that learning to cues in sequences has been shown to be enhanced compared to cues not in sequences (Nissen & Bullemer, 1987), to ensure that any filler cues were of equal status (in terms of learning) to the traffic light cues it would have been necessary to have them in the same sequence runs as the traffic light cues, i.e. Blue 1 would have been the opposite of Green 1, so B+ -> B+. However, given the ubiquitous nature of the traffic light sequence, it was felt best to have single filler cues so as to break up the runs of sequences. This aimed at helping to prevent participants focusing too much on the sequences, and thus guess the nature of the task (this decision does influence the analysis undertaken which is discussed in the Results section). Furthermore, including the ideal control sequences in the design would have substantially increased the length of an already long experiment, exacerbating participant fatigue.

Filler type	Cue	Occurrences per block
Cue G filler	B-	4
	B+	32
Cue R filler	I-	24
	I+	4
Cue RA filler	IP+	24
50:50 filler	P±	16
50:50 filler	YZ±	16
Tracking cue	J±	8

Table 3.2. Filler sequences for Experiment 4. '+' is 100% stop, '±' is 50/50 go, '-' is 100% go. Each letter in the cue column represents a shape drawn randomly from a set pool.

This design does leave the experiment open to issues of repetition priming (which were not an issue in my earlier experiments). This is a separate issue from sequence learning and arises from the influence of past trials in the sequence on responses to the current trial, potentially producing shorter response times if the cue signals the same position/response previously

experienced/required (Soetens, Boer, & Hueting, 1985). Therefore, differences in performance could be driven not only by sequence learning but also by repetition priming effects (F. Jones & McLaren, 2009). However, aspects of the design are likely to minimise any such effects. Firstly, the experiment is within-participants. This means that all participants receive both filler cues and sequence cues helping to dilute repetition priming effects. Additionally, the direction of response and, to some extent, the colour of the cue, are randomised. Even so, there is evidence that repetition priming effects can still occur even in simple designs such as that employed here (M. Jones, Curran, Mozer, & Wilder, 2013). I will use specific analyses to check that there are the associative effects above and beyond any repetition priming effects in my results section, but note that the emphasis now is not so much on what the basis of a speeded or slow response is, but on how people perform on the task under the conditions that prevail during training. Put another way, the emphasis shifts from an interest in mechanism to an interest in the analogy with real-world performance in these experiments.

In the introduction to this chapter, the issue of associative learning to sequences was raised. There is good reason to assume that the work presented here will support the view that sequence learning can occur associatively. Given that the basic design used here is similar to that of Experiment 1 in F. Jones and McLaren (2009), in that it is a two-choice reaction time task and participants will not be informed of the sequences beforehand, it is not a leap of logic to assume the lack of explicit knowledge found in that study will also hold true for the current work. As a check on this assumption, an awareness measure was introduced (see Procedure for details).

### **3.4.1 Method**

#### **3.4.1.1 Participants**

The inclusion criteria and outlier removal process were the same as Experiment 2 (see Results section for details of those removed). A power analysis using the R package SIMR (Green & MacLeod, 2016) was conducted on detecting an effect between a go vs. stop cue (G vs. B, see section 3.4.1.2 for details) at test in reaction times. The effect size for this difference was set at 14.21ms (calculated by averaging the difference in the G vs. B contrast at test across

Experiments 1-3). Overall, the analysis indicated that a sample size of 80 would give sufficient power (82.30%) to detect the effect. However, in view of the stronger learning typically observed when sequences are involved it was decided to test 55 participants and run a new power analysis. The post-hoc power analysis for first sequence trial data at test (which is comparable to Experiments 2 and 3) gave the final sample size of 56 participants a power of 56.30% to detect a difference of 14.21ms between the go vs stop cue (G vs. B) at test. Though, of course, given the use of sequences in the current study, it would be expected that the difference between G and B at test would be greater than in Experiments 2 and 3. This was the case, with the difference being 22.35ms, with the study having a power of 89.30% to detect a 22ms difference. Participants received payment of £5 or one course credit.

#### **3.4.1.2 Design**

Overall, there was one calibration block, six training blocks, and one test block (see Table 3.3 for summary of design) with a 10 second break between each block and a 10 second break halfway through each block (excluding the calibration block), with the proviso that these breaks did not disrupt a sequence. This was achieved by randomly setting the middle trial to be one of the filler cues, with the sequences being randomly ordered around this fixed point. As seen in Table 3.3 each sequence appeared more than once, e.g., the sequence Amber 2 appeared four times. This was in order to give participants enough experience of the sequences to encourage learning. In training, the contingencies for each sequence were those outlined in Table 3.1. For test, all contingencies were 50/50, for example Green 1 was no longer G- -> G- but rather  $G_{\pm}$  ->  $G_{\pm}$ . This 50/50 split was achieved by having half go and half stop sequences, so there were two G- -> G- and two G+ -> G+ in the test block which overall created the  $G_{\pm}$  ->  $G_{\pm}$  contingency. This mean that the sequences of responses over the block were no longer predictive. This approach enabled me to collect enough data on go and nogo trials to analyse both reaction time and commission errors. It should also be noted that due to the total length of the experiment the calibration block was halved to 24 trials.

Phase	Blocks	Individual trials per block	Design
Calibration	1	24	J
Training	6	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16)
Test	1	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16)

Table 3.3. Summary of Experiment 4 design. Letters represent a sequence. The numbers in parentheses indicate how often the sequence will occur per block. At test the contingencies will all be  $\pm$ .

### 3.4.1.3 Procedure

The procedure was identical to that of Experiment 2 except that there was only one test block and there were breaks halfway through the training and test blocks. In the study any sequence randomly followed on from another, thus a run of trials could be RA- G-/ B-/ G+ A+ R+ R+ RA- G- which would be comprised of three sequences, 'Red/Amber', 'Go singleton filler' and 'Green 1'. At the end of the experiment participants were asked to complete an awareness task.

#### 3.4.1.3.1 Awareness measures

The awareness measure was taken from Bowditch (2016). In the task participants were shown the coloured shape cues used in the experiment and asked to rate on a scale of 1-9 (from 'Definitely Not' to 'Definitely') a) How much would you expect to RESPOND to this shape configuration? and b) How much would you expect to WITHOLD your response from this shape configuration? The order of questions was counterbalanced, and the cue order was randomised with the caveat that compound cues (e.g., RA) were always presented first followed by singleton cues (e.g., G). The shapes were presented at the same location and at the same size as in the previous blocks.

### **3.4.2 Analysis and results**

In terms of the statistical approaches, Experiment 4 was conducted in much the same manner as the analysis of training for Experiment 2. In terms of outliers, three participants were removed and replaced for having high omission errors, while two participants were removed and replaced for having high commission errors. Two participants failed to complete the experiment. There were no outliers for reaction times.

However, due to the complexity of the experimental design the analyses differ from that of Chapter 2. For ease the analysis of the training and test phases, as well as the awareness test will be discussed separately, with the results of each section following presentation of each analysis plan.

#### **3.4.2.1 Training phase**

One issue with the experimental design was that comparisons between the cues in the sequences and filler cues are likely to be affected by sequence learning and repetition priming effects. As discussed in the introduction there is evidence to indicate that learning involving sequences is enhanced compared to learning without sequences (Nissen & Bullemer, 1987). Such effects mean it would be inappropriate to conduct tests exploring learning of the feature-positive effect due to the fact that R vs. RA is a sequence contrast, while I vs. IP is not, and as such the feature-positive effect is explored separately in a later section. Furthermore, the effects might also impact the contrasts involving cues G, A and R against B, with the latter not being in the traffic light sequence. However, it was decided to retain these contrasts as the only way to provide contrasts against a clear stop cue, as there is no such cue in the experimental sequence (with R being also involved in RA, a go cue). This approach also fits with the need to try and capture more of the real-world behaviour when experiencing these contingencies at traffic lights, as they will be embedded in sequences like these. A further issue is the repetition priming effects in play for the traffic light sequences, thus findings between cues could be explained by referring to this effect, rather than an associative learning account. However, it is possible to run contrasts in the training phase that control for repetition priming and provide evidence that associative learning is taking place in this phase.

To summarise, contrasts were run on both training reaction time and commission error data to demonstrate that learning had taken place within the training phase. These contrasts were followed by the traffic light analysis (as in Experiment 2) to explore what responses cues might promote when the sequences experienced at traffic lights (and thus the inherent sequence learning and repetition priming) are considered (see descriptive statistics in Table 3.4).

#### **3.4.2.1.1 Reaction times**

For the training data, the best fitting model had a conditional  $R^2$  of 0.87 (see Table 3.5 for AICs).

##### **3.4.2.1.1.1 Evidence of associative learning**

To provide evidence of associative learning for reaction times at training it is possible to compare performance between R vs. I and R vs. B. In effect, these contrasts are a manipulation check (and separate from the main analysis), and thus the alpha level was only corrected to .025. In Table 3.1 the only occurrence of R- (a go R cue) is in the sequence Red 2. Therefore, given that reaction times are calculated by averaging correct responses times for a cue, the only data that will be considered for the average reaction time for cue R will come from this trial. Crucially, R- is the first trial of the Red 2 sequence, and thus will not be affected by repetition priming effects (assuming that the preceding cue was 50% stop across all participants). Cue I is the filler cue for cue R and so it is a natural comparison cue for this contrast. Noting that cue R is overall a stop cue, and cue I a go cue, if participants were learning about the contingencies it would be expected that cue R would have significantly slower reaction times than cue I. The R vs. I contrast was significant,  $t(440) = -3.22$ ,  $p = .001$ , 95% CI [-15.27, -3.72],  $d = -0.31$ , with I being faster than R, confirming that participants were indeed learning the contingencies present in the design. The contrast R vs. B enables me to compare cue R to an out and out stop cue, with this contrast also being significant,  $t(440) = 3.74$ ,  $p = < .001$ , 95% CI [5.23, 16.78],  $d = 0.36$ , with faster reaction times to cue R than B, suggesting cue R did not promote as much a stopping response as cue B.

### 3.4.2.1.1.2 Traffic light contrasts

Noting the above, I can be confident that the following results arise as a consequence of associative learning as well as repetition priming. As G vs. B is no longer a pure manipulation check, the alpha level was corrected to .008. The G vs. B contrast was highly significant,  $t(440) = 11.64$ ,  $p < .001$ , 95% CI [28.52, 40.07],  $d = 1.11$ , with G being faster than B, thus confirming that participants were responding in accordance with the contingencies present in the design. In terms of the experimental contrasts, A vs. B was also highly significant,  $t(440) = 9.35$ ,  $p < .001$ , 95% CI [21.77, 33.32],  $d = 0.89$ , with faster responses to cue A. The contrast A vs. G was not significant at the reduced alpha,  $t(440) = -2.29$ ,  $p = .023$ , 95% CI [-12.52, -0.97],  $d = -0.22$ , but there was a trend for slower responses to cue A than to cue G. The contrast for A vs. R was highly significant,  $t(440) = 5.61$ ,  $p < .001$ , 95% CI [10.76, 22.31],  $d = 0.54$ , with faster reaction times to cue A. In terms of the R contrasts, the R vs. B contrast was significant,  $t(440) = 3.74$ ,  $p < .001$ , 95% CI [5.23, 16.79],  $d = 0.36$ , with cue R having faster responses, suggesting it was not as much of a stop cue as B. The R vs. G contrast was highly significant,  $t(440) = -7.90$ ,  $p < .001$ , 95% CI [-29.06, -17.51],  $d = -0.75$ , with responding in the presence of R being slower, indicating that R was not as go as cue G.

### 3.4.2.1.2 p(respond)

For this measure, the best model was a model that included the main effects of cue with a Gamma family and inverse link and random intercept (see Table 3.5).

#### 3.4.2.1.2.1 Evidence of associative learning

Using the same logic outlined for the training reaction time data, contrasts for commission errors for G vs. B and G vs. I were run to provide evidence of associative learning for this measure. Looking at Table 3.1, Cue G is the only traffic light cue that has a stop contingency that only occurs at the start of a sequence (Green 2). Given that cue B is a stop cue, from an associative learning perspective it would be expected that cue G would have more errors than cue B and this was the case,  $z = 4.84$ ,  $p < .001$ , with more errors for cue G, thus confirming that the participants were learning the contingencies present in the design. The G vs. I contrast was run to establish if cue G was a go cue of similar power to a filler go cue. The contrast was not significant,  $z = 0.88$ ,  $p =$

.379, indicating that learning was similar between these cues (supported by the numerical values in Table 3.4).

### 3.4.2.1.2.2 Traffic light contrasts

Given the result for G vs. B, I can be confident that the following results arise through associative learning as well as repetition priming. For this set of analyses, the alpha level was corrected to .008. The G vs. B contrast was significant,  $z = 4.84$ ,  $p = < .001$ , with more errors for cue G, thus confirming that the participants were learning the contingencies present in the design. The A vs. B contrast was not significant,  $z = 0.85$ ,  $p = .396$ . However, the A vs. R contrast was significant,  $z = 3.39$ ,  $p = < .001$ , with more errors for A than R, suggesting that A was more of a go cue than R. The contrast for A vs. G was also significant,  $z = -4.16$ ,  $p = < .001$ , with more errors for cue G suggesting that A was not as strong a go cue as G. Focusing on the R contrasts, the R vs. B contrast was nearly significant,  $z = -2.60$ ,  $p = .009$ , with more errors for cue B than cue R, suggesting that cue R was a strong stop cue. The R vs. G contrast was highly significant,  $z = -6.69$ ,  $p = < .001$ , with more errors for G than R, suggesting that R did not prime a go response in a similar fashion to G.

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
A-/+	358.43	40.77	0.02	0.01	0.01	0.01
B+	385.98	42.93	0.01	0.01	0.02	0.02
G-	351.69	38.57	0.03	0.03	0.01	0.01
I-	365.48	39.54	0.03	0.03	0.01	0.01
IP+	NA	NA	0.02	0.01	NA	NA
J±	369.23	40.22	0.02	0.02	0.02	0.02
P-/+	381.92	43.85	0.01	0.01	0.02	0.02
R+	374.97	41.35	0.01	0.005	0.01	0.01
RA-	377.26	41.78	NA	NA	0.01	0.01
YZ-/+	379.59	45.39	0.02	0.01	0.01	0.01

Table 3.4. Training descriptive statistics for Experiment 4. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.



### 3.4.2.1.3 Training phase summary

Overall, the training data for reaction time and  $p(\text{respond})$  analysis suggest that when considering some of the sequences experienced at traffic lights, in addition to associative learning, cue G is a go cue, cue A is being experienced as a weak go cue (with the commission error result suggesting it promotes significantly less of a go response than G), and cue R is experienced as a weak stop cue, being significantly faster than cue B for reaction times, but almost prompting a significantly greater stopping response for  $p(\text{respond})$ .

Of course, the fundamental question is whether associative learning is involved. For example, it is quite possible that the speeded reactions to G on go trials owe something to G being preceded by another go trial most of the time. This will not as often be the case for B, and so some of the advantage can be explained in terms of the facilitation in responding brought about by not having to stop on the preceding trial. However, the comparison of R vs B for reaction times is not susceptible to this explanation. The only go trials for R occur at the start of a sequence, and so the preceding trial for R and B are, on average, matched, and hence it can be trusted that R is actually faster than B (and slower than I), and so is not as good a stop cue when evaluated in this way. The point is, of course, that it may well do better when it is embedded in the sequences it is part of, and the  $p(\text{respond})$  data hint at that, but it is impossible to disentangle the effects of sequence learning from the effects of repetition priming there. Turning now to G vs. B for  $p(\text{respond})$ , it is clear that there are more errors to G than to B (and about the same as to I) which establishes that it is a go cue, because once again stop trials for G only occur at the beginning of a sequence and so are not contaminated by repetition priming. Again, this is possibly underestimating how effective G is as a go cue when encountered 'in sequence'. Nevertheless, the conclusion must be that the associatively-mediated learning effects expected based on Experiments 2 and 3 are occurring here.

Model	Training	1 <sup>st</sup> Sequence	Other sequence
Reaction time models			
Main effects of cue	5200.76	6108.90	5987.15
Main effects of cue with random intercept	<b>4404.53</b>	<b>5674.90</b>	<b>5453.67</b>
p(respond)			
Main effects of cue with Gamma family and inverse link	-3283.32	-3417.62	NA
Main effects of cue with Gamma family and inverse link and random intercept	<b>-3336.36</b>	<b>-3670.07</b>	NA

Table 3.5. AICc scores for models for Experiment 4 run on reaction time and p(respond) data at training and test. Bold are the models chosen.

### 3.4.2.2 Test analysis

Due to the design of the test phase whereby, overall, the phase had 50/50 contingencies, but this occurred through runs of certain responses, it is likely that repetition priming effects were having a sizeable impact upon performance. Therefore, it was not possible to undertake the analysis performed in Experiment 2. This is because repetition priming effects on their own could lead to cue G (almost always at the end of a sequence) being faster than R and A, which tend to occur at the start or middle of a sequence. The solution is to undertake two separate analyses. The first focuses on performance on the first trial of every sequence and as such is free from repetition priming effects (the trials just before these trials will, on average, have the same distribution of go and no-go responses). Given the lack of sequential information this is similar to the test analyses conducted in Experiments 2 and 3. The second will focus on performance to cues contained within a sequence (that is cues in any  $n^{\text{th}}$  position bar the 1<sup>st</sup>), and allows for the benefit of learning arising from cues being embedded in sequences to be explored. Therefore, the first set of

analyses are immune from repetition priming effects, whilst the second will, inherently, be affected, but equally for each of the cues embedded in these sequences (see Table 3.6 for descriptive statistics).

#### **3.4.2.2.1 First trial analyses**

The analyses conducted on the first trial of each sequence were a G vs. B contrast (which is once again a manipulation check) and A vs. B, A vs. G, A vs. R, R vs. B and R vs. G. The alpha level for these analyses was corrected to .010 (see Table 3.5 for AICc's).

##### **3.4.2.2.1.1 Reaction times**

The model chosen had a conditional  $R^2$  of 0.63. The G vs. B contrast was significant,  $t(495) = 3.44$ ,  $p = .001$ , 95% CI [9.61, 35.09],  $d = 0.31$  indicating that associative learning had taken place by the test phase. The A vs. B contrast was significant,  $t(495) = 3.12$ ,  $p = .002$ , 95% CI [7.54, 33.03],  $d = 0.28$  suggesting that cue A was not a stop cue. The A vs. G contrast was not significant,  $t(495) = -0.32$ ,  $p = .751$ , 95% CI [-14.81, 10.67],  $d = 0.03$ . The R vs. B contrast was not significant,  $t(495) = 0.36$ ,  $p = .722$ , 95% CI [-10.43, 15.06],  $d = 0.03$ , and cue R had significantly slower reaction times than cue G,  $t(495) = -3.08$ ,  $p = .002$ , 95% CI [-32.78, -7.29],  $d = -0.28$  suggesting that cue R was not a go cue. The contrast A vs. R was also significant,  $t(495) = 2.76$ ,  $p = .006$ , 95% CI [5.22, 30.71],  $d = 0.25$  suggesting that cue A primed more of a go response than cue R.

##### **3.4.2.2.1.2 p(respond)**

For commission errors there was a non-significant difference for G vs. B,  $z = 0.67$ ,  $p = .500$ . The contrast A vs. B was marginally significant at a standard alpha level,  $z = 1.78$ ,  $p = .076$ , and hints at a trend for more errors to A than B. The difference between A vs. G was not significant,  $z = 1.15$ ,  $p = .251$ . For cue R, the contrast against cue B was not significant,  $z = 0.67$ ,  $p = .500$ , as was the contrast against G,  $z = 0.00$ ,  $p = 1.00$ . The contrast for A vs. R was not significant,  $z = 1.15$ ,  $p = .251$ .

##### **3.4.2.2.2 Other sequence trial analyses**

These focused on comparing the traffic light cues within the sequences. Therefore, contrasts undertaken were A vs. G, A vs. R, and R vs. G. The alpha level was corrected to .017. Due to unequal variances between cues for

commission errors the use of mixed-effects models was inappropriate and instead standard  $t$ -tests were run for this performance measure. However, to increase ease of comparability between experiments these contrasts were run on the transformed data.

#### **3.4.2.2.2.1 Reaction times**

The model chosen had a conditional  $R^2$  of 0.72. The A vs. G contrast was not significant,  $t(495) = 0.20$ ,  $p = .839$ , 95% CI [-9.20, 11.32],  $d = 0.02$ , nor was the A vs. R contrast,  $t(495) = 0.77$ ,  $p = .444$ , 95% CI [-6.25, 14.27],  $d = 0.07$ , and neither was the R vs. G contrast,  $t(495) = -0.56$ ,  $p = .574$ , 95% CI [-13.21, 7.31],  $d = -0.05$ .

#### **3.4.2.2.2.2 p(respond)**

The A vs. G contrast was not significant,  $t(55) = -0.12$ ,  $p = .903$ , 95% CI [-0.01, 0.01],  $d = -0.02$ . Due to the variance of R being equal to 0 it was not possible to conduct the A vs. R or R vs. G contrasts. This lack of variance needs explaining, especially considering the data presented in Table 3.6. How can the mean error for R be 0.01 but the variance 0? This suggests that all participants made the same amount of error to cue R which seems unlikely. The strange results can be explained by reference to the transformation applied to the data to enable the mixed-effects models to be run. These models cannot deal with data of zero, and so as discussed in Chapter 2 I applied a transformation to the data to shift all data from zero. This has the effect of shifting all error rates by the same amount whilst leaving variance unaffected, and so leads to the 'on paper' results of cue R having an error rate of 0.01, when in fact the true error rate for the cue was 0 (and hence why zero variance).

#### **3.4.2.2.3 Test phase summary**

The first sequence trial analyses indicate that at a purely associative learning level, participants had learnt that G was a go cue (having significantly faster reaction times than cue B), cue A was a go cue (having similar reaction times to G and marginally significantly more commission errors than B) and that cue R was a stop cue, having similar reaction times to B and priming significantly less going behaviour for reaction times compared to G. The other sequence trial analyses indicate weaker effects overall and suggests that learnt behaviour for traffic light cues within sequences was weak. One explanation for this could be

that performance is being overwhelmed by repetition priming effects and expectancies generated by sequences of trials of a given type. By this I mean that if one experiences a run of go trials then one would learn to expect another go trial. Conversely if exposed to a run of stop trials then one would learn to expect another stop trial.

Type of cues	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Filler cues						
B+	384.06	51.09	0.02	0.02	0.02	0.03
I-	368.22	47.19	0.02	0.02	0.02	0.02
IP+	388.23	60.23	0.02	0.04	0.02	0.03
J±	364.88	50.80	0.03	0.06	0.02	0.05
P-/+	369.14	48.98	0.02	0.04	0.01	0.02
YZ-/+	375.09	50.37	0.01	0.02	0.01	0.02
1st sequence						
trials						
experimental						
cues						
A	363.77	59.28	0.02	0.06	0.01	0.03
G	361.71	63.90	0.02	0.05	0.02	0.06
R	381.74	62.56	0.02	0.05	0.01	0.03
RA	365.51	61.70	0.02	0.05	0.01	0.03
Other sequence						
trials						
experimental						
cues						
A	342.25	51.58	0.01	0.03	0.01	0.00
G	343.32	47.12	0.01	0.02	0.01	0.02
R	346.26	50.19	0.01	0.00	0.01	0.02
RA	335.85	42.45	0.02	0.03	0.02	0.04

Table 3.6. Test descriptive statistics for Experiment 4. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 3.4.2.3 Expectancies

Awareness scores were calculated by subtracting a participants' expectancy to withhold a response from their expectancy to respond for each cue. Thus, scores ranged from -8 (highly likely to withhold a response) to +8 (highly likely to respond).

To assess evidence of awareness the analysis undertaken for the reaction time data was repeated on the awareness questionnaire (see Table 3.7 for summary statistics). For this analysis the alpha level was corrected to .008. The best fitting model included the main effects of cue with a random intercept of participant ( $R^2 = 0.40$ ). The G vs. B contrast was highly significant,  $t(605) = 5.64$ ,  $p < .001$ , 95% CI [2.47, 5.10],  $d = 0.46$ , with participants more likely to rate G as a go cue than B. The A vs. B contrast was not significant,  $t(605) = 1.17$ ,  $p = .242$ , 95% CI [-0.53, 2.10],  $d = 0.10$ . The A vs. G contrast was significant,  $t(605) = -4.47$ ,  $p < .001$ , 95% CI [-4.31, -1.69],  $d = -0.46$ , with participants more likely to rate cue G as a go cue than cue A. The A vs. R contrast was not significant,  $t(605) = 0.64$ ,  $p = .523$ , 95% CI [-0.89, 1.74],  $d = 0.05$ . Focusing on cue R, the contrast for R vs. G was highly significant,  $t(605) = -5.11$ ,  $p < .001$ , 95% CI [-4.74, -2.11],  $d = -0.42$ , with participants more likely to rate cue G as go than cue R. The R vs. B contrast was not significant,  $t(605) = 0.53$ ,  $p = .595$ , 95% CI [-0.96, 1.67],  $d = 0.04$ . Overall, this analysis shows that participants had some awareness of the required response for certain cues. This seems to focus on cue G, with participants showing consistent awareness that this cue was more likely to involve responding than others.

Contrasting the two 50/50 filler cues that were completely neutral (that is, not involved with any other cues) against G and B enabled me to see if B was indeed a neutral cue as Table 3.7 would suggest, or if there was bias in the ratings. By this I mean the bias inherent in the use of scales, for example, the tendency of people to avoid giving extreme scores (Albaum, 1997) or differences between people in how they construe the scale. I contrasted J against G and B, and YZ against G and B. This analysis was separate from that conducted above and so an alpha level of .013 was applied. The J vs. B contrast was not significant,  $t(605) = 0.00$ ,  $p = 1.00$ , 95% CI [-1.31, 1.31],  $d = -0.00$ . The J vs. G contrast was highly significant,  $t(605) = -5.64$ ,  $p < .001$ , 95% CI [-5.10, -2.47],  $d = -0.46$ , with participants more likely to rate G as go than J. The contrast YZ vs. B, was not significant,  $t(605) = -0.45$ ,  $p = .651$ , 95% CI [-1.62, 1.01],  $d = -0.04$ . The YZ vs. G contrast was highly significant,  $t(605) = -6.09$ ,  $p < .001$ , 95% CI [-5.40, -2.77],  $d = -0.50$ , with participants more likely to rate G as go than J. These results suggest that participants expected to make very similar responses to cue B (a stop cue) and cues J and YZ (neutral cues). This suggests that cue B seems to be somewhat neutral in terms of awareness.

This neutrality is dissociated from the response time results which indicate that B is certainly not a neutral cue, but is in fact a stop cue in comparison to a certain go cue, cue G. Of course, it might seem desirable to compare the reaction times for cue B against those of the neutral cues J and YZ, but such comparisons would be invalid. This is because YZ is a compound cue, and J the tracking cue and so receives special treatment during the learning phase.

Cue	Expectancy Rating	
	Mean	SD
A-/+	1.16	4.32
B+	0.38	4.21
G-	4.16	3.75
I-	2.07	4.36
IP+	-1.66	3.94
J±	0.38	4.24
P-/+	0.32	4.1
R+	0.73	4.44
RA-	2.2	4.01
Y-/+	-2.16	4.26
YZ-/+	0.07	4.24
Z-/+	-2.29	4.65

Table 3.7. Descriptive statistics for expectancy ratings for Experiment 4.

#### 3.4.2.4 Correlational analysis

To investigate whether it is the awareness of cues that is causing the differences in performance, one can analyse the correlations between learning and expectancy scores (the zero correlation criterion, see Dienes, Altmann, Kwan, & Goode, 1995; Dienes & Scott, 2005). If awareness is impacting upon performance then I would expect a correlation between the two measures, such that cues rated with a higher likelihood to respond should have faster reaction times. However, if the correlation is zero this would indicate that awareness was not linked to performance.

This analysis was undertaken using the expectancy and performance measure scores from each pair to create two difference scores. These were calculated by subtracting the score for the stimulus most often paired with stopping from that



most often paired with responding, e.g., RA-R. Thus, if a participant's awareness of the contingencies and task performance were related, I would expect the degree of awareness to correlate with greater response time differences and greater commission error differences. This analysis was only performed where there was evidence from the expectation scores that participants were aware of the differences between the cue, so G vs. B, G vs. A, and G vs. R. For training the whole dataset was used, whilst for test only data from first sequence trials was included.

Focusing on training correlations, as seen in the top three panels for Figure 3.1, all three correlations for reaction times and expectancy scores were not significant (even at the standard alpha level before correction). For commission errors, only the G - R correlation was significant, but this became non-significant when a Bonferroni correction was applied, although this result does hint at a trend for awareness to influence performance. The same analysis was performed for reaction times and  $p(\text{respond})$  for first sequence trials (see Figure 3.2), where once again there were no significant results (even at a standard alpha level).

Generally, Figures 3.1 and 3.2 show that there was no relationship, or only a weak correlation, between awareness and performance measures at training and test. Although the result for G - R might be taken to indicate that awareness is likely to be involved at some level in performance, it is not reliable after adjustment for multiple comparisons. Overall, the results suggest that awareness did not play a significant role in creating the difference in performance between cues, and that learning was occurring through associative processes. Of course, one caveat is that perhaps there is a small causal link that the analysis lacks the power to detect.

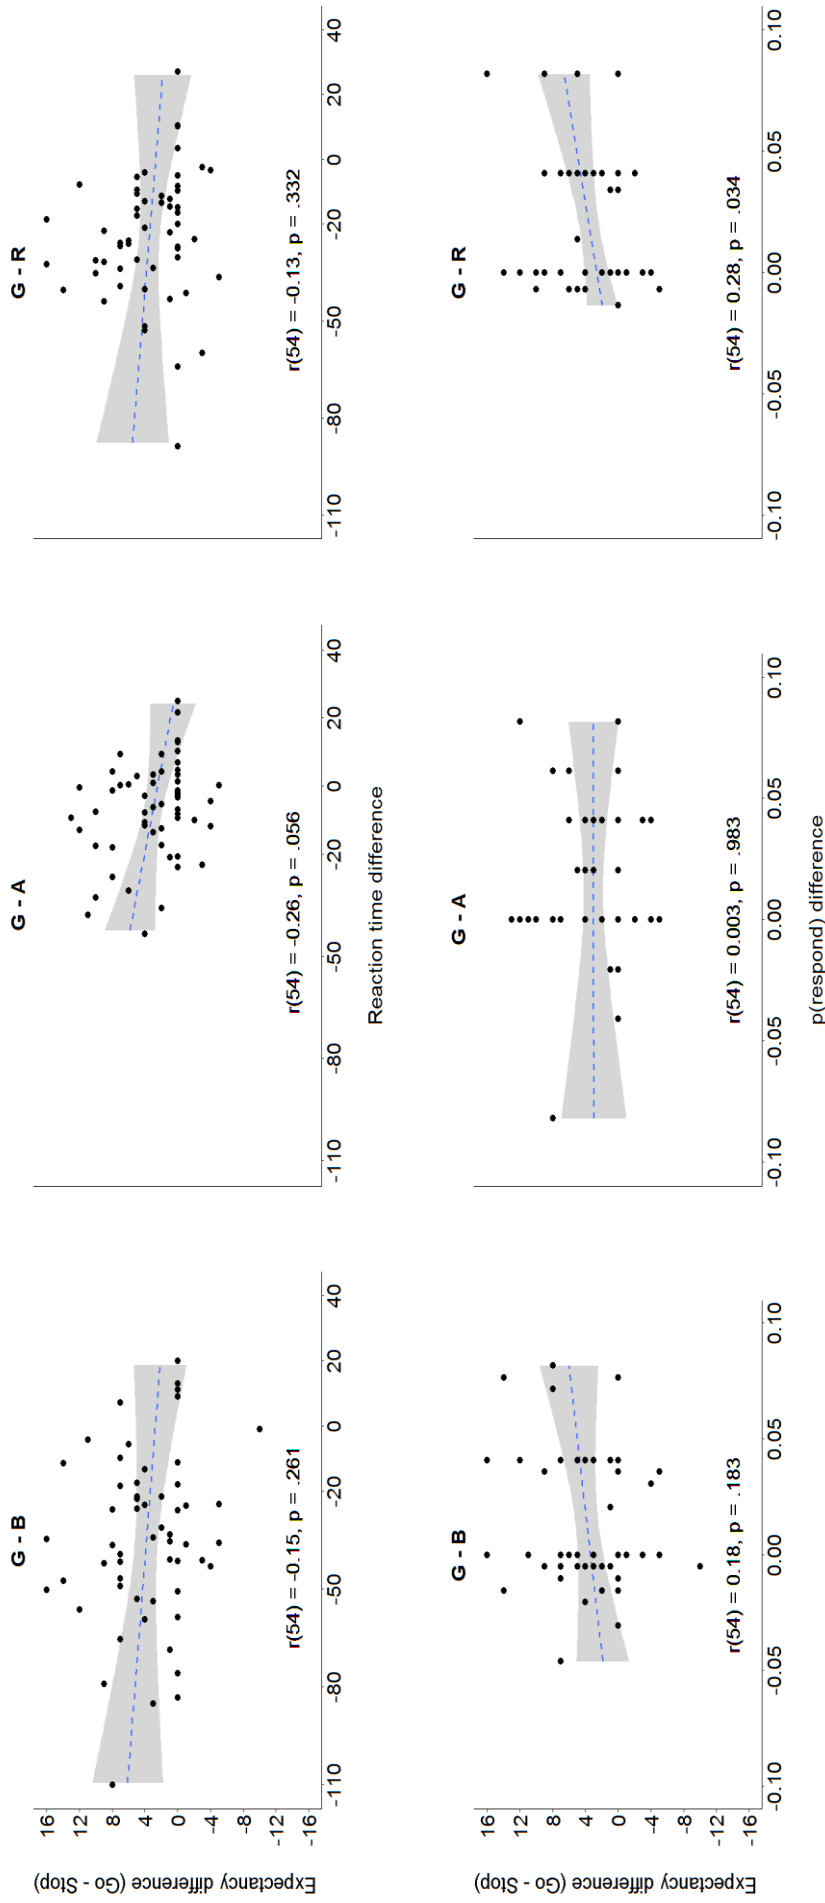


Figure 3.1. Experiment 4 training phase correlations between expectancies and reaction times (top panel) or commission errors (bottom panel) each contrast G vs. B, G vs. A and G vs. R. Scores were calculated by taking the stimuli that was more contingent with stopping from the one that was most contingent with going (i.e. go-stop). Dashed line depicts linear model of the data, shaded area represents the 95% confidence interval of that model. R values refer to Pearson's product-moment correlation.

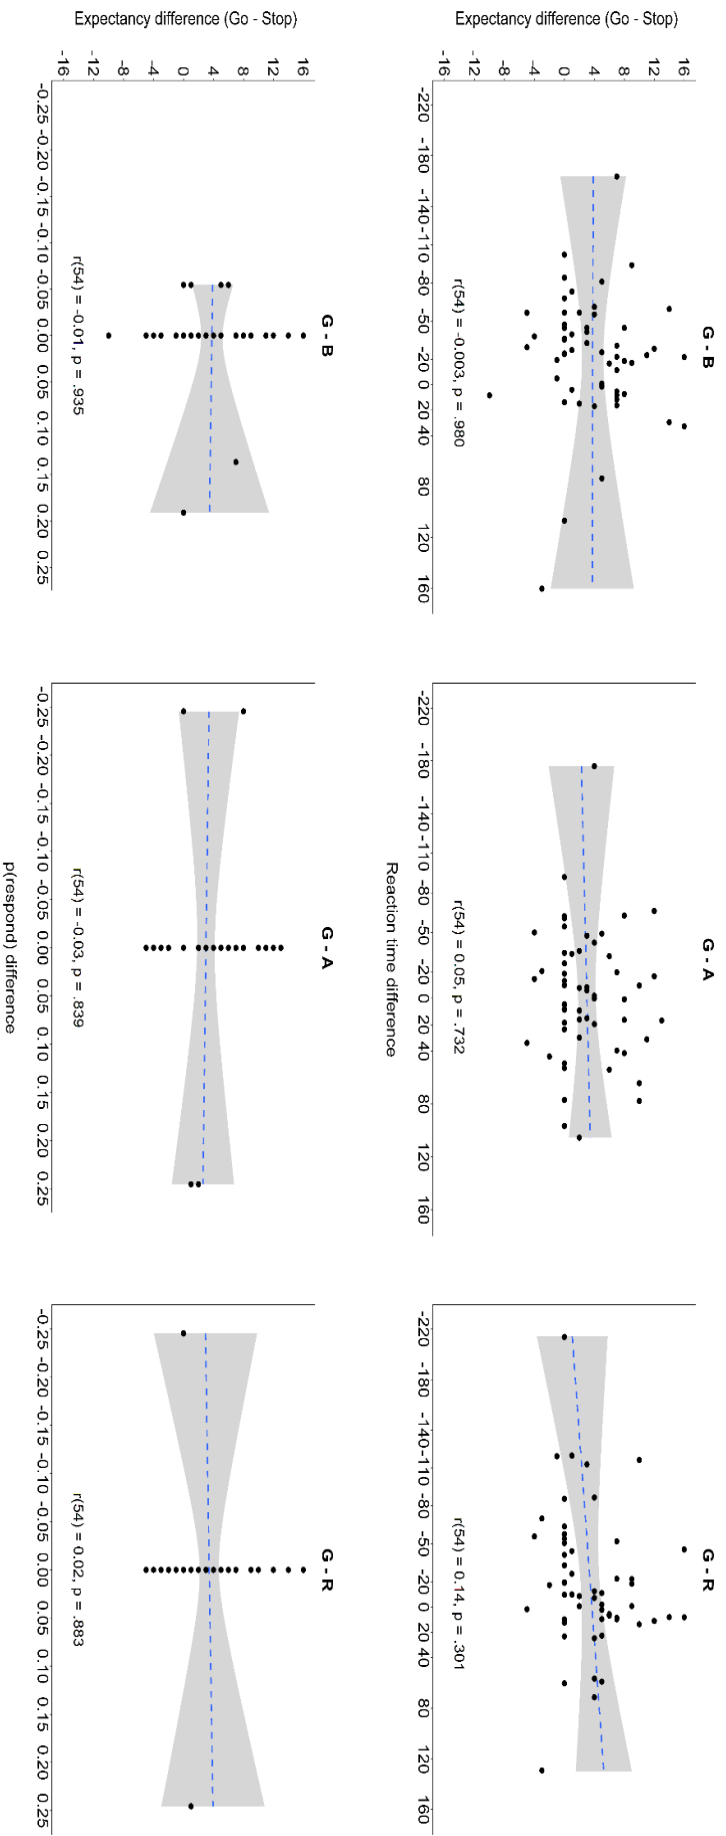


Figure 3.2. Experiment 4 test phase correlations between expectancies and reaction times (top panel) or commission errors (bottom panel) during test for each contrast G vs. B, G vs. A and G vs. R. Scores were calculated by taking the stimuli that was more contingent with stopping from the one that was most contingent with going (i.e. go-stop). Dashed line depicts linear model of the data, shaded area represents the 95% confidence interval of that model. R values refer to Pearson's product-moment correlation.

### 3.4.3 Summary

Experiment 4 investigated associative learning at traffic lights under conditions when some of the sequences involved at traffic lights were in play. During training on both measures the data showed that cue G primed more of a go response than cue B. Additionally, there was evidence to suggest that the resulting effects were due to a combination of repetition priming effects and associative learning.

For cue A the reaction time data at training suggests it primed a reasonably strong go response, being significantly faster than cue B and not significantly slower (at the corrected alpha level) than cue G. However, the result from the G vs. A contrast does suggest that cue A did not prime as much going as cue G and numerically this is borne out. The results for commission errors indicate that cue A was a stop cue. Errors were not significantly greater than cue B (though numerically it had more errors than B) and A had significantly less errors than cue G. Overall it seems that cue A was experienced as something of a weak go cue at training. Focusing on cue R, the fact that at training it was significantly slower and had significantly less commission errors compared to cue G indicates that R was not a go cue. For reaction times it does not seem to be that strong a stop cue either, having significantly faster reaction times than cue B. However, R had almost significantly fewer commission errors than B so it seems that R was a weak stop cue, being slightly more go than cue B based on the reaction time result.

For the first sequence test trials, as expected the contrast G vs. B was significant for reaction times suggesting that learning had occurred. The results suggest that cue A primed a go response. Cue A was significantly faster than cue B at reaction times, but for commission errors cue A only primed marginally significantly more errors at the standard alpha level. Cue A was also similar in performance to cue G for reaction time data. For cue R, the data indicates it was somewhat of a stop cue. It had significantly slower response times than cue G but was similar in performance for cue B.

In summary, when in a task set that has stopping as the effective outcome and sequences are in play cue A seemed to prime a weak go response, with cue R seeming to promote a weak stop response. At test, in an analysis where

repetition priming effects were controlled for, cue A seemed to be linked with a go response and cue R a stop response.

The pattern of results is not dissimilar to those found in Experiment 2, where R was a weak stop cue and cue A a weak go cue. This suggests that the basic effects already established for stop task sets hold true, despite the slight change in go/stop ratio for the cues. Of course, overall R is now perhaps more of a stop cue, and A more of a go cue. This shift for both cues could arise because they were embedded within sequences at training. It could be that the exposure to traffic light sequences enhanced the learning to the cues that seemed to occur in the 'base' version of the design in Experiment 2. Another aspect of the results that should be highlighted is that the cues embedded in sequences have produced markedly enhanced effects compared to those observed in Experiment 2. This is evidenced by the bigger difference in reaction times between G and B, in fact the difference is more double that of Experiment 2 for training. This confirms the well-known sequence learning effect (Nissen & Bullemer, 1987). The results here do not show the unexpected result seen in Experiment 2, where R had significantly more commission errors than G at test, suggesting that this was random noise rather than a true effect.

One issue touched on in the discussion of the design of Experiment 4 is that of comparing cues within sequences to those outside (e.g., G vs. B). Cue G might be expected to be faster than B, both due to the beneficial nature of sequences to learning and because of repetition priming. While it is true that G is more of a go cue than it was in Experiment 2 (with faster reaction times to it) this logic would also hold for cue B, in that if it was in a sequence then one would expect it to become even more of a stop cue than it is now. Thus, the current contrasts between G vs. B (and more generally sequence cue vs. non-sequence cue) are, in some sense, likely to be conservative. Nevertheless, even with repetition priming effects controlled for, the comparison of G and B on p(respond) at training and the results for the first sequence trials analysis makes it clear that G is a strong go cue even without additional help from being tested within a sequence.

While the experiment provides clear evidence of learning, it also speaks to the debate surrounding the nature of this effect and whether it can ever be associatively driven. The lack of significant correlations between awareness and

task performance satisfies the zero-correlation criterion of Dienes et al. (1995), suggesting that the observed performance derived from an associative system that captured the contingencies between cues and outcomes, rather than an explicit system. Of course, this conclusion is not uncontroversial (see Newell & Shanks, 2014) and will be considered further in the general discussion.

## **3.5 EXPERIMENT 5**

Having seen how the inclusion of sequences impacts the learning of traffic light cues under a stop task set, the question now becomes what is learnt under a go task set. Therefore, Experiment 5 introduces the sequences described above into the methodology of Experiment 3. That is, the task set was go, with participants responding to coloured circles, and the instructions changed as before.

### **3.5.1 Method**

#### **3.5.1.1 Participants**

The sample size, inclusion criteria and outlier removal process were the same as in Experiment 4 (see Results section for details of those removed). Of the final sample, 47 were female, with an overall mean age of 20.34 (SD = 3.19).

#### **3.5.1.2 Design**

The design was identical to Experiment 4. However, as going is expected to be the effective outcome, + is now go and – stop. This means that the sequence Green 1 is now G+ -> G+, and the sequence Red 1 is now R- ->R- ->RA+ ->G+ (see Table 3.8 for the sequence runs used in Experiment 5).

Light type	Sequence
Green 1	G+ -> G+
Green 2	G- -> A- ->R- ->R- ->RA+ ->G+
Amber 1	A+ ->A+
Amber 2	A- ->R- ->R- -> RA+ ->G+
Red 1	R- ->R- ->RA+- ->G+
Red 2	R+ ->RA+ ->G+
Red/Amber	RA+ ->G+

Table 3.8. Sequence runs for Experiment 5 ‘-’ is 100% stop, ‘+’ is 100% go.

The filler cues were also amended to reflect the change in task set between Experiment 4 and 5, e.g., IP is now IP- rather than IP+ (see Table 3.9 for full design of filler cues).

Filler type	Cue	Occurrences per block
Cue G filler	B+	4
	B-	32
Cue R filler	I+	24
	I-	4
Cue RA filler	IP-	24
50:50 filler	P±	16
50:50 filler	YZ±	16
Tracking cue	J±	8

Table 3.9. Filler sequences for Experiment 5. ‘-’ is 100% stop, ‘±’ is 50/50 go, ‘+’ is 100% go. Each letter in the cue column represents a shape drawn randomly from a set pool.

### 3.5.1.3 Procedure

The procedure was identical to that of Experiment 4.

## 3.5.2 Analysis and results

The analysis for Experiment 5 was conducted in the same manner as for Experiment 4. In terms of outliers, nine participants were removed and replaced for having high omission errors, while three participants were removed and replaced for having high commission errors. There were no outliers for reaction times.

### 3.5.2.1 Training phase

Descriptive statistics for this phase are contained in Table 3.10.

#### 3.5.2.1.1 Reaction times

For the training data, the best fitting model had a conditional  $R^2$  of 0.87 (see Table 3.11 for AICs).

##### 3.5.2.1.1.1 Evidence of associative learning

As in Experiment 4 it was possible to provide evidence of associative learning by comparing performance between R vs. I and R vs. B. In effect, these contrasts are a manipulation check (and separate from the main analysis), and thus the alpha level was only corrected to .025. The R vs. I contrast was significant,  $t(440) = -2.41$ ,  $p = .017$ , 95% CI [-12.23, -1.25],  $d = -0.23$ , with I being faster than R, confirming that participants were indeed learning the contingencies present in the design. The contrast R vs. B enables me to compare cue R to an out and out stop cue, with this contrast being significant,  $t(440) = 4.63$ ,  $p = < .001$ , 95% CI [7.47, 18.45],  $d = 0.44$ , with cue R having faster responses, suggesting it was experienced as less of a stop cue than B.

##### 3.5.2.1.1.2 Traffic light contrasts

Noting the above results, I can be confident that the following results arise because of associative learning as well as repetition priming effects. As G vs. B is no longer a pure manipulation check, the alpha level was corrected to .008. The G vs. B contrast was highly significant,  $t(440) = 11.01$ ,  $p = < .001$ , 95% CI [25.34, 36.32],  $d = 1.05$ , with G being faster than B, thus confirming that participants were learning the contingencies present in the design. In terms of the experimental contrasts, A vs. B was highly significant,  $t(440) = 8.57$ ,  $p = < .001$ , 95% CI [18.52, 29.50],  $d = 0.82$ , with faster responses to cue A. The contrast A vs. G was not significant at the reduced alpha,  $t(440) = -2.44$ ,  $p = .015$ , 95% CI [-12.31, -1.33],  $d = -0.23$  but there was a trend for slower



responses to cue A than to cue G. The contrast for A vs. R was significant,  $t(440) = 3.94$ ,  $p = < .001$ , 95% CI [5.56, 16.54],  $d = 0.38$ , with faster reaction times to cue A. In terms of the R contrasts, the R vs. B contrast was significant,  $t(440) = 4.63$ ,  $p = < .001$ , 95% CI [7.47, 18.45],  $d = 0.44$ , with cue R having faster responses, suggesting it was experienced as less of a stop cue than B. The R vs. G contrast was highly significant,  $t(440) = -6.38$ ,  $p = < .001$ , 95% CI [-23.36, -12.38],  $d = -0.61$ , with responding in the presence of R being slower, indicating that cue G was being seen as more of a go cue.

### **3.5.2.1.2 p(respond)**

For this measure, the best model was a model that included the main effects of cue with a Gamma family and inverse link and random intercept (see Table 3.11).

#### **3.5.2.1.2.1 Evidence of associative learning**

Using the same logic outlined in Experiment 4, contrasts for commission errors for G vs. B and G vs. I were run to provide evidence of associative learning for this measure. Given that cue B is a stop cue, from an associative learning perspective it would be expected that cue G would have more errors than cue B and this was the case,  $z = 4.98$ ,  $p = < .001$ , confirming that the participants were learning the contingencies present in the design. The G vs. I contrast was run to establish if cue G was a go cue of similar power to a filler go cue. The contrast was almost significant,  $z = 2.23$ ,  $p = .026$ , with more errors to cue G than cue I indicating that cue G was certainly a go cue.

#### **3.5.2.1.2.2 Traffic light contrasts**

Given the result for G vs. B, I can be confident that the following results arise through associative learning as well as repetition priming effects. For this analysis, the alpha level was corrected to .008. The G vs. B contrast was significant,  $z = 4.98$ ,  $p = < .001$ , with more errors for cue G, thus confirming that participants were learning the contingencies present in the design. The A vs. B contrast was significant,  $z = 3.47$ ,  $p = < .001$ , with more errors for A than B, indicating that A was not responded to as a stop cue. The contrast for A vs. G was not significant at the reduced alpha level,  $z = -2.04$ ,  $p = .041$ , but numerically there were more errors for cue G than cue A, suggesting a tendency to experience cue A as not a strong go cue. The A vs. R contrast was

significant,  $z = 4.50$ ,  $p = < .001$ , with more errors for A than R, suggesting that A prompted more of a go response than cue R. Focusing on R, the R vs. B contrast was not significant,  $z = -1.26$ ,  $p = .208$ , but the R vs. G contrast was highly significant,  $z = -5.81$ ,  $p = < .001$ , with more errors for G than R, suggesting that R is not a go cue.

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
A-/+	377.00	38.83	0.03	0.02	0.01	0.01
B-	401.01	43.49	0.02	0.01	0.01	0.02
G+	370.18	39.59	0.04	0.05	0.01	0.00
I+	381.31	41.53	0.03	0.03	0.01	0.01
IP-	NA	NA	0.02	0.01	NA	NA
J±	388.07	43.42	0.03	0.04	0.01	0.02
P-/+	392.68	37.82	0.02	0.02	0.01	0.01
R-	388.05	34.03	0.01	0.01	0.01	0.01
RA+	381.04	41.67	NA	NA	0.01	0.00
YZ-/+	386.44	41.7	0.03	0.02	0.01	0.01

Table 3.10. Training descriptive statistics for Experiment 5. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 3.5.2.1.3 Training phase summary

Experiment 5 allowed for the exploration of associative learning in an experiment that used a go task set and considered some of the sequences experienced at traffic lights. The analysis suggests that cue G primed a go response, cue A was somewhat of a go response (seemingly promoting more of a go response than cue B on both measures but not as strong a go response as cue G), and cue R somewhat of a stop cue, having slower reaction times and less errors than cue G but also faster reaction times than cue B. Additionally,

the results for the associative learning contrasts demonstrate that associative learning is occurring in this phase.

Model	Training	1 <sup>st</sup> Sequence	Other sequence
Reaction time			
Main effects of cue	5168.32	6050.18	6004.79
Main effects of cue with random intercept	<b>4356.62</b>	<b>5704.71</b>	<b>5572.67</b>
p(respond)			
Main effects of cue with Gamma family and inverse link	-2923.04	-3174.77	-3538.25
Main effects of cue with Gamma family and inverse link and random intercept	<b>-3061.52</b>	<b>-3277.10</b>	<b>-3636.21</b>

Table 3.11. AICc scores for models run for Experiment 4 run on reaction time and p(respond) data at training and test. Bold are the models chosen.

### 3.5.2.2 Test analysis

Similarly, to Experiment 4, two separate analyses were undertaken for the test phase. The first focused on first trials of the sequences in conjunction with filler cues, while the second only focused on cues within sequences (see Table 3.12 for descriptive statistics).

#### 3.5.2.2.1 First trial analyses

The analyses conducted on the first trial of each sequence were a G vs. B contrast (which is now once again a manipulation check) and A vs. B, A vs. G, A vs. R, R vs. B and R vs. G. The alpha level for these analyses was corrected to .010 (see Table 3.11 for AICc's).

### 3.5.2.2.1.1 Reaction times

The model chosen had a conditional  $R^2$  of 0.56. The G vs. B contrast was significant,  $t(495) = 2.71$ ,  $p = .007$ , 95% CI [5.09, 31.67],  $d = 0.24$  indicating that learning had taken place by the test phase. The A vs. B contrast was significant,  $t(495) = 3.88$ ,  $p < .001$ , 95% CI [13.03, 39.62],  $d = 0.35$ , suggesting that cue A was not a stop cue. The A vs. G contrast was not significant,  $t(495) = 1.17$ ,  $p = .242$ , 95% CI [-5.35, 21.24],  $d = 0.11$ . The R vs. B contrast was not significant,  $t(495) = -0.02$ ,  $p = .982$ , 95% CI [-13.45, 13.14],  $d = -0.002$ , and cue R had significantly slower reaction times than cue G,  $t(495) = -2.73$ ,  $p = .006$ , 95% CI [-31.83, -5.24],  $d = -0.25$  suggesting that cue R was not a go cue. The contrast A vs. R was also significant,  $t(495) = 3.90$ ,  $p < .001$ , 95% CI [13.19, 39.78],  $d = 0.35$  suggesting that cue A primed more of a go response than cue R.

### 3.5.2.2.1.2 p(respond)

For commission errors there was a significant difference for G vs. B,  $z = 3.16$ ,  $p = .002$ , with more errors for G than B suggesting that learning had occurred for this measure. The contrast A vs. B was not significant,  $z = 0.63$ ,  $p = .526$ . The A vs. G contrast was significant,  $z = -2.67$ ,  $p = .008$ , with more errors for cue G than A indicating that cue A was not a go cue. For cue R, the contrast against cue B was marginally significant at a standard alpha level,  $z = 1.67$ ,  $p = .077$ , with a trend for more errors to cue R than B. The contrast of R vs. G approached being marginally significant at a standard alpha level,  $z = -1.65$ ,  $p = .100$ , with a trend for more errors to cue G than R. This suggests that cue R for commission errors was a neutral or weak go cue. The contrast for A vs. R was not significant,  $z = -1.17$ ,  $p = .244$ .

### 3.5.2.2.2 Other sequence trial analyses

These focused on comparing the traffic light cues within the sequences. Therefore, contrasts undertaken were A vs. G, A vs. R, and R vs. G. The alpha level was corrected to .017 (see Table 3.11 for AICc's).

#### 3.5.2.2.2.1 Reaction times

The model chosen had a conditional  $R^2$  of 0.66. The A vs. G contrast was marginally significant at a standard alpha,  $t(495) = -1.71$ ,  $p = .088$ , 95% CI [-21.75, 1.47],  $d = -0.15$ , with faster reaction times to cue G than A. The contrast A vs. R was not significant,  $t(495) = -0.79$ ,  $p = .432$ , 95% CI [-16.27, 6.96],  $d = -$

0.07. The comparison between R vs. G was also not significant,  $t(495) = -0.93$ ,  $p = .355$ , 95% CI [-17.09, 6.13],  $d = -0.08$ .

#### **3.5.2.2.2 p(respond)**

The A vs. G contrast was not significant,  $z = -0.31$ ,  $p = .760$ , as was the contrast between A and R,  $z = 1.47$ ,  $p = .143$ . The R vs. G contrast was marginally significant at a standard alpha level,  $z = -1.76$ ,  $p = .079$ , with a trend for more errors to cue G than cue R.

#### **3.5.2.2.3 Test phase summary**

The first sequence trial analyses indicate that at a purely associative learning level, participants had learnt that G was a go cue (having significantly faster reaction times and more errors to cue G than cue B). Cue A seemed to prime an overall neutral response, priming a go response for reaction times (significantly faster than B but similar to G) but a stop response for commission errors (significantly less errors than G and similar levels to B). Cue R seemed to prime a weak stop response overall, being roughly neutral for commission errors but significantly slower than cue G in terms of response times.

The other sequence trial analyses show weaker effects overall. The limited evidence suggests that cues A and R promoted more of a stop response compared to cue G, with cue A having marginally significantly slower reaction times and cue R marginally significantly less errors.

Type of cues	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Filler cues						
B+	401.03	45.47	0.02	0.02	0.01	0.01
I-	380.95	51.34	0.03	0.04	0.02	0.02
IP+	396.19	45.58	0.02	0.04	0.01	0.02
J±	392.51	59.52	0.01	0.03	0.01	0.03
P-/+	402.58	61.21	0.02	0.04	0.01	0.00
YZ-/+	381.36	46.05	0.01	0.02	0.01	0.02
1st sequence						
trials						
experimental						
cues						
A	374.70	43.75	0.02	0.05	0.01	0.00
G	382.65	63.78	0.03	0.08	0.02	0.05
R	401.18	55.08	0.02	0.06	0.02	0.05
RA	371.73	54.55	0.07	0.11	0.01	0.03
Other sequence						
trials						
experimental						
cues						
A	369.72	60.23	0.01	0.03	0.01	0.00
G	359.58	42.79	0.01	0.02	0.02	0.05
R	365.06	45.10	0.01	0.01	0.02	0.05
RA	359.48	48.01	0.02	0.04	0.01	0.03

Table 3.12. Test descriptive statistics for Experiment 5. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

### 3.5.2.3 Expectancies

I conducted the expectancy analysis for Experiment 5 in the same manner as described in Experiment 4 (see Table 3.13 for descriptive statistics). The alpha level was corrected to .008. The best model fitted the main effects of cue with random intercept ( $R^2 = 0.25$ ). One participant did not provide any ratings and as

such the degrees of freedom are smaller for this analysis than for its companion in Experiment 4. The G vs. B contrast was significant,  $t(594) = 4.16$ ,  $p < .001$ , 95% CI [1.39, 3.88],  $d = 0.34$ , with participants more likely to rate G as a go cue than B. The A vs. B contrast was not significant at the reduced alpha,  $t(594) = 2.24$ ,  $p = .026$ , 95% CI [0.18, 2.66],  $d = 0.18$ , though there was a trend for participants to rate A as more go than B. The A vs. G contrast was marginally significant at a standard alpha,  $t(594) = -1.92$ ,  $p = .055$ , 95% CI [-2.46, 0.02],  $d = -0.16$  hinting at a trend for cue G to be rated as go over cue A. The A vs. R contrast was not significant,  $t(594) = 0.26$ ,  $p = .796$ , 95% CI [-1.08, 1.41],  $d = 0.02$ . Focusing on cue R, R vs. G was not significant at the reduced alpha,  $t(594) = -2.18$ ,  $p = .030$ , 95% CI [-2.62, -0.14],  $d = -0.18$ , though there was a tendency for participants to rate cue G as more go than cue R. The R vs. B contrast was also not significant at the reduced alpha,  $t(594) = 1.98$ ,  $p = .048$ , 95% CI [0.01, 2.50],  $d = 0.16$ , though numerically participants did rate R as more go than B. Unlike in Experiment 4 the awareness ratings are less clear. Overall, it seems that awareness is centred on G vs. B, with perhaps the participant's awareness of this contrast driving the rest of the pattern of results. There was some evidence that participants were aware that stopping was more likely in response to A and R than to G, and some awareness that stopping was even more likely to B.

As in Experiment 4, comparing J and YZ against B and G allowed me to see if B was experienced as a relatively neutral cue or not. This analysis was separate from that conducted above and so an alpha level of .013 was applied. The contrast J vs. G was significant,  $t(594) = -2.90$ ,  $p = .004$ , 95% CI [-3.08, -0.59],  $d = -0.24$ , with participants more likely to rate G as go over J. The contrast J vs. B was not significant,  $t(594) = 1.26$ ,  $p = .207$ , 95% CI [-0.44, 2.04],  $d = 0.10$ . The contrast YZ vs G was significant,  $t(594) = -2.81$ ,  $p = .005$ , 95% CI [-3.02, -0.54],  $b = -0.23$  with cue G being rated as more go than cue YZ. The contrast for YZ vs. B was not significant,  $t(594) = 1.35$ ,  $p = .179$ , 95% CI [-0.39, 2.10],  $b = 0.11$ . These results suggest that participants expected to make very similar responses to cue B (a stop cue) and cues J and YZ (neutral cues). Overall, the analysis for B vs. J and YZ suggests that the ratings are very close, and that B seems to be somewhat neutral in terms of awareness. This neutrality is

dissociated from the response time results which indicate that B is certainly not a neutral cue, but is in fact a stop cue in comparison to the go cue, cue G.

Cue	Expectancy Rating	
	Mean	SD
A-/+	0.84	3.54
B-	-0.58	4.04
G+	2.05	3.71
I+	0.87	3.52
IP-	-0.09	3.54
J±	0.22	3.99
P-/+	0.02	3.88
R-	0.67	3.15
RA+	0.84	3.48
Y-/+	-1.25	4.01
YZ-/+	0.27	3.62
Z-/+	-1.31	4.19

Table 3.13. Descriptive statistics for expectancy ratings for Experiment 5.

### 3.5.2.4 Correlational analysis

As in Experiment 4 it is possible to undertake a correlational analysis to explore if awareness is driving performance. For this experiment I will only report the correlation for G vs. B, as this is the contrast that expectancies significantly differed on. Thus, the analyses performed were for reaction times and commission errors at training and for the first sequence trial data. As seen in Figure 3.3 correlations were close to zero, indicating that awareness did not drive performance. This suggests that the behaviour was caused by associatively-mediated learning, rather than propositional processes, though with some caveats which will be discussed later.



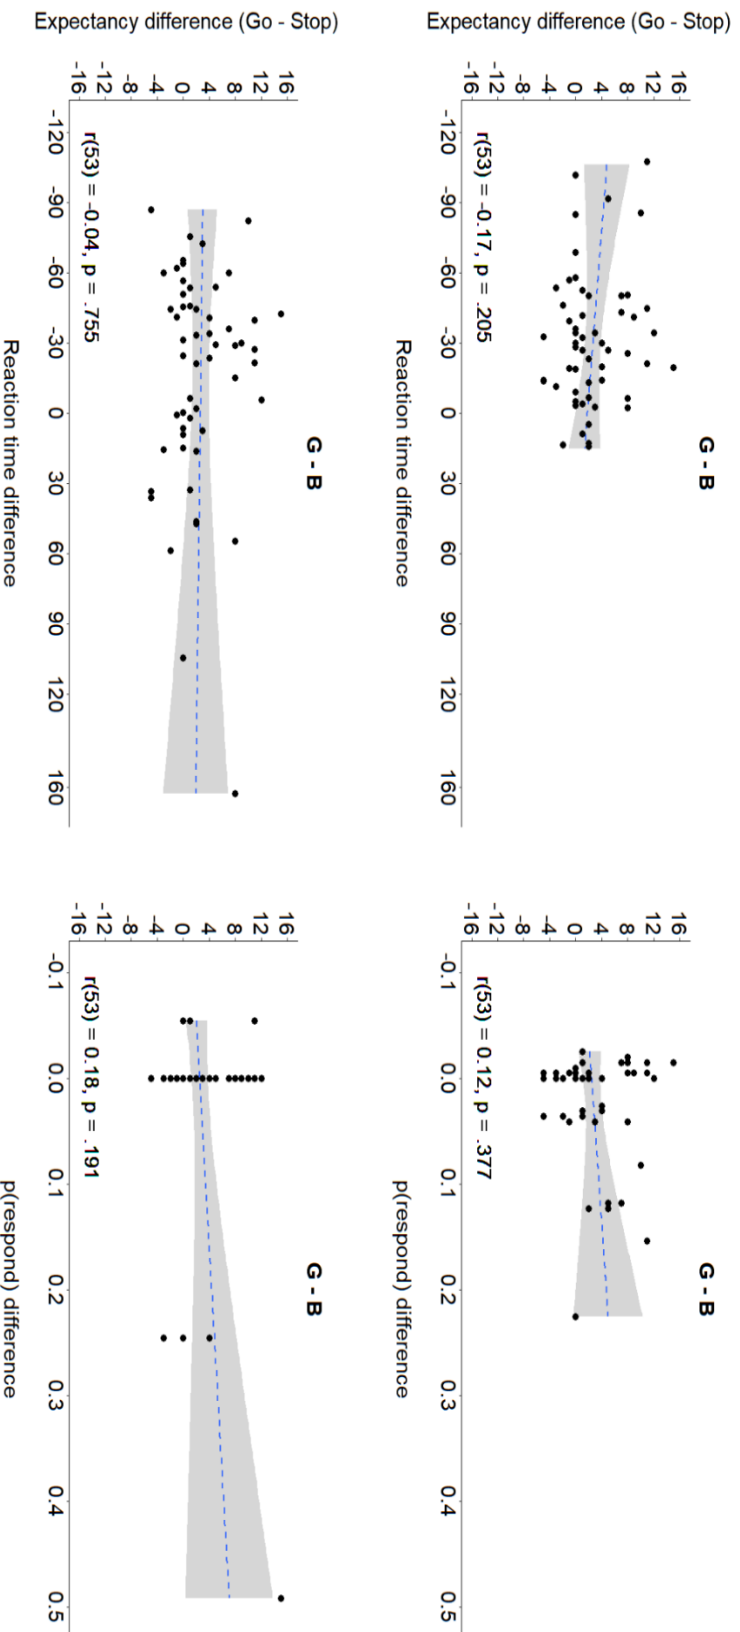


Figure 3.3. Experiment 5 correlations between reaction times and commission errors for training (top panel) and test (bottom panel) for the comparison G vs. B. Scores were calculated by taking the stimuli that was more contingent with stopping from the one that was most contingent with going (i.e. go-stop). Dashed line depicts linear model of the data, shaded area represents the 95% confidence interval of that model. R values refer to Pearson's product-moment correlation.

### 3.5.3 Summary

Experiment 5 explored how the sequences added in Experiment 4 would affect learning in a design that has go as the effective outcome. During training on both measures the data showed that cue G primed more of a go response than cue B – thus learning was occurring as expected. Additionally, there was evidence of behaviour occurring through learning, as well as through repetition priming effects.

For cue A the reaction time data at training suggests it primed a weak go response. The cue was significantly faster than cue B yet, while not significant at the corrected alpha, there was a strong trend for faster reaction times to cue G than A. The results for commission errors indicate that cue A was again a weak go cue. Errors were greater than cue B, but errors were also significantly less than cue G at a standard alpha level. Overall, it seems that cue A was experienced as something of a weak go cue at training. Focusing on cue R, the fact that at training it was significantly slower and had significantly less commission errors compared to cue G indicates that R was not a go cue. For reaction times it does not seem to be that strong a stop cue either having significantly faster reaction times than cue B. However, R had similar levels of errors compared to B. Overall, it would be fair to describe cue R as a weak stop cue.

For the first sequence test trials, as expected the contrast G vs. B was significant for both performance measures suggesting that learning had occurred. The results indicate that cue A primed a fairly neutral response. It was significantly different for cue B at reaction times but not different for commission errors. This pattern was reversed when compared to cue G; the two cues were similar in reaction time speeds, but for  $p(\text{respond})$  cue A had significantly fewer errors. For cue R, the data indicates it was somewhat of a stop cue. It had significantly slower response times than cue G, and the difference at  $p(\text{respond})$  was marginally significant at a standard alpha level (with more errors for G). Compared to cue B, response times were similar, but cue R did have marginally significantly more errors than cue B at a standard alpha level. In summary, when in a task set that has going as the effective outcome and sequences are in play cue A seemed to prime a weak go response, with cue R seeming to promote a weak stop response. At test, in an analysis where repetition priming

effects were controlled for, cue A seemed to be linked with a neutral response and cue R a stop response.

The pattern of results is similar to the companion study Experiment 3. In Experiment 3 cue R primed a stop response and cue A a neutral to weak stop response. While comparatively cue R is perhaps less of a stop cue, and A more of a go cue the rough pattern for the placement of the cues along a go/stop continuum in a go task set holds. It should be noted that the pattern of results is easier to interpret in Experiment 4 (especially for commission errors), with this likely being a product of the enhanced learning generated by sequential learning.

Finally, the results again support the proposition that sequence learning can occur through associatively-mediated processes.

### **3.6 JOINT ANALYSIS OF EXPERIMENT 4 AND EXPERIMENT 5**

As Experiments 4 and 5 are mirror opposites in much the way that Experiments 2 and 3 were, it is possible to combine the two and undertake a between-participants analysis to investigate changes across the studies. While in Chapter 2 I presented a joint analysis of the traffic light cues, to begin to capture a more complete understanding of task set, given the similarity of the experiments presented here to their comparison experiments in Chapter 2 such an analysis will not be undertaken. However, I do report an analysis of the feature-positive effect between Experiments 4 and 5. As these analyses are manipulation checks a standard alpha level was applied.

#### **3.6.1 Joint feature-positive analysis**

In Chapter 2 I conducted a *t*-test looking at R vs. RA against IP vs. I, i.e. of the feature-positive vs feature-negative across Experiments 2 and 3. However, an equivalent analysis is not possible here. This is because IP vs. I is a non-sequence contrast, while R vs. RA is. Instead, the approach used here is to compare the difference in R vs. RA across the two experiments. This approach also controls for repetition priming effects, as both experiments were exposed to the same sequence runs but with opposite task sets. In Experiment 5, R vs. RA is the feature-positive contrast (R- RA+), while in Experiment 4 it is the feature-negative pair (R+, RA-). Based on my past findings, I would expect learning of

these discriminations to be better in Experiment 5 than 4. As I was not able to perform such an analysis within experiments, and the evidence from Experiments 2 and 3 was that the effect was fairly difficult to detect, the contrast was undertaken for both dependent measures at training and test.

The result from response times at training were significant,  $t(110) = -2.32$ ,  $p = .022$ , 95% CI [-17.25, -1.35],  $d = -0.44$ , with enhanced learning to RA compared to R in Experiment 5 (mean difference of 7.01, SD = 22.43) compared to Experiment 4 (M = -2.29, SD = 19.93). This demonstrates that the changes made between the two experiments successfully altered the nature of the discriminations experienced by participants in that the effective outcome changed from stopping to going. However, the results from test were not significant,  $t(110) = -1.12$ ,  $p = .266$ , 95% CI [-36.67, 10.22],  $d = -0.21$ . It was not possible to conduct this analysis on p(respond) in training since there were no errors of commission for RA in either experiment because RA was 100% go. However, at test this contrast was significant,  $t(110) = 2.32$ ,  $p = .023$ , 95% CI [0.007, 0.08],  $d = 0.44$ , with once again better learning of the R vs. RA contrast in Experiment 5 (M = -0.04, SD = 0.12) compared to Experiment 4 (0.00, SD = 0.07).

Therefore, these analyses provide good evidence that the contrast R-RA was significantly different between Experiment 4 (outcome is stop) and Experiment 5 (outcome is go) at both training for the response time data and in commission errors at test. These findings support my results from Chapter 2 and those of Bowditch (2016), and provide further evidence that the manipulations employed in the paradigm are effective in changing task set.

### **3.7 GENERAL DISCUSSION**

Fundamentally, the results presented here support those of Chapter 2 and a wider body of literature (Best et al., 2016; Bowditch et al., 2016) which demonstrate that a cue that is paired consistently with a particular outcome promotes that outcome, even when the primed action is no longer required. Such support was found in reaction time performance for the first sequential trials analysis for both experiments, and for commission errors in Experiment 5. The key aspect of this chapter was the addition of sequences and the subsequent discussion will explore the effects of sequences.

### **3.7.1 Sequence learning**

At the introduction to this chapter I discussed how sequences can enhance the learning of cues. Given this, it might be assumed that the inclusion of sequences would lead to quantitatively different results from those reported in Chapter 2 (bigger effect sizes for example), but that qualitatively the actual pattern of the results would resemble the companion experiments in Chapter 2. There is work consistent with this notion. For example, Gotler, Meiran, and Tzelgov (2003) used a task-switching paradigm with two task sets; task 1 required participants to respond if the target was up or down, while task 2 required participants to respond if the target was right or left. These tasks were presented in sequence blocks and random blocks. Of note for the current discussion, the results demonstrated that the type of block (sequence or random) did not interact with switching costs (the loss in performance caused by the switch from one task to another; Rogers & Monsell, 1995). That is, though sequences did enhance learning overall, with faster reaction times in sequence blocks compared to random blocks, this effect was separate from task-switching costs, with the enhanced learning granted by sequences being equal for switch and non-switch trials. Broadly, the results here support this position, with the two experiments in this chapter having similar findings to their companion experiments in Chapter 2. This of course raises the question as to the need for Experiments 4 and 5. However, as highlighted in the introduction sequences are a key part of human and learning, and as discussed in Chapter 1 the traffic light sequence is a ubiquitous feature of life. Therefore, to design a paradigm with the express aim of capturing the contingencies experienced at UK traffic lights without including the sequences involved would be to ignore a key feature of what makes traffic lights, traffic lights.

### **3.7.2 Feature-positive effect**

To my knowledge the work presented here is the first to explore learning of the feature-positive effect in the context of sequence learning. A priori, one would expect that sequences would enhance rather than diminish the feature-positive effect. However, the contrasts undertaken in Chapter 2 and 3 are not identical and do not allow for direct testing of this assumption. Yet it is possible to undertake the feature-positive analysis performed in this Chapter using the data in Chapter 2. This analysis (presented in Appendix C) supports the argument

that sequences enhance learning of the feature-positive effect, with bigger effect sizes for the sequence versions of the contrasts.

### **3.7.3 Learning and awareness**

Both the experiments conducted in this chapter indicated that awareness of the contingencies was not significantly correlated with task performance. However, some of the correlations were non-zero. This could be taken to indicate the presence of some conscious content (Seth, Dienes, Cleeremans, Overgaard, & Pessoa, 2008), but these correlations were never significant. It seems that cues within sequences can be learnt about associatively, as well as via conscious processes. The lack of correlation seen between performance and awareness would be difficult to explain through a single learning system and so supports the notion of dual-route models (McLaren et al., 2014). However, it must be noted that there are caveats to this claim. The first is that the awareness test was given at the end of the experiment and so falls prey to the ‘immediacy’ criterion for assessing awareness by Newell and Shanks (2014). The authors argue that assessments should be made online (as in the Perruchet experiments) to prevent forgetting or interference from subsequent trials. Indeed, as the 50/50 test phase is, in effect, an extinction phase, it is likely that it would have degraded the contingencies with participants perhaps forgetting the contingencies from training. As it stands it is not possible to say that those who made more accurate ratings (i.e. rated cue G as a go cue) were simply better at remembering cues, compared to those who were less accurate, or if the ratings measure true awareness. In defence of the awareness measure used it does meet the sensitivity criterion, as the same cues were used to assess awareness as seen in the task.

Another issue with the awareness test used is that there are other forms of meta-knowledge it does not measure. For example it does not measure what Dienes and Perner (1996, 1999) refer to as ‘content explicitness’, defined as knowing that one is in possession of knowledge. This is distinct from the meta-knowledge assessed in the expectancy rating task which requires participants to represent content towards a cue as knowledge about that cue rather than confabulation (Dienes & Altmann, 1997), so called ‘attitude explicitness’ (Dienes & Perner, 1996). As Reber (as cited in Dienes & Altmann, 1997) posited, participants “may know that they know something, even though they may not

know what it was that they know” (p.136). Therefore, to provide a more definitive case for associative learning it would be necessary to design an awareness test that addresses both these forms of metacognition.

### **3.7.4 Conclusion**

To conclude, the experiments presented here continue the development of a paradigm exploring contingency learning at UK traffic lights. The addition of the sequences seems to be a key enhancer of task performance, and results in larger learning effects. Crucially, the pattern of results in both stopping and going task sets matches that already established in Chapter 2. With cue A being a weak go cue in a stop task set, and a neutral/weak stop cue in a go task set. This suggests that, depending on the task set in play, associative learning at traffic lights could be such as to encourage dangerous driving. This supports the idea raised in Chapter 1 of the need to develop interventions to target these maladaptive consequences of associative processes. Having seen how sequences can enhance the learning experienced at traffic lights, the next question is whether the between-experiment manipulation of task set used so far captures the fullest experience of traffic lights. In real-world driving, task set will change depending on the light shown. To reflect this, the next experiment aims to incorporate both task sets into one experimental design to complete my investigation of the associative consequences of experiencing the contingencies between traffic light signals and response outcomes.





### ASSOCIATIVE PROCESSES III: THE ROLE OF TASK SET IN CONTINGENCY LEARNING

Throughout this thesis I have highlighted the importance of ‘task set’. This chapter defines exactly what I mean by this and summarises the relevant research, before discussing how the existing literature could inform the design of a within-participants experiment manipulating task set based on the traffic light paradigm developed thus far.

#### 4.1 TASK SET

A ‘task set’ is a configuration of cognitive resources that are required and maintained to complete a task (Monsell, 1996, 2003). In many experiments, including the ones presented in this thesis, task set is established through instructions. As Sakai (2008) p. 219, notes “participants heed the instructions and prepare for the experiment. The participants may remember the instructions by verbally rehearsing them, but after practice for several trials, the task information is maintained as a configuration of perceptual, attentional, mnemonic, and motor processes necessary to perform the task”. In understanding how task sets come to affect behaviour, Meiran (2000) argues this happens through four steps. Firstly, the task set must be configured (e.g., through instructions). Secondly, the information is applied as a mental representation. Next, a process the author refers to as ‘similarity matching’ occurs, whereby the target stimuli facing a participant are compared with the representation of the responses. Through this process each response (e.g., two different button presses in a 2-choice reaction time task) acquires ‘potency’, which is determined by the degree of similarity between the task set response and the stimulus. The final stage is called ‘response decision’. Here the potencies between each response are compared and the response with the highest potency is selected. More generally, the formation of a task set can be seen as two distinct processes. Firstly, there is a preparation stage where the rules of a task set need to be activated, and secondly interference from competing task sets needs to be inhibited (Mayr, Diedrichsen, Ivry, & Keele, 2006; Kiesel et al., 2010).

## 4.2 THE TASK-CUEING PARADIGM

A full review of the many paradigms used in the task set literature is outside the scope of this thesis. Rather, the discussion focuses on the task-cueing paradigm (Sudevan & Taylor, 1987; Meiran, 1996; Monsell & Mizon, 2006) given the parallels with the issues this chapter wishes to address. In experiments based on this paradigm, participants are informed at the start about the rules of the task and the cues that signal the need to implement these rules. Then each trial is preceded by a cue which indicates the type of response required by participants to the subsequent stimulus. Participants must respond quickly and accurately, as in a typical reaction time task. The task continues in this manner with cues preceding stimuli. Typically, there are two tasks to perform with a different cue indicating a different task set, with participants having to switch task sets throughout the experiment. Performance is dependent on participants applying the correct task set based on the cue to the stimuli presented on that trial (Li, Li, Liu, Lages, & Stoet, 2019; McLaren et al., 2019). The exact nature of the task set can vary amongst experiments, from requiring participants to classify rectangles by height or width depending on whether the preceding cue was *h* or *w* respectively (Altmann, 2004), or to classify a single digit as either odd or even depending on the colour of a preceding disk (Monsell, Sumner, & Waters, 2003).

One design of a task-cueing experiment that provides a starting point for the discussion of Experiment 6 can be found in the work by Meier, Lea, Forrest, Angerer, and McLaren (2013). Because the authors wanted to compare the behaviour of human participants against that of pigeons, they used different colours to indicate task set rather than language-based cues. Thus, in the experiment a blue or yellow circle in the centre of the display required participants to respond with task A, while red or green circles required participants to respond with task B. In this way, the colour of the circle informed participants of the correct response towards the subsequent stimuli. The results and exact purpose of the work by Meier et al. (2013) are not relevant to this discussion. It is sufficient to say that the human participants were able to perform the task - highlighting how simple coloured shapes are enough to cue task-sets.

### **4.3 PRESENT EXPERIMENTS**

Throughout this thesis task set has been an important factor in the design of the experiments. Indeed, given that different traffic lights are likely to be experienced within different task sets, this factor is fundamental to the experience of traffic lights. Yet, the designs employed thus far have only enabled comparisons of task set at between participants' levels. Experiment 6 explores task set using a within-participants design. As seen in the introduction, the task-cueing paradigm can allow the changing of task set within an experiment using only simple shapes. To my knowledge, these experiments will be the first to apply such a paradigm to the learning environment experienced at UK traffic lights. Note that the experiments reported so far in this thesis did prepare the ground for the development of the paradigm to capture more real-world learning; in Experiments 4 and 5, there were two sequences per traffic light. For example, in Green 1, G was + mimicking the situation when the lights have recently changed to green, and thus it is likely a driver would have time to cross the junction. Whereas in Green 2, G was -, modelling the situation when the lights have been green for some time and so a driver could be expecting them to change to amber and so might be more cautious around them.

### **4.4 EXPERIMENT 6**

Capturing the task set used for each individual traffic light is an important step in obtaining a fuller picture of contingency learning at UK traffic lights. If a driver is approaching traffic lights which are currently on green, they will know they have some (unpredictable) time to cross the junction, thus their default will be go and their task set (that is, the signalled response they are looking for) will be stop. However, if the driver is waiting at a red light then they will expect the light to change to red and amber soon, and so their default will be stop (with a task set of go). The challenge for Experiment 6 was to capture this combination of traffic light changes and the appropriate task set in the laboratory. My previous experiments have demonstrated how switching the required response to the coloured circles can lead to a change in task set, so this manipulation was indicated. Informed by the design of Meier et al. (2013), a cued conditional feature was added to indicate the appropriate response to coloured circles. While using the approach taken by Meier et al. (2013) and having coloured

circles appear in the middle of the screen would build upon this past work, it would not be consistent with the design of the paradigm developed thus far, as it would disrupt presentation of the fixation bar and the shapes. Therefore, the approach taken here was to add another signal on the screen to let participants know the correct response whilst not interfering with the existing display. A square outline around the centre display was used to cue the task. A dark grey square (RGB code 64, 64, 64) surrounding the cue and circle indicated a go task (respond to coloured circles) whereas a light grey square (RGB code 191, 191, 191) designated a stop task (do not respond to coloured circles). Therefore, in Figure 4.1 participants were required to respond to the right-hand image and withhold a response to the left-hand image.




Figure 4.1. Conditional cues used in Experiment 6.

As well as task set, Chapter 3 showed that the sequence of stimuli experienced at traffic lights was itself an important source of learning, and therefore Experiment 6 used the same sequences, but with the addition of the aforementioned conditional cue. Therefore, the response to each trial depended on the colour of the enclosing square, which correlated with the sequence. Table 4.1 represents the same sequences as Table 3.1 in Chapter 3, except that the cues are now colour coded to reflect the colour of the square that will appear with them. In addition, now that task set is decoupled from response, + equals go and - stop. In Table 4.1, **bold dark grey** means that participants will see a dark grey square (and thus will respond to coloured circles, which means the default will be stop and task set is go) while **bold light grey** means that participants will see a light grey square (and thus will be in a stop task set as participants will have to not respond to coloured circles and respond to white circles).

Light type	Sequence
Green 1	G+ -> G+
Green 2	G- -> A- -> R- -> R- -> RA+ -> G+
Amber 1	A+ -> A+
Amber 2	A- -> R- -> R- -> RA+ ->G+
Red 1	R- -> R- -> RA+ ->G+
Red 2	R+ -> RA+ -> G+
Red/Amber	RA+ -> G+

Table 4.1. Sequence runs for Experiment 6 '-' is 100% stop, '+' is 100% go.

As can be seen in Table 4.1, green traffic lights are always accompanied by a light grey square (default go, task set stop). This is because it is assumed that initially green traffic lights will always induce a default go response, and thus drivers will be looking for stop signs. This helps to correct an issue with the design of the experiments in Chapter 3, whereby the sequence Green 2 did not have a preceding green cue before the green stop cue to mimic the shift in driver's perception of G being a go cue to a stop cue. However, the interaction between task set and go/no-go response sets up a situation where participants see a green cue with the default to respond but in fact are required to withhold their response. The opposite holds true for R+ in the Red 2 sequence. In this sequence R+ is designed to mimic the situation where a driver has come around a corner, seen the lights are on red and approached while the lights are still red. The driver knows the lights must change to red and amber shortly and might be assumed not to brake in the approach towards the light. In this instance R acts a go cue. However, red is likely to still engage a default stop response and so is accompanied by a dark grey square.

As before the filler cues need to balance out the above design. The addition of the conditional cueing means that this also needs to be controlled. This raises another issue for the experimental design: is it more important that the experiment mimics real-life or that it is a balanced design? If the former is key, the filler cues would need to have the opposite conditional cueing to cues in the

sequence, e.g., cue B would have a dark grey square cue. This would mean that the conditional cue (grey outline square) would also signal what type of trial (+ or -) was likely to happen, with the light grey square (that cue linked to G) indicating that a trial is very likely to be a go trial. This design would create a situation where one is being cued for situations where one is mostly likely to go, or to stop – which is what traffic lights in some way do. If on the other hand a balanced design is most important, then cue B would need to be accompanied by a light grey square. In other words, each filler cue has the same conditional cue as its main counterpart. The conditionality then cannot indicate the response that is likely to be needed. To ensure that Experiment 6 was consistent with my previous experiments, it was decided to use a design that was balanced (see Table 4.2). It will be noted that cues J and YZ appear equally in a stop and go task set. This was to prevent a cue becoming linked to a specific colour circle. For example, if cue J was not split by the conditionality feature and response as shown, then a feasible design would be J- accompanied by a dark grey conditional cue, and J+ accompanied by a light grey conditional cue. In this case, while the overall response would be balanced (4 go and 4 stop), participants would only see white circles when this cue appeared. This would be because with J- (stop trials) white circles would be the stop cue, while for J+ (change in task set) white circles would require a go response (the same would also hold true for YZ).

Cue type	Cue	Occurrences per block
Cue G filler	B+	4
	B-	32
Cue R filler	I+	24
	I-	2
Cue RA filler	IP-	24
50:50 filler	P-	8
	P+	8
50:50 filler	YZ-	4
	YZ+	4
	YZ-	4
	YZ+	4
Tracking cue	J-	2
	J+	2
	J-	2
	J+	2

Table 4.2. Filler sequences runs for Experiment 6. '-' is 100% stop, '+' is 100% go.

#### 4.4.1 Method

##### 4.4.1.1 Participants

The inclusion criteria and outlier removal process were the same as Experiment 2 (see Results section for details of those removed). A power analysis using the R package SIMR (Green & MacLeod, 2016) and the data from Experiments 4 and 5 suggested that 50 participants would give a 84.30% chance of detecting a 20.37ms difference between a go vs. stop cue (G vs. B, see section 4.4.1.2 for details) for the first sequence trial analysis for response times (the average difference from G vs B for those two experiments). This sample size is also in line with other recent similar experiments, such as Forrest et al. (2014). Of the final sample, 44 were female with an overall mean age of 20.00 (SD = 3.09). Participants received payment of £5 or one course credit.

##### 4.4.1.2 Design

The design (see Table 4.3) was similar to that of Experiment 4 and 5 except for conditionality, as discussed above, and the calibration block. In past

experiments the calibration block comprised 24 trials. However, in order to allow participants practice at the conditional nature of the current task the calibration block was extended to include 36 trials, split into three runs of six trials for each task set with these runs ordered randomly. To reflect the new nature of the first block in Table 4.3 it is renamed as a Learning Phase. The decision to split J by conditionality and response for training and test blocks also meshes with the change to the calibration block. If the cue had not been modulated by these factors, then participants would have experienced a split design in the learning phase but not for the training or test phases. This could have led to participants changing their performance to cue J in the subsequent blocks, which in turn could have disrupted the tracking procedure. Another feature of the design worth highlighting is that in more traditional cued task-switching designs, trials (and therefore responses) are randomly ordered. However, in the design employed for Experiment 6 the order in each sequence is fixed, but the order of the sequences was random.

Phase	Blocks	Individual trials per block	Design
Learning phase	1	36	J
Training	6	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16)
Test	1	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16)

Table 4.3. Summary of Experiment 6 design. Letters represent a sequence. The numbers in parentheses indicate how often the sequence will occur per block. At test the contingencies were  $\pm$ .



#### 4.4.1.3 Procedure

The procedure was identical to that of Experiments 4 and 5 apart from changes to the instructions and the addition of the cueing square. The task was designed so that a cueing square was always on screen, appearing with the fixation bar for that trial and then changing (or not) with presentation of the next fixation bar (see Figure 4.2 for schematic).

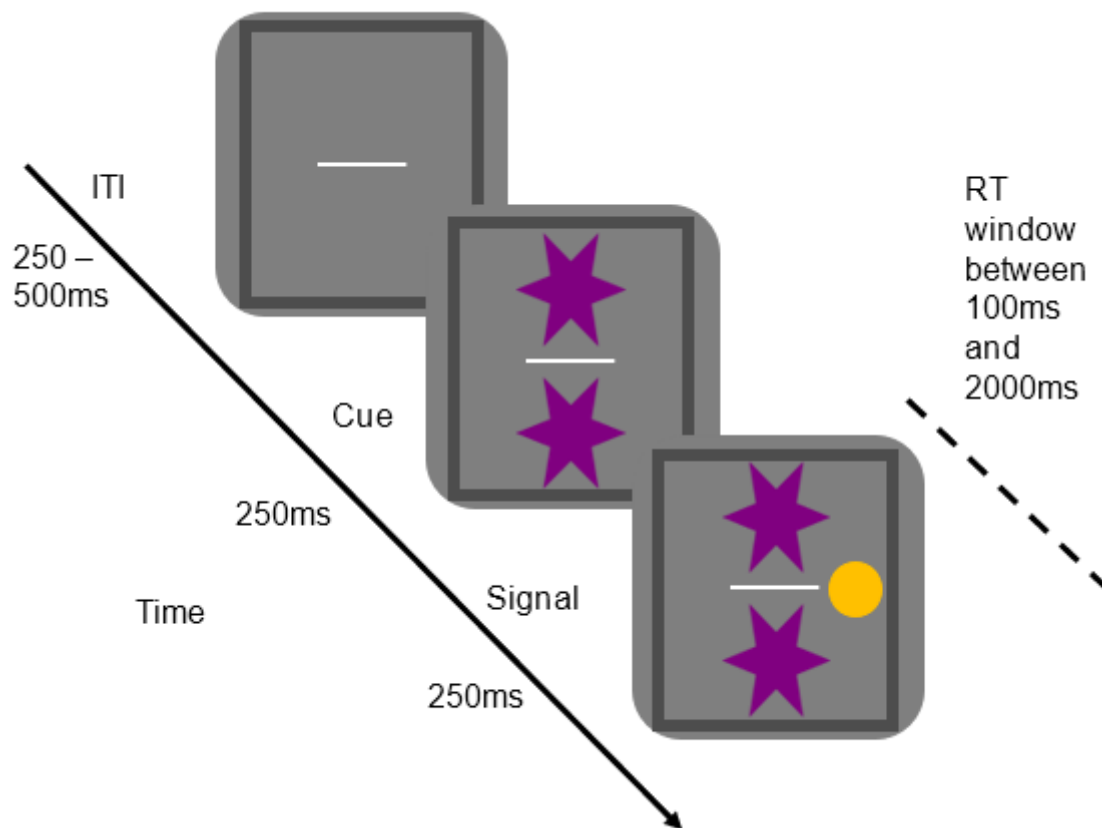


Figure 4.2. Schematic of a trial for Experiment 6. As the cueing square is dark grey this would indicate to participants that they need to respond to coloured circles (go is the task set), so the correct response would be to respond.

The instructions were amended to reflect the addition of the conditional feature. Instructions for responses to a dark grey square (respond to coloured circles, withhold to white circles) were presented first, followed by the instructions for the light grey square (withhold to coloured circles, respond to white). This was then reinforced with the final instructions screen telling participants “if you see a dark grey rectangle then respond to coloured circles. If you see a light grey rectangle do not respond to coloured circles.” The instructions were written with a focus on a response to the coloured circles in order to reinforce the change in task requirements and to ensure that participants were in the correct task set.

By starting with responses to white circles (as in past experiments) it might have led to participants not fully appreciating the change in the task (as the response to colour is the more salient of responses). One potential issue with this is that the second line is an example of a reversed instruction, telling participants what not to do, which can be harder to interpret. Additionally, the fact that the go task set instructions were presented first might lead participants to assume this is the default task and so shift performance to trials with the stop task set. One approach could have been to counterbalance the order of the instructions across participants. However, this would have meant that instructions would have focused on not responding to coloured circles first. It was felt this might have been difficult for participants to understand, given that instructions in most psychological experiments prime what to do first.

#### **4.4.1.3.1 Awareness measures**

As with Experiments 4 and 5 an awareness measure was included.

### **4.4.2 Analysis and results**

The analysis for Experiment 6 was conducted in a similar manner to that for Experiments 4 and 5 and so the particulars will not be repeated here. In terms of outliers, two participants were removed and replaced for having high omission errors, while one participant was removed and replaced for having high commission errors. Two participants failed to complete the experiment. There were no outliers for reaction times. One change is that due to the variation of task set there is no longer a basis for undertaking the feature-positive effect analysis. Additionally, although this chapter has discussed the task switching literature, the analysis still focuses on the learning that takes place to the cues, and so analyses that one would normally see with some task set paradigms (e.g., switch costs) were not investigated here.

#### **4.4.2.1 Training phase**

As with the analyses reported in Chapter 3, performance is likely to be affected by repetition priming effects and so contrasts were undertaken to demonstrate learning in situations where such effects were not in play. I then present the traffic light analysis to explore performance to the cues of interest (Table 4.4 contains the relevant descriptive statistics).

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
A-/+	504.08	76.55	0.07	0.04	0.02	0.02
B-	537.24	84.81	0.05	0.02	0.04	0.04
G+	477.56	87.32	0.05	0.04	0.02	0.01
I+	540.07	75.61	0.05	0.04	0.03	0.02
IP-	NA	NA	0.04	0.02	NA	NA
J±	546.52	76.05	0.08	0.06	0.03	0.03
P-/+	530.46	84.73	0.06	0.03	0.03	0.03
R-	564.59	87.22	0.03	0.01	0.04	0.04
RA+	542.7	82.35	NA	NA	0.02	0.02
YZ-/+	558.39	71.68	0.06	0.04	0.04	0.04

Table 4.4. Training descriptive statistics for Experiment 6. Reaction times means are calculated using are raw data, but mean p(respond) and p(miss) use transformed data.

#### 4.4.2.1.1 Reaction times

For the training data, the best fitting model had a conditional  $R^2$  of 0.86 (see Table 4.5 for AICs).

##### 4.4.2.1.1.1 Evidence of associative learning

As with Chapter 3, the contrasts that provide evidence of associative learning in isolation are R vs. I and R vs. B. Importantly, the contrast R vs. I also controls for the effect of task set (as they both have the same conditional response). The alpha level was corrected to .025. The contrast R vs. I was significant,  $t(392) = -3.86$ ,  $p = < .001$ , 95% CI [-36.97, -12.07],  $d = -0.39$ , with slower responses to R than I indicating that cue R was not a go cue. The comparison R vs. B was also significant,  $t(392) = -4.31$ ,  $p = < .001$ , 95% CI [-39.80, -14.90],  $d = -0.43$ , with faster reaction times to cue B than R suggesting that participants learnt that cue R was a stop cue. Of course, it must be noted that B and R have different conditional cues. This, then, may in part be why R is coming out as a stronger stop cue than B. But that is not to in any way minimise the result. Under these conditions, and the assumed conditions experienced at UK traffic lights, R is a strong stop cue.

#### 4.4.2.1.1.2 Traffic light contrasts

Noting the above results, I can be confident that the following results arise because of associative learning as well as repetition priming. As G vs. B is no longer a pure manipulation check, the alpha level was corrected to .008. The G vs. B contrast was highly significant,  $t(392) = 9.40$ ,  $p < .001$ , 95% CI [47.24, 72.14],  $d = 0.95$ , with G having faster response times than B indicating that learning was occurring. The A vs. B contrast was also highly significant,  $t(392) = 5.22$ ,  $p < .001$ , 95% CI [20.71, 45.61],  $d = 0.53$ , with cue A being faster than B, indicating that A was not experienced as a stop cue. The A vs. G contrast was significant,  $t(392) = -4.18$ ,  $p < .001$ , 95% CI [-38.97, -14.07],  $d = -0.42$ , with faster response for G than A, indicating that while A primed responding more than B it was not as much a go cue as G. Cue R was significantly slower than B,  $t(392) = -4.31$ ,  $p < .001$ , 95% CI [-39.80, -14.90],  $d = -0.43$  suggesting that it was more of a stop cue than B. Cue R was also highly significantly slower than cue G,  $t(392) = -13.70$ ,  $p < .001$ , 95% CI [-99.49, -74.59],  $d = -1.38$ , suggesting that R was not a go cue. The A vs. R contrast was highly significant,  $t(392) = 9.73$ ,  $p < .001$ , 95% CI [48.06, 72.96],  $d = 0.96$ , suggesting that A was more of a go cue than R.

#### 4.4.2.1.2 p(respond)

For this measure, the best model was a model that included the main effects of cue with a Gamma family and inverse link and random intercept (see Table 4.5).

##### 4.4.2.1.2.1 Evidence of associative learning

For this measure the contrasts G vs. B and G vs. I provide the conditions to assess for the occurrence of associative learning, with the contrast G vs. B controlling for the effect of task set. The G vs. B contrast was not significant,  $z = -0.28$ ,  $p = .781$ , nor was the G vs. I contrast,  $z = -0.72$ ,  $p = .473$ . Overall, these two contrasts demonstrate that associative learning was weak for this measure and leads to the conclusion that this measure is less sensitive than reaction times, probably because of the low error rate. It is worth noting that cue B has the same conditional cue as G, and so that is the most important comparison.

#### 4.4.2.1.2.2 Traffic light contrasts

Given the above results it is likely that performance for commission errors is driven more by repetition priming effects than associative learning. For this analysis, the alpha level was corrected to .008. The G vs. B contrast was not significant,  $z = -0.28$ ,  $p = .781$ . The A vs. B contrast was not significant at the corrected alpha,  $z = 2.44$ ,  $p = .015$ , although this does suggest a trend for more errors to A than B which would be consistent with A priming go responses. Looking at Table 4.1 it could be that cue A is not overly affected by repetition priming. This is because when the cue does appear in a sequence trial position greater than 1 it is equally balanced to appear after a go or a stop cue (in Green 2 it occurs after a stop cue and itself is stop, while in Amber 1 the second cue is a go cue A and is preceded by another go cue A). The nearly significant result for A vs. B could be indicative of some associative learning. The A vs. G contrast was highly significant,  $z = 6.24$ ,  $p = .007$ , with surprisingly more errors for A than G, indicating that cue A was also more likely to prime errors of commission than cue G. Given that cue G is more often than not preceded by another go cue (see Red 1 for example), one would expect repetition priming effects to push cue G to become a go cue on this measure. The fact it does not suggests that cue A is a strong go cue. The contrast R vs. B was significant,  $z = -3.08$ ,  $p = .002$ , with unexpectedly more errors for cue B than R, which supports the contrasts for reaction times that suggested cue R was a stronger stop cue than B. However, noting the results in section 4.4.2.1.2.1 the role of repetition priming cannot be ignored. This is especially the case for this contrast where cue R is often embedded in a run of stop trials which is likely to enhance the stopping behaviour exhibited to this cue. The subsequent analysis on test data will help determine if this is the case. The R vs. G contrast was significant,  $z = -2.83$ ,  $p = .005$ , with more errors for G than R, again suggesting that R was not a go cue, yet for the reasons discussed above this difference is not unexpected. The A vs. R contrast was highly significant,  $z = 5.10$ ,  $p < .001$ , demonstrating that cue A was more of a go cue than cue R, and matching the results for reaction times. Nevertheless, these results could also be a consequence of repetition priming.

#### 4.4.2.1.3 Training phase summary

Overall, the reaction time data provides evidence that associative learning was occurring. However, given the lack of such evidence for commission errors it seems prudent to assume that performance on  $p(\text{respond})$  could be, in a large part, driven by repetition priming effects. In summary it seems that, in an experiment that combines go and stop task sets and some of the sequences experienced at traffic lights, cue G is a go cue, cue A is slower than G (and faster than B) but leads to more errors so is perhaps similar to G, and cue R a strong stop cue.

Model	Training	1 <sup>st</sup> Sequence	Other sequence
Reaction time models			
Main effects of cue	5242.19	6045.60	6024.71
Main effects of cue with random intercept	<b>4551.45</b>	<b>5624.26</b>	<b>5529.98</b>
$p(\text{respond})$			
Main effects of cue with Gamma family and inverse link	-1856.54	-2165.71	-2297.90
Main effects of cue with Gamma family and inverse link and random intercept	<b>-1871.31</b>	<b>-2181.90</b>	<b>-2322.91</b>

Table 4.5. AICc scores for models for Experiment 6 run on reaction time and  $p(\text{respond})$  data at training and test. Bold are the models chosen.

#### 4.4.2.2 Test analysis

The same issues within the test phase highlighted in Chapter 3 are also relevant for Experiment 6. Thus, the same two separate analyses will be undertaken as in Chapter 3. The first analysis focuses on filler trials and first trial in a sequence, while the second explores performance on trials embedded within sequences. See Table 4.6 for descriptive statistics.

#### 4.4.2.2.1 First trial analyses

The analyses conducted on the first trial of each sequence were a G vs. B contrast (which is now once again a manipulation check) and A vs. B, A vs. G, A vs. R, R vs. B and R vs. G. The alpha level for these analyses was corrected to .010 (see Table 4.5 for AICc's).

##### 4.4.2.2.1.1 Reaction times

The model chosen had a conditional  $R^2$  of 0.66. The G vs. B contrast was significant,  $t(441) = 4.95$ ,  $p < .001$ , 95% CI [36.48, 84.23],  $d = 0.47$  indicating that learning had taken place by the test phase. The A vs. B contrast approached significance at a standard alpha level,  $t(441) = 1.88$ ,  $p = .061$ , 95% CI [-1.02, 46.72],  $d = 0.18$ , suggesting a weak trend for participants to make faster responses to cue A than B. The A vs. G contrast was significant,  $t(441) = -3.08$ ,  $p = .002$ , 95% CI [-61.38, -13.64],  $d = -0.29$ , suggesting that cue G primed more of a go response than cue A. The R vs. B contrast was not significant,  $t(441) = -1.18$ ,  $p = .241$ , 95% CI [-38.19, 9.56],  $d = 0.11$ , and cue R had significantly slower reaction times than cue G,  $t(441) = -6.13$ ,  $p < .001$ , 95% CI [-98.54, -50.79],  $d = -0.58$  suggesting that cue R was not a go cue. The contrast A vs. R was also significant,  $t(441) = 3.05$ ,  $p = .002$ , 95% CI [13.28, 61.03],  $d = 0.29$  suggesting that cue A primed more of a go response than cue R. These results clearly indicate that G is a go cue, R a stop cue, and A somewhere in between.

##### 4.4.2.2.1.2 $p(\text{respond})$

For commission errors there was a non-significant difference for G vs. B,  $z = -0.19$ ,  $p = .848$ , suggesting that performance was not strong for this contrast. The contrast A vs. B was significant at a standard alpha level,  $z = 2.13$ ,  $p = .033$ , and hints at a trend for more errors to A than B. The difference between A vs. G was significant at a standard alpha level,  $z = 2.29$ ,  $p = .022$ , suggesting an unexpected trend for more errors to cue A than G. For cue R, the contrast against cue B was not significant,  $z = -0.82$ ,  $p = .410$  (this suggest that the significant difference in training was largely due to repetition priming), nor was the contrast against G,  $z = -0.64$ ,  $p = .525$ . The contrast for A vs. R was significant,  $z = 2.77$ ,  $p = .006$ , with more errors to cue A than R. These results suggest that A is as much of a go cue (perhaps more so) than G and confirm that R is a stop cue.

#### **4.4.2.2.2 Other sequence trial analyses**

This analysis focused on comparing the traffic light cues within the sequences. Therefore, contrasts undertaken were A vs. G, A vs. R, and R vs. G. The alpha level was corrected to .017.

##### **4.4.2.2.2.1 Reaction times**

The model chosen had a conditional  $R^2$  of 0.73. The A vs. G contrast was significant,  $t(441) = -3.93$ ,  $p < .001$ , 95% CI [-64.39, -21.56],  $d = -0.37$ , with faster responses to cue G than A suggesting G primed a greater go response. The A vs. R contrast was significant,  $t(441) = 3.85$ ,  $p < .001$ , 95% CI [20.61, 63.43],  $d = 0.37$ , with faster responses to cue A than R. Lastly, the R vs. G contrast was highly significant,  $t(441) = -7.78$ ,  $p < .001$ , 95% CI [-106.40, -63.58],  $d = -0.74$ , with faster reaction times to cue G than R indicating that R was not a go cue. Once again, the reaction time data suggest that G is a strong go cue, R a strong stop cue and A is in between.

##### **4.4.2.2.2.2 p(respond)**

The A vs. G contrast was not significant,  $z = 1.24$ ,  $p = .215$ . The A vs. R contrast was significant at a standard alpha level,  $z = 2.34$ ,  $p = .019$ , suggesting a weak trend for cue A to promote more of a go response than cue R. The R vs. G contrast was not significant,  $z = -1.23$ ,  $p = .219$ . I note that the analysis of the sequence trials is now producing useful results. Yet this is qualified by the change in task set that will occur during sequences.

##### **4.4.2.2.3 Test phase summary**

The first sequence trial analyses indicates that at a purely associative learning level participants had learnt that G was a go cue, with this having significantly faster reaction times than cue B. Cue A seems to be a fairly weak go cue, with hints at being more going than cue B for both reaction times and commission errors, and having significantly more errors than G but slower reaction times. The cue was also significantly more go on both measures when compared to cue R. For cue R learning suggests it primed a stop response, having significantly slower reaction times than cue G but similar performance on both measures to B (and similar levels of error rates to G).

This general pattern is supported by the other sequence trial analyses. Once again cue A seemed to be a fairly weak go cue. The cue had significantly



slower reaction times than cue G but similar commission errors, and it primed significantly more going behaviour than cue R for both measures. For reaction times cue R promoted significantly more stopping behaviour than cue G, suggesting it was a stop cue.

Type of cues	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Filler cues						
B+	539.21	84.33	0.04	0.04	0.02	0.04
I-	537.72	96.37	0.04	0.05	0.03	0.04
IP+	571.41	109.1	0.04	0.06	0.03	0.04
J±	556.54	108.24	0.06	0.12	0.03	0.07
P-/+	527.29	88.52	0.05	0.08	0.04	0.07
YZ-/+	562.16	104.44	0.06	0.08	0.03	0.06
1st sequence						
trials						
experimental						
cues						
A	516.37	91.62	0.06	0.11	0.01	0.03
G	478.86	109.77	0.03	0.07	0.01	0.00
R	553.53	117.28	0.03	0.07	0.05	0.11
RA	531.39	94.17	0.04	0.09	0.02	0.05
Other sequence						
trials						
experimental						
cues						
A	488.02	107.21	0.05	0.09	0.01	0.03
G	445.05	92.87	0.04	0.05	0.01	0.02
R	530.04	97.54	0.03	0.05	0.01	0.02
RA	498.02	96.26	0.02	0.03	0.01	0.00

Table 4.6. Test descriptive statistics for Experiment 6. Reaction time means are calculated using raw data, but mean p(respond) and p(miss) use transformed data.

#### 4.4.2.3 Expectancies

As in Chapter 3 it was possible to contrast expectancy scores between cues to assess the evidence of awareness. This analysis was undertaken in the same manner as the reaction time data (see Table 4.7 for descriptive statistics). An alpha level of .008 was applied to this analysis. The best model fitted the main effects of cue with random intercept ( $R^2 = 0.74$ ). The G vs. B contrast was highly significant,  $t(539) = 6.34$ ,  $p = < .001$ , 95% CI [3.03, 5.73],  $b = 0.55$ , with participants more likely to classify G as a go cue than B. The A vs. B contrast was not significant at the reduced alpha,  $t(539) = 2.00$ ,  $p = .046$ , 95% CI [0.03, 2.73],  $d = 0.17$ , although there was a trend for participants to rate cue A as more go than cue B. The A vs. G contrast was significant,  $t(539) = -4.34$ ,  $p = < .001$ , 95% CI [-4.35, -1.65],  $d = -0.37$ , with participants more likely to rate G as a go cue than cue A. The A vs. R contrast was marginally significant at a standard alpha,  $t(539) = 1.68$ ,  $p = .094$ , 95% CI [-0.19, 2.51],  $d = 0.14$ , with a slight trend for higher go ratings for A than R. Focusing on cue R, R vs. G was highly significant,  $t(539) = -6.03$ ,  $p = < .001$ , 95% CI [-5.51, -2.81],  $d = -0.52$ , with participants more likely to rate cue G as go than cue R. The R vs. B contrast was not significant,  $t(539) = 0.32$ ,  $p = .750$ , 95% CI [-1.13, 1.57],  $d = 0.03$ .

As in Chapter 3 it was possible to use J and YZ as baseline cues to see if awareness for B was neutral or biased in some way. This analysis was separate from that reported above and so an alpha level of .013 was applied. The J vs. G contrast was significant,  $t(539) = -4.95$ ,  $p = < .001$ , 95% CI [-4.77, -2.07],  $d = -0.43$ , with participants more likely to classify G as a go cue than J. The YZ vs. G contrast was highly significant,  $t(539) = -6.08$ ,  $p = < .001$ , 95% CI [-5.55, -2.85],  $d = -0.52$ , again with participants more likely to see G as a go cue than YZ.

Thus, participants seem to have been aware that cue G was a go cue compared to neutral cues. Against B, the contrast for J was not significant,  $t(539) = 1.39$ ,  $p = .164$ , 95% CI [-0.39, 2.31],  $d = 0.12$ , nor was the contrast for YZ vs. B,  $t(539) = 0.26$ ,  $p = .794$ , 95% CI [-1.17, 1.53],  $d = 0.02$ . This indicates that participants expected to make very similar responses to cue B (a stop cue) and cues J and YZ (neutral 50:50 cues). Of course, contrasts against cue R might support the argument regarding participant's lack of awareness best given that R is now a strong stop cue. Indeed, neither the J vs. R contrast,  $t(539) = 1.07$ ,  $p = .284$ , 95% CI [-0.61, 2.09],  $d = 0.09$ , nor the YZ vs. R

contrast,  $t(539) = -0.06$ ,  $p = .954$ , 95% CI [-1.39, 1.31],  $d = -0.00$  was significant. Thus, though cue R seems to be a strong stop cue behaviourally, it is numerically similar to J and YZ in terms of expectancy.

Overall, the analysis shows that for some cue's participants were aware of the required response. This awareness seems to focus on cue G, with participants showing consistent awareness of this cue relative to others.

Cue	Expectancy Rating	
	Mean	SD
A-/+	1.98	3.35
B-	0.60	4.69
G+	4.98	3.40
I+	2.66	3.49
IP-	-0.92	3.92
J±	1.56	4.26
P-/+	1.86	3.73
R-	0.82	4.09
RA+	-0.52	3.60
Y-/+	-0.28	4.11
YZ-/+	0.78	3.65
Z-/+	-0.28	4.53

Table 4.7. Descriptive statistics for expectancy ratings for Experiment 6.

#### 4.4.2.4 Correlational analysis

As in Chapter 3 it is possible to compare performance scores and expectancy ratings to see if there is a correlation between them. Therefore, correlations comparing the significant contrasts from the expectancy ratings are presented, that is, G vs. B, G vs. A, and G vs. R. For training the whole dataset was used, whilst for test only data from first sequence trials was included. Looking at Figure 4.3, once the alpha level correction is applied (to  $p = .017$ ), it seems that awareness did not affect performance at training – though the result for G vs. B for response times might be taken to indicate that awareness did have some influence upon performance. Regarding the test data (Figure 4.4) it seems performance and awareness were not correlated. Thus, although participants appear to notice cue G, and to be able to articulate this knowledge in their

ratings at training, the correlation between the advantage of G over B and the difference in awareness ratings never becomes significant after correction. Of course, one caveat is that perhaps there is a small causal influence that this analysis lacks the power to detect.

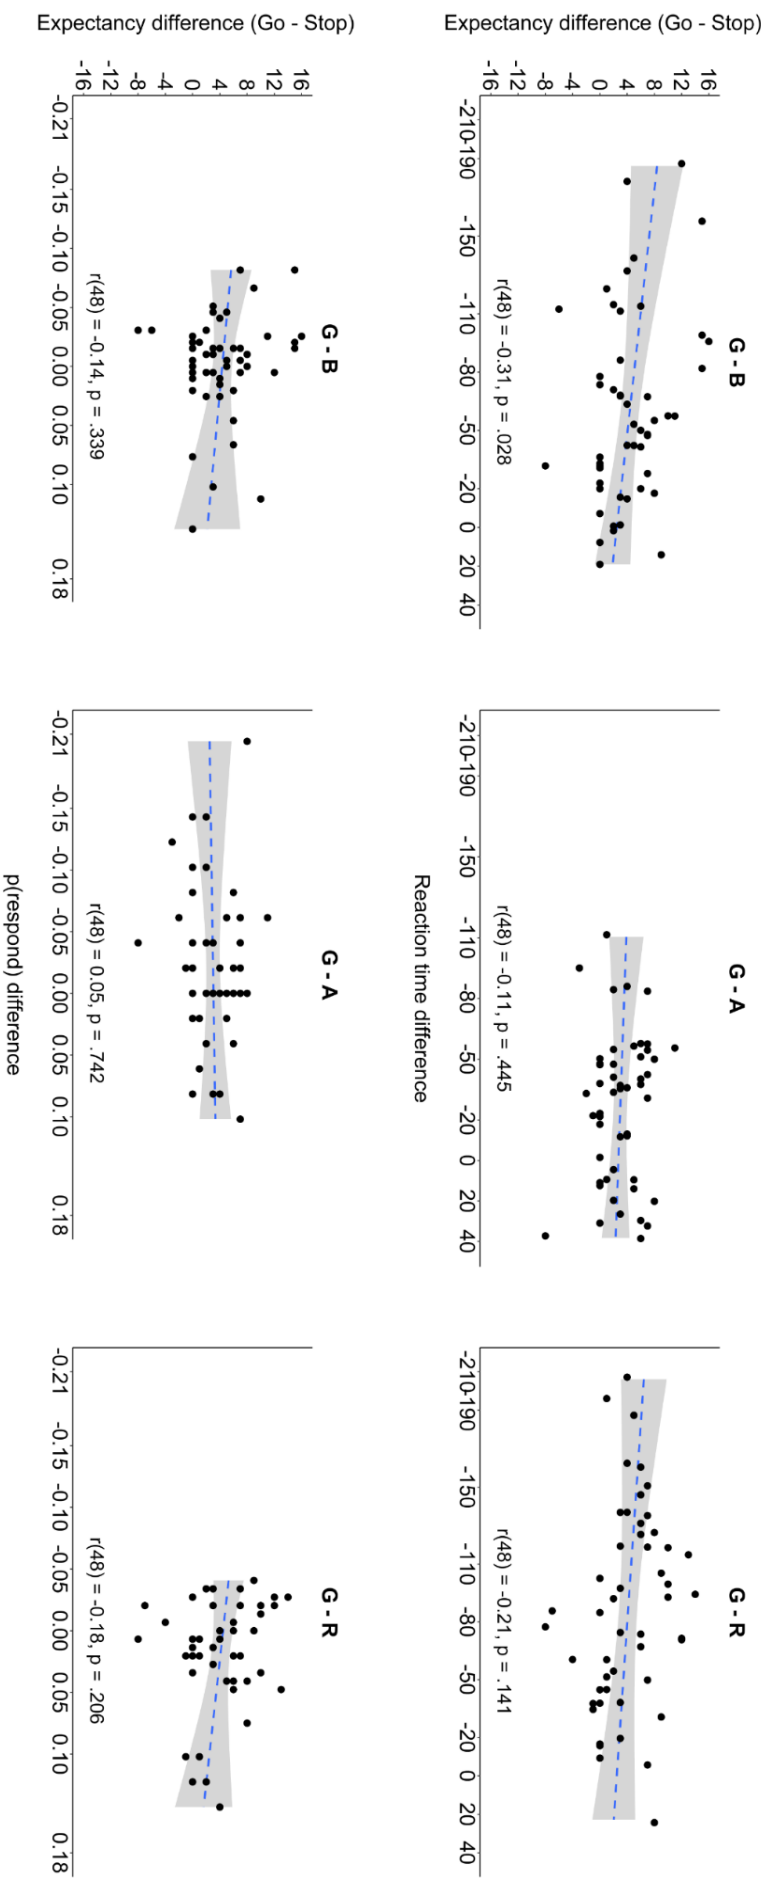


Figure 4.3. Experiment 6 training phase correlations between expectancies and reaction times (top panel) or p(response) (bottom panel) for each contrast G vs. B, G vs. A and G vs. R. Scores were calculated by taking the stimuli that was more contingent with stopping from the one that was most contingent with going (i.e. go-stop). Dashed line depicts linear model of the data, shaded area represents the 95% confidence interval of that model. R values refer to Pearson's product-moment correlation.

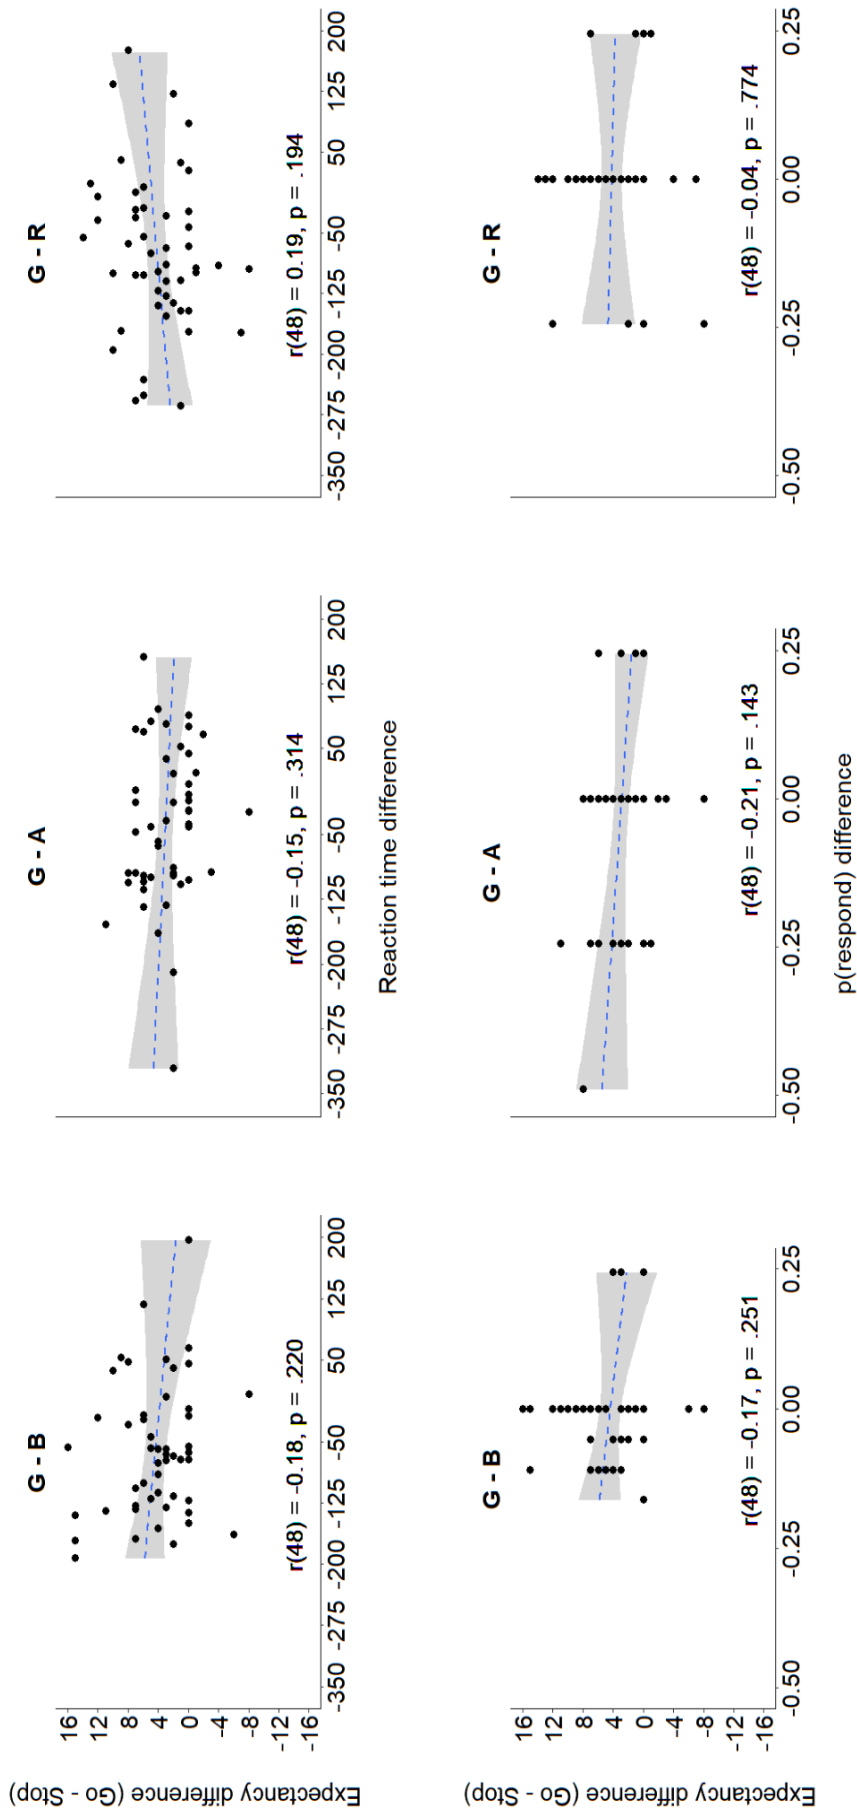


Figure 4.4. Experiment 6 test phase correlations between expectancies and reaction times (top panel) or commission errors (bottom panel) for each contrast G vs. B, G vs. A and G vs. R. Scores were calculated by taking the stimuli that was more contingent with stopping from the one that was most contingent with going (i.e. go-stop). Dashed line depicts linear model of the data, shaded area represents the 95% confidence interval of that model. R values refer to Pearson's product-moment correlation.

### 4.4.3 Summary

Experiment 6 was an attempt to combine two factors highlighted in preceding chapters that seem to influence learning of the contingencies experienced at UK traffic lights, or more accurately, within the paradigm developed here to capture learning at UK traffic lights. Using a within-participants design, Experiment 6 exposed participants to both the task sets and sequences likely to be experienced at traffic lights.

Firstly, it is worth noting that participants were able to cope with this paradigm. The (by now) typical finding of G being significantly different from B was obtained, with significantly faster response times to G than to B in both the training phase and in the first sequence trials analysis, though cue G did not produce significantly more errors for  $p(\text{respond})$  in either phase. This illustrates that at a fundamental level learning to Experiment 6 was still occurring as expected despite the substantial changes that had been made.

Turning to the two key traffic light cues of A and R, at training there is evidence that cue A was not a stop cue, having significantly faster response times than B and a strong trend (though not significant at the corrected alpha level) for more commission errors than B. Of course, caveats apply for the A vs. B contrast for commission error as there was little evidence that associative learning was driving performance for this measure. Compared to G, cue A was significantly slower at training but had significantly more commission errors. The commission contrast is intriguing, being contrary to expected performance based on repetition priming effects and indicating a role for associative learning. Overall, focusing on the reaction times (where I can be confident performance is, in part, an outcome of associative learning) it seems that cue A is a neutral cue. However, the results for  $p(\text{respond})$  do hint that cue A may be more of a weak go cue than a neutral cue. In the case of R, it was significantly slower in terms of reaction times for both cue G and cue B, indicating that cue R is a strong stop cue. For  $p(\text{respond})$ , cue R also had significantly less errors than B or G, supporting the argument that cue R was a stop cue. However, it must be highlighted that for both contrasts for errors there is a clear logic supporting the effects of repetition priming upon performance.

For the first sequence test trials the results suggest that cue A primed a quite weak go response. Reaction times were faster for A than B, but this difference only approached significance at a standard alpha level, and cue A had significantly slower response times compared to G. Yet, for error rates cue A had significantly more errors at a standard alpha level than both B and G, suggesting that for errors there was a trend for cue A to promote a go response. For cue R, the data indicates it was a stop cue. The cue had significantly slower response times than cue G but was similar in performance for cue B. For commission errors learning seemed to be weak, with error rates not being significantly difference between B, G or R. Overall, the results suggest that on a stop-go continuum, cue G was a go cue, cue A was quite a weak go cue, cue B a stop cue, and cue R a stop cue, slightly more so than cue B based on the reaction time data.

In terms of the results for cue A, Experiment 6 is consistent with other results presented in this thesis and shows how an amber light could become linked to a go response, albeit a weak one. For cue R, this experiment contrasts quite sharply with the general conclusion of the previous experiments that suggested R to be a weak stop cue. However, R has been consistently found to promote stopping (to varying degrees) and so the shift is not quite as marked as might at first appear. The difference could well be due to the different task sets in play. Looking back to the experiments in Chapters 2 (and ignoring the enhanced learning generated by sequences) it would seem that what I have called a stop task set (where the default is go) is one that potentiates learning to go to cue A. Whereas, in a go task set (where the default is stop) learning potentiates stopping to R. In some sense, then, the labelling of these task sets can be seen as giving the wrong impression (even though logically they are labelled entirely correctly). Of course, the task sets themselves are confounded with particular sequences and cues that have different contingent relationships to responding, so it is difficult to be unequivocally sure about this conclusion. The take home message is that whilst one may now have more confidence that a red traffic light will indeed be learnt as a stop cue due to driver's experience at traffic lights, it still remains the case that an amber light might promote something of a go response, and certainly not the relatively stop cue envisaged in the Highway Code.



The results for the awareness correlations continue to support the view that participants are learning about the task through associatively-mediated processes, with correlations not being significant after corrections. However, the fact that the G vs. B correlation was significant at a standard alpha level for training does suggest that awareness might have some role to play in performance, and that these null results could partly be due to a lack of power. It might be fairer to acknowledge that participants do seem to pick up on G more often than not, realising that a response will be required to this cue, but that the other cues all seem to be similar in regard to the expectancy ratings.

In past chapters the first experiment has been replicated, but with a switch in task set. This approach is not appropriate here, but Experiment 7 does replicate the training phase of Experiment 6, and begins the process of development of a paradigm that will allow us to change the learnt contingencies to cue A.

## **4.5 EXPERIMENT 7**

Throughout the experiments presented in this thesis, it has become clear that amber traffic lights may not be associated with the weak stop/stop contingency that the Highway Code prescribes, but rather seem to be associated with a weak go response (at least in the task developed here). This suggests that the associative system is cueing people for a go response at amber traffic lights, and that either in tandem with more conscious decision making (e.g., If I jump the light I could get home early), or, worse, without the input of conscious control (McLaren et al., 2019) could lead to people crossing amber lights. Given the role associative learning might have in motivating people to jump amber traffic lights, the next question becomes: is it possible to address this behaviour by retraining learning, so that amber is associated with stop rather than go? This brings to mind the work on associatively-mediated inhibition such as that by N. S. Lawrence, O'Sullivan, et al. (2015) reviewed in Chapter 1, and I will return to this in the discussion. Experiment 7 begins the development of this inhibition training, but it is also a direct replication of Experiment 6 (the inhibition training aspect will be discussed in more detail in Chapter 6).

The existing paradigm has two parts; in the training phase participants learn the contingencies of the task (which reflect some of the contingencies at play in UK traffic lights) and then in the test phase I investigate what learning has taken

place. By the end of the training phase, but before the test phase, participants seem to learn that A is a fairly weak go cue. Therefore, placing an inhibition task whereby cue A is linked to stopping before the test phase should retrain participants to exhibit a stop response to cue A. In the subsequent test phase one would expect those who had this training to have slower responses and less commission errors to cue A than those who received a similar task.

In terms of the participant's experience of the inhibition phase, it needed to be similar to the existing design for learning to transfer, therefore a warning cue (the shape) had to appear on the screen followed by another signal denoting the response required. The responses used so far have been go or stop and these responses were used again. In addition, go signals have previously appeared on either side of the shape and so this will be another constraint on the inhibition task. As task set was not manipulated during this phase, the cueing square was removed. Given the intervention is about learning to stop, this phase only had two cues, a prepotent go cue and a less frequently seen stop cue. Therefore there was no tracking cue, with the experiment having a set response window of 1000ms in line with similar work (Bowditch, 2016).

Regarding the cues used in this inhibition phase, clearly A has to be the stop cue for the experimental group. While it would seem sensible at face value to use cue G as the go cue, the extra exposure to cue G might affect learning at the test phase. Additionally, the inhibition task will not use sequences, and so using a novel go cue was felt best in order to leave comparisons of G vs. B at test unaltered. Therefore, a new cue, cue X, was employed as the go cue for the inhibition phase. For the control condition, in order to match learning as closely as possible, participants were exposed to the go cue X but for the stopping response saw another novel cue, cue O, instead of cue A.

One last change was made: rather than respond (or not) to white or coloured circles participants had to respond (or not) to black arrows. This change was made to ensure that the phase was not too similar to the past blocks, as if it was identical then the phase would not be inhibition per se, but rather additional training, with admittedly changed contingencies. However, to ensure that learning from the inhibition phase did transfer to test it was important to ensure that the phase did not feel too dissimilar either, and so responding to arrows rather than circles was felt to strike the optimum balance. Black was chosen as

the colour for the arrows to ensure that any learning to colours resulting from the training phase did not impact learning for the inhibition phase, with participants never responding (or not) to black circles.

Figure 4.5 shows a go trial for the inhibition phase. What should be obvious is how similar it is to the designs already used (especially to the designs used in Experiment 1-5). In fact, the only visual change from those experiments is the replacement of the signal circle with an arrow pointing to which side key needed pressing (these keys remained the same throughout experiments).




Figure 4.5. Schematic of a go trial for the inhibition phase of Experiment 7.

Figure 4.6 shows a no-go trial. Several other adjustments were made for these trials, primarily to maintain task engagement. This was a concern as the phase was relatively slow in the absence of a tracking cue. Also, the phase was relatively straightforward, with only two cues, and came towards the end of the overall experiment when enthusiasm might be flagging. To address this, instead of having a coloured or white arrow to denote no-go, firstly a black go arrow appeared followed at one of three intervals by a black X superimposed on the arrow to indicate no-go. In effect this is a stop rather than a no-go trial. Finally, to reinforce learning to the stop cue, commission errors led to a two second timeout screen displaying the message “Error! Timeout!”

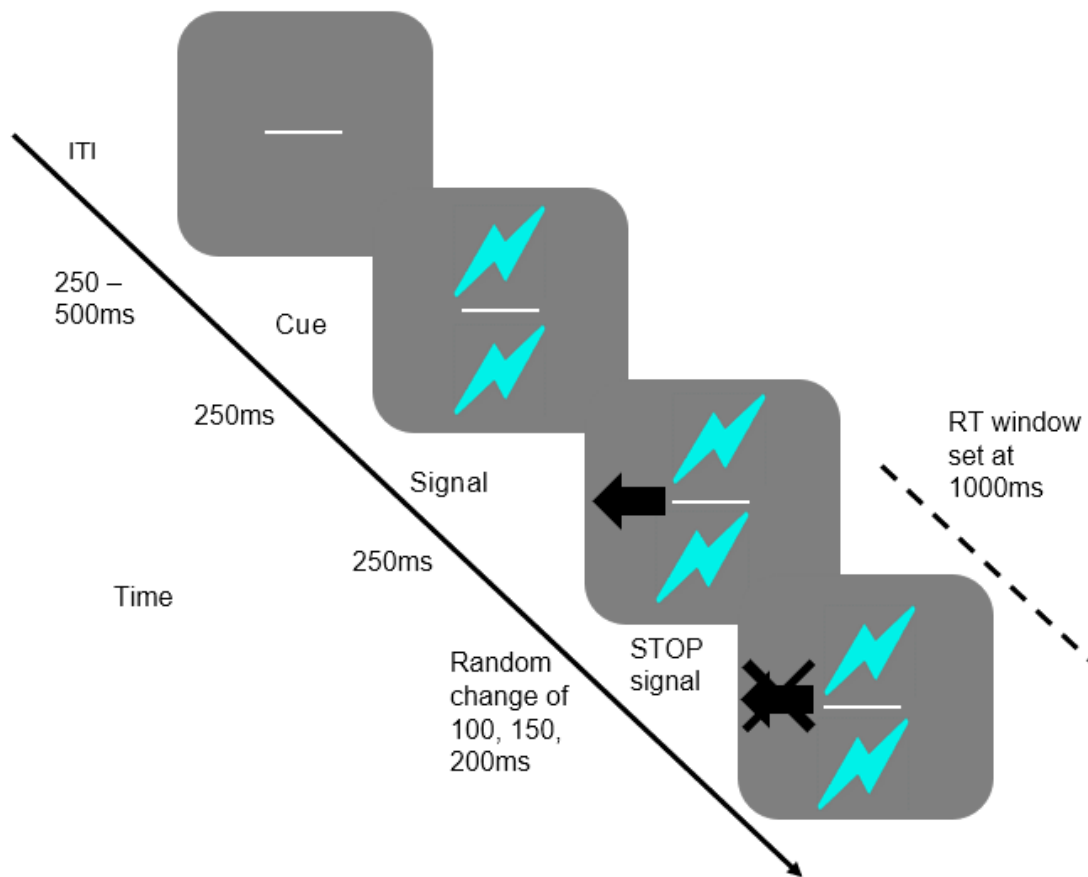


Figure 4.6. Schematic of a stop trial for the inhibition phase of Experiment 7.

## 4.5.1 Method

### 4.5.1.1 Participants

The inclusion criteria and outlier removal process were the same as Experiment 2 (see Results section for details of outliers). Experiment 7 had a similar sample size as Experiment 6 but was split into two groups (26 in each). In the final sample, 8 in the experimental group and 17 in the control group were female, with a mean age of 19.81 (SD = 1.60) and 20.69 (SD = 2.07) respectively. Participants received payment of £10 or one and a half course credits.

### 4.5.1.2 Design

Although Experiment 7 had a between participant's feature (the type of inhibition training), this was only introduced in the inhibition phase, and all participants experienced the learning and training phases identically; in fact, these phases were a direct replication of Experiment 6. The design of Experiment 7 can be seen in Table 4.8. Here, another feature of the design is made clear. While the importance of prepotent responses for the inhibition task

has been noted, it is still important to have data regarding commission errors, and so each cue had eight trials that were the opposite of its majority response (this also served to keep participants focused on the task). One issue with this approach was that while there were three possible stop delays for the stop cue in the inhibition task, there were 22 go trials. To square this mathematical circle, each participant received 7 of each of the delays with the extra delay being chosen from the three groups at random. While there was a 10 second break half-way through each block in the training and test phase, this was not the case for the inhibition block and participants only received 10 second breaks at the end of each block. No awareness measures were taken.

#### **4.5.1.3 Procedure**

The only change in procedure from that of Experiment 6 were the instructions participants received. Participants received the same instructions as in Experiment 6 at the start of the experiment and at the start of the test phase. In the case of the inhibition task there was only one way of presenting the instructions, with participants first being informed how to respond to the arrows and then being instructed not to respond if you see a black cross. To some extent this can be seen as promoting a stop task set. Additionally, the final set of instructions for the inhibition phase instructed participants to “not respond if you see a cross appear on the black arrow”.

Phase	Blocks	Trials per block	Design
Learning Phase	1	36	J
Training	6	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16) J(8)
Inhibition training	2	100	Experimental group: X (70, 8 STOP, 62 GO), A (30 STOP, 8 GO).  Control group: X (70, 8 STOP, 62 GO), O (30 STOP, 8 GO).
Test	1	232	G <sub>1</sub> (4), G <sub>2</sub> (4), A <sub>1</sub> (4), A <sub>2</sub> (4), R <sub>1</sub> (4), R <sub>2</sub> (4), RA(8), B(36), I(28), IP(24), P(16), YZ(16) J(8)

Table 4.8. Summary of Experiment 7 design. Letters represent a sequence. The numbers in parentheses indicate how often the sequence will occur per block. At test the contingencies were  $\pm$ .

#### 4.5.2 Analysis and results

The analysis for Experiment 7 is split into three sections. As the training phase was experienced in the same manner for all participants, this was analysed in the same way as Experiment 6 (to enable assessment of the replication of that experiment). The inhibition phase was analysed in a way that would determine whether participants were learning about the cues contained therein. The test phase was analysed in a similar manner as in Experiment 6, except for the addition of a between-participants factor based on the inhibition task that participants experienced. Two participants were replaced to due to withdrawing from the experiment; two were replaced for being commission outliers, and two for being omission outliers. There were no reaction time outliers.

### **4.5.2.1 Training phase**

As with the analysis of Chapter 3 performance is likely to be affected by repetition priming effects and so contrasts were undertaken to demonstrate learning in situations where such effects were not in play. I then present the traffic light analysis to explore performance to the cues of interest (Table 4.9 contains the relevant descriptive statistics).

#### **4.5.2.1.1 Reaction times**

For the training data, the best fitting model had a conditional  $R^2$  of 0.88 (see Table 4.10 for AICs).

##### **4.5.2.1.1.1 Evidence of associative learning**

As with Experiment 6 the contrasts that provide evidence of associative learning in isolation are R vs. I and R vs. B. Importantly, the contrast R vs. I also controls for the effect of task set (as they both have the same conditional response). The alpha level was corrected to .025. The contrast R vs. I was significant,  $t(408) = -3.61$ ,  $p < .001$ , 95% CI [-36.92, -10.94],  $d = -0.36$ , with slower responses to R than I indicating that cue R was not a go cue. The comparison R vs. B was also significant,  $t(408) = -3.27$ ,  $p = .012$ , 95% CI [-34.64, -8.66],  $d = -0.32$ , with faster reaction times to cue B than R suggesting that participants learnt that cue R was a stop cue.

##### **4.5.2.1.1.2 Traffic light contrasts**

Noting the above results, I can be confident that the following results arise as a consequence of associative learning as well as repetition priming. As G vs. B is no longer a pure manipulation check, the alpha level was corrected to .008. The G vs. B contrast was highly significant,  $t(408) = 11.32$ ,  $p < .001$ , 95% CI [62.04, 88.03],  $d = 1.12$ , with G having faster response times than B, indicating that learning was occurring and replicating the effect seen in Experiment 6. The A vs. B contrast was also highly significant,  $t(408) = 8.06$ ,  $p < .001$ , 95% CI [40.43, 66.42],  $d = 0.80$ , with A being faster than cue B, indicating that cue A was not experienced as a stop cue. The A vs. G contrast was significant,  $t(408) = -3.26$ ,  $p = .001$ , 95% CI [-34.60, -8.61],  $d = -0.32$ , with faster response for G than A, indicating that while A was more go than B it was not as strong at priming responding as G. Cue R was significantly slower than cue B,  $t(408) = -3.27$ ,  $p = .001$ , 95% CI [-34.64, -8.66],  $d = -0.32$ , suggesting that it was more of

a stop cue than B. Cue R was highly significantly slower than cue G,  $t(408) = -14.58$ ,  $p = < .001$ , 95% CI [-109.68, -83.69],  $d = -1.44$ , indicating it was not a go cue. The A vs. R contrast was highly significant,  $t(408) = 11.33$ ,  $p = < .001$ , 95% CI [62.08, 88.07],  $d = 1.12$ , suggesting that A was more of a go cue than R.

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
A-/+	483.67	84.14	0.07	0.05	0.02	0.02
B-	537.09	97.71	0.05	0.03	0.03	0.03
G+	462.06	96.6	0.06	0.06	0.02	0.01
I+	534.82	92.9	0.09	0.06	0.03	0.02
IP-	NA	NA	0.05	0.03	NA	NA
J±	548.13	93.6	0.08	0.06	0.03	0.04
P-/+	527.05	92.83	0.07	0.04	0.03	0.02
R-	558.75	90.74	0.04	0.03	0.03	0.03
RA+	533.35	93.82	NA	NA	0.02	0.02
YZ-/+	554.23	91.77	0.07	0.04	0.04	0.03

Table 4.9. Descriptive statistics for the training phase of Experiment 7. Reaction times means are calculated using are raw data, but mean p(respond) and p(miss) use transformed data.

#### 4.5.2.1.2 p(respond)

For this measure, the best model was a model that included the main effects of cue with a Gamma family and inverse link and random intercept (see Table 4.10 for AICs).

##### 4.5.2.1.2.1 Evidence of associative learning

For this measure the contrasts G vs. B and G vs. I provide the conditions to assess for the occurrence of associative learning, with the contrast G vs. B controlling for the effect of task set. The alpha level was corrected to .025. The G vs. B contrast was not significant,  $z = 1.43$ ,  $p = .154$ , but, the G vs. I contrast was,  $z = -2.39$ ,  $p = .017$ , with more errors to I than G indicating that for this measure G did not promote going to the same extent as I.

##### 4.5.2.1.2.2 Traffic light contrasts

Given the above results it is likely that performance for commission errors is driven more by repetition priming effects than by associative learning. For this



analysis, the alpha level was corrected to .008. The G vs. B contrast was not significant,  $z = 1.43$ ,  $p = .154$ . The A vs. B contrast was not significant at the corrected alpha,  $z = 2.37$ ,  $p = .018$ , although this suggests a trend for more errors to A than B. As discussed in Experiment 6 this provides some evidence for associative learning impacting performance as repetition priming should act to prevent errors to A 'in sequence'. The A vs. G contrast was not significant,  $z = 0.98$ ,  $p = .326$ . The R vs. B contrast was not significant,  $z = -1.35$ ,  $p = .176$ . However, the R vs. G contrast was significant,  $z = -2.72$ ,  $p = .007$ , with more errors for G than R suggesting that G was a go cue compared to R. This contrast can be explained by reference to the repetition priming effects likely to be in play within sequences. The A vs. R contrast was significant,  $z = 3.59$ ,  $p = < .001$ , demonstrating that cue A was more of a go cue than cue R. Nevertheless, these results could also be a consequence of repetition priming.

#### **4.5.2.1.3 Training phase summary**

Overall, what is striking about these results is how similar they are to those of Experiment 6. In fact, the reaction time results are only quantitatively different. There is a slight difference for the p(respond) data where A is not significantly different from G, whereas in Experiment 6 cue A had significantly more errors. The R vs. B contrast is also not significant, where in Experiment 6 it was (with more errors for B than R). However, the general picture of results for p(respond) do not suggest different conclusions to those drawn for Experiment 6. To sum, for Experiment 7 at training it seems that cue G is a go cue, cue A quite a weak go cue, cue R a stop cue and cue B a stop cue.

Model	Training
Reaction time models	
Main effects of cue	5579.53
Main effects of cue with random intercept p(respond)	<b>4801.18</b>
Main effects of cue with Gamma family and inverse link	-1719.24
Main effects of cue with Gamma family and inverse link and random intercept	<b>-1760.74</b>

Table 4.10. AICc scores for training models for Experiment 7 run on reaction time and p(respond) data. Bold are the models chosen.

#### 4.5.2.2 Inhibition phase

In order to explore learning for the inhibition phase within each condition, contrasts comparing cue GO (cue X) to cue STOP (cue A or O in the experimental and control conditions respectively) were run in order to assay evidence for this discrimination. I also ran a between-participants analysis to see if the difference between cue GO vs cue STOP was different between conditions. Due to the lack of sequences for this phase, repetition priming effects will not impact performance, and based on findings in Chapter 2 it is assumed performance is largely due to associative learning (see Table 4.11 for descriptive statistics).

Cue	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Experimental						
X	487.46	48.71	0.21	0.16	0.03	0.02
A	493.60	51.46	0.23	0.14	0.04	0.04
Control						
X	477.43	47.34	0.23	0.14	0.03	0.02
O	497.59	46.25	0.18	0.09	0.05	0.06

Table 4.11. Descriptive statistics for the inhibition phase of Experiment 7.

Reaction times means are calculated using are raw data, but mean p(respond) and p(miss) use transformed data.

The reaction time contrasts were run using a mixed effects model with participant as a random intercept (see Table 4.12 for AICs). An alpha level of .025 was applied to this analysis. The X vs. A contrast (conditional  $R^2$  of 0.84) was not significant,  $t(25) = 1.12$ ,  $p = .274$ , 95% CI [-4.61, 16.89],  $d = 0.05$ , indicating that participants in the experimental condition did not learn the contingencies. However, the same contrast (X vs. O) for the control condition (conditional  $R^2$  of 0.80) was significant,  $t(25) = 3.42$ ,  $p = .002$ , 95% CI [8.61, 31.71],  $d = 1.37$ , suggesting that participants in the control inhibition contrast had learnt the discrimination with faster response times to the go cue.

In terms of commission errors, the best fitting model was a model that included a gamma link but no random terms (see Table 4.12 for AICs). An alpha level of .025 was applied to this analysis. In this model, the contrast X vs. A was not significant,  $z = 0.45$ ,  $p = .651$ . The X vs. O contrast was also not significant,  $z = 1.60$ ,  $p = .109$ . This suggests that for both conditions the discriminations were not well learnt for this measure.

Model	Experimental condition	Control condition
Reaction time models		
Main effects of cue	559.09	552.00
Main effects of cue with random intercept	<b>517.50</b>	<b>516.97</b>
p(respond)		
Main effects of cue with Gamma family and inverse link	<b>-69.03</b>	<b>-83.73</b>
Main effects of cue with Gamma family and inverse link and random intercept	-68.96	-82.81

Table 4.12. AICc scores for inhibition models for Experiment 7 run on reaction time and p(respond) data. Bold are the models chosen.

Given the null findings for p(respond) I did not run a between-participant's analysis for this measure, but I did for the response times. A *t*-test found a marginally significant difference (at a standard alpha level) between the differences for the go cue minus the stop cue between the two conditions for response times,  $t(50) = 1.74$ ,  $p = .089$ , 95% CI [-2.15, 30.19],  $d = 0.48$ , with a greater difference between GO vs. STOP in the control condition ( $M = -20.15$ ,  $SD = 30.05$ ) compared to the experimental condition ( $M = -6.14$ ,  $SD = 27.97$ ). This indicates that there was a weak trend for differences between the two conditions. It should be noted that the learning experience between the conditions was not equal. This was because participants in the control condition were exposed to two novel stimuli, while cue A was a familiar cue for the experimental condition participants. The marginally significant difference in favour of controls could be caused by differences in learning occurring through the different designs, e.g., better learning for the control condition due to the new, and more salient, cues. Moreover, for the experimental condition cue A already seems to be a weak go cue, so that for it to prime stopping it first needs to prime not going, whereas for the control condition cue O is a novel cue and so has no prior learning attached to it. Overall, the results show that learning

was generally poor, though the contrast for reaction times indicates that those in the control condition were learning the discrimination between a go and a stop cue.

#### **4.5.2.3 Test phase**

The analyses conducted for this test phase differ from those undertaken for Experiments 4 – 6. Whereas in these experiments two sets of analyses were conducted, in this experiment the focus is subtly different. The focus is now not on understanding what learning occurs generally or exploring the effect of sequences, but rather seeing if cue A is more go or stop between groups. However, given that the test phase for Experiment 7 was identical to Experiment 4 – 6 repetition priming effects are still an issue. As such, the analysis for Experiment 7 at test used data from cues that were the first trial of a sequence to investigate differences in behaviour towards cue A between groups.

It was important to ascertain that learning had occurred in this phase, and so the G vs. B contrasts are reported first. Next, contrasts were undertaken to explore the effect of the inhibition training upon learning to cue A. These compared A against G and against B. If the training was effective, then it would be expected that the experimental group would see cue A as more stop than those in the control group. As before a standard alpha level was used for the G vs. B contrast, but other contrasts used a corrected alpha level of .017 (see Table 4.13 for descriptive statistics). Separate models were run for each condition for each performance measure (see Table 4.14 for AICs).

##### **4.5.2.3.1 Contrasts for G vs. B**

##### **4.5.2.3.2 Reaction times**

For the experimental condition the contrast for G vs. B (conditional  $R^2 = 0.72$ ) was highly significant,  $t(225) = 4.80$ ,  $p < .001$ , 95% CI [43.57, 103.64],  $d = 0.64$ , with faster response times to cue G than B. For the control condition (conditional  $R^2 = 0.70$ ) the contrast was also significant,  $t(225) = 2.38$ ,  $p = .018$ , 95% CI [7.17, 74.15],  $d = 0.32$ , again with faster responses to G than B. These contrasts indicate that for both conditions associative learning had taken place by the test phase. An interaction model (conditional  $R^2 = 0.71$ ) including both

conditions found a non-significant difference between G vs. B between groups,  $t(450) = -1.44$ ,  $p = .152$ , 95% CI [-77.93, 12.04],  $d = -0.14$ .

#### **4.5.2.3.3 p(respond)**

For p(respond), the contrast in the experimental condition was not significant,  $z = -0.14$ ,  $p = .160$ . However, the same contrast for the control condition was significant,  $z = 2.45$ ,  $p = .014$ , with more errors in the presence of cue G than B. The interaction model was significant,  $z = 2.81$ ,  $p = .005$ , indicating better learning of the G vs. B contrast in the control condition against that of the experimental condition.

#### **4.5.2.3.4 Contrasts for G vs. A**

##### **4.5.2.3.5 Reaction times**

For the experimental condition the G vs. A contrast was significant,  $t(225) = 3.12$ ,  $p < .002$ , 95% CI [17.83, 77.90],  $d = 0.42$ , with faster response times to cue G than A, suggesting A was not as much a go cue as G. For the control condition the contrast was also significant,  $t(225) = 3.28$ ,  $p = .001$ , 95% CI [22.56, 89.55],  $d = 0.44$ , again with faster response to G than A. If the training had been effective, then it would be expected that the difference between G vs. A would be greater in the experimental condition than the control condition, yet an interaction model found no significant differences between groups,  $t(450) = 0.36$ ,  $p = .721$ , 95% CI [-36.79, 53.18],  $d = 0.04$ .

##### **4.5.2.3.6 p(respond)**

For the experimental condition the contrast G vs. A was significant,  $z = -2.53$ ,  $p = .011$ , with more errors to A than G. However, for the control condition the contrast was not significant,  $z = -1.11$ ,  $p = .266$ . The interaction between the two conditions was significant,  $z = 1.49$ ,  $p = .016$ , with a greater difference between G vs. A in the experimental condition compared to the control condition. However, this difference was in the opposite than expected direction (based on the assumption that the inhibition training would cause A to become a stop cue in the experimental group), with cue A being more of a go cue than cue G.

Condition	Reaction Time		p(respond)		p(miss)	
	Mean	SD	Mean	SD	Mean	SD
Training						
A	492.56	114.49	0.08	0.11	0.02	0.00
G	444.7	95.03	0.04	0.07	0.03	0.05
R	520.67	112.59	0.05	0.08	0.02	0.00
RA	494.62	96.27	0.06	0.09	0.02	0.00
B	518.3	105.44	0.05	0.05	0.02	0.01
I	513.35	91.89	0.04	0.04	0.02	0.02
IP	532.86	114.11	0.03	0.03	0.03	0.04
J	538.97	107.56	0.04	0.07	0.03	0.05
P	483.54	83.89	0.05	0.09	0.02	0.02
YZ	540.94	87.66	0.04	0.05	0.03	0.04
Control						
A	517.31	116.78	0.09	0.13	0.02	0.00
G	461.25	121.28	0.07	0.12	0.03	0.05
R	527.08	102.97	0.07	0.12	0.02	0.00
RA	512.38	124.44	0.07	0.10	0.03	0.05
B	501.92	109.53	0.03	0.03	0.02	0.01
I	511.54	106.79	0.06	0.07	0.04	0.04
IP	541.3	93.13	0.03	0.03	0.04	0.04
J	511.66	122.86	0.07	0.10	0.03	0.05
P	506.6	94.13	0.06	0.06	0.02	0.02
YZ	531.79	102.94	0.07	0.07	0.03	0.03

Table 4.13. Descriptive statistics for test by condition for Experiment 7. Reaction times means are calculated using are raw data, but mean p(respond) and p(miss) use transformed data. All data is from first sequence trials.

#### 4.5.2.3.7 Contrasts for A vs. B

#### 4.5.2.3.8 Reaction times

The difference between A and B for the experimental group was marginally significant at a standard alpha level,  $t(225) = 1.68$ ,  $p = .094$ , 95% CI [-4.29, 55.78],  $d = 0.22$ , with faster response times to A hinting at a trend for cue A not to promote a stopping response. For the control condition the difference was not significant,  $t(225) = -0.90$ ,  $p = .369$ , 95% CI [-48.88, 18.10],  $d = -0.12$ . If the

training was effective the difference between A vs. B would be expected to be significantly smaller in the experimental condition compared to the control condition. The difference was only marginally significant at a standard alpha level,  $t(450) = -1.79$ ,  $p = .074$ , 95% CI [-86.12, 3.85],  $d = -0.17$ , with the difference between A and B being greater in the experimental condition than the control condition - with faster responses to cue A than cue B.

#### **4.5.2.3.9 p(respond)**

For commission errors, in the experimental condition the contrast was not significant,  $z = 1.43$ ,  $p = .153$ , with little difference in error rates between the two cues. However, the control condition had a significant difference between A and B,  $z = 3.10$ ,  $p = .002$ , with more errors for cue A than B suggesting that in the control condition A was a go cue. The interaction was significant (but only at the standard alpha),  $z = 2.16$ ,  $p = .031$ , with a greater difference between A vs. B in the control condition compared to the experimental condition. This could be taken as evidence of those in the control condition having less inhibition to cue A. However, the means show that cue A was experienced in a similar fashion for both groups, and rather, the difference lies in cue B. Those in the control condition have lower errors than those in the experimental condition. Thus, the interaction effect is seemingly caused by cue B prompting more stopping in the control condition, rather than less stopping to cue A.

#### **4.5.2.4 Test phase summary**

In conclusion, with regards to the test phase there is clear evidence (in the response times) that both conditions were learning the discrimination G vs. B and that this was in the expected direction. However, the commission error data indicates that those in the control condition learnt the discrimination better than the experimental group. The data suggests that inhibition training did not affect learning to cue A. For the G vs. A contrast both conditions had significantly slower responses to A than G, suggesting that A was not a strong go cue. For commission errors, the results suggest that those in the experimental condition had learnt more about the difference between G vs. A than the control condition (perhaps expected due to the increased exposure to cue A from the inhibition training). However, the direction of the effect indicated that cue A promoted more going than cue G. For the A vs. B contrast, the experimental condition had marginally faster responses to A than B, giving a hint of a trend for cue A to



promote going over B for this group. The difference for the control condition was not significant, with the interaction model supporting the view that those in the experimental condition had learnt more about A vs. B than the control condition, with this learning, albeit only marginally significant at a standard alpha level, priming going to A rather than B. Interestingly the pattern was reversed for commission errors, with the control condition showing significantly greater errors to A than B (suggesting A was more of a go cue), with the experimental contrast being non-significant. However, while the interaction was significant, the means indicated that the difference could be accounted for by the greater stopping witnessed to cue B in the control condition compared to the experimental condition. Overall, there is no evidence to support the conclusion that inhibition training was effective.

Model	Experimental condition	Control condition
Reaction time models		
Main effects of cue	3152.98	3195.10
Main effects of cue with random intercept	<b>2853.45</b>	<b>2906.01</b>
p(respond)		
Main effects of cue with Gamma family and inverse link	-1102.97	-980.18
Main effects of cue with Gamma family and inverse link and random intercept	<b>-1164.66</b>	<b>-1021.74</b>

Table 4.14. AICc scores for test models for Experiment 7 run on reaction time and p(respond) data. Bold are the models chosen.

### 4.5.3 Summary

Experiment 7 served to replicate the findings of Experiment 6 in the training phase and began to investigate the merits of using inhibition training to change learnt behaviour towards amber traffic lights, or in this context, cue A. On the

first objective Experiment 7 was successful, but it was unsuccessful for the second.

In terms of replicating Experiment 6, the experiment found broadly similar findings for the training phase. This gives confidence that the findings of Experiment 6 (notably the strong stop response to cue R) are in fact due to the experimental manipulation rather than error variance. It is clear that the most complete design of the paradigm presented in this thesis, in terms of capturing several aspects of the contingencies experienced at UK traffic lights, suggests that cue A is a weak go cue and R a strong stop cue.

Regarding the inhibition training the picture is less positive. The results from the inhibition phase itself suggest that participants did not learn much during this task, with only the control condition for the GO vs. STOP contrast for response times suggesting any learning had taken place. This could be due to the phase only having two blocks. While increasing the number of blocks is an obvious suggestion to increase the chance of learning occurring, it must be balanced against the overall length of the experiment - it would be undesirable for it to become too long. Given these results, the outcome of the test phase was perhaps a foregone conclusion. Overall, there is little evidence to suggest that those in the experimental group had shifted their response to cue A towards stopping. Indeed, the evidence seems to suggest that, compared to the control group, the experimental group was more inclined to treat cue A as a go cue. Following the trend in the inhibition phase, the control condition seems to be better learners of a basic Go vs Stop contrast, seemingly learning the G vs. B contrast better than the experimental group.

At test, there is evidence of the expected significant G vs. B contrast - with both conditions having faster responses to cue G than B. It is worth reflecting on this. The results show that despite completing a five-minute filler task (the inhibition phase), where participants did not see cue G or B, participants were still retaining the earlier learning about G and B. These results show the strength of learning to the cues resulting from training. They also speak to the null findings resulting from the inhibition phase and indicate that more, or a different version of, inhibition training would be needed to overcome learnt responding arising from the training phase.

## **4.6 GENERAL DISCUSSION**

This chapter has presented research where the ideas from Chapters 2 and 3 were combined in an experimental design that captured both some of the sequence information and (assumed) task sets in operation at UK traffic lights. The results from the two experiments presented here support my past work and those of others (Verbruggen & Logan, 2008b) in showing how pairing a certain response with a cue can lead to that cue priming the response even when the contingencies no longer support this learning. These results should not be dismissed as a foregone conclusion. The designs were complex, with task set (and thus responses) changing frequently. Additionally, the fact that the G vs. B contrasts at test for reaction times were significant in both experiments clearly demonstrates that participants were learning about the cues within the task. This strengthens the validity of the experiment. It also gives confidence in the conclusion that participants are truly learning about the contingencies in play, and that the results are not a product of participants guessing the responses, or applying some arbitrary rule, such as “I’ll respond to every fourth trial” in order to complete the task. At the same time, it must be acknowledged that another consistent finding is weaker (though still reliable) learning overall at test. This may well be due to the fact that it is, in essence, an extinction phase.

### **4.6.1 Learning of traffic light cues**

While the work discussed prior to this chapter can be generally summarised as indicating that cue R is a neutral/weak stop cue, in the work presented here cue R seems to have become a strong stop cue. This is a notable shift from the past work and has at least two possible explanations. The first is that the combination of task set and the embedding of cues in a sequence enabled more effective learning, resulting in cue R coming to effectively prime stopping. The second is that the finding is a result of some unknown issue with the internal validity of both Experiments 6 and 7. Given the results found so far in the thesis it might be tempting to dismiss the findings presented in this chapter as perhaps arising due to participants not understanding the task, or just chance variation in responding. However, they do fit into the research narrative. In Chapter 2, R was a weak stop cue when the task set was stop but was a stronger stop cue when in a go task set. In Experiment 6 and the training phase of Experiment 7, R was also presented in a go task set and seems to be a very

strong stop cue. This line of argument does ignore the results of Experiment 5, where R was fairly weak stop cue, but this could be a consequence of chance fluctuations rather than anything systematic. Overall, if one holds the assumption that Experiment 6 represents a design that is most true to the contingencies experienced at UK traffic lights, then it seems clear that a red traffic light will prime a stopping response, and an amber traffic light a weak go response. In other words, the associative learning for red traffic lights in the designed paradigm conforms to the rules of the Highway Code but that for amber lights does not. This suggests that in addition to training focusing on conscious control, training that targets the associatively-mediated 'amber – weak go' link needs to be developed. This was the motivation for the intervention used in Experiment 7.

#### **4.6.2 Inhibition training**

The inhibition training was not successful, and as discussed earlier, this is likely to be partly due to not enough training to retrain amber as stop. What does this initial foray into inhibition training tell us? Firstly, training needs to be longer to encourage more learning. Secondly, transfer effects might be an issue. This is whether training on one task will generalise to performance on other tasks. In the inhibition task participants responded, or not, to black arrows, while in the test phase they responded to white or coloured circles. Therefore, the tasks are not quite the same, particularly in terms of their surface features, and it may be the case that, even if learning did occur in the inhibition phase, it would not transfer to performance on the test phase. This argument is more formally explored by Simons et al. (2016) who summarises research into inhibition and transfer effects by concluding that inhibition training can improve performance for a practiced task or near identical one, but that the range of transfer is limited, even for related tasks (see also Noack, Lövdén, Schmiedek, & Lindenberger, 2009). These issues will be key considerations for my design of an inhibition intervention aimed at increasing stopping to amber in the next chapter.

#### **4.6.3 Conclusion**

To conclude, the experiments presented in this chapter mark the final development of the laboratory-based paradigm designed to capture contingency learning at UK traffic lights. Overall, this final iteration, one assumed to be the

truest to life, suggests that a red traffic light (or at least the cue that has the same contingencies) develops an associatively based stop response, while an amber traffic light becomes linked to a weak go response. So far, the thesis has focused on exploring *what* is learnt at an associative level. However, Experiment 7 also marked a shift from exploring learning of the contingencies experienced at UK traffic lights to investigating if these learnt, associatively-mediated, contingencies can be changed. Although this attempt at intervention was ultimately unsuccessful, it provides a starting point for ideas further explored in the next chapter, where I present three experiments that used the concept of response inhibition training in a bid to bring about behaviour change to amber traffic lights in a more realistic setting.



### Inhibition Training I: The STOP-CHANGE Paradigm

Chapter 4 introduced the initial testbed for using inhibition training to shift learning of the cue representing amber from go to stop. It was clear that the training needed improving. In this chapter I work towards a design of an inhibition task that shows some promise in changing people's driving behaviour.

#### 5.1 COGNITIVE CONTROL IN DRIVING

As discussed in Chapter 1, there have been calls for interventions aimed at improving driving behaviour to consider associatively-mediated, as well as conscious processes. Decision making is often conceptualised as an interaction between goal-directed deliberate processing and more automatic processes (see the concepts of cold and hot cognition by Kahneman, 2011), and departures from rational decision making have been attributed to cognitive (Tversky & Kahneman, 1973), affective (De Martino, Kumaran, Seymour, & Dolan, 2006) and learnt associatively-mediated processes shared between humans and monkeys (M. K. Chen, Lakshminarayanan, & Santos, 2006; Lakshminarayanan, Keith Chen, & Santos, 2008).

Driving is a complex activity that requires the intermixing and appropriate deployment of sensory, motor, and cognitive abilities (Anstey, Wood, Lord, & Walker, 2005). It is of no surprise that there is considerable evidence to suggest that cognitive control is needed to maintain safe driving, with poorer cognitive abilities being associated with riskier driving (Walshe, Ward McIntosh, Romer, & Winston, 2017). Of relevance to this thesis is the link between poor response inhibition and risky driving. For example, low levels of inhibition have been shown to be strongly correlated with self-reported driving violations (Tabibi, Borzabadi, Stavrinou, & Mashhadi, 2015). Fischer, Barkley, Smallish, and Fletcher (2007) found that those with attention deficit/hyperactivity disorder have significantly more real-world traffic violations compared to matched normal controls, while O'Brien and Gormley (2013) showed that traffic offenders had worse response inhibition compared to non-offenders, though the results in this study were not completely conclusive. Additionally, Brown et al. (2016) found that, compared to drivers without driving convictions, those convicted for

speeding had significantly impaired performance on a task measuring inhibitory control (and see Hatfield, Williamson, Kehoe, & Prabhakaran, 2017). To summarise, the research indicates that poor response inhibition can contribute to dangerous driving, suggesting that targeting inhibition could lead to safer driving.

## **5.2 TRAINING CONTROL**

Given that executive control seems to be important to safe driving it is logical to ask whether improving control, namely inhibition, can bring about safer driving. Such training involves pairing arbitrary cues with stop responses to increase overall inhibition, and thus can be characterised as strengthening the inhibition 'muscle' to bring about global improvements. As noted in Chapter 1 such training has generally been ineffective. For example, N. S. Lawrence, Verbruggen, et al. (2015) found no effects of such training upon food consumption after training in a student population. The study consisted of three groups. Two groups received non-cue specific training whereby participants inhibited responses to random images, or to specific categories of non-food images. In the third condition participants had to make double responses to specific categories of non-food images. This training was designed to increase impulsivity, and thus result in increased consumption, see Guerrieri et al. (2012). Food intake was strikingly similar across all conditions, suggesting that non-cue specific training did not decrease food consumption compared to a group primed for general disinhibition. Yet, work by Berkman, Graham, and Fisher (2012) suggests that such general inhibitory training might work better in children than adults by supporting the development of abstract rule use (Munakata, Snyder, & Chatham, 2012). Overall, it seems that training non-cue specific inhibition would be an unsuccessful route to improving driving behaviour. However, behaviour change has been more successful in an associative context.

## **5.3 LEARNT CONTROL**

As discussed throughout this thesis there is strong evidence that pairing a cue with a certain response leads this cue to become able to prime that response. Furthermore, there is evidence that associative learning within the task



designed in this thesis to mirror the contingencies at UK traffic lights promotes a go response to amber lights, whereas the Highway Code dictates a stop response. Combining these two ideas, Experiment 7 explored whether it would be possible to train an associatively-mediated stop response to amber lights. The evidence supporting this idea has already been covered in Chapters 1 and Chapter 4 and so will not be repeated here. Generally, the research has shown limited promise in training learnt responses to specific cues – though the work in reducing food consumption has shown more success (N. S. Lawrence, O'Sullivan, et al., 2015). Yet, there is good reason to support the application of specific inhibition training to driving. The work by Briggs et al. (2017) mentioned in Chapter 1 showed that whilst under a cognitive load drivers' behaviour was impaired (and see House of Commons Transport Committee, 2019), with drivers seemingly relying on past experiences and schemas. The notion here is that control processes are disrupted by heavy demands, while associatively-mediated processes are relatively unaffected, and that for most people driving is relatively automated (Shinar, Meir, & Ben-Shoham, 1998). Therefore, addressing these processes could result in desirable behaviour change.

## **5.4 PRESENT EXPERIMENTS**

The work presented in this chapter aims to go beyond the traffic light paradigm designed thus far, and rather than investigate what participants learnt at UK traffic lights at an associative level, start to change this learning. Due to this shift in motivation the designs are substantially different from preceding experiments in this thesis. However, the fundamental task of training a cue to become linked to a certain response, and then exploring learning in a test phase, remains the same.

One change from past experiments is that those reported in this chapter will not use go/no-go training to train responses, but rather make use of another paradigm, the STOP-CHANGE task. As outlined in Chapter 1 this is where the required response changes mid-trial, rather than being set at the start as in go/no-training. The paradigm requires participants to inhibit an ongoing go response and replace it with an alternative response. This paradigm was chosen for the important reason that it is more realistic in the context of driving. The work here aims to develop a training task that could be used in the real-

world and therefore to increase the likelihood of the training transferring to real-world driving it needs to match as best as possible to motor responses used in driving. So far, I have written about participants stopping or going at traffic lights, yet this does not capture the full range of motor responses required to stop. To actually stop a car, one needs to stop accelerating (one response) and move the foot to depress the brake pedal (another response). Therefore, while going and stopping are perfectly adequate ways to describe the macro actions at traffic lights, the STOP-CHANGE paradigm is more analogous to the micro actions involved in stopping at traffic lights. Regarding the specific mechanisms tapped by the paradigm, as noted in Chapter 1, there is evidence that inhibition does occur in these types of task (Verbruggen & Logan, 2009a). Indeed, pigeons are capable of performing the task with similar performance to a stop-signal task (Meier, Lea, & McLaren, 2018) I am therefore confident that the switch in task will not, by itself, impact learning about the contingencies in the experiment (though note inhibition training has been found to be less effective when using stop signal, from which the STOP-CHANGE task is adapted from, rather than go/no-go tasks, Veling et al., 2017).

## **5.5 EXPERIMENT 8**

This experiment aimed to provide initial evidence of the benefits of inhibition training to promote stopping at amber traffic lights. The work built upon that of Experiment 7 by specifically addressing the transfer issues discussed in relation to that experiment.

One important change is that while the experiments presented in this thesis so far have focused on behaviour in a simple laboratory task, to assay evidence of behaviour change (and in line with other inhibition work, e.g., N. S. Lawrence, O'Sullivan, et al., 2015), a more real-world measure was used. Therefore, behaviour change was measured using performance on a driving computer game where participants 'drive' through a series of junctions and must decide whether or not to stop at the lights (see section 5.5.1.3.2 for further details).

## 5.5.1 Method

### 5.5.1.1 Design

A mixed-measures 4 x 2 design was used with two independent variables: condition (experimental vs. control 1 vs. control 2 vs. control 3) and time (pre- and post-training performance).

In terms of the groups for the experiment, the desired outcome is that the group that receives inhibition training should develop an amber-stop associative link, and, following the analysis just given, this group will receive cue specific inhibition training. That is, participants will see amber circles in the change trials (requiring participants to change motor response mid-trial, see section 5.5.1.3.1 for details). However, the control conditions are more complicated. Simons et al. (2016) argues that one issue in the inhibition field is the use of passive control groups, with any differences between this control and the experimental group not necessarily being due to the training but rather other factors, such as motivation. To avoid this issue, the first control group will be an active control group, with the task closely matched to that of the experimental group. Therefore, control 1 will receive STOP-CHANGE training to purple circles, with the cue still being specific but irrelevant (as purple circles are not linked to any rules regarding UK traffic lights). This design opens up the issue that rather than the training per se, it is exposure to amber circles, through a process of directing attention towards these cues, which could account for any behaviour change. To counter this, control 2 will not receive inhibition training and instead the STOP-CHANGE trials will be replaced by simple presentation of amber circles, that is a relevant cue in a driving context. Control 3 will receive purple circles instead of the inhibition trials, which can be deemed irrelevant with respect to driving. In a sense, control 2 is a control for the experimental group and control 3 a control for the control 1 group. If the inhibition training is the key determinant of behaviour change then the greatest improvement should be seen in the experimental group, followed by control 1 (non-specific inhibition training effect), then control 2 and control 3 (which would be expected to be roughly equal). However, if exposure to amber lights is also a factor then the order could be experimental group (both effects contributing), control 1 and control 2, then finally control 3. By improvement I mean the biggest uplift in stops to amber traffic lights after training compared to pre-training performance.

Thus, the four conditions were: cue-specific inhibition training (experimental group), cue irrelevant inhibition training (control 1), relevant single trials (control 2), and irrelevant single trials (control 3). Condition was manipulated between-participants and time was manipulated within-participants. Participants were randomly allocated to one of the four conditions. Responses were obtained on one dependent measure: number of stops at amber traffic lights in the driving game.

### **5.5.1.2 Participants**

At the time of recruitment, this experiment was the first to undertake inhibition training to change driving behaviour (now see Hatfield et al., 2018). Due to this I felt it would be inappropriate to assume an effect size, therefore I decided to test 30 in each group, a total of 120 participants. This sample size is similar to other intervention studies (e.g., Porter et al., 2018). One-hundred and thirty-eight participants participated in exchange for payment of £5 or one course credit (see Results section for details on the outlier removal process). The inclusion criteria were that participants had to be 18-65 years old, hold a full or learners driving licence (of any nationality), have normal or corrected vision, and not be colour blind.

### **5.5.1.3 Materials**

#### **5.5.1.3.1 Inhibition task**

The STOP-CHANGE task was designed to train an associatively-mediated stop-and-change response to amber traffic lights. The task involved participants having to stop their primary task response and replace it with a secondary task response (see Figure 5.1). Throughout the task there were only two responses, either a left-hand response (the left-ctrl key) or a right-hand response (the right-ctrl key) which participants were instructed to make with their left or right index fingers respectively on a standard QWERTY keyboard.

The primary task was to respond with the right key when a green circle (70% of trials) was presented, and with a left key response when a red circle (15% of trials) was displayed. The remaining 15% of trials introduced the experimental manipulations. For those in the experimental and control 1 conditions these trials were STOP-CHANGE trials in which a green circle was displayed but changed to amber for the experimental condition or to purple for those in control

1, with participants in both conditions being required to suppress their initial green right hand response and instead execute the appropriate secondary left hand response. For the control 2 and 3 these trials were not stop-change but either displayed a single amber (control 2) or purple (control 3) circle, requiring a left-hand response. Thus, the main difference with the other conditions is that the amber/purple appeared immediately, and the task became a simple 3-choice task because no change in response was required.

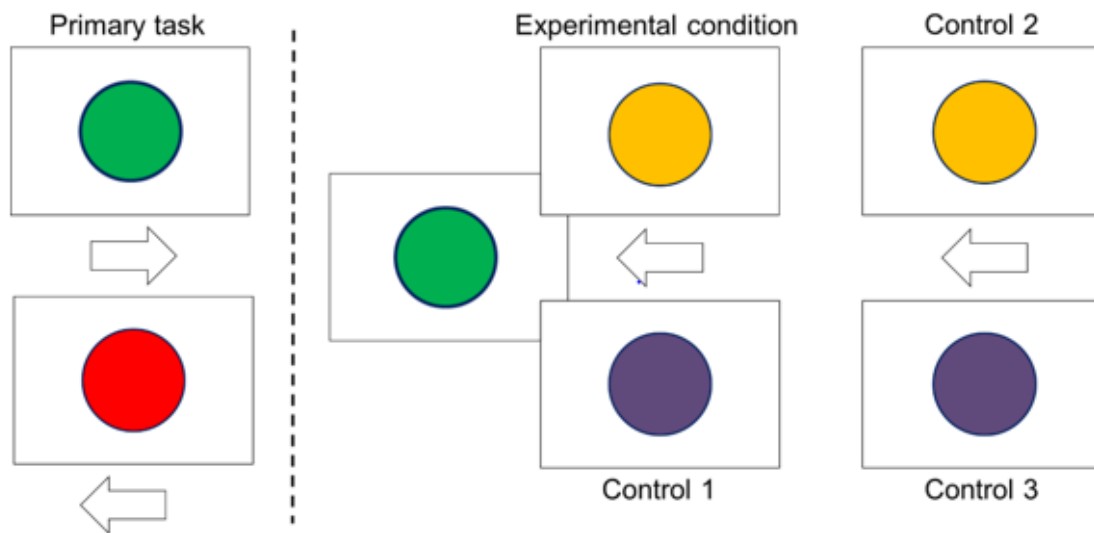


Figure 5.1. Schematic of the inhibition training task for Experiment 8. All groups received two blocks of 100 trials. Participants saw 70 green trials and 15 red trials per block. The remaining 15 trials depended on the condition (see right-hand side of the Figure).

The direction of the responses mapped onto the pedals used whilst driving. The green circle which required a right-arrow response mapped onto the accelerator pedal, while the other colours required a left-key response mapping onto the brake pedal. The task was designed so that experimental STOP-CHANGE trials simulated UK traffic light signals changing from green to amber, with the required response of changing from a right hand to a left hand response imitating the motor actions needed to stop at such traffic lights, i.e. lift the right foot off the accelerator pedal and depress the left brake pedal. The resulting learning should be that participants learn an amber = inhibit right pedal, change to left pedal response, which should transfer to the real-world as amber – stop accelerating and brake. Thus, the task is training participants to make a certain

response to amber cues that is the opposite (in some sense) to that made to green cues, and this should result in braking whilst driving.

The colours of the red, green and amber circles corresponded to the RGB codes of UK traffic lights, measured  $6.7\text{cm}^2$  and were presented in the centre of a 19-inch PC monitor on a white background. The task required participants to press and hold their response for 500ms and would only continue when the correct response was made. The task consisted of two blocks of 100 trials, with a 15 second break between each block. While Experiment 7 did indicate the importance of having lengthy training, in order to keep training as short as possible, given the desire to implement it in the real-world, it was decided to retain the same length of training as in Experiment 7.

Trial order was randomised within each block and there was a fixed 1500ms inter-trial-interval. The change feature of STOP-CHANGE trials happened after a random delay of 100ms, 150ms or 200ms following the presentation of the green circle, with these timings being randomised across the trials but occurring equally within blocks. Participants received no feedback during the task.

In one sense the inhibition task for this experiment is similar to that of Experiment 7, in that participants respond (or not) to shapes. However, one change is that now the task has cue and signal combined, unlike Experiment 7 where participants saw a cue and then responded to the signal. This change was made to better reflect traffic lights, where a driver only sees one stimulus (the light) which acts as both cue and signal.

#### **5.5.1.3.2 Driving task**

Participants completed a driving computer game before and after the training. The Stop Light Task (adapted from Chein, Albert, O'Brien, Uckert, & Steinberg, 2011) involved participants 'driving' a car from the first person point of view along a straight road at a set speed with the goal of reaching the destination within the time limit (8 minutes). Participants were instructed to press and hold the right-ctrl key when they wanted to accelerate the car and to release that key and press the space bar when they wished to brake. The right-ctrl key had no impact upon the game but was used to create the same motor responses required when driving and braking that is of stopping one response and changing to another (and to match motor responses in the inhibition task). In the

game participants crossed 30 traffic light-controlled junctions. For nineteen of the junctions the lights turned from green to amber as the car approached. Participants then had to decide whether to brake the car (by pressing the space bar) and wait three seconds for the traffic light to turn green, or to cross on the amber light (make no response) and risk a crash, which incurred a six second penalty. However, unbeknown to participants, the game was programmed so that it was impossible to crash (this was to avoid any emotive reactions influencing task performance). The remaining ten junctions displayed an equal mix of red and green traffic lights.

The distance between junctions varied from between 10 and 16 seconds. For amber light trials, the light turned amber between 2 and 4 seconds before the participant reached the junction, and 8 seconds for red trials. The order of traffic lights displaying red, green and amber was randomised, and thus were unpredictable to participants, replicating a natural driving experience. A vehicle sometimes crossed the junction ahead to create the feeling of a busy highway. The type of vehicle was randomised to prevent one vehicle type being associated with a specific colour traffic light. Performance on this task was measured by the number of stops at amber traffic lights.

#### **5.5.1.4 Questionnaires**

To ensure that the four groups were well matched on characteristics that could have influenced task performance, participants completed two questionnaires following the second driving task (at the end of the experiment).

##### **5.5.1.4.1 Driving experience**

Participants were asked what type of licence they had and if they had either been involved in a road traffic collision either as a passenger or a driver. Participants with learner licences were asked if they had any driving experience. This questionnaire was of my own design.

##### **5.5.1.4.2 Impulsivity and sensation-seeking**

High scores on these traits have been linked to an increased propensity to engage in red traffic light-jumping (Burgess, 2003) and commit other traffic violations (Curran, Fuentres, Alfonso, & Hennessy, 2010). Therefore to account for individual differences in this regard, participants completed the Impulsive Sensation Seeking Scale from the Zuckerman-Kuhlman Personality

Questionnaire (Zuckerman, Kuhlman, Joireman, Teta, & Kraft, 1993; Questionnaire taken from Zuckerman and Aluak, 2014).

#### **5.5.1.5 Procedure**

Participants first performed the driving game which lasted up to eight minutes. Participants received instructions telling them to drive normally and to imagine that they were driving to a party (to help create a more naturalistic scenario). They were told to arrive as quickly as possible, within the eight-minute time limit, but also safely. It should be noted that the time was more than adequate to complete the task and that all participants did complete it within the timeframe. Next, participants were randomly assigned to a training condition and completed the inhibition training task (which lasted around 15 minutes). Participants then completed another driving game. Following this, participants completed the questionnaires.

#### **5.5.2 Analysis and results**

Data was processed and analysed using R (R Core Team, 2018). Data was analysed using, at least in cognitive psychology, the more traditional generalized linear models (computed through the R package Ez; M. A. Lawrence, 2016) rather than the mixed effects models used up to now. This was to allow for more appropriate, in a statistical sense, comparisons between the results reported here and those of past intervention studies. In terms of reported effect sizes, Cohen's D value has been corrected in line with the formulas provided in Lakens (2013) and, when reporting ANOVAs, generalised eta squared have been provided, rather than partial eta squared, as they provide greater comparability between within- and between-participants designs (Bakeman, 2005).

Prior to analysis, the data of the 138 participants from Experiment 8 were combined with the 80 participants from my masters' experiment (which had the same design as this experiment, this allowed for increased power) and the exclusion criteria applied to this combined dataset. Eighteen participants were removed due to experimental issues which meant that they could not complete the study. No participants were removed for having less than 60% accuracy for green trials. The 60% accuracy threshold is a standard threshold in inhibition studies and ensures that participants are responding above chance to the go



stimuli. Two participants were removed for having response times to green trials of greater than 3 IQRs (i.e.  $2 \times 1.5\text{IQR}$ ; the same exclusion criteria as applied in the other work in this thesis). Eight participants were removed due to having greater than 10% of a trial type in the training task with response times under 150ms (this threshold was chosen as responses faster than 150ms suggest participants were responding before the cue appeared on screen). Nine participants were removed for having no driving experience and 85 participants were removed for not stopping at all red traffic lights in the first driving game. There is an argument to be made that removing those who did not stop at all red traffic lights is akin to removing those who most need the training, and where it might be expected to deliver the largest effects. However, it can be assumed that most drivers do stop at red traffic lights in the real-world, and so those who did not stop at these lights are likely not to be obeying the task instructions in the driving game (this also supports the issue of ecological validity raised later in the discussion of this experiment). The final sample was 96 participants, 38 from my Masters work and 58 from Experiment 8. Conditions were similar sizes, with 27, 24, 20, and 25 participants for the experimental condition, control 1, control 2, and control 3 respectively. Post-hoc analysis using G\*Power 3.1.9.2 (Faul, Erdfelder, Lang, & Buchner, 2007) indicated that the experiment had a 91% chance of detecting an effect size of  $f = 0.20$  for the interaction between time and group at an alpha level of 0.05. The effect size of 0.20 was chosen as it represents a small to medium effect size, which was deemed appropriate for such training to have practical implications for real life.

#### **5.5.2.1 Subject characteristics**

As with many surveys, missing data was an issue. Upon inspection, the missing data was confined to a few participants. Given that answers were opt in, the non-answering is perhaps more likely to be caused by participants skipping questions due to time demands rather than not wishing to give answers due to sensitivity reasons. Missing data was replaced where possible through a by-participant mean replacement strategy. In these cases, the mean for a participant on a particular scale without the missing data was calculated and this mean was then entered into the empty cell(s). In cases where this was not possible, e.g., age, the missing data was removed from analysis as long as the total percent of missing data was not more than 20% of the overall sample (See

Appendix D more detail on the particulars for each scale). What is clear is that missing data overall was relatively low, and therefore is unlikely to materially impact upon conclusions. Randomisation checks revealed that the four conditions were well matched (see Appendix E). There was a significant difference between the four conditions in crash history, yet given the small sample size, especially across the four conditions this is likely to be a result of chance variation, and thus can be taken as a spurious finding (Button et al., 2013). Indeed, applying a Bonferroni correction meant that none of the contrasts were significant. One issue to note is that due to a bug in the driving game it was possible to crash the car at red traffic lights if a participant pressed the spacebar too late resulting in the car stopping across the junction. Removing these participants and running the main analysis only produced a quantitative difference and so these participants were retained, with only the full model being reported.

The analysis is split into two sections. Firstly, analysis to assess behaviour in the training task is presented, followed by the analysis of behaviour change between the two driving games.

### **5.5.2.2 Response inhibition task**

#### **5.5.2.2.1 First reaction times to correct GO trials**

Mean responses to GO trials were calculated by excluding trials where an incorrect key was pressed and trials that had response times under 150ms. A mixed measures ANOVA with condition as the between participants factor and the two blocks as the repeated measure found no significant effect of block,  $F(1, 92) = 0.34$ ,  $p = .561$ ,  $\eta_G^2 = 0.0003$ , indicating that response times were similar across the two blocks, with block 1 having a mean of 409.91ms (SD = 80.39) and block 2 a mean reaction time of 407.42ms (SD = 76.58). The main effect of condition was also not significant,  $F(3, 92) = 0.17$ ,  $p = .919$ ,  $\eta_G^2 = 0.005$ , showing that response times were similar across conditions (see Table 5.1 for condition by block descriptive statistics). The interaction term in the ANOVA was not significant,  $F(3, 92) = 0.25$ ,  $p = .862$ ,  $\eta_G^2 = 0.0006$ .

#### **5.5.2.2.2 First response accuracy rates to correct GO trials**

Accuracy rates to GO trials were calculated by excluding trials where an incorrect key was pressed and for trials that had response times under 150ms.

A mixed measures ANOVA with condition as the between participants factor and the two blocks as the repeated measure found no significant effect of block,  $F(1, 92) = 2.68, p = .105, \eta^2 = 0.008$ , indicating that accuracy did not differ across blocks. This was possibly due to ceiling effects as accuracy was high across both blocks, with block 1 having an accuracy rate of 94.38% (SD = 2.77) and block 2 a rate of 94.97% (SD = 2.27). The main effect of condition was also not significant,  $F(3, 92) = 1.37, p = .257, \eta^2 = 0.031$ , showing that accuracy was similar across conditions (see Table 5.1 for condition by block descriptive statistics). The interaction term in the ANOVA was not significant,  $F(3, 92) = 0.545, p = .653, \eta^2 = 0.005$ .

Condition	GO trials block 1		GO trials block2	
	RT	A%	RT	A%
Experimental	417.56 (85.53)	93.97 (2.90)	416.13 (75.55)	93.97 (2.37)
Control 1	408.70 (85.49)	93.33 (3.02)	408.74 (78.87)	94.36 (2.18)
Control 2	404.45 (71.38)	94.43 (2.45)	405.08 (72.62)	95.71 (2.19)
Control 3	407.17 (80.75)	95.43 (2.58)	398.60 (82.00)	95.29 (2.29)

Table 5.1. Descriptive statistics for GO trials for Experiment 8. SD given in parenthesis. RT = response times. A% = percentage accuracy.

### 5.5.2.2.3 First reaction times to correct CHANGE trials

Response times to CHANGE trials were defined as the difference between the time for a response on a trial and the change delay for the change signal (including any computer lag time, often around 17ms). Error trials, trials where a response was made before the change signal or where an incorrect key was pressed in response to the change signal, and trials with response times under 150ms were removed. Participants with missing values in either blocks were

removed for the analysis (nine in total). As only the experimental and control 1 groups experienced CHANGE trials this analysis was limited to these groups. I ran a mixed-measures ANOVA on these data with condition as the between factor variable and the two blocks as the repeated measure component. The main effect of block was not significant  $F(1, 40) = 2.16, p = 0.150, \eta^2 = 0.013$  with similar response times between blocks ( $M = 530.94, SD = 273.76$  for block 1, and  $M = 482.14, SD = 140.45$  for block 2). The effect of condition was also non-significant,  $F(1, 40) = 1.44, p = .237, \eta^2 = 0.03$ , indicating that response times were similar between conditions (see Table 5.2 for condition by block descriptive statistics). The interaction term was also not significant,  $F(1,40) = 1.164, p = .208, \eta^2 = 0.010$ .

#### **5.5.2.2.4 First response accuracy rates to correct CHANGE trials**

For this measure accuracy is defined as participants making their first response to the change trial. As in the case of the GO cue analysis, incorrect key presses and trials with response times under 150ms were removed. Given that 0% (i.e. no accurate responses) is a possible outcome, there were no missing values for this analysis. I ran a mixed-measures ANOVA on these data with condition as the between factor variable and the two blocks as the repeated measure component. The main effect of block was not significant  $F(1, 49) = 0.21, p = 0.652, \eta^2 = 0.0005$  with similar accuracy rates between blocks ( $M = 40.13\%, SD = 3.89$  for block 1, and  $M = 41.33\%, SD = 4.35$  for block 2). The effect of condition was not significant,  $F(1, 49) = 0.32, p = .564, \eta^2 = 0.006$  (see Table 5.2 for condition by block descriptive statistics). The interaction term was also not significant,  $F(1, 49) = 0.08, p = .781, \eta^2 = 0.0002$ .

Overall, the results suggest that training was experienced in a similar manner for all participants. However, it seems that learning was weak without the expected improvement (faster and more accurate responses) over time. The fact that there were non-significant differences between both blocks of GO trials for response times and accuracy might suggest learning was limited, but certainly for accuracy there might be a ceiling effect with accuracy in the first block already being high. In terms of the CHANGE trials, both measures suggest that both conditions were similar, which is reassuring given I wished to have an active control.

Condition	CHANGE trials block 1		CHANGE trials block2	
	RT	A%	RT	A%
Experimental	579.46 (346.69)	38.53 (3.64)	493.82 (164.91)	39.00 (5.85)
Control 1	466.26 (101.91)	41.93 (4.22)	466.57 (101.66)	43.87 (4.60)

Table 5.2. Descriptive statistics for CHANGE trials for Experiment 8. SD given in parenthesis. RT = response times. A% = percentage accuracy.

### 5.5.2.3 Driving game analysis

#### 5.5.2.3.1 Effect of training on stopping at amber traffic lights

The data was analysed with a mixed 4 x 2 (condition x amber stops) ANOVA, with condition as the between-participant factor and the number of amber lights stopped at in each of the two driving task sessions as the dependent variable. The analysis revealed a non-significant main effect of condition on the number of stops at amber traffic lights,  $F(3, 92) = 1.80, p = .152, \eta^2 = 0.05$ . A significant effect of time was found,  $F(1, 92) = 10.67, p = .002, \eta^2 = 0.01$ , with participants stopping at more amber traffic lights pre-training ( $M = 10.22, SD = 5.68$ ) than post-training ( $M = 8.99, SD = 7.22$ ). The interaction between time and condition was not significant  $F(3, 92) = 0.02, p = .964, \eta^2 = 0.0003$ , indicating that the training had no effect upon behaviour at amber traffic lights (see Table 5.3 for descriptive statistics). Overall, these results suggest that the training did not impact upon stopping at amber traffic lights. Further planned analysis revealed that pre-training there was a marginally significant difference in the number of stops across conditions,  $F(3, 92) = 2.37, p = .076, \eta^2 = 0.072$ , suggesting that before training there were already differences in the groups. It is likely that these differences would have weakened any training effect and thus make it hard to come to a clear conclusion about the effectiveness of the inhibition task. At test there were no significant differences between groups,  $F(3, 92) = 1.20, p = .314, \eta^2 = 0.038$ , and suggesting that the training did not positively impact upon task performance.

Game	Condition	Mean	SE
Pre-training Game	Experimental	10.22	1.11
	Control 1	7.83	1.09
	Control 2	10.95	1.21
	Control 3	11.92	1.13
Post-training Game	Experimental	9.15	1.35
	Control 1	6.58	1.38
	Control 2	9.52	1.76
	Control 3	10.28	1.45

Table 5.3. Descriptive statistics for stops at amber lights for Experiment 8. Stops out of 19.

### 5.5.3 Summary

Experiment 8 investigated whether inhibition training could encourage stopping at amber traffic lights within a computerised driving game. However, no evidence was found to support the effectiveness of the training. In fact, all conditions recorded a decline in stopping at amber traffic lights in the second driving game. This might suggest the training was affecting behaviour, but that it reduced rather than increased stopping. However, it is more likely that the decline in stopping occurs because of experience of the driving game. Over the course of the two games, participants are likely to have learnt that it was not possible to crash at amber lights. Therefore, the most efficient behaviour, especially given the notional time constraints, was to cross all traffic lights. In a sense this learning arguably mimics the development of real-world behaviour. If a driver crosses an amber light and suffers no costs (i.e. does not crash) then they are likely to repeat this behaviour subsequently (as the behaviour is positively reinforced by time saved). Through repeated pairing of accelerating or taking no action when faced with amber lights, the behaviour becomes

associatively-mediated: the prepotent response to amber lights becomes to go, rather than stop. Such a decline indicates that in its current form the training was unable to counteract the association between amber traffic lights and go that may have been learnt in the driving game. The decline could also indicate that participants were not engaging fully with the task for the duration of the experiment. The fact that so many participants failed to meet the inclusion criteria relating to stopping at red traffic lights does suggest task engagement was an issue from the first driving game. Additionally, the driving game does not realistically portray driving (having no sound, ability to steer, or control of the speed of the car beyond braking) and therefore participants' driving behaviour may not be indicative of real-life.

In terms of the inhibition task the overall non-significant effects for blocks is unexpected and goes against the normal improvement in performance (faster responses, better accuracy) with practice. Focusing on go trials, it is likely this lack of improvement over time is due to ceiling effects. Yet, performance for the STOP-CHANGE trials does indicate weak learning (an issue discussed in the general discussion) with the low accuracy (less than 50%) indicating that inhibition was not often successfully employed. Increasing the length of the training, and thus the chance of learning to occur, could lead to more successful inhibition; this has been shown to be important for inhibition training tasks to be effective (A. Jones et al., 2016).

Another issue to consider is transfer. For learning in the inhibition task to transfer, and thus affect behaviour in the second driving task, the two tasks need to be similar. However, this is not obviously the case. For one, the cues were completely different. In the driving game participants saw traffic light shapes while in the inhibition task they responded to simple circles in the middle of the screen. Secondly, the response keys did not match up, with participants using left- and right-Ctrl keys in the inhibition task and spacebar and right-Ctrl in the driving game. In fact, if one assumes that participants learnt that the right-Ctrl key had no impact upon driving (and stopped using it as instructed), then the driving game requires participants to make an active response (press the spacebar) to stop the car, rather than stopping and changing a response. Therefore, the driving and inhibition tasks are not motorically or cognitively

similar. To improve transfer, it will be necessary to reduce the transference gap between the two tasks.

A final issue is the limited sample. Given that the experiment was investigating driving behaviour it would be beneficial to have a range of driving experience. Indeed, it could be argued that those with a longer driving history would be more likely to see amber as a go cue (through years of developing amber – go links), therefore the task might be more effective for more experienced drivers.

In summary, Experiment 8 did not successfully bring about behaviour change. However, several improvements are possible, and Experiment 9 was designed with these changes in mind.

## **5.6 EXPERIMENT 9**

Experiment 9 made several substantial changes to the design above, both in terms of the driving game and the inhibition task. In terms of the driving game, rather than a simple computer game, the experiment used a more realistic simulator. Such simulators have been found to have more validity in correspondence to real-world driving (Aksan et al., 2016; Fors, Ahlstrom, & Anund, 2018), and can predict behaviour at a five year follow up (Hoffman & McDowd, 2010). Regarding the inhibition task, an important change was now that cues were traffic light images rather than coloured circles. This was done to make the task more ecologically valid and thus reduce the transfer gap between the inhibition training and driving games. It should be noted that while the experimental condition in Experiment 8 was designated as cue-specific, this is rather a misnomer as the cues were substantially different in the two tasks. Furthermore, responses were made on a foot pedal response box rather than a keyboard to increase the motor response match between the driving game and the inhibition task. In addition, an extra block of training was added to the inhibition task to improve learning. To speed up recruitment, conditions were collapsed to two groups.

### **5.6.1 Method**

#### **5.6.1.1 Participants**

Given the null findings for Experiment 8 it was again hard to pick an appropriate effect size. An a priori power calculation using G\*Power 3.1.9.2 (Faul et al.,



2007) indicated that a total sample of 84 participants would have a 95% chance of detecting an effect size of  $f = 0.20$  for the interaction between time and group at an alpha level of 0.05. Additionally, recruitment was aimed at older students and staff members. This was to increase the range of driving experience. The recruitment media placed an emphasis on having several years of driving history. Furthermore, the decision was made to limit recruitment to UK drivers only. This was because not all countries use the same traffic light sequences and so it could be the case that non-UK drivers experienced the task in a different manner to those UK drivers.

### **5.6.1.2 Design**

A mixed-measures 2 x 2 design was used with two independent variables: condition (experimental vs. control) and time (pre- and post-training performance). Condition was manipulated between-participants and time was manipulated within-participants. The two conditions were cue-specific inhibition training (the experimental group) and cue irrelevant but specific inhibition training (the control group). Participants were randomly allocated to one of the two conditions. Responses were obtained on one dependent measure: number of stops at amber traffic lights in the driving game.

### **5.6.1.3 Materials**

#### **5.6.1.3.1 Inhibition task**

The same basic paradigm was used from Experiment 8. However, there were several changes (see Figure 5.2). Rather than respond to images of coloured circles participants responded to images of traffic lights. These images were taken from the simulator and so there was a direct match between the inhibition task and the driving task. Participants also responded with foot pedals, though the direction remained the same, e.g., green still required a right-side response. The number of trials per block were the same but with an additional block (total of three blocks). Instead of responding to a purple CHANGE cue, participants in the control condition responded to a blue traffic light – this change was made as it was felt that purple was too near red and blue was a more neutral colour, while still being an irrelevant cue in a driving context. In addition, the response for the CHANGE cue for controls was modified from a left-hand response to a space bar press. This was felt to minimise any rule learning on the part of

control participants around changing response when green lights change, as this could lead participants to change their response (i.e. brake) in the simulator when the traffic lights changed from green to amber. Thus, those in the experimental group had to make a left foot pedal response to CHANGE trials, whilst those in the control group had to press the spacebar. However, many of the features of the training were still controlled for in the control condition. For example, control participants still needed to attend to a change in signal and alter their response accordingly. Finally, the change delay for CHANGE trials was amended to 50ms, 150ms, or 250ms to encourage faster responding to GO trials and to prevent participants guessing the switch times.




Figure 5.2. Schematic of the inhibition training task for Experiments 9. All groups received three blocks of 100 trials. Participants saw 70 green trials and 15 red trials per block. The remaining 15 trials depended on the condition (see right-hand side of the Figure). Arrows indicate direction of foot response.

#### 5.6.1.3.2 Driving task

The simulator comprised a single PC equipped with three 19-inch LED backlit monitors, each with a 1280x1024 screen resolution, a frame rate of 60Hz and an aspect ratio of 5:4, two Logitech speakers, a G27 Logitech steering wheel with force feedback and pedals, and a standard office chair without wheels. The

system was designed by XPI Simulation and included realistic sounds, including braking and accelerating. All participants completed two driving scenarios. The first was a practice scenario designed to familiarise participants with the simulator controls. The task involved participants driving a car from the first-person point of view in a city environment. Halfway through the scenario the car in front of the participants stopped and turned right. At this point participants had to brake and come to a complete halt (if not, then they crashed and the scenario ended), moving away once the car ahead had fully turned. In total this scenario lasted 1.40 minutes. No data was recorded for this scenario.

The second scenario was the main driving task and involved participants driving a car from the first-person point of view along a straight 5.1km road with 30 traffic light-controlled junctions. For twenty of the junctions the traffic light turned from green to amber as the car approached. Participants then had to decide whether to stop the car, by releasing the accelerator pedal and depressing the brake pedal, or to cross the junction by keeping the accelerator depressed. As in real-life, braking would be the safest option but would also result in a time penalty. If a crash did occur the game ended, although the game was designed so that it was impossible to crash if a participant crossed an amber light (but it was possible to crash if participants crossed at red traffic lights). The remaining ten junctions displayed an equal mix of red and green traffic lights. The amber traffic lights were set so that they changed from green to amber at different distances as the car approached the junction - either 10m away from the junction, 20m, 30m, 40m, or 50m, four for each distance. Importantly, as there were no vehicles behind the participants' car (the 'road' was clear in the mirrors) it would be suitable (and possible) to brake at all amber traffic lights if travelling at the speed limit. The game also featured other vehicles, both crossing at junctions but also approaching on the other side of the road. Data was collected from this scenario. Forward, side and rear views (see Figure 5.3) were displayed. Participants were not given any feedback or score for their performance in the scenarios. The fact that the game was designed so that it was impossible to crash if a participant crossed an amber light may have given rise to learning effects, with participants realising there was no penalty for not stopping at amber traffic lights and, as a result, stopping less in the post-training driving game (as seen in Experiment 8). I would argue that this learning effect

reflects the development of similar real-world behaviour and so, if anything, increases the ecological validity of the scenario. Participants' overall average speed and lane position were taken as a measure of the realism of the simulator, with the expectation that participants would drive within the speed limit and maintain a central lane position.



Figure 5.3. Photo of the driving simulator used for Experiment 9. Task displayed is the first trial of the main driving game.

#### **5.6.1.4 Questionnaires**

A wider battery of behavioural characteristics was explored compared to Experiment 8, focused on real-world driving and experience of the simulator.

##### **5.6.1.4.1 Driving experience**

Participants were asked what type of licence they had and if they had been involved in a road traffic collision either as a passenger or a driver. Participants with learner licences were asked if they had any driving experience. This questionnaire was of my own design.

##### **5.6.1.4.2 Driving behaviour**

To assess if participants differ in their real-world driving styles they were asked to complete the extended Manchester Driver Behaviour Questionnaire (Lajunen, Parker, & Summala, 2004). This measures four driving behaviours: ordinary violations, aggressive driving, lapses, and errors.

##### **5.6.1.4.3 Impulsivity and sensation-seeking**

These scores were collected in the same manner as Experiment 8.

#### **5.6.1.4.4 Motion sickness**

Nausea and motion sickness can occur when using simulators and so it was important to ensure that the experience of the simulator was equal for the two conditions. Therefore, participants completed the simulator sickness questionnaire (Kennedy, Lane, Berbaum, & Lilienthal, 1993; Questionnaire taken from Stone III, 2017).

#### **5.6.1.5 Procedure**

The procedure was undertaken as described in Experiment 8 with the addition of the practice scenario before the main phase of the experiment.

### **5.6.2 Analysis and results**

Data was analysed in the manner described in Experiment 8. For Experiment 9 a total of 98 participants were tested. Two participants were unable to complete the experiment due to technical issues. Three participants withdrew due to suffering motion sickness. Finally, one participant met the exclusion criteria of not having stopped at all red traffic lights pre-training. This compares to the 85 participants in Experiment 8 and indicates that the changes to the driving simulator had improved task engagement. The final sample size was 46 participants per condition. Post-hoc analysis using G\*Power 3.1.9.2 (Faul et al., 2007) indicated that the experiment had a 97% chance of detecting an effect size of  $f = 0.20$  for the interaction between time and group at an alpha level of 0.05.

Due to the nature of the driving task where participants were always in control of stopping and starting, the stop data was noisier than that of Experiment 8 and defining a specific event as a stop or go was more complicated. For red traffic light junctions, in addition to crossing on red which represents a clear failure to stop, there was the potential for participants to cross when the traffic light displayed red and amber. In these instances, participants may have come to a halt but then have been too quick in accelerating away from the junction. Given that participants would have stopped, these instances were coded as a stop when calculating frequency of crosses at red traffic lights. Of course, crosses on green traffic lights, whereby the participant would had to have waited for the entire sequence were counted as stops at red lights. For amber traffic light trials the possible events were: a participant crossed while the light was on amber;

misjudged the light and by the time the car was crossing the junction the light had changed to red; braked and waited for the light to cycle round to green; or came to a halt and then crossed before the light displayed green, instead of going on a red light or a red and amber light. Crosses on either a green or red and amber light indicated that the participant had stopped at the lights and waited for the cycle of traffic lights, and as such these two events were coded as stop. Crosses on amber were coded as go. Crosses on red could mean that the participant tried to cross when the light displayed amber but was not quick enough and therefore actually crossed the stop line on a red light, or the participant could have crossed on a red light regardless. Due to the nature of the game if participants did wait at a red light and then crossed, they would have crashed into other cars entering the junction, thus ending the game. Given this, crosses on a red light at amber trials indicated that participants did not stop (see Table 5.4 for tabular version of the coding decisions).

Traffic light	Coding	Reasoning
Amber	Go	As amber will only appear at the start of the sequence any recorded amber cross is a failure to stop.
Red	Go	Recording a red cross on an amber trial means participants will have to 1) misjudge the light and cross the junction on red or 2) brake, wait, but cross the junction while the red light is still on. However, if participants do this then they would crash into the path of oncoming vehicles. In the case of 1) even if participants stop, they will have crossed the stop line on an amber traffic light and so this will count as a cross.
Red and Amber	Stop	To record this event participants must stop at the amber traffic light but accelerate too quickly, crossing the junction on red and amber rather than green. In these cases, while not correct driving, participants did stop at the amber traffic light.
Green	Stop	It is not possible to record a green cross without waiting for the full traffic light cycle because on amber trials the light will always change to amber before participants cross the junction.

Table 5.4. Coding of crosses at traffic lights experienced at amber trials for Experiment 9.

### 5.6.2.1 Participants characteristics

As with Experiment 8 there was some missing data for the questionnaires in Experiment 9. The missing data seemed to be non-systematic and seemingly occurred due to participants missing a question by accident rather than deliberate non-answer. The specific instances and methods taken to deal with the missing data are displayed in Appendix F. For most questionnaires a mean replacement strategy was used, but for the Simulator Sickness questionnaire missing values were replaced with 0, as advised by the scale author's (Kennedy

et al., 1993). Overall, missing data was low and therefore unlikely to affect the conclusions drawn about the sample. Randomisation checks revealed that the two conditions were well matched across variables (see Appendix G).

### **5.6.2.2 Response inhibition task**

#### **5.6.2.2.1 First reaction times to correct GO trials**

Mean responses to GO trials were calculated as in Experiment 8. A mixed measures ANOVA with condition as the between participants factor and the three blocks as the repeated measure found no effect of condition  $F(1, 90) = 0.002$ ,  $p = .957$ ,  $\eta^2 = 0.00003$ , showing that response times were similar between the two conditions (experimental condition:  $M = 717.55$ ,  $SD = 161.78$ , control condition:  $M = 719.46$ ,  $SD = 172.88$ . See Table 5.5 for condition by block descriptive statistics). There was a main effect of block,  $F(2, 180) = 28.69$ ,  $p < .001$  (Huynh-Feldt corrected),  $\eta^2 = 0.064$ , with mean reaction times decreasing over the course of the experiment ( $M = 779.44$   $SD = 175.58$ , for block 1,  $M = 716.03$ ,  $SD = 220.73$  for block 2, and  $M = 660.08$ ,  $SD = 162.60$  for block 3). Thus, participants displayed the expected learning. The interaction term in the ANOVA was not significant,  $F(2, 180) = 0.797$ ,  $p = .452$ ,  $\eta^2 = 0.002$ .

#### **5.6.2.2.2 First response accuracy rates to correct GO trials**

Accuracy rates to GO trials were calculated as in Experiment 8. A mixed measures ANOVA with condition as the between participants factor and the three blocks as the repeated measure found no significant effect of block,  $F(2, 180) = 1.70$ ,  $p = .171$ ,  $\eta^2 = 0.01$ , suggesting that accuracy rates did not change over the course of the training. This is likely to be due to ceiling effects as accuracy rates were high ( $M = 99.26\%$   $SD = 0.82$ , for block 1,  $M = 99.36\%$ ,  $SD = 0.73$  for block 2, and  $M = 99.50\%$ ,  $SD = 0.64$  for block 3; Table 5.5 for descriptive statistics). There was a significant effect of condition,  $F(1, 90) = 5.07$ ,  $p = .027$ ,  $\eta^2 = 0.027$ , with accuracy being higher in the experimental condition ( $M = 99.54\%$ ,  $SD = 0.39$ ) than in the control condition ( $M = 99.20\%$ ,  $SD = 0.60$ ). The interaction between condition and block was not significant  $F(2, 180) = 2.69$ ,  $p = .070$ ,  $\eta^2 = 0.015$ .



Condition	GO trials block 1		GO trials block 2		GO trials block 3	
	RT	A%	RT	A%	RT	A%
Experimental	768.06 (164.40)	99.60 (0.62)	716.07 (192.91)	99.41 (0.78)	668.53 (170.71)	99.63 (0.65)
Control	790.81 (187.22)	98.91 (0.92)	715.94 (247.62)	99.31 (0.78)	651.62 (155.50)	99.39 (0.62)

Table 5.5. Descriptive statistics for GO trials for Experiment 9. SD given in parenthesis. RT = response times. A% = percentage accuracy.

### 5.6.2.2.3 First reaction times to correct CHANGE trials

Response times to CHANGE trials were calculated as in Experiment 8. Participants with no data in a block were removed, with two participants excluded on this basis. I ran a mixed-measures ANOVA on the data with condition as the between factor variable and the three blocks as the repeated measure component. The main effect of condition was significant,  $F(1, 88) = 23.98$ ,  $p < .001$ ,  $\eta^2 = 0.158$ , with response times being slower in the experimental group ( $M = 850.87$ ,  $SD = 196.82$ ) than in the control condition ( $M = 672.11$ ,  $SD = 145.70$ ). See Table 5.6 for condition by block descriptive statistics). There was a significant effect of block,  $F(2, 176) = 5.71$ ,  $p = .006$  (Huynh-Feldt corrected),  $\eta^2 = 0.02$ , with mean response times decreasing but then increasing slightly over the blocks ( $M = 802.55$ ,  $SD = 262.56$  for block 1,  $M = 734.56$ ,  $SD = 193.82$  for block 2, and  $M = 747.36$ ,  $SD = 217.78$  for block 3). Further analysis revealed that the difference between block 1 and 2 was significant,  $t(89) = 3.36$ ,  $p = .001$ , 95% CI [27.83, 108.15],  $d = 0.35$ , with faster reaction times in block 2 than 1. The difference between 1 and 3 was also significant (at a uncorrected alpha level),  $t(89) = 2.19$ ,  $p = .031$ , 95% CI [5.19, 105.19],  $d = 0.23$ , with faster reaction times to block 3 than 1, while the difference between block 2 and 3 was not significant,  $t(89) = -0.72$ ,  $p = .474$ , 95% CI [-48.16, 22.57],  $d = -0.08$ . This indicates that whilst initially participants learnt about the task such learning plateaued. The interaction between block and condition was not significant,  $F(2, 176) = 0.12$ ,  $p = .891$ ,  $\eta^2 = 0.0004$ .

#### 5.6.2.2.4 First response accuracy rates to correct CHANGE trials

Accuracy rates were calculated as in Experiment 8. I ran a mixed-measures ANOVA on this data with condition as the between factor variable and the three blocks as the repeated measure component. There was a non-significant effect of condition,  $F(1, 90) = 0.345$ ,  $p = .557$ ,  $\eta_G^2 = 0.003$ , with accuracy rates being similar in the experimental condition ( $M = 73.87\%$ ,  $SD = 3.52$ ) and in the control condition ( $M = 71.13\%$ ,  $SD = 3.07$ ). There was a significant effect of block,  $F(2, 180) = 4.62$ ,  $p = .012$  (Huynh-Feldt corrected),  $\eta_G^2 = 0.005$ , with accuracy decreasing over time ( $M = 74.47\%$   $SD = 3.26$  for block 1,  $M = 72.67\%$ ,  $SD = 3.61$  for block 2, and  $M = 70.33\%$ ,  $SD = 3.55$  for block 3). The interaction term was not significant,  $F(2, 180) = 1.20$   $p = .303$   $\eta_G^2 = 0.001$  (see Table 5.6 for condition by block descriptive statistics).

Condition	CHANGE trials block 1		CHANGE trials block 2		CHANGE trials block 3	
	RT	A%	RT	A%	RT	A%
Experimental	891.89 (296.51)	74.60 (3.65)	829.10 (200.54)	74.80 (3.70)	831.63 (229.20)	72.20 (3.66)
Control	713.21 (187.25)	74.33 (2.86)	640.03 (132.23)	70.60 (3.54)	663.08 (169.92)	68.53 (3.46)

Table 5.6. Descriptive statistics for CHANGE trials Experiment 9. SD given in parenthesis. RT = response times. A% = percentage accuracy.

The results for the inhibition training phase are less clear than Experiment 8. In terms of GO trials, the significant decrease in response times between blocks shows that participants were learning about the task. However, due to the fact that response times were slower to begin with, compared to Experiment 8, it is hard to state if this learning effect was driven by the additional training block, or if it was always present and that the lack of ceiling effects gives more scope to detect such learning. As seen in Experiment 8 accuracy for GO trials was high. There was a significant effect of condition, yet in both groups' accuracy was still very high for GO trials.

Regarding the CHANGE trials the pattern of response times over the experiment demonstrates that, though learning was occurring, it was variable.

One issue is that there was a significant effect of condition for reaction times, with those in the experimental group having slower reaction times. Clearly there is an effect of the cues (amber vs blue traffic lights). The fact that accuracy rates were similar between the groups suggest that changing to an amber light was much more difficult for the experimental participants. This type of behaviour would be expected if amber lights were indeed priming a going response. However, the two conditions differ in terms of response requirements for the CHANGE trials and it is likely that these differences are behind the effect. In the experimental condition, participants must move their right foot to the left-pedal, while those in the control bar just needed to press the spacebar with their right index finger. Thus, the response required for the control condition in change trials would be, *prima facie*, easier than that in the experimental condition.

Focusing on the accuracy rates for CHANGE trials, the decreasing rates across blocks shows that participants were displaying less inhibition as the training progressed. This could either be due to fatigue, with participants responding with the most likely response (the GO response) in order to complete the task quickly. The effect could also be explained by reference to the task design. Due to the task being designed to assess reactive inhibition, the declining accuracy rates could indicate that, across the experiment, GO trials were eliciting a strong default go response, supported by the reducing reaction times to GO trials, and that participants were unable to inhibit responding. Another interesting feature of the accuracy rates in Experiment 9 is that they are much higher than those in Experiment 8. This could be evidence of a speed-accuracy trade-off (note the slower responses time in Experiment 9 than Experiment 8 for these trials), perhaps arising as an artefact of the changed delay times for the CHANGE trials.

### **5.6.2.3 Driving game analysis**

In this section I present the analysis of the driving game data. First, I present data showing that within the driving simulator both groups maintained safe driving throughout the experiment. These analyses were in effect manipulation checks to ensure that the experience of the simulator did not differ between groups. Next, I present a mixed-measures ANOVA to show the omnibus effects of the training. If the training was effective it would be expected that the interaction would be significant, with more stops post-training than pre- for the

experimental group. Finally, I present a more focused analysis on the pre- and post-training performance.

#### **5.6.2.3.1 Performance within the driving simulator**

The average z-axis (perpendicular to the road) lane position was -47.94 (SD = 0.18) and -47.85 (SD = 0.42) for pre- and post-training scenarios respectively indicating that participants maintained a safe left-hand lane position, with values between -46 and -50 denoting a safe lane position. The average speed for the task was 15.48mph for the pre-training scenario (SD = 0.98) and 16.35mph for the post-training scenario (SD = 1.28), this increase in speed was significant,  $F(1, 90) = 13.21, p < .001, \eta^2 = 0.03$ . Those in the experimental condition recorded a significant difference in speed between the two driving games,  $t(45) = -2.55, p = .014, 95\% \text{ CI } [-0.64, -0.07], d = -0.38$ , increasing from a mean of 15.15mph (SD = 0.91) to 15.95mph (SD = 0.97). The same effect was witnessed in the control condition,  $t(45) = -2.61, p = .012, 95\% \text{ CI } [-0.76, -0.10], d = -0.38$ , increasing from a mean of 15.82mph (SD = 1.03) to 16.78mph (SD = 1.52). There was no effect of condition upon speed,  $F(1, 90), 2.61, p = .109, \eta^2 = 0.023$ .

#### **5.6.2.3.2 Effect of training on stopping at amber traffic lights**

The data was analysed with a mixed 2 x 2 (condition x amber stops) ANOVA, with condition as the between-participant factor and the 'number of stops at amber traffic lights' as the dependent variable. There was a significant effect of time,  $F(1, 90) = 57.00, p < .001, \eta^2 = 0.08$ , with participants stopping at more amber traffic lights post-training (M = 12.54, SD = 2.65) than pre-training (M = 11.13, SD = 2.32). There was a marginally significant main effect of condition,  $F(1, 90) = 3.67, p = .058, \eta^2 = 0.03$ , with those in the experimental condition stopping at significantly more traffic lights than controls (see Table 5.7 for descriptive statistics). The interaction between time and condition was not significant  $F(1, 90) = 1.94, p = 0.166, \eta^2 = 0.003$ .

Further planned analysis showed that both conditions recorded a significant increase in stops at amber traffic lights, for the experimental group,  $t(45) = -6.43, p < .001, 95\% \text{ CI } [-2.20, -1.15], d = -0.95$  and for the control condition,  $t(45) = -4.28, p < .001, 95\% \text{ CI } [-1.69, -0.61], d = -0.63$ . As can be seen in Table 5.7, irrespective of type of training received, stops at amber traffic lights

increased in the second driving game compared to the first. Analysis found a non-significant difference between the experimental and control conditions for pre-training amber stops,  $t(90) = 1.35$ ,  $p = .179$ , 95% CI[-0.31, 1.61],  $d = 0.28$ , but the difference was significant post-training,  $t(90) = 2.17$ ,  $p = .033$ , 95% CI[0.10, 2.25],  $d = 0.45$ , where those in the experimental condition stopped more often compared to controls. This suggests that there was some effect of training, with the experimental group recording a greater increase in stopping than the control condition.

Game	Condition	Mean	SE
Pre-training Game	Experimental	11.46	0.39
	Control	10.80	0.29
Post-training Game	Experimental	13.13	0.38
	Control	11.96	0.38

Table 5.7. Descriptive statistics for stops at amber lights for Experiment 9. Stops out of 20.

### 5.6.3 Summary

Experiment 9 found some rather weak evidence to suggest that inhibition training can increase stopping at amber traffic lights. Of course, the training was not completely successful with the interaction between condition and time not being significant. However, it is possible to argue that the results from Experiment 9 do show promise, particularly given the post-training results.

Compared to the pattern of stopping seen in Experiment 8 the present results are quite different. In Experiment 8 all groups recorded a decline, but in the current study all groups witnessed an increase in stopping at amber traffic lights. The decline in Experiment 8 was hypothesised to occur as a result of participants learning that there was no penalty for crossing amber traffic lights, and so having less reason to stop in the second driving scenario. It seems reasonable to assume that participants did learn about this contingency in Experiment 9; after all, it was still impossible to crash if a participant crossed an amber traffic light. The fact that this learning was not witnessed in task

performance (i.e. as a decrease in stops at amber lights at test) indicates that there were factors in the experiment that had a stronger influence upon behaviour than the in-game learning of amber = go.

The lack of a decline suggests that the both versions of the training were able to counteract the learning in the game. This speaks to the general vs. specific inhibition training debate discussed in the Chapter 1 and the introduction to this chapter. Both groups received response inhibition training (as a component of the overall stop-change response), but for the experimental group this was cue-specific, while for the control condition it was not specific. It may be the case that both forms of training had some effect e.g., priming 'cautious behaviour' (in a similar fashion to the gambling training of Stevens et al., 2015) but that the effects were greater in the cue-specific experimental group. Of course, it must be acknowledged that other accounts, that do not rely on learnt inhibition, could contribute to the witnessed effects. For example, it could be that the training task was priming rule-based knowledge of traffic lights. This argument would explain the significant post-training difference, with participants who received amber CHANGE trials having increased activation of the rules around amber lights than those who received blue CHANGE trials. Additionally, the difference post-training could be explained through demand characteristics, with those who saw amber CHANGE trials assuming that the experiment was focused on behaviour at amber traffic lights and so adjusted their behaviour accordingly. However, it can be questioned how different the amber CHANGE and blue CHANGE trials are. Given the ubiquitous nature of the traffic light sequence, even participants who saw blue CHANGE trials are likely to guess the nature of the task and to have rule-based knowledge primed. Thus, the difference between conditions post-training might reflect the small bonus granted by the cue-specific inhibition training to amber lights specifically, i.e. the retrieval of specific stimulus-response associations at test in addition to general rule priming. Overall, while it is likely that other factors confounded the results, the involvement of amber-specific response inhibition training cannot be entirely dismissed.

In terms of the response times for the CHANGE trials the results do indicate some learning did occur. However, it was variable and the significant difference between conditions suggests task artefacts were influencing performance. The

accuracy rates for CHANGE trials are still low and this could explain the limited behaviour change (see discussion in section 5.7.3). Indeed, the results demonstrate that participants were displaying less inhibition over time. However, it is unclear if this was due to fatigue or arose from GO trials eliciting a strong go response. Of course, it could be a combination of both factors.

One issue with the task is the driving simulator. While the results suggest that participants were engaging with the driving task, supported by the fact that one participant failed to stop at all red lights pre-training, it is clearly not the same as driving, and it might not be the most ecologically valid measure of driving behaviour. It could be the case that while the inhibition task has changed responses to amber traffic lights, such learning is not being expressed in the driving task. Given that the idea is to train associatively-mediated inhibition, for any behaviour change to be witnessed the behaviour measure must allow for the expression of such learning. It could be that due to the relative ease of the simulator participants are deploying propositional learning and thus preventing the expression of the associative learning (see McLaren et al., 2019). Therefore, to address this issue Experiment 10 is a replication of Experiment 9 but with an improved driving simulator that takes the experiment closer to real-world driving.

## **5.7 EXPERIMENT 10**

To increase the validity and to authentically recreate the 'feel' of driving, simulators need to encode multiple senses in a natural manner (Chalmers, Howard, & Moir, 2009). Specifically to increase the experiential validity of a simulator (Pinto, Cavallo, & Ohlmann, 2008) it is important to use large screens (Kemeny & Panerai, 2003) and to have participants in a natural body position (Melo, Rocha, Barbosa, & Bessa, 2016). Therefore, for Experiment 10 the simulator was updated to include larger screens and a car seat allowing the participant to assume a normal driving position in the scenarios.

### **5.7.1 Method**

#### **5.7.1.1 Participants**

As Experiment 10 was effectively a replication of Experiment 9 the same sample size was used. Again, recruitment focused on drivers, but rather than

staff this experiment was aimed at students who had driven in the last two weeks. This decision was made to overcome recruitment issues but also to see if the training was appropriate for younger drivers (who arguably would be of greater need of such interventions given their higher chance of being involved in a traffic collision; International Transport Forum, 2018). Recruitment was still limited to UK drivers only.

### **5.7.1.2 Design**

The design of Experiment 10 was identical to that of Experiment 9.

### **5.7.1.3 Materials**

#### **5.7.1.3.1 Inhibition task**

The inhibition task was the same as in Experiment 9.

#### **5.7.1.3.2 Driving task**

All aspects of the driving task remained the same between Experiment 9 and 10 bar the changes outlined below. The simulator now comprised of three 28-inch LED backlit monitors, each with a 1920x1080 screen resolution, a frame rate of 60hz and an aspect ratio of 16:9. The same two Logitech speakers, G27 Logitech steering wheel with force feedback and pedals where used as in Experiment 9. The simulator now included a RS driving rig that enabled participants to adopt a driving position with the distance between the seat and the pedals being adjustable (see Figure 5.4).



Figure 5.4. Photograph of the driving simulator used in Experiment 10.



#### **5.7.1.4 Questionnaires**

The same questionnaires outlined in Experiment 9 were used in the experiment.

#### **5.7.1.5 Procedure**

The procedure was undertaken as described in Experiment 9.

### **5.7.2 Analysis and results**

Data was analysed in the manner described in Experiment 9. For Experiment 10, a total of 108 participants were tested. Fifteen participants were unable to complete the experiment due to technical issues. Three participants withdrew due to motion sickness. Three participants jumped red traffic lights in the pre-training driving task. Finally, two participants were outliers for mean GO trial response times. The final sample consisted of 43 participants in the control condition and 42 in the experimental condition. While the experiment had aimed for 46 per condition, recruitment difficulties meant it was decided to end the experiment early, and the sample size is still near to that of Experiment 9. Post-hoc analysis using G\*Power 3.1.9.2 (Faul et al., 2007) indicated that the experiment had a 95% chance of detecting an effect size of  $f = 0.20$  for the interaction between time and group at an alpha level of 0.05.

#### **5.7.2.1 Subject characteristics**

Missing data appeared to be non-systematic and seemingly occurred due to participants missing a question by accident rather than being the product of a deliberate non-answer. The specific instances and methods taken to deal with the missing data are displayed in Appendix H. Overall, missing data was low and therefore unlikely to affect the conclusions drawn about the sample. Randomisation checks revealed that two conditions were well matched across a range of variables (see Appendix I).

#### **5.7.2.2 Response inhibition task**

##### **5.7.2.2.1 First reaction times to correct GO trials**

Mean responses to GO trials were calculated as in Experiment 8. A mixed measures ANOVA with condition as the between participants factor and the three blocks as the repeated measure found no effect of condition  $F(1, 83) = 0.34$ ,  $p = .562$ ,  $\eta^2 = 0.003$ , showing that response times were similar between the two conditions (experimental condition:  $M = 657.26$ ,  $SD = 161.14$ , control

condition:  $M = 676.46$ ,  $SD = 142.62$ . See Table 5.8 for condition by block descriptive statistics). There was a main effect of block,  $F(2, 166) = 25.81$ ,  $p < .001$  (Huynh-Feldt corrected),  $\eta^2 = 0.045$ , with mean reaction times decreasing over the course of the experiment ( $M = 713.91$ ,  $SD = 188.90$  for block 1,  $M = 657.93$ ,  $SD = 155.18$  for block 2, and  $M = 629.08$ ,  $SD = 146.60$  for Block 3). Thus, participants displayed the expected learning. The interaction term in the ANOVA was not significant,  $F(2, 166) = 2.30$ ,  $p = .103$ ,  $\eta^2 = 0.004$ .

#### **5.7.2.2.2 First response accuracy rates to correct GO trials**

Accuracy rates to GO trials were calculated as in Experiment 8. A mixed measures ANOVA with condition as the between participants factor and the three blocks as the repeated measure found no significant effect of condition,  $F(1, 83) = 0.61$ ,  $p = .438$ ,  $\eta^2 = 0.003$ , showing that accuracy rates were similar across conditions (experimental condition:  $M = 98.77\%$ ,  $SD = 0.89$ , control condition:  $M = 98.54\%$ ,  $SD = 0.98$ ). There was a significant main effect of block,  $F(1, 83) = 3.11$ ,  $p = .047$ ,  $\eta^2 = 0.020$ , with accuracy varying across time ( $M = 98.36\%$ ,  $SD = 1.49$  for block 1,  $M = 99.04$ ,  $SD = 1.02$  for block 2, and  $M = 98.57\%$ ,  $SD = 1.63$  for block 3). The interaction term was significant,  $F(2, 166) = 3.11$ ,  $p = .047$ ,  $\eta^2 = 0.020$ , with accuracy changing as a function of condition. Those in the control condition displayed an increase in accuracy between block 1 and 2 and then levelled off for block 3, while those in the experimental condition had an increase in accuracy between blocks 1 and 2, but then recorded a decline in accuracy rates for block 3 (see Table 5.8 for condition by block descriptive statistics). As accuracy was still high overall no further analyses were run on this interaction.

Condition	GO trials block 1		GO trials block2		GO trials block3	
	RT	A%	RT	A%	RT	A%
Experimental	693.01 (186.48)	98.77 (1.28)	645.10 (163.33)	99.26 (0.86)	633.67 (163.99)	98.30 (1.78)
Control	733.33 (191.19)	97.97 (1.64)	670.45 (147.64)	98.84 (1.14)	624.60 (129.19)	98.84 (1.45)

Table 5.8. Descriptive statistics for GO trials for Experiment 10. SD given in parenthesis. RT = response times. A% = percentage accuracy.

#### 5.7.2.2.3 First reaction times to correct CHANGE trials

Response times to CHANGE trials were calculated in the same manner as Experiment 8. Participants with no data in a block were removed, which resulted in five participants being removed. A mixed-measures ANOVA was performed on this data with condition as the between factor variable and the three blocks as the repeated measure component. As in Experiment 9, there was a significant effect of condition,  $F(1, 78) = 7.14, p = .009, \eta^2 = 0.06$ , with faster responses in the control condition ( $M = 629.43, SD = 157.73$ ) then in the experimental condition ( $M = 743.81, SD = 221.22$ ), which, as before, is likely due to the different task requirements for the two conditions. The effect of block was marginally significant,  $F(2, 156) = 2.40, p = .094, \eta^2 = 0.009$ , with a trend for faster responses over time ( $M = 707.79, SD = 229.44$  for block 1,  $M = 692.52, SD = 278.24$  for block 2, and  $M = 655.26, SD = 194.51$  for block 3). The interaction was also marginally significant,  $F(2, 156) = 2.59, p = .078, \eta^2 = 0.01$ , with the control condition showing a U-shape pattern of mean response times across the blocks, while the experimental condition had similar levels for blocks 1 and 2 and then a decline for block 3 (see Table 5.9 for condition by block descriptive statistics).

#### 5.7.2.2.4 First response accuracy rates to correct CHANGE trials

Accuracy rates to CHANGE trials were calculated in the same manner as Experiment 8. I ran a mixed-measures ANOVA on these data with condition as

the between factor variable and the three blocks as the repeated measure component. There was a marginally significant effect of condition,  $F(1, 83) = 3.06$ ,  $p = .084$ ,  $\eta^2 = 0.030$ , with a trend for greater accuracy in the control condition ( $M = 72.87\%$ ,  $SD = 2.81$ ) than in the experimental condition ( $M = 64.47\%$ ,  $SD = 3.79$ ). There was a significant effect of block,  $F(2, 166) = 11.72$ ,  $p < .001$  (Huynh-Feldt corrected),  $\eta^2 = 0.02$ , with accuracy increasing over time ( $M = 63.73\%$ ,  $SD = 4.19$  for block 1,  $M = 70.13\%$ ,  $SD = 3.49$  for block 2, and  $M = 72.27\%$ ,  $SD = 3.26$  for block 3). The interaction between block and condition was not significant,  $F(2, 166) = 0.09$ ,  $p = .911$ ,  $\eta^2 = 0.0002$ .

Condition	CHANGE trials block 1		CHANGE trials block 2		CHANGE trials block 3	
	RT	A%	RT	A%	RT	A%
Experimental	773.59 (253.63)	59.07 (4.71)	775.60 (351.28)	66.00 (3.94)	682.23 (193.49)	68.27 (3.63)
Control	645.19 (186.07)	68.40 (3.53)	613.50 (150.12)	74.13 (2.90)	629.61 (193.35)	76.13 (2.77)

Table 5.9. Descriptive statistics for CHANGE trials for Experiment 10. SD given in parenthesis. RT = response times. A% = percentage accuracy.

Overall, the results suggest that the task was learnt by participants, but performance was somewhat variable. The response times for GO trials indicate that both conditions learnt the task equally well. The results for GO accuracy suggest that learning was variable, with accuracy not consistently improving over time as expected, and the significant interaction term reinforces this conclusion. However, it is worth noting that accuracy was still high overall. For the CHANGE trials, the significantly faster response times for the control condition compared to the experimental condition matches that found in Experiment 9 and is likely to be due to task differences. The fact that there was only a marginally significant difference in blocks suggests that learning was not strong. The interaction term also suggests that learning was variable. For accuracy rates the trend for higher accuracy in control conditions compared to

the experimental condition is also likely to be a task artefact, as it is easier to accurately press a spacebar one can see than move a foot between pedals under a desk. Unlike Experiment 9 accuracy significantly improved over time, suggesting inhibition was improving over time. However, accuracy rates were lower than in Experiment 9.

### **5.7.2.3 Driving game analysis**

The analysis presented here was undertaken in the manner described in Experiment 9.

#### **5.7.2.3.1 Performance within the driving simulator**

The average z-axis (perpendicular to the road) lane position was -47.99 (SD = 0.18) and -48.00 (SD = 0.36) for pre- and post-training scenarios respectively; indicating that participants maintained a safe left-hand lane position. The average speed for the task was 15.77mph for the pre-training scenario (SD = 1.13) and 16.51mph for the post-training scenario (SD = 0.95), this increase in speed was significant,  $F(1, 83) = 1.29, p = < .001, \eta_G^2 = 0.03$ . Those in the experimental condition recorded a marginally significant difference in speed between the two driving games,  $t(41) = -1.87, p = .068, 95\% \text{ CI } [-0.61, 0.02] d = -0.29$ , increasing from a mean of 15.79mph (SD = 1.34) to 16.46mph (SD = 1.08). Those in the control condition witnessed a significant increase in speed,  $t(42) = -3.67, p = < .001, 95\% \text{ CI } [-0.58, -0.17], d = -0.56$ , increasing from a mean of 15.73mph (SD = 0.88) to 16.55mph (SD = 0.81). There was no effect of condition upon speed,  $F(1, 83) = 0.0006, p = .981, \eta_G^2 = < .0001$ .

#### **5.7.2.3.2 Effect of training on stopping at amber traffic lights**

The data were analysed with a mixed 2 x 2 (condition x amber stops) ANOVA, with condition as the between-participant factor and 'the number of amber lights stopped at' as the dependent variable. The interaction between time and condition was not significant  $F(1, 83), = 1.52, p = .224, \eta_G^2 = 0.002$ , indicating that the training had a non-significant effect upon behaviour at amber traffic lights (see Table 5.10 for descriptive statistics). The analysis revealed a non-significant main effect of condition on the number of stops at amber traffic lights,  $F(1, 90) = 0.11, p = .738, \eta_G^2 = 0.001$ , suggesting that the groups were similar. A significant effect of time was found,  $F(1, 83) = 31.22, p = < .0001, \eta_G^2 = 0.04$ ,

with participants stopping at more amber traffic lights post-training ( $M = 12.38$ ,  $SD = 3.11$ ) than pre-training ( $M = 11.13$ ,  $SD = 2.88$ ).

Further planned analysis revealed that both conditions recorded a significant increase in stops at amber traffic lights, for the experimental group,  $t(41) = -4.51$ ,  $p < .001$ , 95% CI [-2.21, -0.84],  $d = -0.70$ , and for the control condition,  $t(42) = -3.33$ ,  $p = .002$ , 95% CI [-1.57, -0.39],  $d = -0.51$ . As shown in Table 5.10, irrespective of type of training received, stops at amber traffic lights increased in the second driving game compared to the first. Analysis found a non-significant difference between conditions for pre training amber stops,  $t(83) = -0.11$ ,  $p = .915$ , 95% CI [-1.32, 1.18],  $d = -0.02$  and a non-significant difference at post training as well,  $t(83) = 0.71$ ,  $p = .480$ , 95% CI [-0.87, 1.83],  $d = 0.15$ .

Game	Condition	Mean	SE
Pre-training Game	Experimental	11.10	0.44
	Control	11.16	0.45
Post-training Game	Experimental	12.62	0.40
	Control	12.14	0.54

Table 5.10. Descriptive statistics for stops at amber lights for Experiment 9. Stops out of 20.

### 5.7.3 Summary

Unlike Experiment 9, the results for Experiment 10 do not demonstrate anything that can be termed a significant improvement in stopping at amber traffic lights as a result of the training. However, numerically, there is weak evidence for the direction of effect supporting that in Experiment 9. Given that the training was identical between Experiment 9 and 10 the lack of a stronger effect is puzzling. One reason might be due to the type of participants. Experiment 9 used older more experienced drivers, whereas Experiment 10 used a student sample. Perhaps the training is more effective for those with more driving experience. With more driving experience the link between left-pedal and stop would be assumed to be stronger and this may have helped to increase transference between the inhibition training and the driving game. It could also be that the

existing associative learning of amber=GO is stronger in more experienced drivers, leading to greater effects of training (i.e. as there is a stronger response to change). It would not be possible to tease out any effects from the samples here, but it is a thought for future research (see general discussion for further). In terms of the CHANGE trials performance the results suggests learning was more variable than in Experiment 9 and this could explain the poor training effects, with A. Jones et al. (2016) showing that for training to be effective learning must occur. Additionally, another cause of the lack of strong training effects might be because the CHANGE success was low, ranging between 60-75% in Experiment 10. In the food inhibition literature successful inhibition is often around 90% or more, and A. Jones et al. (2016) has shown that tasks that have higher successful inhibition rates have stronger training effects. Indeed, Jones et al. argued that training needs to be highly accurate and employ consistent stopping for cue-specific training to be effective. It could be the case that the STOP-CHANGE task is too difficult for learning to occur sufficiently to affect behaviour change.

#### **5.7.4 Bayesian meta-analysis**

Given that the three experiments used similar conditions in terms of the experimental and control 1 group it is possible to combine the experiments and undertake a Bayesian meta-analysis to assay the overall evidence for the effect of inhibition training. To enable comparisons between the experiments, I focused on performance between the experimental group and the first control group. I used the process by Zoltan Dienes on his website (see here: [http://www.lifesci.sussex.ac.uk/home/Zoltan\\_Dienes/inference/Bayes.htm](http://www.lifesci.sussex.ac.uk/home/Zoltan_Dienes/inference/Bayes.htm)) to create the prior for the Bayes analysis. This involved working out the mean difference score, and associated standard error of this mean, for the experimental and control 1 condition for the three experiments. Using a meta-analytical approach, I computed a prior based on performance in Experiments 9 and 10. I then undertook a Bayes analysis whereby I compared this posterior to the results from Experiment 8 (I used the calculator by Anupam Singh: <https://medstats.github.io/bayesfactor.html>). This analysis resulted in a Bayes Factor of 2.28. Bayes Factors range from 0 to infinity, with Factors above 3 indicating support for a theory and Factors under 0.33 representing support for the null (Dienes, 2011). Thus, the meta Bayes Factor demonstrates only weak

evidence in favour of the theory in this case, i.e. that inhibition training increases the tendency to stop at amber traffic lights. This approach compares two experiments against the first experiment, another way to calculate the Bayes factor would be to calculate a Bayes Factor between each experiment (so Experiment 8 vs. 9, and Experiment 9 vs. 10) and multiply the two outputs to compute a combined Bayes Factor. Such an approach generated a Bayes Factor of 2.58.

## **5.8 GENERAL DISCUSSION**

Inhibition training designed to train associatively-mediated stopping to amber traffic lights was assayed across three experiments. Experiment 8 found no evidence of an effect, with all conditions being numerically similar. Experiments 9 and 10 showed some promise, with the experimental condition having numerically (though rarely significant) positive improvements in stopping at amber lights following training relative to the control condition. Combining the studies into a meta Bayes analysis indicated that overall, the training only had a weak effect.

### **5.8.1 Inhibition training in driving**

As mentioned in Chapter 1 there is only one other study looking at inhibition within driving, that of Hatfield et al. (2018). So far, discussion of this work has been conspicuous by its absence in this thesis. This was because the work reported here was developed separately from that of Hatfield et al. (2018) and therefore it felt logical to present my findings here before placing them into the context of the (limited) literature. Compared to the work here Hatfield et al found weaker effects of the training, but there are two key differences between their work and mine.

Firstly, the inhibition training used images from the driving simulator to help overcome issues of transference between the training task and the simulator in Experiments 9 and 10. Hatfield et al. (2018) argued against such a design given their aim (and mine) of developing a task to address real-world driving.

However, this approach is setting oneself up to fail as the experiment does not take place in the real-world but rather in a laboratory. Therefore, deliberately not matching the inhibition training as closely as possible to the simulated driving is



akin to showing computerised images in the inhibition task and then assessing real-world driving. Of course, once one moves to assessing real-world driving then it would be appropriate to use real-world images. On a related concept, in Experiments 9 and 10 participants used foot pedals in both the driving and inhibition tasks (though only those in the experimental condition used the foot pedal for CHANGE trials). This was to increase the likelihood of transfer between the two tasks. However, if one wishes to design a training task to be used in the real-world then it must be questioned how likely is it that people have access to foot pedals. Using keyboard (or even touchscreen) responses would more readily fit into people's daily interactions with technology and might not be such a transference barrier as first thought. The use of such motor responses in online training has led to real-world reduced food intake and weight loss, indicating some transfer of a button-press or touchscreen response to the rather different act of picking up and buying or eating a food item (N. S. Lawrence, O'Sullivan, et al., 2015; Beurden, Smith, Lawrence, Abraham, & Greaves Van Beurden, 2019). However, specific motor responses are more important in driving, and it would be interesting to compare the effectiveness of the training developed here against similar training using keyboard presses.

Secondly, as in Experiment 8 and 10, Hartfield et al. (2018) used a student sample. However, the results from Experiment 9 hint at stronger effects for participants that represent the wider driving population. While inhibition work in food has found success in students (N. S. Lawrence, O'Sullivan, et al., 2015) it can be argued that these populations already have a strong approach tendency to cake (that is, most people like to eat cake). Young drivers might not have such a strong amber – go link. Therefore, testing more experienced drivers (who one assumes have a strong go prime to amber) would bring the driving domain of inhibition training into line with the wider literature and might lead to strong effects, though of course ultimately any intervention would ideally improve driving of young people. Such an argument has support from work which has found stronger effects of go/no-go training in those with stronger impulses to go to food (Houben & Jansen, 2011; Veling et al., 2013a; yet see Z. Chen, Veling, Dijksterhuis, & Holland, 2018).

### **5.8.2 Inhibition training across domains**

The experiments presented here indicate that response inhibition techniques might not fulfil the promise of easy and effective behaviour change. The fact that the Bayes Factor provided only weak evidence in favour of the alternative hypothesis despite including three large studies does sit at odds with the work into food and alcohol which suggest that, to varying degrees, the training is effective. However, the results do tie into recent findings by Bos et al. (2019). This study was a randomised control trial investigating the use inhibition training to reduce smoking. Those in the experimental condition were trained to no-go to smoking images, whilst those in the control condition saw control images. Similar to the findings reported here, both groups had a reduction in smoking and the differences between groups at each time point (post intervention, 1 month, 3 months follow ups) were not significant. This suggests that there is some difference between tasks within the inhibition literature, in that either the mechanism underlying the effects, or the measurements used to assess behaviour change, are more valid/sensitive in the consumption research.

### **5.8.3 Future directions**

While the measure of driving behaviour, the simulator, was refined across the experiments it is still far from capturing the experience of real-world driving. It might be the case that the training did bring about changes in driving behaviour, but that the simulator did not give participants the correct environment to deploy the learning. That is, the training might be effective in the real-world but not in the simulator. As already discussed, driving is cognitively demanding and so is a rich environment for associatively-mediated processes to guide behaviour. Whereas in the laboratory the effort required to 'drive' is likely to be less than that for real-world driving (in a sense this is a similar argument to that of McLaren et al., 2019 on the need for the appropriate procedure to detect associative learning). Informally, one could compare driving in the simulator and driving in the real-world as being equivalent to, on the one hand, driving along an empty motorway and on the other, driving on your own in the middle of London. It's easy to see that in each condition the resources required are likely to be markedly different. Thus, performance in the simulator might be influenced by propositional rather than associatively-mediated processes. Furthermore, in the real-world there are genuine reasons (at least in the eyes of the driver) for

committing traffic violations and strong motivations to drive within the law (e.g., to avoid a fine). This behaviour algebra was missing from the experiments presented here undermining the realism of the simulator environment. Therefore, an obvious next step would be to explore the effects of the training in the real-world.

Although participants seemed to learn about the inhibition task over time there are at least two ways the task could be improved. Firstly, the task could include multiple images of traffic lights to improve generalisation. Many inhibition experiments use multiple training images, e.g., Camp and Lawrence (2019) used eight different images of meat as no/go cues in their training task. Therefore, would seem sensible to redesign the current inhibition task to display multiple different images of traffic lights. Of course, given that traffic lights can be seen from many angles, in all lighting/weather conditions and in a myriad of physical contexts, this does suggest another avenue of enquiry surrounding how close to the real-world context would these images need to be. The second change would be to make the task adaptive. While many experiments in this field use fixed timing, including a tracking procedure such as that used in Experiments 1-7 would help tailor the inhibition training to each participant's subjective level of task difficulty. An example of such a tracking procedure being applied to inhibition training comes from Johnstone et al. (2012), who used a tracking procedure in a go/no-go task with children with attention deficit/hyperactivity disorder. Of course, as noted earlier, success on inhibition trials predicts effect size of training and therefore it would be important to ensure that the tracking procedure was one that did not make the task harder.

The task was developed in order to tackle associatively-mediated learning at UK traffic lights. It is also likely that there are instances where propositional learning guides behaviour (e.g., if there is a police car besides you at the lights). Therefore, it would be useful to combine such training as presented here with more traditional educational interventions to tackle both routes. Such an approach was taken by Veling et al. (2014) who in their research into food consumption combined go/no-go training and implementation intention (if-then planning) training; finding no added benefit of combining the interventions. Still, implementation intentions have been shown to have some success in changing real-world driving behaviour (Elliott & Armitage, 2006), and designing a program

that addresses both propositional and associatively-mediated learning could prove fruitful.

#### **5.8.4 Conclusions**

From the perspective of developing an intervention to address real-world driving the experiments here do not indicate strong enough evidence to be of significance to road safety bodies. However, the discussion has highlighted several lines of enquiry that need to be explored before the use of inhibition in a driving context can be dismissed entirely. Therefore, it seems worthwhile for future research to explore ways to increase the effectiveness of the training – until driverless cars become common place that is.

The experiments in this thesis have implications for the study of associative learning, response inhibition, and most crucially, the development of driving safety programmes. In this final chapter I briefly summarise the preceding work and then draw out conclusions relating to the theories discussed in Chapter 1 and make suggestions for directions of future research.

## 6.1 THE EXPERIMENTAL WORK

### 6.1.1 Chapter 2: Associative learning at traffic lights

In the first trio of experiments I sought to begin to understand the role of associative learning at UK traffic lights. Human factors have been found to play a significant role in road traffic incidents (U.S. Department of Transportation, 2015), with non-compliance at traffic light signals making up a large segment of these incidents (Retting et al., 1995). Research has shown that associative learning happens in the background and can influence behaviour when control processes are weak (McLaren et al., 2019), and work in the driving literature has shown how cognitively taxing driving is (Walshe et al., 2017). This increases the likelihood that the associative learning that takes place in response to traffic lights will manifest itself upon driving behaviour. Therefore, it was important to explore the behaviour that associative learning might prime. The experiments in Chapter 2 also introduced the notion of the feature-positive effect and applied this to investigate how learning at lights was affected by task set. At a fundamental level, the act of driving can be seen as a dichotomy between going and stopping, and so it was important for the experiments to capture this aspect of the activity. A between experiments analysis confirmed that the change in task set between Experiments 2 and 3 had been successful, with the difference in the R vs. RA and I vs. IP discriminations shifting in line with the feature-positive effect. The effect on task set was also to modulate the learning of cue A. In circumstances where participants were using a stop task set, cue A seemed to be a go cue, yet when a go task set was in force, cue A was more of a stop cue. However, setting aside Experiment 1 (due to its poor

design), a between experiment analysis, that was a first step in capturing the effects of both task sets, showed that overall cue A was a go cue. This runs counter to the rules of driving in the UK which dictate that amber is a neutral to stop cue.

### **6.1.2 Chapter 3: The role of sequences**

The experiments in Chapter 2 were designed in the style typical of cognitive psychology experiments in associative learning, rather than within a framework focussed on capturing the essence of the experience of UK traffic lights.

Chapter 2 marked a shift towards to a greater focus on drivers' experience at traffic lights, specifically with regards to the role of sequences. The importance of sequences in learning is well documented (Ashe et al., 2006), and indeed, they are crucial to the traffic light system in the UK. Having found evidence to suggest that the contingency learning for amber lights might be inappropriate in a relatively simple design, the next step was to explore learning in a design that more closely matched that experienced in the real-world. The two experiments in Chapter 2 still retained the necessary cues to enable analysis of the feature-positive effect, again finding good evidence to support the notion that the between experiment manipulations was having the desired shift in effective task outcome. Regarding the addition of sequences upon learning, while learning itself was noticeably stronger, the overall pattern of results was similar to that of Chapter 2. These results confirmed that associative learning at traffic lights could well be such as to encourage dangerous driving. The experiments also gave credence to the claim that the paradigm was addressing associative rather than propositional learning, with no evidence of participants' ratings correlating significantly with learning about the cues. Though of course, the issues raised by Shanks and St. John (1994) and others apply to such claims.

### **6.1.3 Chapter 4: The effect of task set**

Until Chapter 4 the learning for each task set, that is if participants (and thus drivers) were looking for go or stop cues, was addressed at a between participants' level. However, drivers are likely to shift task set depending on the signal the traffic light is displaying. Therefore, Experiment 6 implemented a design that embedded task set switching in a within participant's design, enabling the experience of real-life traffic lights to be factored into the paradigm.

Overall, the results for Experiment 6 (and the training phase of Experiment 7) suggested that cue A was again a weak go cue and that, reassuringly (both in terms of road safety and the validation of the paradigm) cue R, which represented the contingences of a red traffic light, was a strong stop cue.

With these results in mind, Experiment 7 introduced a form of inhibition training aimed to shift learnt responses to cue A from go to stop. However, the training did not prove successful, with the experimental group not showing the hypothesised shift compared to a control group.

#### **6.1.4 Chapter 5: Effecting behaviour change**

Chapter 5 marked another transition in the focus of the thesis. Rather than exploring learning at traffic lights, this chapter focused on continuing the work of Experiment 7 and investigated whether any such pre-existing associative links could be changed. Given that my experiments undertaken so far indicated that amber lights primed a go response, and that driving seems to be a rich environment for such associative learning to take place, it seemed only natural to explore whether an intervention could be designed to amend the associative learning at amber traffic lights.

While this chapter continued to draw upon associative theories, it adapted the experimental design to take into account the findings from applied work into reducing alcohol and food consumption and to better fit the real-life circumstances of driving. Across three experiments the inhibition task was developed and refined, along with the driving simulator used to assay evidence of behaviour change. Bayesian analyses found that across the three experiments there was weak evidence for the effectiveness of the training. As a result, future developments of the paradigm were suggested that might elicit strong behavioural shifts.

#### **6.1.5 Summary**

To summarise the empirical work undertaken within this thesis; across a series of experiments I developed a paradigm that sought to capture an increasingly complete picture of the factors influencing contingency learning at UK traffic lights. These results suggest that associative processes are likely to prime a go response to amber traffic lights. With this knowledge, my attention then focused on investigating if a response inhibition training task could be developed to

adjust this associative learning to be less go and more stop. In this I was only partially successful, with further development clearly required. Having summarised the experiments in this thesis, this chapter now concerns itself with the broader theoretical and applied implications of the work.

## **6.2 ASSOCIATIVELY-MEDIATED LEARNING**

One clear outcome from the experiments in Chapters 2 and 4 is that evidence of associatively-mediated learning was observed to some extent. The go/no-go training resulted in slower reaction times and less probability of committing commission errors for stimuli paired with a stop response compared to those stimuli linked to a go response. Of course, what the experiments cannot do, nor were they designed to, is speak to the mechanisms that underpin this learning. However, they do speak to the ongoing debate in Psychology on the nature of human learning.

### **6.2.1 Dual process theories of human behaviour**

The findings from the test phases sit in direct contrast to the single propositional account of human learning (e.g., Mitchell et al., 2009), and the thesis can be seen as further evidence for a dual-process account of human learning. The 50/50 test blocks allowed for the decoupling of associatively-mediated and propositional learning that could have developed in the training blocks. In training it is entirely possible that subjects had developed propositional rules regarding cues and responses in addition to associatively-mediated learning. Yet, at test, given that the blocks used 50/50 contingencies, performance based on previous learning would be uncertain, and so it is unlikely that participants were able to deploy any propositional learning. On this basis, the test blocks allow for the assessment of raw associative learning (similar logic was posited by McLaren et al., 2019). Therefore, the findings at test where cue G was often significantly faster than cue B suggests that associatively-mediated learning was affecting participant responses. Of course, results at test were often weaker than at training, with this likely arising as a fact that the test blocks were essentially extinction phases. Over time both associative learning and any propositional learning would lead to participants responding to cues uniformly given the 50/50 nature of the blocks.



I mention that propositional learning could develop through the paradigm and influence performance, and while this is true little evidence was found to suggest such learning did develop. Looking at the results for the awareness tests conducted in Experiments 4, 5 and 6 it is clear that while participants were able to learn that some cues denoted a go response more often than other cues, the correlational analysis demonstrated little evidence for participants using this knowledge to influence responding. Of course, such tests have many issues (see discussion in Chapter 3 7.3).

One final comment on the nature of learning within the experiments. Though I would argue that the results support an associative account of human learning, whether this occurs in the traditional sense of the phrase as used by McLaren et al. (2014), or through the two processes, one system of McLaren et al. (2019) cannot be fully ascertained.

## **6.2.2 Applying associatively-mediated learning paradigms**

One implication of the experiments in Chapters 2, 3 and 4 is that they demonstrate how it is possible to use existing learning paradigms and apply them to begin to understand associatively-mediated learning in a real-life setting. Though it is clear that the paradigm developed in this thesis does not capture all facets of the experience of traffic lights, the thesis shows how it is possible to use a relatively simple experimental design to start to highlight the role of associatively-mediated learning in specific human behaviours. The design could be applied to explore other traffic behaviours. For example, it is likely that there is a large associatively-mediated learning component in the response to the National Speed Limit sign in the UK. This sign indicates to drivers they are leaving an environment where speed restrictions were in place and can now drive up to 60 m.p.h. Given the strong GO response often made to this sign (placing one's foot on the accelerator) it is likely that over time the sign could come to automatically cue a strong go response, even in circumstances when such a response would be inappropriate (such as traffic up ahead). Further insights into the strength or otherwise of associatively-mediated learning for a particular behaviour is likely to aid the development of more effective interventions through addressing those facets of learning that lead to the behaviour.

### **6.3 IMPLICATIONS FOR ROAD SAFETY**

The evidence presented in this thesis indicates that associative learning at amber traffic lights is contrary to that expected by the Highway Code. This has several implications for road safety practitioners. Firstly, it suggests that a focus on using the Theory of Planned Behaviour (Ajzen, 1992; see section 1.5) as a model to bring about behaviour change in the driving domain is not likely to be the most effective choice, given its focus on conscious thought processes to the exclusion of associatively-mediated learning. This conclusion is supported by work discussed earlier, such as Steinmetz et al. (2016), who found only small effect sizes for interventions using the Theory of Planned Behaviour to change traffic behaviour. Therefore, it is recommended that work is undertaken to develop new models that can address both associative and conscious routes to change driving behaviour.

Secondly, the work highlights how associative learning can prime dangerous responses and thus the need for the road safety literature to consider the role of associative learning in other driving contexts. Indeed, this thesis demonstrates the need for interventions to be rooted to scientific work and the importance of using laboratory research to support and develop more effective interventions. More work aimed at uncovering the impact of associative learning on driving behaviours is needed to develop more holistic interventions. One aspect worth noting is that learning is specific to the contingencies in play within a particular system. For example, in New South Wales, Australia, the traffic light sequence is green – amber – red – green and therefore any associatively-mediated learning arising from experience of this sequence is likely to be quite different to that at UK traffic lights. Furthermore, the paradigm designed in Chapter 2 through to Chapter 5 focused on traffic lights, and the learning to Pelican crossings (in which a flashing amber light replaces the Red and Amber signal) could be quite different. Arguably the flashing amber light is a strong stop signal and so the solo amber light at these crossings might be associated with stop to a greater extent than a solo amber at traffic lights. The point here is that it is not a simple case of exploring associative learning in one context and then assuming such learning holds true for conceptually-related situations.

Lastly, as well as illustrating the likely role of associative learning in driving behaviour, this thesis provides a springboard to illustrate how interventions

could be developed to address dangerous driving. The thesis supports calls by Lheureux et al. (2016) and Fylan (2017) to include features in interventions that address associative learning in a driving environment. Associatively-mediated learning principles are often used in real-life for training, for example positive reinforcement techniques are a key part of the training programme for Guide Dogs in the UK (Dopson, 2018), but are largely ignored in driving interventions (Gardner, 2015b).

One aspect of the real-world use of the response inhibition intervention developed here (and similar designs) is that it only targets the specific go response to amber traffic lights. While training amber to become a stop cue is an obvious step in reducing road traffic incidents at traffic lights, if a driver is going too fast, even if they do brake in response to an amber light, they might not be able to stop and so cross through the junction illegally. The point here is that a whole system approach is needed to successfully change behaviour. Successful behaviour change is likely to involve targeting a range of behaviours that form part of a 'behaviour link' that precedes the specific action of interest in a particular study. Therefore, in the driving domain this might entail using a combination of several theories and research domains to target all links in the chain. For traffic lights, in addition to the development of the intervention presented here, work could focus on speed reduction media campaigns or exploring the effects of road calming measures, such as narrowing of roads, or even how different colour tarmac might help reduce drivers speed and increase their preparedness to stop. Such a focus on both the individual, wider societal factors, and the environment, is likely to bring about the greatest behaviour change.

## **6.4 INHIBITION TRAINING**

As well as implications for the road safety domain this thesis also speaks to the response inhibition training literature. The results from Chapter 5 counter the notion in the popular press of the success and effectiveness of 'brain training' apps. The finding that the behaviour change resulting from the training was small suggests that more robust training paradigms are needed to bring about real-world change. The results from Chapter 5 lead to several practical

suggestions for those wishing to use inhibition training to deliver behaviour change.

Firstly, the research in Chapter 5 indicated how the STOP-CHANGE task could be used within an intervention. This task is rarely used in the applied literature. However, it could be more appropriate than the oft used go/no-go training where the focus is on changing, rather than stopping, behaviour. For example, with regards to drug taking, to encourage individuals to take methadone rather than heroin. It could be the case that use of the STOP-CHANGE task rather than the go/no-go task would reduce the transfer gap between the training task and the real-world behavioural responses, and so increase the effectiveness of the intervention. One issue with the use of the STOP-CHANGE task in the current intervention was the low successful inhibition rates, with these likely being a key reason behind the limited success of the intervention developed in this thesis. Using a staircase design to tailor the intervention to individuals could help increase inhibition rates by reducing the difficulty of the task.

Secondly, the experiments in Chapter 5 demonstrated the need to consider the length of training. While it is all too easy to focus on the benefits of shortening training tasks in order to increase their acceptability, this needs to be balanced against the impact on their effectiveness. The results in Chapter 5 demonstrated that learning within the training task was variable, and it is suggested that increasing training length would be beneficial, despite the costs. Of course, it might not be simply the duration of the training that increases its effectiveness (by which I mean participants demonstrating learning within the task). The number and length of breaks, or the number of days over which the training is conducted, are all likely to be factors that can influence the success of the training (see Bakkour et al., 2018 for further).

Thirdly, the importance of using ecologically valid measurements of behaviour change cannot be understated. The changes to the driving simulator introduced for Experiment 9 correlated with improved effectiveness of the intervention. Given the ultimate desire of all inhibition tasks to change behaviour in the real-world it is crucial that the experiments reflect their real-life counterparts as closely as possible. A related issue is the idea discussed in Chapter 5, that the simulator did not give participants the appropriate environment to display the behaviour change. As McLaren et al. (2019) note, procedure is key. Given the

focus of inhibition training to change associatively-mediated learning it is crucial that the methods used to assess behaviour change enable the behaviours arising from associative learning to be expressed.

Lastly, the work in Chapter 5 highlights how elusive the effects of inhibition can be. The fact that across the three experiments varying degrees of (limited) successful behaviour change was found demonstrates that a range of factors beyond the specific training design implemented play a role. In the case of driving, age, experience, past driving history are likely to be key determinants on behaviour. Therefore, it is not merely enough to design the training task, but rather researchers need to consider the characteristics of participants as well. Such a focus is especially important for meta-analysis or review papers where the effects of inhibition training across studies could be explained by differing sample pools. The limited effects of the inhibition task developed in this thesis also speak to a wider issue in the field, that of weak effects. Throughout this thesis one narrative arc has been the uncertainty surrounding the effectiveness of inhibition training, with training seemingly effective for one domain but not another. For example, while research (N. S. Lawrence, O'Sullivan, et al., 2015) has supported the real-world effectiveness of response inhibition upon food consumption, this has not been consistently found for alcohol (A. Jones et al., 2018; yet see Strickland et al., 2019). The work in Chapter 5 can be seen as a microcosm of this issue, with the meta Bayes Factor suggesting the training had some effect, but that this was weak. This conclusion is certainly consistent with the work discussed in Chapter 1. It suggests that as a whole inhibition training is entering a new phase where the focus should be on marginal gains to improve and understand mechanisms behind current training tasks.

## **6.5 CONCLUSIONS**

This thesis has clearly demonstrated the need to consider the role of associative learning within a driving context. Importantly, the thesis takes associative theories and applies them to a novel domain. The thesis also makes an important contribution to developing a paradigm that could help address the maladaptive behaviour that associative learning at traffic lights primes, but clearly these are first steps and much more by way of development is needed. The exact nature of the role of associative learning at UK traffic lights, and how

exactly this manifest itself in the expression of driving behaviour is still unclear; and this is a challenge for future research.

## Appendix A: Non-Significant Reaction Time Contrasts for Test for Experiment 3

In terms of the experimental contrasts, A vs. B was non-significant,  $t(432) = 0.85$ ,  $p = .396$ , 95% CI [-7.99, 20.23],  $d = 0.08$ . The A vs. G contrast was not significant,  $t(432) = 0.08$ ,  $p = .940$ , 95% CI [-13.56, 14.65],  $d = 0.01$ . The A vs. R contrast was also non-significant,  $t(432) = -0.24$ ,  $p = .814$ , 95% CI [-15.80, 12.41],  $d = -0.02$ . For the contrasts against R, R vs. B was not significant,  $t(432) = 1.09$ ,  $p = .278$ , 95% CI [-6.29, 21.93],  $d = 0.10$ . The R vs. G was also non-significant,  $t(432) = 0.31$ ,  $p = .756$ , 95% CI [-11.87, 16.35],  $d = 0.03$ . The difference between R- vs. RA+ and IP- vs. I+ was non-significant,  $t(54) = -1.44$ ,  $p = .156$ , 95% CI [-32.06, 5.26],  $d = -0.19$ .

## **Appendix B: Non-Significant Commission Error**

### **Contrasts for Test for Experiment 3**

The A vs. B was not significant,  $z = 1.23$ ,  $p = .217$ . The A vs. G contrast was not significant,  $z = 1.23$ ,  $p = .217$ . The A vs. R contrast was not significant,  $z = 0.58$ ,  $p = .564$ . Focusing on the R cues, the R vs. B contrast was not significant,  $z = 0.67$ ,  $p = .500$ , nor was the R vs. G contrast,  $z = 0.67$ ,  $p = .500$ .



## Appendix C: R-RA Contrast on Data from Experiments 2 and 3

In Chapter 3 I performed a contrast looking at the difference between R and RA to assess for evidence of a feature-positive effect. This was because in Experiment 5 R vs. RA was the feature-positive contrast (R- RA+), while in Experiment 4 it was the feature-negative pair (R+, RA-). The same logic holds for Experiments 2 and 3 in Chapter 2, where in Experiment 2 the R RA contrast was the feature-negative pair (R+, RA-), while in Experiment 3 the R vs. RA was the feature-positive contrast (R- RA+). Therefore, it is possible to undertake the same analysis presented in Chapter 3 upon the data in Chapter 2. A standard alpha level was applied to this set of analysis.

The result from response times at training were marginally significant at the standard alpha level,  $t(108) = -1.91$ ,  $p = .006$ , 95% CI [-15.97, 0.29],  $d = -0.37$ , with enhanced learning to RA compared to R in Experiment 3 (mean difference of 6.70, SD = 24.37) compared to Experiment 2 (M = -1.14, SD = 18.20). This demonstrates that the changes made between the Experiments 2 and 3 successfully changed the nature of the discriminations experienced by participants, that is the effective outcome changed from stopping to going. However, the results from test were not significant,  $t(108) = 0.14$ ,  $p = .886$ , 95% CI [-18.70, 21.62],  $d = 0.03$ .

As it was not possible to conduct the R RA contrast for training commission errors in Chapter 3, I did not undertake it for Chapter 2 data, but I did perform the analysis for test data. This contrast was not significant,  $t(108) = 1.34$ ,  $p = .183$ , 95% CI [-0.006, 0.033],  $d = 0.26$ .

## Appendix D: Tables of Missing Data Decisions for Experiment 8

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
Gender		No missing data	
Age	2; 1 control 2, 1 control 3	5% for control 2 and 4% for control 3	Took mean of sample for both conditions with missing data dropped
Crash history	2; 1 control 2, 1 control 3	5% for control 2 and 4% for control 3	Missing data dropped from analysis
Type of driving licence	3; 1 control 2, 2 for control 3	5% for control 2 and 8% for control 3	Data included as a third category in the analysis
Length of time full driving licence held	5; 1 control 1, 2 control 2	5.26% for control 1, and 6.67% for control 2.	Missing data dropped from sample.

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
Impulsivity	4; 1 experimental, 2 control 2, 1 control 3	3.70% for the experimental group, 10% for the control 2, and 4% for control 3	Two participants did not complete any questions in the scale and so were removed from the analysis. The other two had missing data replaced with their mean.
Sensation-seeking	2; 1 control 2, 1 control 3	5% for control 2 and 4% for control 3	These participants did not complete any questions in the scale and so were removed from the analysis.

## Appendix E: Group Characteristics Comparisons for Experiment 8

Means are presented with standard deviations in parentheses. *p* values in bold represented fisher exact test values. Those with missing values for a particular characteristic are not included in the descriptive or inferential statistics.

Variable	Experimental condition	Control 1	Control 2	Control 3	Statistical test
Length of time full driving licence held					
< 3 months	2	0	0	0	$\chi^2(9) = 12.78, p = .092$
< 6 months	1	0	0	1	
< 1 year	0	4	3	1	
> 1 year	20	14	12	17	
Impulsivity score	17.96 (5.45)	18.63 (5.73)	17.53 (4.50)	18.04 (3.64)	$F(3, 90) = 0.18, p = .908$
Sensation-seeking score	35.22 (8.48)	34.50 (8.36)	31.84 (11.12)	32.17 (7.93)	$F(3, 90) = 0.83, p = .479$

Variable	Experimental condition	Control 1	Control 2	Control 3	Statistical test
Gender					
Female	16	18	13	21	$F(3, 92) = 1.47, p = .229$
Male	11	6	7	4	
Age	21.14 (3.21)	20.67 (2.39)	22.90 (7.44)	21.42 (3.89)	$F(3, 90) = 1.17, p = .326$
Crash history					
Yes	12	16	5	7	$F(3, 90) = 3.40, p = .021$
No	15	8	14	17	
Type of driving licence					
Full	23	19	16	19	$\chi^2(6) = 4.21, p = .772$
Provisional	4	5	3	4	
Not specified	0	0	1	2	

## Appendix F: Tables of Missing Data Decisions for Experiment 9

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
Gender		No missing data	
Age		No missing data	
Type of driving licence		No missing data	
Length of time full driving licence held		No missing data	
Impulsivity score		No missing data	
Simulator sickness			
Oculomotor	1; control	2.17%	0 replacement
Disorientation	11; 5 experimental, 6 control.	10.87% for experimental, 13.04% for control	0 replacement
Nausea	4; 1 experimental, 3 control experimental	2.17% for experimental, 6.52% for control	0 replacement

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
Driver Behaviour			
Aggressive violations	5; 2 experimental 3, control	4.34% for experimental, 6.52% for control	Two in the control condition dropped. Others mean replaced.
“Ordinary” violations	6; 2 experimental, 4 control	4.34% for experimental, 8.70% for control	Three in the control condition dropped, Others mean replaced.
Errors	7; 3 experimental, 4 control	6.52% for experimental, 8.70% for control	Three in the control condition dropped, Others mean replaced.
Lapses	9; 5 experimental, 4 control	10.90% for experimental, 8.70% for control	Three in the control condition dropped, Others mean replaced.

## Appendix G: Group Characteristics Comparisons for Experiment 9

Means are presented with standard deviations in parentheses. *p* values in bold represented fisher exact test values. Those with missing values for a particular characteristic are not included in the descriptive or inferential statistics.

Variable	Experimental condition	Control condition	Statistical test
Gender			
Female	28	31	$t(90) = -0.65, p = .520$
Male	18	15	
Type of driving licence			
Full	43	41	$t(90) = 0.73, p = .465$
Provisional	3	5	
If a full licence how long have you been driving for?			
Less than one year	1	2	
1-2 years	5	3	
3-5 years	6	9	$\chi^2(5) = 1.75, p = .890$
6-8 years	6	6	
9-10 years	1	1	
10 or more years	24	20	



Variable	Experimental condition	Control condition	Statistical test
Age	32.24 (10.58)	31.74 (10.77)	$t(90) = -0.22, p = .823$
Manchester Driver Behaviour Questionnaire			
Aggressive violations	4.63 (1.60)	4.55 (1.55)	$t(88) = -0.26, p = .798$
“Ordinary” violations	14.93 (4.82)	15.47 (4.23)	$t(87) = 0.55, p = .584$
Errors	11.67 (2.98)	11.88 (2.59)	$t(87) = 0.35, p = .725$
Lapses	16.00 (4.04)	16.93 (4.80)	$t(87) = 0.99, p = .325$
Impulsivity score	17.98 (5.54)	18.28 (6.12)	$t(90) = 0.25, p = .803$
Sensation-seeking score	33.26 (9.41)	33.89 (7.66)	$t(90) = 0.35, p = .725$
Simulator Sickness			
Oculomotor	26.37 (19.09)	33.29 (24.83)	$t(90) = 1.50, p = .138$
Disorientation	26.33 (29.23)	33.29 (35.03)	$t(90) = 1.03, p = .304$
Nausea	23.85 (22.37)	25.92 (27.03)	$t(90) = 0.41, p = .690$
Total Severity	29.43 (22.14)	35.53 (29.91)	$t(90) = 1.11, p = .269$

## Appendix H: Tables of Missing Data Decisions for Experiment 10

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
Gender		No missing data	
Age		No missing data	
Type of driving licence		No missing data	
Length of time full driving licence held	1 in the Experimental condition	2.38%	Missing data dropped from analysis
Impulsivity score	4; 1 experimental, 3 control	2.38% in the experimental group, 6.98% in the control group	Mean replacement
Sensation-seeking scale	6; all control	13.95%	Mean replacement

Scale/ question	Total number of missing cells	% missing cells of group mean	Replacement strategy
<b>Driver Behaviour</b>			
Aggressive violations		No missing data	
“Ordinary” violations	5; 2 experimental, 3 control	4.76% for experimental, 6.98% for control	Mean replacement
Errors	3; all control	6.98%	Mean replacement
Lapses	3; 1 experimental, 2 control	2.38% in the experimental group, 4.65% for control	Mean replacement
<b>Simulator sickness</b>			
Oculomotor	4: 1 experimental, 3 control	2.38% in the experimental group, 6.98% for control	0 replacement
Disorientation	6; 3 in each group	7.14% for experimental, 6.98% for control	0 replacement
Nausea	2; both experimental	4.76%	0 replacement

## Appendix I: Group Characteristics Comparisons for Experiment 10

Means are presented with standard deviations in parentheses. *p* values in bold represented fisher exact test values. Those with missing values for a particular characteristic are not included in the descriptive or inferential statistics.

Variable	Experimental condition	Control condition	Statistical test
Age	22.38 (7.81)	21.79 (6.54)	$t(83) = -0.38, p = .706$
Manchester Driver Behaviour Questionnaire			
Aggressive violations	4.93 (1.94)	4.58 (2.06)	$t(83) = -0.80, p = .427$
“Ordinary” violations	16.24 (5.42)	16.12 (5.33)	$t(83) = -0.10, p = .917$
Errors	12.00 (2.30)	12.58 (3.01)	$t(83) = 1.00, p = .320$
Lapses	17.67 (4.03)	18.56 (4.12)	$t(83) = 1.01, p = .317$
Impulsivity score	19.90 (5.75)	19.23 (5.51)	$t(83) = -0.55, p = .583$
Sensation-seeking score	34.64 (9.29)	34.60 (8.82)	$t(83) = -0.02, p = .985$
Simulator Sickness			
Oculomotor	29.24 (23.77)	27.32 (19.65)	$t(83) = -0.41, p = .874$
Disorientation	29.17 (30.25)	27.19 (30.37)	$t(83) = -0.30, p = .765$
Nausea	21.12 (22.65)	20.41 (18.69)	$t(83) = -0.16, p = .874$
Total Severity	30.54 (25.47)	28.79 (22.76)	$t(83) = -0.34, p = .739$

Variable	Experimental condition	Control condition	Statistical test
<b>Gender</b>			
Female	31	34	$t(83) = -0.57, p = .573$
Male	11	9	
<b>Type of driving licence</b>			
Full	42	42	N/A
Provisional	0	1	
<b>If a full licence how long have you been driving for?</b>			
Less than one year	5	5	$\chi^2(5) = 1.21, p = .992$
1-2 years	22	22	
3-5 years	8	10	
6-8 years	1	1	
9-10 years	1	0	
10 or more years	4	4	



## REFERENCES

---

- Abrahamse, E., Braem, S., Notebaert, W., & Verguts, T. (2016). Grounding cognitive control in associative learning. *Psychological Bulletin*, *142*(7), 693-728.
- Abramson, C. I., Cakmak, I., Duell, M. E., Bates-Albers, L. M., Zuniga, E. M., Pendegraft, L., . . . Wells, H. (2013). Feature-positive and feature-negative learning in honey bees. *The Journal of Experimental Biology*, *216*(2), 224-229.
- Adams, R. C., Lawrence, N. S., Verbruggen, F., & Chambers, C. D. (2017). Training response inhibition to reduce food consumption: Mechanisms, stimulus specificity and appropriate training protocols. *Appetite*, *109*, 11-23.
- af Wåhlberg, A. E. (2018). The effect of driver improvement interventions on crash involvement; has it been under-estimated? *Transportation Research Part F: Traffic Psychology and Behaviour*, *54*, 349-356.
- Albaum, G. (1997). The Likert scale revisited. *International Journal of Market Research*, *39*(2), 1-21.
- Ajzen, I. (1991). The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, *50*(2), 179-211.
- Ajzen, I. (2002a). Perceived behavioral control, self-efficacy, locus of control, and the theory of planned behavior. *Journal of Applied Social Psychology*, *32*(4), 665-683.
- Ajzen, I. (2002b). Residual effects of past on later behavior: Habituation and reasoned action perspectives. *Personality and Social Psychology Review*, *6*(2), 107-122.
- Ajzen, I., & Fishbein, M. (1980). *Understanding attitudes and predicting social behaviour*. Englewood Cliffs, N.J., USA: Prentice-Hall.
- Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control*, *19*(6), 716-723.
- Aksan, N., Hacker, S. D., Sager, L., Dawson, J., Anderson, S., & Rizzo, M. (2016). Correspondence between Simulator and On-Road Drive Performance: Implications for Assessment of Driving Safety. *Geriatrics*, *1*(1), 8.

- Alcorn III, J. L., Pike, E., Stoops, W. S., Lile, J. A., & Rush, C. R. (2017). A pilot investigation of acute inhibitory control training in cocaine users. *Drug and Alcohol Dependence*, *174*, 145-149.
- Allom, V., Mullan, B., & Hagger, M. (2016). Does inhibitory control training improve health behaviour? A meta-analysis. *Health Psychology Review*, *10*(2), 168-186.
- Altmann, E. M. (2004). The preparation effect in task switching: Carryover of SOA. *Memory & Cognition*, *32*(1), 153-163.
- Anderson, M. C., & Green, C. (2001). Suppressing unwanted memories by executive control. *Nature*, *410*(6826), 366-369.
- Anstey, K. J., Wood, J., Lord, S., & Walker, J. G. (2005). Cognitive, sensory and physical factors enabling driving safety in older adults. *Clinical Psychology Review*, *25*(1), 45-65.
- Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *The Neuroscientist*, *13*(3), 214-228.
- Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychiatry*, *69*(12), e55-e68.
- Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, *18*(4), 177-185.
- Ashe, J., Lungu, O. V., Basford, A. T., & Lu, X. (2006). Cortical control of motor sequences. *Current Opinion in Neurobiology*, *16*(2), 213-221.
- Augustine, S. (2006). Confessions (F. J. Sheed, Trans.). In M. P. Foley (Ed.), *Confessions*. Cambridge, MA, USA: Hackett Publishing.
- Baayen, R. H., & Milin, P. (2010). Analyzing reaction times. *International Journal of Psychological Research*, *3*(2), 12-28.
- Bachoo, S., Bhagwanjee, A., & Govender, K. (2013). The influence of anger, impulsivity, sensation seeking and driver attitudes on risky driving behaviour among post-graduate university students in Durban, South Africa. *Accident Analysis & Prevention*, *55*, 67-76.
- Baddeley, A. (1996). Exploring the central executive. *The Quarterly Journal of Experimental Psychology Section A*, *49*(1), 5-28.
- Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, *4*(10), 829-839.



- Bakeman, R. (2005). Recommended effect size statistics for repeated measures designs. *Behavior Research Methods*, 37(3), 379-384.
- Bakkour, A., Botvinik-Nezer, R., Cohen, N., Hover, A. M., Poldrack, R. A., & Schonberg, T. (2018). Spacing of cue-approach training leads to better maintenance of behavioral change. *PloS One*, 13(7), e0201580.
- Baratian-Ghorghi, F., Zhou, H., & Franco-Watkins, A. (2017). Effects of red light cameras on driver's stop/go decision: Assessing the green extension hypothesis. *Transportation Research Part F: Traffic Psychology and Behaviour*, 46, 87-95.
- Bari, A., & Robbins, T. W. (2013). Inhibition and impulsivity: Behavioral and neural basis of response control. *Progress in Neurobiology*, 108, 44-79.
- Barkley, R. A., & Cox, D. (2007). A review of driving risks and impairments associated with attention-deficit/hyperactivity disorder and the effects of stimulant medication on driving performance. *Journal of Safety Research*, 38(1), 113-128.
- Barrett, L. C., & Livesey, E. J. (2010). Dissociations between expectancy and performance in simple and two-choice reaction-time tasks: A test of associative and nonassociative explanations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(4), 864-877.
- Bartsch, A. L., Kothe, E., Allom, V., Mullan, B., & Houben, K. (2016). The Effect of Non-specific Response Inhibition Training on Alcohol Consumption: An Intervention. *Journal of Addiction Research & Therapy*, 7(1), 1000260.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67(1), 1-48.
- Baumeister, R. F., & Heatherton, T. F. (1996). Self-regulation failure: An overview. *Psychological Inquiry*, 7(1), 1-15.
- Berkman, E. T., Burklund, L., & Lieberman, M. D. (2009). Inhibitory spillover: Intentional motor inhibition produces incidental limbic inhibition via right inferior frontal cortex. *Neuroimage*, 47(2), 705-712.
- Berkman, E. T., Graham, A. M., & Fisher, P. A. (2012). Training self-control: A domain-general translational neuroscience approach. *Child Development Perspectives*, 6(4), 374-384.

- Berryessa, C. M. (2017). Attention, reward, and inhibition: Symptomatic features of ADHD and issues for offenders in the criminal justice system. *ADHD Attention Deficit and Hyperactivity Disorders*, 9(1), 5-10.
- Best, M., Lawrence, N. S., Logan, G. D., McLaren, I. P. L., & Verbruggen, F. (2016). Should I stop or should I go? The role of associations and expectancies. *Journal of Experimental Psychology: Human Perception and Performance*, 42(1), 115-137.
- Beuk, J., Beninger, R., & Paré, M. (2014). Investigating a race model account of executive control in rats with the countermanding paradigm. *Neuroscience*, 263, 96-110.
- Beurden, S. B. v., Smith, J. R., Lawrence, N. S., Abraham, C., & Greaves, C. J. Feasibility randomized controlled trial of ImpulsePal: Smartphone App–Based Weight Management Intervention to Reduce Impulsive Eating in Overweight Adults. *JMIR Formative Research*, 3(2), e1156.
- Bıçaksız, P., & Özkan, T. (2016). Impulsivity and driver behaviors, offences and accident involvement: A systematic review. *Transportation Research Part F: Traffic Psychology and Behaviour*, 38, 194-223.
- Biggs, A. T., Cain, M. S., & Mitroff, S. R. (2015). Cognitive training can reduce civilian casualties in a simulated shooting environment. *Psychological Science*, 26(8), 1164-1176.
- Blair, C., & Razza, R. P. (2007). Relating effortful control, executive function, and false belief understanding to emerging math and literacy ability in kindergarten. *Child Development*, 78(2), 647-663.
- Boecker, M., Gauggel, S., & Druke, B. (2013). Stop or stop-change—Does it make any difference for the inhibition process? *International Journal of Psychophysiology*, 87(3), 234-243.
- Bos, J., Staiger, P. K., Hayden, M. J., Hughes, L. K., Youssef, G., & Lawrence, N. S. (2019). A randomized controlled trial of inhibitory control training for smoking cessation and reduction. *Journal of Consulting and Clinical Psychology*, 87(9), 831-843.
- Bowditch, W. A. (2016). *Response inhibition and associative learning: Training stimulus specific response inhibition* (Doctoral dissertation). Retrieved from <http://hdl.handle.net/10871/24923>

- Bowditch, W. A., Verbruggen, F., & McLaren, I. P. L. (2016). Associatively mediated stopping: Training stimulus-specific inhibitory control. *Learning & Behavior*, *44*(2), 162-174.
- Braver, T. S. (2012). The variable nature of cognitive control: A dual mechanisms framework. *Trends in Cognitive Sciences*, *16*(2), 106-113.
- Briggs, G. F., Hole, G. J., & Turner, J. A. J. (2017). The impact of attentional set and situation awareness on dual tasking driving performance. *Transportation Research Part F: Traffic Psychology and Behaviour*, *57*, 36-47.
- Brooks, M. E., Kristensen, K., van Benthem, K. J., Magnusson, A., Berg, C. W., Nielsen, A., . . . Bolker, B. M. (2017). glmmTMB balances speed and flexibility among packages for zero-inflated generalized linear mixed modeling. *The R Journal*, *9*(2), 378-400.
- Brown, T. G., Ouimet, M. C., Eldeb, M., Tremblay, J., Vingilis, E., Nadeau, L., . . . Bechara, A. (2016). Personality, executive control, and neurobiological characteristics associated with different forms of risky driving. *PLoS One*, *11*(2), e0150227.
- Burgess, C. N. W. (2003). *Breaking the rules of the road : Towards a personality model of rule-governed behaviour* (Doctoral dissertation). Retrieved from <http://ethos.bl.uk/OrderDetails.do?uin=uk.bl.ethos.400920>
- Burnham, K. P., & Anderson, D. R. (2002). *Model Selection and Multimodel Inference: A practical information-theoretic approach* (2nd ed.). NY, USA: Springer-Verlag.
- Burnham, K. P., & Anderson, D. R. (2004). Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods & Research*, *33*(2), 261-304.
- Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: "Gambler's fallacy" versus "hot hand". *Psychonomic Bulletin & Review*, *11*(1), 179-184.
- Button, K. S., Ioannidis, J. P., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S., & Munafò, M. R. (2013). Power failure: Why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, *14*(5), 365-376.
- Camalier, C., Gotler, A., Murthy, A., Thompson, K., Logan, G., Palmeri, T., & Schall, J. (2007). Dynamics of saccade target selection: Race model

- analysis of double step and search step saccade production in human and macaque. *Vision Research*, 47(16), 2187-2211.
- Camp, B., & Lawrence, N. S. (2019). Giving pork the chop: Response inhibition training to reduce meat intake. *Appetite*, 14, 104315.
- Carbine, K. A., & Larson, M. J. (2019). Quantifying the presence of evidential value and selective reporting in food-related inhibitory control training: A p-curve analysis. *Health Psychology Review*, 13(3), 318-343.
- Castanier, C., Deroche, T., & Woodman, T. (2013). Theory of planned behaviour and road violations: The moderating influence of perceived behavioural control. *Transportation Research Part F: Traffic Psychology and Behaviour*, 18, 148-158.
- Castro, L., & Wasserman, E. A. (2016). Executive control and task switching in pigeons. *Cognition*, 146, 121-135.
- Chalmers, A., Howard, D., & Moir, C. (2009, April). *Real virtuality: A step change from virtual reality*. Paper presented at the Proceedings of the 25th Spring Conference on Computer Graphics, Budmerice, Slovakia.
- Chamberlain, S. R., Fineberg, N. A., Blackwell, A. D., Robbins, T. W., & Sahakian, B. J. (2006). Motor inhibition and cognitive flexibility in obsessive-compulsive disorder and trichotillomania. *American Journal of Psychiatry*, 163(7), 1282-1284.
- Chambers, C. D., Garavan, H., & Bellgrove, M. A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience & Biobehavioral Reviews*, 33(5), 631-646.
- Chein, J., Albert, D., O'Brien, L., Uckert, K., & Steinberg, L. (2011). Peers increase adolescent risk taking by enhancing activity in the brain's reward circuitry. *Developmental Science*, 14(2), F1-F10.
- Chen, M. K., Lakshminarayanan, V., & Santos, L. R. (2006). How basic are behavioral biases? Evidence from capuchin monkey trading behavior. *Journal of Political Economy*, 114(3), 517-537.
- Chen, Z., Holland, R., Quandt, J., Dijksterhuis, A., & Veling, H. (2019). When Mere Action Versus Inaction Leads to Robust Preference Change. *Journal of Personality and Social Psychology*, 117(4), 721-740.
- Chen, Z., Veling, H., Dijksterhuis, A., & Holland, R. W. (2018). Do impulsive individuals benefit more from food go/no-go training? Testing the role of inhibition capacity in the no-go devaluation effect. *Appetite*, 124, 99-110.

- Cheng, A. S. K., Ng, T. C. K., & Lee, H. C. (2012). Impulsive personality and risk-taking behavior in motorcycle traffic offenders: A matched controlled study. *Personality and Individual Differences, 53*(5), 597-602.
- Chikazoe, J., Jimura, K., Hirose, S., Yamashita, K.-i., Miyashita, Y., & Konishi, S. (2009). Preparation to inhibit a response complements response inhibition during performance of a stop-signal task. *Journal of Neuroscience, 29*(50), 15870-15877.
- Cirillo, J., Cowie, M. J., MacDonald, H. J., & Byblow, W. D. (2017). Response inhibition activates distinct motor cortical inhibitory processes. *Journal of Neurophysiology, 119*(3), 877-886.
- Clegg, B. A., DiGirolamo, G. J., & Keele, S. W. (1998). Sequence learning. *Trends in Cognitive Sciences, 2*(8), 275-281.
- Cohen, A., & Curran, T. (1993). On tasks, knowledge, correlations, and dissociations: Comment on Perruchet and Amorim (1992). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 19*(6), 1431-1437.
- Cohen, J. R., & Lieberman, M. D. (2010). The common neural basis of exerting self-control in multiple domains. In R. Hassin, K. Ochsner, Y. Trope (Eds.), *Self Control in Society, Mind, and Brain* (pp. 141-163). New York, N.Y., USA: Oxford University Press.
- Conner, M., & Sparks, P. (2005). Theory of planned behaviour and health behaviour. In M. Conner & P. Norman (Eds.), *Predicting Health Behaviour* (2nd ed., pp. 170-222). Berkshire, England: Open University Press.
- Conway, C. M., & Christiansen, M. H. (2001). Sequential learning in non-human primates. *Trends in Cognitive Sciences, 5*(12), 539-546.
- Cook, S. A., & Harris, R. E. (1937). The verbal conditioning of the galvanic skin reflex. *Journal of Experimental Psychology, 21*(2), 202-210.
- Coulacoglou, C., & Saklofske, D. H. (2017). Executive Function, Theory of Mind, and Adaptive Behavior. In C. Coulacoglou & D. H. Saklofske (Eds.), *Psychometrics and Psychological Assessment: Principles and Applications* (pp. 91-131). San Diego, USA: Academic Press.
- Coxon, J. P., Stinear, C. M., & Byblow, W. D. (2006). Intracortical inhibition during volitional inhibition of prepared action. *Journal of Neurophysiology, 95*(6), 3371-3383.

- Cribari-Neto, F., & Zeileis, A. (2010). Beta Regression in R. *Journal of Statistical Software*, 34(2), 13-37.
- Curran, M. F., Fuertes, J. N., Alfonso, V. C., & Hennessy, J. J. (2010). The association of sensation seeking and impulsivity to driving while under the influence of alcohol. *Journal of Addictions & Offender Counseling*, 30(2), 84-98.
- Custance, D., Whiten, A., & Fredman, T. (1999). Social learning of an artificial fruit task in capuchin monkeys (*Cebus apella*). *Journal of Comparative Psychology*, 113(1), 13-23.
- Dawson, M. E. (1970). Cognition and conditioning: Effects of masking the CS-UCS contingency on human GSR classical conditioning. *Journal of Experimental Psychology*, 85(3), 389-396.
- Dawson, M. E., & Biferno, M. A. (1973). Concurrent measurement of awareness and electrodermal classical conditioning. *Journal of Experimental Psychology*, 101(1), 55-62.
- De Houwer, J. (2009). The propositional approach to associative learning as an alternative for association formation models. *Learning & Behavior*, 37(1), 1-20.
- De Houwer, J., & Beckers, T. (2003). Secondary task difficulty modulates forward blocking in human contingency learning. *The Quarterly Journal of Experimental Psychology Section B*, 56(4b), 345-357.
- De Houwer, J., Tanaka, A., Moors, A., & Tibboel, H. (2018). Kicking the habit: Why evidence for habits in humans might be overestimated. *Motivation Science*, 4(1), 50-59.
- De Martino, B., Kumaran, D., Seymour, B., & Dolan, R. J. (2006). Frames, biases, and rational decision-making in the human brain. *Science*, 313(5787), 684-687.
- de Wit, S., Kindt, M., Knot, S. L., Verhoeven, A. A., Robbins, T. W., Gasull-Camos, J., . . . Gillan, C. M. (2018). Shifting the balance between goals and habits: Five failures in experimental habit induction. *Journal of Experimental Psychology: General*, 147(7), 1043-1065.
- Department of Health. (2016). *Guide to creating a front of pack (FoP) nutrition label for pre-packed products sold through retail outlets*. London: Department of Health.

- Department of Transport. (2018). *Reported Road Casualties Great Britain: 2017 Annual Report*. London: Department of Transport.
- Der, G., & Deary, I. J. (2006). Age and sex differences in reaction time in adulthood: Results from the United Kingdom Health and Lifestyle Survey. *Psychology and Aging, 21*(1), 62-73.
- Descartes, R. (1989). *Passions of the Soul* (S. H. Voss, Trans.). In S. H. Voss (Ed.), *Passions of the Soul*. Indianapolis, USA: Hackett Publishing.
- Deutsch, R. (2016). Associative and Propositional Processes from the Perspective of the Reflective–Impulsive Model. In R. Deutsch, B. Gawronski, & W. Hofmann (Eds.), *Reflective and Impulsive Determinants of Human Behavior* (pp. 51-69). New York, N.Y., USA: Routledge.
- Diamond, A. (2013). Executive functions. *Annual Review of Psychology, 64*, 135-168.
- Diamond, S., Balvin, R. S., & Diamond, F. R. (1963). *Inhibition and choice: A neurobehavioral approach to problems of plasticity in behavior*. New York, N.Y., USA: Harper & Row Publishers.
- Dickinson, A. (1985). Actions and habits: The development of behavioural autonomy. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences, 308*(1135), 67-78.
- Dickman, S. J. (1990). Functional and dysfunctional impulsivity: Personality and cognitive correlates. *Journal of Personality and Social Psychology, 58*(1), 95-102.
- Dickman, S. J., & Meyer, D. E. (1988). Impulsivity and speed-accuracy tradeoffs in information processing. *Journal of Personality and Social Psychology, 54*(2), 274-290.
- Dienes, Z. (2011). Bayesian Versus Orthodox Statistics: Which Side Are You On? *Perspectives on Psychological Science, 6*(3), 274-290.
- Dienes, Z., & Altmann, G. (1997). Transfer of implicit knowledge across domains: How implicit and how abstract? In D. C. Berry (Ed.), *Debates in psychology. How implicit is implicit learning?* (pp. 107-123). New York, N.Y., USA: Oxford University Press.
- Dienes, Z., Altmann, G., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*(5), 1322-1338.

- Dienes, Z., & Perner, J. (1996). Implicit knowledge in people and connectionist networks. In G. D. M. Underwood (Ed.), *Implicit Cognition* (pp. 227-255). New York, N.Y., USA: Oxford University Press.
- Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22(5), 735-808.
- Dienes, Z., & Scott, R. (2005). Measuring unconscious knowledge: Distinguishing structural knowledge and judgment knowledge. *Psychological Research*, 69(5-6), 338-351.
- Donders, F. C. (1868). On the speed of mental processes. In W. G. Koster (Ed.), *Attention and Performance II* (pp. 412-431). Amsterdam: North-Holland.
- Dopson, J. (2018, March). *The Use of Positive Reinforcement Techniques during Guide Dog Training*. Paper presented at the Associative Learning Symposium 2018, Gregynog, Wales.
- Dostoevsky, F. (2018). Winter notes on summer impression (K. FitzLyon, Trans.). *Winter notes on summer impression*. Surrey, UK: Alma Books.
- Dreisbach, G. (2012). Mechanisms of cognitive control: The functional role of task rules. *Current Directions in Psychological Science*, 21(4), 227-231.
- Dreisbach, G., Goschke, T., & Haider, H. (2007). The role of task rules and stimulus–response mappings in the task switching paradigm. *Psychological Research*, 71(4), 383-392.
- Duque, J., Lew, D., Mazzocchio, R., Olivier, E., & Ivry, R. B. (2010). Evidence for two concurrent inhibitory mechanisms during response preparation. *Journal of Neuroscience*, 30(10), 3793-3802.
- Eagle, D. M., Bari, A., & Robbins, T. W. (2008). The neuropsychopharmacology of action inhibition: Cross-species translation of the stop-signal and go/no-go tasks. *Psychopharmacology*, 199(3), 439-456.
- Elchlepp, H., & Verbruggen, F. (2017). How to withhold or replace a prepotent response: An analysis of the underlying control processes and their temporal dynamics. *Biological Psychology*, 123, 250-268.
- Elgendi, M., Kumar, P., Barbic, S., Howard, N., Abbott, D., & Cichocki, A. (2018). Subliminal priming—state of the art and future perspectives. *Behavioral Sciences*, 8(6), 54.



- Elliott, M. A., & Armitage, C. J. (2006). Effects of implementation intentions on the self-reported frequency of drivers' compliance with speed limits. *Journal of Experimental Psychology: Applied*, *12*(2), 108-117.
- Elliott, M. A., & Thomson, J. A. (2010). The social cognitive determinants of offending drivers' speeding behaviour. *Accident Analysis & Prevention*, *42*(6), 1595-1605.
- Evans, J. S. B. (2008). Dual-processing accounts of reasoning, judgment, and social cognition. *Annual Review of Psychology*, *59*, 255-278.
- Evans, J. S. B., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223-241.
- Faul, F., Erdfelder, E., Lang, A. G., & Buchner, A. (2007). G\*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavioural Research Methods*, *39*(2), 175-191.
- Felicio, G. P., Grepo, L. C., Reyes, V. F., & Yupingkun, L. C. (2015). Traffic Light Displays and Driver Behaviors: A Case Study. *Procedia Manufacturing*, *3*, 3266-3273.
- Ferrier, D. (1886). *The functions of the brain*: London, UK: Smith, Elder.
- Fischer, M., Barkley, R. A., Smallish, L., & Fletcher, K. (2007). Hyperactive children as young adults: Driving abilities, safe driving behavior, and adverse driving outcomes. *Accident Analysis & Prevention*, *39*(1), 94-105.
- Forman, E. M., Manasse, S. M., Dallal, D. H., Crochiere, R. J., Loyka, C. M., Butryn, M. L., . . . Houben, K. (2019). Computerized neurocognitive training for improving dietary health and facilitating weight loss. *Journal of Behavioral Medicine*, 1-12.
- Forrest, C. L. D., Monsell, S., & McLaren, I. P. L. (2014). Is performance in task-cuing experiments mediated by task set selection or associative compound retrieval? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *40*(4), 1002-1024.
- Fors, C., Ahlstrom, C., & Anund, A. (2018). A comparison of driver sleepiness in the simulator and on the real road. *Journal of Transportation Safety & Security*, *10*(1-2), 72-87.
- Fox, J. (2008). *Applied Regression Analysis and Generalized Linear Models* (2nd ed.). CA, USA: Sage Publications.

- Frane, A. V. (2015). Planned Hypothesis Tests Are Not Necessarily Exempt From Multiplicity Adjustment. *Journal of Research Practice*, 11(1), Article P2.
- Fylan, F. (2017). *Using Behaviour Change Techniques: Guidance for the road safety community*. RAC Foundation.
- Gardner, B. (2012). Habit as automaticity, not frequency. *European Health Psychologist*, 14(2), 32-36.
- Gardner, B. (2015a). Defining and measuring the habit impulse: Response to commentaries. *Health Psychology Review*, 9(3), 318-322.
- Gardner, B. (2015b). A review and analysis of the use of 'habit'in understanding, predicting and influencing health-related behaviour. *Health Psychology Review*, 9(3), 277-295.
- Gardner, B., & Tang, V. (2014). Reflecting on non-reflective action: An exploratory think-aloud study of self-report habit measures. *British Journal of Health Psychology*, 19(2), 258-273.
- Ginsburg, S., & Jablonka, E. (2010). The evolution of associative learning: A factor in the Cambrian explosion. *Journal of Theoretical Biology*, 266(1), 11-20.
- Glendon, A. I., McNally, B., Jarvis, A., Chalmers, S. L., & Salisbury, R. L. (2014). Evaluating a novice driver and pre-driver road safety intervention. *Accident Analysis & Prevention*, 64, 100-110.
- Godefroy, O., Cabaret, M., Petit-Chenal, V., Pruvo, J.-P., & Rousseaux, M. (1999). Control functions of the frontal lobes. Modularity of the central-supervisory system? *Cortex*, 35(1), 1-20.
- Gollwitzer, P. M., Sheeran, P., Trotschel, R., & Webb, T. L. (2011). Self-regulation of priming effects on behavior. *Psychological Science*, 22(7), 901-907.
- Gotler, A., Meiran, N., & Tzelgov, J. (2003). Nonintentional task set activation: Evidence from implicit task sequence learning. *Psychonomic Bulletin & Review*, 10(4), 890-896.
- Grant, W. M. (2010). Can a Libertarian Hold that Our Free Acts are Caused by God? *Faith and Philosophy*, 27(1), 22-44.
- Green, P., & MacLeod, C. J. (2016). SIMR: An R package for power analysis of generalized linear mixed models by simulation. *Methods in Ecology and Evolution*, 7(4), 493-498.

- Groom, M. J., van Loon, E., Daley, D., Chapman, P., & Hollis, C. (2015). Driving behaviour in adults with attention deficit/hyperactivity disorder. *BMC Psychiatry*, *15*(1), 175.
- Guerrieri, R., Nederkoorn, C., & Jansen, A. (2012). Disinhibition is easier learned than inhibition. The effects of (dis) inhibition training on food intake. *Appetite*, *59*(1), 96-99.
- Gulberti, A., Arndt, P., & Colonius, H. (2014). Stopping eyes and hands: Evidence for non-independence of stop and go processes and for a separation of central and peripheral inhibition. *Frontiers in Human Neuroscience*, *8*, 61.
- Hagger, M. S., Rebar, A. L., Mullan, B., Lipp, O. V., & Chatzisarantis, N. L. D. (2015). The subjective experience of habit captured by self-report indexes may lead to inaccuracies in the measurement of habitual action. *Health Psychology Review*, *9*(3), 296-302.
- Harrison, X. A., Donaldson, L., Correa-Cano, M. E., Evans, J., Fisher, D. N., Goodwin, C. E., . . . Inger, R. (2018). A brief introduction to mixed effects modelling and multi-model inference in ecology. *PeerJ*, *6*, e4794.
- Hatfield, J., Williamson, A., Kehoe, E. J., Lemon, J., Arguel, A., Prabhakaran, P., & Job, R. S. (2018). The effects of training impulse control on simulated driving. *Accident Analysis & Prevention*, *119*, 1-15.
- Hatfield, J., Williamson, A., Kehoe, E. J., & Prabhakaran, P. (2017). An examination of the relationship between measures of impulsivity and risky simulated driving amongst young drivers. *Accident Analysis & Prevention*, *103*, 37-43.
- Hawkins, R. D., Clark, G. A., & Kandel, E. R. (2006). Operant conditioning of gill withdrawal in *Aplysia*. *Journal of Neuroscience*, *26*(9), 2443-2448.
- Hochman, S., Henik, A., & Kalanithroff, E. (2018). Stopping at a red light: Recruitment of inhibitory control by environmental cues. *PLoS One*, *13*(5), e0196199.
- Hoffman, L., & McDowd, J. M. (2010). Simulator driving performance predicts accident reports five years later. *Psychology and Aging*, *25*(3), 741-745.
- Hofmann, W., Schmeichel, B. J., & Baddeley, A. D. (2012). Executive functions and self-regulation. *Trends in Cognitive Sciences*, *16*(3), 174-180.
- Houben, K., Havermans, R. C., Nederkoorn, C., & Jansen, A. (2012). Beer à No-Go: Learning to stop responding to alcohol cues reduces alcohol

- intake via reduced affective associations rather than increased response inhibition. *Addiction*, *107*(7), 1280-1287.
- Houben, K., & Jansen, A. (2011). Training inhibitory control. A recipe for resisting sweet temptations. *Appetite*, *56*(2), 345-349.
- Houben, K., & Jansen, A. (2015). Chocolate equals stop. Chocolate-specific inhibition training reduces chocolate intake and go associations with chocolate. *Appetite*, *87*, 318-323.
- Houben, K., Nederkoorn, C., Wiers, R. W., & Jansen, A. (2011). Resisting temptation: Decreasing alcohol-related affect and drinking behavior by training response inhibition. *Drug and Alcohol Dependence*, *116*(1-3), 132-136.
- House of Commons Transport Committee. (2019). *Road safety: driving while using a mobile phone*. London: House of Commons.
- Hull, C. L. (1943). *Principles of behavior: An introduction to behavior theory*. Oxford, England: Appleton-Century.
- International Transport Forum. (2018). *Road safety annual report 2018: United Kingdom*.
- Jahfari, S., Stinear, C. M., Claffey, M., Verbruggen, F., & Aron, A. R. (2010). Responding with restraint: What are the neurocognitive mechanisms? *Journal of Cognitive Neuroscience*, *22*(7), 1479-1492.
- Jahfari, S., Verbruggen, F., Frank, M. J., Waldorp, L. J., Colzato, L., Ridderinkhof, K. R., & Forstmann, B. U. (2012). How preparation changes the need for top-down control of the basal ganglia when inhibiting premature actions. *Journal of Neuroscience*, *32*(32), 10870-10878.
- Jami, S. A., Wright, W. G., & Glanzman, D. L. (2007). Differential classical conditioning of the gill-withdrawal reflex in *Aplysia* recruits both NMDA receptor-dependent enhancement and NMDA receptor-dependent depression of the reflex. *Journal of Neuroscience*, *27*(12), 3064-3068.
- Jason, L. A., Neal, A. M., & Marinakis, G. (1985). Altering contingencies to facilitate compliance with traffic light systems. *Journal of Applied Behavior Analysis*, *18*(1), 95-100.
- Jenkins, H. M., & Sainsbury, R. S. (1970). Discrimination learning with the distinctive feature on positive or negative trials. In D. Mostofsky (Ed.),

- Attention: Contemporary Theory and Analysis* (pp. 239-275). New York, N.Y., USA: Appleton-Century-Crofts.
- Jha, A., Nachev, P., Barnes, G., Husain, M., Brown, P., & Litvak, V. (2015). The frontal control of stopping. *Cerebral Cortex*, *25*(11), 4392-4406.
- Jiménez-Fernández, G., Vaquero, J. M., Jiménez, L., & Defior, S. (2011). Dyslexic children show deficits in implicit sequence learning, but not in explicit sequence learning or contextual cueing. *Annals of Dyslexia*, *61*(1), 85-110.
- Joblove, G. H., & Greenberg, D. (1978). *Color spaces for computer graphics*. Paper presented at the Proceedings of the 5th annual conference on Computer graphics and interactive techniques, Atlanta, G.A., USA.
- Johnson, P. C. D., Barry, S. J. E., Ferguson, H. M., & Müller, P. (2015). Power analysis for generalized linear mixed models in ecology and evolution. *Methods in Ecology and Evolution*, *6*(2), 133-142.
- Johnstone, S. J., Roodenrys, S., Blackman, R., Johnston, E., Loveday, K., Mantz, S., & Barratt, M. F. (2012). Neurocognitive training for children with and without AD/HD. *ADHD Attention Deficit and Hyperactivity Disorders*, *4*(1), 11-23.
- Jones, A., Di Lemma, L. C., Robinson, E., Christiansen, P., Nolan, S., Tudur-Smith, C., & Field, M. (2016). Inhibitory control training for appetitive behaviour change: A meta-analytic investigation of mechanisms of action and moderators of effectiveness. *Appetite*, *97*, 16-28.
- Jones, A., & Field, M. (2013). The effects of cue-specific inhibition training on alcohol consumption in heavy social drinkers. *Experimental and Clinical Psychopharmacology*, *21*(1), 8-16.
- Jones, A., Hardman, C. A., Lawrence, N., & Field, M. (2017). Cognitive training as a potential treatment for overweight and obesity: A critical review of the evidence. *Appetite*, *124*, 50-67.
- Jones, A., McGrath, E., Robinson, E., Houben, K., Nederkoorn, C., & Field, M. (2018). A randomized controlled trial of inhibitory control training for the reduction of alcohol consumption in problem drinkers. *Journal of Consulting and Clinical Psychology*, *86*(12), 991-1004.
- Jones, A., Robinson, E., Duckworth, J., Kersbergen, I., Clarke, N., & Field, M. (2018). The effects of exposure to appetitive cues on inhibitory control: A meta-analytic investigation. *Appetite*, *128*, 271-282.

- Jones, F. W., & McLaren, I. P. L. (2009). Human sequence learning under incidental and intentional conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, 35(4), 538-553.
- Jones, M., Curran, T., Mozer, M. C., & Wilder, M. H. (2013). Sequential effects in response time reveal learning mechanisms and event representations. *Psychology Review*, 120(3), 628-666.
- Jongen, E. M., Brijs, K., Komlos, M., Brijs, T., & Wets, G. (2011). Inhibitory control and reward predict risky driving in young novice drivers—a simulator study. *Procedia-Social and Behavioral Sciences*, 20, 604-612.
- Kahneman, D. (2011). *Thinking, fast and slow*. London, UK: Penguin Group.
- Kandel, E. R. (1976). *Cellular basis of behavior: An introduction to behavioral neurobiology*. Oxford, England: W. H. Freeman.
- Kawai, N., & Matsuzawa, T. (2000). Cognition: numerical memory span in a chimpanzee. *Nature*, 403(6765), 39-40.
- Kemeny, A., & Panerai, F. (2003). Evaluating perception in driving simulation experiments. *Trends in Cognitive Sciences*, 7(1), 31-37.
- Kennedy, R. S., Lane, N. E., Berbaum, K. S., & Lilienthal, M. G. (1993). Simulator Sickness Questionnaire: An Enhanced Method for Quantifying Simulator Sickness. *The International Journal of Aviation Psychology*, 3(3), 203-220.
- Ker, K., Roberts, I., Collier, T., Beyer, F., Bunn, F., & Frost, C. (2005). Post-licence driver education for the prevention of road traffic crashes: A systematic review of randomised controlled trials. *Accident Analysis & Prevention*, 37(2), 305-313.
- Kiesel, A., Steinhauser, M., Wendt, M., Falkenstein, M., Jost, K., Philipp, A. M., & Koch, I. (2010). Control and interference in task switching—A review. *Psychological Bulletin*, 136(5), 849-874.
- Kiriakidis, S. (2017). Perceived behavioural control in the theory of planned behaviour: Variability of conceptualization and operationalization and implications for measurement. In A. Kavoura, D. P. Saksas, & P. Tomaras (Eds.), *Strategic Innovative Marketing: 4<sup>th</sup> IC-SIM* (pp. 197-202). Switzerland: Springer Proceedings in Business and Economics.
- Kleiman, E. (2017). EMAtools: Data Management Tools for Real-Time Monitoring/Ecological Momentary Assessment Data. R package version 0.1.3. <https://CRAN.R-project.org/package=EMAtools>

- Kuznetsova, A., Brockhoff, P. B., & Christensen, R. H. B. (2016). lmerTest: Tests in Linear Mixed Effects Models. R package version 2.0-33.
- Labrecque, J. S., & Wood, W. (2015). What measures of habit strength to use? Comment on Gardner (2015). *Health Psychology Review*, 9(3), 303-310.
- Ladefoged, P., Silverstein, R., & Papçun, G. (1973). Interruptibility of speech. *The Journal of the Acoustical Society of America*, 54(4), 1105-1108.
- Lajunen, T., Parker, D., & Summala, H. (2004). The Manchester Driver Behaviour Questionnaire: A cross-cultural study. *Accident Analysis & Prevention*, 36(2), 231-238.
- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4, 863.
- Lakshminaryanan, V., Keith Chen, M., & Santos, L. R. (2008). Endowment effect in capuchin monkeys. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 363(1511), 3837-3844.
- Lappin, J. S., & Eriksen, C. W. (1966). Use of a delayed signal to stop a visual reaction-time response. *Journal of Experimental Psychology*, 72(6), 805-811.
- Lavric, A., Pizzagalli, D. A., & Forstmeier, S. (2004). When 'go' and 'nogo' are equally frequent: ERP components and cortical tomography. *European Journal of Neuroscience*, 20(9), 2483-2488.
- Lawrence, M. A. (2016). ez: Easy Analysis and Visualization of Factorial Experiments. R package version 4.4-0.
- Lawrence, N. S., O'Sullivan, J., Parslow, D., Javaid, M., Adams, R. C., Chambers, C. D., . . . Verbruggen, F. (2015). Training response inhibition to food is associated with weight loss and reduced energy intake. *Appetite*, 95, 17-28.
- Lawrence, N. S., Verbruggen, F., Morrison, S., Adams, R. C., & Chambers, C. D. (2015). Stopping to food can reduce intake. Effects of stimulus-specificity and individual differences in dietary restraint. *Appetite*, 85, 91-103.
- Lea, S. E., Chow, K., Meier, C., McLaren, I. P. L., & Verbruggen, F. (2019). Pigeons' performance in a tracking change-signal procedure is consistent with the independent horse-race mode. *Journal of*

- Experimental Psychology: Animal Learning and Cognition*, 45(4), 464-473.
- Lea, S. E., & Wills, A. J. (2008). Use of multiple dimensions in learned discriminations. *Comparative Cognition & Behavior Reviews*, 3, 115-133.
- Leiva, A., Parmentier, F. B., Elchlepp, H., & Verbruggen, F. (2015). Reorienting the mind: The impact of novel sounds on go/no-go performance. *Journal of Experimental Psychology: Human Perception and Performance*, 41(5), 1197-1202.
- Lenartowicz, A., Verbruggen, F., Logan, G. D., & Poldrack, R. A. (2011). Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. *Journal of Cognitive Neuroscience*, 23(11), 3388-3399.
- Leocani, L., Cohen, L. G., Wassermann, E. M., Ikoma, K., & Hallett, M. (2000). Human corticospinal excitability evaluated with transcranial magnetic stimulation during different reaction time paradigms. *Brain*, 123(6), 1161-1173.
- Lheureux, F., Auzoult, L., Charlois, C., Hardy-Massard, S., & Minary, J. P. (2016). Traffic Offences: Planned or Habitual? Using the Theory of Planned Behaviour and habit strength to explain frequency and magnitude of speeding and driving under the influence of alcohol. *British Journal of Psychology*, 107(1), 52-71.
- Li, X., Li, B., Liu, X., Lages, M., & Stoet, G. (2019). Task-switching costs disappear if non-Chinese participants respond to Chinese characters. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 45(11), 2051–2071.
- Littman, R., & Takács, Á. (2017). Do all inhibitions act alike? A study of go/no-go and stop-signal paradigms. *PLoS One*, 12(10), e0186774.
- Livesey, E. J., & Costa, D. S. J. (2014). Automaticity and conscious control in single and choice response time versions of the Perruchet effect. *Quarterly Journal of Experimental Psychology*, 67(4), 646-664.
- Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, 95(4), 492-527.
- Logan, G. D. (1994). On the ability to inhibit thought and action: A users' guide to the stop signal paradigm. In D. Dagenbach & T. H. Carr (Eds.),



- Inhibitory processes in attention, memory, and language* (pp. 189-239).  
San Diego, CA, USA: Academic Press.
- Logan, G. D., & Burkell, J. (1986). Dependence and independence in responding to double stimulation: A comparison of stop, change, and dual-task paradigms. *Journal of Experimental Psychology: Human Perception and Performance*, 12(4), 549-563.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91(3), 295-327.
- Logan, G. D., & Irwin, D. E. (2000). Don't look! Don't touch! Inhibitory control of eye and hand movements. *Psychonomic Bulletin & Review*, 7(1), 107-112.
- Logan, G. D., Van Zandt, T., Verbruggen, F., & Wagenmakers, E.-J. (2014). On the ability to inhibit thought and action: General and special theories of an act of control. *Psychological Review*, 121(1), 66-95.
- Lotz, A., Uengoer, M., Koenig, S., Pearce, J. M., & Lachnit, H. (2012). An exploration of the feature-positive effect in adult humans. *Learning and Behavior*, 40(2), 222-230.
- Lovibond, P. F. (2003). Causal beliefs and conditioned responses: Retrospective reevaluation induced by experience and by instruction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(1), 97-106.
- Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, 28(1), 3-26.
- Lum, J. A., Ullman, M. T., & Conti-Ramsden, G. (2013). Procedural learning is impaired in dyslexia: Evidence from a meta-analysis of serial reaction time studies. *Research in Developmental Disabilities*, 34(10), 3460-3476.
- MacLeod, C. M. (2007). The concept of inhibition in cognition. In D. S. Gorfein & C. M. MacLeod (Eds.), *Inhibition in cognition* (pp. 3-23). Washington, DC, US: American Psychological Association.
- MacLeod, C. M., Dodd, M. D., Sheard, E. D., Wilson, D. E., & Bibi, U. (2003). In opposition to inhibition. In B. H. Ross (Ed.), *The Psychology of Learning and Motivation: Advances in research and theory*, Vol. 43, pp. 163-214). New York, NY, US: Elsevier Science.

- Macmillan, M. (1996). The concept of inhibition in some nineteenth century theories of thinking. *Brain and Cognition*, 30(1), 4-19.
- Majid, D. S. A., Cai, W., George, J. S., Verbruggen, F., & Aron, A. R. (2011). Transcranial Magnetic Stimulation Reveals Dissociable Mechanisms for Global Versus Selective Corticomotor Suppression Underlying the Stopping of Action. *Cerebral Cortex*, 22(2), 363-371.
- Mäntylä, T., Karlsson, M. J., & Marklund, M. (2009). Executive control functions in simulated driving. *Applied Neuropsychology*, 16(1), 11-18.
- Manuel, A. L., Bernasconi, F., & Spierer, L. (2013). Plastic modifications within inhibitory control networks induced by practicing a stop-signal task: An electrical neuroimaging study. *Cortex*, 49(4), 1141-1147.
- Marteau, T. M., Hollands, G. J., & Fletcher, P. C. (2012). Changing human behavior to prevent disease: The importance of targeting automatic processes. *Science*, 337(6101), 1492-1495.
- Matzke, D., Verbruggen, F., & Logan, G. D. (2018). The Stop-Signal Paradigm. In E.-J. Wagenmakers, & J. T. Wixted (Eds.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience: Volume 5: Methodology* (pp. 383-429). New York, N.Y., USA: John Wiley & Sons, Inc.
- Mayr, U., Diedrichsen, J., Ivry, R., & Keele, S. W. (2006). Dissociating task-set selection from task-set inhibition in the prefrontal cortex. *Journal of Cognitive Neuroscience*, 18(1), 14-21.
- McAndrew, A., Jones, F. W., McLaren, R. P., & McLaren, I. P. L. (2012). Dissociating expectancy of shock and changes in skin conductance: An investigation of the Perruchet effect using an electrodermal paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*, 38(2), 203-208.
- McDaniel, M. A., & Scullin, M. K. (2010). Implementation intention encoding does not automatize prospective memory responding. *Memory & Cognition*, 38(2), 221-232.
- McEachan, R. R. C., Conner, M., Taylor, N. J., & Lawton, R. J. (2011). Prospective prediction of health-related behaviours with the theory of planned behaviour: A meta-analysis. *Health Psychology Review*, 5(2), 97-144.

- McLaren, I. P. L., Forrest, C. L. D., McLaren, R. P., Jones, F. W., Aitken, M. R. F., & Mackintosh, N. J. (2014). Associations and propositions: The case for a dual-process account of learning in humans. *Neurobiology of Learning and Memory*, *108*, 185-195.
- McLaren, I. P. L., Green, R. E. A., & Mackintosh, N. J. (1994). Animal learning and the implicit/explicit distinction. In N. C. Ellis (Ed.), *Implicit and explicit learning of languages* (pp. 313-332). London, UK: Academic Press.
- McLaren, I. P. L., Jones, F. W., McLaren, R. P., & Yeates, F. (2013). *Cue Competition in Human Incidental Learning*. Paper presented at the Proceedings of the 1st Reinforcement Learning and Decision Making Meeting, Princeton, New Jersey, USA.
- McLaren, I. P. L., McAndrew, A., Angerer, K., McLaren, R. P., Forrest, C., Bowditch, W. A., . . . Verbruggen, F. (2019). Association and Cognition: Two Processes, One System. *Quarterly Journal of Experimental Psychology*, *72*(2), 98-117.
- McLaren, I. P. L., & Verbruggen, F. (Eds.). (2016). Association, inhibition, and action. In R. Murphy & R. Honey (Eds.), *The Wiley Handbook on the Cognitive Neuroscience of Learning* (pp. 489–514). Chichester, England: Wiley-Blackwell.
- Meier, C., Lea, S. E., Forrest, C. L., Angerer, K., & McLaren, I. P. L. (2013). Comparative evidence for associative learning in task switching. In M. Knauff, N. Sebanz, M. Pauen, & I. Wachsmuth (Eds.), *Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 1020-1026). Austin, TX: Cognitive Science Society.
- Meier, C., Lea, S. E., & McLaren, I. P. L. (2016). Task-switching in pigeons: Associative learning or executive control? *Journal of Experimental Psychology: Animal Learning and Cognition*, *42*(2), 163-176.
- Meier, C., Lea, S. E., & McLaren, I. P. L. (2018). Pigeons in control of their actions: Learning and performance in stop-signal and change-signal tasks. *Journal of Experimental Psychology: Animal Learning and Cognition*, *44*(1), 82-94.
- Meier, C., Pant, S. R., van Horik, J. O., Laker, P. R., Langley, E. J., Whiteside, M. A., . . . Madden, J. R. (2017). A novel continuous inhibitory-control task: Variation in individual performance by young pheasants (*Phasianus colchicus*). *Animal Cognition*, *20*(6), 1035-1047.

- Meiran, N. (1996). Reconfiguration of processing mode prior to task performance. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(6), 1423-1442.
- Meiran, N. (2000). Modeling cognitive control in task-switching. *Psychological Research*, 63(3-4), 234-249.
- Melo, M., Rocha, T., Barbosa, L., & Bessa, M. (2016). *The Impact of Body Position on the Usability of Multisensory Virtual Environments: Case study of a virtual bicycle*. Paper presented at the Proceedings of the 7th International Conference on Software Development and Technologies for Enhancing Accessibility and Fighting Info-exclusion, Portugal.
- Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, 24(1), 167-202.
- Milner, B. (1963). Effects of different brain lesions on card sorting: The role of the frontal lobes. *Archives of Neurology*, 9(1), 90-100.
- Milner, B., & Petrides, M. (1984). Behavioural effects of frontal-lobe lesions in man. *Trends in Neurosciences*, 7(11), 403-407.
- Mitchell, C. J., De Houwer, J., & Lovibond, P. F. (2009). The propositional nature of human associative learning. *Behavioral and Brain Sciences*, 32(2), 183-198.
- Mitchell, C. J., Wardle, S. G., Lovibond, P. F., Weidemann, G., & Chang, B. P. (2010). Do reaction times in the Perruchet effect reflect variations in the strength of an associative link? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(2), 567-572.
- Miyake, A., Friedman, N. P., Emerson, M. J., Witzki, A. H., Howerter, A., & Wager, T. D. (2000). The unity and diversity of executive functions and their contributions to complex "frontal lobe" tasks: A latent variable analysis. *Cognitive Psychology*, 41(1), 49-100.
- Moffitt, T. E. (2017). Adolescence-limited and life-course-persistent antisocial behavior: A developmental taxonomy. In K. B. Beaver (Ed.), *Biosocial Theories of Crime* (pp. 69-96). Oxford, UK: Routledge.
- Moffitt, T. E., Arseneault, L., Belsky, D., Dickson, N., Hancox, R. J., Harrington, H., . . . Caspi, A. (2011). A gradient of childhood self-control predicts health, wealth, and public safety. *Proceedings of the National Academy of Sciences of the United States of America*, 108(7), 2693-2698

- Monsell, S. (1996). Control of mental processes. In V. Bruce (Ed.), *Unsolved mysteries of the mind: Tutorial essays in cognition* (pp. 93-148). East Sussex, UK: Taylor & Francis.
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7(3), 134-140.
- Monsell, S., & Driver, J. (2000). *Control of cognitive processes: Attention and performance XVIII*. Cambridge, M.A, USA: MIT Press.
- Monsell, S., & Mizon, G. A. (2006). Can the task-cuing paradigm measure an endogenous task-set reconfiguration process? *Journal of Experimental Psychology: Human Perception and Performance*, 32(3), 493-516.
- Monsell, S., Sumner, P., & Waters, H. (2003). Task-set reconfiguration with predictable and unpredictable task switches. *Memory & Cognition*, 31(3), 327-342.
- Munakata, Y., Snyder, H. R., & Chatham, C. H. (2012). Developing cognitive control: Three key transitions. *Current Directions in Psychological Science*, 21(2), 71-77.
- Nachev, P., Wydell, H., O'neill, K., Husain, M., & Kennard, C. (2007). The role of the pre-supplementary motor area in the control of action. *Neuroimage*, 36, T155-T163.
- Nakagawa, S., Johnson, P. C., & Schielzeth, H. (2017). The coefficient of determination R<sup>2</sup> and intra-class correlation coefficient from generalized linear mixed-effects models revisited and expanded. *Journal of the Royal Society Interface*, 14(134), 20170213.
- Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R<sup>2</sup> from generalized linear mixed-effects models. *Methods in Ecology and Evolution*, 4(2), 133-142.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264-336). Hillsdale, NJ, US: Lawrence Erlbaum Associates, Inc.
- Nemme, H. E., & White, K. M. (2010). Texting while driving: Psychosocial influences on young people's texting intentions and behaviour. *Accident Analysis & Prevention*, 42(4), 1257-1265.
- Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *Behavioral and Brain Sciences*, 37(1), 1-19.

- Newman, J., Wolff, W. T., & Hearst, E. (1980). The feature-positive effect in adult human participants. *Journal of Experimental Psychology: Human learning and memory*, 6(5), 630-650.
- Nicholson, W.G., Verbruggen, F., & McLaren, I.P.L. (2018). What can Associative Learning do for Driving? In T.T. Rogers, M. Rau, X. Zhu, & C. W. Kalish (Eds.), *Proceedings of the 40th Annual Conference of the Cognitive Science Society* (pp. 2152 - 2158). Austin, TX: Cognitive Science Society.
- Nieuwenhuis, R., te Grotenhuis, H., & Pelzer, B. (2012). Influence. ME: Tools for detecting influential data in mixed effects models. *The R Journal*, 4(2), 38-47.
- Nigg, J. T. (2000). On inhibition/disinhibition in developmental psychopathology: Views from cognitive and personality psychology and a working inhibition taxonomy. *Psychological Bulletin*, 126(2), 220-246.
- Nigg, J. T. (2001). Is ADHD a disinhibitory disorder? *Psychological Bulletin*, 127(5), 571-598.
- Nigg, J. T. (2017). Annual Research Review: On the relations among self-regulation, self-control, executive functioning, effortful control, cognitive control, impulsivity, risk-taking, and inhibition for developmental psychopathology. *Journal of Child Psychology and Psychiatry*, 58(4), 361-383.
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1-32.
- Noack, H., Lövdén, M., Schmiedek, F., & Lindenberger, U. (2009). Cognitive plasticity in adulthood and old age: Gauging the generality of cognitive intervention effects. *Restorative Neurology and Neuroscience*, 27(5), 435-453.
- Noel, X., Brevers, D., Hanak, C., Kornreich, C., Verbanck, P., & Verbruggen, F. (2016). On the automaticity of response inhibition in individuals with alcoholism. *Journal of Behavior Therapy and Experimental Psychiatry*, 51, 84-91.
- Norman, D. A., & Shallice, T. (1986). Attention to action. In G. E. Schwartz, D. Shapiro & R. J. Davidson (Eds.), *Consciousness and Self-regulation* (pp. 1– 18). Boston, MA: Springer.

- O'Brien, F., & Gormley, M. (2013). The contribution of inhibitory deficits to dangerous driving among young people. *Accident Analysis & Prevention*, *51*, 238-242.
- Ouellette, J. A., & Wood, W. (1998). Habit and intention in everyday life: The multiple processes by which past behavior predicts future behavior. *Psychological Bulletin*, *124*(1), 54-74.
- Özyurt, J., Colonius, H., & Arndt, P. A. (2003). Countermanding saccades: Evidence against independent processing of go and stop signals. *Perception & Psychophysics*, *65*(3), 420-428.
- Palat, B., & Delhomme, P. (2016). A simulator study of factors influencing drivers' behavior at traffic lights. *Transportation Research Part F: Traffic Psychology and Behaviour*, *37*, 107-118.
- Palmer, J., Huk, A. C., & Shadlen, M. N. (2005). The effect of stimulus strength on the speed and accuracy of a perceptual decision. *Journal of Vision*, *5*(5), 376-404.
- Parkinson, J., & Haggard, P. (2014). Subliminal priming of intentional inhibition. *Cognition*, *130*(2), 255-265.
- Pascual-Leone, A., Valls-Solé, J., Wassermann, E. M., & Hallett, M. (1994). Responses to rapid-rate transcranial magnetic stimulation of the human motor cortex. *Brain*, *117*(4), 847-858.
- Pearce, J. M., & Bouton, M. E. (2001). Theories of associative learning in animals. *Annual Review of Psychology*, *52*(1), 111-139.
- Peck, R. C. (2011). Do driver training programs reduce crashes and traffic violations? — A critical examination of the literature. *IATSS Research*, *34*(2), 63-71.
- Pek, J., & Wu, H. (2015). Profile likelihood-based confidence intervals and regions for structural equation models. *Psychometrika*, *80*(4), 1123-1145.
- Perret, E. (1974). The left frontal lobe of man and the suppression of habitual responses in verbal categorical behaviour. *Neuropsychologia*, *12*(3), 323-330.
- Perruchet, P. (1985). Expectancy for airpuff and conditioned eyeblinks in humans. *Acta Psychologica*, *58*(1), 31-44.
- Perruchet, P. (2015). Dissociating conscious expectancies from automatic link formation in associative learning: A review on the so-called Perruchet

- effect. *Journal of Experimental Psychology: Animal Learning and Cognition*, 41(2), 105-127.
- Perruchet, P., & Amorim, M.-A. (1992). Conscious knowledge and changes in performance in sequence learning: Evidence against dissociation. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(4), 785-800.
- Perruchet, P., Cleeremans, A., & Destrebecqz, A. (2006). Dissociating the effects of automatic activation and explicit expectancy on reaction times in a simple associative learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(5), 955-965.
- Pinto, M., Cavallo, V., & Ohlmann, T. (2008). The development of driving simulators: Toward a multisensory solution. *Le travail humain*, 71(1), 62-95.
- Polders, E., Cornu, J., De Ceunynck, T., Daniels, S., Brijs, K., Brijs, T., . . . Wets, G. (2015). Drivers' behavioral responses to combined speed and red light cameras. *Accident Analysis & Prevention*, 81, 153-166.
- Poppelaars, A., Scholten, H., Granic, I., Veling, H., Johnson-Glenberg, M. C., & Luijten, M. (2018). When winning is losing: A randomized controlled trial testing a video game to train food-specific inhibitory control. *Appetite*, 129, 143-154.
- Porter, L., Bailey-Jones, C., Priudokaite, G., Allen, S., Wood, K., Stiles, K., . . . Lawrence, N. (2018). From cookies to carrots; the effect of inhibitory control training on children's snack selections. *Appetite*, 124, 111-123.
- Poulter, D. R., & McKenna, F. P. (2010). Evaluating the effectiveness of a road safety education intervention for pre-drivers: An application of the theory of planned behaviour. *British Journal of Educational Psychology*, 80(2), 163-181.
- Quine, L., Rutter, D. R., & Arnold, L. (2001). Persuading school-age cyclists to use safety helmets: Effectiveness of an intervention based on the Theory of Planned Behaviour. *British Journal of Health Psychology*, 6(4), 327-345.
- R Core Team. (2018). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>



- Ramautar, J., Kok, A., & Ridderinkhof, K. (2004). Effects of stop-signal probability in the stop-signal paradigm: The N2/P3 complex further validated. *Brain and Cognition*, *56*(2), 234-252.
- Reberg, D., & Leclerc, R. (1977). A feature positive effect in conditioned suppression. *Animal Learning & Behavior*, *5*(2), 143-147.
- Reed, J., & Johnson, P. (1994). Assessing implicit learning with indirect tests: Determining what is learned about sequence structure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(3), 585-594.
- Renner, W., & Anderle, F.-G. (2000). Venturesomeness and extraversion as correlates of juvenile drivers' traffic violations. *Accident Analysis & Prevention*, *32*(5), 673-678.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and non reinforcement. In A.H. Black & W.F. Prokasy (eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York, USA: Appleton-Century-Crofts.
- Retting, R. A., Williams, A. F., Preusser, D. F., & Weinstein, H. B. (1995). Classifying urban crashes for countermeasure development. *Accident Analysis & Prevention*, *27*(3), 283-294.
- Richardson, W. K., & Massel, H. K. (1982). The feature-positive effect in adult humans: Within-group design. *The American Journal of Psychology*, *95*(1), 125-138.
- Roberts, I., & Kwan, I. (2001). School based driver education for the prevention of traffic crashes. *Cochrane Systematic Review* (3), CD003201.
- Robertson, E. M. (2007). The Serial Reaction Time Task: Implicit Motor Skill Learning? *The Journal of Neuroscience*, *27*(38), 10073-10075.
- Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*(2), 207-231.
- Rose, J., & Colombo, M. (2005). Neural correlates of executive control in the avian brain. *PLoS Biology*, *3*(6), e190.
- Sakai, K. (2008). Task set and prefrontal cortex. *Annual Review of Neuroscience*, *31*, 219-245.

- Sani, S. R. H., Tabibi, Z., Fadardi, J. S., & Stavrinou, D. (2017). Aggression, emotional self-regulation, attentional bias, and cognitive inhibition predict risky driving behavior. *Accident Analysis & Prevention, 109*, 78-88.
- Satiennam, W., Satiennam, T., Triyabutra, T., & Rujopakarn, W. (2018). Red light running by young motorcyclists: Factors and beliefs influencing intentions and behavior. *Transportation Research Part F: Traffic Psychology and Behaviour, 55*, 234-245.
- Schachar, R., Logan, G. D., Robaey, P., Chen, S., Ickowicz, A., & Barr, C. (2007). Restraint and cancellation: Multiple inhibition deficits in attention deficit hyperactivity disorder. *Journal of Abnormal Child Psychology, 35*(2), 229-238.
- Schmidt, R., Leventhal, D. K., Mallet, N., Chen, F., & Berke, J. D. (2013). Canceling actions involves a race between basal ganglia pathways. *Nature Neuroscience, 16*(8), 1118-1124.
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychology Review, 84*(1), 1-66.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology, 80*(1), 1-27.
- Seabrooke, T., Wills, A. J., Hogarth, L., & Mitchell, C. J. (2019). Automaticity and cognitive control: Effects of cognitive load on cue-controlled reward choice. *Quarterly Journal of Experimental Psychology, 72*(6), 1507-1521.
- Sechenov, I. (1863). Etudes physiologiques sur les centres modérateurs des mouvements reflexes dans le cerveau de la grenouille. *Annales es Sciences Naturelles, ser. 4*(19), 109–134.
- Senderecka, M., Grabowska, A., Szewczyk, J., Gerc, K., & Chmylak, R. (2012). Response inhibition of children with ADHD in the stop-signal task: An event-related potential study. *International Journal of Psychophysiology, 85*(1), 93-105.
- Seth, A. K., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L. (2008). Measuring consciousness: Relating behavioural and neurophysiological approaches. *Trends in Cognitive Sciences, 12*(8), 314-321.
- Shallice, T., & Burgess, P. W. (1991). Deficits in strategy application following frontal lobe damage in man. *Brain, 114*(2), 727-741.

- Shanks, D. R. (2009). Learning: From Association to Cognition. *Annual Review of Psychology*, 61(1), 273-301.
- Shanks, D. R., & Darby, R. J. (1998). Feature-and rule-based generalization in human associative learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 24(4), 405-415.
- Shanks, D. R., & St. John, M. F. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences*, 17(3), 367-395.
- Sheeran, P., Gollwitzer, P. M., & Bargh, J. A. (2013). Nonconscious processes and health. *Health Psychology*, 32(5), 460-473.
- Sherman, J. W., Gawronski, B., & Trope, Y. (2014). *Dual-process theories of the social mind*. New York, N.Y., USA: Guilford Publications.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2), 127-190.
- Shinar, D., Meir, M., & Ben-Shoham, I. (1998). How automatic is manual gear shifting? *Human Factors*, 40(4), 647-654.
- Simons, D. J., Boot, W. R., Charness, N., Gathercole, S. E., Chabris, C. F., Hambrick, D. Z., & Stine-Morrow, E. A. (2016). Do "Brain-Training" Programs Work? *Psychological Science in the Public Interest*, 17(3), 103-186.
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143(2), 534-547.
- Smith, J. L., Dash, N. J., Johnstone, S. J., Houben, K., & Field, M. (2017). Current forms of inhibitory training produce no greater reduction in drinking than simple assessment: A preliminary study. *Drug and Alcohol Dependence*, 173, 47-58.
- Smith, R. (1992). *Inhibition: History and meaning in the sciences of mind and brain*. Berkeley, Los Angeles, USA: University of California Press.
- Smith, T. R., & Beran, M. J. (2018). Task switching in rhesus macaques (*Macaca mulatta*) and tufted capuchin monkeys (*Cebus apella*) during computerized categorization tasks. *Journal of Experimental Psychology: Animal Learning and Cognition*, 44(3), 229-246.

- Smithson, M., & Verkuilen, J. (2006). A better lemon squeezer? Maximum-likelihood regression with beta-distributed dependent variables. *Psychological Methods, 11*(1), 54-71.
- Smyth, S., Barnes-Holmes, D., & Barnes-Holmes, Y. (2008). Acquired equivalence in human discrimination learning: The role of propositional knowledge. *Journal of Experimental Psychology: Animal Behavior Processes, 34*(1), 167-177.
- Sniehotta, F. F., & Pesseau, J. (2011). The habitual use of the self-report habit index. *Annals of Behavioral Medicine, 43*(1), 139-140.
- Soetens, E., Boer, L. C., & Hueting, J. E. (1985). Expectancy or automatic facilitation? Separating sequential effects in two-choice reaction time. *Journal of Experimental Psychology: Human Perception and Performance, 11*(5), 598-616.
- Spiegel, R., & McLaren, I. P. L. (2001). Human Sequence Learning: Can Associations Explain Everything? In J. D. Moore, & K. S. Stenning (Eds.), *Proceedings of the 23rd Annual Conference of the Cognitive Science Society* (pp. 1004 - 1010). Austin, TX: Cognitive Science Society.
- Spiegel, R., & McLaren, I. P. L. (2006). Associative sequence learning in humans. *Journal of Experimental Psychology: Animal Behavior Processes, 32*(2), 156-163.
- Stacy, A. W., & Wiers, R. W. (2010). Implicit cognition and addiction: A tool for explaining paradoxical behavior. *Annual Review of Clinical Psychology, 6*, 551-575.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate? *Behavioral and Brain Sciences, 23*(5), 645-665.
- Stead, M., Tagg, S., MacKintosh, A. M., & Eadie, D. (2004). Development and evaluation of a mass media Theory of Planned Behaviour intervention to reduce speeding. *Health Education Research, 20*(1), 36-50.
- Steinmann, T. P., Andrew, C. M., Thomsen, C. E., Kjær, T. W., Meintjes, E. M., Molteno, C. D., . . . Sorensen, H. B. (2011). *An auditory Go/No-Go study of event-related potentials in children with fetal alcohol spectrum disorders*. Paper presented at the 2011 Annual International Conference of the IEEE Engineering in Medicine and Biology Society.

- Steinmetz, H., Knappstein, M., Ajzen, I., Schmidt, P., & Kabst, R. (2016). How effective are behavior change interventions based on the theory of planned behavior? A three-level meta-analysis. *Zeitschrift für Psychologie, 224*(3), 216-233.
- Stevens, T., Brevers, D., Chambers, C. D., Lavric, A., McLaren, I. P. L., Mertens, M., . . . Verbruggen, F. (2015). How does response inhibition influence decision making when gambling? *Journal of Experimental Psychology: Applied, 21*(1), 15-36.
- Stice, E., Lawrence, N. S., Kemps, E., & Veling, H. (2016). Training motor responses to food: A novel treatment for obesity targeting implicit processes. *Clinical Psychology Review, 49*, 16-27.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review, 8*(3), 220-247.
- Strack, F., & Deutsch, R. (2015). The duality of everyday life: Dual-process and dual system models in social psychology. In M. Mikulincer, P. R. Shaver, E. Borgida, & J. A. Bargh (Eds.), *APA handbooks in psychology. APA handbook of personality and social psychology, Vol. 1. Attitudes and social cognition* (pp. 891-927). Washington, DC, USA: American Psychological Association.
- Strickland, J. C., Hill, J. C., Stoops, W. W., & Rush, C. R. (2019). Feasibility, acceptability, and initial efficacy of delivering alcohol use cognitive interventions via crowdsourcing. *Alcoholism: Clinical and Experimental Research, 43*(5), 888-899.
- Stone III, W. B. (2017). *Psychometric evaluation of the Simulator Sickness Questionnaire as a Measure of cybersickness* (Doctoral dissertaion). Retrieved from: <https://lib.dr.iastate.edu/etd/15429>
- Sudevan, P., & Taylor, D. A. (1987). The cuing and priming of cognitive operations. *Journal of Experimental Psychology: Human Perception and Performance, 13*(1), 89-103.
- Sutherland, R. J., & Rudy, J. W. (1989). Configural association theory: The role of the hippocampal formation in learning, memory, and amnesia. *Psychobiology, 17*(2), 129-144.

- Swick, D., Ashley, V., & Turken, U. (2011). Are the neural correlates of stopping and not going identical? Quantitative meta-analysis of two response inhibition tasks. *Neuroimage*, *56*(3), 1655-1665.
- Tabibi, Z., Borzabadi, H. H., Stavrinou, D., & Mashhadi, A. (2015). Predicting aberrant driving behaviour: The role of executive function. *Transportation Research Part F: Traffic Psychology and Behaviour*, *34*, 18-28.
- Terry, D. J., & O'Leary, J. E. (1995). The theory of planned behaviour: The effects of perceived behavioural control and self-efficacy. *British Journal of Social Psychology*, *34*(2), 199-220.
- Thrifty. (2011). Amber gamblers: the truth about Britain's hazardous drivers. Retrieved from <http://www.thrifty.co.uk/thrifty/pdfs/news/021861311256771.pdf>
- Tolma, E. L., Reininger, B. M., Evans, A., & Ureda, J. (2006). Examining the theory of planned behavior and the construct of self-efficacy to predict mammography intention. *Health Education & Behavior*, *33*(2), 233-251.
- Tremblay, R. E., Pihl, R. O., Vitaro, F., & Dobkin, P. L. (1994). Predicting early onset of male antisocial behavior from preschool behavior. *Archives of General Psychiatry*, *51*(9), 732-739.
- Turton, R., Bruidegom, K., Cardi, V., Hirsch, C. R., & Treasure, J. (2016). Novel methods to help develop healthier eating habits for eating and weight disorders: A systematic review and meta-analysis. *Neuroscience & Biobehavioral Reviews*, *61*, 132-155.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, *5*(2), 207-232.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, *185*(4157), 1124-1131.
- U.S. Department of Transportation. (2015). *Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey*. (DOT HS 812 115). Washington, D.C., USA: National Highway Traffic Safety Administration.
- van der Schoot, M., Licht, R., Horsley, T. M., & Sergeant, J. A. (2005). Effects of stop signal modality, stop signal intensity and tracking method on inhibitory performance as determined by use of the stop signal paradigm. *Scandinavian Journal of Psychology*, *46*(4), 331-341.

- van Gaal, S., De Lange, F. P., & Cohen, M. X. (2012). The role of consciousness in cognitive control and decision making. *Frontiers in Human Neuroscience*, 6, 121.
- van Gaal, S., Lamme, V. A., Fahrenfort, J. J., & Ridderinkhof, K. R. (2011). Dissociable brain mechanisms underlying the conscious and unconscious control of behavior. *Journal of Cognitive Neuroscience*, 23(1), 91-105.
- van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *Journal of Neuroscience*, 28(32), 8053-8062.
- van Gaal, S., Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. A. (2010). Unconscious activation of the prefrontal no-go network. *Journal of Neuroscience*, 30(11), 4143-4150.
- van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P., & Lamme, V. A. (2009). Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered response inhibition in the stop-signal task. *Journal of Experimental Psychology: Human Perception and Performance*, 35(4), 1129-1139.
- Vandierendonck, A., Liefoghe, B., & Verbruggen, F. (2010). Task switching: Interplay of reconfiguration and interference control. *Psychological Bulletin*, 136(4), 601-626.
- Veling, H., Aarts, H., & Papies, E. K. (2011). Using stop signals to inhibit chronic dieters' responses toward palatable foods. *Behaviour Research and Therapy*, 49(11), 771-780.
- Veling, H., Aarts, H., & Stroebe, W. (2013a). Stop signals decrease choices for palatable foods through decreased food evaluation. *Frontiers in Psychology*, 4, 875.
- Veling, H., Aarts, H., & Stroebe, W. (2013b). Using stop signals to reduce impulsive choices for palatable unhealthy foods. *British Journal of Health Psychology*, 18(2), 354-368.
- Veling, H., Chen, Z., Huaiyu, L., Quandt, J., & Holland, R. W. (2019). Updating the p-curve analysis of Carbine & Larson with results from preregistered experiments. *Health Psychology Review*, 1-13.
- Veling, H., Lawrence, N. S., Chen, Z., van Koningsbruggen, G. M., & Holland, R. W. (2017). What is trained during food go/no-go training? A review

- focusing on mechanisms and a research agenda. *Current Addiction Reports*, 4(1), 35-41.
- Veling, H., van Koningsbruggen, G. M., Aarts, H., & Stroebe, W. (2014). Targeting impulsive processes of eating behavior via the internet. Effects on body weight. *Appetite*, 78, 102-109.
- Verbruggen, F., Adams, R., & Chambers, C. D. (2012). Proactive motor control reduces monetary risk taking in gambling. *Psychological Science*, 23(7), 805-815.
- Verbruggen, F., Adams, R. C., van't Wout, F., Stevens, T., McLaren, I. P.L., & Chambers, C. D. (2013). Are the effects of response inhibition on gambling long-lasting? *PLoS One*, 8(7), e70155.
- Verbruggen, F., Best, M., Bowditch, W. A., Stevens, T., & McLaren, I.P.L. (2014). The inhibitory control reflex. *Neuropsychologia*, 65, 263-278.
- Verbruggen, F., & De Houwer, J. (2007). Do emotional stimuli interfere with response inhibition? Evidence from the stop signal paradigm. *Cognition and Emotion*, 21(2), 391-403.
- Verbruggen, F., & Logan, G. D. (2008a). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, 12(11), 418-424.
- Verbruggen, F., & Logan, G. D. (2008b). Automatic and controlled response inhibition: Associative learning in the go/no-go and stop-signal paradigms. *Journal of Experimental Psychology: General*, 137(4), 649-672.
- Verbruggen, F., & Logan, G. D. (2008c). After-effects of goal shifting and response inhibition: A comparison of the stop-change and dual-task paradigms. *Quarterly Journal of Experimental Psychology*, 61(8), 1151-1159.
- Verbruggen, F., & Logan, G. D. (2009a). Models of Response Inhibition in the Stop-Signal and Stop-Change Paradigms. *Neuroscience and Biobehavioral Reviews*, 33(5), 647-661.
- Verbruggen, F., & Logan, G. D. (2009b). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3), 835-854.
- Verbruggen, F., & Logan, G. D. (2009c). Automaticity of cognitive control: Goal priming in response-inhibition paradigms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35(5), 1381-1388.



- Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the control homunculi in studies of action control and behavior change. *Perspectives on Psychological Science, 9*(5), 497-524.
- Verbruggen, F., Schneider, D. W., & Logan, G. D. (2008). How to stop and change a response: The role of goal activation in multitasking. *Journal of Experimental Psychology: Human Perception and Performance, 34*(5), 1212-1228.
- Verbruggen, F., Stevens, T., & Chambers, C. D. (2014). Proactive and reactive stopping when distracted: An attentional account. *Journal of Experimental Psychology: Human Perception and Performance, 40*(4), 1295-1300.
- Verplanken, B., Aarts, H., Knippenberg, A., & Moonen, A. (1998). Habit versus planned behaviour: A field experiment. *British Journal of Social Psychology, 37*(1), 111-128.
- Verplanken, B., & Orbell, S. (2003). Reflections on past behavior: A self-report index of habit strength. *Journal of Applied Social Psychology, 33*(6), 1313-1330.
- Verplanken, B., & Wood, W. (2006). Interventions to break and create consumer habits. *Journal of Public Policy & Marketing, 25*(1), 90-103.
- Wadley, V. G., Okonkwo, O., Crowe, M., Vance, D. E., Elgin, J. M., Ball, K. K., & Owsley, C. (2009). Mild cognitive impairment and everyday function: An investigation of driving performance. *Journal of Geriatric Psychiatry and Neurology, 22*(2), 87-94.
- Walshe, E., Ward McIntosh, C., Romer, D., & Winston, F. (2017). Executive function capacities, negative driving behavior and crashes in young drivers. *International Journal of Environmental Research and Public Health, 14*(11), 1314-1330.
- Watson, L., & Watson, P. (2014). *Juvenal: Satire 6*. Cambridge, United Kingdom: Cambridge University Press.
- Weidemann, G., Broderick, J., Lovibond, P. F., & Mitchell, C. J. (2012). Both trace and delay conditioned eyeblink responding can be dissociated from outcome expectancy. *Journal of Experimental Psychology: Animal Behavior Processes, 38*(1), 1-10.
- Weidemann, G., McAndrew, A., Livesey, E. J., & McLaren, I. P. L. (2016). Evidence for multiple processes contributing to the Perruchet effect:

- Response priming and associative learning. *Journal of Experimental Psychology: Animal Learning and Cognition*, 42(4), 366-379.
- Weidemann, G., Tangen, J. M., Lovibond, P. F., & Mitchell, C. J. (2009). Is Perruchet's dissociation between eyeblink conditioned responding and outcome expectancy evidence for two learning systems? *Journal of Experimental Psychology: Animal Behavior Processes*, 35(2), 169-176.
- Weinstein, N. D., & Klein, W. M. (1996). Unrealistic optimism: Present and future. *Journal of Social and Clinical Psychology*, 15(1), 1-8.
- Whitely, P. L., & Blankfort, G. (1933). The influence of certain prior conditions upon learning. *Journal of Experimental Psychology*, 16(6), 843-853.
- Whiten, A., Custance, D. M., Gomez, J.-C., Teixidor, P., & Bard, K. A. (1996). Imitative learning of artificial fruit processing in children (*Homo sapiens*) and chimpanzees (*Pan troglodytes*). *Journal of Comparative Psychology*, 110(1), 3-14.
- Willingham, D. B., Greeley, T., & Bardone, A. M. (1993). Dissociation in a serial response time task using a recognition measure: Comment on Perruchet and Amorim (1992). *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 19(6), 1424-1430.
- Willingham, D. B., Nissen, M. J., & Bullemer, P. (1989). On the development of procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 15(6), 1047-1060.
- Wills, A. J., Graham, S., Koh, Z., McLaren, I. P. L., & Rolland, M. D. (2011). Effects of concurrent load on feature-and rule-based generalization in human contingency learning. *Journal of Experimental Psychology: Animal Behavior Processes*, 37(3), 308-316.
- Wood, W., & Neal, D. T. (2007). A new look at habits and the habit-goal interface. *Psychological Review*, 114(4), 843-863.
- Wood, W., Quinn, J. M., & Kashy, D. A. (2002). Habits in everyday life: Thought, emotion, and action. *Journal of Personality and Social Psychology*, 83(6), 1281-1297.
- Wood, W., & Runger, D. (2016). Psychology of habit. *Annual Review of Psychology*, 67, 289-314.
- Yeates, F., Jones, F. W., Wills, A. J., Aitken, M., & McLaren, I. P. L. (2013). Implicit learning: A demonstration and a revision to a novel SRT paradigm. In M. Knauff, N. Sebanz, M. Pauen, & I. Wachsmuth (Eds.),

- Proceedings of the 35th Annual Conference of the Cognitive Science Society* (pp. 3829 - 3835). Austin, TX: Cognitive Science Society.
- Yeates, F., Jones, F. W., Wills, A. J., McLaren, R. P., & McLaren, I. P. L. (2013). Modeling human sequence learning under incidental conditions. *Journal of Experimental Psychology: Animal Behavior Processes*, 39(2), 166-173.
- Zandbelt, B. B., Bloemendaal, M., Neggers, S. F., Kahn, R. S., & Vink, M. (2013). Expectations and violations: Delineating the neural network of proactive inhibitory control. *Human Brain Mapping*, 34(9), 2015-2024.
- Zuckerman, M., & Aluja, A. (2014). Measures of Sensation Seeking. In G.J. Boyle, D.H. Saklofske, & G. Matthews (Eds.), *Measures of Personality and Social Psychological Constructs* (pp. 352-380). USA: Academic Press.
- Zuckerman, M., Kuhlman, D. M., Joireman, J., Teta, P., & Kraft, M. (1993). A comparison of three structural models for personality: The Big Three, the Big Five, and the Alternative Five. *Journal of Personality and Social Psychology*, 65(4), 757-768.

