TOP-DOWN AND BOTTOM-UP INFLUENCES ON RESPONSE INHIBITION

Submitted by Maisy Jane Best to the University of Exeter
as a thesis for the degree of
Doctor of Philosophy in Psychology
In September 2016

I certify that all material in this thesis which is not my own work has been identified and that no material has previously been submitted and approved for the award of a degree by this or any other University.

Signature:...

ABSTRACT

Following exposure to consistent stimulus–stop mappings, response inhibition can become automatised with practice. What is learned is less clear, even though this has important theoretical and practical implications. The main contribution of this thesis is to investigate how stimulus-stop associations are acquired and the conditions under which they influence behaviour. To this end, this thesis addressed several outstanding issues concerning the associative architecture of stop learning, the role of expectancies, and the specificity of learning in inhibition tasks. Experiments 1-4 provide evidence that participants can acquire direct associations between specific stimuli and the stop goal without mediation via a single representation of the stop signal. However, these experiments also suggest that the influence of stimulus-stop associations on behaviour depends on top-down attentional settings: if participants begin to ignore the stop-associated stimuli, the effects of stop learning are diminished or eliminated entirely. Across eight experiments, this thesis provides evidence that participants generate expectancies during stop learning that are consistent with the stimulus-stop contingencies in play. However, Experiments 5-6 indicate that there may be some differences in the relationships between stimulus-stop expectancies and task performance under instructed and uninstructed conditions; stimulus-stop associations that are acquired via task instructions or via task practice have similar effects on behaviour, but seem to differ in how they trigger response slowing for the stop-associated items. Experiments 7-8 investigated the role of signal detection processes during the acquisition of stimulus-stop associations. To distinguish between stimulus-stop learning and stimulus-signal learning, the contingencies between specific stimuli and the stop goal and the contingencies between specific stimuli and the spatial location of the stop signal were independently manipulated. Although these experiments showed evidence of stop/go (goal) learning, there was no evidence that participants acquired the stimulus-signal associations. Across four experiments, this thesis investigated the specificity of stop learning. Experiments 9-10 compared the effects of training on behavioural performance in inhibition (go/no-go) and non-inhibition (two-choice) tasks. The results of these experiments revealed that learning in inhibition and non-inhibition tasks could arise through similar associative mechanisms, but suggest that the effects of training in these tasks could also depend on top-down response settings and general non-associative processes. Experiments 11-12 investigated the neural specificity of stop learning. These experiments also revealed similar effects of training across the go/no-go and two-choice tasks adding weight to the claim that training in inhibition tasks primarily influences task-general processes. Combined, the overall conclusion of this thesis is that bottom-up control can influence response inhibition but what is learned depends on top-down factors. It is therefore important to consider bottom-up factors and top-down factors as dependent, rather than independent, influences on response inhibition.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my first supervisor Frederick Verbruggen. Not only is Frederick a great scientist, he is also a great supervisor. Throughout my PhD, he has always been available to provide invaluable expertise, encouragement, and advice whenever needed.

I would also like to thank my second supervisor, Fraser Milton, for his helpful comments during the preparation of this thesis. Thanks also to the other members of the cognition group at Exeter for providing an engaging and inspiring research environment in which to complete a PhD. In particular, I would like to acknowledge the contribution of Heike Elchlepp who provided useful suggestions and comments on the event-related potential experiments presented in Chapter 6. I learnt a great amount from my conversations with Heike.

During the second year of my PhD, I spent six highly enjoyable and informative weeks in Geoff Woodman's electrophysiology lab at Vanderbilt University. Thanks to Geoff and the members of his lab for making me feel so welcome.

Special thanks to the past and present 'residents' of 230: Fayme Yeates, Amy McAndrew, Will Bowditch, Kathryn Carpenter, Lorna Hardy, and Lucy Porter. I couldn't have wished for a better group of friends and colleagues to share this journey with. I can't wait to see where our lives take us next.

Last but not least, I would like to thank my family for reminding me that there is more to life than running 'yet another experiment'. And finally, to my wonderful Desmond for his unconditional love. I couldn't have done this without you.

The research presented in this thesis was supported by a studentship from the Economic and Social Research Council (ES/J50015X/1).

TABLE OF CONTENTS

Abstract 2

Acknowledgements 3

Table of contents 4

List of tables 11

List of figures 14

List of appendices 16

Author’s declaration 17

Chapter 1: General Introduction 19

- The ‘Homunculus Problem’ 19
- What is ‘Response Inhibition’? 20
- Response Inhibition in the Laboratory 21
  - Independent Horse-Race Model 22
  - Variants of the Independent Horse-Race Model 23
- The Neural Mechanisms of Response Inhibition 25
- Top-Down and Bottom-Up Forms of Response Inhibition 27
  - Response Inhibition as a ‘Prepared Reflex’ 27
  - Response Inhibition as a ‘Learned Reflex’ 29
    - Sequential after-effects 30
    - Response Inhibition as an ‘Automatic Act of Control’ 31
- What is ‘Automaticity’? 33
- How do Stimulus-Stop Associations Influence Responding? 34
  - Theoretical Explanations 34
  - Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding? 34
  - The Neural Mechanisms of Stimulus-Stop Learning 35
- Can ‘Top-down’ Forms of Response Inhibition be Trained? 37
- Practical Applications 38
- Main Outstanding Issues 39
  - What is Learned? 39
  - The Role of Expectancies 40
  - The Specificity of Stop Learning 41
- Thesis Outline 42

Chapter 2: Associations and Expectancies in Response Inhibition Tasks 45
- Are Associations Between Stimuli and Stopping Direct? 45
- What is the Role of Expectancy in Stimulus-Stop Learning? 47
- Experiment 1 48
  - Method 48
    - Participants 48
    - Apparatus and stimuli 48
    - Procedure 49
    - Analyses 51
    - Training phase 52
    - Test phase 53
    - Expectancy ratings 55
  - Interim Discussion 57
- Experiment 2 60
  - Method 60
    - Participants 60
    - Apparatus, stimuli, procedure, and analyses 60
  - Results 60
    - Training phase 60
    - Test phase 61
    - Expectancy ratings 61
  - Interim Discussion 62
- Experiment 3 63
  - Method 63
    - Participants 63
    - Apparatus, stimuli, procedure, and analyses 64
  - Results 64
    - Training phase 64
    - Test phase 64
    - Expectancy ratings 64
  - Interim Discussion 65
- Experiment 4 66
  - Method 67

Participants 67
Apparatus, stimuli, procedure, and analyses 67
Results 67
Training phase 67
Test phase 70
Interim Discussion 71
Chapter Discussion 73
Are Associations Between Stimuli and Stopping Direct? 73
What is the Role of Expectancy in Stimulus-Stop Learning? 74
Does Attention to the Stop Items Affect the Extent to Which Stimulus-Stop Learning Influences Behaviour? 76
Conclusion 77
Appendix A 78
Appendix B 79
Appendix C 79
Chapter 3: The Effects of Instructed and Acquired Contingencies in Response Inhibition Tasks 89
The Present Chapter 90
Experiments 93
Method 94
Participants 94
Apparatus and stimuli 95
Procedure 95
Analyses 96
Results 97
No-signal analyses 97
No-signal RTs 97
Probability of a missed go response 103
Probability of a correct go response 104
Signal analyses 106
Expectancy analyses 108
Chapter Discussion 112
Appendix D 118
Appendix E 135

Appendix F 136
Appendix G 147
Chapter 4: The Role of Signal Detection in Stop/Go Learning 150
The Present Chapter 153
Experiment 7 153
Method 154
Participants. 154
Apparatus, stimuli, and procedure 154
Analyses. 157
Results 158
Go analyses. 158
Stop analyses. 160
Expectancy ratings analyses. 160
Interim Discussion 161
Experiment 8 162
Method 163
Participants. 163
Apparatus, stimuli, and procedure. 163
Analyses. 167
Results 168
Go analyses. 168
Stop analyses. 170
Eye movement analyses 171
Word interval. 171
Signal interval 173
Expectancy ratings analyses. 177
Chapter Discussion 178
Conclusion 181
Appendix H 182
Appendix I 183
Chapter 5: Behavioural Practice-Effects in Inhibition and Non-Inhibition tasks 184
The Present Chapter 187
Experiment 9 188

      Participants. 188
      Apparatus, stimuli, and procedure. 189
      Analyses. 190
    Results and Discussion 191
      Training phase. 191
      Test phase. 193
      Training vs. test phase. 196
  Experiment 10 198
      Participants 199
      Stimuli, apparatus, procedure, and analyses. 199
    Results and Discussion 200
      Training phase. 200
      Test phase. 203
      Training vs. test phase. 206
  Chapter Discussion 208
    The Role of Non-Associative Processes 208
  Conclusion 212
  Appendix J 213
Chapter 6: The Electrophysiology of Practice-Effects in Inhibition and Non-Inhibition Tasks 215
  The Present Chapter 219
  Experiment 11 220
    Method 221
      Participants 221
      Apparatus, stimuli, and procedure. 221
      EEG data acquisition 222
      Analyses 223
        Behavioural analyses 224
        ERP analyses 224
    Results 227
      Behavioural Results 227
      ERP Results 228
        N2 range 229
          Mean amplitudes. 229

Peak latencies. 229
P3a range. 232
Mean amplitudes. 232
Peak latencies. 232
P3b range. 233
Mean amplitudes. 233
Peak latencies. 234
Response-locked analyses. 236
Interim Conclusions 237
Experiment 12 239
Method 239
Participants 240
Stimuli, apparatus, and procedures. 240
Analyses. 240
Results 240
Behavioural Results 240
ERP results 241
N2 range 241
Mean amplitudes 241
Peak latencies 244
P3a range 245
Mean amplitudes 245
Peak latencies 245
P3b range 246
Mean amplitudes 246
Peak latencies 247
Response-locked analyses 249
Chapter Discussion 249
The N2 Component 251
The P3 Component(s) 252
The frontal P3a component 252
The parietal P3b component 253
Conclusion 254

Appendix K 256

Appendix L 269

Chapter 7: General Discussion 272

What is Learned? 272

Can Participants Acquire Direct (Stimulus-Stop) Associations? 272

Can Participants Acquire Indirect (Stimulus-Signal) Associations? 272

Is Stop Learning Special? 272

From Associations to Behaviour 278

How do Stimulus-Stop Associations Influence Behaviour? 278

What Role Do Expectancies Play In Stop Learning? 278

Top-Down and Bottom-Up Influences on Response Inhibition 278

Stimulus Detection 282

Action Selection 284

Implications for Applied Research 287

Avenues for Future Research 288

Alternative Routes 288

Mediation via a Category Representation 289

Transfer to a New Stimulus Set 291

Concluding Remarks 292

Appendix M 293

References 296

LIST OF TABLES

Table 2.1. Proportion of stop-signal trials 50
Table 2.2. Probability of a missed go response 52
Table 2.3. Overview of ANOVA performed to compare go and stop training phase performance 55
Table 2.4. Overview of ANOVA to compare go and stop test phase performance 57
Table 2.5. Go reaction times (in ms) in the test phase 71
Table 3.1. Overview of the no-signal data 99
Table 3.2. Overview of ANOVA on the no-signal data 102
Table 3.3. No-signal RT Bayesian analysis 104
Table 3.4. Overview of ANOVA on the no-signal data in the instructed condition only 105
Table 3.5. Overview of ANOVA on the no-signal data in the uninstructed condition only 106
Table 3.6. Overview of the relevant stop-signal data 107
Table 3.7. Overview of ANOVA on the stop-signal data 108
Table D.1. Overview of the no-signal data in Experiment 5 119
Table D.2. Overview of the relevant stop-signal data in Experiment 5 120
Table D.3. Overview of ANOVA on the no-signal data in Experiment 5 121
Table D.4. No-signal RT Bayesian analysis in Experiment 5 122
Table D.5. Overview of ANOVA on the stop-signal data in Experiment 5 123
Table D.6. Overview of ANOVA on the instructed condition only in Experiment 5 124
Table D.7. Overview of ANOVA on the uninstructed condition only in Experiment 5 125
Table D.8. Overview of the no-signal data in Experiment 6 127
Table D.9. Overview of the relevant stop-signal data in Experiment 6 128
Table D.10. Overview of ANOVA on the no-signal data in Experiment 6 129
Table D.11. No-signal RT Bayesian analysis in Experiment 6 130
Table D.12. Overview of ANOVA on the stop-signal data in Experiment 6 131
Table D.13. Overview of ANOVA on the instructed condition only in Experiment 6 132
Table D.14. Overview of ANOVA on the uninstructed condition only in Experiment 6 133
Table F.1. Overview of the number of fixations and average distance in the word interval 139
Table F.2. Overview of the number of fixations and average distance in the stop-signal interval 140
Table F.3. Overview of the proportion of the fixations in the word interval 141
Table F.4. Overview of the proportion of the fixations in the stop-signal interval 142

Table F.5. Overview of ANOVA performed to compare number of fixations and the average distance in the word interval 143

Table F.6. Overview of ANOVA performed to compare number of fixations and the average distance in the stop-signal interval 144

Table F.7. Overview of ANOVA performed to compare number of fixations and the average distance in the instructed condition only 145

Table F.8. Overview of ANOVA performed to compare number of fixations and the average distance in the uninstructed condition only 146

Table G.1. Overview of the go data 148

Table G.2. Overview of ANOVA 149

Table G.3. No-signal RT Bayes analysis 149

Table 4.1. The proportion of cue-types in Experiment 7 156

Table 4.2. Overview of ANOVA in Experiment 7 161

Table 4.3. The proportion of cue-types in Experiment 8 166

Table 4.4. Overview of ANOVA in Experiment 8 170

Table 4.5. Overview of the number of fixations and average distance in Experiment 8 171

Table 4.6. Overview of ANOVA performed to compare number of fixations and the average distance in the word interval 172

Table 4.7. Overview of the number of fixations and average distance in the signal interval 174

Table 4.8. Overview of ANOVA performed to compare number of fixations and the average distance in the signal interval 175

Table 4.9. Overview of ANOVA performed to compare number of fixations and the average distance for the stop-associated cues 176

Table 4.10. Overview of ANOVA performed to compare number of fixations and the average distance for the go-associated cues 177

Table 5.1. Overview of the analyses of variance on the training phase data in Experiment 9 192

Table 5.2. Overview of the ANOVA on the test phase data in Experiment 9 195

Table 5.3. Bayesian analysis for the test phase data in Experiment 9 196

Table 5.4. Bayesian analyses for the difference between performance at the end of training vs. in the test phase in Experiment 9 198

Table 5.5. Overview of the ANOVA on the training phase data in Experiment 10 202

Table 5.6. Overview of the ANOVA on the test phase data in Experiment 10 204

Table 5.7. Bayesian analysis for the test phase data in Experiment 10 205

Table 6.1. Overview of go/no-go training studies with behavioural and N2 and/or P3 outcome measures. 218
Table 6.2. Overview of the average number of artifact-free ERP segments 225
Table 6.3. Overview of the behavioural data in Experiment 11 and Experiment 12. 228
Table 6.4. Overview of the ANOVA of the reaction time data in Experiments 11 and 12 228
Table 6.5. Overview of the ANOVA of the stimulus-locked mean amplitudes in Experiment 11 230
Table 6.6. Overview of the ANOVA of the stimulus-locked peak latencies in Experiment 11 235
Table 6.7. Overview of the ANOVA of the response-locked amplitudes in Experiment 11 237
Table 6.8. Overview of the ANOVA of the stimulus-locked mean amplitudes in Experiment 12 242
Table 6.9. Overview of the ANOVA of the stimulus-locked peak latencies in Experiment 12 244
Table 6.10. Overview of the ANOVA of the response-locked amplitudes in Experiment 12 249
Table K.1. Overview of the average number of artifact-free ERP segments 257
Table K.2. Overview of the behavioural data in the between-subjects EEG pilot experiment 258
Table K.3. Overview of the ANOVA of the reaction time data 259
Table K.4. Overview of the ANOVA of the training phase mean amplitudes 260
Table K.5. Overview of the ANOVA of the training phase stimulus-locked peak latencies 262
Table K.6. Overview of the ANOVA of the test phase mean average amplitudes 267
Table K.7. Overview of the ANOVA of the test phase stimulus-locked peak latencies 268
Table K.8. Overview of the ANOVA on the response-locked average amplitudes 268
Table L.1. Overview of the behavioural data of the control experiment 270
Table L.2. Overview of the ANOVA in the control experiment 270
Table L.3. Bayesian ANOVA for the reaction time data in the control experiment 271

LIST OF FIGURES

Figure 1.1. Trial structure in the go/no-go and stop-signal tasks 22
Figure 1.2. (a) Visual depiction of the independent race model (b) Inhibition function 25
Figure 1.3. Architecture of the associative stop system .... -PAGE-
Figure 2.1. Example go/stop trial sequence 51
Figure 2.2. Reaction times for the correct go trials in Experiment 1 53
Figure 2.3. Reaction times for the correct go trials and *p*(respond|stop) in Experiment 2 62
Figure 2.4. Reaction times for the correct go trials and *p*(respond|stop) in Experiment 3 65
Figure 2.5. Reaction times for the correct go trials, *p*(respond|stop) and expectancy ratings in Experiment 4 .... -PAGE-
Figure B.1. Go RT distributions in Experiment 1 80
Figure B.2. Go RT distributions in Experiment 2 81
Figure B.3. Go RT distributions in Experiment 3 82
Figure B.4. Go RT distribution in Experiment 4 83
Figure C.1. Expectancy/RT correlations in Experiment 1 84
Figure C.2. Expectancy/RT correlations in Experiment 2 85
Figure C.3. Expectancy/RT correlations in Experiment 3 86
Figure C.4. Expectancy/RT correlations in Experiment 4 87
Figure 3.1. The distractor stop task 93
Figure 3.2. Behavioural data in Experiments 5-6 101
Figure 3.3. Expectancy/RT correlations 111
Figure F.1. Analysed eye movement regions 137
Figure 4.1. Trial structures in Experiment 7 155
Figure 4.2. Go performance data in Experiment 7 .... -PAGE-
Figure 4.3. Stop performance data in Experiment 7 161
Figure 4.4. Trial structures in Experiment 8 165
Figure 4.5. Go and stop data in Experiment 8 169
Figure 5.1. Stimuli and trial structure in Experiments 9-10 188
Figure 5.2. Reaction times in Experiment 9 192
Figure 5.3. The probability of correct responses in Experiment 9 .... -PAGE-
Figure 5.4. The probability of responding on no-go trials in Experiment 9 194
Figure 5.5. Reaction times in Experiment 10 201

Figure 5.6. The probability of correct responses in Experiment 10 203
Figure 5.7. The probability of responding on no-go trials in Experiment 10 206
Figure J.1. RT distributions for Experiments 9-10 214
Figure 6.1. Stimulus sets used in (a) Experiment 11 and (b) Experiment 12 222
Figure 6.2. Electrode arrangement 223
Figure 6.3. Stimulus-locked waveforms in Experiment 11 231
Figure 6.4. Topographies for the beginning minus end difference in Experiment 11 233
Figure 6.5. Mean average amplitudes and peak latencies in Experiment 11 237
Figure 6.6. Lateralised readiness potentials in Experiment 11 237
Figure 6.7. Stimulus-locked waveforms in Experiment 12 243
Figure 6.8. Topographies for the beginning minus end difference in Experiment 12. 246
Figure 6.9. Mean average amplitudes and peak latencies in Experiment 12 248
Figure 6.10. Lateralised readiness potentials in Experiment 12 249
Figure K.1. Stimulus-locked waveforms in the ERP pilot experiment 261
Figure K.2. Topographies in the ERP pilot experiment. 263
Figure K.3. Mean average amplitudes in the ERP pilot experiment 264
Figure K.4. Peak latencies in the ERP pilot experiment 265
Figure K.5. Lateralised readiness potentials in the ERP pilot experiment 266
Figure 7.1. An visual depiction of top-down control and bottom-up learning 286
Figure M.1. Go RT distributions in the instructed condition of Experiments 5-6 294
Figure M.2. Go RT distributions in the uninstructed condition of Experiments 5-6 294
Figure M.3. Go RT distributions in Experiment 8 295

LIST OF APPENDICES

Appendix A. A full list of the rules and the stimuli used in Experiments 1-4 78
Appendix B. RT Percentiles 79
Appendix C. Expectancy/RT correlation plots 84
Appendix D. Separate Analyses of Experiment 5 and Experiment 6 118
Appendix E. Go Stimuli 135
Appendix F. Analyses of the Eye Data 136
Appendix G. Pilot Experiment 147
Appendix H. Words used in Experiment 7 182
Appendix I. Words used in Experiment 8 183
Appendix J. RT Distributions 213
Appendix K. Between-Subjects Pilot ERP Experiment.. 256
Appendix L. Behavioural Control Experiment... 269
Appendix M. RT Percentiles 293

AUTHOR'S DECLARATION

The research reported in this thesis was carried out at the University of Exeter at the University of Exeter between October 2013 and July 2016, and was supervised by Prof. Frederick Verbruggen and Dr Fraser Milton.

This thesis has not been submitted, in whole or in part, for any other degree, diploma or qualification at any university. Several of the experiments reported in this thesis have been accepted or are being prepared for publication. Chapter 2 is published in Journal of Experimental Psychology: Human Perception and Performance by Best, M., Lawrence, N.S., Logan, G.D., McLaren, I.P.L., and Verbruggen, F. Chapter 3 has been prepared for submission by Best, M., McLaren, I. P. L., and Verbruggen, F. In both cases, conducted the experiments and analyses, wrote the first draft and prepared the figures and tables. My co-authors edited the manuscript.

Maisy Best
Exeter, August 2016.

CHAPTER 1

GENERAL INTRODUCTION

Modern day life is awash with a multitude of choices, options, and competing courses of action. As such, the ability to behave flexibly in response to a rapidly changing environment is one of the most powerful commodities available to us. In cognitive psychology, flexible and goal-driven behaviour is often attributed to an 'executive' or 'cognitive' control system (e.g. Logan & Gordon, 2001; Miller & Cohen, 2001; Monsell & Driver, 2000; Norman & Shallice, 1986). The term 'executive control' is typically used to refer to 'the functions of the cognitive system that allow people to regulate their behaviour according to higher-order goals or plans' (Verbruggen, McLaren, & Chambers, 2014, p. 497). On a daily basis, we rely on executive control to rapidly configure actions when required, resist temptations, and overcome habits (e.g. Verbruggen, McLaren, et al., 2014). Indeed, impaired executive control is linked to a variety of psychopathological and impulse-control disorders, such as attention deficit/hyperactivity disorder, obsessive-compulsive disorder, substance abuse, pathological gambling, and eating disorders (Bechara, Noël & Crone, 2006; Crews & Boettiger, 2009; de Wit, 2009; Fernie, Peeters, et al., 2013; Garavan & Stout, 2005; Nigg, 2001; Noël, Brevers & Bechara, 2013). Taken together, it would not be an overstatement to say that we would be 'doomed' without executive control.

The 'Homunculus Problem'

Although executive control has been the subject of extensive research over the past century, many theories still attribute executive control to an ill-defined set of 'executive' control functions[1] or, even worse, an intelligent agent (a 'homunculus') without explaining how control is achieved (Verbruggen, McLaren, et al., 2014). This is problematic because the outcomes of cognitive control (e.g. to inhibit a prepotent response) are too readily confused with the mechanisms through which control is achieved (e.g. 'response inhibition'). To address the 'homunculus problem', Verbruggen and colleagues recently suggested that one of the first steps in 'banishing the control homunculus' from theories of executive control is to consider how control is achieved across different time-scales.

When researchers describe the time-scale of executive control processes, they commonly draw a contrast with 'automatic' processes (e.g. Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). Research has often characterised control processes as dichotomous to 'automatic' processes

---

[1] Correlational research has shown that there are at least three executive control functions: switching between tasks ('shifting'); updating representations stored in working memory ('updating'); and suppressing irrelevant information and stopping irrelevant or inappropriate actions ('inhibition'; Miyake, Friedman, et al., 2000).

(e.g. Miller & Cohen, 2001; Norman & Shallice, 1986). 'Executive' processes are considered to be slow, effortful, rule-based, intentional, and 'top-down' whereas 'automatic' processes are considered to be fast, associative, effortless, and 'bottom-up'[2]. Even beyond the action control literature, the characterisation of controlled and automatic processes as discrete influences on behaviour is common[3]. However, in recent years, it has been suggested that executive and automatic processes may be better characterised as a continuum in which executive control begins as 'top-down', deliberate, intentional, and goal-driven but becomes more 'bottom-up' with training (e.g. Verbruggen & Logan, 2008a; Verbruggen, McLaren, et al., 2014). Indeed, Verbruggen and colleagues argue that by not acknowledging the bottom-up factors that influence control processes 'we generate an inherently limited perspective on the cognitive mechanisms behind action control' (Verbruggen, McLaren, et al., 2014, p. 498).

To this end, the research presented in this thesis contributes to the action control literature by considering the interplay between top-down and bottom-up influences on control. In particular, this thesis considers the bottom-up influences on a function of the control system that is often considered to a hallmark of 'top-down' executive control; 'response inhibition' (e.g. Miyake, Friedman, et al., 2000).

What is 'Response Inhibition'?

The term 'inhibition' has been linked to a variety of phenomena across several research domains (for an overview, see Aron, 2007; Bari & Robbins, 2013). The concept of 'inhibition' is commonly used in psychology to describe suppression of memories, actions, and attention. However, the exact role of inhibition in several paradigms is debated (see e.g. MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003). The main reason for this debate is that inhibition at the behavioural level mostly cannot be clearly operationalised or observed (Aron, 2007). Nevertheless, there is one exception; most researchers accept that some kind of inhibition at the behavioural level is required to stop a motor response (see MacLeod, Dodd, Sheard, Wilson, & Bibi, 2003). In this case, 'response inhibition' refers to the ability to quickly stop and replace responses that are no longer relevant or that are inappropriate in the current task environment (e.g. Verbruggen & Logan, 2008b). Note, that although the term 'inhibition' is frequently used in the neuroscience domain, there is no clear relationship between 'inhibition' at the cellular level and the kind of behavioural inhibition that is the subject of research presented in this

---

2. Throughout this thesis, 'top-down control' refers to control processes that are 'guided by internal states or intentions' (Miller & Cohen, 2001, p. 168) whereas 'bottom-up control' refers to control processes that are guided by lower-level processes in the absence of goals and/or intentions.
3. For example, several behaviour change interventions are based on the idea that automatic and controlled (intentional) processes reflect discrete or dichotomous influences on behaviour (for an overview, see Webb & Sheeran, 2006).

thesis (see, e.g. Aron, 2007; Bari & Robbins, 2013)[4].

Response inhibition can occur 'reactively' following the detection of an external 'stop signal' and 'proactively' in advance of stop-signal detection. The majority of research has focussed on the mechanisms of reactive inhibition, but more recently it has been shown that when people expect a stop signal in the near future, they preemptively adjust attentional settings, increase response thresholds, or proactively suppress all motor output to enhance detection of the stop signal and to prevent premature go responses (e.g. Aron, 2011; Elchlepp, Lavric, Chambers, & Verbruggen, 2016; Jahfari, Stinear, Claffey, Verbruggen, & Aron, 2010; Verbruggen & Logan, 2009a; Verbruggen, Stevens, & Chambers, 2014; Zandbelt, Bloemendaal, Neggers, Kahn, & Vink, 2012). Both reactive and proactive response inhibition play a key role in enabling optimal and goal-directed behaviour; without the ability to stop habitual or no-longer relevant actions, we would very rapidly find ourselves responding to any stimulus that presented itself in our environment.

Response Inhibition in the Laboratory

Popular paradigms used to study response inhibition in healthy and clinical populations are the go/no-go task (Donders, 1868/1969) and the stop-signal task (Logan & Cowan, 1984; Verbruggen & Logan, 2008b). Figure 1.1 presents typical trial structures for these tasks.

In the go/no-go paradigm, participants are presented with a series of stimuli and are instructed to respond when a 'go' stimulus is presented and to withhold their response when a 'no-go' stimulus is presented (e.g. non-living words = go; living words = no-go; Figure 1.1). The primary index of inhibitory control in the go/no-go task is the probability of responding on no-go trials (i.e. commission errors) which becomes lower as no-go performance improves.

In the stop-signal paradigm, participants are required to perform a speeded choice reaction time task (e.g. living words = left keypress; non-living words = right keypress; Figure 1.1). On a random selection of trials, a stop signal is presented (e.g. an additional visual stimulus or an auditory tone) after a variable delay (stop-signal delay; SSD) instructing participants to stop their go response. The stop-signal delay is usually varied using an adaptive tracking procedure based on the participants' individual levels of stop performance. Although inhibitory control in the stop-signal task can be indexed via the probability of responding on stop trials, manipulating the SSD allows researchers to estimate the covert latency of the inhibitory control process in the stop-signal task[5] (i.e. the stop-signal reaction time; SSRT; Logan, 1981; Logan & Cowan, 1984; Logan, Van Zandt, Verbruggen, &

---

[4] 4. Indeed, inhibition at the cellular (neural) level is even involved in response execution (for an overview, see Aron, 2007).

[5] 5. Note that for reasons that will be discussed in Chapter 2, I do not use the SSRT as a measure of inhibitory control in the experiments presented in this thesis. This approach is consistent with similar experiments in the stop learning literature (e.g. Bowditch, Verbruggen, & McLaren, 2016; Noël, Brevers, et al., 2016; Verbruggen, Best, Bowditch, Stevens, & McLaren, 2014).

Wagenmakers, 2014; but see Verbruggen, Chambers, & Logan, 2013 for a cautionary note about estimation methods).

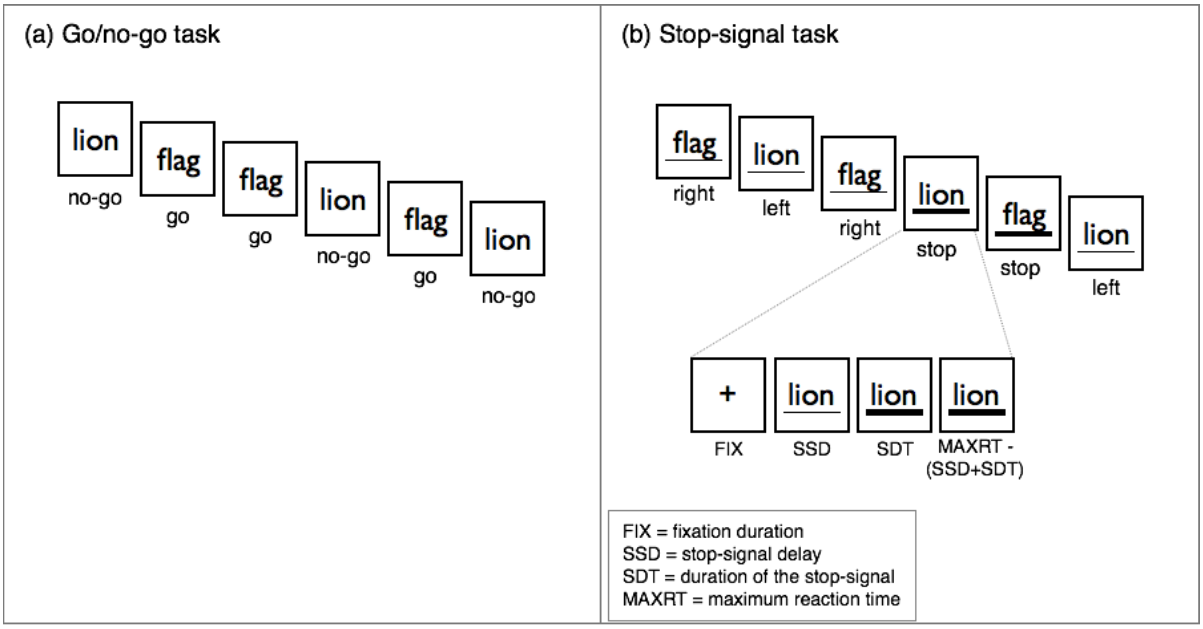

Figure 1.1. Trial structure in the (a) go/no-go task and (b) stop-signal task. In the go/no-go task, participants are required to respond to one class of stimuli (e.g. non-living words = go) and to withhold their response to another class of stimuli (e.g. living words = no-go). In the stop-signal task, participants perform a choice reaction time task (e.g. non-living words = right keypress; living words = left keypress). On a selection of trials, an additional stimulus is presented (in this example, the line beneath the word becoming thicker) after a variable SSD to instruct participants to stop their go response. Participants are instructed to respond as quickly and as accurately as possible on no-signal (i.e. go) trials, but to stop their response if a stop signal is presented. On no-signal and stop-signal trials, the stimulus remains on the screen until the maximum reaction time (MAXRT) has elapsed.

Modelling Performance in Response Inhibition Tasks

Independent Horse-Race Model

Performance in the go/no-go and stop-signal tasks can be modelled as a 'race' between a 'go' process triggered by the presentation of the go stimulus (i.e. a 'go goal') and a 'stop' process triggered by the presentation of the no-go stimulus or the stop signal (i.e. a 'stop goal'; Logan & Cowan, 1984; Logan, Van Zandt, et al., 2014; for a review, see Verbruggen & Logan, 2009b).

According to the independent race model, response inhibition depends on the relative finishing times of the go and stop processes [Logan & Cowan, 1984; Figure 1.2(a)]. When the stop process finishes before the go process [i.e. go RT > (SSRT + SSD)] response inhibition is successful and the go response is withheld (i.e. signal-inhibit). On successful stop trials, the inhibition of the go response typically occurs within 200-300 ms of stop-signal presentation (Logan & Cowan, 1984).

When the go process finishes before the stop process [i.e. go RT < (SSRT + SSD)], response inhibition is unsuccessful and the go response is executed (i.e. signal-respond)[6]. The independent race model assumes that the SSD will influence the relative finishing time of the stop process; when the SSD is increased the stop process will start later and will be more likely to finish *after* the go process (i.e. stopping is unsuccessful); when the SSD is decreased the stop process will start earlier and will be more likely to finish *before* the go process [i.e. stopping is successful; for a visual depiction of this idea, see Figure 1.2(b)].

Variants of the Independent Horse-Race Model

Although the independent race model can account for performance in most response inhibition tasks, there are some situations where it cannot. Therefore, some variants of the independent race model have been proposed to address its shortcomings. In this section, I briefly focus discuss two of these variants, the interactive race model and the blocked-input model, that are relevant to the research presented in this thesis.

A central assumption of Logan and Cowan's (1984) model is that the go and stop processes race *independently* against each other. Most research supports the independence assumption but some behavioural studies have shown that the independence assumption is often not met when the SSD is short (e.g. Logan & Buskell, 1986; Logan, Yamaguchi, Schall, & Palmeri, 2015; Logan, Van Zandt, et al., 2014; Verbruggen & Logan, 2015). Furthermore, the independence assumption is inconsistent with some neuroimaging research (Aron, Durston, et al., 2007; Hanes, Patterson, & Schall, 1998). Thus, an 'interactive race model' (Boucher, Palmeri, Logan, & Schall, 2007) was devised to account for the dependence between the go and stop processes. The 'interactive race model' assumes that the go and stop processes are independent during an initial afferent delay period following the presentation of the stop signal, but that they briefly interact when the stop unit becomes active. On no-signal (go) trials, activation in a go unit (i.e. the go 'goal') begins to accumulate following the presentation of the go stimulus and the go response is executed when an activation threshold is reached. On stop-signal trials, activation of the go unit begins to accumulate in much the same way as on no-signal trials; the initiation of the stop process begins following the presentation of the stop signal and activation of the stop unit (i.e. the stop 'goal') begins after this initial delay. When response inhibition is successful, the stop unit is activated before activation in the go unit reaches threshold; when response inhibition is unsuccessful, activation in the go unit reaches threshold before the stop unit is activated. Thus, according to the 'interactive race model' the go and stop processes remain

---

[6] In support of the race idea, it is often the case that RTs on unsuccessful stop trials are faster than RTs on no-signal trials because only the fastest go responses can escape inhibition [e.g. Logan & Cowan, 1984; see also Bissett & Logan, 2011; see also Figure 1.2(b)].

independent for the majority of their latencies, before they interact towards the end (for reviews, see Verbruggen & Logan, 2008b, 2009b).

Another shortcoming of the independent race model is that whilst it provides a good explanation of whether and when a response is inhibited, it does not address *how* the go response is inhibited. In other words, the independent race model (Logan & Cowan, 1984) describes well the race between the go and stop processes, but does not describe what happens after the stop process has won the race. Boucher and colleagues' interactive race model addressed this shortcoming with the suggestion that the activation of the stop unit slows the accumulation of activation in the go unit (i.e. the go response; Boucher, Palmeri, et al., 1998). Recently, however, an alternative model was proposed which assumes that response inhibition is achieved by blocking input to the go unit instead of inhibiting the accumulation of activation in the go unit (Logan, Yamaguchi, et al., 2015). In sum, the blocked-input model assumes that go responding can be inhibited by disabling the link between the perceptual system and the motor system (Logan, 1983) or by deleting the goals that allow for the go response (Logan & Cowan, 1984).

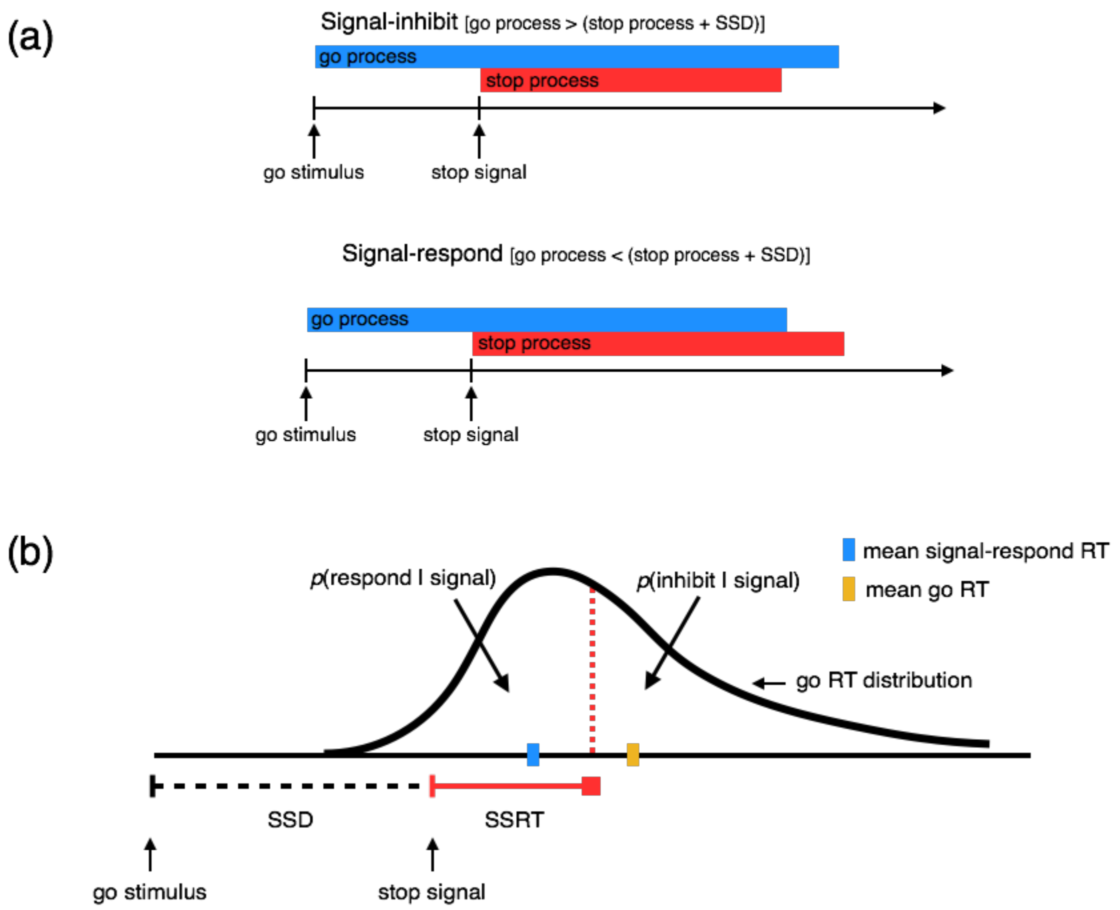

Figure 1.2. (a) Visual depiction of the independent race model (Logan & Cowan, 1984). When response inhibition is successful, the stop process finishes before the go process (i.e. signal-inhibit); when the response inhibition is unsuccessful, the go process finishes before the stop process (i.e. signal-respond) (b) Inhibition function based on the independent race model (Logan & Cowan, 1984). The area under the curve to the left of the red dashed line is the probability of a signal-respond trial; the area to the right of the red dashed line is the probability of a signal-inhibit trial. This figure also shows the mean signal-respond RT and the mean go RT. This shows why signal-respond RTs are faster than mean go RTs: signal-respond RTs are calculated on the basis of the fastest RTs that escaped inhibition (i.e. the RTs to the left of the red dashed line) whereas go RTs are calculated based on the whole RT distribution. This figure also illustrates that the probability of responding on stop trials depends on three factors: the SSD, the SSRT, and the go RT distribution. SSD = stop-signal delay; SSRT = stop-signal reaction time.

The Neural Mechanisms of Response Inhibition

An extensive body research has been conducted into the neural mechanisms of response inhibition. To avoid repetition with forthcoming sections of this thesis, I primarily focus on neuroimaging and brain stimulation (transcranial magnetic stimulation; TMS[7]) studies here. A comprehensive review of electrophysiology research on response inhibition can be found in Chapter 6.

Response inhibition depends on a 'stopping' network of cortical and subcortical regions comprising the right inferior frontal gyrus (IFG; ventrolateral prefrontal cortex), the pre-supplementary motor area (pre-SMA), and the basal ganglia (e.g. Aron & Poldrack, 2006; Rubia, Taylor, et al., 2001; Zandbelt & Vink, 2010; for reviews, see e.g. Aron, 2007; Chambers, Garavan, & Bellgrove, 2009). In particular, research has especially highlighted the critical role of the right IFG for successful response inhibition. For example, the use of TMS has provided causal evidence for the involvement of the right IFG in successful response inhibition (Chambers, Bellgrove, et al., 2006, 2007). Furthermore, neuroimaging research has shown that when response inhibition is successful, activation in the right IFG negatively correlates with the latency of the stop process (i.e. the SSRT; Aron & Poldrack, 2006; Chevrier, Noseworthy, & Schachar, 2007). The most common interpretation is that the right IFG is an 'inhibitory hub' specialised for response inhibition (e.g. Aron, Robbins, & Poldrack 2004; van Boxtel, van der Molen, Jennings, & Brunia, 2001). However, alternative accounts have suggested that the right IFG is involved in more general processes, such as stop signal detection and/or action updating (i.e. replacing the action plan to 'go' with the action plan to 'stop'; Dodds, Morein-Zamir, & Robbins, 2011; Hampshire, Chamberlain, Monti, Duncan, & Owen, 2010; Mars, Piekema, Coles, Hulstijn, & Toni, 2007). In support of the latter interpretation, the use repetitive TMS has shown that the right IFG can be functionally segregated into a dorsal region responsible for the detection of visual changes in the environment, and a ventral region responsible for action updating (Verbruggen, Aron, Stevens, & Chambers, 2010; see also, Chikazoe, Jimura, et al., 2009; Levy & Wagner, 2011)[8].

Alongside, the right IFG, research has also focussed on discerning the relative contributions of the pre-SMA and the basal ganglia regions. Common interpretations are that the pre-SMA is involved in updating current action plans and monitoring for response conflict (e.g. Ridderinkhof, Ullsperger, Crone, & Nieuwenhuis, 2004) and the basal ganglia is involved in the selection of response alternatives, (e.g. to go or to stop) based on prior experiences (Aron & Poldrack, 2006; Chambers, Garavan, et al., 2009; Cohen & Frank, 2009; Graybiel, 2005).

Current research suggests that there are three pathways through which these regions could interact in order to suppress motor output; the direct pathway (also involved in going), the indirect pathway, and the hyperdirect pathway (Aron, Durston, et al., 2007; Nambu, Tokuno, & Takada, 2002;

---

7. TMS generates a transient electrical current in the brain, meaning that neural activity can be disrupted in a given region. The advantage of TMS over neuroimaging techniques is that it allows researchers to establish the causal contribution of a given brain region to the cognitive process of interest (e.g. O'Shea & Walsh, 2007).

8. Note, however, that recent research suggests that there could be something 'special' about response inhibition after all, but it remains unclear exactly 'what' is special about response inhibition (see, e.g. Maizey, unpublished PhD thesis).

Wiecki & Frank, 2013). An in-depth overview of these pathways is beyond the scope of this introduction (but for an overview, see Jahanshani, Obeso, Rothwell, & Obeso, 2015). However, for the purposes of this thesis, it is important to highlight that these different pathways could support both global and specific forms of response inhibition. For example, it has been suggested that the indirect pathway could support a relatively slow but selective inhibition of specific responses (e.g. inhibit left keypress; Aron, 2011; Aron & Verbruggen, 2008) whereas the hyperdirect pathway could support a fast, but global inhibition of the motor system (e.g. inhibit all responses; see Wiecki & Frank, 2013). I will return to this distinction later in this chapter (see *Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding?*).

Top-Down and Bottom-Up Forms of Response Inhibition

Most research in the action control literature attributes response inhibition to ‘top-down’ control. However, in the past decade, growing evidence has suggested that response inhibition is less ‘executive’ than originally thought. In this section, I review two lines of research in support of this idea showing that response inhibition can become ‘a prepared reflex’ readily triggered by information in the environment or a ‘learned reflex’ triggered by the retrieval of previously acquired associations between stimulus and stopping (Verbruggen, Best, et al., 2014).

Response Inhibition as a ‘Prepared Reflex’

Many cognitive control theories attribute successful response inhibition to the *intentional* activation of a stopping process (e.g. Logan & Gordon, 2001; Miller & Cohen, 2001). However, recent research suggests that when processes involved in action control are prepared in advance, the actual initiation of action control may not require much ‘control’ at all (Elchlepp, Lavric, et al., 2016; Verbruggen, McLaren, et al., 2014). In other words, top-down control could become a ‘prepared reflex’. The term ‘prepared reflex’ here refers to idea that ‘stimuli can reflexively trigger the corresponding action based on instructed or planned stimulus–response (S–R) information, even without any prior practice’ (Cohen-Kdoshay & Meiran, 2009, p. 128). More generally, the idea that response inhibition can become a ‘prepared reflex’ is consistent with the ‘implementation intentions’[9] idea that preparing controlled behaviours in advance means that the actual execution of top-down control can be rapid and unintentional (Gollwitzer, 1999; for a discussion, see Verbruggen, Best, et al., 2014).

In support of the ‘prepared reflex’ idea, several studies have demonstrated that response inhibition can be triggered by the presentation of primes in go/no-go and stop-signal tasks. For example, in a series of experiments, van Gaal and colleagues demonstrated that responding on go trials was slower when the go stimulus was preceded by the presentation of a low-visibility (masked)

---

[9] ‘Implementation intentions’ refer to the formation of ‘if-then’ plans that specify goal-directed behaviours in response to situations cues (e.g. ‘if placed in situation X, I will execute behaviour Y’; Gollwitzer, 1999).

no-go stimulus or stop signal than when the go stimulus was not preceded by a no-go stimulus or stop signal (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)[10]. In addition to these behavioural results, van Gaal and colleagues demonstrated that the presentation of a low-visibility no-go stimulus or stop signal activated regions in the 'inhibitory control network' (the inferior frontal gyrus and the pre-SMA; van Gaal, Ridderinkhof, Scholte, & Lamme, 2010) and elicited event-related potential components thought to reflect processes involved in successful response inhibition (the N2 and P3 components; van Gaal, Lamme, Fahrenfort, & Ridderinkhof, 2010; for similar findings, see Hughes, Velmans, & De Fockert, 2009), even though a go response was required. Importantly, the degree of activation in the inhibitory control network and the amplitude of the event-related potential components correlated with the magnitude of the observed slowing for the go responses on low-visibility prime trials. Thus, these neuroimaging and electrophysiology findings support the conclusion that the slowing to the go stimuli observed on prime trials was due, at least in part, to the initiation of stopping by the masked no-go and stop-signal stimuli. Taken together, van Gaal and colleagues' findings provide evidence for the 'prepared reflex' idea because they show that the stop goal can be triggered by the presentation of low-visibility primes even when the explicitly instructed task goal is to respond.

In the van Gaal experiments, the primes were task-relevant as they instructed whether participants should stop on high-visibility stop-signal or no-go trials. However, research has also demonstrated that the activation of a stop goal can be primed by task-irrelevant, high-visibility cues (Verbruggen & Logan, 2009c). In a series of experiments, Verbruggen and Logan (2009c) presented the primes 'GO', '###', and 'STOP' inside shapes (circles or squares). In a first experiment, participants were required to respond to the shape (e.g. circle = left keypress; square = right keypress, or vice versa) but were required to stop this response if an additional auditory tone (the stop signal) was presented. Participants were instructed to ignore the primes presented in the centre of the shapes (i.e. the primes were always task-irrelevant). In support of the 'prepared reflex' idea, responding on go trials in which the 'STOP' prime was presented was slower than on go trials in which the '###' prime or the 'GO' prime was presented. Verbruggen and Logan (2009c) replicated these findings in a subsequent experiment where the same primes were used but as stop signals. Similar findings have also been demonstrated in a go/no-go task where the go stimulus was flanked by other stimuli that were irrelevant to the task; the results showed that when the flankers were presented in the same colour as the no-go stimulus they suppressed motor responding (Anderson & Folk, 2012, 2014).

---

[10] van Gaal and colleagues interpreted these priming effects as reflecting the 'unconscious' activation of stopping. However, these claims about (un)consciousness should be interpreted caution. Although a body of research supports the idea of 'unconscious' or automatic priming (for an overview, see Papies & Aarts, 2016), Newell and Shanks (2014) argued that the detection tasks used by van Gaal, et al. are not sensitive enough to accurately assess awareness.

A key feature of these priming effects is that they appear to depend on top-down control. In the Verbruggen and Logan (2009c) experiments, the 'STOP' prime only activated the stop goal in tasks where stopping was relevant to task performance (e.g. in stop-signal or go/no-go tasks) but not in a task where stopping was irrelevant to task performance (e.g. participants were always required to go). Furthermore, Chiu and Aron (2014) showed that the van Gaal priming effects (i.e. RT slowing following the presentation of a masked low-visibility no-go stimulus) only influenced responding in a go/no-go task (i.e. in the original van Gaal task) but not on a go-only task condition (that was not included in the original van Gaal experiments)[11]. This is entirely consistent with the 'prepared reflex' idea; when top-down control is prepared in advance, the actual initiation of control can operate in a 'reflex' like manner following the presentation of 'trigger' stimuli in the environment (i.e. you need some 'control' before you can observe the reflex).

Combined, these priming findings show that response inhibition can be triggered by 'unconscious' or task-irrelevant cues in the environment. Another feature of a 'prepared reflex' is that a behaviour can be triggered by stimuli based on instructed (S-R) contingencies that have not received any prior practice (Cohen-Kdoshay & Meiran, 2009). Recent research suggests that similar effects to those observed for stimulus-response instructed contingencies can also be obtained for instructed but not overtly practiced stimulus-stop contingencies (Liefooghe, Degryse, & Theeuwes, 2016). In the Liefooghe, Degryse, et al. task, participants were provided with explicit go/no-go task instructions (e.g. "if Q, press left", "if E, do not press"). However, prior to performing the go/no-go task participants performed a number of trials of a two-choice task in which they had to respond to the orientation of a series of letters (e.g. upright = left keypress; italic = right keypress). The main finding was that responding was slower in the two-choice task to stimuli that were assigned to the no-go mapping in the go/no-go task instruction (e.g. "E") than for stimuli assigned to the go mapping in the go/no-go task instructions (e.g. "Q"), even though response inhibition was not required in the two-choice task. Note that the key feature of this experiment is that the go/no-go task instructions were not overtly practiced prior to performance of the two-choice task. Thus, when response inhibition is prepared on the basis of task instructions, these findings suggest that response inhibition can be triggered even when the instructed contingencies were irrelevant to the task and were not previously practiced.

Response Inhibition as a 'Learned Reflex'

The 'prepared reflex' idea suggests that response inhibition can depend on an interplay between top-down and bottom-up factors. However, it is possible that after some practice, top-down control may

---

[11] Note, however, that some concerns have been raised about the priming manipulation used by Chiu and Aron (2014); Lin and Murray (2015) showed that differences in the priming effect between the go/no-go and the go only conditions reported by Chiu and Aron (2014) could arise from differences in awareness of the primes rather than differences in the executive setting (as originally proposed).

not be required at all. When there is a consistent mapping between specific stimuli and the stop goal or response, stimulus-stop associations or 'bindings' could form. With sufficient training, the retrieval of these associations could activate the stop goal[12] in a bottom-up manner, eliminating the need for top-down control entirely (e.g. Logan, 1988; Schneider & Shiffrin, 1977; Shiffrin & Schneider, 1977). In this section, I will review evidence that response inhibition can become a 'learned reflex' triggered by the retrieval of stimulus-stop associations from memory.

Sequential after-effects. Research using modified versions of the go/no-go and stop-signal tasks has demonstrated that there are both short-term and long-term after-effects of stopping (e.g. Bissett & Logan, 2011; Enticott, Bradshaw, Bellgrove, Upton, & Ogloff, 2009; Rieger & Gauggel, 1999; Verbruggen, Logan, Liefooghe, & Vandierendonck, 2008; Verbruggen & Logan, 2008c). The typical finding is that responding to a stimulus is slowed following a stop-signal trial, especially if the go stimulus presented on the preceding stop trial is repeated (Bissett & Logan, 2011; Verbruggen, Logan, et al., 2008). A compelling hypothesis proposed for the observed slowing is that the stimulus presented on a stop-signal trial becomes associated with the stop goal (e.g. see Bissett & Logan, 2011). When the stop goal is 'bound' or associated to the stimulus on trial *n*-1 and this stimulus is repeated on trial *n*, the stop goal is retrieved and this slows responding (Rieger & Gauggel, 1999; Verbruggen, Logan, et al., 2008; Verbruggen & Logan, 2008c). The stimulus-stop association idea in the sequential after-effects literature overlaps with the 'do-not-respond' tag account of negative priming[13]. According to the 'do-not-respond' tag account, a 'do-not-respond' or 'ignore-it' tag becomes associated with the old distractor stimulus; repeating the distractor as a target stimulus on a subsequent trial retrieves this tag which slows go responding (Neill, Valdes, Terry & Gorfein, 1992; Neill & Valdes, 1992). However, it is perhaps no surprise that there are similarities between the 'do-not-respond' tag idea and the stimulus-stop association idea as both accounts are based on the principles of memory-retrieval theories (e.g. Instance theory; Logan, 1988; for details, see *What is 'Automaticity'?*).

Before proceeding, it is important to note that the sequential after-effects demonstrated in the aforementioned studies are assumed to reflect the retrieval of stimulus-stop associations rather than the inhibition of an acquired stimulus-go association. For example, it might be possible to argue that

---

[12] The use of the term 'stop goal' is consistent with the original work on 'automatic inhibition' conducted by Verbruggen and Logan (2008a). However, it is important to note that in the context of the research presented in this thesis, there is no functional difference between 'stop goal' and 'stop response'. Nevertheless, I use 'stop goal' throughout primarily to denote that there is no overt response executed on stop trials (for a similar rationale, see Giesen & Rothermund, 2014). Furthermore, the term 'stop goal' is consistent with research suggesting that the retrieval of stimulus-stop associations has a global effect on all go responding rather than a selective effect in stopping a specific go response (for a discussion, see *Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding?*). Other broadly synonymous terms used in the literature are 'stop unit', 'stop tag', 'stop centre', and 'stop representation'.

[13] The negative priming effect refers to the finding that performance is generally impaired for target stimuli that previously appeared as a distractor in previous trials (e.g. Tipper, 1985).

successful response inhibition on the preceding trial(s) might weaken the association between the target stimulus and the corresponding go response. However, whilst the inhibition of a stimulus-go association explanation may initially seem parsimonious, it cannot explain sequential effects where the stimulus associated with stopping is task-irrelevant. In a series of experiments, Giesen and Rothermund (2014) used a sequential priming paradigm in which there was a ‘prime’ trial and a ‘probe’ trial that occurred in succession. On each prime and probe trial, a word was presented and participants were required to press the spacebar as soon as the word appeared. On a selection of prime and probe trials, an additional auditory tone (i.e. the stop signal) was presented to instruct participants to stop their response. The meaning of the word was always task-irrelevant. Crucially, the word could either be repeated between the prime and probe trials or replaced with a new word. The results showed that if participants were required to stop their response on a prime trial, responding was slower on the subsequent probe trial if the word identity was repeated. This argues against the ‘inhibition of a stimulus-go association’ idea because it shows that the retrieval of an stimulus-stop association can influence go responding even when the stimulus identity associated with stopping was not associated with any go response (i.e. there was no stimulus-response association to be inhibited or weakened). Thus, this experiment suggests that the irrelevant stimulus identity (word meaning) became associated with the stop goal which was retrieved when the stimulus identity was repeated.

Response Inhibition as an ‘Automatic Act of Control’. The sequential after-effects research suggests the presentation of a stop signal can affect performance up to 20 trials following stop-signal presentation (Verbruggen & Logan, 2008c). Thus, the relatively long-lasting nature of stimulus-stop associations reported in the sequential after-effects literature suggests that an automatic form of response inhibition could develop if the stimulus-stop associations were reinforced within a training phase. After all, as will be reviewed in more detail below (see *‘What is ‘Automaticity’?*), the formation of associations between specific stimuli and responding may be the first step towards automaticity (Logan, 1989, 1990, 1992).

In a series of experiments, Verbruggen and Logan (2008a) examined the idea that inhibitory control in go/no-go and stop-signal tasks can be triggered automatically via the retrieval of stimulus-stop associations from memory. In one experiment, participants performed a go/no-go task in which a word category determined if a participant should respond or not (e.g. living word referents = go; non-living word referents = no-go; Verbruggen & Logan, 2008a, Experiment 1). After a training phase, the go/no-go mapping was reversed in a test phase. It was found that responding to the old stop stimuli was slower than responding to new stimuli that were not previously presented during training (consequently, these new stimuli were not associated with going or stopping). Similar differences were also found in a go/no-go task in which the judgments changed between the training and test phases; participants made living/non-living judgments in the training phase and small/large judgments in the test phase (or vice versa; Experiment 2). Responding was slower for the inconsistent items (i.e. no-go in the training phase task but go in the test phase task) than for the consistent items (i.e. go in both the training and test phases).

Based on these findings, Verbruggen and Logan (2008a) proposed the 'automatic inhibition hypothesis'; following training, 'automatic inhibition' occurs when old no-go items retrieve the stop goal, even when go responding is required. In its traditional form, the stop-signal task does not allow for the formation of stimulus-stop associations because there is no consistent mapping between specific stimuli and stopping (Figure 1.1). However, in a final experiment Verbruggen and Logan (2008a) demonstrated that similar effects to those observed in the go/no-go task could occur in the stop-signal task when the mappings between specific stimuli and stop signal presentation are manipulated (Experiment 5). In this experiment, there was a training phase in which there was a subset of stimuli that was consistently presented on stop-signal trials, a subset of stimuli that was consistently presented on no-signal trials, and another subset that was presented on both stop-signal and no-signal trials (i.e. these stimuli were not consistently associated with either stopping or going). In a test phase, the mappings were reversed such that the stop-associated stimuli were consistently presented on no-signal trials. Consistent with the go/no-go experiments, responding was slowed for the stimuli that were consistently presented on stop-signal trials during the training phase. The finding that responding is slowed for stop-associated stimuli has been subsequently replicated using similar paradigms in healthy (e.g. Bowditch, Verbruggen, et al., 2016; Lenartowicz, Verbruggen, Logan, & Poldrack, 2011) and clinical populations (Noël, Brevers, et al., 2016; Woolard, Kose, et al. 2010). Combined, these findings suggest that response inhibition may not always require top-down 'executive' control and can become a 'learned reflex' when the mappings between specific stimuli are consistent throughout task practice.

Before proceeding, it is important to note that the slowing observed for the old stop-associated stimuli cannot be accounted for by non-specific sequential effects. For example, due to the stimulus-stop manipulation and the fixed overall probability of stop signals, an old stop item (which was always 'go' in test) would more likely to follow a stop trial than the control items (which were 'go' or 'stop') in the test phase. In a recent experiment, participants performed a stop learning task in which there was a training (acquisition) phase followed by a test (reversal) phase (Verbruggen, Best, et al., 2014, Experiment 1). In an experimental condition, a subset of stimuli were consistently presented on stop-signal trials during training, but were presented no-signal trials in the test phase ('stop-then-go' items) and another subset were presented on stop-signal and no-signal trials with equal probability during training (50%) but the trial order was otherwise random. The test trial for these items was also always a no-signal trial ('stop/go-then-go'). The trials in the control group followed the same sequence as participants in the experimental group but there were no consistent mappings between specific stimuli and stopping. Consistent with the ‘automatic inhibition’ account, the results showed that responding in the test phase was slower to ‘stop-then-go’ items than to the ‘stop/go-then-go’ items in the experimental condition. Importantly, there was no reliable difference between these stimulus-types in the control condition. This difference was supported by a reliable condition by stimulus-type interaction. The absence of a difference between the stimulus-types in the control group indicates that the slowing observed for the stop-then-go items in the experimental group was unlikely to be due to general (non-stimulus specific) sequential after-effects of stopping (e.g. goal priorities, proactive adjustments; Bissett & Logan, 2011).

In summary, responding to items that are associated with stopping is slowed compared with responding to items that are inconsistently associated with stopping. This slowing has been attributed to the retrieval of stimulus-stop associations from memory; when there is a consistent mapping between specific stimuli and stopping, associations can form between the stimulus and the stop goal over practice. This could allow automatic processing to develop (Shiffrin & Schneider, 1977).

What is ‘Automaticity’?

The formation of stimulus-stop bindings may be to the first step towards automaticity (Logan, 1990). Memory-retrieval accounts of automatisation assume that every time people respond to a stimulus, processing episodes are stored as ‘instances’ (Logan, 1988) or ‘event files’ (Hommel, 1998, 2004) in memory. These instances or event files may contain information about the stimulus (e.g. the word), the interpretation given to the stimulus (e.g. ‘non-living’), the task goal (e.g. ‘go’) and the response (e.g. ‘left key press’). These episodes are retrieved when a stimulus is repeated, and will influence responding. Logan’s (1988) Instance theory postulates that action selection can be construed as a race between an algorithmic response-selection process and a memory-retrieval process; the process that finishes first determines which action is selected. The Instance theory assumes that when people begin to perform a novel task they rely on general algorithms (i.e. rules) that are sufficient to perform the task, but as they get more experienced they begin retrieve the instances stored in memory. When the memory-retrieval process wins the race, the decision is said to be automatic; when the algorithmic process wins the race, the decision is said to be deliberate or intentional (Logan, 1988, 1990, 1992).

When the stimulus-response or stimulus-stop mapping is the same throughout practice, multiple instances are formed and automatic processing can develop (Logan, 1988, 1990, 1992; Shiffrin and Schneider, 1977). Behaviour is considered to be fully automatised when there is ‘single-step direct access’ of past solutions from memory (Logan, 1988). At this point, the algorithm can be abandoned as there are enough instances in memory to respond on the basis of the retrieved instances on every trial. In effect, the algorithm races against the fastest instance that can be retrieved from memory; as training progresses, the algorithm is more likely to lose the race because the finishing time of the algorithmic process remains the same whereas the finishing time of the memory-retrieval process becomes faster. Thus the shift to the memory-retrieval process will either be based on the statistical properties of the race between these processes or based on a strategic decision to abandon the algorithmic process entirely (Logan, 1988). Either way, during training, the Instance theory predicts that reaction times will decrease as a power function of the number of instances stored in memory (Logan, 1992; for an example in response inhibition tasks, see Verbruggen & Logan, 2008a)[14]. Verbruggen and Logan (2008a) showed that the effects of stop learning emerged after just four stimulus presentations.

In summary, memory-retrieval accounts suggest that, with sufficient training, response inhibition can become automatised, such that the requirement for top-down control could theoretically be eliminated entirely.

How do Stimulus-Stop Associations Influence Responding?

Theoretical Explanations

Verbruggen and Logan (2008a) suggested that the activation of the stop goal following the retrieval of stimulus-stop associations from memory can be best explained by the interactive race model (Boucher, Palmeri, et al., 2007; see *Modelling performance in response inhibition tasks*); following training, the go goal can be activated via the retrieval of an association between the go stimulus and going and the stop association can be activation via the retrieval of an association between the go stimulus and stopping. When a stop-associated stimuli is presented, activation in the stop unit slows the rate of accumulation in the go unit. Thus, go RTs are slowed for the stop-associated stimuli relative to the control stimuli that were not particularly associated with stopping or going. Verbruggen and Logan (2008a) assumed that the activation of the stop unit is weaker when it is incorrectly activated on go trials than when it is correctly activated on no-go or stop-signal trials, which explains why go responding is slowed following the retrieval of a stimulus-stop association but is only occasionally completely inhibited.

Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding?

In two experiments, Giesen and Rothermund (2014) investigated whether the retrieval of the stop goal had a specific effect on responding (e.g. stop left keypress) or a global effect on responding (e.g. stop all responses). In these experiments, words were presented in either green or red and participants were required to make a left keypress response if the word was presented in green and a right keypress response if the word was presented in red. On a selection of trials, an auditory stop signal was

---

[14] Note that the Instance theory makes no assumptions about the numbers of trials required to reach automaticity. Consequently, it is possible that performance on some trials is automatic whereas performance on others is (more) algorithmic (Logan, 1988).

presented to instruct participants to stop their planned response. Consistent with their previous research (see *Sequential after-effects*, Chapter 1), when the word was repeated on the probe trial responding was delayed if the word was presented with a stop signal on the preceding prime trial, regardless of the word colour. Importantly, the response slowing on the probe trial was not influenced by the response that was stopped on the prime trial and the correct response on the probe trial. Although speculative, the global effect on responding observed by Giesen and Rothermund could indicate that stimulus-stop associations operate through the hyperdirect neural pathway (see *The Neural Mechanisms of Response Inhibition*, Chapter 1; for similar conclusions, see *The Neural Mechanisms of Stimulus-Stop Learning*, Chapter 1).

The Neural Mechanisms of Stimulus-Stop Learning

To date, a small number of studies have investigated the neural mechanisms of stimulus-stop learning. Recent neuroimaging research indicates that the presentation of an old stop-associated stimulus on a go trial may activate regions in the inhibitory control network (Lenartowicz, Verbruggen, et al., 2011). In a study conducted by Lenartowicz and colleagues, there were two experimental sessions: on day 1, participants performed a stop-signal task in which some face stimuli were paired with a stop signal more often than other face stimuli; on day 2, participants performed the stop learning task in the MRI scanner. The first two blocks in the scanner were training blocks to ensure that participants had not forgotten the mappings, then the mappings were reversed in a test phase, such that the old stop-associated faces were always presented on go trials. Consistent with the predictions of the 'automatic inhibition' hypothesis, the behavioural data showed that stopping performance improved during the training phase for the faces that were consistently paired with stopping compared with the faces that were not particularly associated with either stopping or going (these faces presented on stop trials on 50% of presentations). Following the reversal of the mappings in the test phase, go performance was slower for the old stop-associated faces. Importantly, it was found that the right IFG was activated in the test phase following the presentation of the old stop-associated faces (even though a go response was required). As discussed above (see *The Neural Mechanisms of Response Inhibition*, Chapter 1), several studies have suggested that the right IFG plays a key role in successful stopping (e.g. Aron, Robbins, et al., 2004, 2014; Chambers, Garavan, et al., 2009). This suggests that in addition to being activated by explicit stop signals, key regions of the inhibitory control network could be activated following the retrieval of stimulus-stop associations associations (cf. van Gaal, Ridderinkhof, Scholte, & Lamme, 2010). This is consistent with the idea that the presentation of old stop-associated items can activate the stop goal, even when go responding is required. However, although the stimulus-stop explanation for these results is plausible, it is important to highlight that the right IFG has been linked with a variety of other functions, such as attentional orienting and action selection (e.g. Dodds, Morein-Zamir, et al., 2011; Hampshire, Chamberlain, et al., 2010) meaning that the activation observed by Lenartowicz, Verbruggen, et al. cannot be conclusively attributed to the retrieval of stimulus-stop associations[15]. Nevertheless, on the face of it, these initial findings suggest that there is an overlap between the mechanisms of 'top-down' and 'bottom-up' forms of response inhibition.

Consistent with the idea that the presentation of old no-go or stop-associated stimuli suppress the accumulation of activation in the 'go unit', a TMS study has demonstrated that presentation of old no-go stimuli influenced the excitability of the motor cortex (Chiu, Verbruggen, & Aron, 2012). Following a training phase, TMS was applied in that study to the cortralateral motor cortex of the go response hand and motor-evoked potentials[16] were measured after word presentation. Chiu, Verbruggen, et al. (2012) showed that motor excitability was suppressed within 100 ms following the presentation of old no-go words. The rapid timing of this effect suggests that, at the very least, 'automatic inhibition' reflects a fast, efficient form of inhibition[17]. The rapid timing of the motor suppression is consistent with event-related potential research which showed that following performance of an auditory go/no-go task, the topographies of auditory-evoked potentials were modulated within ~100 ms of presenting stimuli previously consistently associated with stopping (see also Manuel, Grivel, Bernasconi, Murray, & Spierer, 2010). Interestingly, however, there was no reduction of corticospinal excitability for no-go stimuli in a consistent condition where the go/no-go mappings were not reversed in the test phase. Chiu, Verbruggen, et al. suggested that the effects of go/no-go learning on the motor cortex could be context dependent. According to this idea, the retrieval of stimulus-stop associations can suppress motor output but these effects appear to be dependent on an top-down 'executive setting' (a similar conclusion was reached about the response inhibition priming work, see *Response Inhibition as a 'Prepared Reflex'*, Chapter 1).

However, it is important to highlight that in addition to the 'automatic inhibition' account, there is an alternative explanation that could explain why motor excitability is suppressed following the reversal of the go/no-go mappings. It was recently suggested that the decreased motor excitability for the old no-go items in the inconsistent condition could have been driven, at least in part, by conflict between the trained goal or response representation (i.e. no-go) and the required response (i.e. 'go'; see Verbruggen, Best, et al., 2014). Research has suggested that when conflict between responses is detected, a global braking mechanism is triggered that prevents premature responding

---

[15] For example, the stimulus-stop mappings were reversed only when participants entered the scanner. This is potentially problematic as participants could have re-learnt the new mappings in the scanner, reducing the influence of the previously learnt mappings, and introducing the possibility that the observed right IFG activation reflects re-learning, not the effect of retrieving old stimulus-stop associations from memory (e.g. Passingham, & Toni, 2001; Wallis, Dias, Robbins, & Roberts, 2001).

[16] Motor-evoked potentials provide a measure of corticospinal excitability; reduced corticomotor excitability corresponds to inhibition of the motor cortex (e.g. Valls-Solé, Pascual-Leone, Wassermann, & Hallett, 1992).

[17] The rapid timing of this effect could correspond to the fast hyperdirect pathway detailed above (for a similar claim, see *Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding?*, Chapter 1; for an overview of the hyperdirect pathway, see *The Neural Mechanisms of Response Inhibition*, Chapter 1).

(Frank, 2006; Wiecki & Frank, 2013). Thus, although the automatic inhibition and the conflict detection accounts both predict slowing for the old no-go items, they propose different mechanisms through which the suppression arose. According to the 'automatic inhibition' account, it is the direct activation of the stop response; according to the the conflict detection account, it is the conflict between the indirect activation of the braking mechanism (i.e. after conflict has been detected). In their paper, Chiu, Verbruggen, et al. cast doubts on the conflict account because the observed effects on motor excitability occurred very rapidly (within 100 ms of stimulus presentation), leaving very little time for conflict to be detected. However, some support for the conflict account is provided by a recent event-related potential experiment showing that sequential after-effects of stopping do not always result from the 'automatic' retrieval of the stop goal but rather interference between stop and go goals (Oldenburg, Roger, Assecondi, Verbruggen, & Fias, 2012).

Can 'Top-down' Forms of Response Inhibition be Trained?

Before proceeding, it is important to highlight that some training effects have been reported in response inhibition tasks where there is no consistent stimulus-stop mapping (Berkman, Kahn, & Merchant, 2014; Chavan, Mouthon, Draganski, van der Zwaag, & Spierer, 2015; Ditye, Jacobson, Walsh, & Lavidor, 2012; Hartmann, Sallard & Spierer, 2015; Logan & Burkell, 1986; Manuel, Bernasconi, & Spierer, 2013; but see Cohen & Poldrack, 2008). However, unlike the aforementioned stimulus-stop training effects, the absence of consistent stimulus-stop associations in these tasks mean that training-effects in these studies are unlikely to reflect the transition from top-down control to bottom-up control. Instead, it is seems likely that training-effects in 'top-down' response inhibition tasks reflect the transition from top-down reactive control to top-down proactive control (for an overview of reactive and proactive control, see *What is 'Response Inhibition'?*, Chapter 1). In support of this suggestion, a recent neuroimaging study suggests that stop-signal training could induce strategic (deliberate) changes in performance (Berkman, Kahn, et al., 2014). In Berkman, Kahn, et al.'s study, participants performed a stop-signal task across 10 sessions in a three-week period. Neural activation was measured at the beginning and at the end training. The neuroimaging results indicated that after training, activity in right IFG following stop-signal presentation decreased whereas activity prior to stop-signal presentation increased (relative to a control group who competed the task but without any stop signals). The observed shift in the timing of activation suggests that training encouraged a proactive response strategy, as participants started 'controlling' their responses before the presentation of the stop signal (for similar claims about the involvement of proactive control in top-down training-effects see Verbruggen, Adams, & Chambers, 2012; Verbruggen, Adams, et al., 2013). Thus, top-down control can be trained, but this seems to reflect in strategic changes to task performance rather than the development of automatic forms of inhibition.

Practical Applications

Over the past few years, a growing body of research has investigated whether stop learning can be used as a behaviour change intervention. Recent studies suggest that training on item-specific response inhibition tasks could be an effective way to reduce engagement in impulsive behaviors, such as excessive food (Houben, 2011; Houben & Jansen, 2011; Lawrence, Verbruggen, Morrison, Adams, & Chambers, 2015; Lawrence, O'Sullivan, et al., 2015a; van Koningsbruggen, Veling, Stroebe, & Aarts, 2014; Veling, Aarts, & Papies, 2011; Veling, Aarts, & Stroebe, 2013a, 2013b; Veling, van Konningsbruggen, Aarts, & Stroebe, 2014) and alcohol (Bowley, Faricy, et al., 2013; Houben, Havermans, Nederkoorn, & Jansen, 2012; Houben, Nederkoorn, Wiers, & Jansen, 2011; Jones & Field, 2013) consumption in healthy and clinical populations (for meta-analyses see Allom, Mullan, & Hagger, 2015; Jones, Di Lemma, et al., 2016). The majority of these studies used paradigms in which no-go or stop signals were superimposed over, or presented around, images of unhealthy foods or alcohol. The effects of training in these have been assessed in terms of intake following task completion (e.g. Houben, 2011; Houben & Jansen, 2011) and outside of the lab for up to 6 months following food inhibition training (Allom & Mullan, 2015, Study 1; Lawrence, Verbruggen, et al., 2015b) or up to one week following alcohol inhibition training (e.g. Bowley, Faricy, et al., 2012; Houben, Havermans, Nederkoorn, & Jansen, 2012; Houben, Nederkoorn, Wiers, & Jansen, 2011; Jones & Field, 2013). The effects of training on unhealthy food consumption were strongest in individuals who are particularly prone to overeating ('restrained eaters'; Houben & Jansen, 2011; Lawrence, Verbruggen, et al., 2015; Veling, Aarts, & Papies, 2011), in individuals with a high appetite, and in individuals who regularly consume the unhealthy food items used in the training task (Veling, Aarts, & Stroebe, 2013a, 2013b) suggesting that these training tasks could be effective in 'boosting' inhibitory control in individuals with low top-down inhibitory control.

Training-effects on consumption appear to be strongest in tasks where there is a consistent mapping between the food- and alcohol-related images and stopping (Allom, Mullan, et al., 2015); evidence for general training effects in top-down response tasks (i.e. where the stimulus-stop mappings are inconsistent) are weak (e.g. Guerrieri, Nederkoorn, & Jansen, 2012). This suggests that the retrieval of stimulus-stop associations could play a role in the observed effects on food and alcohol consumption. The challenge in this applied research is to develop inhibitory control training tasks that have a long-term effect on behaviour. In order to develop effective training protocols, it is important to understand the mechanisms through which stimulus-stop associations are acquired and the factors that influence the effect of stimulus-stop learning on behaviour. Therefore, alongside being of theoretical interest, the research presented in this thesis could also be of interest to researchers in the applied inhibitory control training domain.

Main Outstanding Issues

In light of the research discussed in this chapter, I propose that there are three main outstanding issues that need to be addressed in order to understand how and when response inhibition can become automatised: (1) what is learned?; (2) the role of expectancies; and (3) the specificity of stop learning. The research presented in this thesis addressed each of these outstanding issues. Therefore, I give a brief overview of each issue below.

In light of the research discussed in this chapter, I propose that there are three main outstanding issues that need to be addressed in order to understand how and when response inhibition can become automatised: (1) what is learned?; (2) the role of expectancies; and (3) the specificity of stop learning. The research presented in this thesis addressed each of these outstanding issues. Therefore, I give a brief overview of each issue below.

What is Learned?

In a recent review article, it was proposed that there are at least two pathways through which learning in response inhibition tasks could influence behaviour (Verbruggen, Best, et al., 2014; for a simplified version of this associative architecture, see Figure 1.3). First, there is the *direct* pathway in which a link between a stimulus and the 'stop goal' or stop representation is formed. The direct pathway[18] was assumed in Verbruggen and Logan's (2008a) original 'automatic inhibition' hypothesis; it is via the direct pathway that the repeated presentations of stop-associated stimuli can begin to automatically activate the stop goal and slow responding on go trials when the mappings are reversed. Second, there is the *indirect* pathway in which a link is formed between a stimulus and a representation of whatever stop signal or no-go category was used in the given response inhibition task. The *indirect* pathway exploits the consistent association between the stimulus and the stop signal in response inhibition tasks. Verbruggen and colleagues suggested that the acquisition of stimulus-signal associations would have important consequences for our understanding of what is learned in response inhibition tasks; importantly, under these circumstances the presentation of a stop-associated stimulus could just activate the representation of the stop signal rather than the stop goal. This would result in the absence of (or severely weakened) stop learning relative to circumstances where participants acquired stimulus-stop associations. In support of this these pathways, recent research has shown that the slowing observed for old stop-associated stimuli was more pronounced in a condition where the the stop-associated stimuli were paired with multiple coloured stop signals than when the stop-associated stimuli were paired with a single coloured stop signal (Bowditch, Verbruggen, et al., 2016). This could indicate that participants in the 'multiple signal' group acquired more direct associations between the stop items and the stop goal compared with participants in the 'single signal' group. Thus, there currently exists some (indirect) evidence suggesting that the effects of stop learning are

---

[18] Note that there is no direct mapping between this direct learning pathway and the direct neural pathway outlined earlier (see *The Neural Mechanisms of Response Inhibition*, Chapter 1).

weaker when stimuli are paired with a single representation of the stop signal than when stimuli are paired with multiple representations of the stop signal (Bowditch, Verbruggen, et al., 2016). However, it remains unclear whether participants can acquire direct associations between stimuli and the stop goal that are not mediated via a representation of the stop signal or indirect stimulus-signal associations when the contingencies between specific stimuli and features of the stop signal are manipulated.

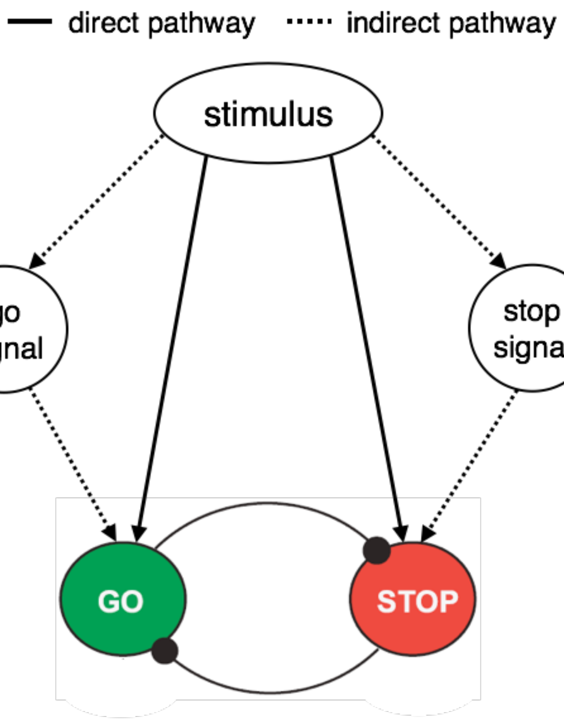

Figure 1.3. Overview of the architecture of the associative stop system (for a more detailed overview, see Verbruggen, Best, et al., 2014). There are two associative routes to activating the stop-goal; a direct pathway between the stimulus and the go/stop goal (solid lines), or an indirect pathway between the stimulus or cue and the go/stop goal that is mediated via a representation of the go/stop signal (dashed lines).

The Role of Expectancies

Previous studies in the stop learning literature can be placed into two groups: those in which the stimulus-specific stop associations are acquired (in modified stop-signal tasks; e.g. Lenartowicz, Verbruggen, et al., 2011; Verbruggen & Logan, 2008a, Experiment 5) and those in which a stimulus-category rule is explicitly instructed (e.g. in the go/no-go task; Chiu, Aron, et al., 2012; Verbruggen & Logan, 2008a, Experiments 1-4). This raises the issue of whether the influence of stop learning under instructed and uninstructed conditions on behaviour operate through similar mechanisms. Research in the associative learning literature indicates that people often generate expectancies about the characteristics of upcoming trials, even under conditions of uncertainty (McAndrew, Yeates, Verbruggen, & McLaren, 2013; Perruchet, Cleeremans, & Destrebecqz, 2006). However, expectancies have not been examined in stop learning studies. Yet, evidence of expectancies could have important implications for our understanding of what is learned in response inhibition tasks as it could indicate that training induces strategic (deliberate) changes in performance, as participants could learn to control their responses before a stop signal is presented (i.e. proactive control), rather than after the occurrence of the stop signal (i.e. reactive control). Thus, it is important to investigate the role of expectancies during stop learning as evidence for expectancies could undermine the idea that item-specific stop training causes a shift from top-down control to bottom-up control (i.e. the 'automatic inhibition hypothesis'; Verbruggen & Logan, 2008a). Instead, it would indicate that after practice, there is a shift from top-down reactive control to top-down proactive control as seems to be the case in top-down response inhibition training (see *Can 'Top-down' Forms of Response Inhibition be Trained?*, Chapter 1).

The Specificity of Stop Learning

It has been suggested that response inhibition may be similar to response execution in many ways (e.g. McLaren & Verbruggen, 2014; Verbruggen, McLaren, et al., 2014; see also Mostofsky & Simmonds, 2008). Indeed, it has recently been suggested that various forms of action control depend on three basic cognitive processes: signal detection, action selection, and action execution (Verbruggen, McLaren, et al., 2014). These basic cognitive processes can account for performance in inhibition and non-inhibition tasks (e.g. Logan & Bundesen, 2003; Monsell & Mizon, 2006; Ratcliff & Smith, 2004; Schall, 2001; Smith & Ratcliff, 2004). In the action control literature, however, Verbruggen and colleagues suggest that researchers often attribute control to an ill-defined set of 'black-box homunculi' that are assumed to perform control functions, such as 'response inhibition', rather than considering how control is achieved (Verbruggen, McLaren, et al., 2014).

Consistent with this suggestion, many researchers in stop learning studies explain the effects of training in terms of strengthening or improving 'inhibitory control' (see, e.g. Spierer, Chavan, & Manuel, 2013). This could give the impression that the processes influenced by training are specific to response inhibition tasks. However, this seems unlikely. After all, the idea that specific stimuli can become associated with stopping shares many similarities with stimulus-response bindings observed in other cognitive control paradigms, such as in the task-switching paradigm. A typical finding in the task-switching paradigm is that performance is impaired (responding is slower and accuracy is lower) when the task changes compared with when it stays the same; a 'switch cost' (e.g. Monsell. 2003). Allport and colleagues have suggested that the switch cost can be explained in terms of the longer-term effects of interference via the retrieval of previous stimulus-task associations. In other words, they suggest that prior task performance results in the formation of 'bindings' between stimulus attributes and the related task-set that hinders performance on switch trials (e.g. Koch & Allport, 2006; Waszak, Hommel & Allport, 2003, 2004, 2005). Furthermore, similar effects of training to those reported by Verbruggen and Logan (2008a) have also been shown in non-inhibition tasks (i.e. go only tasks; see Horner & Henson, 2009). However, the effects of training in inhibition and non-inhibition tasks have not yet been compared in the same experimental design. Thus, claims about the specificity of stop learning are currently based on indirect cross-experiment comparisons. Therefore, it remains an outstanding issue whether the effects of learning are specific to inhibition tasks or whether they represent task-general (non-inhibition) effects.

Thesis Outline

The goal of this thesis is to increase understanding of how stimulus-stop associations are acquired and to investigate the conditions under which stop learning influences behaviour.

In Chapter 2 (Experiments 1-4), I focus on two of the main outstanding issues highlighted above (see *Main Outstanding Issues*, Chapter 1): (1) can participants acquire stimulus-stop associations that are not mediated via a single representation of the stop signal (i.e. is there support for the *direct* pathway?; see above); and (2) what is the role of expectancies during stimulus-stop learning? To this end, participants performed a task that combined features of the go/no-go and stop-signal tasks in which the stop-signal rule changed at the beginning of each block. Furthermore, expectancies were measured following task completion (Experiments 1-3) or measured within the task (Experiment 4).

In Chapter 3 (Experiments 5-6), I take a closer look at the role of expectancies during the acquisition of item-specific stop learning. To investigate whether stimulus-stop learning could induce strategic (deliberate) changes in performance, I adapted a paradigm that was specifically developed to examine the effects of proactive control in the stop-signal task. Furthermore, to examine the contributions of explicit expectancies and learning, I introduced a between-subjects expectancy manipulation in which one group of participants were explicitly informed about the stimulus-stop contingencies and another group did not receive any information about the stimulus-stop contingencies.

In Chapter 4 (Experiments 7-8), I directly investigate whether participants can acquire stimulus-signal associations when the representation of the signal remains the same throughout training (i.e. the *indirect* pathway; Figure 1.3). Participants performed a go/stop learning task in which the contingencies between specific stimuli and the spatial location of the stop signal were manipulated, in addition to the contingencies between specific stimuli and the stop goal.

In Chapter 5 (Experiments 9-10), I concentrate on another of the outstanding issues highlighted above: the specificity of learning in response inhibition tasks. Currently, very little is known about whether training in response inhibition tasks influences specific inhibition processes or more general processes involved in the performance of inhibition and non-inhibition tasks. I aimed to advance on this position by comparing the effects of learning on behaviour in the go/no-go task with the two-choice task.

In Chapter 6 (Experiments 11-12), I investigate the neural mechanisms of stop learning. Previous electrophysiology studies have suggested that the N2 and P3 event-related potential components are sensitive to training in the go/no-go task. However, no study has investigated the specificity of these effects. Therefore, I followed-up on the behavioural findings reported in Chapter 5 by comparing the effects of training (within-subjects) on the latencies and amplitudes of the N2 and P3 components in inhibition (go/no-go) and non-inhibition (two-choice) tasks.

In Chapter 7, I summarise the main findings of the research presented in this thesis. Furthermore, I discuss the contribution of this thesis to our understanding of action control, highlight the limitations of the experiments presented, and provide some suggestions for future research.

CHAPTER 2

ASSOCIATIONS AND EXPECTANCIES IN RESPONSE INHIBITION TASKS

Current research in the stop-learning literature appears to provide strong support for the ‘automatic inhibition’ account that postulates that stimuli can become associated with the act of stopping. However, as outlined in Chapter 1, it is still unclear exactly what is learned in these tasks and how learning influences performance. The present chapter was designed to address two of the main outstanding issues that were highlighted in Chapter 1 (similar issues were also recently raised in the context of S-R bindings; Henson, Eckstein, Waszak, Frings & Horner, 2014): (1) are associations between stimuli and stopping direct, and (2) to what extent does expectancy play a role?

Are Associations Between Stimuli and Stopping Direct?

The automatic inhibition account assumes that people learn direct associations between a stimulus and the act of stopping in go/no-go tasks and modified versions of the stop-signal task. However, the results of a recent experiment are inconsistent with this account (Verbruggen, Best, et al., 2014). In that experiment, participants made speeded semantic categorisations (living/non-living) of a series of words. On some trials (stop-signal trials) an additional visual signal was presented below the word, instructing participants to withhold their planned response. Certain words were consistently presented on stop-signal trials, whereas other words were presented on go and stop-signal trials with equal probability. Verbruggen, Best, et al. found that the probability of responding on stop-signal trials was lower for the consistent words than for the inconsistent words in the training phase, indicating that learning had occurred. However, there was no evidence of a go RT difference between the old stop words and the inconsistent words when the stimulus-stop mapping was subsequently reversed in the test phase. In other words, learning influenced stop performance on signal trials in the training phase, but it did not influence go performance on no-signal trials in the test phase. Verbruggen, Best, et al. proposed that this pattern of results indicates that participants learned stimulus-signal associations rather than stimulus-stop associations. Such associations between the stop words and the stop signal (i.e. the line turning bold) will prime the representation of the stop signal rather than the stop goal. Signal detection plays a critical role in successful stopping (e.g. Verbruggen, Stevens & Chambers, 2014), and computational work suggests that a considerable proportion of the stopping latency is occupied by perceptual or afferent processes (Boucher, Palmeri, et al., 2007; Logan, Van Zandt, et al., 2014; Logan, Yamaguchi, et al., 2015; Salinas & Stanford, 2013). Thus, by priming the representation of the stop signal, learning could lead to improvements in stopping performance on stop-signal trials without influencing responding on go trials in the test phase.

The idea that participants could learn stimulus-signal associations is also consistent with a range of research on learning and conditioning in humans and other animals indicating that stimulus detection can itself become conditioned (McLaren, Wills, & Graham, 2010) and, of course, that links between perceptual stimuli can be established. As an illustrative (and rather basic) example, in a classic autoshaping paradigm with pigeons, the presentation of a conditioned stimulus (e.g. a keylight) and an unconditioned stimulus (e.g. the delivery of food) usually co-occur. With practice, the presentation of the conditioned stimulus alone can come to elicit the conditioned response (e.g. pecking at this key). The conditioned stimulus can activate this response via two routes; either indirectly via the CS-US link, or more directly, via a CS-R link (Hall, 2002). Thus, it seems plausible that learning can also influence perception of the no-go or stop signal in response-inhibition paradigms.

The potential for stimulus-signal associations has important implications for the interpretation of previously reported behavioural effects in the stop-learning literature. Previous studies that have used no-go or stop signals to manipulate stimulus-stop learning cannot distinguish between stimulus-goal and stimulus-signal learning. It is therefore possible that previously observed RT effects and neural activations (Lenartowicz , Verbruggen, et al., 2011; Manuel, Grivel, et al., 2010) could be mediated by a link between the stimulus, the stop signal, and stopping (see Figure 1.3, Chapter 1). Similarly, in go/no-go experiments in which the go/no-go rules are explicit (e.g. living = go, non-living is no-go), the stimulus-stop association could be mediated via the go/no-go category (e.g. ‘desk = non-living -> non-living = no-go’, instead of ‘desk = no-go’)[19].

In addition to being of theoretical interest, the idea of stimulus-stop associations also has implications for applied stop-training research as stimulus-signal learning could reduce or even eliminate the effect of stop training on behaviour (for an overview of the applied inhibition training literature, see *Practical Applications*, Chapter 1). For example, if participants acquired stimulus-signal associations the effects of inhibitory control training on food- or alcohol-consumption (if any) would only occur when in the presence of the stop signal. Thus, it is important to establish whether there is any evidence for the original idea (i.e. as suggested by Verbruggen & Logan, 2008a) that direct associations can be acquired between a stimulus and the stop goal, without mediation via a representation of the stop signal (or no-go category). To discourage the formation of stimulus-signal associations, the stop signal and the task rules changed at the beginning of each block. As previous studies in the stop learning literature have all used a constant signal representation, the demonstration of response slowing for consistent stop items in the present study would provide the strongest

---

19. Mediation via a category representation may actually have an adaptive advantage because it would reduce the memory-load of saving all individual instances in memory (for a similar idea, see Arrington & Logan, 2004).

evidence to date that, at least in some situations, participants can acquire direct stimulus-stop associations.

What is the Role of Expectancy in Stimulus-Stop Learning?

In the associative-learning literature, there is an on-going debate surrounding the involvement of explicit and implicit processes in the acquisition of stimulus-action associations (Mitchell, De Houwer, & Lovibond, 2009). To make a broad distinction, ‘explicit’ processes are assumed to be controlled, intentional, effortful and rule-based; by contrast, ‘implicit’ processes are assumed to be automatic, effortless, and associative (e.g. McLaren, Green, & Mackintosh, 1994; for a recent discussion of the distinction between associative and propositional processes, see McLaren, Forrest, et al., 2014)[20]. Expectancy ratings have been used to dissociate between the two processes (e.g. McLaren, Forrest, et al., 2014; Newell & Shanks, 2014). In the context of stop-learning, this dissociation between rule-based processes and associative (S-S or S-R) processes has important theoretical implications. After all, expectancy of a stop signal for old stop items could indicate that the response slowing observed for old stop items is due to proactive inhibitory control, rather than ‘automatic inhibition’. When a cue indicates that a stop signal is likely to occur on the following trial(s), participants proactively increase response thresholds or suppress motor activation (e.g. Jahfari, Verbruggen, et al., 2012; Ramautar, Kok, & Ridderinkhof, 2004; Verbruggen & Logan, 2009a; Zandbelt, Bloemendaal, et al., 2013). Stimuli associated with stopping could act as such cues (e.g. ‘if stimulus X then p(stop) is high’), and participants would adjust their response strategies accordingly. In other words, slowing for old stop items could be due to proactive control (which may be conceived as another ‘algorithmic’ process; cf. Logan, 1988), rather than to the direct activation of the stop response via memory retrieval. The role of expectancy-driven processes is also relevant for the applied stop-training research. Indeed, the extent to which training effects like these reflect implicit or associative effects has been called into question. For example, Boot and colleagues (2013) argued that many ‘control’ training effects could be due to changes in expectations and demand characteristics. The involvement of expectancies would have implications for the longevity of these inhibitory control training effects and the variability of training efficacy across individuals (cf. Boot, Simons, Stothart, & Stutts, 2013).

In the present chapter the role of expectancy in stimulus-stop learning was investigated via the inclusion of an additional dependent variable that was sensitive enough (Newell & Shanks, 2014)

---

[20] Note that the 'consciousness’ is also often referred to in this distinction. However, the unconscious/conscious characteristics of explicit/implicit processes are subject to intense debate (see, e.g. Shanks & St John, 1994). This is primarily because it very difficult to assess whether a given process was performed unconsciously. However, for the purposes for the present chapter, it is important to highlight that conscious awareness of a stimulus-stop association does not necessarily indicate that stop learning was *not* automatic (for a discussion, see *Is ‘automatic inhibition’ truly automatic?*, Chapter 1).

to detect stimulus-stop learning following task completion (Experiments 1-3) or within the task (Experiment 4).

Experiment 1

In Experiment 1, features of a go/no-go task and a stop-signal task were combined. The go/stop task used was based on those used in studies examining the effects of no-go training effects on food and alcohol consumption (see *Practical Applications*, Chapter 1). Similar to picture-word Stroop tasks (see e.g. MacLeod, 1991), go and stop signals were superimposed over forty neutral images. There was no delay between the presentation of the images and the signals (i.e. zero ms). A subset of the images was consistently associated with stop signals (75% stop), another subset was consistently associated with go signals (100% go), and the remaining images were control images (25% stop). After twelve training blocks (80 trials per block), the mappings were reversed, and participants had to respond to the stop-associated images. Participants were not informed about the mappings, but they were told at the beginning of each block what the go and stop signals were. To discourage the formation of stimulus-signal or stimulus-category associations, the representation of the go and stop signals was varied at the beginning of each block. It was predicted that this change manipulation would encourage the formation of image-stop associations (cf. Verbruggen & Logan, 2008a) instead of image-signal associations (i.e. S-R rather than S-S learning). Learning during task performance was indexed via two measures. The first index was the probability of responding on the stop trials, *p*(respond|stop), which was predicted to be lower for stop-associated images than for the control images. The second index was RT on go trials, which was predicted to be longer for the stop-associated images than for the control images. To examine the role of expectancy in stop learning, participants were asked to rate the extent to which they expected to withhold their response for each of the images presented in the task at the end of the experiment.

Method

Participants. Participants. Thirty-one students from the University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 19.43 years, *SD* = 1.70 years, 17 females, 27 right-handed). Two participants were excluded because they incorrectly executed a response on ≥ 30% of the stop-signal trials (there was no delay between the presentation of the image and the stop signal; consequently, *p*(respond|stop) was expected to be low). The target sample and exclusion criteria were determined before data collection.

Apparatus and stimuli. The experiment was run on an Apple iMac using Psychtoolbox (Brainard, 1997). The stimuli were presented on a 20-in monitor (screen size: 1680 × 1050 pixels). The experimental paradigm consisted of a go/stop task in which the go/stop rule changed at the beginning of each block. The go and stop signals (a full list of the signals used is presented in Appendix A) were superimposed over forty task-irrelevant neutral images (size: 250 × 250 pixels), which were presented in the centre of the screen on a white background. Each image was presented twice per block. In each block, two go signals (e.g. the vowels 'a' or 'e') and two stop signals were used (e.g. the consonants 't' or 'n'). Participants responded on go trials by pressing the spacebar on a keyboard with their right index finger; they were instructed to withhold their response on stop trials. The signals and the go/stop mapping were shown on the screen at the beginning of each block for a minimum of 5 seconds and participants had to press a key to start the first trial. The order of the task rules was randomised across the blocks and the response-rule category was counterbalanced across participants (e.g. ‘go = vowels, stop = consonants’ vs. ‘go = consonants, stop = vowels’).

Procedure. Unbeknown to the participants, there were two phases in the experimental paradigm that determined the image-stop/go mappings; the first 12 blocks of 80 trials comprised the ‘training phase’ and the final two blocks of 80 trials comprised the ‘test phase’. Participants were verbally instructed to read the task rule screen carefully before starting each block. There was a 15 second break between each block.

There were three image types (Table 2.1). First, stop-associated images were paired with a stop signal on 75% of presentations in the training phase; in the test phase, they were always paired with a go signal. Second, go-associated images were always paired (100%) with a go signal in the training phase, but they could occur on stop trials in the test phase (eight old go-associated images were paired with a stop signal on 75% of presentations; eight old go-associated images were never paired with a stop signal). Third, control images were paired with a stop signal on 25% of presentations in the training and test phases. The control images were mostly paired with a go signal during training to ensure that the overall probability of a stop trial [$p(\text{stop}) = 0.25$] was the same in the training and the test phases (stopping performance is sensitive to minor variations in signal probability, e.g. see Bissett, & Logan, 2011).

	# images	% stop-signal trials	
		Training phase	*Test phase*
Experiment 1			
stop-associated	8	75%	0%
go-associated	16	0%	8 images: 75%; 8 images: 0%
control	16	25%	25%
Experiment 2			
stop-associated	10	100%	0%
go-associated	30	0%	20 images: 0%; 10 images: 100%
Experiment 3 & Experiment 4			
stop-associated	8	75%	0%
go-associated	16	0%	4 images: 0%; 12 images: 50%
control	16	25%	8 images: 0%; 8 images: 50%

Table 2.1. Proportion of stop-signal trials as a function of experiment, image type, and phase. The overall *p*(stop) both across experiments and within the experimental phases was 0.25.

All trials began with the concurrent presentation of the image and a go/stop signal (Figure 2.1), instructing participants to execute (go) or withhold (stop) the spacebar response. After 750 ms (regardless RT), the images and go/stop signal were replaced by a feedback message (‘correct’, ‘incorrect’, or ‘too slow’ in case the participant did not respond before the end of the trial) which remained on the screen for 500 ms. The feedback message was presented to encourage fast and accurate responding. Following the feedback message, there was a blank screen for 250 ms, after which the next trial started.

Following completion of the experimental task, each image was again presented on the screen. The order of the images was randomised anew for each participant. Participants were asked to rate ‘how much do you expect to withhold your response when this image is presented?’ on a scale between 1 ('I definitely do not think this image indicates that I have to withhold my response') and 9 ('I definitely think this image indicates that I have to withhold my response'). As a manipulation check, I also asked participants to rate how much they expected to respond (i.e. go) to each of the images (the order of the respond/withhold ratings was counterbalanced across participants). These go ratings were consistent with the stop expectancy ratings so are not reported further.

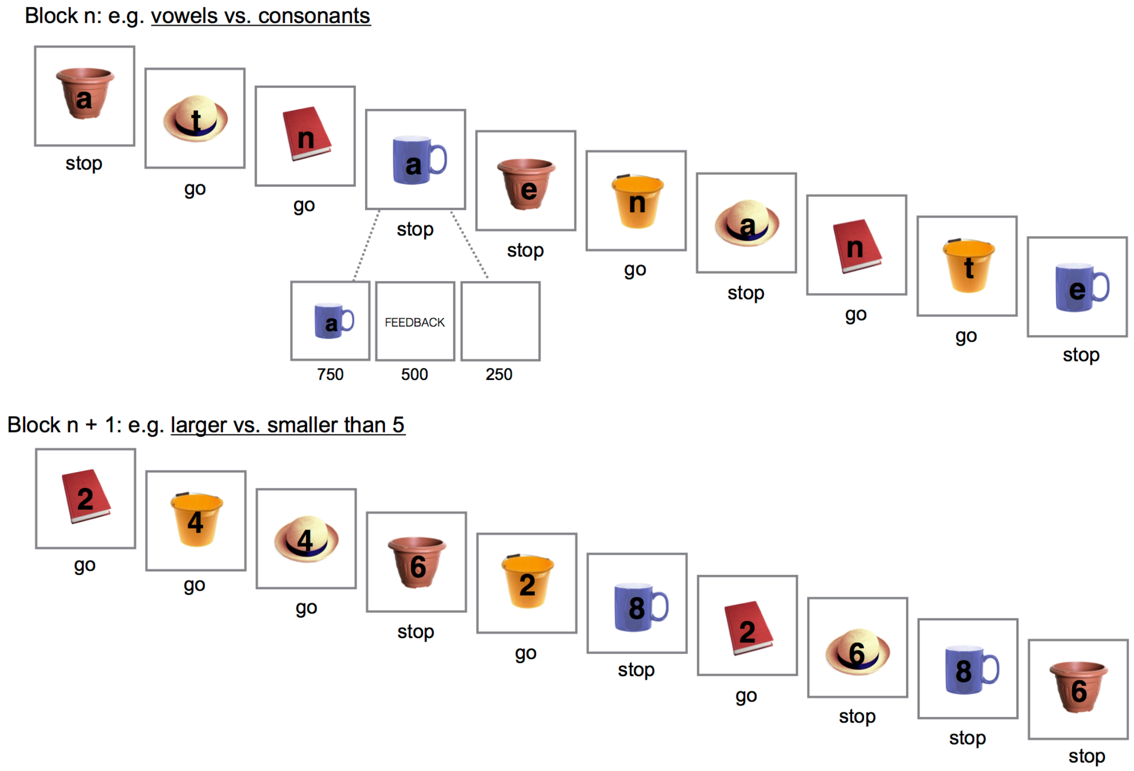

Figure 2.1. Example go/stop trial sequence. The task rule changed at the beginning of each block (e.g. Block *n*: vowel = stop; consonant = go, Block *n* + 1: > 5 = stop; < 5 = go). In Experiments 1-3, the go/stop signals were superimposed on top of the image (as shown). In Experiment 4, the signals were presented in one of the four corners of the image (top-left, bottom-left, top-right, bottom-right).

Analyses. All data processing and analyses were completed using R (R Development Core Team, 2014). The training and test phase trials were analysed separately using Analyses of Variance (ANOVA) with image type and block as within-subjects factors. Performance was assessed in terms of average RT for correct go responses, the probability of a missed go response [*p*(miss)] and the probability of responding on a stop trial [*p*(respond|stop)][21]. RTs < 1 ms were removed prior to analyses. I did not analyse *p*(miss) further as values were very low (Table 2.2). Table 2.3 provides an overview of the ANOVAs. For pairwise comparisons, Hedge's g_{av} is the reported effect size measure (Lakens, 2013). All data files and R scripts used for the analyses are deposited on the Open Research Exeter data repository (http://hdl.handle.net/10871/17735).

---

21. The design was optimised to examine stimulus-stop learning for the stop-associated items but this design made it suboptimal for the estimation of SSRTs; therefore, SSRT were not estimated or analysed. However, I analysed *p*(respond|signal) to determine if learning influenced stopping performance (see also e.g. Bowditch, Verbruggen, et al., 2016; Noël, Brevers, et al., 2016). It was predicted that the *p*(respond|signal) would become lower for the stop-associated images than for the other image types during the training phase.

Results

Training phase. The main effect of image type on go RTs was reliable ($p < 0.001$); planned comparisons revealed that responding to the stop-associated images (on the relevant 25% of trials) was slower (414 ms) than to the go-associated images (403 ms), $t(28) = -4.93$, $p < 0.001$, $g_{av} = 0.44$, and to the control images (406 ms), $t(28) = -3.26$, $p = 0.002$, $g_{av} = 0.33$. There was a marginally reliable difference between the go and the control images, $t(28) = -1.99$, $p = 0.055$, $g_{av} = 0.11$ (Figure 2.2; Table 2.3). In line with my predictions, the *p*(respond|stop) was lower for the stop-associated images (0.13) than for the control images (0.15), $p = 0.019$ (Figure 2.2). Thus, performance on go and stop trials suggests that participants acquired the image-stop associations. The effect of block and the interaction between block and image type did not reach significance suggesting that the effect of image type was present in most blocks (Table 2.3). This is consistent with previous work which indicates that the effect of stop learning emerges after a single trial presentation and that it then quickly asymptotes (Verbruggen & Logan, 2008a, 2008c). The absence of an overall practice effect is most likely due to the introduction of a novel go/stop rule at the beginning of each block; consistent with this idea, a post-hoc test confirmed that participants responded faster in the second half of a block than in the first half, $t(28) = 3.99$, $p < 0.001$, $g_{av} = 0.32$.

	Training phase		Test phase	
	M	***SD***	***M***	***SD***
Experiment 1				
stop-associated	0.020	0.071	0.013	0.033
go-associated	0.015	0.024	0.014	0.024
control	0.016	0.028	0.017	0.036
Experiment 2				
stop-associated	-	-	0.024	0.043
go-associated	0.016	0.025	0.013	0.031
Experiment 3				
stop-associated	0.028	0.083	0.038	0.060
go-associated	0.023	0.034	0.037	0.062
control	0.018	0.031	0.036	0.052
Experiment 4				
stop-associated	0.028	0.088	0.018	0.033
go-associated	0.020	0.032	0.007	0.018
control	0.021	0.032	0.025	0.036

Table 2.2. Probability of a missed go response [*p*(miss)] as a function of experiment, image type and experimental phase. *P*(miss) is the ratio of the number of omitted responses to the total number of no-stop-signal trials: *p*(miss) = missed/ (correct + missed). *M* = mean; *SD* = standard deviation.

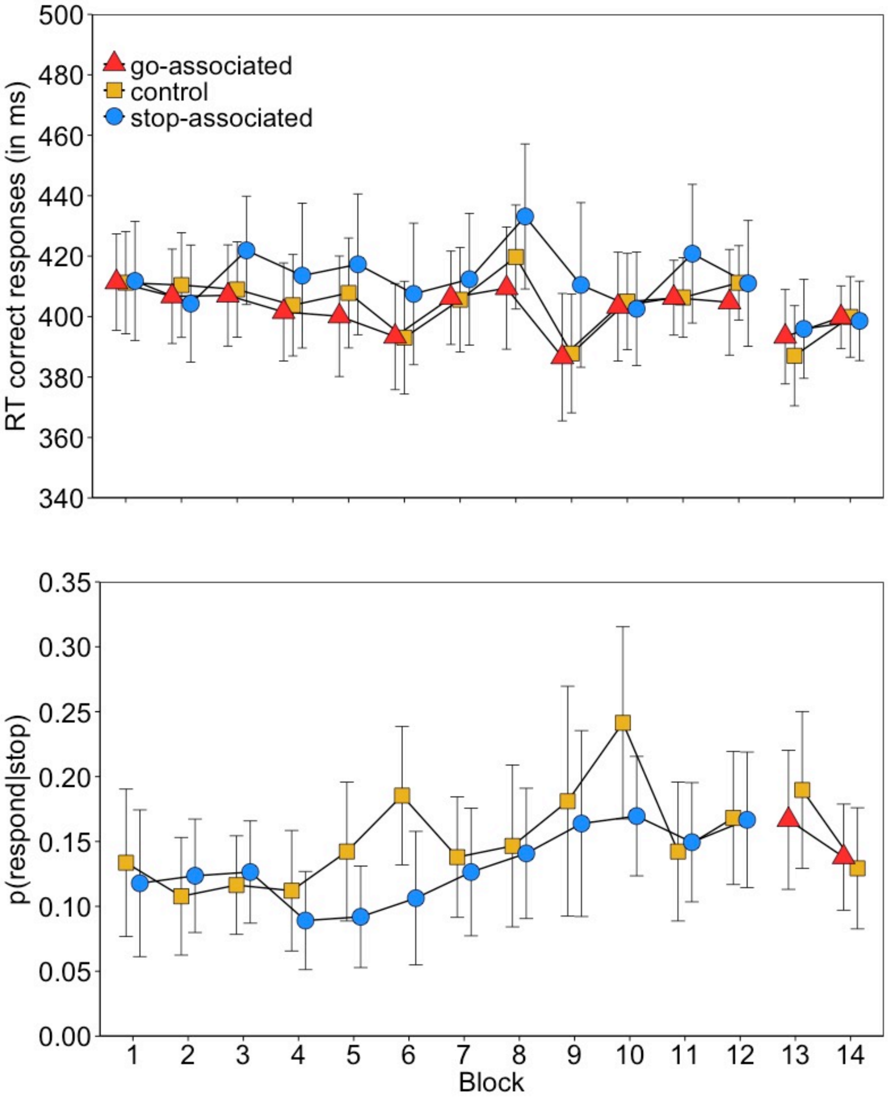

Figure 2.2. Reaction times for the correct go trials (upper panel) and *p*(respond|stop) data (lower panel) for the three image types (stop; go; control) as a function of the block (blocks 1-12 = training phase; blocks 13-14= test phase) in Experiment 1. Error bars are 95% confidence intervals.

Test phase. In the test phase, the stop-associated images were always paired with a go signal, the control images were paired with a stop signal on 25% of the trials (i.e. the control images remained the same in the training and test phases), and the go-associated images were mostly paired with a stop signal (8 images: 75% stop; 8 images: 0% stop; Table 2.1). Based on the automatic inhibition hypothesis, I predicted that responding on go trials would be slower for the stop-associated images than for the go-associated images and for the control images. Furthermore, *p*(respond|stop) should be higher for the go-associated images than for the control images. However, image type did not influence RT nor the *p*(respond|stop) in the test phase (*p'* ≥ 0.557; Table 2.4). It is possible that the absence of the test phase effect is due to differences in the overall RT (as RTs were faster in the test phase than in the training phase). To investigate this possibility, I plotted RT percentiles for the training and test phases. This revealed that the overall test phase RT cannot account for the absence of the predicted image-stop learning effects (see Appendix B).

	df 1	***df 2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Experiment 1							
Go Reaction Time							
image type	2	56	21980	41878	14.70	**< 0.001**	0.009
block	11	308	43427	1438376	0.85	0.575	0.017
image type by block	22	616	17475	431981	1.13	0.331	0.007
p(respond|stop)							
image type	1	28	0.071	0.323	6.17	**0.019**	0.005
block	11	308	0.547	7.616	2.01	**0.043**	0.040
image type by block	11	308	0.154	3.384	1.28	0.238	0.011
Experiment 2							
Go Reaction Time							
block	11	319	27502	405062	1.97	**0.039**	0.037
p(respond|stop)							
block	11	319	0.199	3.846	1.50	0.136	0.037
Experiment 3							
Go Reaction Time							
image type	2	60	5589	47836	3.51	0.058	0.006
block	5	150	64257	437275	4.41	**< 0.001**	0.061
image type by block	10	300	16048	162531	2.96	0.005	0.016
p(respond|stop)							
image type	1	30	0.010	0.200	1.48	0.232	0.002
block	5	150	0.053	2.688	0.60	0.703	0.009
image type by block	5	150	0.055	1.780	0.93	0.461	0.009
Experiment 4							
Go Reaction Time							
image type	2	54	6816	87254	2.11	0.153	0.009
block	5	135	75823	192102	10.66	**< 0.001**	0.092
image type by block	10	270	13700	203293	1.82	0.094	0.018
p(respond|stop)							
image type	1	27	0.087	0.317	7.45	**0.011**	0.015
block	5	135	0.109	1.857	1.58	0.173	0.019
image type by block	5	135	0.180	2.363	2.05	0.077	0.031

Table 2.3. Overview of repeated Analyses of Variance performed to compare go and stop training phase performance. Image type (Experiments 1, 3, & 4: stop-associated, go-associated, control) and block (Experiments 1 & 2: 1-12; Experiments 3 & 4: 1-6) are the within-subjects factors. I did not analyse *p*(miss) because values were low. *p*s < 0.05 are highlighted in bold.

Expectancy ratings. Due to technical reasons, one participant in Experiment 1 did not complete the expectancy ratings task (note that this participant was included in the training and test phase analyses). The results of the test phase raise some doubts about whether participants learned long-term image-stop associations. However, the analysis of the expectancy ratings obtained following task completion revealed a main effect of image type, *F*(2, 54) = 10.06, $p < 0.001$, *gen.* $\eta^2 = 0.075$. Consistent with the stimulus-stop contingencies during training, participants expected to withhold their response more when the stop-associated images were presented (4.83) than when the go-associated images (3.91) and the control images (4.26) were presented; *t*(27) = -3.46, $p = 0.001$, $g_{av} = 0.65$, and *t*(27) = -2.74, $p = 0.010$, $g_{av} = 0.40$, respectively. The difference between the control and the go-associated images was also reliable, *t*(27) = -2.89, $p = 0.007$, $g_{av} = 0.27$. Thus, participants could distinguish between the images on the basis of their association with the stop and go goals. The ‘stop minus control image’ expectancy difference correlated with the corresponding RT difference in the test phase, *r*(26) = 0.437, $p = 0.019$: participants who expected to withhold their response more during the presentation of the old stop-associated images slowed more when they had to respond to these images in the test phase (see Figure C.1, Appendix C). This suggests that expectancies generated on the basis of the acquired image-stop mappings may contribute to the manifestation of an ‘automatic’ inhibition effect in the test phase. However, there was no reliable correlation between the ‘stop minus control’ expectancy difference and the corresponding RT in the training phase, *r*(26) = 0.010, $p = 0.961$. There was also no reliable correlation between the RT and expectancy differences for the stop- and the go-associated images in the training phase, *r*(26) = - 0.040, $p = 0.841$, or the test phase, *r*(26) = 0.272, $p = 0.161$ (note that uncorrected *p*s are reported).

	df 1	*df 2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Experiment 1							
Go Reaction Time							
image type	2	56	471	22425	0.59	0.557	0.002
block	1	28	2318	53598	1.21	0.281	0.010
image type by block	2	56	771	13040	1.66	0.200	0.003
p(respond|stop)							
image type	1	28	0.001	0.268	0.16	0.695	< 0.001
block	1	28	0.058	0.380	4.24	**0.048**	0.028
image type by block	1	28	0.007	0.316	0.64	0.429	0.004
Experiment 2							
Go Reaction Time							
image type	1	29	160	10621	0.44	0.513	< 0.001
block	1	29	390	47352	0.24	0.629	0.002
image type by block	1	29	60	7087	0.25	0.624	< 0.001
p(respond|stop)							
block	1	29	0.020	0.337	1.73	0.198	0.019
Experiment 3							
Go Reaction Time							
image type	2	60	711	28880	0.74	0.479	0.005
p(respond|stop)							
image type	1	30	0.51	0.416	3.73	0.062	0.043
Experiment 4							
Go Reaction Time							
image type	2	54	7329	28228	7.01	**0.004**	0.064
p(respond|stop)							
image type	1	27	0.050	0.473	2.83	0.104	0.034

Table 2.4. Overview of repeated Analyses of Variance performed to compare go and stop test phase performance. Image type (Experiments 1, 3, & 4: stop-associated, go-associated, control; Experiment 2: stop-associated, go-associated) and block (Experiments 1 & 2: 13-14; Experiments 3 & 4: 7) are the within-subjects factors. I did not analyse *p*(miss) because values were low. *p*s < 0.05 are highlighted in bold.

Interim Discussion

In Experiment 1, I investigated two questions highlighted in Chapter 1: (1) can participants learn direct associations between stimuli and stopping; and (2) what is the role of expectancy in stimulus-stop learning? The results provide some answers to both questions. Task performance during the training phase showed that participants could acquire direct stimulus-stop associations when the rules (and consequently, signals) constantly changed throughout the task. This indicates that the learning effects were not mediated via signal representations (as each image was only presented twice per block and there were only two stop signals and two go signals per block). Furthermore, the expectancy data obtained following task completion showed that participants generated expectancies that were consistent with the stimulus-stop contingencies acquired during training.

However, the results of Experiment 1 raised a new question: why did stimulus-stop associations not influence performance in the test phase? I found an associative effect on behaviour that appeared early in training but then disappeared again in the later training blocks and in the test phase (Figure 2.3; for similar results in another action control paradigm, see Gaschler & Nattkemper, 2012), even though the expectancy data measured at the end of the experiment indicated that the associations were not forgotten. I attribute this to an interaction between attention and learning. The role of attention in stimulus-stop learning has not yet been considered. In previous studies demonstrating stimulus-stop learning (e.g. Verbruggen & Logan, 2008a), the go/stop items were task-relevant as they determined the required response; consequently, optimal task performance in these studies depended on participants attending to the stop items (as opposed to the signals). In the present chapter, I adapted a paradigm frequently used in applied research (e.g. Houben & Jansen, 2011) whereby go/stop signals were superimposed on a series of images. This was advantageous as it allowed me to vary the representation of the go/stop signals throughout the task whilst independently manipulating the image-stop contingencies. However, a consequence of this procedure is that optimal task performance does not depend on attending to the stop-associated images. Initially, the task-irrelevant images may have captured attention because they were novel, allowing the effects of learning to emerge. But habituation to the images and reduced salience may have reduced attentional capture, and consequently, weakened or even eliminated the effects of stop-learning on behaviour in later blocks.

The hypothesised role of attention in the acquisition of stimulus-stop associations is consistent with the associative learning literature. For example, a review by Kruschke (2003) indicates that attention is crucial in explaining associative learning phenomena. Following the principles first enunciated by Mackintosh (1975), he argued that attending to informative cues whilst ignoring irrelevant cues will accelerate learning. Furthermore, the amount of attention that is paid to the cues will determine the influence of acquired associations on behaviour. In a similar vein, the Instance theory assumes that attention determines what is learned and what is retrieved (Logan & Etherton, 1994; Logan, 1988). But attention can also be influenced by learning. For example, the learned predictability of the outcome relative to other concurrently presented cues may influence the extent to which cues are considered informative or salient, and consequently, the extent to which participants attend to them (see Mackintosh, 1975). Consistent with this suggestion, Livesey & McLaren (2007) demonstrated that stimuli that were better predictors of an outcome became relatively more salient than stimuli that were worse predictors of the outcome over practice (see also

Le Pelley & McLaren,2003)[22]. In other words, previous research indicates that attention and associative learning go hand in hand.

In Experiment 1, the stop-associated images could be considered relatively worse predictors of the stop goal when presented with a stop signal. After all, the stop-associated images were associated with the stop goal (i.e. the outcome in this case) on 75% of the trials, whereas any given stop signal (e.g. the consonants 't' or 'n') was associated with the stop goal on 100% of presentations. Similarly, control images could occur on both go and stop trials. Therefore, attentional accounts of associative learning predict that the images would decrease in salience with exposure; consequently, their contribution to performance would also diminish with increased image exposure (see Le Pelley, Suret & Beesley, 2009). The suggestion that the relative salience of the images diminished during training is also consistent with conflict monitoring accounts (e.g. Botvinick, Braver, Barch, Carter & Cohen, 2001). These accounts predict decreased attention to the images due to response conflict triggered by the inconsistency in the predictability of these images. For instance, Egner & Hirsch (2005) have demonstrated that when response conflict is detected, task-relevant information is amplified. Hence, conflict detection accounts predict that participants should increase their attention to the go/stop signals relative to the task-irrelevant images. Thus, in this regard, the main difference between the associative learning and conflict monitoring accounts is the detailed mechanism by which the cognitive system adjusts attentional settings. The conflict account requires conflict to drive this change in attention whereas the associability account does not. All the latter requires is that one stimulus (in this case the stop signal itself) has a greater associative strength to the outcome (stopping) than the other stimulus present (the image).

In sum, the findings of Experiment 1 show that participants can acquire direct associations between specific stimuli and the stop goal. However, despite reliable learning effects in the training phase and in expectancy ratings obtained following task completion, I found no evidence of learning in the test phase when the stimulus-stop mappings were reversed. I hypothesise that attention plays a role in determining the influence of stimulus-stop learning on behaviour. This idea could put important constraints on current theories of the automaticity of control processes. Therefore, I conducted three more experiments to replicate and extend the findings of Experiment 1, and to explore the role of attention in the influence of stimulus-stop associations on behaviour.

---

22. There is some evidence suggesting that inconsistent reinforcement can increase attention to, and the motivation salience of, conditioned stimuli. For example, the Pearce-Hall (1980) model suggests that associability is maintained for stimuli that are followed by unpredictable outcomes. However, despite animal data in support of this effect (e.g. Anselme, Robinson & Berridge, 2013), there is relatively little data showing this effect in humans (see Hogarth, Dickinson, Austin, Brown & Duka, 2008). The weight of evidence using humans participants is in favour of the Mackintosh (1975) model outlined above (but for a combination of both algorithms in one model, see Pearce & Mackintosh, 2010).

Experiment 2

In Experiment 1, I hypothesised that habituation and the predictability of the signal-stop contingency relative to the image-stop contingency decreased the amount of attention that was paid to the stop-associated images over practice. To investigate the predictability hypothesis, in Experiment 2, I manipulated the contingency between the images and stopping, to ensure that the stop-associated images were paired with a stop signal and were predictive of the stop goal on 100% of presentations during training (cf. 75% of presentations in Experiment 1). This should prevent conflict driving down attention, but it would not abolish any associability effects as the stop signal would still tend to be the stimulus with the strongest connection to stopping. All that an associability theory requires for the images to lose attention is that they are worse predictors of the outcome relative to the stop signal(s). This will occur when the stop signal(s) always predicts the outcome whereas the images only predict the stop goal on the trials on which they occur. As a result, image associability will be driven down in a block, and will not have time to recover when the stop signal changes at the beginning of each block.

Method

Participants. Thirty students from the University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 19.97 years, *SD* = 2.81, 23 females, 27 right-handed). No participants were excluded.

Apparatus, stimuli, procedure, and analyses. The apparatus, stimuli and procedure were identical to those of Experiment 1, except for the following changes: the stop-associated images (10 images) were paired with a stop signal on 100% of trials during the training phase, and were never paired with a stop signal in the test phase; the go-associated images (30 images) were never paired with a stop signal in the training phase, but some of these images were paired with a stop signal in the test phase (20 old go-associated images were never paired with a stop signal; 10 old go-associated images were paired with a stop signal on 100% of the trials). The analyses were identical to those of Experiment 1, except that the contingencies meant that, for obvious reasons, I could not examine the effect of image type on go RTs or *p*(respond|stop) in the training phase of this experiment (see Table 2.1).

Results

Training phase. In the training phase, the RT for the go-associated images reliably decreased as a function of block (*p* = 0.038). This suggests that participants acquired the stimulus-go associations during the training. The *p*(respond|stop) for the stop-associated images did not reliably decrease as a function of practice (Figure 2.3, Table 2.3), which could be due to a floor effect.

Test phase. Contrary to the predictions of the automatic inhibition hypothesis, go RT was not influenced by image type in the test phase when the image-stop mappings were reversed (Table 2.4). As in Experiment 1, the absence of an effect in the test phase cannot be accounted for by the overall speeding of RTs (for RT distributions, see Appendix B).

Expectancy ratings. Despite the absence of an effect of image-stop learning in the test phase, expectancy ratings obtained following task completion revealed a main effect of image type: participants expected to withhold their response more for the stop-associated images (5.99) than for the go-associated images (3.86), $t(29) = -5.17$, $p < 0.001$, $g_{av} = 1.44$. This suggests that participants had learned the image-stop contingencies during training, even though these contingencies did not significantly influence performance in the test phase. The ‘stop minus go’ image expectancy difference did not significantly correlate with the RT difference in the test phase, $r(28) = 0.262$, $p = 0.162$ (see Figure C.2, Appendix C). Note that the ‘stop minus go’ expectancy difference was larger in Experiment 2 than in Experiment 1 (in which stop items could occur on 25% of go trials in the training phase), $t(49) = -2.47$, $p = 0.017$, Cohen’s $d = 0.64$. In other words, this between-experiment comparison indicates that the image-stop contingency (100% in Experiment 2 relative to 75% in Experiment 1) influenced expectancy ratings but it did not influence performance during the test phase.

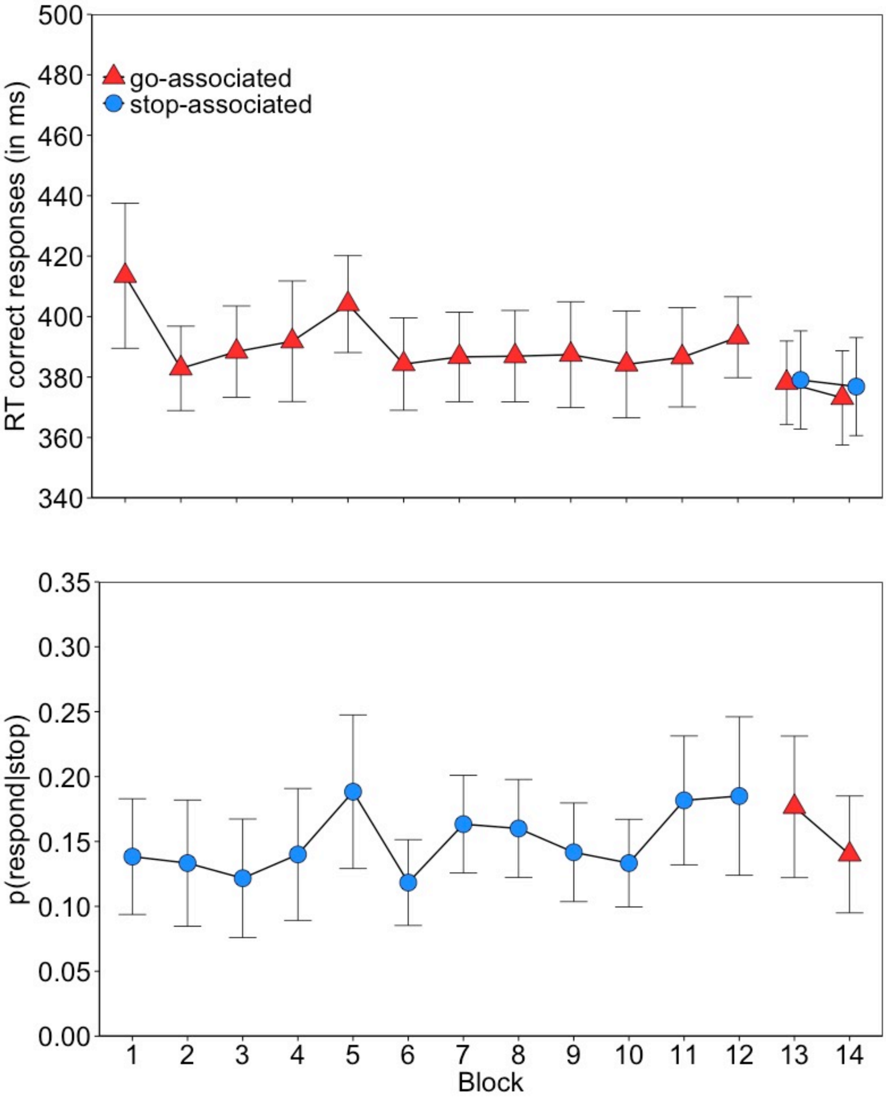

Figure 2.3. Reaction times for the correct go trials (upper panel) and *p*(respond|stop) data (lower panel) for the two image types (stop; go) as a function of the block (blocks 1-12 = training phase; blocks 13- 14 = test phase) in Experiment 2. Error bars are 95% confidence intervals.

Interim Discussion

In Experiment 2, I investigated whether the relative predictability of the stop-associated images influenced the extent to which the acquired stimulus-stop associations influenced task performance when these mappings were reversed. Therefore, the stop-associated images were paired with stopping on 100% of presentations during training (cf. 75% of presentations in Experiment 1).

Consistent with Experiment 1, the decrease in go RT for the go-associated images shows that participants acquired the stimulus associations during training (i.e. they associated the go-associated images with responding), and the expectancy ratings obtained following task completion show that participants expected to stop their responses more for the stop-associated images than for the go-associated images. Furthermore, these expectancy ratings were sensitive to the increased predictability of the stop-associated images as the expectancy difference between stop-associated and go-associated images was larger in Experiment 2 than in Experiment 1. However, as in Experiment 1, RTs were comparable for the old stop-associated images and the old go-associated images in the test phase, which indicates that the acquired associations did not influence performance in the test phase when the image-stop mappings reversed. On the face of it, these results do not support the conflict account of attentional modulation (e.g. Botvinick, Braver, et al., 2001). However, it is possible that participants quickly learned to ignore the images in the test phase when the mapping had reversed. Consistent with this idea, participants were slower to respond to the stop-associated images (382 ms) than to the go-associated images (376 ms) in the first half of block 13, but this was in the opposite direction in the second half of block 13 (stop-associated images: 374 ms, go-associated images: 380 ms; this reversal could be due to an increased error signal in the first half of the test phase). This suggests that participants may have quickly re-learnt the new mappings in the test phase[23]. Note that I did not conduct any inferential statistics on this difference due to low numbers of trials (≤ 20 trials per cell). An alternative possibility is that participants habituated to the images and stopped paying attention to them because the images were less novel. I tested the habituation hypothesis in Experiment 3.

Experiment 3

The aim of Experiment 2 was to investigate whether the relative predictiveness of the stop-associated images influenced the extent to which the stimulus-stop mappings acquired during training influenced task performance in the test phase. However, even though participants acquired the stimulus-stop mappings, these mappings did not modulate performance in the test phase. It is possible that this manipulation did not prevent participants 'tuning-out' attention to these images over practice because the images still became less novel. Therefore, in Experiment 3 I investigated whether stimulus exposure influenced the extent to which participants attended to the stop-associated images. To this end, I halved the number of stimulus presentations in the training phase, such that there were 12 presentations prior to the test phase (cf. 24 presentations in Experiments 1 & 2).

Method

Participants. Thirty-two students from the University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 19.19 years, *SD* = 1.49, 26 females, 29 right-handed). One participant was excluded because they incorrectly executed a response on ≥ 30% of stop trials.

---

[23] More generally, the idea that participants could rapidly re-learn the new mappings is also consistent with the categorisation literature; when people are learning a rule-based category structure and the stimulus-response mappings are switched, participants can use rules to reset the mappings with minimal impairment (see, e.g. Ashby, Ell, & Waldron, 2003).

Apparatus, stimuli, procedure, and analyses. The apparatus, stimuli and procedure were identical to those of Experiments 1 and 2, except for the following changes: each image was presented once per block (i.e. 14 presentations in total). To ensure that the overall *p*(stop) was the same as in Experiments 1 and 2, the reduced number of image presentations meant that the stimulus-stop contingencies for the go and the control images in the test phase had to be altered (for the specific contingencies, see Table 2.1). As in Experiment 1, the stop-associated images were paired with a stop signal on 75% of presentations during the training phase to provide an index of image-stop learning during training. For comparison with Experiments 1 and 2, in the analyses the blocks were collapsed to ensure that the number of observations per cell was comparable.

Results

Training phase. In the training phase, the main effect of image type on go RTs was marginally significant ($p = 0.058$); planned comparisons revealed marginally significant differences between the stop-associated images (428 ms) and the go-associated images (422 ms), $t(30) = -1.99$, $p = 0.055$, $g_{av} = 0.23$, and between the stop-associated images and the control images (422 ms), $t(30) = -1.92$, $p = 0.064$, $g_{av} = 0.24$. There was no reliable difference between the control and the go-associated images, $t(30) = 0.28$, $p = 0.777$, $g_{av} = 0.02$**.** However, Figure 2.4 shows that RTs were longer for the stop-associated images than for the control and the go-associated images in blocks 1-3, but this difference disappeared from block 4 onwards. This conclusion was supported by a reliable interaction between image type and block ($p = 0.005$). The overall main effect of block was reliable, suggesting that participants improved as a function of task practice ($p < 0.001$). There were no reliable differences in *p*(respond|stop).

Test phase. As in Experiments 1 and 2, there was no main effect of image type on go RT in the test phase ($p = 0.479$; Table 2.4). However, the difference in *p*(respond|stop) between the go-associated images (0.18) and the control images (0.13) was marginally significant, $p = 0.062$, suggesting that the image-go associations did influence test phase performance to some extent (Table 2.4).

Expectancy ratings. Consistent with the previous experiments, image type influenced expectancy ratings, $F(2, 60) = 11.44$, $p < 0.001$, *gen.* $\eta^2 = 0.136$. Expectancy ratings were greater for the stop-associated images (5.54) than for the go-associated images (4.57), $t(30) = -3.50$, $p = 0.001$, $g_{av} = 0.85$, and the control images (4.76), $t(30) = -3.44$, $p = 0.001$, $g_{av} = 0.69$. There was no reliable difference between the control images and the go-associated images, $t(30) = -1.84$, $p = 0.075$, $g_{av} = 0.20$. However, the expectancy differences did not correlate with the corresponding RT differences (*r*s ≤ 0.136, *ps* ≥ 0.464; see Figure C.3, Appendix C).

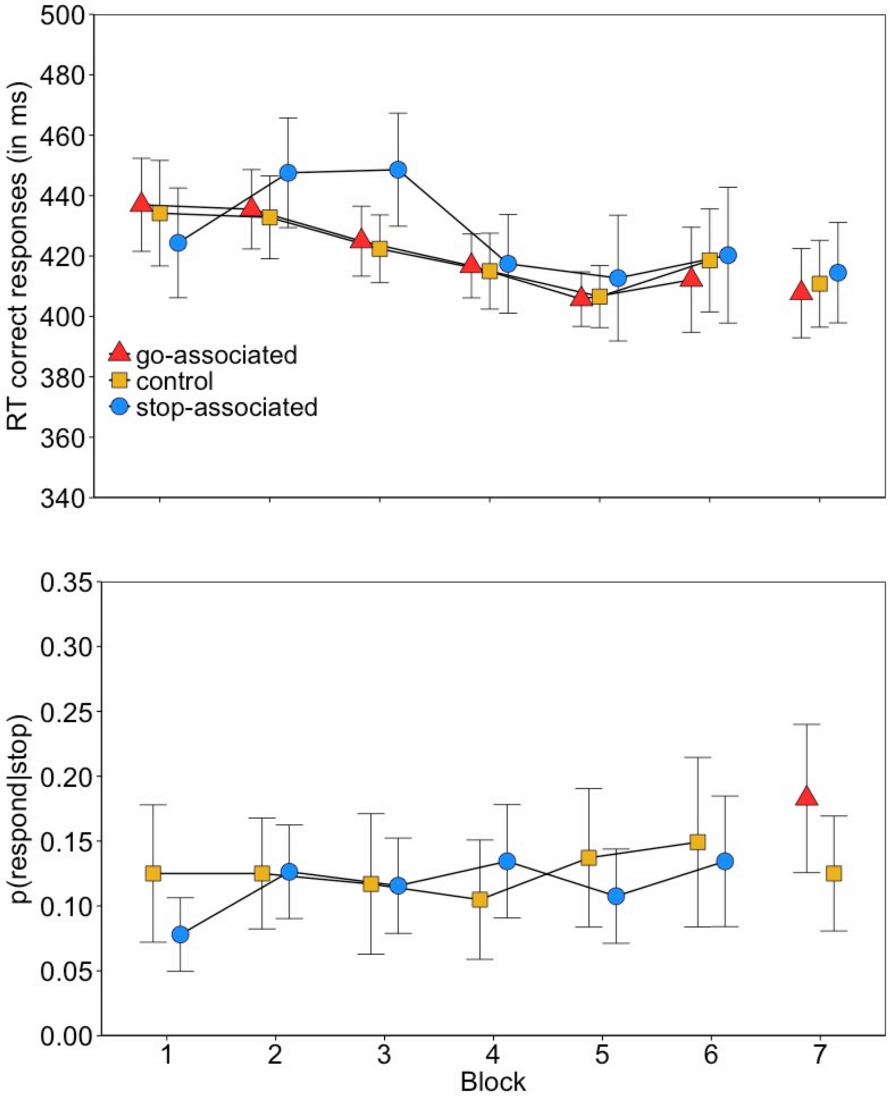

Figure 2.4. Reaction times for the correct go trials (upper panel) and *p*(respond|stop) data (lower panel) for the three image types (stop; go; control) as a function of the block (blocks 1-6 = training phase; block 7 = test phase) in Experiment 3. Error bars are 95% confidence intervals.Figure 2.4. Reaction times for the correct go trials (upper panel) and p(respond|stop) data (lower panel) for the three image types (stop; go; control) as a function of the block (blocks 1-6 = training phase; block 7 = test phase) in Experiment 3. Error bars are 95% confidence intervals.

Interim Discussion

In Experiment 3, I investigated whether the amount of exposure to the stop-associated images influenced the extent to which the stimulus-stop mappings acquired during training affected task performance in the test phase (when the stimulus-stop mappings were reversed).

Consistent with Experiment 1-2, the results of Experiment 3 indicate that participants acquired the stimulus-stop mappings during training; participants were slower to respond to the stop-associated images than to the go images and the control images. However, this effect appeared and then disappeared again throughout practice; this conclusion was supported by a significant interaction between block and image type. This is consistent with the (numerically) diminished learning effect observed at the end of the training phase in Experiment 1. Furthermore, participants were not slower to respond to the stop-associated images than to the go-associated images and to the control images in the test phase (although I observed a marginally significant difference between go and control images). This suggests that the amount of habituation to the images cannot entirely account for the absence of the test phase effect. This leaves an associability mechanism controlling attention to the stimuli as the most plausible explanation for the results of the experiments so far.

As in Experiments 1-2, I find clear evidence that participants acquired the stimulus-stop contingencies in the expectancy ratings obtained following task completion; participants expected to stop their response more for the stop-associated images than for the go-associated images and the control images. This suggests that participants did not forget the stimulus-stop contingencies, despite the disappearance of the learning effect on task performance towards the end of the training phase and during the test phase.

Experiment 4

In the final experiment of this chapter, I presented the image before the go and stop signals, and asked participants to rate whether they expected to stop or not. Furthermore, I presented the go and stop signals around the image, at one of four possible locations (one of four corners of the image; for a similar procedure see Houben & Jansen, 2011). These manipulations served two purposes. First, the results of Experiments 1-3 suggested that participants stopped paying attention to the task-irrelevant images. I tried to increase attention to the images by making them perfect predictors of the outcome (Experiment 2) or by decreasing image habituation (Experiment 3). These manipulations were only moderately effective: some behavioural indices indicate that the manipulation influenced learning, but the effect of learning on test performance still disappeared over training. By presenting the images before the go and stop signals, and asking participants to rate their stop expectancy, participants were less likely to ignore the images in Experiment 4 (however, subjects were not explicitly informed to attend to the images so as to keep the image-stop mappings uninstructed as in Experiments 1-3). Furthermore, the images initially did not have the stop signal present as a competitor driving their associability down. If the proposed attentional account is correct, I should observe the effects of stop training in the later blocks of the training phase and in the test phase. Second, in Experiments 1-3, I found that participants generated expectancies based on the image-stop associations acquired during training. In Experiment 1, expectancy correlated with some aspects of performance in the test phase, but I could not replicate this finding in Experiments 2-3. It is possible that obtaining the expectancy ratings following task performance meant that these expectancies were contaminated by the re-learning of the new (inconsistent) mappings in the test phase. Therefore, in Experiment 4, I further investigated the role of expectancy in stimulus-stop learning by obtaining expectancy ratings during task performance (for a similar procedure, see e.g. McAndrew, Jones, McLaren, & McLaren, 2012; Perruchet, Cleeremans, et al., 2006).

Method

Participants. Thirty-two students from the University of Exeter participated for partial course credit (*M* = 18.47 years, *SD* = 0.62 years, 27 females, 31 right-handed). Four participants were excluded because they incorrectly executed a response on ≥ 30% of stop trials.

Apparatus, stimuli, procedure, and analyses. The apparatus, stimuli and procedure were identical to those of Experiment 3, except for the following changes: All trials began with the presentation of the image in the centre of the screen. The word 'RATING' was presented above and below the image to instruct participants to rate 'how much do you expect to withhold your response?'. Participants inputted their ratings on a scale between 1 ('I definitely do not think that I will have to withhold my response') and 9 ('I definitely think that I will have to withhold my response') using the number keys of the keyboard with their right index finger (latency rating response: *M* = 969 ms; *SD* = 681 ms). After participants made their expectancy rating, a go/stop signal appeared at one of four locations on the screen (top-left, bottom-left, top-right, or bottom-right corner of the image). The delay between the expectancy response and the presentation of the go/stop signals varied randomly between 500 and 1250 ms. Participants responded on go trials by pressing the spacebar on a keyboard with their left index finger. To allow for the presentation of the signals at each location on the screen, task rules used in Experiments 1-3 that were based on signal location (e.g. 'X on the left/right of the image') or signal shape (e.g. 'shape bigger/smaller than a fifty pence piece') were excluded and, of the remaining rules, seven rules were selected on the basis of response latencies in Experiments 1-3 using a non-parametric box and whisker method (Tukey, 1977). A full list of the signals used appears in Appendix A. The expectancy ratings data in the training and test phase trials were analysed separately using ANOVAs with image type and block as within-subjects factors.

Results

Training phase. In the training phase, responding became numerically slower for the stop-associated images than for the go-associated images and the control images in the second half of the training phase (see Figure 2.5). However, the interaction between image type and block did not reach significance ($p = 0.094$) and there was also no reliable main effect of image type ($p = 0.153$). Nevertheless, analyses of the RT percentiles in the training phase showed that responding was slower for the stop-associated images than for the go-associated and control images in the fastest RTs (Appendix B). This was supported by a reliable two-way interaction between image type and percentile in the training phase ($p < 0.001$; Appendix B). Furthermore, the overall main effect of block was reliable suggesting that performance improved throughout the training phase ($p < 0.001$). Analyses of performance on stop trials showed that the p(respond|stop) was reliably lower for the stop-associated images (0.15) than for the control images (0.19) ($p = 0.011$). The interaction between image type and block in the p(respond|stop) was not reliable. However, visual inspection of Figure 2.5 shows that this difference began

The analysis of the online expectancy ratings also revealed a reliable image type by block interaction ($p = 0.005$), reflecting higher stopping expectancies for the stop-associated images in the second half of the training phase (blocks 4-6; see Figure 2.5). There was also a reliable main effect of block on the expectancy ratings ($p = 0.012$): overall mean expectancy ratings decreased with task practice, which is consistent with the overall p(stop) of 0.25 (note, the increase in expectancy ratings across block for the stop-associated images was not reliable, $p = 0.261$). Combined, these findings indicate that participants were generating appropriate expectancies during the acquisition of the stimulus-stop mappings. Importantly, the overall ‘stop minus go’ expectancy ratings difference reliably correlated with the corresponding RT difference in the training phase, $r(26) = 0.575$, $p = 0.001$; the overall ‘stop minus control’ expectancy ratings difference also correlated with the corresponding RT difference, $r(26) = 0.498$, $p = 0.006$. Taken together, these findings provide some support for the idea that participants acquired the stimulus-stop associations in the training phase.

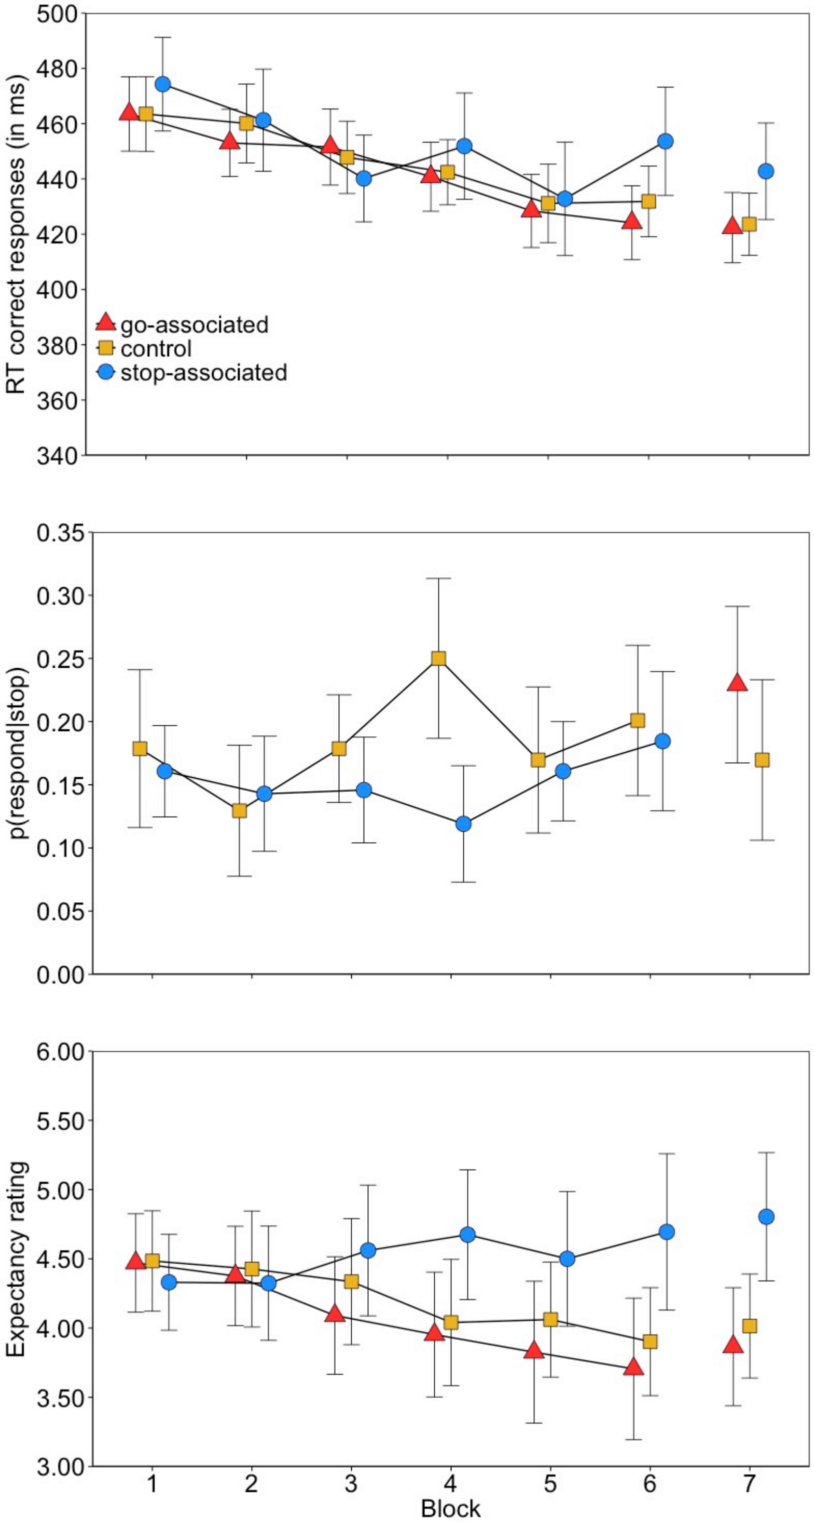

Figure 2.5. Reaction times for the correct go trials (upper panel), *p*(respond|stop) data (middle panel), and expectancy ratings (lower panel) for the three image types (stop; go; control) as a function of the block (blocks 1-6 = training phase; block 7 = test phase) in Experiment 4. Error bars are 95% confidence intervals.

Test phase. Unlike in Experiments 1-3, I found a main effect of image type on go RTs in the test phase ($p = 0.004$). Planned comparisons revealed that responding to the old stop-associated images was slower (443 ms) than to the go-associated images (422 ms), $t(27) = -2.84$, $p = 0.008$, $g_{av} = 0.52$, and to the control images (424 ms), $t(27) = -2.87$, $p = 0.007$, $g_{av} = 0.54$. There was no reliable difference between the go and control images, $t(27) = -0.31$, $p = 0.756$, $g_{av} = 0.04$, (Figure 2.5; Table 2.3). Image type did not reliably influence p(respond|stop) in the test phase (however, the means were in the predicted direction, see Figure 2.5; Table 2.4).

There was also a reliable main effect of image type on test phase expectancies ($p = 0.002$); planned comparisons revealed that participants expected to stop more for the old stop-associated images (4.80) than for the go-associated images (3.86), $t(27) = -2.65$, $p = 0.013$, $g_{av} = 0.81$, and the control images (4.01), $t(27) = -2.83$, $p = 0.008$, $g_{av} = 0.72$. There was no reliable difference between the go-associated and the control images, $t(27) = -1.37$, $p = 0.181$, $g_{av} = 0.14$. As in the training phase, I found that the 'stop minus go' expectancy ratings difference reliably correlated with the corresponding RT difference, $r(26) = 0.624$, $p < 0.001$; the 'stop minus control' expectancy ratings difference also correlated with the corresponding RT difference, $r(26) = 0.653$, $p < 0.001$ (see Figure C.4, Appendix C). Hence, participants who had a stronger expectancy to stop their response when the stop-associated images were presented displayed greater response slowing for these images than for the go-associated images and for the control images upon signal presentation.

To further investigate to what extent the expectancy to stop determined response slowing for the stop-associated images, I conducted an exploratory median-split analysis on the expectancy ratings of the test phase (I could not perform a similar analysis in the training phase because there were not enough trials in each block). I calculated the median for each image type and participant separately. Ratings greater than the median were classified as a 'stop' expectancy whereas ratings less than or equal to the median were classified as a 'go' expectancy. Four participants were excluded from these analyses as they always entered the same expectancy rating for one or more of the image types (consequently, I could not perform a median split). I analysed the data with a 2 (expectancy: stop vs. go) by 3 (image type) ANOVA. Consistent with previous work on proactive control (see e.g. Verbruggen & Logan, 2009a), responding was slower for trials on which participants expected a stop signal (445 ms) compared with trials on which participants expected a go signal (420 ms), $F(1, 23) = 13.96$, $p = 0.001$, *gen.* $\eta^2 = 0.088$. As discussed above, image type also had a reliable main effect on performance. Importantly, the effects of stimulus-stop learning and expectancy were additive; i.e. the two-way interaction between expectancy and image type was not reliable, $F(2, 46) = 0.08$, $p = 0.915$, *gen.* $\eta^2 < 0.001$ (for descriptive statistics, see Table 2.5). Thus, the slowing for the stop-associated images is unlikely to reflect an entirely strategic, expectancy-driven effect.

	Stop expectancy		Go expectancy	
	M	***SD***	***M***	***SD***
Stop-associated	453	56	429	51
Go-associated	437	43	411	36
Control	440	40	411	33

Table 2.5. Go reaction times (in ms) in the test phase as a function of expectancy (go, stop) and image type (stop-associated, go-associated, control) in Experiment 4. *M* = mean; *SD* = standard deviation.

Interim Discussion

Consistent with the results of Experiments 1-3, I find evidence that participants acquired the stimulus-stop associations. In the training phase, responding became slower for the stop-associated images than for the go-associated images and the control images with task practice, and the *p*(respond|stop) was lower for the stop-associated images than for the control images. In addition, the expectancy ratings showed that participants generated expectancies that were consistent with the trained stimulus-stop contingencies in the second half of the training phase. These expectancies correlated with task performance in the training phase: participants who expected to withhold their response more to the stop-associated images responded more slowly to these images than to the go-associated images and to the control images during training. Unlike in Experiments 1-3, I also find that learning influenced performance in the test phase: participants were slower to respond to the stop-associated images than to the go-associated images and the control images during the test phase.

These results suggest that presenting the images before the go/stop signals and asking participants to rate their expectancy on each trial increased the extent to which participants attended to these images. In order to ensure that attention to the task-irrelevant images was maximised, I combined these manipulations in the same procedure. Consequently, it was not possible to determine the relative contributions of these manipulations to the observed slowing for the stop-associated images in the test phase. One could speculate that the observed slowing reflects an *entirely* strategic, expectancy-driven effect, rather than the implicit retrieval of the acquired stimulus-stop associations (as predicted by the automatic inhibition account). I argue that an *entirely* explicit explanation is unlikely for several reasons. First, previous studies have demonstrated stop-learning effects using procedures in which the stop-associated stimuli are presented prior to stop-signal onset but, unlike the present experiment, without expectancy ratings on each trial. For example, in a recent study the stop-associated stimuli were presented as 'warning cues' for a variable duration prior to the presentation of the stop signal, and observed stop learning effects during the training and test phases (Bowditch, Verbruggen, et al., 2015). Similarly, Veling and colleagues have conducted two experiments using go/no-go designs in which food images were presented 100 ms (Veling, van Koningsbruggen, et al., 2014) or 500 ms (Veling, Aarts, et al., 2013a) prior to the onset of the go/no-go signal. They found that when the food images were consistently presented on no-go trials, subsequent choice of the food items was reduced (Veling, Aarts & Stroebe, 2013a) and weight loss was facilitated (Veling, van Koningsbruggen, Aarts & Stroebe, 2014).

Second, the median split analysis on expectancy ratings in the test phase shows that the slowing for the stop-associated images occurred even when stop signal expectancy was relatively low. This result suggests that expectancy ratings cannot account for the whole data pattern. However, it is important to note that whilst the median split analyses provided some insight into the present data, median split analyses can have limitation (e.g. a loss of statistical power and a reduction in effect sizes; Moreau, Kirk, & Waldie, 2016). Furthermore, it might be possible to argue that the very act of asking participants to provide an expectancy ratings on each trial could have had extraneous effects on how they approached the task. Therefore, in Chapter 3 I followed-up on these findings by investigating the role of expectancies using a methodology that did not rely on median split analyses or trial-by-trial expectancy ratings (the results of Chapter 3 are consistent with the results reported here).

Finally, research in the wider action control literature is consistent with the pattern of findings in the present chapter. For example, Frings and Moeller (2012) found that associations between old distractor stimuli and the previously required target response only interfered with responding when the distractors were presented prior to the target stimuli. When combined with the findings of these studies, the present findings strongly suggest that presenting the task-irrelevant image before the go or no-go signal increases attention to the images, and consequently, the probability that the image-stop association are retrieved.

To conclude, the presence of a learning effect in the test phase is consistent with the proposed hypothesis that attention to the images determines whether acquired stimulus-stop associations influence behaviour in the test phase. Now that the images are more task-relevant and associability is no longer driven down for the images by virtue of their competition for attention with the stop signal, there is a strong effect on test phase go RTs. Furthermore, the test-phase expectancy ratings show that participants continued to generate expectancies consistent with the image-stop mappings acquired during training, despite the reversal of these mappings. As in the training phase, these expectancies reliably correlated with task performance: participants who expected to withhold their response more for the stop-associated images responded more slowly to these images than to the go-associated images and to the control images in the test phase. However, the median split also suggested a contribution of implicit (non-expectancy related) processes.

Chapter Discussion

In the present chapter, I investigated three outstanding issues relating to the mechanisms of stimulus-stop learning. The first two issues were highlighted in Chapter 1: (1) are associations between stimuli and stopping direct, and (2) what is the role of expectancy in stimulus-stop learning? Based on the results of Experiment 1, Experiments 2-4 also investigated a third issue: (3) does attention to the stop items affect the extent to which stimulus-stop learning influences behaviour? Based on the present findings, I can answer each of these questions.

Are Associations Between Stimuli and Stopping Direct?

Across four experiments where the specific stop signals and rules were always changing, I provide strong evidence for the idea that participants can learn direct stimulus-stop associations (Verbruggen & Logan, 2008a). During training, I found that responding was reliably slower (Experiments 1, 3, and numerically slower in Experiment 4; in Experiment 2, I could not compare stop and go-associated images in the training phase) and the *p*(respond|stop) was lower (Experiment 1 & Experiment 4) for images that were consistently associated with stopping than for images associated with going and for control images that were not particularly associated with stopping or going.

In recent experiments, it was observed that learning can influence the *p*(respond|stop) but not response latencies on go trials (see e.g. Experiment 2 in Verbruggen, Best, et al., 2014). Based on previous findings in the conditioning literature (for a review, see Hall, 2002), Verbruggen, Best, et al. (2014) hypothesised that participants in these experiments learned an association between an item and a representation of a no-go or stop signal. Hence, when the item was repeated, it primed the signal so that it was detected sooner on stop-signal trials, resulting in improved response inhibition and, consequently, a lower *p*(respond|stop). The signal priming idea explains why it can be that learning influences the probability of stopping on signal trials without influencing response latencies on go trials. In the present chapter, both RTs and *p*(respond|stop) were influenced even though the go/stop signals and task rules constantly changed (and there were two go signals and two stop signals in each block). This indicates that learning was not (solely) mediated via image-signal associations. The most parsimonious account is that the effects in the present chapter reflect the direct association of the stop-associated images with a stop goal, rather than the association of the stop-associated images with the representation of a single stop signal. Therefore, the present chapter provides strong support for the original hypothesis outlined by Verbruggen and Logan (2008a) that participants can acquire direct (i.e. un-mediated) associations between specific stimuli and the stop goal. In situations where the task rules do not constantly change, it is likely that individuals will acquire both stimulus-goal and stimulus-signal associations (indeed, research in the conditioning literature suggests that the acquisition of multiple associations is the norm; Hall, 2002). It is possible that experimental factors, such as the perceptual properties of the stop signal, will influence which association dominates behaviour. In Chapter 4, I investigate the role of signal detection during stop learning.

It is important to note that the learning effects demonstrated in the present chapter are assumed to reflect the acquisition of stimulus-stop associations rather than the absence of stimulus-go learning on stop trials. Whilst the ‘absence of go learning’ explanation may initially seem parsimonious, it cannot account for several findings previously reported in the stop-learning literature. First, Verbruggen and Logan (2008a) have previously demonstrated that responding to old stop items is slowed compared with novel items that were not presented during training (hence, these items were not associated with going or stopping; Verbruggen & Logan, 2008a, Experiment 1). Second, neuroimaging work has shown that the presentation of old stop items activates the neural inhibitory control network (Lenartowicz, Verbruggen, et al., 2011; but see *What is the Role of Expectancy in Stimulus-Stop Learning?* below). Third, brain stimulation studies have shown that even when the probability of go and no-go signals is equal (i.e. 50/50), motor-evoked potentials are below baseline 200-300 ms following no-go stimulus presentation (indicating that responding is suppressed; Leocani, Cohen & Wassermann, 2000). In other words, successful performance on a no-go trial requires the activation of a no-go or stop response, and not just the absence of a go response. Fourth, short-term after-effects of stopping further support the idea that participants can learn stimulus-stop associations that can have a (global) inhibitory effect on responding (Giesen & Rothermund, 2014; for an extended discussion of this experiment, see *Sequential after-effects,* Chapter 1). Finally, in the present chapter, response latencies decrease for go and control images but I observe an initial increase in response latencies for stop-associated images over practice (Experiment 1). In Experiments 3, this conclusion is further supported by a reliable interaction between image type and block. Finally, the comparison of expectancy ratings in Experiments 1 and 2 revealed that expectancy ratings were altered when the image-stop consistencies had changed (even though the image-go contingencies did not change). Therefore, previous results and the findings reported in the present chapter are consistent with the idea that participants can learn go associations on go trials and stop associations on stop trials (which interfere with responding).

What is the Role of Expectancy in Stimulus-Stop Learning?

In the present chapter, I show that participants generated expectancies that were consistent with the stimulus-stop mappings acquired during training: participants expected to withhold their responses more when stop-associated images were presented than when go and control images were presented. Furthermore, these expectancy ratings were sensitive to the specific contingencies in play: participants expected to withhold their responses more for the stop-associated images that were reinforced on 100% of presentations (Experiment 2) than for the stop-associated images that were reinforced on 75% of presentations (Experiment 1). Finally, I found that these expectancies correlated with task performance both during the acquisition of the stimulus-stop mappings in the training phase (Experiment 4) and following the reversal of these mappings in the test phase (Experiment 1 & Experiment 4).

The role of expectancies in stimulus-stop learning has not been previously investigated. Therefore, the present chapter provides initial evidence that stimulus-stop learning could be partly mediated via explicit knowledge of the stimulus-stop contingencies in play. Although the median split analysis and the absence of significant correlations in some of the experiments indicate that implicit processes could play a role as well. However, it is not possible to conclusively attribute a role of implicit processes as the absence of reliable correlations could be due to issues of sensitivity because I assessed expectancies at the end of task performance in those experiments. If expectancies did play a role, this could indicate that the response slowing observed for the stop-associated images is caused by top-down control processes. First, the slowing could be partly due to proactive control. According to this proactive control account of stimulus-stop learning, stop items could become predictive cues (e.g. if image X then p(stop) is high) that indicate that participants should adjust their response strategies accordingly. If this were the case, this would suggest that earlier findings that have demonstrated response slowing and neural activation of the inhibitory control network by old stop items (Lenartowicz, Verbruggen, et al., 2011) could be due to proactive control (i.e. another algorithmic process), rather than the direct activation of the stop response via memory-retrieval (i.e. picture X = stop). Therefore, whilst the *retrieval* of the stimulus-stop association may still be automatic, the subsequent *slowing* observed following the reversal of the stimulus-stop mapping would be due to a top-down control process (rather than a bottom-up process as is currently assumed). The role of proactive control during stop learning is investigated in Chapter 3. Second, stop items could effectively become a new stop signal (the direct stopping account). In other words, the only difference between the stop items and an external stop signal is that the association with stopping is acquired via learning in the case of the stop items, whereas it is acquired via instructions in the case of the stop signal. Thus, in both cases, response inhibition is a deliberate act of control. But the advantage of the former form of control is that the go and stop processes in stop-signal tasks could be initiated simultaneously and, therefore, start the race at the same time (Logan & Cowan, 1984); consequently, response inhibition is more likely to succeed.

It is important to note, however, that the proactive control route and the direct stopping route are both compatible with the idea that associative learning plays a key role in response inhibition paradigms; indeed, both accounts still assume that stimulus-specific learning influences stop performance. Learning offers participants another route to control their behaviour. The key difference between these two top-down accounts and the ‘automatic’ inhibition account is the nature of the process that occurs following the retrieval of the stimulus-stop association; either this association directly activates the stop goal via an stimulus-response based link (in the automatic stopping account) or this association indirectly activates the stop goal via a top-down (algorithmic and deliberate) control process. In Chapter 3, I directly compare the proactive control and automatic inhibition accounts.

Does Attention to the Stop Items Affect the Extent to Which Stimulus-Stop Learning Influences Behaviour?

In Experiments 1-3, the acquired stimulus-stop associations did not influence performance in the test phase, despite effects of learning on task performance in the training phase and on expectancies following task completion (suggesting that participants had not forgotten the stimulus-stop associations).

A potential explanation for this finding is that the images used in the present chapter were task-irrelevant so participants may have begun to ignore the images as they became less novel and as they learned that they were less predictive. In Experiments 1-3, task performance did not require participants to attend to the stop-associated stimuli (unlike in previous work; see e.g. Verbruggen & Logan, 2008a), so participants may have started ignoring all the images over time. In line with this possibility, the effect of image type reliably interacted with block (in Experiment 3) and visual inspection of the data shows that the influence of image-stop learning on performance began to disappear at the end of the training phase (Experiment 1). Since there were no differences between the image types in the final block of the training phase, this may explain why I did not find any effect of image-stop learning in the test phase[24]. Several associative learning accounts suggest that the reduced predictiveness of the images relative to the go/stop signals (in Experiments 1 & 3) may have decreased the extent to which they were considered informative or salient and, consequently, the extent to which participants attended to them and the extent to which they can influence performance (Mackintosh, 1975). Effects that point to this conclusion have been previously observed in animals (see e.g. Sutherland & Mackintosh, 1971 for a review of this literature) and, importantly, also in humans (Le Pelley & McLaren, 2003; Livesey & McLaren, 2007; Suret & McLaren, 2005). For example, Le Pelley and McLaren (2003) showed that foods that were worse predictors of an outcome than other foods present on a trial in an allergy discrimination task became less salient, resulting in slower learning of a new association to these stimuli in a later training phase (cf. learned irrelevance; Mackintosh, 1975). Note that the majority of the results in Experiments 1-3 are also consistent with conflict monitoring accounts (e.g. Botvinick, Braver, et al., 2001), which predict that participants will

---

[24] Note that I can rule out the possibility that the absence of a test phase effect in Experiments 1-3 is due to the use of images or the frequent rule switching; in a word version of the go/stop task with a single rule, it was similarly found that reversing the word-go/stop mapping in the test phase did not influence performance either (For experimental details see Best, Lawrence, Logan, McLaren, & Verbruggen, 2016, Appendix B).

ignore task-irrelevant information that produces response conflict or choice errors. However, unlike the associative learning accounts, these conflict monitoring accounts do not easily explain the absence of a learning effect in the test phase found in Experiment 2 when conflict should have been minimised by the use of 100% contingencies.

Importantly, I found a clear effect of stimulus-stop learning on test phase performance when attention to the images was increased in Experiment 4 (as a result of presenting them before the go/stop signals and the requirement to make an online expectancy rating on each trial). This finding is consistent with the Instance theory (Logan, 1988; Logan & Etherton, 1994) and other theories of associative learning. For example, Instance theory suggests that processing episodes will only be stored and retrieved from memory when participants attend to each stimulus presentation (Logan, 1988; Logan & Etherton, 1994). Thus, by encouraging subjects to attend to the image in Experiment 4, the image-stop associations were more likely to be retrieved, and performance was influenced in the test phase. Therefore, the present chapter strongly indicates that the influence of image-stop learning on behaviour is likely to be determined by the interplay of both attentional control and associative learning systems (see also Logan, 1988; Verbruggen, McLaren, et al., 2014).

Conclusion

In sum, the present chapter indicates that participants can learn direct associations between stimuli and a stop goal when the go/stop rule changes at the beginning of each block. Exposure to the image-stop associations influenced task performance during training, and expectancies following task completion. However these results also suggest that attention to stimulus attributes is key for retrieval of processing episodes; if participants do not attend to the stop stimulus then the previously acquired stimulus-stop associations will not influence behaviour. The results presented in this chapter are consistent with the Instance theory and other attentional accounts of associative learning.

Appendix A

A full list of the rules and the stimuli used in Experiments 1-4

In Experiments 1-3 I used fourteen go/stop rules. In Experiment 4, I used rules 1-7 only. The signals used in rules 1-11 were presented in Arial (font size = 50). The sizes of the signals used in rules 12-14 are provided in pixels below (screen resolution: 1680 × 1050).

1. Vowels ('a' or 'e') vs. consonants ('t' or 'n').
2. Symbols that are the same ('@@' or '&&') vs. symbols that are different ('@&' or '&@').
3. Uppercase letters ('H' or 'R') vs. lowercase letters ('h' or 'r').
4. Long symbol strings ('£%£%' or '%£%£') vs. short symbol strings ('£%' or '%£').
5. Curved letters ('S' or 'C') vs. angled letters ('K' or 'W').
6. Digits smaller than 5 ('2' or '4') vs. digits bigger than 5 ('6' or '8').
7. Curly brackets ('{' or '}') vs. square brackets ('[' or ']').
8. Words that refer to animals ('horse' or 'sheep') vs. words that refer to fruit ('lemon' or 'apple').
9. Symmetric letter strings ('UYYU' or 'YUUY') vs. asymmetric letter strings ('YYUU' or 'UUYY').
10. Crosses on the left of image vs. crosses on the right of the image relative to the centre (crosses appeared at the top and bottom of the image).
11. Asterisks on the top of the image vs. asterisks on the bottom of the image relative to the centre (asterisks appeared on the left and right of the image).
12. Horizontal lines (lines appeared across the top or bottom of the image relative to the centre) [width: 240 pixels] vs. vertical lines (lines appeared along the left or right of the image relative to the centre) [height: 240 pixels].
13. Shapes bigger than a fifty pence piece (square or circle) [100 × 100 pixels] vs. shapes smaller than a fifty pence piece (square or circle) [40 × 40 pixels].
14. Lines thicker than a matchstick vs. lines thinner than a matchstick (lines appeared horizontally [width = 240 pixels] or vertically [height = 240 pixels] about the centre of the image).

Appendix B

RT Percentiles

To investigate the possibility that the absence of an effect of image type in the test phase of Experiments 1-3 is due to response latencies (responding was faster in the test phase than in the training phase) I plotted RT percentiles for the training and test phases. These RT percentiles revealed that the overall response latency cannot account for the absence of a learning effect in the test phase.

Furthermore, in Experiments 1-3, visual inspection of the percentile plots suggests that the slowing for the stop-associated images emerges in the slow end of the RT distribution. This conclusion is supported by a reliable two-way interaction between image type (stop; go; control) and percentile in the training phase of Experiment 1, $F(16, 448) = 23.12$, $p < 0.001$, $gen.\ \eta^2 = 0.030$. However, in Experiment 4, the slowing for the stop-associated images emerges in the fast end of the RT distribution. This conclusion is also supported by a reliable two-way interaction between image type and percentile in the training phase, $F(16, 432) = 56.50$, $p < 0.001$, $gen.\ \eta^2 = 0.097$[25].

In Experiments 1-3, processing the image could slow overall RT; but for stop-associated items, processing the image would also lead to retrieval of the stop associations, and consequently, automatic inhibition of the response. Alternatively, only on slower trials, the stimulus-stop associations could be retrieved in time and affect performance. In Experiment 4, attention to the images prior to signal presentation meant that there was more time for the acquired stimulus-stop associations to be retrieved and thus influence performance.

---

[25] Note that for consistency with the percentile analyses reported later in this thesis (see Appendix M), the 'percentile' factor comprised the 10th-90th percentiles (as shown in Figures B.1-B.4). However, to reduce the number of levels in the percentile factor, this analysis was previously presented in Best, Lawrence, et al. (2016) with only the 25th, 50th, and 75th percentiles included. However, it is important to note that there were no meaningful differences between the percentile findings presented here and those presented by Best, Lawrence, et al. (2016).

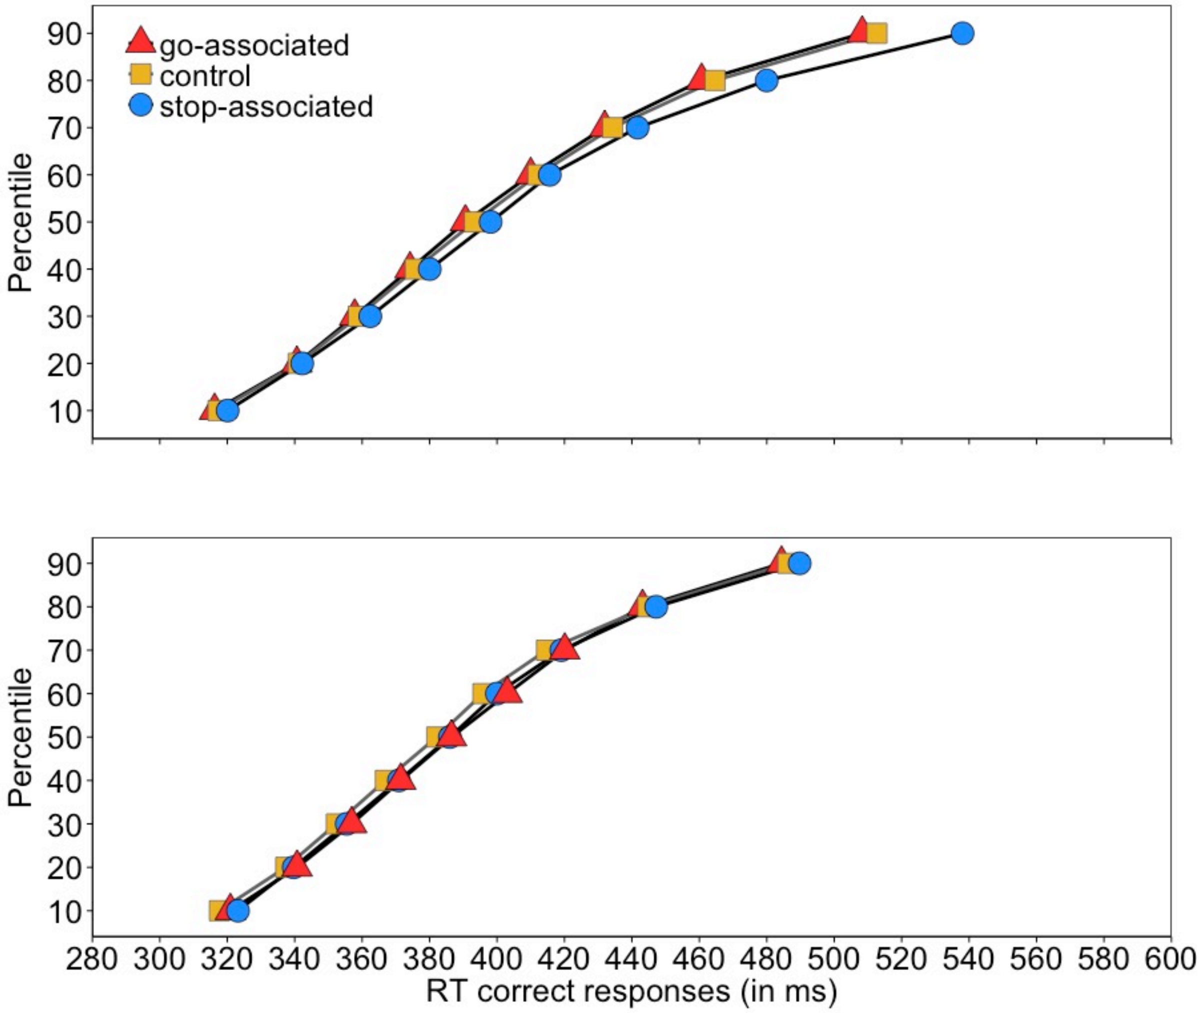

Figure B.1. Go RTs (in ms) in the training phase (blocks 1-12; upper panel) and the test phase (blocks 13-14; lower panel) for the three image types (stop-associated; go-associated; control) as a function of percentile in Experiment 1.

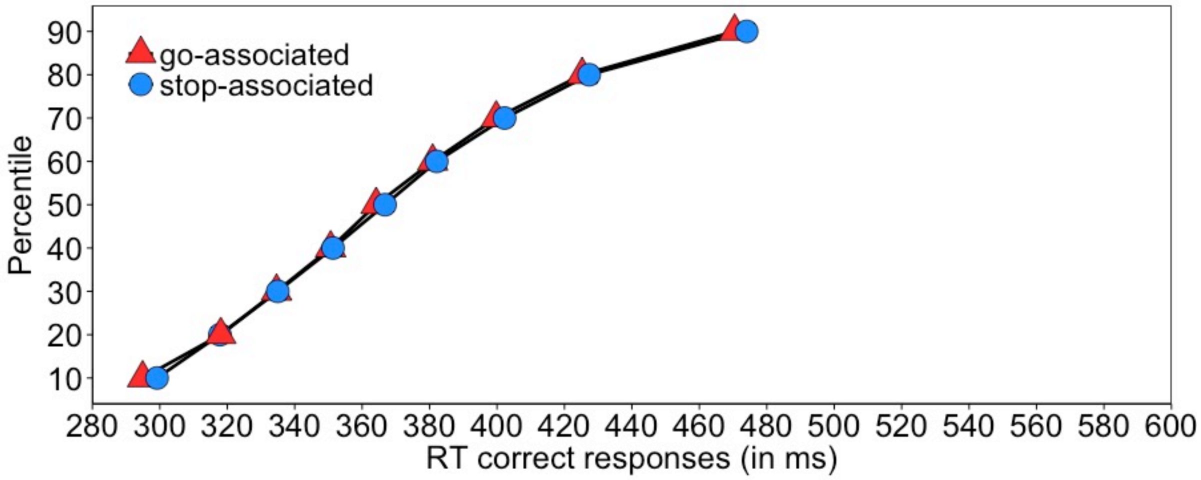

Figure B.2. Go RTs (in ms) in the test phase (blocks 13-14) for the two image types (stop-associated; go-associated) as a function of percentile in Experiment 2. Due to the stimulus-stop contingencies used, I could not plot RT percentiles for the training phase data.

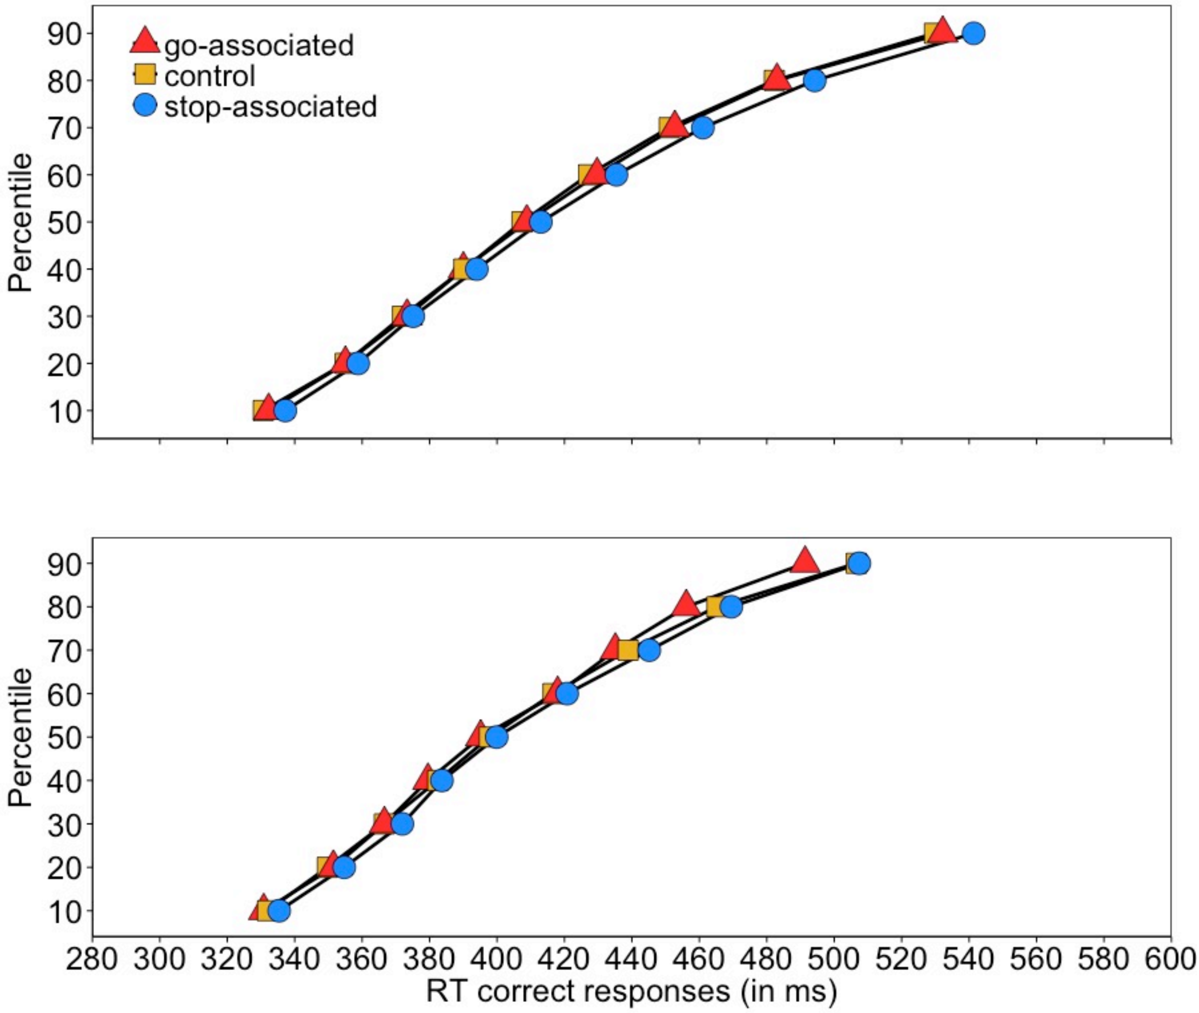

Figure B.3. Go RTs (in ms) in the training phase (blocks 1-6; upper panel) and the test phase (block 7; lower panel) for the three image types (stop-associated; go-associated; control) as a function of percentile in Experiment 3.

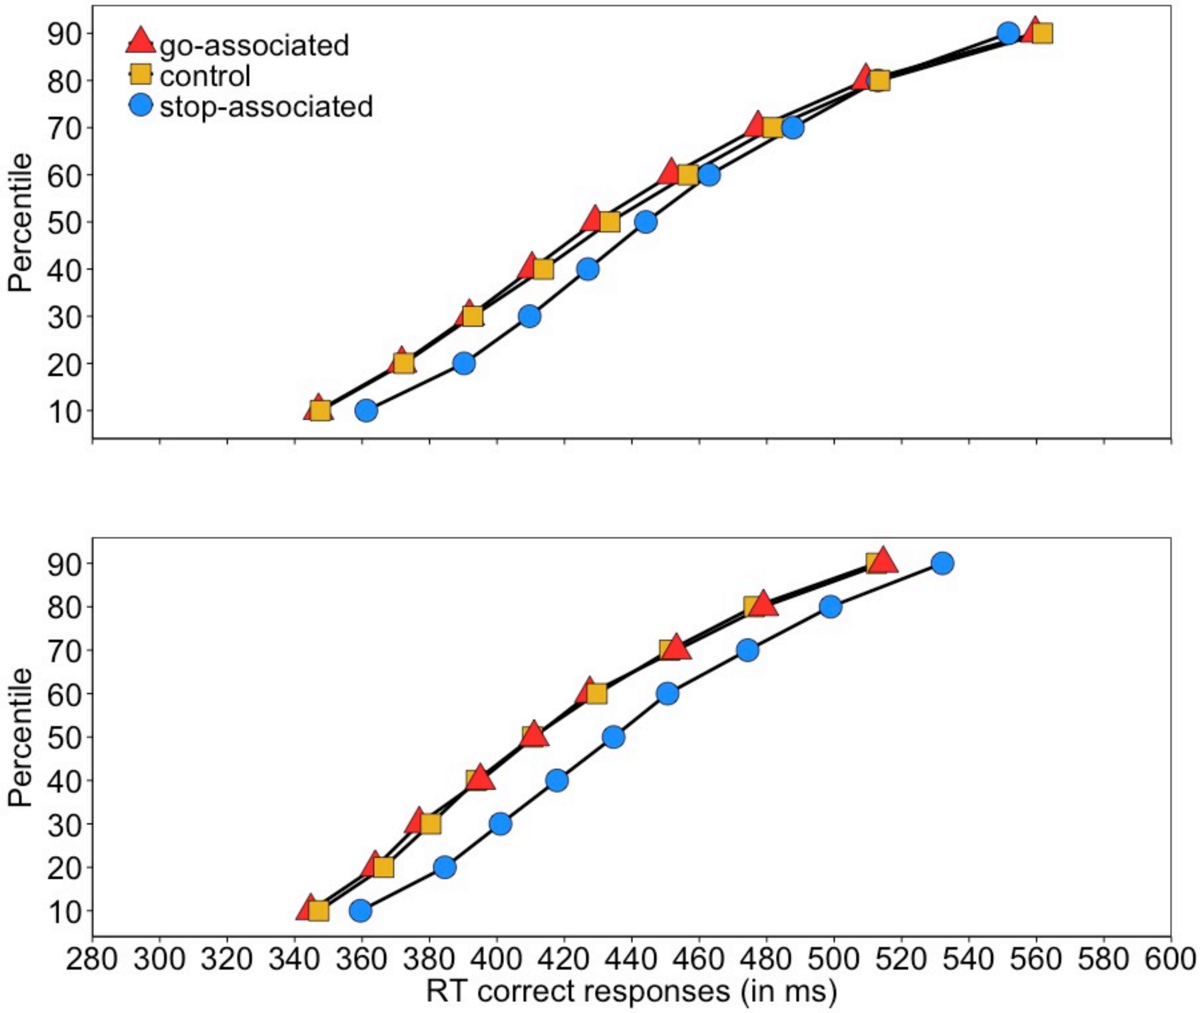

Figure B.4. Go RTs (in ms) in the training phase (blocks 1-6; upper panel) and the test phase (block 7; lower panel) for the three image types (stop-associated; go-associated; control) as a function of percentile in Experiment 4.

Appendix C

Expectancy/RT correlation plots

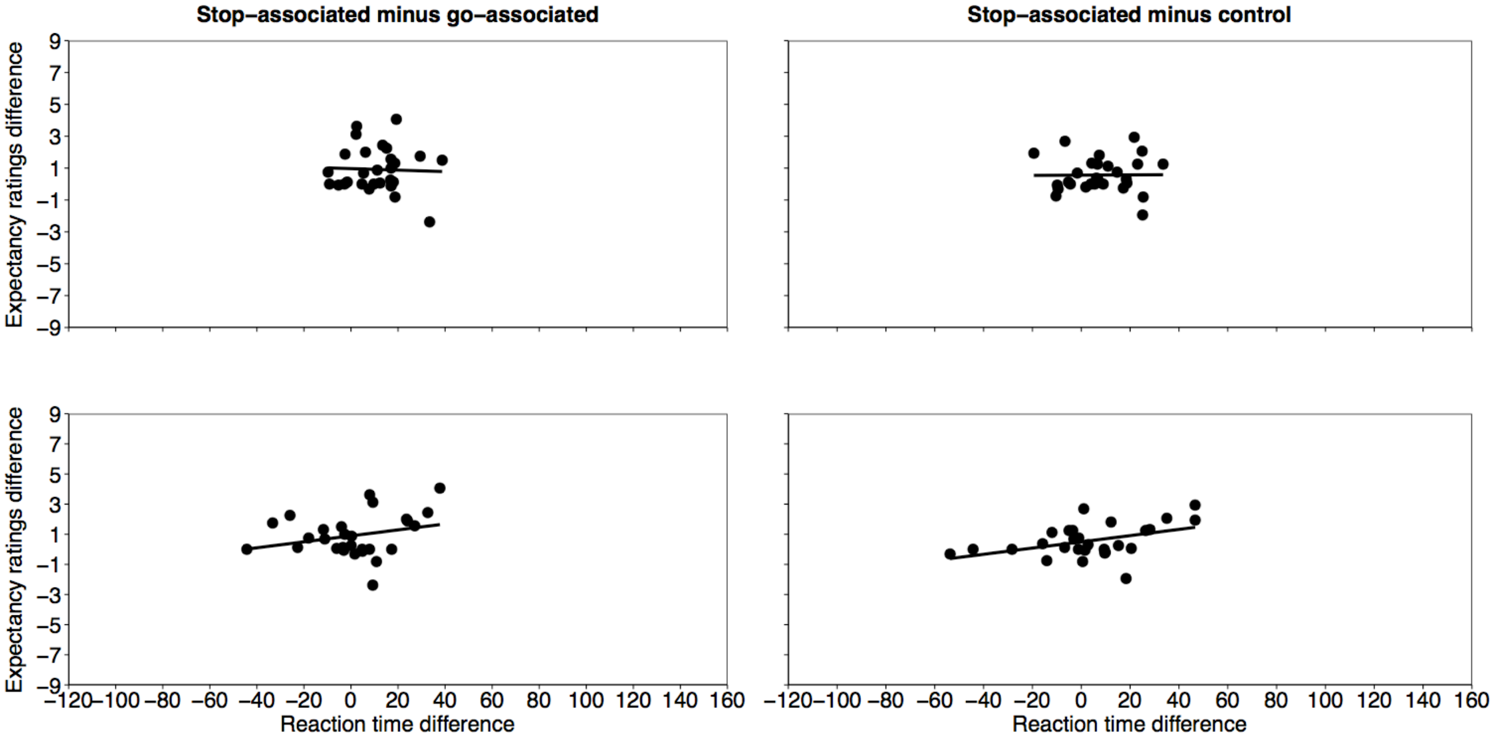

Figure C.1. Expectancy/RT correlations in the training phase (blocks 1-12; upper panels) and the test phase (blocks 13-14; lower panels) in Experiment 1. Note, 'stop-associated minus control image' expectancy difference reliably correlated with the corresponding RT difference in the test phase, $r(26) = 0.437$, $p = 0.019$. All other correlations were not reliable (rs ≤ 0.272, ps ≥ 0.161).

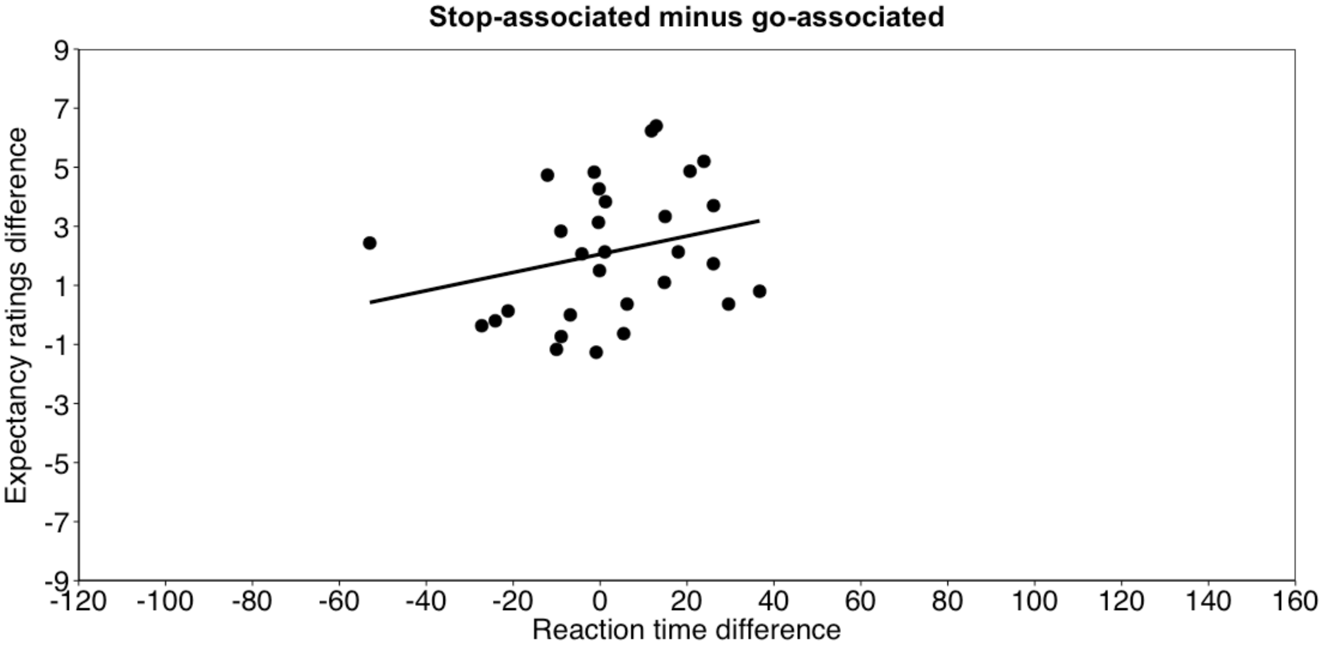

Figure C.2. Expectancy/RT correlations in the test phase (blocks 13-14; lower panel) in Experiment 2. Due to the stimulus-stop contingencies used, I could not run these correlations on the training phase data. Note, the stop-associated minus go-associated correlation was not reliable (*r*(28) = 0.262, *p* = 0.162).

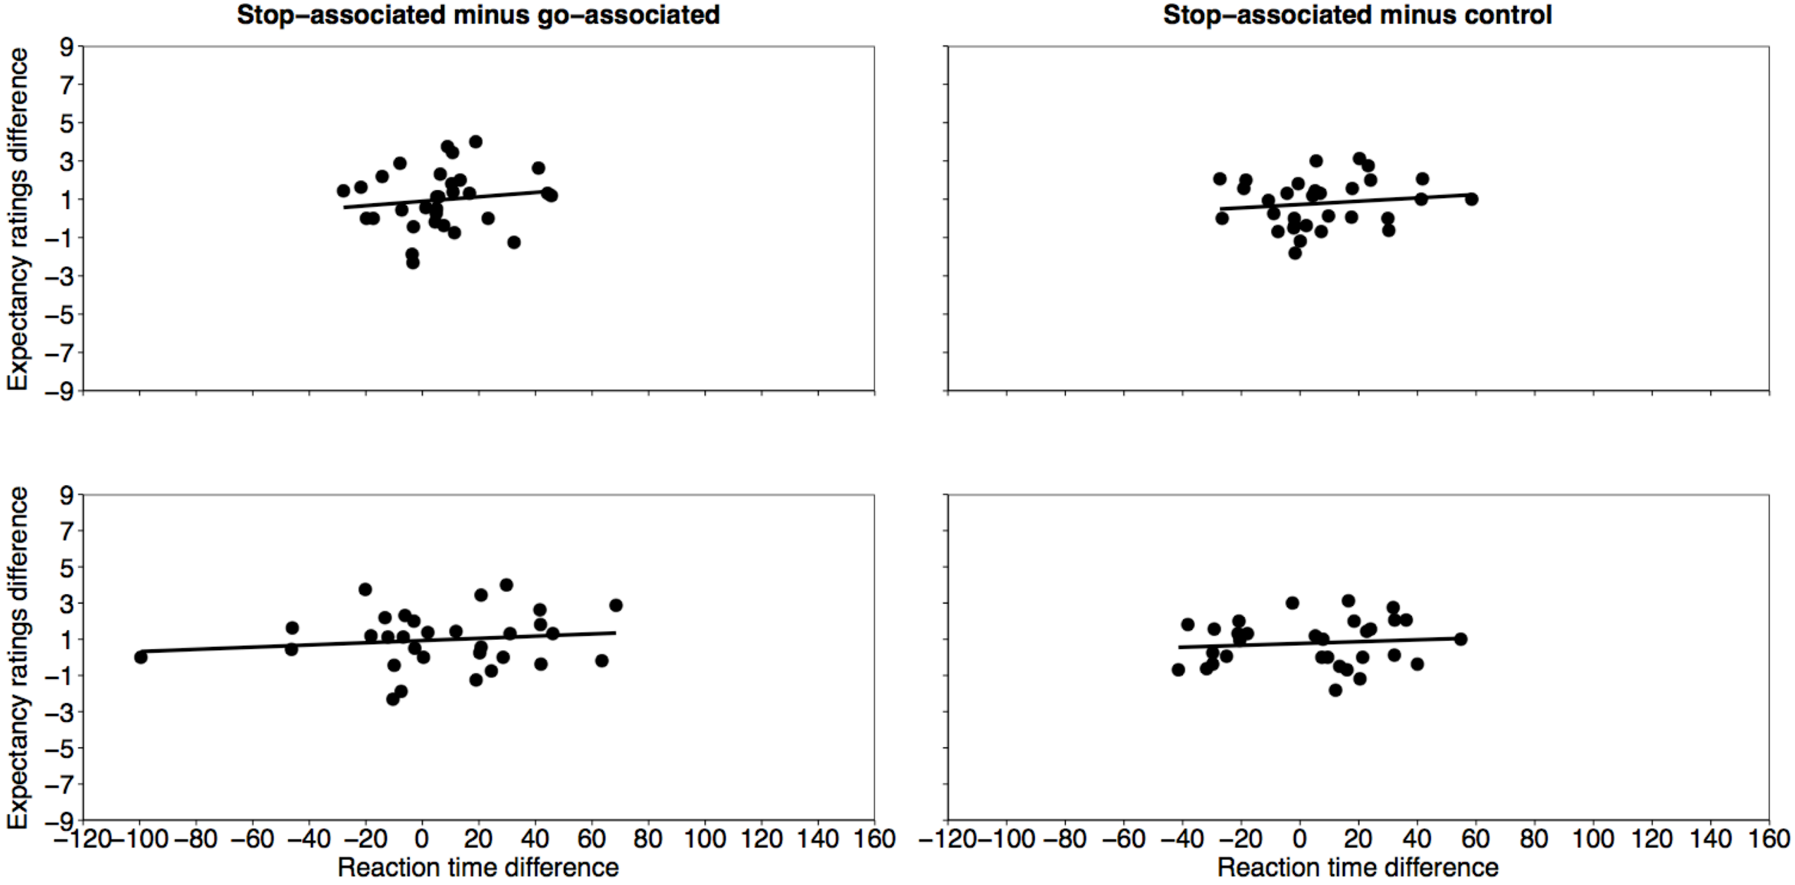

Figure C.3. Expectancy/RT correlations in the training phase (blocks 1-6; upper panels) and the test phase (block 7; lower panels) in Experiment 3. Note, all correlations were not reliable (*r*s ≤ 0.136, *ps* ≥ 0.464).

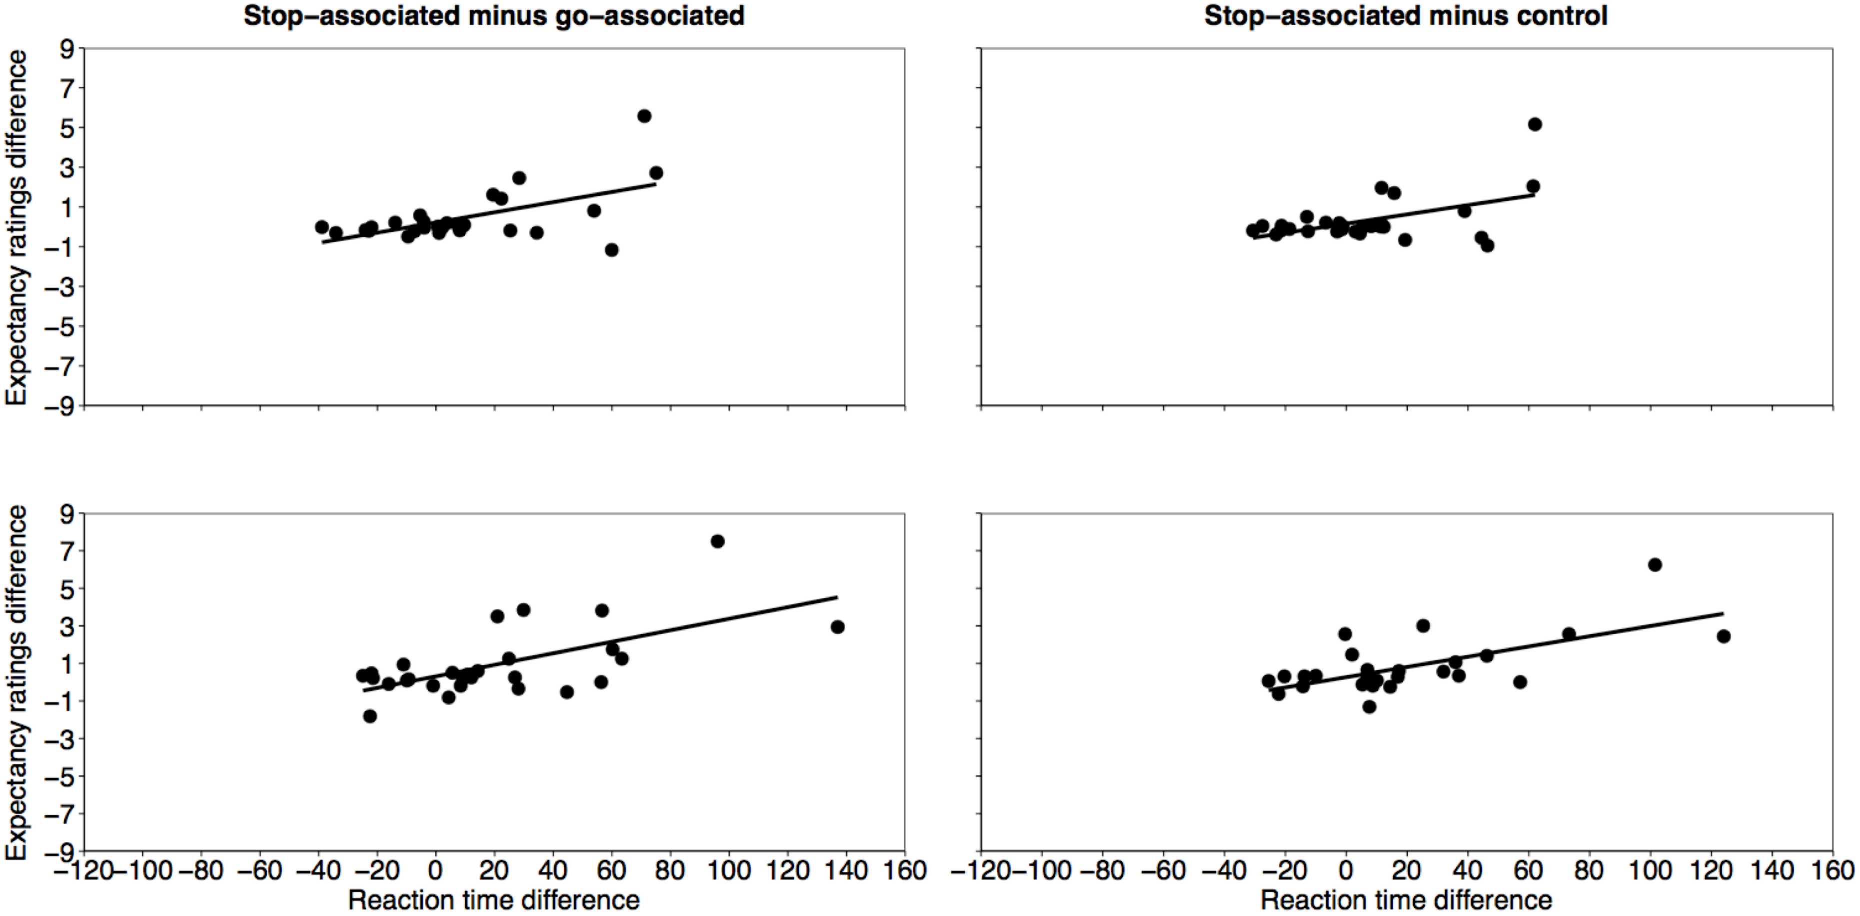

Figure C.4. Expectancy/RT correlations in the training phase (blocks 1-6; upper panels) and the test phase (block 7; lower panels) in Experiment 4. Note, all correlations were reliable (*r*s ≥ 0.498, *p*s ≤ 0.006).

CHAPTER 3

THE EFFECTS OF INSTRUCTED AND ACQUIRED CONTINGENCIES IN RESPONSE INHIBITION TASKS

The ‘automatic inhibition’ account states that after stimulus-specific stop training, the need for top-down executive control is reduced and may eventually disappear altogether as control is triggered in a bottom-up manner following stimulus presentation (see also Chein & Schneider, 2012). Consistent with this account, previous work has demonstrated that the effects of automatic inhibition can occur very rapidly (within ~100 ms) following the presentation of old no-go stimuli (Chiu, Aron, et al., 2012; Manuel, Grivel, et al., 2010), suggesting that, at the very least, a fast, efficient form of control is trained. However, in some situations, slowing for old stop items could also be caused by anticipatory control processes. One of the main functions of the executive control system is biasing competition between stimulus or response options on the basis of expectancy or task rules. For example, when the control system predicts a certain action, it will pre-activate the motor network, biasing action selection and reducing the latency of the anticipated action (e.g. Bestmann, 2012; Meyer & Kieras, 1997; but see Verbruggen, McAndrew, Weidemann, Stevens, & McLaren, 2016).

In Chapter 2, I found that item-specific stop learning may have been mediated via expectancies of the stimulus-stop contingencies in play. Participants had to rate the extent to which they expected to withhold their response for each of the stimuli presented in the task at the end of each experiment. I found correlations between expectancy ratings and the ‘automatic inhibition’ effect in Experiment 1 and Experiment 4. This suggests that expectancies generated on the basis of the acquired stimulus-stop mappings may contribute to the manifestation of an ‘automatic’ inhibition effect: when participants expect a no-go or stop signal in the near future, they proactively adjust attentional settings, increase response thresholds, or proactively suppress all motor output to enhance detection of no-go or stop signals and to prevent premature go responses (e.g. Aron, 2011; Elchlepp, Lavric, et al., 2016; Jahfari, Stinear, et al., 2010; Verbruggen & Logan, 2009a; Verbruggen, Stevens, et al., 2014; Zandbelt, Bloemendaal, et al., 2012). For example, previous work found that participants slowed their responses when an instructional cue, such as ‘p(stop-signal) = 0.75’, at the beginning of the trial indicated that a stop signal was likely to occur on the following trial(s) (Verbruggen & Logan, 2009a). In other words, participants became more cautious when stop-signal presentation was likely. The expectancy results in Chapter 2 suggest that stop-associated stimuli could become such cues (for example ‘if stimulus X then p(stop) is high’). In other words, the slowing observed for old stop items could reflect deliberate ‘proactive’ inhibition, another algorithmic process, rather than the direct activation of the stop response via memory-retrieval (e.g. stimulus X = stop) as assumed by the automatic inhibition account.

Although evidence for proactive control has been found in traditional ‘top-down’ stop-signal tasks, the role of proactive control adjustments during stop learning has not been previously examined. Yet, evidence for the involvement of anticipatory control adjustments in stop-training tasks would have two important implications. First, it would contradict the idea that inhibitory control was fully automatised in stop-training tasks where participants made anticipatory control adjustments. Second, it would have important practical implications for work in the applied inhibition-training as it could indicate that the effects of inhibition training are more ‘top-down’ than initially thought (for an overview of the applied inhibition-training work, see *Practical Applications,* Chapter 1).

The automatic inhibition and expectancy accounts both predict increased RTs for old stop items. One way to assess whether participants proactively adjust response strategies for old stop items is to question them at the end of the experiment (as in Chapter 2). However, measuring knowledge of rules, expectancies, or strategies at the end of the experiment has limitations (Newell & Shanks, 2014; Shanks, 2010). For example, Newell and Shanks (2014) proposed four criteria for the assessment of awareness: (1) *reliability* (the assessment should not be affected by factors that did not affect behavioural performance); (2) *relevance* (assessments should only target information relevant to the behaviour in question); (3) *sensitivity* (assessments should be made under optimal conditions, such as using the same cues as in the task); and (4) *immediacy* (the assessment should be made during behaviour or immediately afterwards). Whilst the assessment of expectancies at the end of task performance meets the first three criteria, it is sub-optimal that expectancies were obtained after learning had taken place. However, as discussed in Chapter 2, asking participants to provide an expectancy rating on each trial (as in Experiment 4) could change the nature of the task. For example, if participants are performing a task more associatively and are then asked to provide expectancy ratings, this may ‘push’ participants towards a more deliberative process. Therefore, I used a complementary approach in the present chapter that did not require participants to provide a trial-by-trial expectancy rating; in addition to measuring expectancy ratings at the end of the experiment, I introduced a between-subjects condition in which participants were informed about the stimulus-stop contingencies at the beginning of the experiment. By contrasting performance in the instructed (explicit) condition with performance in an ‘uninstructed’ condition in which participants were not told about the stimulus-stop contingencies, I could further explore how expectancies and learning influenced performance in a stop-signal task.

The Present Chapter

In the present chapter, I used a learning variant of the stop-signal paradigm developed by Verbruggen, Stevens, et al. (2014). The paradigm used by Verbruggen, Stevens, et al. (2014) was designed to investigate the effects of expectancies on proactive control. In the Verbruggen, Stevens, et al. (2014) study, participants responded to centrally presented words on go trials. In some blocks, a stop signal was presented on a random 33% of the trials. There were three types of blocks: central-signal blocks (in which a visual stop signal could occur in the centre of the screen), non-central signal blocks (in which a visual stop signal could occur in the periphery), and no-signal blocks (in which no stop signals could occur). On a random 50% of the go and stop-signal trials, visual distractors were presented (see Figure 3.1 for an example of a distractor trial). I predicted a trade-off between stop-signal detection and interference control in non-central blocks. On the one hand, participants try to widen their attentional focus to detect stop signals in the periphery; on the other hand, they try to narrow their focus to avoid processing distractors. These opposing demands were expected to result in a larger distractor effect[26] on no-signal trials in non-central signal blocks than in the two other block types without the opposing attentional demands. Their results supported this prediction. In other words, when participants knew that a stop signal could appear in the periphery, they proactively widened their attentional focus to detect the signal, but this resulted in an increased distractor effect on no-signal trials (Verbruggen, Stevens, et al., 2014).

Thus, by introducing a perceptual distractor manipulation, I could examine whether participants adjusted attentional settings (reflected by an increased distractor effect when participants expected a stop signal in the periphery), response settings (reflected by a general RT increase when participants expected a stop signal in the periphery), or both, when they expected a stop signal. In the present chapter, on each trial, a large square and a fixation signal were presented in the middle of the screen (Figure 3.1). After a delay, a single word appeared in the centre of the square. Participants had to decide whether the word referred to a natural or human-made object. On stop-signal trials, the lines of the surrounding square became thicker, instructing the participants to withhold their response. There were three word types in this experiment: *80%-stop words* (on 80% of the trials, a stop signal was presented); *20%-stop words* (on 20% of the trials, a stop signal was presented); and *0%-stop words* (these words could only occur on no-signal trials). To distinguish between the effects of explicit, instructed stimulus-stop contingencies and the effects of practice-induced, uninstructed stimulus-stop contingencies I directly manipulated contingency knowledge via task instructions (between-subjects). In the *instructed* condition, participants were presented with a list of the 80%-stop words at the beginning of each block; in the *uninstructed* condition, participants received no information about the stimulus-stop contingencies. Note that participants in the *instructed* condition received no information about the 20%-stop words or the 0%-stop words. The 80%-stop words remained the same throughout the whole experiment. On 50% of no-signal and stop-signal trials, 2-letter distractor strings appeared. Participants were instructed to ignore them; however, based on the

---

[26] Performance on distractor trials minus performance on no-distractor trials.

previous work by Verbruggen, Stevens, et al., I expected that these distractors would interfere with performance on no-signal trials. If item-specific expectancy effects induce attentional shifts (cf. Verbruggen, Stevens, et al., 2014), the distractor effect on no-signal trials should be larger for the 80%-stop words than for the two other word types: if participants expected a stop signal after having identified an 80%-stop word they would monitor the periphery for the stop signal to occur (as observed by Verbruggen, Stevens, et al., 2014), making them more vulnerable to the distractors. The findings presented in Chapter 2 introduced the possibility that stop-associated stimuli could become such explicit cues after stimulus-stop learning. The idea that the stop-associated items could become explicit cues is contrary to the idea of automatic inhibition which supposes that the stop-associated items become (implicitly) associated with the stop goal following training. Although it was predicted that the overall effect of the stimulus-stop contingencies would be greater in the instructed condition than in the uninstructed condition, the effect of the distractor stimulus allowed me to discern whether the stop-associated stimuli became explicit cues in the uninstructed condition. If this were the case, there should be a larger distractor effect of the 80%-stop words than for the other item-types and, critically, this distractor effect should be qualitatively, and possibly quantitatively, similar across the item-types in the instructed and uninstructed conditions.

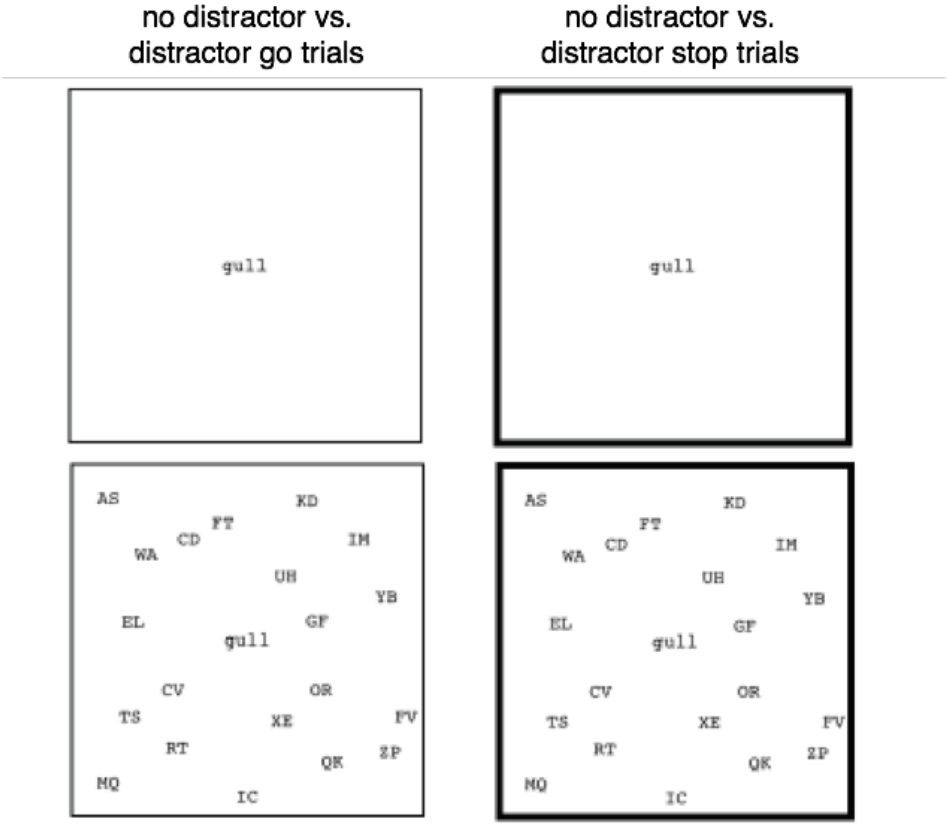

Figure 3.1. The distractor stop task. A word was presented in the middle of a square. Participants had to judge whether the word referred to a human-made or a natural object. On half of the trials (distractor trials), random two-letter strings appeared at random locations every 100 ms. On some trials (stop-signal trials), the square turned bold after a variable delay from the onset of the word, instructing the participants to withhold their response. For display purposes, foreground and background colours are switched in this figure. However, in Experiment 5-6, white stimuli appeared against a black background.

Experiments

There were two experiments. In Experiment 5, I encouraged fast responding using a strict response deadline (1250 ms). Furthermore, I obtained eye movement data as an additional exploratory dependent variable to detect within-trials shifts of attentional focus. However, the probability of missed no-signal responses (*M* = 0.050, *SD* = 0.066) in Experiment 5 was slightly higher than in the previous study by Verbruggen, Stevens, et al., including for words that were not associated with stopping. Therefore, I ran another experiment in which I extended the response deadline (2000 ms) to reduce the probability of missed responses. In Experiment 5, the eye-tracker was calibrated at the beginning of each block and lasted approximately 2 minutes per calibration; to ensure that the overall task duration was comparable across experiments, I therefore did not obtain eye movement data in

Experiment 6. After all, the eye movement data of Experiment 5 did not substantially add to the overall pattern of behavioural data (see Appendix F; note that the eye-movement data similarly did not add much to the pattern of data reported by Verbruggen, Stevens, et al., 2014).

Initial analyses revealed that whilst the probability of missed no-signal responses was reduced in Experiment 6 (and the no-signal RTs were numerically increased) the overall pattern of results (i.e. differences between word types, distractor types, and instruction conditions) was consistent across both experiments. This was revealed by analyses of no-signal RTs, go accuracy and the probability of responding on stop-signal trials. Separate analyses for Experiment 5 and Experiment 6 can be found in Appendix D. These separate analyses produced similar results; in both experiments, there is reliable evidence of stimulus-stop learning on go and stop performance, and a reliable distractor effect on go and stop performance in both experiments. Furthermore, there was reliable evidence in both experiments that the differences between the word types were larger in the instructed condition than in the uninstructed condition. Thus, both experiments find (statistically) clear evidence that the stimulus-stop contingencies influenced performance, and that the effects of these contingencies were sensitive to the instructional manipulation. Therefore, I analysed the data of both experiments together (total *N* = 120), and consequently will report this study as a single experiment. Although the individual studies were sufficiently powered to detect small to medium-sized within-between interactions they were not sufficiently powered to detect direct medium-sized between-subjects comparisons. This combined approach was beneficial as it increased the power to detect medium-sized effects in the between-subjects comparisons and allowed me to run sufficiently powered correlational analyses on the expectancy and task performance data.

Method

Participants. 120 volunteers (48 in Experiment 5 and 72 in Experiment 6) from the University of Exeter participated for monetary compensation (£8) or partial course credit (*M* = 19.88 years, *SD* = 3.05, 98 females, 109 right-handed). Five participants were replaced: two participants were replaced because the percentage of correct go trials was below 70%[27] (one in Experiment 5 and one in Experiment 6; this exclusion criterion was based on an earlier pilot study presented in Appendix G) and three further participants in Experiment 5 were replaced due to poor calibration of the eye-tracker.

All experiments in the present chapter were approved by the local research ethics committee at the School of Psychology, University of Exeter. Written informed consent was obtained after the nature and possible consequences of the studies were explained. The target sample size and exclusion criteria for each experiment were decided in advance of data collection. Participants were randomly assigned to between-subjects groups.

---

[27] This exclusion criteria was set on the basis of the pilot study presented in Appendix G.

Apparatus and stimuli. The experiments were run using Psychtoolbox (Brainard, 1997). The stimuli were presented on a 17-in CRT monitor (screen size: 1024 × 768 pixels) in Experiment 5 and a 21-in iMac (screen size: 1920 × 1080) in Experiment 6. In Experiment 5, an EyeLink 1000 Desktop Mount camera system (SR Research, Ottawa, Canada), calibrated before each block, tracked the gaze position of the right eye during the whole block.

The stimuli consisted of a large square (12 cm diameter: 350 × 350 pixels) and a word in white lowercase font (Courier 16 point) on a black background (Figure 3.1). The stimuli were presented in the centre of the screen. I created a list of 50 four-letter words (Appendix E). The word could refer to a natural or a human-made object. The experiment consisted of three parts, with 5 blocks per part. Each word was presented twice per block (1x with and 1x without distractors). There were 3 different word types: *80%-stop words* (10 words; on 80% of the trials, a stop signal was presented); *20%-stop words* (30 words; on 20% of the trials, a stop signal was presented); and *0%-stop words* (10 words; these words could only occur on no-signal trials). There were more 20%-stop words than 80%-stop words to keep the overall probability of stop-signal trials low (0.28). There is no reason to assume that this would influence the acquisition of stimulus-stop or stimulus-go associations as each word was presented an equal number of times per block. Words were counterbalanced over conditions between participants.

Procedure. On each trial, the square and a fixation signal were presented in the middle of the screen. After 250 ms, the word appeared in the centre of the square (replacing the fixation signal). Participants had to decide whether the word referred to a natural or a human-made object. Half of the participants had to press the ‘c’ key (with their left index finger) when the word referred to a natural object, and the ‘m’ key (with their right index finger) when the word referred to a human-made object. This mapping was reversed for the other participants. On 50% of the trials twenty 2-letter randomly generated uppercase strings were displayed as distractors (distractor trials; Figure 3.1). They appeared in random locations within the square. To avoid overlap between the distractors and words, the centre of the distractors was outside a smaller central region (100 × 50 pixels). New distractors were presented in different locations every 100ms throughout the duration of the trial. Participants were instructed to ignore the distractors.

On stop-signal trials, the outer square turned bold (1 to 3 pixels) after a variable SSD, instructing participants to withhold their response. The stop signal occurred equally often on distractor and on non-distractor trials. The SSD for 20%-stop words was initially set at 500 ms and was continuously adjusted to a tracking procedure to obtain a probability of successful stopping of 0.50: the SSD decreased by 50 ms following an unsuccessful stop trial, but increased by 50 ms following a successful stop. I used two separate one-up/one-down tracking procedures to set the SSD for the 20%-

stop words (one for distractor trials and one for no-distractor trials). The SSD for the 80%-stop words was yoked to the SSD for the corresponding 20%-stop words.

After the response deadline elapsed (1250 ms in Experiment 5; 2000 ms in Experiment 6) I presented feedback (on no-signal trials: ‘correct’, ‘incorrect’, or ‘not quick enough’ in case the participant did not respond before the end of the trial; on stop-signal trials: ‘correct stop’ or ‘failed stop’), which remained on the screen for 500 ms. The feedback was presented to encourage fast and accurate responding. The next trial started immediately after the feedback.

Stimulus-stop contingency knowledge was manipulated between-subjects. Participants in the *instructed* condition were presented with a list of the 80%-stop words on the screen at the beginning of each block and were instructed ‘*For certain words, the lines of the surrounding square will become thicker (indicating that you have to withhold your response) more often than for other words. These words are listed below*’. The stimulus-stop contingencies remained the same throughout the whole experiment so the same list was presented at the beginning of each block. No information about the 20%-stop words or and 0%-stop words was provided. Participants were instructed to remember as many of the words as possible. Once participants had done so, participants were required to press the ‘s’ key to move on to the next screen. There was no deadline on the word list screens. Participants in the *uninstructed* condition were not provided with any information about the stimulus-stop contingencies. All participants were required to press the ‘s’ key to start the first trial. At the end of each block, I presented as feedback to the participant their mean RT on no-signal trials, the number of no-signal errors and the number of missed no-signal responses, and the percentage of failed stops.

Following completion of the experimental task, each word was again presented on the screen. The order of the words was randomised anew for each participant. Participants were asked to rate ‘how much do you expect to withhold your response when this word is presented?’ on a scale between 1 ('I definitely do not think this word indicates that I have to withhold my response') and 9 ('I definitely think this word indicates that I have to withhold my response'). There was no response deadline for the expectancy ratings.

Analyses. All data processing and analyses were completed using R (R Development Core Team, 2014). All data files and R scripts are deposited in Dropbox (https://www.dropbox.com/sh/obhjddhyqmth6ng/AABW3NcPOhM9Xms1Cxxc1zKVa?dl=0).

Proactive strategy adjustments could result in a higher percentage of missed responses in addition to higher no-signal accuracy (Verbruggen & Logan, 2009a) so I distinguished between the proportion of correct no-signal trials and the proportion of missed no-signal trials. To determine if learning influenced stopping performance, I analysed the *p*(respond|signal); see also Chapter 2; Bowditch, Verbruggen, et al., 2016; Noël, Brevers, et al., 2016).

ANOVAs were performed on the correct no-signal RTs, the probability of missed responses on no-signal trials, the probability of correct no-signal trials, and on the probability of responding on stop-signal trials. Performance was analysed as a function of 'part' (there were 5 blocks per part: 1 = blocks 1-5; 2 = blocks 6-10; 3 = blocks 11-15), word type (80%-stop words, 20%-stop words, and 0%-stop words), distractor type (distractor vs. no distractor), and instruction condition (instructed vs. uninstructed). Where appropriate, I applied the Huyhn-Feldt correction for violations of sphericity. For pairwise comparisons, Hedge's g_{av} is the reported effect size measure (Lakens, 2013).

For the RT analyses, I also calculated Bayes factors for all main effects and interaction contrasts in the ANOVA design (Rouder, Morey, Speckman, & Province, 2012). I calculated these with the BayesFactor package in R, using the default prior (0.707; Morey, Rouder, et al., 2015). As I had four factors (word type, distractor type, part, and condition), I used a top-down approach to reduce the number of model comparisons. Top-down model comparisons investigate the effect of removing each fixed factor and interaction from the overall model, such that the removal of meaningful factors or interactions will have a deleterious effect on the model fit whereas the removal of non-meaningful factors or interactions will not.

Tables 3.1 and 3.6 presents an overview of the descriptive statistics and Tables 3.2 and 3.3 present an overview of the inferential statistics (separate analyses for the instructed and uninstructed conditions are presented in Tables 3.4-3.5). I present exploratory analyses of the eye-tracking data of Experiment 5 in Appendix F.

Results

No-signal analyses.

No-signal RTs. Figure 3.2 (upper panel) shows correct no-signal RTs for each condition (instructed or uninstructed), word type (80%-, 20%-, 0%-stop), and distractor type (distractor or no-distractor). I found a main effect of word type ($p < 0.001$; Table 3.2), and an interaction between condition and word type ($p < 0.001$; Table 3.2): Consistent with the stimulus-stop expectancy manipulation, the word type effect was more pronounced in the instructed condition than in the uninstructed condition (Figure 3.2). Follow-up tests (Tables 3.4 & 3.5) revealed that the main effect of word type on no-signal RTs was significant in the instructed condition ($p < 0.001$) and in the uninstructed condition ($p < 0.001$; Table 3.4). In the instructed condition, additional planned comparisons revealed that no-signal RTs were longer for the 80%-stop words (933 ms) than for the 0%-stop words (863 ms), $t(59) = 10.17$, $p < 0.001$, $g_{av} = 0.34$ (one-tailed directional t-test: $p < 0.001$[28]), and for the 20%-stop words (874 ms), $t(59) = 9.34$, $p < 0.001$, $g_{av} = 0.29$ (one-tailed directional t-test:

---

[28] The 'automatic inhibition' account makes strong predictions about the direction of this effect. No differences between word types or differences in the opposite direction (i.e. shorter RTs or lower p(miss) for 80% words than for 0% and 20% words) would argue against stimulus-stop associations. Therefore, I report both the two-tailed p-values and the p-values of one-directional t-tests.

$p < 0.001$)[29]. The difference between the 0%-stop words and the 20%-stop words was also reliable, $t(59) = -5.10$, $p < 0.001$, $g_{av} = -0.05$ (one-tailed directional t-test: $p < 0.001$). In the uninstructed condition, planned comparisons revealed that no-signal RTs were longer for the 80%-stop words (921 ms) than for the 0%-stop words (900 ms), $t(59) = 4.87$, $p < 0.001$, $g_{av} = 0.10$ (one-tailed directional t-test: $p < 0.001$), and for the 20%-stop words (902 ms), $t(59) = 4.71$, $p < 0.001$, $g_{av} = 0.09$ (one-tailed directional t-test: $p < 0.001$). The difference between the 0%-stop words and the 20%-stop words was not reliable, $t(59) = - 1.14$, $p = 0.261$, $g_{av} = -0.01$ (one-tailed directional t-test: $p = 0.130$).

---

29. To explore the difference between the 0%-stop words and the 20%-stop words I repeated the main analysis excluding the 80%-stop words. This showed that the two-way interaction between condition and word type was reliable, $F(1, 118) = 7.12$, $p = 0.009$, *gen.* $\eta^2 < 0.001$, reflecting the larger RT differences between the 0%-stop words and the 20%-stop words in the instructed condition than in the uninstructed condition. Furthermore, the three way interaction between condition, word type and part was also reliable, $F(2, 236) = 3.07$, $p = 0.047$, *gen.* $\eta^2 < 0.001$: the numerical difference between the 0%-stop words and the 20%-stop words became larger as a function of part in the instructed condition but became smaller as a function of part in the uninstructed condition.

	***p*(correct)**				***p*(miss)**				**no-signal RT**			
	Distractors		***No distractors***		***Distractors***		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed												
Part 1												
0%-stop	0.93	0.07	0.92	0.06	0.029	0.033	0.026	0.031	860	146	818	149
20%-stop	0.93	0.05	0.92	0.06	0.034	0.037	0.026	0.029	870	147	823	144
80%-stop	0.90	0.13	0.89	0.13	0.085	0.120	0.050	0.100	914	175	859	149
Part 2												
0%-stop	0.93	0.06	0.94	0.06	0.021	0.033	0.025	0.034	883	232	835	219
20%-stop	0.94	0.05	0.94	0.05	0.018	0.023	0.020	0.026	896	235	848	221
80%-stop	0.91	0.10	0.91	0.11	0.042	0.085	0.035	0.071	955	250	924	250
Part 3												
0%-stop	0.93	0.06	0.93	0.07	0.029	0.039	0.025	0.039	916	252	863	248
20%-stop	0.93	0.06	0.93	0.06	0.029	0.046	0.023	0.033	923	252	881	249
80%-stop	0.91	0.10	0.84	0.15	0.053	0.098	0.043	0.081	995	290	954	296
Uninstructed												
Part 1												
0%-stop	0.93	0.06	0.92	0.07	0.032	0.039	0.023	0.029	882	149	825	144
20%-stop	0.93	0.07	0.92	0.07	0.038	0.039	0.025	0.030	893	150	831	143
80%-stop	0.92	0.09	0.92	0.10	0.033	0.066	0.038	0.069	904	151	845	160
Part 2												
0%-stop	0.94	0.07	0.94	0.07	0.036	0.053	0.020	0.030	935	262	891	258
20%-stop	0.94	0.06	0.94	0.06	0.035	0.044	0.027	0.031	942	271	886	253
80%-stop	0.95	0.10	0.91	0.10	0.035	0.086	0.032	0.070	962	286	920	272
Part 3												
0%-stop	0.93	0.08	0.93	0.07	0.037	0.051	0.027	0.044	958	277	907	275
20%-stop	0.93	0.07	0.93	0.06	0.042	0.058	0.028	0.034	960	280	901	269
80%-stop	0.92	0.13	0.92	0.12	0.033	0.060	0.030	0.053	974	281	923	279

Table 3.1. Overview of the no-signal data. Probability of an accurate go response [*p*(correct)], probability of a missed go response [*p*(miss)] and average no-signal reaction time (RT) as a function of condition, word type, distractor type, and part. Accuracy is the ratio of correct no-signal trials to the number of correct and incorrect no-signal trials (missed trials are excluded). *P*(miss) is the ratio of omitted responses to the total number of no-signal trials. *M* = mean; *SD* = standard deviation.

I found a main effect of distractor type ($p < 0.001$; Table 3.2), but there was no reliable difference in the size of the distractor effect between conditions ($p = 0.270$; Table 3.2). Follow-up tests showed that the presentation of distractors slowed no-signal RTs in the instructed condition (no-distractor: 867 ms, distractor: 912 ms; $p < 0.001$; Table 3.4) and in the uninstructed condition (no-distractor: 881 ms, distractor: 934 ms; $p < 0.001$; Table 3.5). Crucially, there was no reliable two-way interaction between word type and distractor type ($p = 0.349$; Table 3.2). Furthermore, the three-way interaction between condition, word type, and distractor type was not reliable ($p = 0.447$; Table 3.2), and follow-up tests showed that there was no reliable two-way interaction between word type and distractor type in the instructed condition ($p = 0.665$; Table 3.4) nor in the uninstructed condition ($p = 0.195$; Table 3.5). As discussed above, in the previous study by Verbruggen, Stevens, et al. (2014) it was found that proactive adjustments in non-central stop-signal blocks led to an increased distractor effect on no-signal RTs. Therefore, the absence of an interaction between word type and distractor type in the present chapter indicates that item-based expectancy effects differ from the kind of proactive adjustments of attentional settings observed in block-based designs. Instead, the absence of an interaction between word type and distractor supports the idea that the stop-associated words primed stopping. However, the presence of a reliable two-way interaction between condition and word type suggests that this priming effect must have been greater in the instructed condition than in the uninstructed condition. I will return to this issue in the Chapter Discussion.

I found an interaction between part and word type in the main analysis (Table 3.2) and in the analysis of the instructed condition (p = 0.005; Table 3.4). However, the interaction between word type and part did not reach significance in the uninstructed condition (p = 0.323; Table 3.5), suggesting that the effect of word type was present in all parts (see Table 3.1). This is consistent with previous work which indicates that the effect of stop learning emerges after a single trial presentation, and that it then quickly asymptotes (Verbruggen & Logan, 2008a, 2008c).

Table 3.3 shows the outcome of the Bayesian analyses. As can be seen, the results are largely consistent with the ANOVAs reported in Tables 3.2. Table 3.3 shows that dropping word type, distractor type, part, and the two-way interaction between condition and word type had a deleterious effect on the model. All other factors or interactions could be dropped (Table 3.3 also shows the separate models for the instructed and uninstructed conditions). These analyses provide further support for the conclusion that the distractor effect was similar for the different word types.

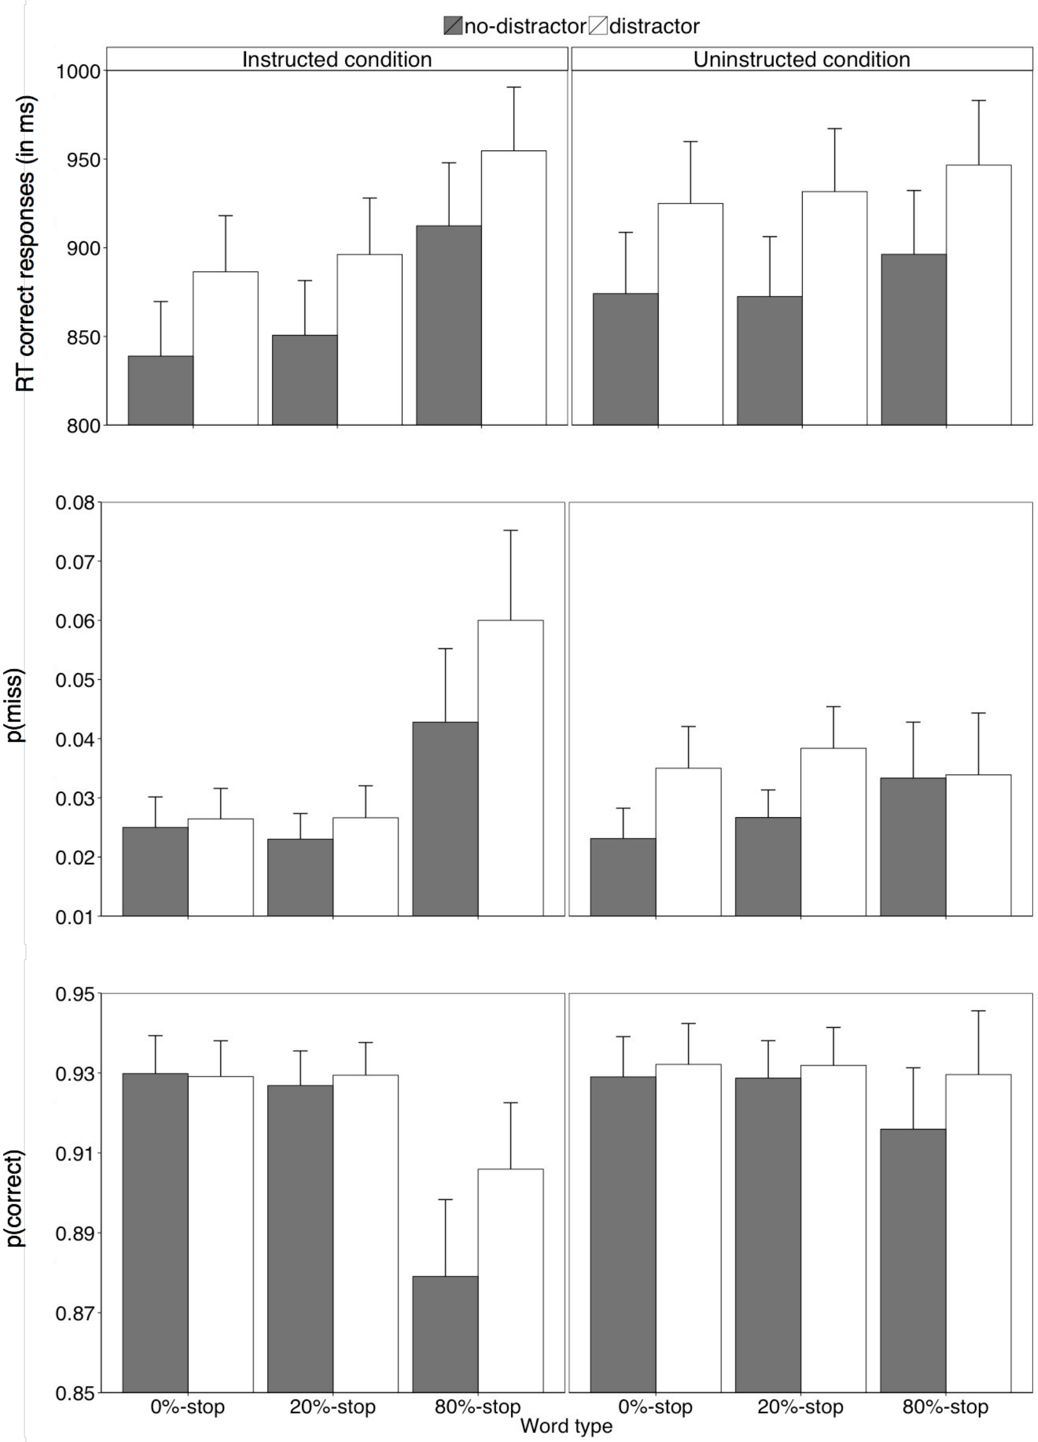

Figure 3.2. Reaction times for the correct no-signal trials (upper panel), the probability of missed no-signal responses (middle panel), and the probability of correct no-signal responses (lower panel) for the instructed condition (left panels) and for the uninstructed condition (right panels) as a function of word type (0%-stop, 20%-stop or 80%-stop words) and distractor type (distractor, no-distractor). Error bars reflect 95% confidence intervals.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Go trials: go RT							
Condition	1	118	171051.50	94794587.10	0.21	0.645	0.002
Part	2	236	1783555.0	13010265.90	16.18	**< 0.001**	0.016
Type (0%-stop, 20%-stop, 80%-stop)	2	236	904922.00	966318.40	110.50	**< 0.001**	0.008
Distract	1	118	1309571.0	896684.60	172.33	**< 0.001**	0.011
Condition by part	2	236	64644.43	13010265.90	0.59	0.492	0.001
Condition by type	2	236	246795.90	966318.40	30.14	**< 0.001**	0.002
Condition by distract	1	118	9326.65	896684.60	1.23	0.270	< 0.001
Part by type	4	472	41534.56	1219546.70	4.02	**0.011**	< 0.001
Part by distract	2	236	7617.38	664742.90	1.35	0.261	< 0.001
Type by distract	2	236	3299.20	387996.50	1.00	0.349	< 0.001
Condition by part by type	4	472	28190.83	1219546.70	2.73	0.052	< 0.001
Condition by part by distract	2	236	812.94	664742.90	0.14	0.857	< 0.001
Condition by type by distract	2	236	2398.44	387996.50	0.73	0.447	< 0.001
Part by type by distract	4	472	7405.20	1006301.70	0.87	0.443	< 0.001
Condition by part by type by distract	4	472	3603.93	1006301.70	0.42	0.705	< 0.001
Go trials: *p*(correct)							
Condition	1	118	0.07	8.00	0.99	0.321	0.004
Part	2	236	0.10	1.13	9.95	**< 0.001**	0.006
Type	2	236	0.23	1.38	19.83	**< 0.001**	0.015
Distract	1	118	0.04	0.40	10.57	**0.001**	0.002
Condition by part	2	236	0.01	1.13	0.69	0.499	< 0.001
Condition by type	2	236	0.10	1.38	8.39	**0.002**	0.007
Condition by distract	1	118	0.00	0.40	0.34	0.562	< 0.001
Part by type	4	472	0.02	1.43	1.55	0.204	0.001
Part by distract	2	236	0.01	0.70	1.35	0.260	0.001
Type by distract	2	236	0.04	0.64	7.35	**0.003**	0.003
Condition by part by type	4	472	0.01	1.43	0.94	0.417	0.001
Condition by part by distract	2	236	0.04	0.70	6.33	**0.002**	0.003
Condition by type by distract	2	236	0.01	0.64	1.35	0.257	< 0.001
Part by type by distract	4	472	0.04	1.24	3.56	**0.019**	0.002
Condition by part by type by distract	4	472	0.03	1.24	3.18	**0.030**	0.002
Go trials: *p*(miss)							
Condition	1	118	0.00	2.56	0.13	0.724	< 0.001
Part	2	236	0.02	1.09	2.38	0.098	0.003
Type	2	236	0.10	0.66	18.03	**< 0.001**	0.015
Distract	1	118	0.03	0.23	16.71	**< 0.001**	0.005
Condition by part	2	236	0.02	1.09	2.02	0.138	0.003
Condition by type	2	236	0.07	0.66	11.92	**< 0.001**	0.010
Condition by distract	1	118	0.00	0.23	0.03	0.869	< 0.001
Part by type	4	472	0.02	1.07	1.89	0.146	0.003
Part by distract	2	236	0.00	0.28	1.22	0.296	< 0.001
Type by distract	2	236	0.00	0.37	0.14	0.760	< 0.001
Condition by part by type	4	472	0.01	1.07	0.83	0.453	0.001
Condition by part by distract	2	236	0.01	0.28	3.34	**0.039**	0.001
Condition by type by distract	2	236	0.02	0.37	6.54	**0.007**	0.003
Part by type by distract	4	472	0.00	0.55	0.59	0.585	< 0.001
Condition by part by type by distract	4	472	0.01	0.55	1.37	0.254	0.001

Table 3.2. Overview of repeated measures Analyses of Variance on the no-signal data. Condition is included as a between-subjects factor, all other factors are within-subjects. In the no-signal RT analysis, incorrect, and missed no-signal trials were removed. *p*s < 0.05 are highlighted in bold.

Probability of a missed go response [p(miss)]. Figure 3.2 (middle panel) shows the probability of a missed go response for each condition (instructed or uninstructed), word type (80%-stop, 20%-stop, 0%-stop), and distractor type (no-distractor or distractor). I found a main effect of word type ($p < 0.001$; Table 3.2) and a reliable two-way interaction between condition and word type ($p < 0.001$; Table 3.2). The main effect of word type on the probability of missed responses was reliable in the instructed condition ($p < 0.001$; Table 3.4) but not in the uninstructed condition ($p = 0.224$; Table 3.5). Follow-up comparisons revealed that the p(miss) in the instructed condition was higher for the 80%-stop words (0.051) than for the 0%-stop words (0.026), $t(59) = 4.43$, $p < 0.001$, $g_{av} = 0.58$ (one-tailed directional t-test: $p < 0.001$), and for the 20%-stop words (0.025), $t(59) = 4.43$, $p < 0.001$, $g_{av} = 0.61$ (one-tailed directional t-test: $p < 0.001$). There was no reliable difference between the 0%-stop words and the 20%-stop words, $t(59) = 0.675$, $p = 0.502$, $g_{av} = 0.04$ (one-tailed directional t-test: $p = 0.749$).

The p(miss) was higher on distractor trials than on no-distractor trials in both the instructed and uninstructed conditions (ps ≤ 0.012; Tables 3.4-3.5; main effect of distractor: $p < 0.001$; Table 3.2). The three-way interaction between condition, word type, and distractor type was reliable ($p = 0.007$; Table 3.2). Follow-up tests showed that the two-way interaction between word type and distractor type was reliable in the instructed condition ($p = 0.047$; Table 3.4): the distractor effect was numerically larger for the 80%-stop words than the other word types (Table 3.1). The two-way interaction between word type and distractor type was not reliable in the uninstructed condition ($p = 0.091$; Table 3.5).

Combined, these results suggest that the instructed stimulus-stop associations increased outright stopping following the presentation of stop-associated words (even when a go response was required). The two-way interaction between word type and distractor type in the instructed condition could indicate that participants relied more on the instructed stimulus-stop contingency when monitoring for the stop signal was difficult (on distractor trials); when monitoring was easier (on no-distractor trials), the decision to stop (or not) was primarily influenced by the actual presentation of the stop signal (and hence, there were fewer missed responses following the presentation of 80%-stop words). This conclusion is supported by the eye movement data; as can be seen in Appendix F, participants made very few eye-movements in this task. Nevertheless, there were some small differences in the number of eye movements between the word types. Importantly, however, I found no evidence that the distractors influenced eye-movements and little evidence to support the view that participants moved their eyes to enable stop-signal detection following the presentation of the 80%-stop words compared with the other word types. I discuss the implications of these findings for the involvement of anticipatory control adjustments in the Chapter Discussion.

Omitted factor(s)	Bayes Factor	Confidence interval
Main analysis		
Part	< 0.00	±104.9%
Distract	< 0.00	±66.33%
Type	< 0.00	±89.32%
Condition by type	0.01	±98.86%
Condition	6.76	±65.95%
Condition by part	68.63	±103.7%
Condition by distract by part by type	268.60	±57.32%
Condition by distract by type	387.73	±91.84%
Condition by distract	702.78	±108.94%
Distract by part by type	796.09	±63.71%
Part by type	878.80	±88.37%
Distract by type	901.37	±88.94%
Condition by part by type	4387.12	±79.46%
Distract by part	4814.11	±97.42%
Condition by distract by part	88222.89	±108.29%
Instructed condition only		
Type	< 0.00	±44.47%
Part	< 0.00	±55.46%
Distract	< 0.00	±44.69%
Part by type	28.38	±43.48%
Distract by part	66.71	±47.56%
Distract by type	87.66	±53.96%
Distract by part by type	178.86	±57.8%
Uninstructed condition only		
Part	< 0.00	±51.14%
Distract	< 0.00	±52.73%
Type	0.15	±48.48%
Distract by part	12.84	±52.35%
Distract by type	23.79	±50.26%
Distract by part by type	43.38	±44.85%
Part by type	67.32	±45.79%

Table 3.3. No-signal RT Bayesian analysis. Bayes factors < 1 indicate that the removal of the factor or interaction had a deleterious effect on the model, whereas Bayes factors > 1 indicate that the factor or interaction could be removed without impairing the fit much. Note that ‘participant’ was included as a factor for all models, but this factor is not added to the model descriptions in the tables to reduce the amount of text. Note: distract = distractor type (no-distractor, distractor); type = word type.

Probability of a correct go response. Figure 3.2 (lower panel) shows the probability of a correct go response for each condition (instructed or uninstructed), word type (80%-stop, 20%-stop, 0%-stop), and distractor type (no-distractor or distractor). I found a reliable main effect of word type ($p < 0.001$; Table 3.2) and a reliable two-way interaction between condition and word type ($p < 0.001$; Table 3.2). Follow-up analyses revealed that the main effect of word type on no-signal accuracy was reliable in the instructed condition ($p < 0.001$; Table 3.4) but not in the uninstructed condition ($p = 0.204$; Table 3.5). Accuracy was lower in the instructed condition for the 80%-stop words (0.89) than for the 0%-stop words (0.93), $t(59) = 4.73$, $p < 0.001$, $g_{av} = 0.53$, and for the 20%-stop words (0.93),

$t(59) = 4.88$, $p < 0.001$, $g_{av} = 0.51$. There was no reliable difference between the 0%-stop words and the 20%-stop words, $t(59) = -0.44$, $p = 0.661$, $g_{av} = -0.03$.

There was a small but statistically significant difference between the no-distractor trials and the distractor trials ($p = 0.001$; Table 3.2). This difference was significant in both conditions (ps $\leq$ 0.029). It is plausible that the decreased accuracy on the no-distractor trials reflects a speed-accuracy trade-off due to the faster responding on no-distractor trials than on distractor trials. I also found a reliable two-way interaction between word type and distractor type in the instructed condition (p = 0.008); the difference between the no-distractor and distractor trials was numerically larger for the 80%-stop words than for the other word types. I discuss the implications of these results in the Chapter Discussion.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Go trials: go RT							
Part	2	118	752724.90	5564658.50	7.98	**0.002**	0.014
Type	2	118	1048297.00	667567.70	92.65	**< 0.001**	0.020
Distract	1	59	548932.50	536161.00	60.41	**< 0.001**	0.010
Part by type	4	236	60508.50	751780.30	4.75	**0.005**	0.001
Part by distract	2	118	1735.85	334462.00	0.31	0.732	< 0.001
Type by distract	2	118	1253.55	233353.70	0.32	0.665	< 0.001
Part by type by distract	4	236	10025.40	556483.70	1.06	0.360	< 0.001
Go trials: *p*(correct)							
Part	2	118	0.06	0.62	5.54	**0.005**	0.007
Type	2	118	0.32	0.87	21.34	**< 0.001**	0.039
Distract	1	59	0.02	0.26	5.69	**0.020**	0.003
Part by type	4	236	0.03	0.84	2.01	0.123	0.004
Part by distract	2	118	0.04	0.37	6.37	**0.002**	0.005
Type by distract	2	118	0.04	0.37	6.43	**0.008**	0.005
Part by type by distract	4	236	0.05	0.69	4.10	**0.012**	0.006
Go trials: *p*(miss)							
Part	2	118	0.04	0.75	3.15	**0.047**	0.010
Type	2	118	0.16	0.50	19.17	**< 0.001**	0.040
Distract	1	59	0.01	0.13	6.80	**0.012**	0.004
Part by type	4	236	0.02	0.63	2.07	0.132	0.006
Part by distract	2	118	0.01	0.14	4.10	**0.019**	0.003
Type by distract	2	118	0.01	0.20	3.80	**0.047**	0.003
Part by type by distract	4	236	0.01	0.28	1.54	0.217	0.002

Table 3.4. Overview of repeated measures Analyses of Variance on the no-signal data in the *instructed* condition only with part, word type, and distract included as within-participants factors. type = word type. ps < 0.05 are highlighted in bold.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen.*** η^2
Go trials: go RT							
Part	2	118	1095475.00	7445607.40	8.68	**0.002**	0.018
Type	2	118	103420.60	298750.70	20.42	**< 0.001**	0.002
Distract	1	59	769965.60	360523.60	126.01	**< 0.001**	0.013
Part by type	4	236	9216.89	467766.40	1.16	0.323	< 0.001
Part by distract	2	118	6694.47	330280.90	1.20	0.300	< 0.001
Type by distract	2	118	4444.09	154642.80	1.70	0.195	< 0.001
Part by type by distract	4	236	983.74	449817.90	0.13	0.925	< 0.001
Go trials: *p*(correct)							
Part	2	118	0.04	0.51	5.04	**0.009**	0.006
Type	2	118	0.01	0.51	1.65	0.204	0.002
Distract	1	59	0.01	0.14	5.02	**0.029**	0.002
Part by type	4	236	0.00	0.58	0.14	0.943	< 0.001
Part by distract	2	118	0.01	0.33	0.99	0.372	0.001
Type by distract	2	118	0.01	0.27	1.45	0.239	0.001
Part by type by distract	4	236	0.02	0.55	2.44	0.072	0.003
Go trials: *p*(miss)							
Part	2	118	0.00	0.34	0.12	0.873	< 0.001
Type	2	118	0.00	0.16	1.53	0.224	0.001
Distract	1	59	0.02	0.10	10.41	**0.002**	0.006
Part by type	4	236	0.00	0.45	0.37	0.716	0.001
Part by distract	2	118	0.00	0.14	0.36	0.690	< 0.001
Type by distract	2	118	0.01	0.16	2.76	0.091	0.003
Part by type by distract	4	236	0.00	0.27	0.41	0.711	0.001

Table 3.5. Overview of repeated measures Analyses of Variance on the no-signal data in the *uninstructed* condition only with part, word type, and distract included as within-subjects factors. Note: distract = distractor type (no-distractor, distractor); type = word type. type = word type. *p*s < 0.05 are highlighted in bold.

Signal analyses. Tables 3.6-3.7 present an overview of the descriptive and inferential statistics. Separate analyses for the instructed and uninstructed conditions are also presented in Table 3.7. Analyses of probability of responding on stop-signal trials revealed a main effect of word type (20% vs. 80% words), $p < 0.001$ (Table 3.7). Consistent with the stimulus-stop instructional manipulation, there was a reliable two-way interaction between condition and word type ($p < 0.001$) reflecting the larger difference between the 80%-stop words and the 20%-stop words in the instructed condition than in the uninstructed condition. Follow-up tests revealed that the probability of responding was lower for the 80%-stop words than for the 20%-stop words in the instructed condition (80%-stop words: 0.36, 20%-stop words: 0.47; $p < 0.001$) and in the uninstructed condition (80%-stop words: 0.43, 20%-stop words: 0.46; $p = 0.001$). Note that the observed difference in the probability of responding between the 80%-stop words and the 20%-stop words is presumably due to response speed (i.e. participants are slower to respond to the 80%-stop words than to the 20%-stop words) but, importantly, response speed was influenced by the stimulus-stop associations. Thus, the probability of responding on stop trials benefited from the stimulus-stop associations in the instructed and uninstructed conditions. Furthermore, as shown in Table 3.6, the SSD values were similar for the

80%-stop and 20%-stop words (due to the yoked tracking procedure; for details see *Procedure*), showing that the difference in the probability of responding was not due to differences in the SSD.

	***p*(respond)**				**SSD**			
	Distractors		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed								
Part 1								
20%-stop	0.47	0.11	0.40	0.11	486	297	438	167
80%-stop	0.37	0.17	0.32	0.16	490	300	441	169
Part 2								
20%-stop	0.50	0.11	0.48	0.05	463	166	435	196
80%-stop	0.36	0.16	0.36	0.13	466	169	436	199
Part 3								
20%-stop	0.50	0.07	0.50	0.05	448	188	493	269
80%-stop	0.37	0.15	0.35	0.16	452	188	494	270
Uninstructed								
Part 1								
20%-stop	0.42	0.09	0.38	0.10	439	252	439	166
80%-stop	0.44	0.15	0.35	0.11	440	254	441	168
Part 2								
20%-stop	0.49	0.07	0.48	0.05	511	187	435	207
80%-stop	0.44	0.11	0.45	0.13	515	189	435	210
Part 3								
20%-stop	0.51	0.08	0.49	0.06	487	222	455	229
80%-stop	0.46	0.12	0.46	0.12	491	224	457	230

Table 3.6. Overview of the relevant stop-signal data. Probability of responding on a stop trial [*p*(respond| signal)] and average SSD as a function of condition, word type, distractor type, and part. *M* = mean; *SD* = standard deviation.

The probability of responding was lower on no-distractor trials than on distractor trials ($p < 0.001$). Follow-up tests showed that the presence of distractors increased the probability of responding in the instructed condition (no-distractor: 0.40, distractor: 0.43; $p = 0.001$) and in the uninstructed condition (no-distractor: 0.43, distractor: 0.46; $p = 0.001$). However, there was no reliable two-way interaction between word type and distractor type (interaction effect: $p = 0.947$; instructed condition: $p = 0.571$; uninstructed condition: $p = 0.633$).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Main analysis							
Condition	1	118	0.36	5.35	7.94	**0.006**	0.019
Part	2	236	1.04	4.51	27.13	**< 0.001**	0.053
Type (20%-stop, 80%-stop)	1	118	1.95	2.27	101.58	**< 0.001**	0.095
Distract	1	118	0.21	0.99	24.50	**< 0.001**	0.011
Condition by part	2	236	0.09	4.51	2.31	0.109	0.005
Condition by type	1	118	0.68	2.27	35.21	**< 0.001**	0.035
Condition by distract	1	118	0.00	0.99	0.07	0.792	< 0.001
Part by type	2	236	0.12	1.85	7.59	**0.001**	0.006
Part by distract	2	236	0.22	1.60	16.03	**< 0.001**	0.012
Type by distract	1	118	0.00	0.73	0.00	0.947	< 0.001
Condition by part by type	2	236	0.00	1.85	0.20	0.834	< 0.001
Condition by part by distract	2	236	0.00	1.60	0.17	0.840	< 0.001
Condition by type by distract	1	118	0.00	0.73	0.55	0.459	< 0.001
Part by type by distract	2	236	0.03	1.36	2.26	0.106	0.001
Condition by part by type by distract	2	236	0.04	1.36	3.15	**0.045**	0.002
Instructed condition only							
Part	2	118	0.26	2.08	7.46	**0.001**	0.023
Type	1	59	2.47	1.50	97.20	**< 0.001**	0.183
Distract	1	59	0.11	0.55	12.20	**0.001**	0.010
Part by type	2	118	0.07	1.01	3.99	**0.023**	0.006
Part by distract	2	118	0.10	0.92	6.31	**0.004**	0.009
Type by distract	1	59	0.00	0.37	0.33	0.571	< 0.001
Part by type by distract	2	118	0.01	0.61	1.13	0.326	0.001
Uninstructed condition only							
Part	2	118	0.86	2.43	20.94	**< 0.001**	0.102
Type	1	59	0.17	0.77	12.61	**0.001**	0.021
Distract	1	59	0.09	0.44	12.39	**0.001**	0.012
Part by type	2	118	0.05	0.84	3.77	**0.026**	0.007
Part by distract	2	118	0.12	0.68	10.53	**< 0.001**	0.016
Type by distract	1	59	0.00	0.36	0.23	0.633	< 0.001
Part by type by distract	2	118	0.05	0.76	3.97	**0.021**	0.007

Table 3.7. Overview of repeated measures Analyses of Variance on the stop-signal data. Condition is included as a between-subjects factor, all other factors are within-subjects. type = word type. *p*s < 0.05 are highlighted in bold.

Expectancy analyses. Analyses of the expectancy data revealed a main effect of word type ($p < 0.001$). There was also a reliable interaction between word type and condition ($p < 0.001$), indicating that the differences between the word types were larger in the instructed condition ($F(2, 118) = 138.81$, $p < 0.001$, *gen.* $\eta^2 = 0.58$) than in the uninstructed condition ($F(2, 118) = 13.28$, $p < 0.001$, *gen.* $\eta^2 = 0.05$). Consistent with the stimulus-stop contingencies, planned comparisons revealed that in the instructed condition participants expected to stop more for the 80%-stop words (6.66) than for the 0%-stop words (3.70), $t(59) = -12.98$, $p < 0.001$, $g_{av} = 2.62$, and for the 20%-stop words (4.22), $t(59) = -11.27$, $p < 0.001$, $g_{av} = 2.35$. The difference between the 0%-stop words and the 20%-stop words in the uninstructed condition was also reliable, *t*(59) = -5.47, $p < 0.001$, g_{av} = -0.46. Similarly, in the uninstructed condition participants expected to stop more for the 80%-stop words (5.22) than for the 0%-stop words (4.65), *t*(59) = -4.26, $p < 0.001$, g_{av} = 0.50, and for the 20%-stop words (4.75), *t*(59) = -4.01, $p < 0.001$, g_{av} = 0.44. The difference between the 0%-stop words and the 20%-stop words in the uninstructed condition was not reliable, *t*(59) = -1.03, $p = 0.307$, g_{av} = -0.09. Thus, participants in both the instructed condition and the uninstructed condition could distinguish between the 80%-stop words and the other word-types on the basis of their association with stopping.

To investigate the relationship between the expectancy ratings and task performance, I also investigated whether the expectancy ratings correlated with task performance (see Figure 3.3). In the instructed condition, I found that the 80%-stop minus 0%-stop expectancy difference reliably correlated with the corresponding RT difference, *r*(58) = 0.465, $p < 0.001$; similarly, the 80%-stop minus 20%-stop expectancy difference reliably correlated with the corresponding RT difference, *r*(58) = 0.494, $p < 0.001$. In the uninstructed condition, all correlations were not significant (*p*s ≥ 0.501; note that uncorrected *p*s are reported).

To investigate whether the differences between the correlations in the instructed and uninstructed conditions were different, I compared the correlation coefficients in the instructed and uninstructed conditions using the Fisher's r-to-z transformation. This confirmed that the difference between the instructed and uninstructed conditions was reliable for the 80%-stop minus 0%-stop correlations ($Z = 3.59$, $p < 0.001$) and for the 80%-stop minus 20%-stop correlations (Z = 3.58, $p < 0.001$). Furthermore, Bayesian regression analyses supported the null hypothesis that the expectancy ratings did not correlate with the RT slowing in the uninstructed condition (Bayes factor 80%-stop words minus 0%-stop words: 0.28; Bayes factor 80%-stop words minus 20%-stop words: 0.32). Consistent with the ANOVAs, the Bayesian analyses for the instructed condition supported the alternative hypothesis (Bayes factor 80%-stop words minus 0%-stop words: 2701.52; Bayes factor 80%-stop words minus 20%-stop words: 1023.85). Note that inspection of Figure 3.3 shows that the difference between these correlations is unlikely to be entirely due to differences the range of the RT and/or expectancy values between the instructed and uninstructed conditions. Furthermore, this difference cannot be easily explained by the differences in the acquisition rates of the stimulus-stop contingencies between the instructed and uninstructed conditions; RT/expectancy correlations performed on only part 3 (i.e. the final part of training) showed a similar pattern as the aforementioned correlations with all parts of training included. In the instructed condition, I found that the 80%-stop minus 0%-stop expectancy difference reliably correlated with the corresponding RT difference, *r*(58) = 0.419, $p < 0.001$; similarly, the 80%-stop minus 20%-stop expectancy difference reliably correlated with the corresponding RT difference, *r*(58) = 0.341, $p = 0.008$. In the uninstructed condition, all correlations were not significant (*p*s ≥ 0.223). Thus, it seems that although the behavioural effects are the same (i.e. response slowing for stop-associated words), the underlying mechanisms are somewhat different between the instructed and uninstructed conditions.

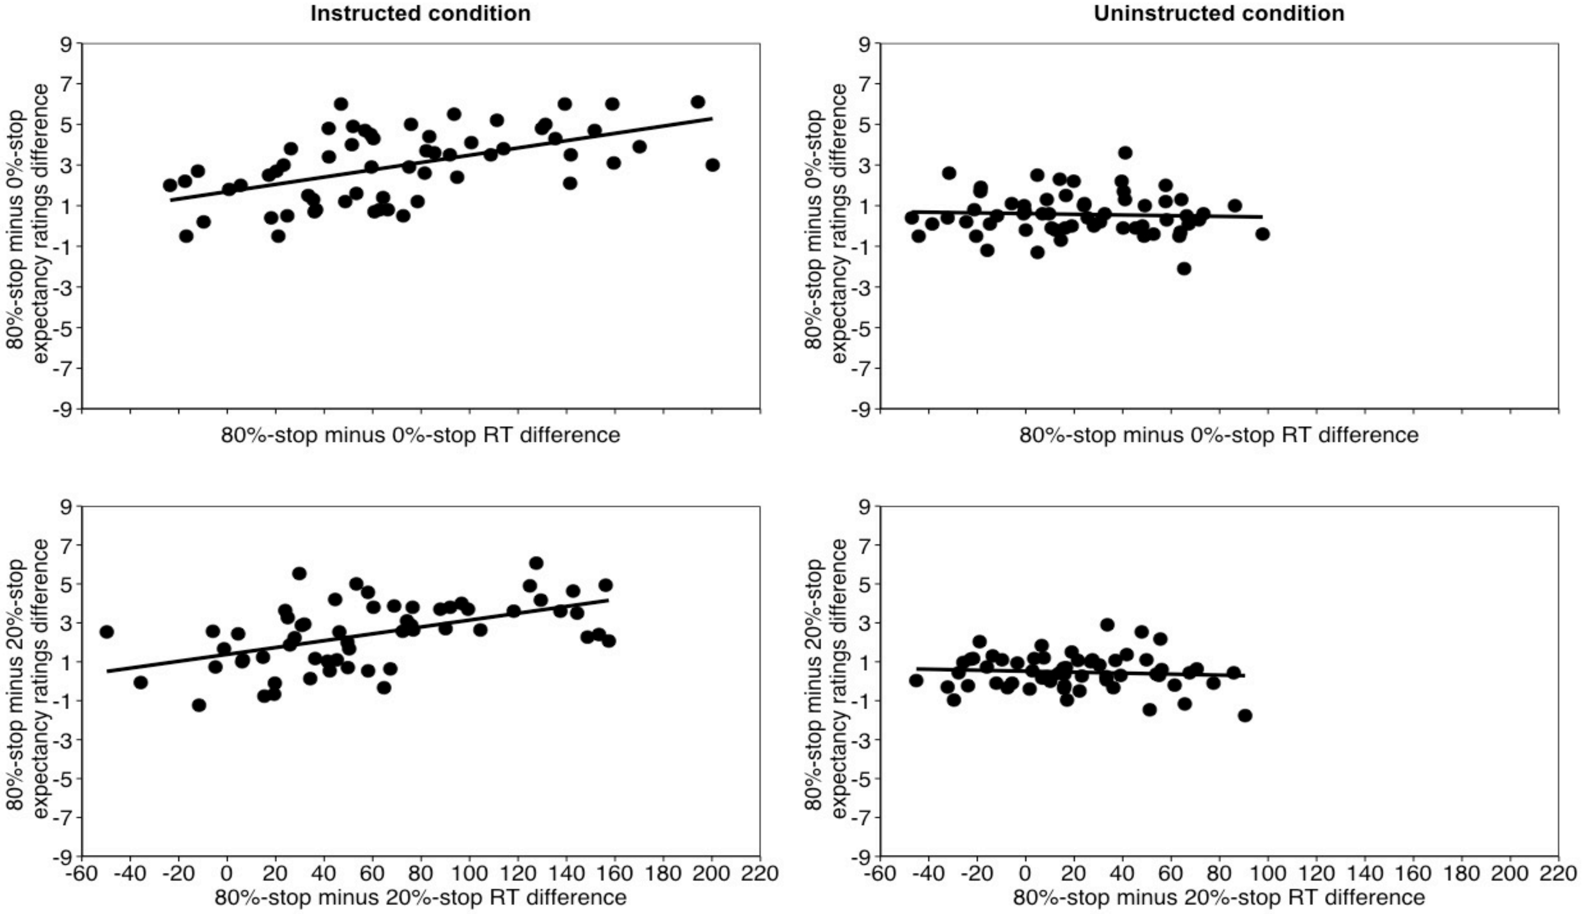

Figure 3.3. Expectancy/RT correlations for the 80%-stop and 0%-stop words (upper panels) and the 80%-stop and 20%-stop words (lower panels) for the instructed condition (left panels) and for the uninstructed condition (right panels). The correlations in the instructed condition were reliable ($ps \leq 0.001$) but the correlations in the uninstructed condition were not reliable ($ps \geq 0.501$).

Chapter Discussion

The findings reported in Chapter 2 show that slowing for old stop words was mediated by expectancies of the stimulus-stop contingencies acquired during training. Previous work has demonstrated that when an instructional cue [e.g. '*p*(stop-signal) = 0.75'] indicates that a stop signal is likely to occur on the following trial(s), participants slow down by proactively adjusting their processing strategies (e.g. Elchlepp, Lavric, et al., 2016; Verbruggen & Logan, 2009a). Therefore, if participants used the stop-associated words as similar cues, slowing for old stop words could reflect anticipatory control adjustments driven by expectancies of the stimulus-stop contingencies in play. The primary purpose of the present chapter was to investigate how expectancies influence performance in an word-specific stop learning task. In addition to measuring expectancy ratings following task completion (as in Chapter 2), I also introduced a between-subjects expectancy manipulation whereby one group of participants were informed about the stimulus-stop contingencies (*instructed* condition) and another group of participants were not told about the stimulus-stop contingencies (*uninstructed* condition). If expectancies play a role in item-specific stop learning, performance in the uninstructed condition should be qualitatively similar to performance in the instructed condition.

The results show that the stimulus-stop contingencies influenced task performance in both the instructed and the uninstructed conditions. As predicted, RTs on no-signal trials were higher and the probability of responding on stop-signal trials was lower for the 80%-stop words than for the 0%-stop words and for the 20%-stop words. As expected, the effects of word type on task performance were greater in the instructed condition compared with the uninstructed condition. Furthermore, participants in both conditions generated expectancies that were consistent with the stimulus-stop contingencies in play: participants expected to stop their responses more when 80%-stop words were presented than when 0%-stop words or 20%-stop words were presented. But, as seen in the task performance data, the differences between word types on the expectancy ratings were greater in the instructed condition than in the uninstructed condition (as predicted). It is important to highlight that I also found a reliable difference in expectancies between the 0%-stop words and the 20%-stop words in the instructed condition even though I did not tell participants which words were the 0%-stop words or the 20%-stop words. Similarly, there was a reliable difference between the 0%-stop words and the 20%-stop words in the task performance data. In other words, extra learning occurred in the instructed condition. However, these differences may not necessarily reflect implicit learning; in the uninstructed condition, the difference between 0% and 20% was not significant (and the difference in the instructed condition was significantly greater than the difference in the uninstructed condition). It is possible that informing participants in the instructed condition about the 80%-stop words encouraged them to look for further stimulus-stop contingencies and identify that the probability of stopping was greater for the 20%-stop words than for the 0%-stop words. Combined, these results suggest that stimulus-stop associations acquired through explicit instructions and stimulus-stop associations acquired through task practice have qualitatively similar effects on stop-signal task performance (i.e. both induce slowing on no-signal trials and reduce the probability of responding on stop-signal trials). However, as predicted, the effects are much stronger when the stop contingencies are instructed.

It is important to highlight that I also found some differences between the instructed and uninstructed conditions. In the instructed condition, I found that the expectancy ratings correlated with task performance: participants who expected to withhold their response more for the 80%-stop words responded more slowly to these words than to the 0%-stop words and to the 20%-stop words. In contrast, the expectancy ratings in the uninstructed condition did not reliably correlate with task performance. This conclusion was further supported by Bayesian regression analyses. In other words, participants did generate some expectancies but they did not use these expectancies to influence task performance. In Chapter 2, I also found that the stop-associated words slowed responses, even when participants did not expect a stop signal for those words. I suggested that when the stimulus-stop contingencies are uninstructed, stimulus-stop learning is mediated via an interplay of explicit and implicit processes. The findings reported in the present chapter provide additional support for this conclusion: it seems that, unlike in the instructed condition, the response slowing observed in the uninstructed condition reflects at least *some* contribution of bottom-up (i.e. implicit) processes independent of top-down expectancies. Indeed, the absence of a correlation in the uninstructed condition satisfies the *zero correlation criterion* which states that under situations of unconscious (i.e. implicit) learning there should be no relationship between confidence judgments (i.e. expectancies) and task performance (Dienes, Altmann, Kwan & Goode, 1995; Dienes, 2007). The idea that response inhibition can become automatised independent of top-down expectancy is consistent with the Instance theory which construes automaticity as the single-step direct retrieval of past solutions from memory (Logan, 1988). Furthermore, research in the associative learning literature has shown that performance on various implicit and explicit measures are separable, and not strongly correlated (Gebauer & Mackintosh, 2007; Yang & Li, 2012; see also McLaren, Forrest, et al., 2014). Therefore, it is possible that performance in the uninstructed condition is primarily influenced by (implicit) associations, but that the induction of expectancies is explicit and hence expectancies not correlated with performance.

However, it is important to highlight that in Experiment 1 (Chapter 2) I found some evidence that expectancies obtained following task completion correlated with the magnitude of the response slowing following the reversal of the acquired stimulus-stop mappings in a test phase following training (Experiment 1). Thus, it is possible that although expectancies do not drive the acquisition of slowing, they constitute a measure of the strength of stimulus-stop learning on performance. As such, they could reliably predict the magnitude of slowing for the old stop-associated items in the test phase. Furthermore, in Experiment 4 (Chapter 2) when expectancies were obtained on each trial, the expectancies ratings correlated with task performance during the training and test phases. As discussed in Chapter 2, it could be that asking participants to rate their expectancies on each trial provided a more sensitive measure of the acquisition of the stimulus-stop expectancies. Alternatively, it could be that asking participants to rate their expectancies on each trial changed the nature of the task and encouraged a more deliberate form of processing. Note, however, that the latter explanation cannot account for all of the stop learning effects observed in that experiment (for a discussion, see Chapter 2).

In the Verbruggen, Stevens, et al. study, expectancy was manipulated in a block-based fashion (i.e. participants were informed at the beginning of a block if signals could occur in the periphery or in the centre of the screen), and proactive control adjustments increased the distractor effect in no-signal RTs. In the present chapter, word-specific stop learning did not interact with the distractor effect in no-signal RTs. I observed this pattern of results in both the instructed and uninstructed conditions (even though I observed greater stop learning effects in the instructed condition). Note that I also replicated the pattern of results observed in the uninstructed condition in an earlier pilot experiment (the results of this pilot experiment are reported in Appendix G). There appear to be three plausible explanations for the absence of an interaction between expectancy and the distractor manipulation in no-signal RTs. First, it is possible that participants simply cannot adjust their attentional settings on a trial-by-trial basis (see Strayer & Kramer, 1994, for a similar idea). Recently, I tested this idea in a different (unpublished) experiment where there were ‘pure’ blocks in which the location of the signal stayed the same (non-central signal blocks, central signal blocks, and no-signal blocks as in Verbruggen, Stevens, et al. 2014) and ‘mixed’ blocks in which the location of the signal changed on a trial-by-trial basis. In both block types, the location of the upcoming signal was instructed by way of a cue presented at the beginning of each trial. I did not find any differences in the size of the distractor effect across the signal locations in the pure blocks or in the mixed blocks. It is possible that the presentation of the cue on each trial encouraged participants to attend to the centre of the screen in that experiment. Alternatively, it could be that signal detection processes during stopping play less of a role in the current task than originally thought. If it turns out that participants cannot make attentional adjustments on a trial-by-trial basis in this task, these findings would conflict with previous research that has demonstrated that participants can make trial-by-trial control adjustments in other control paradigms. For example, evidence in the task-switching literature has demonstrated that participants can reconfigure their attentional settings when a cue indicates that a task switch is required (see, e.g. Longman, Lavric, & Monsell, 2013; Longman, Lavric, Munteanu & Monsell, 2014; Rushworth, Passingham, & Nobre, 2002).

Second, it is possible that in the previous study by Verbruggen, Stevens, et al. (2014), the presence of distractors primarily interfered with the detection or analysis of the go words rather than response selection. This could explain the difference between block-based and word-based expectancy manipulations because unlike block-based expectancies, word-based expectancies can only influence processing *after* the word has been identified (assuming a serial architecture e.g. stimulus X -> stop word -> stop). Hence, the distractor manipulation and the word manipulation should have additive effects.

Third, it is possible that participants formed direct associations between the stop-associated words and stop response (without a mediating role for the stop signal). If this were the case, stop-associated words would not induce attentional shifts and no increased distractor effect is expected. However, it is not possible to distinguish between these explanations in the present chapter.

The RT analysis showed that that word type did not influence the distractor effect. However, I did find that the distractor effect was larger for the 80%-stop words than the other words types in the probability of missed responses in the instructed condition. However, this is unlikely to reflect proactive adjustments of attentional settings. First, the absence of an interaction between word type and distractor type in no-signal RTs argues against the explanation that participants allocated attention to the stop signal and consequently got timed out, especially in Experiment 6 where I extended the response deadline to 2000 ms (considerably longer than overall mean RTs). Second, the eye-movement data (Appendix F) suggests that participants made *fewer* eye movements following the presentation of 80%-stop words on distractor trials than on no-distractor trials. These data support the idea that participants used the instructed stimulus-stop contingency more on distractor trials (when it was difficult to monitor for the stop signal) than on no-distractor trials (when it was easier to monitor for the stop signal). Hence, on no-distractor trials the decision to stop or not was primarily influenced by the actual presentation of the stop signal rather than the stimulus-stop contingencies. Finally, the probability of missed responses was low suggesting that, even if we assume that these differences reflect anticipatory adjustments, this happened on only a small minority of no-signal trials.

The absence an interaction between the distractor manipulation and the word type manipulation also sheds further light on what is learned in stop-signal tasks. In Chapters 1 and 2, I discussed the idea that in stop-learning tasks that use a single representation of the stop signal (e.g. the lines of the surrounding square becoming thicker) participants could acquire associations between the stop-associated words and the stop signal; retrieval of these associations will prime the representation of the stop signal rather than the stop goal or stop response (i.e. the ‘indirect’ pathway; see also Bowditch, Verbruggen, et al., 2016; Verbruggen, Best, et al., 2014). This signal priming account predicts that stop learning will influence the probability of stopping in training without influencing responding on no-signal trials. In the paradigm used in the present chapter, the signal-detection component may be more important than in other variants of the stop-signal paradigm. Nevertheless, I found that both go (no-signal) and stop performance were influenced by word-specific learning. Furthermore, if participants learned a stimulus-signal association, the distractors effect should be larger for 80%-stop words than for the 20%-stop words due to the increased priming of stop-signal detection for the stop-associated words. As can be seen in Tables 3.3-3.6, distractors and word type had additive effects on *p*(respond|signal). This is inconsistent with the stop-signal priming idea.

Finally, the analyses of task performance on no-signal trials revealed an unexpected finding. In addition to the predicted stop learning effect on no-signal RTs, I found that learning influenced accuracy in the instructed condition: accuracy was lower for the 80%-stop words than for the 0%-stop words and the 20%-stop words. Note, missed trials were excluded from the analyses of accuracy rates so these findings cannot be accounted for by the observed differences in the probability of missed responses on go trials. I did not find any evidence of this accuracy effect in the uninstructed condition. This accuracy effect cannot be accounted for by a speed-accuracy tradeoff as no-signal RTs were reliably slower for the stop-associated words than for the other word types in both conditions. Instead, these accuracy effects may be explained by the blocked-input model of response inhibition (Logan, Yamaguchi, et al., 2015). According to this model, response inhibition can be achieved by blocking input to go units. It is possible that the retrieval of strong stimulus-stop associations in the instructed condition had a similar effect. This could explain why accuracy was lower for stop-associated words in the instructed condition: when input to the relevant go unit is reduced (as specified by the word itself), the probability of an incorrect response (and indeed the probability of a missed response) may increase. However, future research is required to test this blocked-input idea (for an extended discussion of the implications of this blocked-input idea, see *How do Stimulus-Stop Associations Influence Behaviour?*, Chapter 7) .

Conclusion

In conclusion, these results suggest that slowing for stimuli than are consistently associated with stopping can arise from explicit, top-down expectancies or from learning through experience during task practice. Given that I found effects of the stimulus-stop contingencies on expectancies in the instructed and uninstructed conditions, it is tempting to argue that expectancies are required for word-specific stop learning. However, Bayesian analyses showed that expectancies correlated with response slowing on no-signal trials in the instructed condition but not in the uninstructed condition. In other words, the slowing for the 80%-stop words observed in the instructed condition can be accounted for primarily by the instructed expectancies; on the other hand, the slowing for the 80%-stop words observed in the uninstructed condition can be accounted for primarily by implicit learning processes (the expectancies observed were neither necessary nor sufficient to elicit the response slowing and can be considered a by-product of learning). Thus, stimulus-stop learning may well be mediated via both explicit and implicit processes.

Appendix D

Separate Analyses of Experiment 5 and Experiment 6

In Appendix D, the separate analyses for Experiment 5 and Experiment 6 are presented. There were three differences between these experiments: (1) sample sizes ($N = 48$ in Experiment 5, $N = 72$ in Experiment 6), (2) response deadlines (1250 ms in Experiment 5, 2000 ms in Experiment 6), and (3) dependent measures obtained (with eye-tracking in Experiment 5, without eye-tracking in Experiment 6).

Experiment 5. Descriptive statistics for the task performance data are presented in Tables D.1 and D.2, inferential statistics are presented in Table D.3 and D.5 (separate analyses for the instructed and uninstructed conditions are presented in Tables D.6 and D.7).

Task performance. I found a reliable main effect of word type on no-signal RTs ($p < 0.001$), on the probability of correct responses ($p = 0.001$), the probability of missed responses ($p < 0.001$), and the probability of responding on stop-signal trials ($p < 0.001$). The main effect of distractor type (distractor, no-distractor) was also reliable in each of the aforementioned measures (ps ≤ 0.001). The interactions between condition (instructed, uninstructed) and word type were reliable (ps ≤ 0.006) reflecting the greater differences between the word types in the instructed condition than in the uninstructed condition. The word type by distractor interactions were not reliable in go RTs, the probability of missed responses, and the probability of responding on stop-signal trials (ps ≥ 0.418) but was reliable in the probability of correct responses ($p < 0.001$). The p(correct) interaction was also reliable in the combined analysis; for possible explanations, see the Chapter Discussion. Table D.4 shows the outcome of the Bayesian analysis. As can be seen, the results are largely consistent with the ANOVA reported in Tables D.3.

	***p*(correct)**				***p*(miss)**				**go RT**			
	Distractors		***No distractors***		***Distractors***		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed												
Part 1												
0%-stop	0.92	0.06	0.90	0.06	0.05	0.04	0.03	0.03	781	69	726	64
20%-stop	0.92	0.06	0.91	0.05	0.06	0.04	0.04	0.03	786	66	730	66
80%-stop	0.89	0.14	0.86	0.15	0.14	0.13	0.07	0.09	831	99	756	79
Part 2												
0%-stop	0.92	0.06	0.94	0.05	0.03	0.03	0.03	0.03	754	83	705	89
20%-stop	0.93	0.05	0.93	0.05	0.02	0.02	0.02	0.02	767	83	717	94
80%-stop	0.90	0.11	0.87	0.11	0.07	0.09	0.04	0.05	825	108	770	120
Part 3												
0%-stop	0.93	0.06	0.92	0.06	0.04	0.04	0.04	0.05	767	94	710	105
20%-stop	0.92	0.06	0.92	0.06	0.05	0.06	0.03	0.04	779	107	725	101
80%-stop	0.92	0.11	0.83	0.14	0.11	0.13	0.07	0.10	835	139	773	145
Uninstructed												
Part 1												
0%-stop	0.90	0.08	0.89	0.09	0.06	0.05	0.04	0.03	815	90	757	80
20%-stop	0.89	0.09	0.88	0.09	0.07	0.04	0.04	0.03	819	86	758	81
80%-stop	0.90	0.11	0.88	0.11	0.07	0.09	0.07	0.08	829	92	757	99
Part 2												
0%-stop	0.91	0.08	0.91	0.09	0.06	0.07	0.03	0.04	809	138	754	131
20%-stop	0.91	0.07	0.91	0.07	0.05	0.05	0.04	0.03	807	129	753	137
80%-stop	0.94	0.12	0.89	0.10	0.04	0.08	0.05	0.08	823	140	785	133
Part 3												
0%-stop	0.89	0.10	0.90	0.07	0.06	0.06	0.04	0.04	810	128	752	134
20%-stop	0.90	0.08	0.91	0.06	0.06	0.07	0.04	0.04	812	130	748	132
80%-stop	0.89	0.15	0.89	0.14	0.04	0.07	0.04	0.06	823	137	756	143

Table D.1. Overview of the no-signal data in Experiment 5. Probability of an accurate go response [*p*(correct)], probability of a missed go response [*p*(miss)], and average go reaction time (RT) as a function of condition, word type, distractor type, and part. Accuracy is the ratio of correct no-signal trials to the number of correct and incorrect no-signal trials (missed trials are excluded). *P*(miss) is the ratio of omitted responses to the total number of no-signal trials. *M* = mean; *SD* = standard deviation.

	***p*(respond)**				**SSD**			
	Distractors		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed								
Part 1								
20%-stop	0.48	0.08	0.41	0.06	366	130	360	114
80%-stop	0.39	0.16	0.29	0.14	369	128	357	114
Part 2								
20%-stop	0.52	0.07	0.48	0.04	373	166	333	113
80%-stop	0.38	0.16	0.39	0.13	376	166	333	112
Part 3								
20%-stop	0.50	0.06	0.51	0.04	351	174	393	111
80	0.37	0.12	0.35	0.12	354	177	393	111
Uninstructed								
Part 1								
20%-stop	0.46	0.08	0.43	0.09	375	145	329	85
80%-stop	0.48	0.18	0.41	0.10	377	144	330	87
Part 2								
20%-stop	0.51	0.08	0.49	0.05	366	156	286	136
80%-stop	0.44	0.11	0.46	0.14	371	153	287	137
Part 3								
20%-stop	0.52	0.08	0.49	0.04	347	161	416	131
80%-stop	0.48	0.10	0.44	0.11	349	164	417	131

Table D.2. Overview of the relevant stop-signal data in Experiment 5. Probability of responding on a stop trial [*p*(respond|signal)] and average SSD as a function of condition, word type, distractor type, and part. *M* = mean; *SD* = standard deviation.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Go trials: go RT							
Condition	1	46	123926.20	7639132.60	0.75	0.392	0.012
Part	2	92	6022.55	1209571.42	0.23	0.705	0.001
Type (0%-, 20, 80)	2	92	208930.80	260449.40	36.90	**< 0.001**	0.020
Distract	1	46	720715.70	157581.96	210.39	**< 0.001**	0.067
Condition by part	2	92	7831.28	1209571.42	0.30	0.655	0.001
Condition by type	2	92	79512.53	260449.40	14.04	**< 0.001**	0.008
Condition by distract	1	46	141.16	157581.96	0.04	0.840	< 0.001
Part by type	4	184	13498.73	225370.39	2.76	0.056	0.001
Part by distract	2	92	6570.54	147124.22	2.05	0.134	0.001
Type by distract	2	92	1324.51	80018.08	0.76	0.418	< 0.001
Condition by part by type	4	184	3870.41	225370.39	0.79	0.533	< 0.001
Condition by part by distract	2	92	654.58	147124.22	0.20	0.815	< 0.001
Condition by type by distract	2	92	1174.59	80018.08	0.68	0.512	< 0.001
Part by type by distract	4	184	3844.43	295307.76	0.60	0.588	< 0.001
Condition by part by type by distract	4	184	1284.19	295307.76	0.20	0.865	< 0.001
Go trials: *p*(correct)							
Condition	1	46	0.01	4.07	0.15	0.698	0.002
Part	2	92	0.04	0.35	5.73	**0.005**	0.006
Type (0, 20, 80)	2	92	0.13	0.64	9.48	**0.001**	0.018
Distract	1	46	0.05	0.18	12.73	**0.001**	0.007
Condition by part	2	92	0.00	0.35	0.23	0.792	< 0.001
Condition by type	2	92	0.09	0.64	6.75	**0.006**	0.013
Condition by distract	1	46	0.00	0.18	1.15	0.289	0.001
Part by type	4	184	0.00	0.55	0.31	0.819	0.001
Part by distract	2	92	0.00	0.34	0.27	0.733	< 0.001
Type by distract	2	92	0.06	0.29	8.81	**< 0.001**	0.008
Condition by part by type	4	184	0.01	0.55	0.47	0.703	0.001
Condition by part by distract	2	92	0.02	0.34	3.03	0.061	0.003
Condition by type by distract	2	92	0.01	0.29	1.20	0.293	0.001
Part by type by distract	4	184	0.01	0.63	0.99	0.388	0.002
Condition by part by type by distract	4	184	0.02	0.63	1.26	0.289	0.002
Go trials: *p*(miss)							
Condition	1	46	0.00	1.06	0.04	0.846	< 0.001
Part	2	92	0.07	0.55	5.45	**0.008**	0.019
Type (0, 20, 80)	2	92	0.11	0.26	19.46	**< 0.001**	0.033
Distract	1	46	0.06	0.08	33.63	**< 0.001**	0.018
Condition by part	2	92	0.02	0.55	1.45	0.240	0.005
Condition by type	2	92	0.09	0.26	16.18	**< 0.001**	0.027
Condition by distract	1	46	0.00	0.08	1.48	0.229	0.001
Part by type	4	184	0.02	0.60	1.68	0.184	0.007
Part by distract	2	92	0.01	0.18	2.36	0.107	0.003
Type by distract	2	92	0.00	0.20	0.61	0.480	0.001
Condition by part by type	4	184	0.01	0.60	0.62	0.572	0.002
Condition by part by distract	2	92	0.01	0.18	1.43	0.245	0.002
Condition by type by distract	2	92	0.05	0.20	12.58	**< 0.001**	0.016
Part by type by distract	4	184	0.01	0.35	0.72	0.498	0.002
Condition by part by type by distract	4	184	0.00	0.35	0.30	0.754	0.001

Table D.3. Overview of repeated measures Analyses of Variance on the no-signal data in Experiment 5. Condition is included as a between-subjects factor, all other factors are within-subjects. In the no-signal RT analysis, incorrect and missed no-signal trials were removed. type = word type. *p*s < 0.05 are highlighted in bold.

Omitted factor(s)	Bayes Factor	Confidence interval
Main analysis		
Distract	0.00	±126.46%
Type	0.00	±126.61%
Condition:type	0.00	±100.21%
Condition	0.04	±103.06%
Distract by part	0.23	±100.01%
Condition by distract by part by type	0.44	±99.77%
Part	0.53	±104.66%
Condition by distract by part	0.61	±103.96%
Condition by distract by type	0.66	±112.99%
Part by type	0.92	±101.07%
Distract by type	1.26	±117.18%
Condition by distract	1.31	±131.65%
Condition by part	2.10	±127.69%
Distract by part by type	5.45	±100.85%
Condition by part by type	15.47	±116.17%
Instructed condition only		
Distract	0.00	±20.90%
Type	0.00	±20.70%
Distract:part	18.30	±27.20%
Part	19.16	±18.78%
Distract by type	26.77	±24.48%
Distract by part by type	33.41	±17.65%
Part by type	63.26	±21.04%
Uninstructed condition only		
Distract	0.00	±56.17%
Type	3.17	±30.79%
Distract by part	11.37	±29.60%
Distract by type	20.19	±35.73%
Distract by part by type	26.55	±31.2%
Part	78.87	±43.83%
part by type	87.01	±30.20%

Table D.4. No-signal RT Bayesian analysis in Experiment 5. Bayes factors < 1 indicate that the removal of the factor or interaction had a deleterious effect on the model, whereas Bayes factors > 1 indicate that the factor or interaction could be removed without impairing the fit much. Note that ‘participant’ was included as a factor for all models, but this factor is not added to the model descriptions in the tables to reduce the amount of text. distract = distractor type (no-distractor, distractor); type = word type.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Stop-signal trials: *p*(respond)							
Condition	1	46	0.28	1.52	8.48	**0.006**	0.044
Part	2	92	0.22	1.04	9.63	**< 0.001**	0.034
Type (20%-, 80%-stop)	1	46	0.86	1.01	39.23	**< 0.001**	0.123
Distract	1	46	0.13	0.39	14.87	**< 0.001**	0.020
Condition by part	2	92	0.01	1.04	0.52	0.562	0.002
Condition by type	1	46	0.29	1.01	13.27	**0.001**	0.045
Condition by distract	1	46	0.00	0.39	0.13	0.717	< 0.001
Part by type	2	92	0.04	0.62	3.05	0.052	0.007
Part by distract	2	92	0.09	0.62	6.61	**0.002**	0.014
Type by distract	1	46	0.00	0.28	0.29	0.591	< 0.001
Condition by part by type	2	92	0.01	0.62	0.54	0.586	0.001
Condition by part by distract	2	92	0.01	0.62	1.03	0.363	0.002
Condition by type by distract	1	46	0.00	0.28	0.00	0.975	< 0.001
Part by type by distract	2	92	0.04	0.65	2.96	0.057	0.007
Condition by part by type by distract	2	92	0.00	0.65	0.28	0.759	0.001

Table D.5. Overview of repeated measures Analyses of Variance on the stop-signal data in Experiment 5. Condition is included as a between-subjects factor, all other factors are within-subjects. type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Go trials: go RT							
Part	2	46	11078.02	547919.98	0.47	0.538	0.003
Type (0%-, 20%-, 80%-stop)	2	46	272678.70	185040.03	33.89	**< 0.001**	0.065
Distract	1	23	350341.90	67823.77	118.81	**< 0.001**	0.081
Part by type	4	92	12048.98	112538.93	2.46	0.082	0.003
Part by distract	2	46	1987.93	63177.04	0.72	0.490	0.001
Type by distract	2	46	2376.33	40560.40	1.35	0.265	0.001
Part by type by distract	4	92	933.74	141689.52	0.15	0.914	< 0.001
Go trials: *p*(correct)							
Part	2	46	0.01	0.15	2.20	0.123	0.004
Type (0%-, 20%-, 80%-stop)	2	46	0.22	0.44	11.66	**0.001**	0.066
Distract	1	23	0.04	0.09	10.57	**0.004**	0.013
Part by type	4	92	0.00	0.24	0.28	0.818	0.001
Part by distract	2	46	0.02	0.18	2.15	0.139	0.005
Type by distract	2	46	0.05	0.13	9.55	**0.002**	0.016
Part by type by distract	4	92	0.02	0.35	1.53	0.226	0.007
Go trials: *p*(miss)							
Part	2	46	0.07	0.39	4.17	**0.024**	0.037
Type (0%-, 20%-, 80%-stop)	2	46	0.20	0.20	23.16	**< 0.001**	0.100
Distract	1	23	0.04	0.03	29.88	**< 0.001**	0.024
Part by type	4	92	0.02	0.32	1.50	0.230	0.011
Part by distract	2	46	0.01	0.09	3.46	**0.044**	0.008
Type by distract	2	46	0.04	0.10	9.43	**0.003**	0.021
Part by type by distract	4	92	0.01	0.18	0.70	0.595	0.003
Stop-signal trials: *p*(respond)							
Part	2	46	0.14	0.41	8.10	**0.001**	0.042
Type (20%-, 80%-stop)	1	23	1.07	0.68	36.03	**< 0.001**	0.249
Distract	1	23	0.08	0.15	11.34	**0.003**	0.023
Part by type	2	46	0.02	0.31	1.52	0.229	0.006
Part by distract	2	46	0.07	0.31	5.32	**0.008**	0.022
Type by distract	1	23	0.00	0.15	0.12	0.730	< 0.001
Part by type by distract	2	46	0.03	0.28	2.61	0.088	0.010

Table D.6. Overview of repeated measures Analyses of Variance on the *instructed* condition only in Experiment 5 with part, word type, and distract included as within-participants factors. type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Go trials: go RT							
Part	2	46	2775.81	661651.44	0.10	0.845	< 0.001
Type (0%-, 20%-, 80%-stop)	2	46	15764.69	75409.38	4.81	**0.025**	0.003
Distract	1	23	370514.90	89758.19	94.94	**< 0.001**	0.058
Part by type	4	92	5320.16	112831.46	1.08	0.356	0.001
Part by distract	2	46	5237.20	83947.19	1.43	0.249	0.001
Type by distract	2	46	122.76	39457.68	0.07	0.837	< 0.001
Part by type by distract	4	92	4194.88	153618.24	0.63	0.565	0.001
Go trials: *p*(correct)							
Part	2	46	0.03	0.21	3.54	**0.037**	0.008
Type (0%-, 20%-, 80%-stop)	2	46	0.00	0.20	0.20	0.772	< 0.001
Distract	1	23	0.01	0.09	3.17	0.088	0.003
Part by type	4	92	0.01	0.31	0.47	0.724	0.002
Part by distract	2	46	0.01	0.17	1.11	0.334	0.002
Type by distract	2	46	0.01	0.17	1.55	0.228	0.003
Part by type by distract	4	92	0.01	0.28	0.63	0.588	0.002
Go trials: *p*(miss)							
Part	2	46	0.01	0.16	1.69	0.200	0.008
Type (0%-, 20%-, 80%-stop)	2	46	0.00	0.06	0.64	0.480	0.001
Distract	1	23	0.02	0.05	8.92	**0.007**	0.013
Part by type	4	92	0.01	0.28	0.75	0.508	0.006
Part by distract	2	46	0.00	0.08	0.13	0.858	< 0.001
Type by distract	2	46	0.02	0.10	3.82	**0.046**	0.011
Part by type by distract	4	92	0.00	0.17	0.31	0.776	0.002
Stop-signal trials: *p*(respond)							
Part	2	46	0.09	0.64	3.15	0.066	0.029
Type (20%-, 80%-stop)	1	23	0.08	0.32	5.38	**0.030**	0.025
Distract	1	23	0.05	0.24	5.02	**0.035**	0.018
Part by type	2	46	0.03	0.31	2.06	0.139	0.009
Part by distract	2	46	0.03	0.31	2.32	0.109	0.011
Type by distract	1	23	0.00	0.13	0.17	0.680	< 0.001
Part by type by distract	2	46	0.01	0.38	0.88	0.419	0.005

Table D.7. Overview of repeated measures Analyses of Variance on the *uninstructed* condition only in Experiment 5 with part, word type and distract included as within-subjects factors. Note: distract = distractor type (no-distractor, distractor); type = word type. *p*s < 0.05 are highlighted in bold.

Expectancy analyses. Analyses of the expectancy data revealed a main effect of word type ($p < 0.001$) and a reliable interaction between word type and condition ($p < 0.001$), reflecting the larger difference between the word types in the instructed condition ($p < 0.001$) than in the uninstructed condition ($p = 0.001$). Planned comparisons revealed that in the instructed condition participants expected to stop more for the 80%-stop words (6.74) than for the 0%-stop words (3.61), $t(23) = 9.64$, $p < 0.001$, $g_{av} = 3.08$, and for the 20%-stop words (4.10), $t(23) = 9.21$, $p < 0.001$, $g_{av} = 2.97$. The difference between the 0%-stop words and the 20%-stop words was also reliable, $t(23) = -3.33$, $p = 0.003$, $g_{av} = -0.47$. Similarly, in the uninstructed condition participants expected to stop more for the 80%-stop words (5.08) than for the 0%-stop words (4.35), $t(23) = 2.91$, $p = 0.008$, $g_{av} = 0.53$, and for the 20%-stop words (4.34), $t(23) = 3.79$, $p < 0.001$, $g_{av} = 0.59$. The difference between the 0%-stop words and the 20%-stop words was not reliable, $t(23) = 0.022$, $p = 0.983$, $g_{av} = 0.00$.

Experiment 6. Descriptive statistics for the task performance data are presented in Tables D.8 and D.9, inferential statistics are presented in Table D.10 and D.12 (separate analyses for the instructed and uninstructed conditions are presented in Tables D.13 and D.14).

Task performance. I found a reliable main effect of word type on no-signal RTs ($p < 0.001$), the probability of correct responses ($p = 0.001$), the probability of missed responses ($p = 0.043$), and the probability of responding on stop-signal trials ($p < 0.001$). The main effect of distractor type (distractor, no-distractor) was also reliable for no-signal RTs, the probability of correct responses, and the probability of responding on stop trials (ps ≤ 0.002), but not for the probability of missed no-signal responses ($p = 0.440$). The interactions between condition (instructed, uninstructed) and word type were reliable for no-signal RTs ($p < 0.001$) and the probability of responding on stop trials ($p < 0.001$) reflecting the greater differences between the word types in the instructed condition than in the uninstructed condition. There was no reliable two-way interactions between word type and distractor type (ps ≥ 0.119). Table D.11 shows the outcome of the Bayesian analysis. As can be seen, the results are largely consistent with the ANOVAs reported in Tables D.10.

	_p_(correct)				**_p_(miss)**				**go RT**			
	Distractors		***No distractors***		***Distractors***		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed												
Part 1												
0%-stop	0.93	0.07	0.93	0.06	0.02	0.03	0.02	0.03	913	161	880	158
20%-stop	0.93	0.05	0.92	0.06	0.02	0.03	0.02	0.03	926	159	885	149
80%-stop	0.90	0.13	0.91	0.11	0.05	0.10	0.04	0.10	969	194	927	145
Part 2												
0%-stop	0.93	0.06	0.95	0.07	0.02	0.03	0.02	0.04	968	259	922	237
20%-stop	0.94	0.05	0.94	0.06	0.02	0.03	0.02	0.03	982	264	936	238
80%-stop	0.92	0.09	0.94	0.09	0.03	0.08	0.03	0.08	1041	280	1027	262
Part 3												
0%-stop	0.93	0.06	0.93	0.08	0.02	0.03	0.02	0.03	1016	276	965	263
20%-stop	0.93	0.06	0.93	0.07	0.02	0.02	0.02	0.02	1018	276	985	264
80%-stop	0.91	0.10	0.85	0.16	0.02	0.04	0.03	0.06	1102	316	1075	311
Uninstructed												
Part 1												
0%-stop	0.95	0.04	0.94	0.05	0.01	0.02	0.01	0.02	927	164	870	160
20%-stop	0.95	0.04	0.94	0.03	0.02	0.02	0.01	0.02	943	164	879	156
80%-stop	0.93	0.08	0.95	0.08	0.01	0.03	0.02	0.05	953	162	904	166
Part 2												
0%-stop	0.95	0.05	0.96	0.04	0.02	0.03	0.01	0.02	1018	292	982	283
20%-stop	0.96	0.04	0.96	0.04	0.02	0.03	0.02	0.03	1032	303	974	275
80%-stop	0.96	0.09	0.93	0.09	0.03	0.09	0.02	0.06	1054	320	1010	303
Part 3												
0%-stop	0.95	0.05	0.95	0.06	0.02	0.04	0.02	0.04	1056	307	1010	297
20%-stop	0.95	0.05	0.95	0.06	0.03	0.04	0.02	0.03	1058	311	1003	289
80%-stop	0.94	0.11	0.94	0.10	0.03	0.05	0.02	0.05	1075	308	1035	293

Table D.8. Overview of the no-signal data in Experiment 6. Probability of an accurate go response [*p*(correct)], probability of a missed go response [*p*(miss)], and average go reaction time (RT) as a function of condition, word type, distractor type, and part. Accuracy is the ratio of correct no-signal trials to the number of correct and incorrect no-signal trials (missed trials are excluded). *P*(miss) is the ratio of omitted responses to the total number of no-signal trials. *M* = mean; *SD* = standard deviation.

	***p*(respond)**				**SSD**			
	Distractors		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed								
Part 1								
20%-stop	0.46	0.13	0.39	0.13	465	173	587	214
80%-stop	0.35	0.17	0.33	0.17	465	178	591	215
Part 2								
20%-stop	0.48	0.12	0.47	0.06	513	205	619	221
80%-stop	0.35	0.16	0.34	0.12	519	210	623	222
Part 3								
20%-stop	0.50	0.08	0.49	0.06	493	258	465	162
80%-stop	0.36	0.17	0.36	0.18	495	261	467	167
Uninstructed								
Part 1								
20%-stop	0.39	0.09	0.35	0.09	565	224	613	300
80%-stop	0.41	0.12	0.32	0.10	570	227	615	300
Part 2								
20%-stop	0.48	0.06	0.47	0.06	442	166	622	307
80%-stop	0.43	0.11	0.44	0.12	446	169	625	311
Part 3								
20%-stop	0.50	0.08	0.49	0.07	415	204	553	175
80%-stop	0.44	0.14	0.46	0.13	414	204	556	178

Table D.9. Overview of the relevant stop-signal data in Experiment 6. Probability of responding on a stop trial [*p*(respond|signal)] and average SSD as a function of condition, word type, distractor type, and part. *M* = mean; *SD* = standard deviation.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Go trials: go RT							
Condition	1	70	60763.15	65147721.30	0.07	0.799	0.001
Part	2	140	3157371.00	10406859.00	21.24	**< 0.001**	0.038
Type (0%-, 20%-, 80%-stop)	2	140	730815.40	668837.20	76.49	**< 0.001**	0.009
Distract	1	70	614976.60	708947.80	60.72	**< 0.001**	0.008
Condition by part	2	140	70810.58	10406859.00	0.48	0.555	0.001
Condition by type	2	140	169491.00	668837.20	17.74	**< 0.001**	0.002
Condition by distract	1	70	13219.55	708947.80	1.31	0.257	< 0.001
Part by type	4	280	31759.51	986425.00	2.25	0.092	< 0.001
Part by distract	2	140	2934.04	514823.60	0.40	0.662	< 0.001
Type by distract	2	140	9656.24	297829.80	2.27	0.119	< 0.001
Condition by part by type	4	280	28348.02	986425.00	2.01	0.093	< 0.001
Condition by part by distract	2	140	1066.16	514823.60	0.14	0.856	< 0.001
Condition by type by distract	2	140	3690.91	297829.80	0.87	0.422	< 0.001
Part by type by distract	4	280	4795.34	703878.20	0.48	0.672	< 0.001
Condition by part by type by distract	4	280	8200.87	703878.20	0.82	0.472	< 0.001
Go trials: *p*(correct)							
Condition	1	70	0.18	3.28	3.94	0.051	0.025
Part	2	140	0.06	0.77	5.17	**0.008**	0.008
Type (0%-, 20%-, 80%-stop)	2	140	0.11	0.71	10.40	**0.001**	0.015
Distract	1	70	0.00	0.20	1.38	0.244	0.001
Condition by part	2	140	0.01	0.77	1.07	0.343	0.002
Condition by type	2	140	0.02	0.71	2.41	0.116	0.003
Condition by distract	1	70	0.00	0.20	0.04	0.842	< 0.001
Part by type	4	280	0.02	0.85	1.71	0.175	0.003
Part by distract	2	140	0.01	0.35	2.03	0.135	0.001
Type by distract	2	140	0.00	0.33	1.00	0.369	0.001
Condition by part by type	4	280	0.02	0.85	2.05	0.120	0.004
Condition by part by distract	2	140	0.02	0.35	3.23	**0.042**	0.002
Condition by type by distract	2	140	0.00	0.33	0.46	0.561	< 0.001
Part by type by distract	4	280	0.02	0.60	2.90	**0.043**	0.003
Condition by part by type by distract	4	280	0.02	0.60	2.57	0.064	0.003
Go trials: *p*(miss)							
Condition	1	70	0.00	1.06	0.12	0.726	0.001
Part	2	140	0.00	0.48	0.05	0.954	< 0.001
Type (0%-, 20%-, 80%-stop)	2	140	0.02	0.34	4.01	**0.043**	0.007
Distract	1	70	0.00	0.11	0.59	0.440	< 0.001
Condition by part	2	140	0.01	0.48	2.05	0.133	0.005
Condition by type	2	140	0.01	0.34	1.74	0.191	0.003
Condition by distract	1	70	0.00	0.11	1.71	0.196	0.001
Part by type	4	280	0.01	0.45	0.85	0.440	0.002
Part by distract	2	140	0.00	0.10	0.05	0.951	< 0.001
Type by distract	2	140	0.00	0.13	0.13	0.776	< 0.001
Condition by part by type	4	280	0.01	0.45	1.32	0.270	0.003
Condition by part by distract	2	140	0.00	0.10	2.14	0.122	0.001
Condition by type by distract	2	140	0.00	0.13	0.04	0.895	< 0.001
Part by type by distract	4	280	0.00	0.19	0.27	0.829	< 0.001
Condition by part by type by distract	4	280	0.00	0.19	1.81	0.151	0.002

Table D.10. Overview of repeated measures Analyses of Variance on the no-signal data in Experiment 6. Condition is included as a between-subjects factor, all other factors are within-subjects. In the no-signal RT analysis, incorrect and missed no-signal trials were removed. type = word type. *p*s < 0.05 are highlighted in bold.

Omitted factor(s)	Bayes Factor	Confidence interval
Main analysis		
Part	0.00	±119.21%
Type	0.00	±94.87%
Distract	0.00	±94.69%
Condition by type	0.03	±96.26%
Condition by part	5.03	±104.23%
Condition by distract	11.28	±90.51%
Condition	13.91	±100.7%
Condition by distract by type	101.88	±96.08%
Part by type	117.19	±82.66%
Condition by part by type	123.91	±102.94%
Distract by type	126.63	±80.76%
Condition by distract by part by type	129.53	±94.28%
Distract by part	199.86	±101.93%
Distract:part by type	613.79	±95.13%
Condition by distract by part	796.95	±115.45%
Instructed condition only		
Part	0.00	±48.01%
Type	0.05	±43.9%
Distract	0.92	±45.34%
Distract by type	12.15	±43.3%
Distract by part	18.34	±45.41%
Distract by part by type	30.86	±43.12%
Part by type	71.44	±42.89%
Uninstructed condition only		
Part	0.00	±54.2%
Distract	0.00	±56.11%
Type	0.53	±58.57%
Distract by type	13.18	±55.68%
Distract by part	31.46	±60.1%
Distract by part by type	37.37	±56.58%
Part by type	96.71	±58.44%

Table D.11. No-signal RT Bayesian analysis in Experiment 6. Bayes factors < 1 indicate that the removal of the factor or interaction had a deleterious effect on the model, whereas Bayes factors > 1 indicate that the factor or interaction could be removed without impairing the fit much. Note that 'participant' was included as a factor for all models, but this factor is not added to the model descriptions in the tables to reduce the amount of text. *Note:* distract = distractor type (no-distractor, distractor); type = word type.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Stop-signal trials: *p*(respond)							
Condition	1	70	0.12	3.56	2.30	0.134	0.010
Part	2	140	0.89	3.30	18.86	**< 0.001**	0.069
Type (20%-, 80%-stop)	1	70	1.10	1.26	61.03	**< 0.001**	0.084
Distract	1	70	0.09	0.59	10.31	**0.002**	0.007
Condition by part	2	140	0.17	3.30	3.63	**0.036**	0.014
Condition by type	1	70	0.39	1.26	21.53	**< 0.001**	0.031
Condition by distract	1	70	0.00	0.59	0.00	0.968	< 0.001
Part by type	2	140	0.08	1.21	4.64	**0.014**	0.007
Part by distract	2	140	0.13	0.96	9.70	**< 0.001**	0.011
Type by distract	1	70	0.00	0.45	0.27	0.608	< 0.001
Condition by part by type	2	140	0.00	1.21	0.17	0.828	< 0.001
Condition by part by distract	2	140	0.01	0.96	0.72	0.475	0.001
Condition by type by distract	1	70	0.01	0.45	0.85	0.361	< 0.001
Part by type by distract	2	140	0.01	0.67	0.55	0.578	< 0.001
Condition by part by type by distract	2	140	0.05	0.67	5.42	**0.005**	0.004

Table D.12. Overview of repeated measures Analyses of Variance on the stop-signal data in Experiment 6. Condition is included as a between-subjects factor, all other factors are within-subjects. type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Go trials: go RT							
Part	2	70	1319009.00	4439376.10	10.40	**0.001**	0.035
Type (0%-, 20%-stop, 80%-stop)	2	70	802089.70	456056.60	61.56	**< 0.001**	0.022
Distract	1	35	223933.10	442994.70	17.69	**< 0.001**	0.006
Part by type	4	140	54354.20	633346.70	3.00	**0.039**	0.001
Part by distract	2	70	303.58	270729.30	0.04	0.960	< 0.001
Type by distract	2	70	6998.30	184672.20	1.33	0.269	< 0.001
Part by type by distract	4	140	11504.42	412381.40	0.98	0.397	< 0.001
Go trials: *p*(correct)							
Part	2	70	0.05	0.47	3.97	**0.023**	0.012
Type (0%-, 20%-, 80%-stop)	2	70	0.12	0.41	9.97	**< 0.001**	0.026
Distract	1	35	0.00	0.15	0.32	0.575	< 0.001
Part by type	4	140	0.04	0.59	2.51	0.075	0.009
Part by distract	2	70	0.02	0.19	4.53	**0.015**	0.006
Type by distract	2	70	0.01	0.23	0.99	0.347	0.001
Part by type by distract	4	140	0.03	0.34	2.98	**0.039**	0.006
Go trials: *p*(miss)							
Part	2	70	0.01	0.32	0.61	0.545	0.003
Type (0%-, 20%-, 80%-stop)	2	70	0.03	0.24	3.64	0.060	0.014
Distract	1	35	0.00	0.06	0.12	0.728	< 0.001
Part by type	4	140	0.01	0.29	1.51	0.230	0.007
Part by distract	2	70	0.00	0.05	0.82	0.447	0.001
Type by distract	2	70	0.00	0.08	0.10	0.819	< 0.001
Part by type by distract	4	140	0.00	0.09	1.07	0.362	0.002
Stop-signal trials: *p*(respond)							
Part	2	70	0.14	1.65	3.06	0.061	0.018
Type (20%-, 80%-stop)	1	35	1.40	0.81	60.39	**< 0.001**	0.154
Distract	1	35	0.04	0.39	3.98	0.072	0.006
Part by type	2	70	0.06	0.68	2.90	0.062	0.007
Part by distract	2	70	0.04	0.61	2.03	0.151	0.005
Type by distract	1	35	0.01	0.22	1.06	0.309	0.001
Part by type by distract	2	70	0.01	0.30	1.44	0.245	0.002

Table D.13. Overview of repeated measures Analyses of Variance on the *instructed* condition only in Experiment 6 with part, word type and distract included as within-participants factors. type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Go trials: go RT							
Part	2	70	1909172.00	5967482.90	11.20	**0.001**	0.043
Type (0%-, 20%-, 80%-stop)	2	70	98216.67	212780.60	16.16	**< 0.001**	0.002
Distract	1	35	404263.00	265953.20	53.20	**< 0.001**	0.009
Part by type	4	140	5753.33	353078.30	0.57	0.614	< 0.001
Part by distract	2	70	3696.62	244094.30	0.53	0.545	< 0.001
Type by distract	2	70	6348.86	113157.60	1.96	0.156	< 0.001
Part by type by distract	4	140	1491.79	291496.80	0.18	0.893	< 0.001
Go trials: *p*(correct)							
Part	2	70	0.02	0.30	1.80	0.178	0.006
Type (0%-, 20%-, 80%-stop)	2	70	0.01	0.31	1.66	0.205	0.005
Distract	1	35	0.00	0.05	1.84	0.183	0.001
Part by type	4	140	0.00	0.27	0.49	0.680	0.001
Part by distract	2	70	0.00	0.16	0.35	0.702	0.001
Type by distract	2	70	0.00	0.10	0.13	0.841	< 0.001
Part by type by distract	4	140	0.02	0.26	2.40	0.079	0.007
Go trials: *p*(miss)							
Part	2	70	0.01	0.16	1.89	0.158	0.008
Type (0%-, 20%-, 80%-stop)	2	70	0.00	0.09	0.94	0.364	0.002
Distract	1	35	0.00	0.05	2.61	0.115	0.003
Part by type	4	140	0.00	0.16	0.30	0.711	0.001
Part by distract	2	70	0.00	0.05	1.34	0.268	0.002
Type by distract	2	70	0.00	0.05	0.06	0.842	< 0.001
Part by type by distract	4	140	0.00	0.10	1.02	0.383	0.003
Stop-signal trials: *p*(respond)							
Part	2	70	0.92	1.65	19.44	**< 0.001**	0.176
Type (20%-, 80%-stop)	1	35	0.09	0.45	7.03	**0.012**	0.021
Distract	1	35	0.04	0.20	7.47	**0.010**	0.010
Part by type	2	70	0.03	0.53	1.76	0.181	0.006
Part by distract	2	70	0.11	0.35	10.70	**< 0.001**	0.024
Type by distract	1	35	0.00	0.23	0.08	0.780	< 0.001
Part by type by distract	2	70	0.04	0.37	4.22	**0.019**	0.010

Table D.14. Overview of repeated measures Analyses of Variance on the *uninstructed* condition only in Experiment 6 with part, word type and distract included as within-subjects factors. type = word type. Note: distract = distractor type (no-distractor, distractor); type = word type. *p*s < 0.05 are highlighted in bold.

Expectancy analyses. Analyses of the expectancy data revealed a main effect of word type ($p < 0.001$) and a reliable interaction between word type and condition ($p < 0.001$), reflecting the larger difference between the word types in the instructed condition ($p < 0.001$) than in the uninstructed condition ($p = 0.004$). Planned comparisons revealed that in the instructed condition participants expected to stop more for the 80%-stop words (6.61) than for the 0%-stop words (3.77), $t(35) = 9.05$, $p < 0.001$, $g_{av} = 2.33$, and for the 20%-stop words (4.31), $t(35) = 7.50$, $p < 0.001$, $g_{av} = 2.02$. The difference between the 0%-stop words and the 20%-stop words was also reliable, $t(35) = -4.29$, $p < 0.001$, $g_{av} = -0.44$. Similarly, in the uninstructed condition participants expected to stop more for the 80%-stop words (5.32) than for the 0%-stop words (4.85), $t(35) = -3.12$, $p < 0.001$, $g_{av} = 0.48$, and for the 20%-stop words (5.03), $t(35) = 2.08$, $p = 0.045$, $g_{av} = 0.32$. The difference between the 0%-stop words and the 20%-stop words was not reliable, $t(35) = -1.55$, $p = 0.130$, $g_{av} = -0.19$. Consistent with the combined analyses, participants in both the instructed and uninstructed conditions could distinguish between the words on the basis of their association with stopping.

Appendix E

Go Stimuli

I used 50 words in this study.

Natural: pear, wasp, moth, calf, plum, crow, slug, dove, toad, swan, crab, pony, deer, worm, lamb, goat, frog, hawk, rice, lion, wolf, duck, bull, bear, tree.

Human-made: tram, coil, mast, gong, harp, wand, vase, raft, sofa, drum, fork, sock, coin, jeep, shed, pill, barn, sink, flag, pipe, bowl, belt, shoe, desk, book.

Appendix F

Analyses of the Eye Data

In Experiment 5, an EyeLink 1000 Desktop Mount camera system (SR Research, Ottawa, Canada) was calibrated at the beginning of each block. The gaze position of the right eye was tracked throughout each block (sampling rate: 250 Hz). The EyeLink was calibrated and controlled via Psychtoolbox (Cornelissen, Peters, & Palmer, 2002). Eye movement data were subsequently exported using the Eyelink Data Viewer (SR Research, Ottawa, Canada) for each participant. I generated a file with information about all fixations and a file with trial sequence information, and integrated these files using R for further analyses.

In the analyses of the eye movement data, I excluded participants when no fixation was registered at the beginning of the trial event (i.e. the presentation of the word) on more than 15% of trials, as this could indicate that eye-movement registration was suboptimal. Based on this criterion, I excluded two participants (1 in the instructed condition, 1 in the uninstructed condition). Note that the inclusion of these participants did not substantially alter the overall pattern of behavioural results (not shown). I also excluded all fixations that were off screen (0.1%).

In the analyses of the eye data, I focused on the number of fixations and the fixation location for two intervals: (1) the interval between the word presentation and the response on no-signal trials, and (2) after the stop signal on stop-signal trials. Eye-movements made 400 ms after the presentation of the stop signal were excluded; this value is based on the cut-off previously used in Verbruggen, Stevens, et al. (2014) as the design used in the present study was suboptimal to accurately determine the SSRTs (which were used as the cut-off in Verbruggen, Stevens, et al., 2014).

The descriptive statistics are in Tables F.1-F.4 and the inferential statistics are in Tables F.5-F.7. If the number of fixations for a particular interval = 1, then the participant did not move their eyes during this time. Based on Verbruggen, Stevens, et al. (2014), I analysed the distance between the fixated location and the centre of the screen [distance = sqrt(x-coordinate^2 + y-coordinate^2)]. In Tables F.3 and F.4, I show the proportion of fixations that fell within 4 pre-defined regions: a region around the word in the centre of the screen, a region in which the distractors occurred, a region around the stop-signal, and an outside region; these regions with their coordinates are depicted in Figure F.1. Consistent with Verbruggen, Stevens, et al (2014), participants did not make many fixations during the intervals of interest, suggesting that they mostly fixated on the centre of screen (Tables F.1 and F.2). Furthermore, numerical differences between groups were very small.

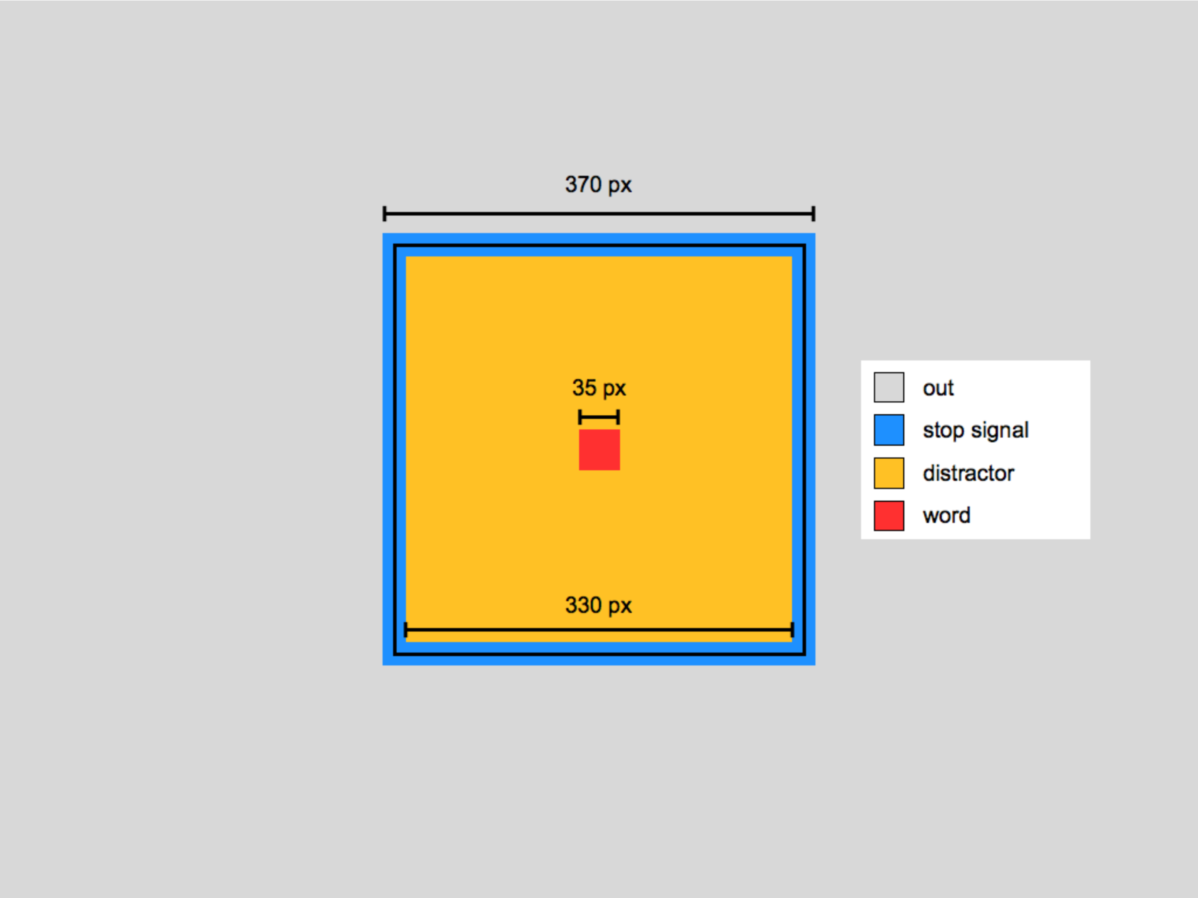

Figure F.1. To analyse fixation location, I predefined 4 regions (squares): a central region around the word, a region with the distractors, a region around the stop signal, and an outside region. The size of each square is in pixels (as pixel coordinates were used for registration of fixation location). Screen size: 1024 x 768 pixels.

Analyses of the number of eye movements on no-signal trials showed that the main effect of word type was reliable ($p = 0.008$). Follow-up analyses showed a reliable main effect of word type in the instructed condition ($p = 0.004$) but was not reliable in the uninstructed condition ($p = 0.579$). This difference was supported by a reliable two-way interaction between condition and word type ($p = 0.011$). Planned comparisons revealed that in the instructed condition participants made more eye movements for the 80%-stop words (1.42) than for the 0%-stop words (1.35), $t(22) = -3.21$, $p = 0.003$, $g_{av} = 0.15$, and for the 20%-stop words (1.36), $t(22) = -2.73$, $p = 0.012$, $g_{av} = 0.11$. The difference between the 0%-stop words and the 20%-stop words was also reliable, $t(22) = -2.46$, $p = 0.022$, $g_{av} = -0.04$.

On stop-signal trials, participants made slightly more eye movements for the 80%-stop words (1.11) than for the 20%-stop words (1.10) ($p = 0.008$). Follow-up analyses revealed that the main effect of word type was reliable in the uninstructed condition ($p = 0.033$) but was only marginally significant in the instructed condition ($p = 0.071$). The two-way interaction between condition and word type was not reliable (*p* = 0.644). Together, these results support the view that the stimulus-stop associations influenced eye movements on no-signal and stop-signal trials. However, the word type differences were very small, and the distance analyses (see below) suggest that participants rarely moved their eyes to the signal location for stop words (for distance analyses, see Tables F.3-F.4). The main effect of distractor type on the number of fixations was not significant on no-signal trials (*p* = 0.965) or on stop-signal trials (*p* = 0.159). The two-way interactions between word type and distractor type were also not significant on no-signal trials and on stop-signal trials (*p*s ≥ 0.542).

Analyses of the average distance between the fixated location and the central word region did not reveal any reliable differences between word types on no-signal (*p*s ≥ 0.395; Table F.5) nor on stop-signal trials (*p*s ≥ 0.364; Table F.6). In the instructed condition, the average distance was greater on distractor trials (43.71) than on no-distractor trials (40.47) on stop-signal trials (*p* = 0.016) but all other differences between distractor types were not reliable (*p*s ≥ 0.246). There was no word type by distractor type interaction on no-signal trials or on stop-signal trials (*p*s ≥ 0.081). Note, I did find a reliable three-way interaction between word type, distractor type, and part on stop-signals trials (*p* = 0.030) but this interaction did not reach significance in the separate analyses of each condition (*p*s ≥ 0.067). Furthermore, post-hoc tests revealed no reliable interaction between distractor type and part for the 20%-stop words (*p* = 0.146) or for the 80%-stop words (*p* = 0.221). Similarly, Bayesian analyses showed that omitting the word type, distractor type, part interaction did not have a negative effect on the model fit (*BF* = 12.06, ±68.67%).

In sum, I find little evidence to support the view that participants moved their eyes to enable stop-signal detection following the presentation of the 80%-stop words compared with the 0%- and 20%-stop words. Furthermore, the eye movement data provide further support for the idea that the effects of word type were similar when the associations were acquired through explicit instructions and task practice.

	Fixation				Distance			
	Distractors		*No distractors*		*Distractors*		*No distractors*	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed								
Part 1								
0%-stop	1.38	0.41	1.35	0.38	40.31	19.39	39.00	19.55
20%-stop	1.41	0.44	1.37	0.39	42.71	21.47	39.83	19.31
80%-stop	1.44	0.53	1.39	0.45	38.45	20.31	39.14	23.03
Part 2								
0%-stop	1.36	0.51	1.35	0.46	40.49	25.89	38.86	25.14
20%-stop	1.38	0.50	1.35	0.48	41.11	26.04	38.67	25.41
80%-stop	1.45	0.63	1.46	0.57	39.67	26.31	40.38	25.67
Part 3								
0%-stop	1.35	0.46	1.29	0.42	39.29	24.55	37.02	23.38
20%-stop	1.35	0.48	1.31	0.45	39.55	24.38	38.69	23.22
80%-stop	1.42	0.54	1.36	0.58	38.83	24.77	38.03	24.76
Uninstructed								
Part 1								
0%-stop	1.26	0.31	1.27	0.29	36.74	21.62	36.92	20.94
20%-stop	1.25	0.30	1.28	0.32	36.55	20.38	37.15	20.13
80%-stop	1.29	0.36	1.33	0.39	37.44	22.25	36.05	21.61
Part 2								
0%-stop	1.20	0.27	1.22	0.26	28.66	16.47	28.44	16.71
20%-stop	1.18	0.22	1.20	0.22	28.49	15.67	27.79	15.59
80%-stop	1.15	0.23	1.19	0.30	25.78	13.83	29.01	19.04
Part 3								
0%-stop	1.13	0.14	1.18	0.20	28.86	11.71	29.30	13.63
20%-stop	1.14	0.16	1.17	0.19	29.36	12.90	29.40	13.75
80%-stop	1.11	0.17	1.19	0.29	31.48	16.36	29.59	14.18

Table F.1. Overview of the number of fixations and average distance between the fixation location and the centre of the screen (in pixels) in the interval between the word presentation and the response for each condition, part, word type, and distractor type. *M* = mean; *SD* = standard deviation.

	Fixation				**Distance**			
	Distractors		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed								
Part 1								
20%-stop	1.16	0.20	1.13	0.15	42.04	27.21	38.32	26.60
80%-stop	1.18	0.22	1.14	0.17	41.53	25.39	38.25	23.62
Part 2								
20%-stop	1.13	0.15	1.11	0.14	45.71	38.34	41.96	37.26
80%-stop	1.14	0.17	1.13	0.17	44.63	39.31	42.32	37.22
Part 3								
20%-stop	1.12	0.18	1.11	0.18	43.01	36.76	42.19	35.62
80%-stop	1.14	0.20	1.11	0.12	45.34	36.20	39.89	34.54
Uninstructed								
Part 1								
20%-stop	1.09	0.15	1.11	0.14	35.38	24.92	35.36	24.07
80%-stop	1.11	0.16	1.13	0.15	37.35	24.99	36.62	24.43
Part 2								
20%-stop	1.06	0.12	1.05	0.06	29.67	18.58	27.29	15.93
80%-stop	1.07	0.11	1.07	0.09	27.72	15.83	29.43	18.33
Part 3								
20%-stop	1.05	0.07	1.06	0.08	29.10	14.20	29.56	14.51
80%-stop	1.04	0.06	1.05	0.08	28.35	11.75	29.37	13.77

Table F.2. Overview of the number of fixations and average distance between the fixation location and the centre of the screen (in pixels) in the interval after the stop signal for each condition, part, word type, and distractor type. *M* = mean; *SD* = standard deviation.

	Word region				**Distractor region**				**Stop-signal region**				**Outside region**			
	D		***ND***		***D***		***ND***		***D***		***ND***		***D***		***ND***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed																
Part 1																
0%-stop	0.72	0.21	0.73	0.21	0.24	0.20	0.24	0.20	0.03	0.04	0.02	0.03	0.02	0.04	0.01	0.03
20%-stop	0.71	0.21	0.73	0.20	0.24	0.20	0.23	0.19	0.03	0.04	0.02	0.03	0.02	0.05	0.01	0.03
80%-stop	0.74	0.23	0.72	0.25	0.23	0.22	0.23	0.24	0.02	0.05	0.03	0.06	0.01	0.03	0.02	0.05
Part 2																
0%-stop	0.75	0.20	0.77	0.19	0.20	0.15	0.18	0.13	0.03	0.06	0.03	0.07	0.02	0.05	0.01	0.03
20%-stop	0.74	0.19	0.77	0.19	0.20	0.14	0.19	0.14	0.04	0.07	0.03	0.06	0.02	0.04	0.02	0.04
80%-stop	0.77	0.20	0.75	0.21	0.18	0.16	0.19	0.17	0.02	0.05	0.03	0.08	0.03	0.07	0.02	0.05
Part 3																
0%-stop	0.74	0.19	0.76	0.18	0.21	0.15	0.20	0.14	0.03	0.06	0.02	0.06	0.02	0.05	0.01	0.04
20%-stop	0.75	0.18	0.75	0.18	0.21	0.14	0.20	0.14	0.03	0.06	0.02	0.05	0.01	0.04	0.02	0.04
80%-stop	0.78	0.19	0.77	0.21	0.17	0.14	0.19	0.17	0.04	0.06	0.02	0.04	0.01	0.04	0.02	0.03
Uninstructed																
Part 1																
0%-stop	0.72	0.20	0.74	0.19	0.24	0.15	0.23	0.15	0.01	0.04	0.02	0.04	0.02	0.07	0.01	0.05
20%-stop	0.73	0.17	0.73	0.19	0.23	0.14	0.24	0.15	0.01	0.03	0.02	0.03	0.02	0.07	0.02	0.06
80%-stop	0.71	0.21	0.74	0.23	0.26	0.18	0.23	0.19	0.01	0.04	0.02	0.04	0.02	0.07	0.02	0.07
Part 2																
0%-stop	0.83	0.17	0.82	0.18	0.15	0.14	0.16	0.16	0.01	0.03	0.01	0.03	0.01	0.04	0.01	0.03
20%-stop	0.84	0.16	0.84	0.16	0.14	0.14	0.14	0.14	0.01	0.02	0.01	0.02	0.01	0.04	0.01	0.03
80%-stop	0.87	0.17	0.84	0.19	0.12	0.16	0.13	0.15	0.00	0.02	0.01	0.04	0.01	0.03	0.01	0.05
Part 3																
0%-stop	0.79	0.20	0.79	0.20	0.20	0.20	0.21	0.19	0.01	0.02	0.00	0.01	0.00	0.01	0.00	0.01
20%-stop	0.79	0.20	0.79	0.20	0.20	0.19	0.20	0.19	0.00	0.01	0.01	0.02	0.01	0.02	0.00	0.01
80%-stop	0.78	0.21	0.81	0.18	0.20	0.20	0.17	0.16	0.01	0.03	0.01	0.04	0.01	0.04	0.01	0.02

Table F.3. Overview of the proportion of the fixations that fell within the 4 pre-defined regions during the interval between the word and the response on no-signal trials in the interval between the word presentation and the response for each condition, part, word type, and distractor type. *D* = distractors; *ND* = No distractors; *M* = mean; *SD* = standard deviation.

	Word region				**Distractor region**				**Stop-signal region**				**Outside region**			
	D		***ND***		***D***		***ND***		***D***		***ND***		***D***		***ND***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Instructed																
Part 1																
20%-stop	0.73	0.23	0.73	0.23	0.21	0.20	0.24	0.21	0.04	0.05	0.02	0.04	0.02	0.06	0.02	0.05
80%-stop	0.71	0.24	0.72	0.25	0.24	0.21	0.24	0.22	0.03	0.06	0.02	0.05	0.02	0.05	0.01	0.04
Part 2																
20%-stop	0.72	0.26	0.74	0.25	0.20	0.15	0.20	0.16	0.06	0.12	0.04	0.09	0.02	0.05	0.02	0.06
80%-stop	0.72	0.26	0.74	0.26	0.21	0.15	0.20	0.17	0.05	0.09	0.04	0.08	0.03	0.06	0.02	0.06
Part 3																
20%-stop	0.73	0.25	0.72	0.26	0.21	0.17	0.23	0.19	0.03	0.07	0.03	0.08	0.03	0.08	0.02	0.05
80%-stop	0.70	0.24	0.75	0.24	0.24	0.16	0.20	0.16	0.04	0.08	0.03	0.07	0.02	0.06	0.02	0.06
Uninstructed																
Part 1																
20%-stop	0.76	0.20	0.74	0.22	0.20	0.17	0.22	0.17	0.01	0.02	0.01	0.03	0.03	0.10	0.02	0.09
80%-stop	0.74	0.19	0.74	0.22	0.23	0.14	0.23	0.17	0.01	0.03	0.01	0.04	0.02	0.08	0.02	0.07
Part 2																
20%-stop	0.83	0.18	0.85	0.15	0.15	0.15	0.13	0.13	0.01	0.02	0.01	0.02	0.02	0.04	0.01	0.05
80%-stop	0.84	0.16	0.82	0.19	0.14	0.14	0.16	0.16	0.01	0.04	0.01	0.03	0.01	0.02	0.01	0.03
Part 3																
20%-stop	0.79	0.22	0.79	0.21	0.20	0.20	0.21	0.20	0.01	0.03	0.00	0.01	0.00	0.01	0.00	0.02
80%-stop	0.78	0.22	0.78	0.20	0.21	0.21	0.20	0.19	0.00	0.01	0.01	0.03	0.00	0.01	0.00	0.01

Table F.4. Overview of the proportion of the fixations that fell within the 4 pre-defined regions during the interval after the stop signal in the interval after presentation of the stop signal for each condition, part, word type, and distractor type. *D* = distractors; *ND* = No distractors; *M* = mean; *SD* = standard deviation.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Word interval							
Number of fixations							
Condition	1	44	5.88	103.44	2.50	0.121	0.046
Part	2	88	1.00	9.13	4.82	**0.019**	0.008
Type (0%-, 20%-, 80%-stop)	2	88	0.20	1.29	6.90	**0.008**	0.002
Distract	1	44	0.00	1.99	0.00	0.965	< 0.001
Condition by part	2	88	0.36	9.13	1.74	0.191	0.003
Condition by type	2	88	0.19	1.29	6.32	**0.011**	0.002
Condition by distract	1	44	0.27	1.99	5.91	**0.019**	0.002
Part by type	4	176	0.01	1.76	0.33	0.771	< 0.001
Part by distract	2	88	0.01	1.14	0.28	0.747	< 0.001
Type by distract	2	88	0.01	1.01	0.47	0.542	< 0.001
Condition by part by type	4	176	0.14	1.76	3.43	**0.025**	0.001
Condition by part by distract	2	88	0.04	1.14	1.59	0.212	< 0.001
Condition by type by distract	2	88	0.01	1.01	0.41	0.576	< 0.001
Part by type by distract	4	176	0.00	1.62	0.11	0.925	< 0.001
Condition by part by type by distract	4	176	0.01	1.62	0.33	0.755	< 0.001
Distance							
Condition	1	44	13071.11	273765.99	2.10	0.154	0.037
Part	2	88	3448.82	43303.67	3.50	0.051	0.010
Type (0%-, 20%-, 80%-stop)	2	88	37.77	2926.34	0.57	0.536	< 0.001
Distract	1	44	70.53	3423.98	0.91	0.346	< 0.001
Condition by part	2	88	2727.41	43303.67	2.77	0.087	0.008
Condition by type	2	88	50.18	2926.34	0.75	0.450	< 0.001
Condition by distract	1	44	78.72	3423.98	1.01	0.320	< 0.001
Part by type	4	176	86.11	6146.05	0.62	0.580	< 0.001
Part by distract	2	88	18.90	3267.00	0.25	0.719	< 0.001
Type by distract	2	88	49.12	2466.27	0.88	0.381	< 0.001
Condition by part by type	4	176	83.39	6146.05	0.60	0.592	< 0.001
Condition by part by distract	2	88	11.40	3267.00	0.15	0.804	< 0.001
Condition by type by distract	2	88	54.12	2466.27	0.97	0.354	< 0.001
Part by type by distract	4	176	132.39	4606.86	1.26	0.290	< 0.001
Condition by part by type by distract	4	176	82.08	4606.86	0.78	0.495	< 0.001

Table F.5. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the word interval. Type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Signal interval							
Number of fixations							
Condition	1	44	0.54	9.33	2.56	0.117	0.044
Part	2	88	0.20	1.31	6.57	**0.008**	0.016
Type (20%-, 80%-stop)	1	44	0.02	0.10	7.82	**0.008**	0.002
Distract	1	44	0.01	0.30	2.05	0.159	0.001
Condition by part	2	88	0.03	1.31	0.84	0.389	0.002
Condition by type	1	44	0.00	0.10	0.22	0.644	< 0.001
Condition by distract	1	44	0.05	0.30	7.62	**0.008**	0.004
Part by type	2	88	0.01	0.24	1.54	0.220	0.001
Part by distract	2	88	0.01	0.31	1.13	0.328	0.001
Type by distract	1	44	0.00	0.13	0.08	0.774	< 0.001
Condition by part by type	2	88	0.01	0.24	1.06	0.350	< 0.001
Condition by part by distract	2	88	0.01	0.31	1.09	0.340	0.001
Condition by type by distract	1	44	0.00	0.13	0.18	0.670	< 0.001
Part by type by distract	2	88	0.01	0.22	1.71	0.192	0.001
Condition by part by type by distract	2	88	0.00	0.22	0.21	0.777	< 0.001
Distance							
Condition	1	44	16193.59	338042.94	2.11	0.154	0.039
Part	2	88	564.24	46859.63	0.53	0.537	0.001
Type (20%-, 80%-stop)	1	44	1.41	686.52	0.09	0.765	< 0.001
Distract	1	44	355.66	2519.06	6.21	**0.017**	0.001
Condition by part	2	88	3415.95	46859.63	3.21	0.061	0.009
Condition by type	1	44	13.46	686.52	0.86	0.358	< 0.001
Condition by distract	1	44	360.40	2519.06	6.30	**0.016**	0.001
Part by type	2	88	21.94	1682.83	0.57	0.566	< 0.001
Part by distract	2	88	12.93	2133.38	0.27	0.762	< 0.001
Type by distract	1	44	1.37	587.69	0.10	0.750	< 0.001
Condition by part by type	2	88	33.15	1682.83	0.87	0.424	< 0.001
Condition by part by distract	2	88	8.22	2133.38	0.17	0.840	< 0.001
Condition by type by distract	1	44	42.92	587.69	3.21	0.080	< 0.001
Part by type by distract	2	88	134.67	1509.96	3.92	**0.030**	< 0.001
Condition by part by type by distract	2	88	58.77	1509.96	1.71	0.192	< 0.001

Table F.6. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the stop-signal interval. Type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Word interval							
Number of fixations							
Part	2	44	0.17	5.05	0.73	0.486	< 0.001
Type (0%-, 20%-, 80%-stop)	2	44	0.38	0.94	8.91	**0.004**	0.004
Distract	1	22	0.14	1.38	2.21	0.151	0.001
Part by type	4	88	0.05	1.10	1.00	0.381	0.001
Part by distract	2	44	0.03	0.78	0.97	0.383	< 0.001
Type by distract	2	44	0.00	0.56	0.00	0.992	< 0.001
Part by type by distract	4	88	0.01	1.12	0.24	0.793	< 0.001
Distance							
Part	2	44	159.59	22782.42	0.15	0.757	0.001
Type (0%-, 20%-, 80%-stop)	2	44	87.19	2072.94	0.93	0.395	< 0.001
Distract	1	22	149.14	2307.32	1.42	0.246	0.001
Part by type	4	88	84.19	4402.15	0.42	0.688	< 0.001
Part by distract	2	44	0.71	1964.50	0.01	0.971	< 0.001
Type by distract	2	44	102.71	1017.55	2.22	0.137	< 0.001
Part by type by distract	4	88	48.04	2695.10	0.39	0.736	< 0.001
Signal interval							
Number of fixations							
Part	2	44	0.05	0.65	1.62	0.216	0.006
Type (20%-, 80%-stop)	1	22	0.01	0.08	3.60	0.071	0.001
Distract	1	22	0.06	0.24	5.50	**0.028**	0.007
Part by type	2	44	0.00	0.16	0.03	0.969	< 0.001
Part by distract	2	44	0.01	0.15	1.13	0.334	0.001
Type by distract	1	22	0.00	0.11	0.01	0.942	< 0.001
Part by type by distract	2	44	0.01	0.16	1.06	0.354	0.001
Distance							
Part	2	44	638.16	26408.78	0.53	0.524	0.002
Type (20%-, 80%-stop)	1	22	3.08	384.79	0.18	0.679	< 0.001
Distract	1	22	716.05	2318.83	6.79	**0.016**	0.002
Part by type	2	44	1.78	1056.91	0.04	0.964	< 0.001
Part by distract	2	44	2.74	1132.79	0.05	0.948	< 0.001
Type by distract	1	22	14.48	391.44	0.81	0.377	< 0.001
Part by type by distract	2	44	121.97	976.44	2.75	0.095	< 0.001

Table F.7. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the *instructed condition* only. Type = word type. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Word interval							
Number of fixations							
Part	2	44	1.19	4.08	6.43	**0.010**	0.041
Type (0%-, 20%-, 80%-stop)	2	44	0.01	0.35	0.41	0.579	< 0.001
Distract	1	22	0.13	0.61	4.64	**0.043**	0.005
Part by type	4	88	0.10	0.66	3.35	**0.024**	0.004
Part by distract	2	44	0.01	0.36	0.84	0.438	< 0.001
Type by distract	2	44	0.02	0.44	1.00	0.343	0.001
Part by type by distract	4	88	0.00	0.50	0.15	0.917	< 0.001
Distance							
Part	2	44	6016.64	20521.25	6.45	**0.010**	0.048
Type (0%-, 20%-, 80%-stop)	2	44	0.76	853.40	0.02	0.952	< 0.001
Distract	1	22	0.11	1116.66	0.00	0.963	< 0.001
Part by type	4	88	85.30	1743.90	1.08	0.362	0.001
Part by distract	2	44	29.59	1302.50	0.50	0.575	< 0.001
Type by distract	2	44	0.53	1448.72	0.01	0.960	< 0.001
Part by type by distract	4	88	166.43	1911.76	1.92	0.148	0.001
Signal interval							
Number of fixations							
Part	2	44	0.17	0.66	5.76	**0.019**	0.050
Type (20%-, 80%-stop)	1	22	0.01	0.03	5.20	**0.033**	0.002
Distract	1	22	0.01	0.06	2.19	0.153	0.002
Part by type	2	44	0.01	0.08	4.07	**0.024**	0.004
Part by distract	2	44	0.01	0.16	1.10	0.342	0.002
Type by distract	1	22	0.00	0.01	1.23	0.280	< 0.001
Part by type by distract	2	44	0.00	0.07	0.72	0.448	0.001
Distance							
Part	2	44	3342.03	20450.85	3.60	0.057	0.034
Type (20%-, 80%-stop)	1	22	11.79	301.72	0.86	0.364	< 0.001
Distract	1	22	0.01	200.22	0.00	0.977	< 0.001
Part by type	2	44	53.31	625.92	1.87	0.169	0.001
Part by distract	2	44	18.40	1000.59	0.40	0.639	< 0.001
Type by distract	1	22	29.81	196.24	3.34	0.081	< 0.001
Part by type by distract	2	44	71.47	533.52	2.95	0.067	0.001

Table F.8. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the *uninstructed condition* only. Type = word type. *p*s < 0.05 are highlighted in bold.

Appendix G

Pilot Experiment

I also ran a pilot version of the stop distract where no stimulus-stop information was provided (i.e. this experiment was a procedural replication of the uninstructed condition).

Method

Participants. Twenty-four students from University of Exeter participated for monetary compensation (£5) or course credit. I excluded one participant because the percentage correct go trials was below 70%. The target sample size and exclusion criteria were decided in advance of data collection.

Apparatus and stimuli, procedure and analyses. The apparatus, stimuli, and procedure were identical to those of Experiment 5, except for the following changes: the experiment was run on a 21-in iMac, I did not track eye gaze position, and I did not obtain expectancy ratings following task completion. Participants were not provided with stimulus-stop contingency information.

Results and Discussion

Analyses of no-signal RTs showed that there was a significant effect of word type ($p = 0.018$). Planned comparisons showed that no-signal RTs were generally longer for the 80%-stop words ($M = 785$ ms) than for the 0%-stop words ($M = 774$ ms) and the 20%-stop words ($M = 776$ ms), $t(22) = 2.26$, $p = 0.034$, $g_{av} = 0.09$ (one-tailed t-test: $p = 0.017$), and $t(22) = 1.99$, $p = 0.058$, $g_{av} = 0.07$ (one-tailed t-test: $p = 0.029$), respectively. Table G.1 suggests that this difference between word types was largest in the final part of the experiment, but the interaction between part and word type was not significant ($p = 0.643$).

The presence of distractors slowed no-signal RTs (no-distractor: 763 ms, distractor: 793 ms; $p < 0.001$; Table G.2). There was no reliable interaction between word type and the presence of distractors ($p = 0.992$); the three-way interaction between word type, distractor type, and part was also not significant ($p = 0.826$).

The outcome of the Bayesian analyses on no-signal RTs are broadly consistent with the ANOVA. As can be seen in Table G.3, removal of the part or distractor factors had a deleterious effect on the model. However, all other factors or interactions could be dropped from the model. Note, the discrepancy between the Bayesian analyses and the ANOVA in terms of word type could be due to the restricted sample size in this pilot study. Therefore, in the main experiment I increased the sample size (and consequently, the power).

In sum, the absence of an interaction between word type and distractor type in the pilot study replicates the findings of the uninstructed condition in the main study. Therefore, these findings provide further support for the idea that, in situations where no stimulus-stop contingency information is explicitly instructed, participants form direct associations between the 80%-stop words and a stop representation via task practice.

	***p*(correct)**				***p*(miss)**				**no-signal RT**			
	Distractors		***No distractors***		***Distractors***		***No distractors***		***Distractors***		***No distractors***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Part 1												
0%-stop	0.89	0.07	0.90	0.06	0.03	0.03	0.04	0.04	764	102	730	96
20%-stop	0.90	0.06	0.90	0.05	0.05	0.04	0.04	0.04	766	100	738	95
80%-stop	0.89	0.11	0.86	0.12	0.04	0.07	0.02	0.05	774	121	741	109
Part 2												
0%-stop	0.93	0.07	0.93	0.06	0.04	0.04	0.02	0.03	788	135	762	125
20%-stop	0.93	0.06	0.93	0.04	0.04	0.03	0.03	0.03	795	140	762	127
80%-stop	0.91	0.12	0.92	0.10	0.03	0.04	0.04	0.08	800	136	765	135
Part 3												
0%-stop	0.93	0.07	0.93	0.06	0.05	0.04	0.03	0.03	815	145	784	131
20%-stop	0.93	0.07	0.92	0.06	0.04	0.04	0.03	0.03	811	144	785	137
80%-stop	0.92	0.11	0.92	0.14	0.03	0.06	0.03	0.06	825	148	805	162

Table G.1. Overview of the go data. Probability of an accurate go response [*p*(correct)], probability of a missed go response [*p*(miss)] and average go reaction time (RT) as a function of word type, distractor type, and part. Accuracy is the ratio of correct go trials to the number of correct and incorrect go trials (missed trials are excluded). *P*(miss) is the ratio of omitted responses to the total number of go trials. *M* = mean; *SD* = standard deviation.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Go trials: go RT							
Type (0%-, 20%-, 80%-stop)	2	44	9832	49364	4.38	**0.018**	0.002
Part	2	44	187212	466213	8.83	**< 0.001**	0.028
Distract	1	22	90251	65262	30.04	**< 0.001**	0.014
Type by part	4	88	2613	91335	0.63	0.643	< 0.001
Type by distract	2	44	23	64422	0.01	0.992	< 0.001
Part by distract	2	44	721	49297	0.32	0.726	< 0.001
Type by part by distract	4	88	1374	80576	0.38	0.826	< 0.001
Go trials: *p*(correct)							
Type (0%-, 20%-, 80%-stop)	2	44	0.02	0.22	2.10	0.134	0.007
Part	2	44	0.09	0.28	7.15	**0.002**	0.031
Distract	1	22	0.00	0.12	0.11	0.748	< 0.001
Type by part	4	88	0.00	0.45	0.18	0.946	0.001
Type by distract	2	44	0.00	0.21	0.15	0.864	< 0.001
Part by distract	2	44	0.00	0.16	0.18	0.837	< 0.001
Type by part by distract	4	88	0.01	0.28	0.74	0.565	0.003
Stop-signal trials: *p*(respond)							
Type	1	22	0.12	0.19	13.88	**0.001**	0.046
Part	2	44	0.21	0.59	7.83	**0.001**	0.079
Distract	1	22	0.05	0.11	10.11	**0.004**	0.021
Type by part	2	44	0.02	0.32	1.23	0.301	0.007
Type by distract	1	22	0.00	0.12	0.07	0.794	< 0.001
Part by distract	2	44	0.06	0.34	3.51	**0.039**	0.022
Type by part by distract	2	44	0.01	0.41	0.47	0.626	0.004

Table G.2. Overview of repeated measures analyses of variance. In the no-signal RT analysis, incorrect and missed no-signal trials were removed. I did not analyse *p*(miss) because values were low. Type = word type. *p*s < 0.05 are highlighted in bold.

Omitted factor(s)	**Bayes Factor**	**Confidence interval**
Part	< 0.00	±32.35%
Distract	< 0.00	±31.28%
Type	6.84	±40.51%
Distract:type	15.66	±32.81%
Distract:part	16.84	±32.03%
Distract:part:type	28.68	±29.54%
Part:type	61.62	±33.05%

Table G.3. No-signal RT Bayes analysis. Bayes factors < 1 indicate that the removal of the factor or interaction had a deleterious effect on the model whereas Bayes factors > 1 indicate that the factor or interaction could be removed. Note that 'participant' was included as a factor for all models, but this factor is not added to the model descriptions in the tables to reduce the amount of text. distract = distractor type (no-distractor, distractor); type = word type.

CHAPTER 4

THE ROLE OF SIGNAL DETECTION DURING STOP/GO LEARNING

It has recently been proposed that, at the most basic level, action control depends on three cognitive processes; signal detection, action selection, and action execution (Verbruggen, McLaren, et al., 2014; For an extended discussion, see Chapter 1). Each of these processes can be modulated via learning. In the present chapter, I tested part of this framework by examining how learning modulates signal detection processes in response inhibition tasks.

Perceptual processes play an important, yet under-researched, role in successful stopping. One of the first stages of stopping a response often involves detecting a stop signal (e.g. noticing the oncoming vehicle when crossing a road or noticing a traffic light turning red). Indeed, computational modelling work suggests that a considerable proportion of the stopping latency is occupied by perceptual or afferent processes (Boucher, Palmeri, et al., 2007; Logan, Van Zandt, et al., 2014; Salinas & Stanford, 2013; for an overview of this work). Consistent with this idea, it has recently been demonstrated that people make adjustments to their attentional focus when they anticipate the presentation of a stop signal in the periphery of the screen (Verbruggen, Stevens, et al., 2014; for an extended discussion of this experiment, see Chapter 3). Similarly, electrophysiology research has also shown that early sensory event-related potential components, such as the N1 and the Selection Negativity, are modulated in contexts where participants are instructed to monitor for a stop signal but not in contexts where participants are instructed to ignore the stop signals (Elchlepp, Lavric, et al., 2016; Langford, Krebs, Talsma, Woldorff, & Boehler, 2016). This pattern of findings suggests that people adjust their attentional settings in situations where they are required to monitor for the signal (Elchlepp, Lavric, et al., 2016). The ‘biased competition’ theory of visual attention suggests that when there is competition between sources of information, attention is biased to specific features of the visual object in order to resolve the competition (e.g. Desimone & Duncan, 1995; Duncan, 2006; Kastner & Ungerleider, 2000). Thus, Elchlepp, Lavric, et al. (2016) speculated that signal detection could be biased during performance of response inhibition tasks by increasing the baseline activity in the sensory neurons that code for specific features of the stop signal. Consequently, those features of the signal would be more likely to win the competition, and stopping would be more likely to be successful. In other words, successful stopping in stop-signal tasks depends on selective attention to the features of the stop signal.

In stop learning tasks, the effects of learning on behaviour depends on participants not only attending to the stop signal, but also to the stop-associated stimuli. As suggested in Chapter 2, if participants begin to ignore the stop-associated stimuli, the effects of stop learning on behaviour are likely to be diminished or even eliminated entirely. Thus, I concluded that attention likely plays an important role in mediating the effects of stop learning on behaviour. However, the relationship between attention and learning is not a one-way street. Attention can also be influenced by learning. In the go/no-go task, electrophysiology research has shown that the amplitude of the visual N1 component[30] decreased (i.e. became less negative) for the no-go stimuli relative to the go stimuli with training (Benikos, Johnstone, & Roodenrys, 2013; see also Benikos, unpublished PhD thesis). In particular, this amplitude difference was only found when the go/no-go task demanded rapid responding to the go stimuli. The change in the N1 amplitude observed by Benikos, Johnstone, et al. could indicate that participants attended more to the go stimuli than to the no-go stimuli with training. However, it is important to highlight that this is not always optimal for the effects of training on behaviour; in the Benikos, Johnstone, et al. experiment, no-go performance became worse with training not better. Thus, response inhibition training can influence attention to go-associated and stop-associated stimuli, but the effects of learning on behaviour depends on attention. Furthermore, Manuel and colleagues showed that following performance of an auditory go/no-go task the topography of auditory-event related potentials was modulated 61-104 ms for stimuli associated with stopping. Manuel, Grivel, et al. attributed this modulation to reduced activity in the decreased activity within left parietal cortices (i.e. no-go; Manuel, Grivel, et al., 2010). It is likely that this effect was due to the effects of learning on the detection of the no-go stimuli (rather than direct activation of an inhibition or no-go network; see Verbruggen, Best, et al., 2014). Just as in the Benikos, et al. experiment, the design of this task emphasised rapid responding on go trials. Thus, when the task places emphasis on go processing, learning may influence selective attention to attributes of the go stimuli.

Consistent with the idea that stop learning may influence attentional processes, Awh and colleagues suggested that attentional selection can be influenced by past experiences (Awh, Belopolsky, & Theeuwes, 2012). However, although the aforementioned studies showed that learning may influence attention to the no-go stimuli, no study conducted to date has directly manipulated attention during stop learning. As such, our understanding of the effects of learning on attention in response inhibition tasks is based on mere speculation. Beyond the response inhibition literature, several studies have demonstrated that attention is sensitive to prior learning. For example, it has been shown that attention is biased towards stimuli that were predictive of high reward in a previous training task (e.g. Anderson, Laurent, & Yantis, 2011; Della Libera & Chelazzi, 2009; Le Pelley, Mitchell, & Johnson, 2013; Le Pelley, Pearson, Griffiths, & Beesley, 2015). Furthermore, studies in

---

[30] The visual N1 component is thought to be sensitive to perceptual features of the stimulus and is modulated by attention (Näätänen and Picton, 1987). Furthermore, it has been found to be larger on trials where stopping was successful indicating that a component of stopping depends on the ability to attend to the stop signal (Bekker, Kenemans, Hoeksma, Talsma, & Verbaten, 2005).

the associative learning literature have shown that attention in human and animals is sensitive to the learned predictiveness of a given stimulus; stimuli that are consistently followed by the same outcome capture attention faster than stimuli that are not consistently followed by the same outcome (e.g. Le Pelley, Vadillo, & Luque, 2013; for a review, see Mitchell & Le Pelley, 2010). Thus, attentional selection is not solely determined by top-down control or by physical properties of the target stimulus, but can also be influenced by associative learning. On the basis of these experiments, it follows that learning in response inhibition tasks where the representation of the stop signal remains the same throughout training may result in attention being directed to the stop signal following the presentation of stop-associated stimuli. In the present chapter, I investigate the effects of learning on attentional selection in response inhibition tasks.

The role of learning on attention would have important implications for our understanding of the associative architecture of stop learning outlined in Chapter 1 (see *What is Learned?*). To recap, I discussed the possibility that there are multiple pathways that will enable a stimulus that is consistently paired with stopping (or ‘not going’ in a go/no-go paradigm) to influence performance. One pathway is that participants could form a direct association between a stop-associated stimulus and the stop goal. This was the original hypothesis outlined by Verbruggen and Logan (2008a). In Chapter 2, I provide strong evidence for the direct pathway (see also Bowditch, Verbruggen, et al., 2016). The other pathway is an association between the stop-associated stimulus and the stop goal that is mediated via a representation of the stop signal. The stimulus-signal learning pathway predicts that attention would become trained to features of the signal. The signal representation could include feature dimensions (e.g. colour, shape) alongside the spatial location and/or spatial configuration of the signal. Thus, unlike stimulus-stop learning, the retrieval of a stimulus-signal association would prime the detection of the stop signal. This could result in a lower probability of responding on stop signal trials for stop-associated items but would have no effect or very little effect on reaction times on no-signal trials[31]. However, to date, there has been mixed evidence in favour of the stimulus-signal route. For example, Bowditch and colleagues showed that reducing the contingencies between specific stimuli and the stop signals resulted in stronger effects of stop learning on task performance than when specific stimuli are paired with a single representation of a stop signal. (Bowditch, Verbruggen, et al., 2016). Furthermore, some indirect support was provided by some recent experiments in which the signal representation remained the same throughout training (the line beneath the stop- and go-associated stimuli became bold) whereby stop learning influenced the probability of responding on stop trials but had no effect on response latencies on go trials (Verbruggen, Best, et al., 2014, Experiment 2). However, in Chapter 3, despite using a single representation of the stop signal throughout training, I did not find any evidence to support the idea

---

[31] A weak effect on RTs would arise from some partial activation of the stop goal via the indirect pathway (e.g. stimulus -> stop signal -> stop goal; see Bowditch, Verbruggen, et al., 2016).

that attention was directed towards the stop signal during stop learning. It is possible that the absence of an effect of stop learning on perceptual processes was due to the distractor manipulation that I implemented in that task. Alternatively, it could be that attention to the signal is not influenced by stop learning after all or that the effects of mediation on performance are much smaller than previously predicted by Verbruggen, Best, et al. (2013; for discussion of this possibility, see also Chapter 3). Therefore, in the present chapter, I adopted a different approach.

The Present Chapter

In the present chapter, I investigate whether participants can acquire associations between specific stimuli and the spatial location of the stop signal. Research in the visual attention literature has shown that the detection of a target stimulus is enhanced by the provision of a cue indicating its spatial location (Posner, 1980). To this end, I combined features of the stop learning tasks used in Chapter 2 with features of spatial cueing tasks. In a typical spatial cueing task, participants are presented with a target stimulus in the periphery of the screen (usually either to the left or to the right of the centre), which is preceded by a ‘cue’ indicating that participants should expect the target to appear in a given signal location (Posner, 1980; Posner & Cohen, 1984). As such, on each trial in the Experiments 7-8 presented in this chapter participants were presented with a cue in the centre of the screen that preceded the presentation of a stop or go signal which could occur on the right or on the left of the screen. I manipulated both the contingencies between the cues and stopping (or going) and the contingencies between the cues and the location of the stop (or go) signal. Participants were not informed about the cue-stop or cue-signal mappings. If participants acquired the associations between the cues and the signal location in this task, the general learning effect would be greater (i.e. a larger difference in go RTs on no-signal trials and the probability of responding on stop trials between the beginning and the end of training) for the cues that were consistently paired with the same signal location than for the cues that were paired with both signal locations with equal probability (Experiment 7). Furthermore, performance on trials in which the signal appeared in the trained location should be better than performance on trials in which the signal appeared in the untrained location (Experiment 8). Evidence that participants acquired the associations between the cues and the signal location would provide some of the strongest evidence to date for the idea that associations can form between stimuli and a representation of the stop signal.

Experiment 7

In Experiment 7, a subset of cues was consistently associated with stop signals and another subset was consistently associated with go signals. I indexed stop/go learning during task performance via two measures. The first index was the probability of responding on stop-signal trials, that was predicted to become lower as a function of training. The second index was RT on correct go-signal trials, that was predicted to decrease as a function of training. Furthermore, I also obtained stop/go expectancy ratings following task completion in which participants were asked to rate the extent to which they expected to withhold their response for each of the images (cues) presented in the task. Note that this experiment was originally designed as a pilot experiment for an ERP experiment. Therefore, to maximise the available trials for ERP analyses, the stop-associated cues were paired with a stop signal and the go-associated cues with a go signal on 100% of presentations. Due to these contingencies and the absence of a test phase, it was not possible to directly compare task performance for the stop-associated and go-associated cues. However, Chapters 2 and 3 indicate that if participants acquired the stimulus-stop associations during task performance they would also generate expectancies consistent with the stimulus-stop contingencies in play. Therefore, I also obtained expectancy ratings following task completion.

To investigate the role of signal detection during stop learning, I varied the location of the go/stop signal on the screen; for a subset of cues, the signal was consistently presented in the same signal location whereas for another subset, the signal was presented in two signal locations with equal probability.

Method

Participants. Thirty-six volunteers from University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 20.39 years, *SD* = 3.10, 30 females, 33 right-handed). Five participants were excluded as the percentage of correct go responses was less than 70%[32]. All experiments of the present chapter were approved by the local research ethics committee at the School of Psychology, University of Exeter. Written informed consent was obtained after the nature and possible consequences of the study were explained. The target sample size and exclusion criteria were decided in advance of data collection to ensure that I had enough power (0.80) to detect medium-sized effects in the within-subject comparisons.

Apparatus, stimuli, and procedure. The experiment was run on a 21-inch iMac (screen size: 1920 × 1080 pixels) using Psychtoolbox (Brainard, 1997). The cues were words presented in black lowercase font (Courier 16 point) and the go and stop signals consisted of coloured dots (size: 20 × 20 pixels) presented on a grey background (RGB: 255 255 255; Figure 4.1). The word was presented in the centre of the screen and a go or stop signal was presented to the left of the word (*x* = -160 pixels, -4.2 cm from the centre) or to the right of the word (*x* = 160 pixels, 4.2 cm from the centre). A control signal was always on the opposite side. I created a list of 112 five- and six-letter words (Appendix H). The word could refer to a natural or a human-made object (56 natural words, 56 human-made words).

---

[32] This exclusion criteria was set on the basis of Experiments 5-6 (Chapter 3) which similarly required participants to categorise words as natural/human-made.

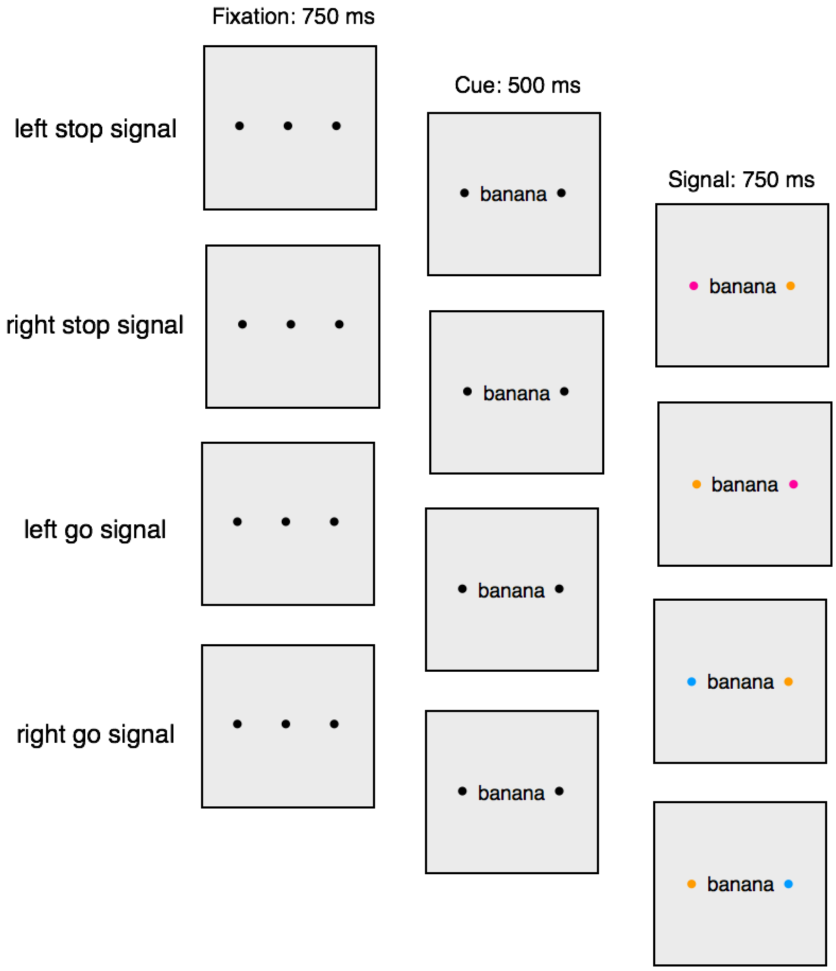

Figure 4.1. The left stop signal, right stop signal, left go signal, and right go signal trial structures in Experiment 7. Each trial began with the presentation of three fixation dots, one in the centre of the screen, one to the left of the centre, and one to the right of the centre for 750 ms. Following the removal of the fixation dots, a word (i.e. the cue) was presented in the centre of the screen. After 500 ms, the go/stop signal replaced the dot to the left or to the right of the word and a control signal (which was presented on go and stop trials) was presented on the opposite side. The word and the signals remained on the screen for 750 ms (hence, MAXRT was 1250 ms). In this example the stop signal is represented as a pink dot, the go signal is represented by a blue dot and the control signal is represented by a yellow dot. Note: the dot colours were fully counterbalanced.

The experiment consisted of three parts, with four blocks per part. Each word was presented once per block. There were six different cue-types (see Table 4.1). First, *stop-left* cues were always paired with a stop signal and the stop signal always appeared to the left of the word. Second, *stop-right* cues were always paired with a stop-signal and the stop signal always appeared to the right of the word. Third, *stop-inconsistent* cues were always paired with a stop signal, but on half of the trials the stop signal appeared to the left of the word and on the other half of trials the stop signal appeared to the right of the word. Fourth, *go-left* cues were always paired with a go signal and the go signal always appeared to the left of the word. Fifth, *go-right* cues were always paired with a go signal and the go signal always appeared to the right of the word. Finally, *go-inconsistent* cues were always paired with a go signal but on half of the trials the go signal appeared to the left of the word and on the other half of trials the go signal appeared to the right of the word. Words were randomised over the stop/go-types across participants. Note that there were no ‘stop/go-inconsistent but location-consistent’ cues.

Cue-type	**No. of words**	**% left stop signal**	**% right stop signal**	**% left go signal**	**% right go signal**
Stop-left	12	100	0	0	0
Stop-right	12	0	100	0	0
Stop-inconsistent	24	50	50	0	0
Go-left	16	0	0	100	0
Go-right	16	0	0	100	0
Go-inconsistent	32	0	0	50	50

Table 4.1. The proportion of left stop-signal, right stop-signal, left go-signal, and right go-signal trials as a function of stop/go-type in Experiment 7. Note that for the analyses, the stop-left and stop-right cues were averaged (*stop-consistent* cues) and the go-left and the go-right cues were averaged (*go-consistent* cues). *Note*: the overall p(stop-signal) was 0.43.

On each trial, three black fixation dots were presented marking the central, left, and right stimulus locations (Figure 4.1). After 750 ms, the central fixation dot was replaced with a word. After 500 ms, the fixation dots on the left and on the right of the word were replaced with coloured dots (aside from the colour change, the size and position of the dots were the same as the fixation). There were three dot colours: yellow (RGB: 255 155 0), pink (RGB: 255 0 155), and blue (RGB: 0 155 255). On go trials, one of the dots on the left or on the right of the word was presented in the go signal colour (e.g. left = blue) and the dot on the other side of the word was presented in the control signal colour (e.g. right = yellow). On go trials, participants had to decide whether the word referred to a natural or human-made object. Half of the participants had to press the ‘c’ key (with their left index finger) when the word referred to a natural object, and the ‘m’ key (with their right index finger) when the word referred to a human-made object. This mapping was reversed for the other half of participants. On stop trials, one of the dots on the left or on the right of the word was presented in the stop-signal colour (e.g. left = pink) and the dot on the other side of the word was presented in a control signal colour (e.g. right = yellow). The colours were fully counterbalanced over the go, stop, and control signals. The control signal was the same on go-signal and stop-signal trials. To increase the likelihood that participants would be required to inhibit a planned response on stop-signal trials, the experimenter instructed participants to begin preparing their go response following the presentation of the word (i.e. the cue). At the end of each block, participants received feedback on their mean RT on go trials, the number of go errors, the number of missed go responses, and the percentage of failed stops. After 15 s had elapsed, participants had to press any key to start the next block.

Following completion of the experimental task, each word was again presented on the screen. The order of the words was randomised anew for each participant. Participants were asked to rate "How much do you expect to withhold your response when this word is presented?" on a scale, ranging from 1 (*I definitely do not think this word indicates that I have to withhold my response)* and 9 (*I definitely think this word indicates that I have to withhold my response).* There was no response deadline for the expectancy ratings.

Analyses. All data processing and analyses were completed using R (R Development Core Team, 2014). All data files and R scripts are deposited in Dropbox (https://www.dropbox.com/sh/wflj1y7r4hb4tg4/AADZXtHg1sHkKilu39que0AHa?dl=0). For all analyses, I collapsed the *stop-left* cues and the *stop-right* cues as both of these cues consistently predicted a given signal location. These are referred herein to as *stop-consistent* cues. For the same reasons, I also collapsed the *go-left* cues and the *go-right* cues. These are referred to as *go-consistent* cues. Note that the delay between the presentation of the cues and the signals was always 500 ms. Consequently, it was not possible to estimate or analyse SSRT. I analysed *p(*respond|signal) to determine if stop/go learning influenced stopping performance (see also e.g. Chapters 2 and 3; Bowditch, Verbruggen, et al., 2016; Noël, Brevers, et al., 2016; Verbruggen, Best, et al., 2014).

ANOVAs were performed on correct RTs and the probability of correct responses on go-signal trials, and on the probability of responding on stop-signal trials. Behavioural performance was analysed as a function of 'part' (there were 4 blocks per part: 1 = blocks 1-4; 2 = blocks 5-8; 3 = blocks 9-12) and cue-type (stop-consistent, stop-inconsistent, go-consistent, go-inconsistent) which were included as within-subjects factors. I also performed an ANOVA on the expectancy ratings with stop/go-type (stop-associated cues, go-associated cues) and location-type (consistent-signal location cues, inconsistent-signal location cues) as within-subjects factors. Where appropriate, I applied the Huyhn-Feldt correction for violations of sphericity. For pairwise comparisons, Hedge's g_{av} is the reported effect size measure (Lakens, 2013).

To examine support for the null hypothesis, I also computed Bayesian t-tests for all pairwise comparisons and Bayesian regressions for the stop/go expectancy-behaviour correlations (Rouder, Morey, et al., 2012). Bayes factors compare the likelihood of the data under the null hypothesis of no difference against the alternative hypothesis of a difference with an effect size that corresponds to the prior. The default prior assumes a large difference between conditions. A Bayes factor less than 0.33

constitutes support for the null hypothesis whereas a Bayes factor of more than 3 constitutes support for the alternative hypothesis. I calculated the Bayes factors with the BayesFactor package in R, using the default prior of 0.707 (Morey, Rouder, et al., 2015).

Results

Go analyses. Analyses of go RTs revealed a main effect of part ($p < 0.001$; Table 4.2) reflecting the decreasing RTs as function of training. However, there was no reliable main effect of cue-type ($p = 0.104$; Table 4.2; Figure 4.2) and the corresponding Bayes factor for the difference between the go-consistent cues and the go-inconsistent cues was 0.67. The two-way interaction between part and cue-type was also not reliable ($p = 0.099$; Table 4.2).

Analyses of the probability of correct go responses showed a similar pattern of results: an increase in the probability of correct go responses as a function of part ($p < 0.001$; Table 4.2; Figure 4.2), but no reliable main effect of cue-type ($p = 0.440$; Table 4.2) and no reliable two-way interaction between part and cue-type ($p = 0.887$; Table 4.2). The Bayesian analyses supported the null hypothesis of no difference between the go-consistent and go-inconsistent cues ($BF = 0.25$).

The probability of missed go responses reliably decreased as a function of part ($p < 0.001$; Table 4.2). The main effect of cue-type did not reach significance ($p = 0.079$; Table 4.2), but the two-way interaction between part and cue-type was reliable ($p = 0.026$; Table 4.2). Follow-up comparisons showed that the main effect of cue-type was reliable in part 1, $t(30) = 2.47$, $p = 0.019$, $g_{av} = 0.26$, $BF = 2.57$, but was not reliable in part 2, $t(30) = -0.80$, $p = 0.430$, $g_{av} = -0.08$, $BF = 0.26$, or in part 3, $t(30) = 1.03$, $p = 0.313$, $g_{av} = 0.62$, $BF = 0.31$. However, as can be seen in Figure 4.2, the difference between the cue-types was very small ($M_{diff} = 0.02$) and the probability of missed responses was numerically higher for the go-consistent cues than for the go-inconsistent cues (i.e. in the opposite direction to that predicted if participants acquired the the word-signal location associations). Thus, if anything, the probability of missed go response data further argues against the hypothesis that participants acquired associations between the cues and the signal locations.

Thus, go performance improved with training but there was no reliable evidence that participants acquired the stimulus-signal associations. Furthermore, the Bayesian analyses provided support for the null hypothesis of no difference between the cue-types in the probability of correct responses.

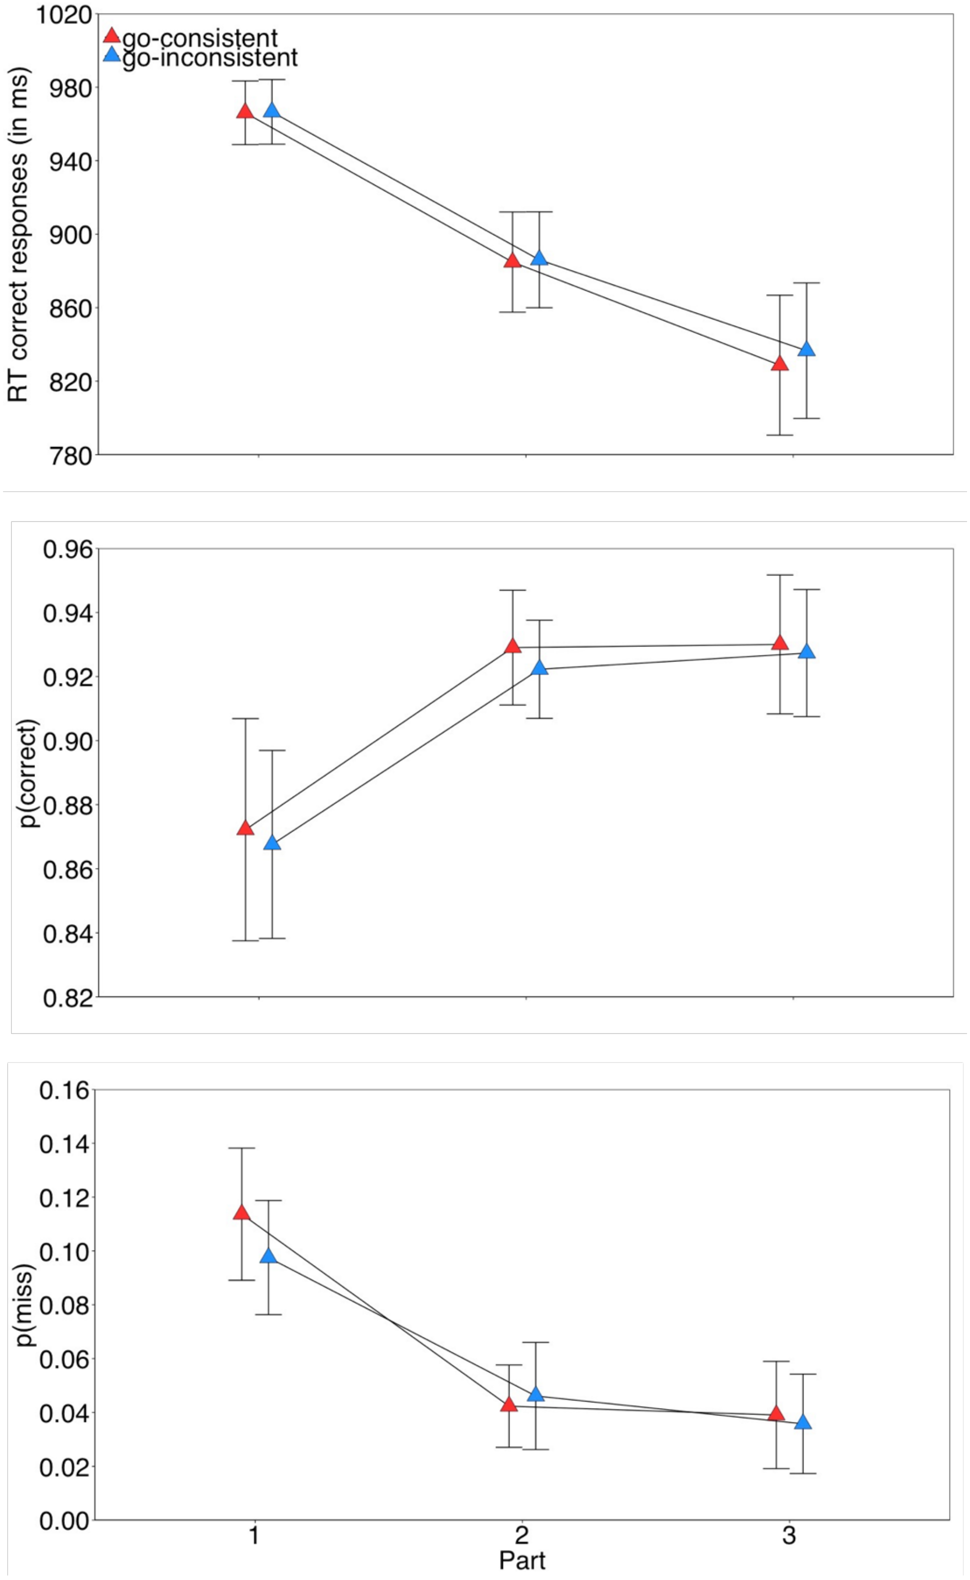

Figure 4.2. Reaction times for the correct go trials (upper panel), the probability of correct go responses (middle panel), and the probability of missed go responses (lower panel) in Experiment 7 as a function of cue-type (go-consistent, go-inconsistent) and part (1-3). Error bars reflect 95% confidence intervals.

Stop analyses. Analyses of the probability of responding revealed a main effect of part (p = 0.027 Table 4.2); as can be seen in Figure 4.3, the p(respond|signal) decreased with practice. There was no main effect of cue-type (p = 0.142; Table 4.2) and no reliable two-way interaction between cue-type and part (p = 0.336; Table 4.2). The Bayes factor for the overall difference between the stop-consistent cues and the stop-inconsistent cues was 0.53.

Expectancy ratings analyses. The expectancy ratings obtained following task completion revealed that participants expected to withhold their response more for stop-associated cues (6.73) than for the go-associated cues (2.58), $F(1, 30) = 60.79$, $p < 0.001$, *gen.* $\eta^2 = 0.628$. The Bayes factor was 1225723. There was no reliable stop/go rating difference between the consistent signal location cues (4.69) and the inconsistent signal location cues (4.62), $F(1, 30) = 0.67$, $p = 0.420$, *gen.* $\eta^2 <$ 0.001. Bayesian analyses supported the null hypothesis of no difference between the consistent signal location cues and the inconsistent signal location cues ($BF = 0.25$). The two-way interaction between stop/go-type and location-type (consistent signal location, inconsistent signal location) was not reliable, $F(1, 30) = 0.70$, $p = 0.783$, *gen.* $\eta^2 < 0.001$. Thus, participants were able to distinguish between the words on the basis of their association with the go and/or stop goal.

To examine the relationship between the expectancy ratings and task performance, I analysed the correlation between the stop-minus-go cue expectancy ratings difference and the part 1-minus-part 2 performance difference. This revealed that the expectancy ratings difference reliably correlated with the RT practice-effect, $r(29) = 0.548$, $p = 0.001$. $BF_{\text{regression}} = 23.16$: participants who expected to withhold their response less for the go-associated cues (and more for the stop-associated cues) became faster throughout training. This suggests that, alongside general practice-effects, the improvement in go performance with training may also reflect the acquisition of associations between the go-associated cues and responding (i.e. going). However, there was no reliable correlation between the stop-minus-go expectancy difference and the difference between part 1-minus-part 3 p(respond|stop) difference, $r(29) = -0.277$, $p = 0.132$, $\text{BF}_{\text{regression}} = 0.84$. Note that uncorrected ps are reported. However, it is important to highlight that whilst the Bayesian regression provided strong evidence for a relationship between the expectancy ratings and the RT practice-effect, it supported neither the null hypothesis of no relationship nor the alternative hypothesis of a relationship between the expectancy ratings and the p(respond|stop) practice-effect.

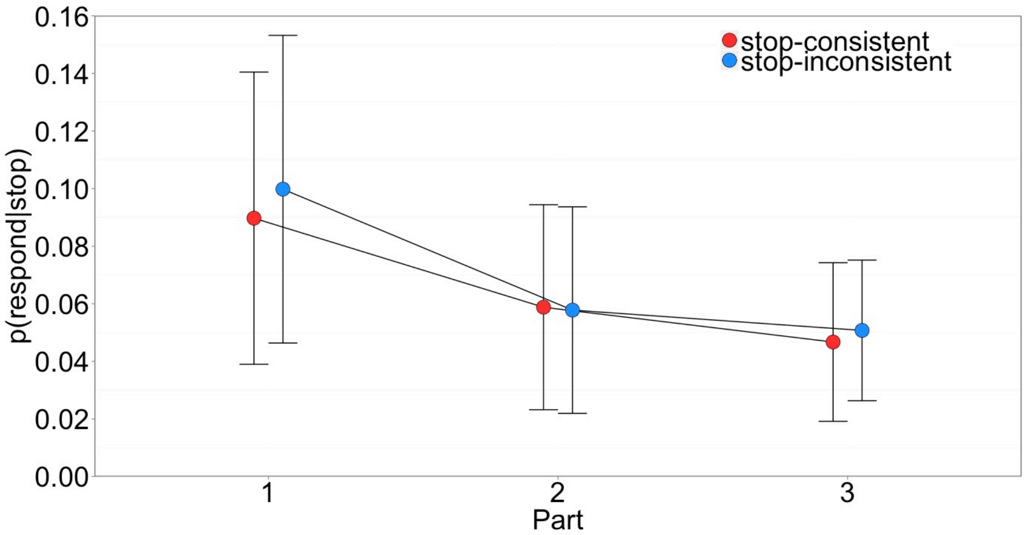

Figure 4.3. The probability of responding on stop trials in Experiment 7 as a function of signal location-type (go-consistent, go-inconsistent) and part (1-3). Error bars reflect 95% confidence intervals.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Go trials: go RT							
Part	2	60	562108.20	342233.78	49.27	**< 0.001**	0.342
Cue-type	1	30	476.52	5086.46	2.81	0.104	< 0.001
Part by cue-type	2	60	513.92	6326.63	2.44	0.099	< 0.001
Go trials: *p*(correct)							
Part	2	60	0.14	0.29	13.91	**< 0.001**	0.149
Cue-type	1	30	0.00	0.05	0.61	0.440	0.001
Part by cue-type	2	60	0.00	0.03	0.12	0.887	< 0.001
Go trials: *p*(miss)							
Part	2	60	0.17	0.18	29.77	**< 0.001**	0.245
Cue-type	1	30	0.00	0.01	3.31	0.079	0.002
Part by cue-type	2	60	0.00	0.02	4.04	**0.026**	0.006
Stop-signal trials: *p*(respond|stop)							
Part	2	60	0.07	0.51	4.34	**0.027**	0.034
Cue-type	1	30	0.00	0.01	2.28	0.142	< 0.001
Part by cue-type	2	60	0.00	0.03	1.11	0.336	< 0.001

Table 4.2. Overview of repeated measures Analyses of Variance in Experiment 7. In the no-signal RT analysis, incorrect and missed no-signal trials were removed. *p*s < 0.05 are highlighted in bold.

Interim Discussion

Go and stop performance improved with task practice. Furthermore, the expectancy ratings obtained following task completion indicate that participants acquired the stimulus-stop/go contingencies in play: for go RTs, these expectancies reliably correlated with the magnitude of the RT training effect. However, despite finding evidence of stop/go learning, there was no support for the idea that participants acquired the associations between the cues and the location of the go/stop signal. The magnitude of the performance improvements were similar for the cues that were consistently associated with a single signal location and for the cues that were associated with two signal locations with equal probability. The Bayesian analyses for the probability of correct responses on go trials provided strong support for the null hypothesis of no difference between the go-consistent and go-inconsistent cues. However, the Bayes factors for the go RTs, the probability of missed go responses, and the probability of responding on stop trials provided only tentative support for the null hypothesis.

It is important to note, however, that there are some limits on the conclusions that can be drawn from this experiment (partly because it was originally designed an an ERP pilot). First, stop-associated cues (and go-associated cues) in Experiment 7 were associated with the stop goal (or go goal for the go-associated cues) on 100% of presentations. Consequently, the design was suboptimal for dissociating the effects of stimulus-stop/go learning and stimulus-signal learning on behavioural task performance as it was not possible to compare the stop-associated and go-associated cues during training. Conclusions about the acquisition of item-specific stop/go associations in this experiment were based entirely on the expectancy ratings obtained following task completion. Thus, it is possible that signal learning has an additive effect on go/stop learning in a way that I could not detect in this experiment. Second, although the use of coloured targets is common in the ERP visual attention literature (e.g. Woodman, 2013), a stimulus is more likely to capture attention if it is highly perceptually salient (e.g. if it is brightly coloured; Folk, Remington, & Johnston, 1992). Thus, it is possible that the use of coloured signals made it too easy to distinguish between the go and stop signals, which could have reduced or diminished the effect of prior learning on selective attention. I address these issues in Experiment 8.

Experiment 8

In Experiment 8, I made several changes to the task design. First, participants did not have to respond to the words, but were instead instructed to attend to the words in anticipation for a recall test later in the experiment. It is possible that responding to the words in Experiment 8 influenced learning. For example, extensive research has shown that responding is faster and more accurate when the stimulus is presented in the same laterality as the response key (i.e. the ‘Simon effect’; Simon & Rudell, 1967). Therefore, the absence of a signal location training effect could be due to the conflict between the associative retrieval of the trained signal location and the response hand[33]. Second, the go signal and the stop signal differed in stimulus orientation (square, diamond) rather than in colour and I increased the distance between the words and the signals. I expected that these modifications would make the

---

[33] Post-hoc analyses revealed some support for this idea: the probability of correct responses on go trials was lower on incompatible trials (where the signal location of the go signal and the response hand were in opposite spatial signal locations; 0.89) than on compatible trials (where the signal location of the go signal and the response hand where in the same spatial signal location; 0.93), *t*(30) = 5.82, $p < 0.001$, $g_{av} = 0.84$. Similarly, was also a reliable difference in the go RTs between the incompatible trials (897 ms) and the compatible trials (884 ms), *t*(30) = -5.85, $p < 0.001$, $g_{av} = -0.20$.

go and stop signals harder to distinguish and thus reduce the influence of stimulus saliency on signal detection (see Awh, Belopolsky, et al., 2012). Third, I manipulated the contingency between the words and stopping and the contingency between the words and the signal location, such that the words were not paired with stopping (or going) and the same stop-signal location on all presentations (i.e. instead of a 100%-contingency mapping, I used a 80%-contingency mapping). It was anticipated that this would allow me to better index the effects of stimulus-stop/go learning and/or stimulus-signal location learning on task performance throughout training. In addition to the stop expectancy ratings obtained in Experiment 7, participants also rated where they expected the stop signal to appear on the screen (i.e. to the left or to the right of word) following task completion. Finally, it was possible that the behavioural measures (i.e. go RTs and the probability of responding on stop trials) were not sensitive enough to capture the effects of signal location learning. Therefore, in Experiment 8 I also obtained eye-movement data. The gaze position of the right eye was tracked throughout task performance to provide an additional dependent variable of shifts of attentional focus within each trial. Previous research has shown that gaze position can be used as of selective attention in learning tasks (e.g. Rehder & Hoffman, 2005). In the present experiment, two measures were obtained; the number of fixations on each trial and the distance (in pixels) between the fixation location and the trained signal location (i.e. to the left or to the right of the cue). If participants acquired the associations between the cues and the signal location, it was predicted that (1) there would be an increased number of fixations on trials in which the signal appeared in the untrained location than on trials in the signal appeared in the trained location and (2) the distance from the fixated location and the trained signal location would decrease with task practice.

Method

Participants. Thirty-two volunteers from the University of Exeter participated for monetary compensation (£10) or partial course credit (*M* = 20.25 years, *SD* = 3.13, 24 females, 31 right-handed). Two participants were excluded and replaced due to technical difficulties with the eye-tracking software; no participants were replaced for performance-related reasons.

Apparatus, stimuli, and procedure. The stimuli were presented on a 17-inch CRT monitor (screen size: 1024 × 768 pixels) using Psychtoolbox (Brainard, 1997). Eye movements were recorded throughout task performance. An EyeLink 1000 Desktop Mount camera system (SR Research, Ottawa, Canada) was calibrated at the beginning of each block (there were 120 trials between calibrations)[34]. The gaze position of the right eye was tracked throughout each block (sampling rate: 500 Hz). The EyeLink was calibrated and controlled via Psychtoolbox (Cornelissen, Peters, et al., 2002).

---

[34] Note that the numbers of trials (and time durations) between calibrations using this system is consistent with previous research (Verbruggen, Stevens, et al., 2014).

The cues consisted of a word in white lowercase font (Courier 16 point) presented on a black background (Figure 4.4). For half the participants, the go signal was a diamond and the stop signal was a square; for the other half, the go signal was a square and the stop signal was a diamond (size: 24 × 24 pixels). Aside from the orientation, the go signal and stop signal were identical. On half of all trials, the go signal or stop signal was presented to the left of the word (x = -256 pixels, -6.8 cm from the centre); on the other half, the go signal or stop signal was presented to the right of the word (x = 256 pixels, 6.8 cm from the centre). The control signal was always a circle, presented on the opposite side to the go/stop signal (Figure 4.3). I randomly created three lists of 12 four-letter words (Appendix I) anew for each participant. Two word lists were used in the main task, and the remaining word list was used in a recognition ratings task following completion of the main task. To encourage participants to attend to the words, the presentation of the two word lists alternated on a block-by-block basis in the main task (e.g. wordlist 1, wordlist 2, wordlist 1, …). The experiment consisted of three parts with four blocks per part. There were 120 trials per block. Each word was presented 10x per block; there were 7 presentations that were paired with a consistent response goal (e.g. stop) and consistent signal location (e.g. left stop signal); 1 presentation that was consistent with the trained response (e.g. stop) but inconsistent with the trained signal location (e.g. right stop signal); 1 presentation that was inconsistent with the trained response (e.g. go) but consistent with the trained signal location (e.g. left go signal); and 1 presentation that was inconsistent with the trained response (e.g. go) and inconsistent with the trained signal location (e.g. right go signal). There were four cue-types in equal proportions (see Table 4.3); *stop-left* cues were mostly presented on left stop signal trials; *stop-right* cues were mostly presented on right stop signal trials; *go-left* cues were mostly presented on left go signal trials; and *go-right* cues were mostly presented on right go signal trials.

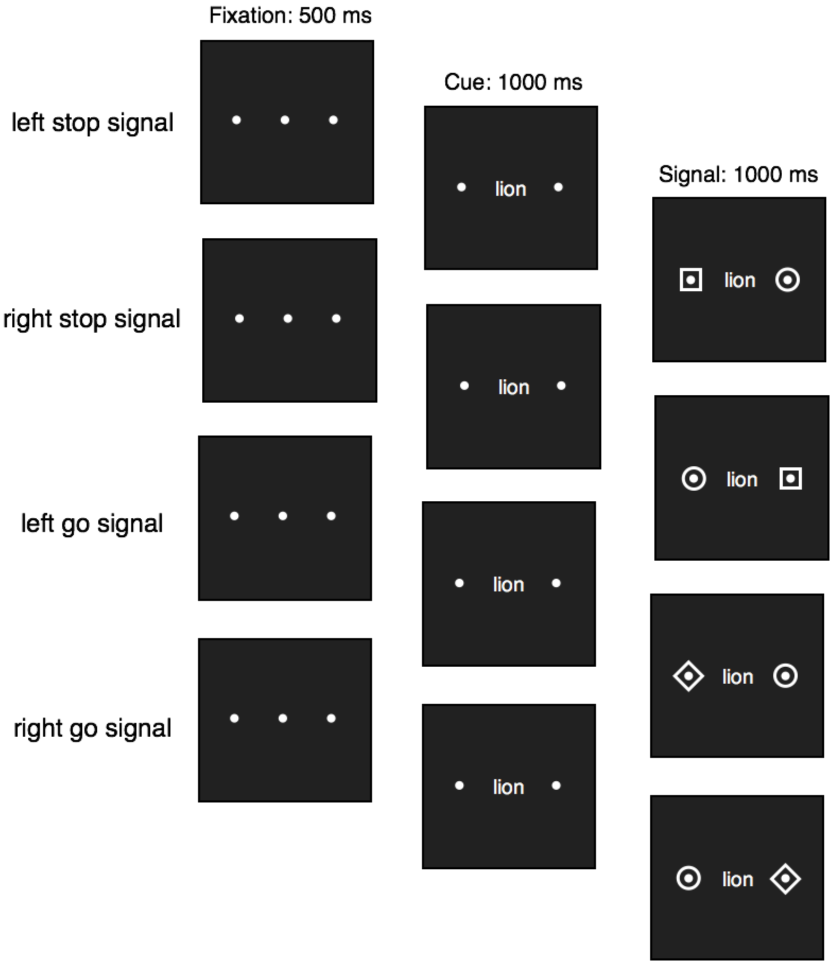

Figure 4.4. The left stop signal, right stop signal, left go signal, and right go signal trial structures in Experiment 8. Each trial began with the presentation of three fixation dots, one in the centre of the screen, one to the left of the centre, and one to the right of the centre for 500 ms. Following the removal of the fixation dots, a four-letter word (i.e. the cue) was presented in the centre of the screen. After 1000 ms, the go/stop signal appeared around the dot to the left or to the right of the word and a control signal (which was presented on go and stop trials) appeared around the dot on the other side. The word and the signals remained on the screen for 1000 ms (hence, MAXRT was 2000 ms). In this example, the stop-signal is represented as a square, and the go signal is represented by a diamond. The control signal was always a circle.

Cue-type	No. of words	% left stop signal	% right stop signal	% left go signal	% right go signal
Stop-left	3	70	10	10	10
Stop-right	3	10	70	10	10
Go-left	3	10	10	70	10
Go right	3	10	10	70	10

Table 4.3. The proportion of left stop-signal, right stop-signal, left go-signal, and right go-signal trials per block as a function of cue-type in Experiment 8. Note that there were 24 words used in the experimental task in total, with two lists of 12 words which alternated on a block-by-block basis (i.e. there were 12 words per block). Note: the overall *p*(stop-signal) was 0.50.

On each trial, three black fixation dots were presented marking the central, left, and right stimulus locations (Figure 4.4). After 500 ms, the central fixation dot was replaced with a word. After 1000 ms, one shape appeared around the fixation dot to the left of the word and one shape appeared around the fixation dot on the right of the word. There were three shapes (square, diamond, and circle): for half of the participants, the go signal was the square and the stop signal was the diamond; for the other half, the go signal was the diamond and the stop signal was the square. The control signal was always the circle. On go trials, participants responded by pressing the spacebar on a keyboard with their right index finger. Participants were informed that the aim of the study was to investigate memory for simple four-letter words, as such they were instructed to attend to the words and try to remember as many as possible for a memory recall test following task completion. At the end of each block, I presented as feedback to the participant their mean RT on go trials, the number of go errors, the number of missed go responses, and the percentage of failed stops. After 15 s had elapsed, participants had to press a key to start the next block.

Following completion of the main task, participants completed two ratings tasks in which each word was again presented on the screen. The order of the words was randomised anew for each participant. First, participants were asked to rate “How much do you expect to withhold your response when this word is presented?” on a scale, ranging from 1 (*I definitely do not think this word indicates that I have to withhold my response)* to 9 (*I definitely think this word indicates that I have to withhold my response).* Second, participants were asked to rate “Where do you expect the stop signal to appear on the screen?” on a scale, ranging from 1 (*I definitely expect the stop signal to appear on the left)* to 9 (*I definitely expect the stop signal to appear on the right).*

Following completion of the main task and the expectancy ratings, I probed memory for the words presented in the task. Before proceeding, however, it is important to highlight these measures were obtained following completion of the stop/go and signal location expectancy ratings tasks meaning that the recall of the words could be higher than if I obtained these measures immediately following completion of the main task. Note, however, that memory performance was not of primary interest in the present chapter. Nevertheless, I report these data because they provide an indirect measure of the extent to which participants attended to the words during the main task and/or expectancy ratings tasks. First, participants performed a free recall test whereby they were instructed to write down as many words as they could remember from the main task in 80 seconds. This procedure was based on previous research on the recall of automatic processes (Snyder, Ashitaka, Shimada, Ulrich, & Logan, 2014). Second, participants were presented with each of the 24 words from the main task and 12 novel words that had not been previously presented. Participants were asked to rate "How do you think that this word was presented in the main task?" on a scale, ranging from 1 (*I definitely think this word was presented in the main task)* and 9 (*I definitely do not think this word was presented in the main task)*[35].

Analyses. The behavioural analyses were identical to Experiment 7, except that behavioural performance data was analysed as a function of stop/go-type (stop-associated cues, go-associated cues), signal-location (congruent with training, incongruent with training), and part. Furthermore, I did not analyse the *p*(miss) as values were very low. Figure 4.4 presents an overview of the descriptive statistics for the go RTs and the *p*(respond|stop) and Table 4.4 presents an overview of the inferential statistics. Note that, unlike in Experiment 7, I could not examine the *p*(correct) as participants responded with only one keypress on go trials (i.e. there was no go choice task).

For the analyses of the eye movement data, the eye data were exported using the Eyelink Data Viewer (SR Research, Ottawa, Canada) for each participant. I generated a file with information about all fixations and a file with trial sequence information, and integrated these files using R for further analyses. I excluded participants from the analyses when no fixation was registered at the beginning of the trial event (i.e. the presentation of the word cue) on more than 15% of trials, as this could indicate that eye-movement registration was suboptimal. Based on this criterion, I excluded four participants. Note that the inclusion of these participants did not substantially alter the overall pattern of behavioural results (not shown). I also excluded all fixations that were off screen (0.2%). In the analyses of the eye data, I focussed on the number of fixations and the fixation signal location for two intervals: (1) the interval between the word (cue) presentation and the presentation of the stop/go signal, and (2) the interval after the presentation of the stop/go signal. The analyses in the word interval (prior to the presentation of the signal) indexes anticipatory 'proactive' shifts in attention whereas the analyses in the signal interval indexes 'reactive' shifts in attention following signal presentation. Eye-movements made 471 ms after the presentation of the stop/go signal were excluded. I based this value on the mean average go RT value (note, this value is also broadly consistent with

---

[35] Participants recalled a mean average of 15 words (*SD* = 5.18) out of 24 words used in the task. A one-sample t-test revealed that this was significantly greater than zero, $t(31) = 16.34$, $p < 0.001$. The recognition ratings task also revealed a highly reliable difference between the old words used in the task (1.17) and the novel words (8.36), $t(31) = -30.44$, $p < 0.001$, $g_{av} = -9.02$. Thus, it seems that participants did attend to the words during task completion, even though the words were task irrelevant.

the cut-off used in Chapter 3). For the location analyses, I analysed the distance between the fixated location and the location of the consistently trained signal [distance = sqrt(x-coordinate^2 + y-coordinate^2)]. Note, if the number of fixations for a particular interval = 1, then the participant did not move their eyes during this time. Eye movement data in the interval between the word presentation and the presentation of the signals was analysed as a function of stop/go-type (stop-associated cues, go-associated cues), location-type (left-associated cues, right-associated cues), and part (1-3). Eye movement data in the interval after the presentation of the stop/go signal was analysed as a function of stop/go-type (stop-associated cues, go-associated cues), signal-location (congruent with training, incongruent with training), signal-type (stop signal, go signal), and part (1-3). The descriptive statistics for the eye movement data can be found in Tables 4.5 and 4.7 and the inferential statistics can be found in Table 4.6 and 4.8 (separate analyses for the stop-associated items and for the go-associated items are presented in Tables 4.9 and 4.10).

Results

Go analyses. Analyses of go RTs revealed that the main effect of part was reliable ($p = 0.005$; Table 4.5; Figure 4.4). However, there was no reliable difference between the stop-associated and go-associated cues ($p = 0.915$; Table 4.4) and no reliable difference between the consistent and inconsistent signal location trials ($p = 0.362$; Table 4.4). The two-way interaction between stop/go-type and the signal-location was not reliable ($p = 0.865$; Table 4.5). All other interactions were not reliable ($ps \geq 0.551$; Table 4.4). Bayesian analyses supported the null hypotheses of no difference between the stop-associated cues and the go-associated cues, $BF = 0.19$, and of no difference between the consistent and inconsistent signal location trials, $BF = 0.28$. Taken together, these analyses suggest that performance benefited from non-specific task practice, but that the stimulus-stop/go contingencies and the stimulus-signal contingencies did not influence go task performance.

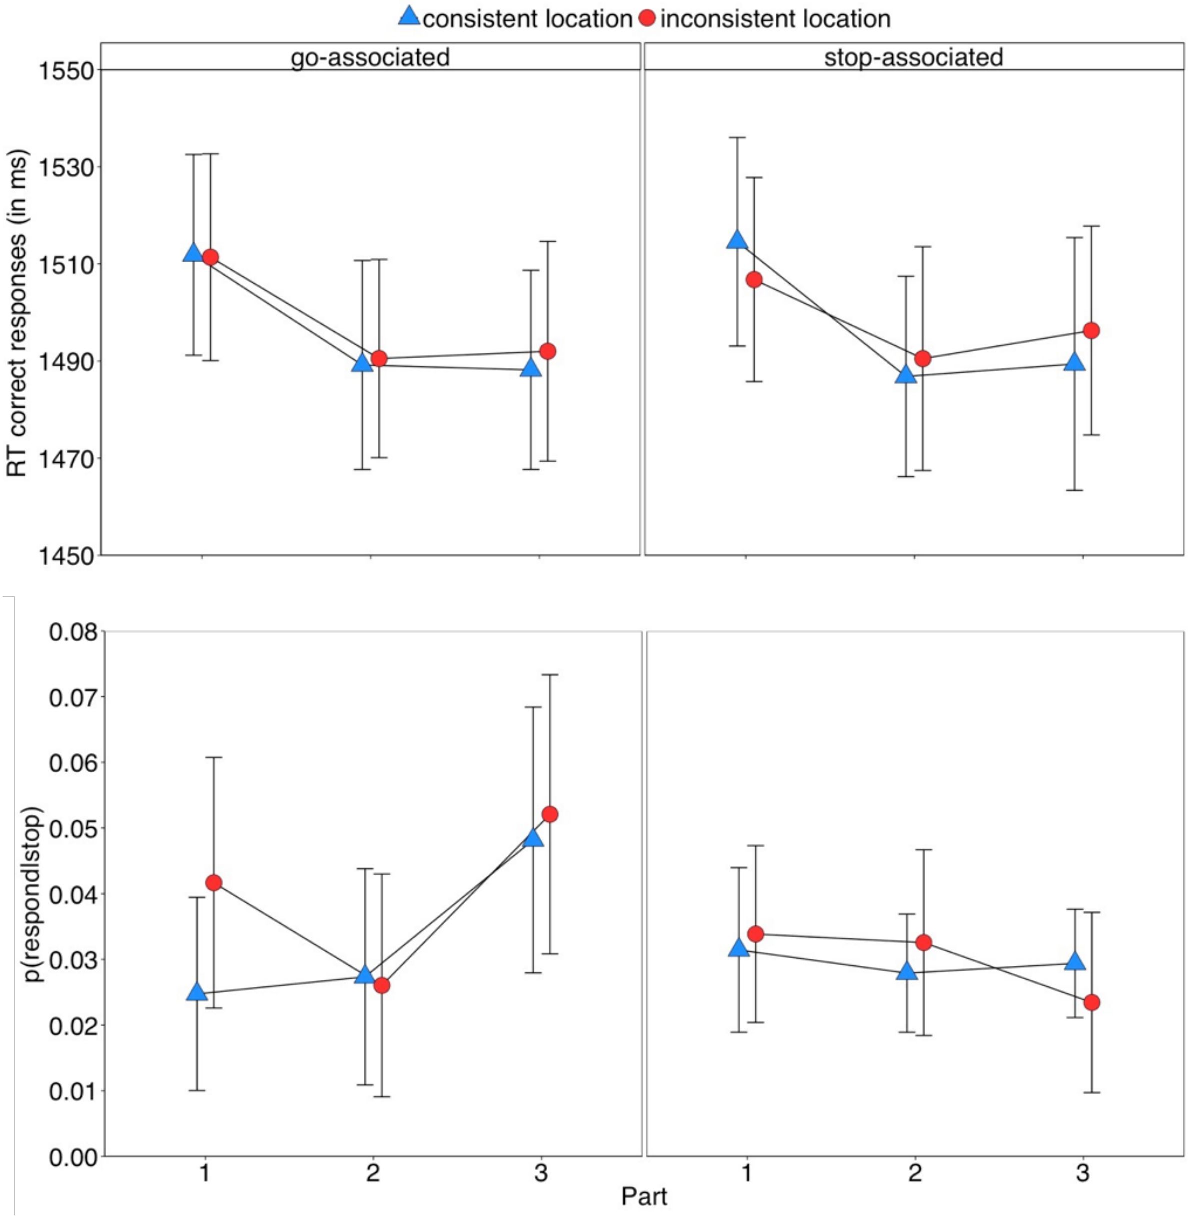

Figure 4.5. Reaction times for the correct go trials (upper panel) and the probability of responding on stop trials (lower panel) in Experiment 8 for the go-associated items (left panels) and for the stop-associated items (right panels) as a function of signal location-type (consistent signal location, inconsistent signal location) and part (1-3). Error bars reflect 95% confidence intervals.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Combined analysis							
Go trials: go RT							
Part	2	62	37140.47	73577.01	15.65	**< 0.001**	0.027
Stop/go-type	1	31	3.83	10330.14	0.01	0.915	< 0.001
Signal-location	1	31	148.01	5362.99	0.86	0.362	< 0.001
Part by stop/go-type	2	62	311.89	16074.06	0.60	0.551	< 0.001
Part by signal-location	2	62	1515.16	19538.16	2.40	0.099	0.001
Stop/go-type by signal-location	1	31	9.63	10247.00	0.03	0.866	< 0.001
Part by stop/go-type by signal-location	2	62	535.26	26840.59	0.62	0.536	< 0.001
Stop-signal trials: *p*(respond|stop)							
Part	2	62	0.01	0.08	2.26	0.112	0.009
Stop/go-type	1	31	< 0.001	0.04	3.47	0.072	0.007
Signal-location	1	31	< 0.001	0.03	1.29	0.264	0.002
Part by stop/go-type	2	62	0.01	0.08	5.54	**0.006**	0.020
Part by signal-location	2	62	0.00	0.06	0.98	0.368	0.003
Stop/go-type by signal-location	1	31	0.00	0.05	0.58	0.451	0.001
Part by stop/go-type by signal-location	2	62	0.00	0.06	0.92	0.402	0.003
Stop-associated cues only							
Go trials: go RTs							
Part	2	62	17459.93	57184.14	9.47	**< 0.001**	0.024
Signal-location	1	31	41.07	10713.65	0.12	0.733	< 0.001
Part by signal-location	2	62	1900.69	37094.95	1.59	0.213	0.003
Stop-signal trials: *p*(respond|stop)							
Part	2	62	0.00	0.04	0.97	0.386	0.006
Signal-location	1	31	0.00	0.02	0.01	0.916	< 0.001
Part by signal-location	2	62	0.00	0.05	0.63	0.534	0.005
Go-associated cues only							
Go trials: go RTs							
Part	2	62	19992.42	32466.92	19.09	**< 0.001**	0.030
Signal-location	1	31	116.57	4896.34	0.74	0.397	< 0.001
Part by signal-location	2	62	149.73	9283.80	0.50	0.609	< 0.001
Stop-signal trials: *p*(respond|stop)							
Part	2	62	0.02	0.12	4.79	**0.012**	0.038
Signal-location	1	31	0.00	0.06	1.10	0.302	0.004
Part by signal-location	2	62	0.00	0.08	1.15	0.323	0.006

Table 4.4. Overview of repeated measures Analyses of Variance in Experiment 8, with stop/go-type (stop-associated, go-associated), signal-location (congruent signal location, incongruent signal location), and part (1-3) as within-subjects factors. In the no-signal RT analysis, incorrect and missed no-signal trials were removed. *p*s < 0.05 are highlighted in bold.

Stop analyses. Analyses of the probability of responding revealed that the main effect of part was not reliable ($p = 0.112$; Table 4.4; Figure 4.5). Numerically, the *p*(respond|signal) was lower for the stop-associated cues (0.03) than for the go-associated cues (0.04). However, this difference did not reach statistical significance ($p = 0.072$; Table 4.4). However, the two-way interaction between part and stop/go-type was reliable ($p = 0.006$; Table 4.4); consistent with the 'automatic inhibition' hypothesis, the *p*(respond|stop) became lower for the stop-associated cues than for the go-associated cues (Figure 4.5). This indicates that the stimulus-stop/go associations influenced stop performance.

However, there was no reliable difference between the consistent signal location trials and the inconsistent signal location trials (*p* = 0.264; Table 4.4) and the two-way interaction between stop/go-type and signal-location was also not reliable (p = 0.451; Table 4.4). The corresponding Bayes factor for the difference between the consistent and inconsistent signal location trials was 0.34. All other interactions were also not reliable (*p*s ≥ 0.381; Table 4.4).

Eye movement analyses.

Word interval. Analyses of the number of eye movements in the interval between the presentation of the word and the presentation of the go/stop signal revealed a reliable main effect of part (*p* < 0.001; Table 4.6) reflecting an increasing number of fixations as a function of task practice (Table 4.5). The mean number of fixations was slightly higher for the cues associated with a right signal (2.07) than for the cues associated with a left signal (2.05; *p* = 0.044; Table 4.6). However, the corresponding Bayesian analysis was inconclusive (*BF* = 1.04). All other main effects and interactions were not reliable (*p*s ≥ 0.201; Table 4.6).

Analyses of the average distance between the fixated signal location and the trained signal location did not reveal any reliable main effects or interactions (*p*s ≥ 0.416; Table 4.6) and Bayesian analyses confirmed that there was no difference between the stop-associated and go-associated cues (*BF* = 0.25) and between the consistent-left and consistent-right cues (*BF* = 0.24).

	Fixation				**Distance**			
	left-associated		***right-associated***		***left-associated***		***right-associated***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Part 1								
Go-associated	1.83	0.37	1.85	0.40	256.12	17.82	257.19	16.55
Stop-associated	1.85	0.41	1.86	0.39	258.93	17.48	255.87	19.65
Part 2								
Go-associated	2.11	0.49	2.13	0.53	259.43	13.64	255.70	13.79
Stop-associated	2.09	0.51	2.16	0.56	258.74	12.05	254.59	13.23
Part 3								
Go-associated	2.22	0.57	2.21	0.58	260.04	15.67	256.70	13.54
Stop-associated	2.21	0.60	2.21	0.59	258.21	11.44	255.08	14.20

Table 4.5. Overview of the number of fixations and average distance between the fixation signal location and the trained signal location (in pixels) in the interval between the word presentation and the stop/go signal presentation for each stop/go-type (go-associated, stop-associated), location-type (left-associated cues, right-associated cues), and part (1-3). *M* = mean; *SD* = standard deviation.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Number of fixations							
Stop/go-type	1	27	0.00	0.25	0.20	0.657	< 0.001
Location-type	1	27	0.03	0.17	4.47	**0.044**	< 0.001
Part	2	54	8.08	9.50	22.95	**< 0.001**	0.089
Stop/go-type by location-type	1	27	0.01	0.14	1.72	0.201	< 0.001
Stop/go-type by part	2	54	0.01	0.57	0.30	0.744	< 0.001
Location-type by part	2	54	0.02	0.51	1.32	0.276	< 0.001
Stop/go-type by location-type by part	2	54	0.01	0.43	0.61	0.548	< 0.001
Distance from trained signal							
Stop/go-type	1	27	32.93	1959.09	0.45	0.506	< 0.001
Location-type	1	27	623.21	46612.73	0.36	0.553	0.008
Part	2	54	14.52	2131.92	0.18	0.833	< 0.001
Stop/go-type by location-type	1	27	43.88	1943.60	0.61	0.442	0.001
Stop/go-type by part	2	54	87.90	2660.46	0.89	0.416	0.001
Location-type by part	2	54	132.60	12750.41	0.28	0.630	0.002
Stop/go-type by location-type by part	2	54	77.13	3471.37	0.60	0.552	0.001

Table 4.6. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the word interval. Stop/go-type (stop-associated, go-associated), location-type (consistent-left, consistent-right), and part (1-3) were within-subjects factors. *p*s < 0.05 are highlighted in bold.

Signal interval. Analyses of the number of eye movements in the interval after the stop/go signal presentation revealed a main effect of part ($p < 0.001$; Table 4.8) reflecting a decreasing number of fixations as a function of task practice (Table 4.7). Thus, participants moved their eyes less in this interval with task practice. Participants made slightly more eye-movements on stop-signal trials (1.74) than on go-signal trials (1.68; $p < 0.001$; Table 4.8). However, there was no reliable difference between the trials in which the signal location was congruent with training and the trials in which the signal location was incongruent with training ($p = 0.473$; Table 4.8). There was also no reliable two-way interaction between stop/go-type and signal-location ($p = 0.897$; Table 4.8). All other main effects and interactions were not reliable (ps ≥ 0.128; Table 4.8). Bayesian analyses supported the null hypotheses of no difference between the stop-associated cues and the go-associated cues ($BF = 0.20$) and the null hypothesis of no difference between the consistent signal location trials and the inconsistent signal location trials ($BF = 0.26$).

Analyses of the average distance measure revealed that the main effect of part was reliable ($p < 0.001$; Table 4.8) reflecting an increasing distance between the fixation and the trained signal as a function of task practice (Table 4.7). The average distance was greater on stop-signal trials (252 pixels) than on go-signal trials (245 pixels; $p < 0.001$; Table 4.8) and the two-way interaction between part and signal was reliable ($p < 0.001$; Table 4.8). Follow-up analyses revealed that, although the increase was numerically greater on the go-signal trials, the main effect of part was reliable on both go signal trials, $F(2, 54) = 72.94$, $p < 0.001$, *gen.* $\eta^2 = 0.483$, and on stop-signal trials, $F(2, 54) = 12.29$, $p < 0.001$, *gen.* $\eta^2 = 0.107$. Thus, this pattern of results indicate a general (i.e. non-specific) effect of training in which participants fixated on the signals less as a function of training.

	Fixation				**Distance**			
	congruent signal location		***incongruent signal location***		***congruent signal location***		***incongruent signal location***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Go signal								
Part 1								
Go-associated	1.84	0.49	1.85	0.48	217.59	31.48	216.02	38.64
Stop-associated	1.87	0.52	1.82	0.45	221.06	32.37	219.40	33.90
Part 2								
Go-associated	1.59	0.31	1.59	0.34	258.93	23.59	256.37	27.63
Stop-associated	1.58	0.34	1.62	0.35	259.32	31.35	253.38	29.22
Part 3								
Go-associated	1.60	0.30	1.56	0.36	256.21	21.59	253.83	29.64
Stop-associated	1.61	0.35	1.57	0.32	261.26	27.53	266.81	29.62
Stop signal								
Part 1								
Go-associated	1.90	0.58	1.92	0.59	243.47	32.62	243.41	29.02
Stop-associated	1.92	0.55	1.91	0.58	242.18	27.58	244.35	33.01
Part 2								
Go-associated	1.70	0.41	1.67	0.35	253.68	31.33	258.17	33.29
Stop-associated	1.66	0.33	1.68	0.35	259.14	20.65	257.89	35.64
Part 3								
Go-associated	1.65	0.35	1.64	0.33	252.25	32.30	261.60	36.34
Stop-associated	1.65	0.32	1.62	0.33	256.15	20.99	255.40	29.06

Table 4.7. Overview of the number of fixations and average distance between the fixation signal location and the trained signal location (in pixels) in the interval after the stop signal for each stop/go-type (stop-associated, go-associated), signal-location (congruent signal location, incongruent signal location), signal-type (go signal, stop signal), and part (1-3). *M* = mean; *SD* = standard deviation.

The difference between the go-associated and the stop-associated cues was not reliable (*p* = 0.170; Table 4.7). The Bayes factor for the corresponding difference was 0.47. There was no reliable two-way interaction between stop/go-type and part (*p* = 0.569; Table 4.8). The three-way interaction between stop/go-type, signal-type, and part did reach significance (*p* = 0.033; Table 4.8). However, the two-way interaction between stop/go-type and part did not reach significance on go-signal trials, *F*(2, 54) = 2.86, *p* = 0.066, *gen.* η^2 = 0.012, or on stop-signal trials, *F*(2, 54) = 0.50, *p* = 0.586, *gen.* η^2 = 0.002. Furthermore, this three-way interaction does not fit with my prior predictions; it is unclear why stop/go learning would influence the fixation distance on go-signal trials more than on stop-signal trials (or vice versa). Thus, it is possible that the reliable three-way interaction reflects a spurious effect (for a discussion of issues surrounding multiple comparisons in ANOVAs, see Cramer, van Ravenzwaaij, et al., 2016).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Number of fixations							
Stop/go-type	1	27	0.00	0.37	0.00	0.968	< 0.001
Signal-type	1	27	0.71	0.78	24.40	**< 0.001**	0.006
Part	2	54	9.54	11.51	22.40	**< 0.001**	0.079
Signal-location	1	27	0.01	0.43	0.53	0.473	< 0.001
Stop/go-type by signal-type	1	27	0.01	0.27	0.62	0.436	< 0.001
Stop/go-type by part	2	54	0.00	0.79	0.04	0.956	< 0.001
Signal-type by part	2	54	0.02	1.25	0.47	0.627	< 0.001
Stop/go-type by signal-location	1	27	0.00	0.44	0.02	0.897	< 0.001
Signal-type by signal-location	1	27	0.00	0.27	0.37	0.549	< 0.001
Part by signal-location	2	54	0.04	0.69	1.48	0.237	< 0.001
Stop/go-type by signal-type by part	2	54	0.00	0.78	0.14	0.867	< 0.001
Stop/go-type by signal-type by signal-location	1	27	0.00	0.41	0.03	0.863	< 0.001
Stop/go-type by part by signal-location	2	54	0.07	0.86	2.17	0.128	0.001
Signal-type by part by signal-location	2	54	0.02	0.71	0.79	0.460	< 0.001
Stop/go-type by signal by signal-location	2	54	0.00	0.78	0.17	0.848	< 0.001
Distance from trained signal							
Stop/go-type	1	27	717.51	9749.93	1.99	0.170	0.003
Signal-type	1	27	8934.07	11420.89	21.12	**< 0.001**	0.036
Part	2	54	105647.00	54349.03	52.48	**< 0.001**	0.307
Signal-location	1	27	34.11	5498.49	0.17	0.686	< 0.001
Stop/go-type by signal-type	1	27	454.35	6955.59	1.76	0.195	0.002
Stop/go-type by part	2	54	319.13	15138.45	0.57	0.569	0.001
Signal-type by part	2	54	26172.32	19125.01	36.95	**< 0.001**	0.099
Stop/go-type by signal-location	1	27	97.64	7357.98	0.36	0.554	< 0.001
Signal-type by signal-location	1	27	590.42	8237.09	1.94	0.176	0.002
Part by signal-location	2	54	552.39	11115.25	1.34	0.269	0.002
Stop/go-type by signal-type by part	2	54	1383.91	10308.01	3.62	**0.033**	0.006
Stop/go-type by signal-type by signal-location	1	27	381.05	5917.42	1.74	0.198	0.002
Stop/go-type by part by signal-location	2	54	226.35	12013.26	0.51	0.603	0.001
Signal-type by part by signal-location	2	54	94.38	14099.98	0.18	0.832	< 0.001
Stop/go-type by signal by signal-location	2	54	796.31	12637.36	1.70	0.193	0.003

Table 4.8. Overview of Analyses of Variance performed to compare number of fixations and the average distance in the signal interval. Stop/go-type (stop-associated, go-associated), signal-location (congruent location, incongruent location), signal-type (stop-signal, go-signal), and part (1-3) were within-subjects factors. *p*s < 0.05 are highlighted in bold.

There was also no reliable difference between trials in which the signal location was congruent with training and trials in which the signal was incongruent with training ($p = 0.473$; Table 4.8). The Bayesian analyses supported the null hypothesis of no difference between these trials ($BF = 0.22$). There was no reliable two-way interaction between stop/go-type and signal-location ($p = 0.897$; Table 4.8). All other main effects and interactions were also not reliable (ps ≥ 0.170; Table 4.8).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Word interval							
Number of fixations							
Location-type	1	27	0.03	0.14	6.62	**0.016**	0.001
Part	2	54	3.83	5.72	18.08	**< 0.001**	0.081
Location-type by part	2	54	0.03	0.45	1.86	0.166	0.001
Distance from trained signal							
Location-type	1	27	498.92	22250.18	0.61	0.443	0.014
Part	2	54	20.61	2325.67	0.24	0.785	0.001
Location-type by part	2	54	10.45	9218.58	0.03	0.919	< 0.001
Stop-signal interval							
Number of fixations							
Signal-type	1	27	0.29	0.59	13.35	**0.001**	0.005
Part	2	54	4.86	6.45	20.35	**< 0.001**	0.082
Signal-location	1	27	0.01	0.53	0.30	0.589	< 0.001
Signal-type by part	2	54	0.01	1.13	0.28	0.736	< 0.001
Signal-type by signal-location	1	27	0.00	0.27	0.34	0.567	< 0.001
Part by location	2	54	0.08	0.73	3.06	0.060	0.002
Signal-type by part by signal-location	2	54	0.01	0.59	0.54	0.585	< 0.001
Distance from trained signal							
Signal-type	1	27	2679.46	6939.97	10.42	**0.003**	0.023
Part	2	54	54460.63	33051.48	44.49	**< 0.001**	0.322
Signal-location	1	27	8.16	8808.22	0.03	0.875	< 0.001
Signal-type by part	2	54	14214.75	17486.43	21.95	**< 0.001**	0.110
Signal-type by signal-location	1	27	11.42	3670.03	0.08	0.774	< 0.001
Part by signal-location	2	54	516.82	9924.06	1.41	0.254	0.004
Signal-type by part by signal-location	2	54	524.36	12363.44	1.15	0.319	0.005

Table 4.9. Overview of Analyses of Variance performed to compare number of fixations and the average distance for the *stop-associated* cues. In the word interval, location-type (consistent-left, consistent-right), and part (1-3) were within-subjects factors. In the stop-signal interval, signal-location (congruent location, incongruent location), signal-type (stop-signal, go-signal), and part (1-3) were within-subjects factors. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Word interval							
Number of fixations							
Location-type	1	27	0.00	0.17	0.41	0.529	< 0.001
Part	2	54	4.25	4.36	26.37	**< 0.001**	0.097
Location-type by part	2	54	0.00	0.49	0.19	0.827	< 0.001
Distance from trained signal							
Location-type	1	27	168.18	26306.15	0.17	0.681	0.004
Part	2	54	81.82	2466.71	0.90	0.405	0.002
Location-type by part	2	54	199.29	7003.19	0.77	0.418	0.005
Stop-signal interval							
Number of fixations							
Signal-type	1	27	0.42	0.47	24.41	**< 0.001**	0.007
Part	2	54	4.68	5.85	21.62	**< 0.001**	0.076
Signal-location	1	27	0.00	0.34	0.23	0.637	< 0.001
Signal-type by part	2	54	0.01	0.90	0.43	0.655	< 0.001
Signal-type by signal-location	1	27	0.00	0.41	0.05	0.823	< 0.001
Part by location	2	54	0.02	0.83	0.81	0.449	< 0.001
Signal-type by part by signal-location	2	54	0.01	0.90	0.41	0.658	< 0.001
Distance from trained signal							
Signal-type	1	27	6708.96	11436.51	15.84	**< 0.001**	0.051
Part	2	54	51505.45	36436.00	38.17	**< 0.001**	0.293
Signal-location	1	27	123.59	4048.26	0.82	0.372	0.001
Signal-type by part	2	54	13341.47	11946.60	30.15	**< 0.001**	0.097
Signal-type by signal-location	1	27	960.05	10484.49	2.47	0.128	0.008
Part by signal-location	2	54	261.93	13204.44	0.54	0.588	0.002
Signal-type by part by signal-location	2	54	366.33	14373.89	0.69	0.507	0.003

Table 4.10. Overview of Analyses of Variance performed to compare number of fixations and the average distance for the *go-associated* cues. Results discussed in the text are identified in bold. In the word interval, location-type (consistent-left, consistent-right), and part (1-3) were within-subjects factors. In the stop-signal interval, signal-location (congruent location, incongruent location), signal-type (stop-signal, go-signal), and part (1-3) were within-subjects factors.

Expectancy ratings analyses. Expectancy ratings analyses. Analyses of the stop/go expectancy ratings showed that participants expected to stop slightly more for the stop-associated cues (5.13) than for the go-associated cues (4.56). This difference was reliable, $t(31) = -2.45$, $p = 0.020$, $g_{av} = 0.65$. However, it is important to note that the Bayesian analyses were inconclusive ($BF = 2.45$). The stop-minus-go expectancy ratings difference did not reliably correlate with the corresponding RT difference, $r(30) = -0.081$, $p = 0.661$, $BF_{regression} = 0.36$, nor the corresponding *p*(respond|stop) difference, $r(30) = 0.16$, $p = 0.374$, $BF_{regression} = 0.46$.

Analyses of the stop/go location expectancy ratings showed no reliable main effect of location-type (left-associated cues, right-associated cues), $F(1, 31) = 0.80$, $p = 0.377$, *gen.* $\eta^2 = 0.007$. There was also no reliable location-rating difference between the mean ratings between the stop-associated cues (4.67) and the go-associated cues (4.94), $F(1, 31) = 3.28$, $p = 0.080$, *gen.* $\eta^2 = 0.027$. However, the two-way interaction between signal location-type and stop/go-type (stop-associated, go-associated) was marginally significant, $F(1, 31) = 3.66$, $p = 0.065$, *gen.* $\eta^2 = 0.026$. As the interaction was marginally significant, I conducted some follow-up analyses. These showed that the difference between the stop-left cues (4.88) and the stop-right cues (4.48) was marginally significant, $t(31) = 2.00$, $p = 0.055$, $g_{av} = 0.50$, $BF = 1.07$ (one-tailed: $p = 0.973$). Note, however, that if participants acquired the stimulus-signal location contingencies, there should be lower ratings for the stop-left cues than for the stop-right cues. Thus, if anything, the expectancy ratings are in the opposite direction. As expected, there was also no difference between the go-left and go-right cues, $t(31) = -0.57$, $p = 0.574$, $g_{av} = -0.14$, $BF = 0.22$.

Chapter Discussion

In the present chapter, I investigated how learning influences attention in a stop/go task. Across two experiments, I found evidence of stop/go learning on task performance and on expectancy ratings obtained following task completion. Therefore, I conclude that participants learned to stop and/or go in both experiments. These findings replicate the overall pattern of results reported in Chapters 2 and 3. However, despite finding evidence of stimulus-stop/go learning, I found no reliable evidence that participants acquired the associations between the stop-associated and go-associated cues and the spatial location of the stop/go signal.

In Experiment 7, I compared cues that were always paired with the presentation of a stop/go signal in the same spatial location (e.g. always to the left of the cue or always to the right of the cue) with cues that were paired with the presentation a stop/go signal in two spatial locations with equal probability (e.g. sometimes to the left and sometimes to the right of the cue). There were no overall differences in stop or go task performance and no difference in the magnitude of the practice-effect between these cues. However, it is important to highlight that although these differences were not statistically significant, the Bayesian analyses provided only tentative support for the null hypothesis. Thus, I cannot rule out the possibility that some signal location learning occurred in this task.

Therefore, in Experiment 8 I made some modifications to the task design to better allow for the comparison of stimulus-stop/go and stimulus-signal learning during task performance. First, I adapted the task such that participants were no longer required to respond to the stimulus. Thus, I could examine the effects of signal location learning without any potential interference relating to (in)congruency between the signal location and the response hand. Second, I modified the contingencies between specific stimuli and the signal location, such that for each stop/go-type there were trials that were consistent with both the trained response goal (e.g. stop) and consistent with the trained signal location (e.g. left), trials that were consistent with the trained response goal (e.g. stop) but inconsistent with the trained signal location (e.g. right), trials that were inconsistent with the trained response goal (e.g. go) but consistent with the trained signal location (e.g. left), and trials that were inconsistent with both the trained response goal (e.g. go) and with the trained signal location (e.g. right). Consistent with the predictions of the 'automatic inhibition' account, a reliable difference emerged between the stop-associated and go-associated cues in the probability of responding on stop trials. However, despite finding evidence for stop/go learning, I found no reliable difference between task performance on trials in which the signal location was congruent with training and task performance on trials in which the signal location was incongruent with training. Bayesian analyses supported the null hypothesis of no difference between the congruent and incongruent trials in go RTs, in the eye-movement data, and in the expectancy ratings following task completion. Therefore, on the basis of these findings, I conclude that participants did not learn to associate specific stimuli with the spatial signal location of the stop or go signal in this paradigm.

There are at least two plausible explanations for the absence of a stimulus-signal learning effect. First, it is possible that participants learned direct associations between the cues with stopping/going. The automatic inhibition account assumes that during stop learning a direct association between the stop-associated stimuli and the stop goal is acquired (for an extended discussion, see Chapter 2). Thus, when participants acquire direct stimulus-stop associations, the presentation of stop-associated cues primes the stop goal and initiates the stop process. Thus, in situations where participants acquire direct stimulus-stop associations, the need for stop-signal detection is diminished. However, although the direct association idea could account for the pattern of results observed in Experiment 7, this is less clear in Experiment 8. The reasons are discussed in detail in Chapter 7 (see *Can Participants Acquire Indirect (Stimulus-Signal) Associations?*). Second, it is possible that the participants detected the signals on the basis attentional selection influences aside from stimulus-signal learning. For example, it is well-known that attentional control can be driven by bottom-up or exogenous factors, such as items that contrast with surrounding stimuli on the basis of a visual feature value (e.g. brightness) are detected faster (e.g. the 'pop-out' effect; for an overview, see Connor, Egeth, & Yantis, 2004). In addition, Awh and colleagues suggested that attentional control can also be influenced by selection history, such that attentional selection for a given feature (e.g. colour) can be primed (Awh, Belopolsky, et al., 2012). For example, it has been demonstrated that in a search reaction time task containing two equally salient colour singletons, participants could only eliminate the inference from the irrelevant singleton when the selected feature matched that of the previous trial (Theeuwes & Van der Burg, 2011). Thus, it is possible that the influence of learning could not compete with other selection influences in the present task. Future research is required to determine whether stimulus-signal associations would influence performance in situations were the go/stop signals are harder to distinguish, and where the trial sequences are controlled to eliminate trial-by-trial attentional selection priming effects. Third, it is possible that the measures used in the present chapter were not sensitive enough to detect the effects of signal location learning. For example, in future research it could be interesting to increase the number of possible signal locations in a search array as this would increase the demand on visual attention (for example arrays, see Woodman, 2013). The N2-posterior-contralateral (N2pc) event-related potential component could provide a useful alternative measure as this component has been previously used as an index of attention to targets in a search array, and is sensitive to the effects of trial repetition (e.g. Woodman & Luck, 2003). Unlike eye-movements which capture overt attention, the benefit of the N2pc is that it also captures covert attention (i.e. attention to objects without movement of the eyes; see Luck, 2009). However, it is important to highlight that the N2pc components requires a minimum of 250 trials per condition per participant (Woodman, 2010) meaning that future researchers must satisfy this requirement whilst ensuring that the training protocol ensures that it is possible to detect some learning across time.

In Chapter 1, I discussed the idea that when the stop-signal representation remains the same throughout training participants could acquire an association between the stop-associated cues and the stop signal. The retrieval of stimulus-signal associations would prime the representation of the stop signal rather than the stop goal or the stop response (see also Chapter 2; Bowditch, Verbruggen, et al., 2016; Verbruggen, Best, et al., 2014). The retrieval of stimulus-signal associations should result in lower probability of responding for the stop-associated cues than for the go-associated cues but, unlike the retrieval of stimulus-stop associations, there should be no (or very little) difference between the stop-associated and go-associated cues in go RTs. The absence of stimulus-signal location learning seems inconsistent with this account. However, I cannot completely rule out an involvement of stimulus-signal learning in the present chapter, especially as the pattern of behavioural performance in Experiment 8 match the predictions of the stimulus-signal learning account; an effect of learning on the probability of responding, but no difference between go RTs for the stop-associated and go-associated cues. There are two possible explanations for these results. First, it is possible that participants acquired associations between the stop-associated (and/or go-associated) cues and representation of the signal that did not include spatial signal location. For example, the stop and go signals were always presented in the same colour in Experiment 7 (e.g. yellow = stop; blue = go, or vice versa) and always in the same shape in Experiment 8 (e.g. diamond = stop; square = go, or vice versa), regardless of the signal location contingencies. Indeed, several studies have provided evidence in support for independent attentional selection of colour and spatial signal location (e.g. Andersen, Muller, & Hillyard, 2009; Kasten & Navon, 2008). Thus, it is possible that learning primed features of the signal representation, independent of spatial signal location. Research in the selective attention literature suggests that, in situations where the spatial signal location is uncertain, colour-based attentional selection is more efficient (Adams & Chambers, 2012). Future research should investigate whether participants can acquire associations between stop- or go-associated cues and other features of the signal representation, such as colour. Second, it is possible that signal detection processes are less important during stop learning than originally thought and that differences in the effects of learning on the probability of responding and go RTs arise due to differences in the sensitivities of these measures. This idea is discussed in detail in Chapter 7 (see *What is Learned?*, Chapter 7). If this is the case, we do not have to assume that participants learned two qualitatively different stimulus-stop and stimulus-signal associations when learning influences stop performance but not go performance.

Finally, at the beginning of this Chapter, I discussed a recent framework that proposed that action control depends on three basic cognitive processes; signal detection, action selection, and action execution that are influenced by development, rule activation and maintenance, monitoring, proactive control and, importantly for the present purposes, associative learning (Verbruggen, McLaren, et al., 2014). If it turns out that attention to the stop signal is not influenced by learning in response inhibition tasks, this would suggest that the influence of associative learning on signal detection processes may be less than previously anticipated. The implications of these findings are discussed further in Chapter 7.

Conclusion

Taken together, the experiments presented in this chapter provide further evidence that participants can acquire item-specific stop/go associations. However, across two experiments, I find no evidence to support the idea that participants acquired the associations between specific stimuli and the spatial location of the stop or go signal. Thus, these findings suggest that, for spatial location features held within representations of the stop signal, the acquisition of stimulus-signal associations is not an inevitable consequence of performing a stop learning task where the representation of the stop signal remains the same throughout training.

Appendix H

Words used in Experiment 7

We used 112 words (56, five-letter, 56 six-letter) in this experiment. Half of the words referred to natural objects, the other half referred to human-made objects. The natural and human-made words were matched on SUBTLEX frequency (*M*= 2.39, *SD* = 0.51).

Natural: sheep, eagle, otter, camel, hippo, grape, shark, snail, raven, chimp, horse, moose, llama, squid, chick, onion, lemon, panda, maple, snake, finch, sloth, apple, acorn, peach, cedar, basil, rhino, monkey, turtle, lizard, peanut, iguana, banana, papaya, orchid, turkey, tomato, donkey, walnut, potato, baboon, carrot, rabbit, falcon, salmon, weasel, pepper, ferret, shrimp, walrus, pigeon, celery, garlic, gibbon, radish.

Human-made: brick, scarf, canoe, cabin, diary, label, whisk, ferry, quill, leash, apron, spoon, shawl, bench, banjo, skirt, penny, torch, arrow, swing, medal, couch, piano, brush, wagon, towel, buggy, badge, brooch, crayon, cement, teacup, cradle, pajama, sandal, seesaw, jigsaw, napkin, paddle, helmet, funnel, statue, sleigh, eraser, jumper, stereo, anchor, cinema, poster, violin, candle, pillow, dinghy, bucket, guitar, corset.

Appendix I

Words used in Experiment 8

There were 36 words in this experiment. Words were randomly allocated to three word lists of 12 words each anew for each participant. Two of the word lists were used in the main task, the remaining word list was used in the recognition ratings following task completion (For more details, see Procedure). Participants were not required to respond to the words in this experiment, so all words referred to natural objects. The words were taken from those used in Experiments 5-6 (Chapter 3).

Words: mule, calf, kiwi, wren, swan, newt, leaf, stag, lime, clam, crow, flea, pear, wasp, moth, rice, plum, germ, slug, dove, toad, boar, crab, pony, deer, worm, lamb, goat, frog, hawk, tree, lion, wolf, duck, bull, bear.

CHAPTER 5

BEHAVIOURAL PRACTICE-EFFECTS IN INHIBITION AND NON-INHIBITION TASKS

It was recently suggested that one of the biggest problems in the cognitive control literature is that researchers often fall into the trap of referring to general constructs, such as 'inhibition' without specifying what the underlying mechanisms are (Verbruggen, McLaren, et al., 2014). Consistent with this suggestion, the effects of training in response inhibition tasks are typically explained in terms of strengthening or improving a general 'inhibitory control' mechanism rather than identifying how training influences the processes underlying action control (see, e.g. Spierer, Chavan, & Manuel, 2013). Referring to general constructs, such as 'inhibitory control', is problematic because it does not explain *how* training influences performance. To address this issue, Verbruggen and colleagues proposed that various forms of action control depend on three basic cognitive processes: signal detection, action selection, and action execution (Verbruggen, McLaren, et al., 2014). Indeed, these processes can account for performance in both inhibition and non-inhibition tasks (e.g. Logan & Bundesen, 2003; Monsell & Mizon, 2006; Ratcliff & Smith, 2004; Schall, 2001; Smith & Ratcliff, 2004). Furthermore, each of these processes can be modulated by learning. In the present chapter, I address the third outstanding issue highlighted in Chapter 1; whether the effects of learning on action selection processes are specific to inhibition tasks or whether they represent task-general (non-inhibition, Chapter 1) effects. In particular, I compare the effects of training on behavioural task performance in the go/no-go task with the effects of training on behavioural task performance in the two-choice task. This comparison could further our understanding of what is learned in response inhibition tasks and how training influences behaviour.

The comparison of the go/no-go task with the two-choice task has been the subject of investigation since the original work conducted by Donders (1868/1969); the typical finding is that go response latencies are shorter and accuracy is greater in the go/no-go task than in the two-choice task (Gomez, Ratcliffe & Perea, 2007). This difference was initially attributed to the presence of a response-selection stage in the two-choice task but not in the go/no-go task[36]. However, more recent work suggests that the action selection demands in these tasks are more similar than previously thought. For example, neuroimaging research has shown that the pre-supplementary motor area (a brain region shown to be critically involved in action selection in choice tasks) is also activated on go and no-go trials in the go/no-go task (Simmonds, Pekar, & Mostofsky, 2008; Mostofsky & Simmonds, 2008). As outlined in Chapter 1, the pre-SMA is a key component of the 'inhibitory control network' involved in updating current action plans and monitoring for response conflict (e.g. Ridderinkhof, Ullsperger, et al., 2004). Similarly, performance of inhibition and non-inhibition tasks has been shown to elicit similar electrophysiological signatures thought to reflect medial prefrontal cortex activity (e.g. the N2 component; Smith, Smith, Provost, & Heathcote, 2009; for an overview, see Chapter 6). Thus, these findings suggest that performance of inhibition and non-inhibition tasks could recruit similar brain regions. Indeed, it has even been suggested that response execution and response inhibition are similar at a neural level (Jasinska, 2013; Mostofsky and Simmonds, 2008). Combined, these findings suggest that response selection/execution and response inhibition are most likely 'two sides of the same coin' that draw on common neural substrates (Mostofsky and Simmonds, 2008).

Consistent with this conclusion, research using diffusion models also suggest that there are similarities between response selection and response inhibition. According to these models, response selection can be explained in terms of the accumulation of activation towards a response threshold; the response that reaches the threshold first is selected and then executed (e.g. Ratcliff & Smith, 2004; Smith & Ratcliff, 2004; for an overview, see Ratcliff, Smith, Brown, & McKoon, 2016). Although these models have primarily been used to account for response selection processes, more recent modelling work shows that these models can also account for response inhibition. For example, a diffusion model with response boundaries for the go and no-go responses provided a better fit for the go/no-go task performance data than a model with no response boundary for the no-go response (Gomez, Ratcliff, et al., 2007; see also Logan, Van Zandt, et al., 2014, for a diffusion version of the independent race model). This modelling work accommodates the shorter RTs and lower accuracy on go trials in the go/no-go task (compared with trials in the two-choice task) as a bias in the starting point of the go activation (Gomez, Ratcliff, et al., 2007; see also Ratcliff, Smith, et al., 2016). This

---

[36] Donders (1868/1969) originally argued that performance in the two-choice task requires the insertion of an action selection stage of processing (e.g. left key press vs. right key press) in addition to the stimulus-categorisation stage (e.g. diamond vs. square) required in both the two-choice and go/no-go tasks (for an overview, see Ulrich, Mattes & Miller, 1999). Indeed, there is some evidence that suggests that an additional response selection process is inserted in the two-choice task that is not present in the go/no-go task (Vidal, Burle, Grapperon & Hasbroucq, 2011). However, this has only been observed using a go/no-go task in which participants had to decide between two responses ion go trials in addition to the no-go response (e.g. digit 3 = 'right key press'; digit 4 = 'left key press'; digits 7 and 8 = no-go). This could indicate that there may be some hierarchy of response selection (e.g. digit -> go or no-go -> [if go: left keypress or right keypress]). Note, however, that in tasks similar to those used in the present chapter, the weight of evidence is in favour of the view that the action selection processes required in the go/no-go and two-choice tasks are more similar than different.

suggests that, just as in the two-choice task, performance of the go/no-go task is best considered as a choice between alternative responses (e.g., Band, Ridderinkhof, & van der Molen, 2003; Donkers & van Boxtel, 2004; Gomez, Ratcliff, et al., 2007). Logan, Van Zandt, et al. (2014) also argued that go responses and stop responses in a stop-signal paradigm differ primarily in the conditions that trigger them and the outcome rather than in selection mechanisms; in other words, they differ in content rather than form. Combined, these studies indicate that similar selection mechanisms are likely to be involved in response-inhibition and choice tasks[37].

To date, some indirect evidence exists suggesting that learning in these tasks may recruit similar mechanisms. For example, practice-induced improvements in both tasks have been attributed to the retrieval of associations between a given stimulus and a given (motor or inhibitory) response. Memory-retrieval accounts (such as the Instance theory; Logan,1988; for an overview, see *Modelling Performance in Response Inhibition Tasks*, Chapter 1) can account for training-effects in both inhibition and non-inhibition tasks. For example, previous research has demonstrated that, just as in the go/no-go task, there is a performance cost in choice tasks when acquired stimulus-response (S-R) associations are reversed (Horner & Henson, 2009; 2011; see also Race, Shanker & Wagner, 2009). Henson and Horner (2009; 2011) attributed this cost to the automatic retrieval of instances containing information at various processing stages, including response selection (cf. Verbruggen, Best, et al., 2014). Response-congruency effects that can be attributed to the retrieval of stimulus-response associations have also been observed in several other (non-inhibition) tasks, including task-switching (e.g. Wylie & Allport, 2000; Koch & Allport, 2006; Waszak, Hommel, & Allport 2003; Schneider, 2015; Waszak, Pfister, & Kiesel, 2013; Wendt & Kiesel, 2008) and interference control (e.g. Cohen-Kdoshay & Meiran, 2009; Wenke, De Houwer, De Winne & Liefooghe, 2014) paradigms. For example, in some task-switching designs, conflict between S-R associations can result in slower responding, increased errors, and a larger task-switch cost[38] for incongruent stimuli (that are associated with a different response in each task) than for congruent stimuli (that are associated with the same response in each task; e.g., Sudevan & Taylor, 1987; Rogers & Monsell, 1995). Thus, there is good evidence to suggest that in various tasks (including the two-choice task) practice can lead to delayed responding when there is a conflict with established S-R associations. Thus, it seems that practice-related effects on response selection are central to performance of a range of tasks.

On the basis of these findings, it seems plausible that action control processes are modulated by task practice in a similar fashion in both non-inhibition and inhibition tasks. However, the effects

---

[37] It is important to note, however, that some recent studies indicate that there might be differences as well (Bender, Filmer, Garner, Naughtin, & Dux, 2016; Rae, Hughes, Weaver, Anderson, & Rowe, 2014). For example, Bender, Filmer, et al. (2016) found in a correlational study that response-inhibition tasks and response-selection tasks loaded on different factors. Thus, there might be both unity and diversity in these tasks.

[38] The decrement in performance (longer RTs and higher errors rates) when the task changes compared to when the task stays the same (Monsell, 2003).

of training in inhibition and non-inhibition tasks have not been compared in the same experimental design meaning that claims about the specificity of learning in response inhibition tasks are currently based on cross-experimental comparisons. This is problematic because similarities (and indeed differences) between these tasks could be masked by the various inconsistencies between these experiments. Therefore, the first step in investigating whether response inhibition training influences a general action selection mechanism is to compare the go/no-go and two-choice tasks using the same experimental design.

The Present Chapter

The aim of the present chapter was to provide an initial investigation into the specificity behavioural effects of task practice on action selection processes in the go/no-go task. For the purposes of the research presented in this thesis, the present chapter also served as an exploratory behavioural pilot for the event-related potential experiments reported in Chapter 6. I used a modified version of the task used by Verbruggen and Logan (2008a; Experiment 2) in their original demonstration of automaticity in the go/no-go task. This task had a training (acquisition) phase, followed by a test phase where I introduced a new judgment task. Half of all stimuli were associated with different responses in the training and test phases (incongruent stimuli) whereas the other half of stimuli were associated with the same response in the training and test phases (congruent stimuli). The introduction of a new judgment task in the test phase meant that performance could no longer benefit from the retrieval of the acquired stimulus-category associations. The novel contribution of Experiments 9-10 was that I also included a two-choice task condition which was identical to the go/no-go task except that participants had to make a keypress response on every trial. Participants were randomly allocated to the two-choice task condition or to the go/no-go task condition (between-subjects).

For the purposes of the present study, mediation of the stimulus-response or stimulus-stop (no-go) associations via a category representation was undesirable because associations between categories and responding and/or categories and stopping could differentially contribute to the effects of training on performance in the go/no-go and two-choice tasks. Therefore, I made two main modifications to the original Verbruggen and Logan (2008a) task design to minimise the contribution of category-related costs; I used a smaller stimulus set (four chequerboards; see Figure 5.1) than in the Verbruggen and Logan task (60 words). Research suggests that participants are more likely to rely on stimulus-response associations when the stimulus-set is smaller (e.g. Forrest, Monsell, & McLaren, 2014). Furthermore, participants were required to perform a visual judgment task in Experiments 9-10 rather than a word semantic judgment (Verbruggen & Logan, 2008a, Experiment 2).

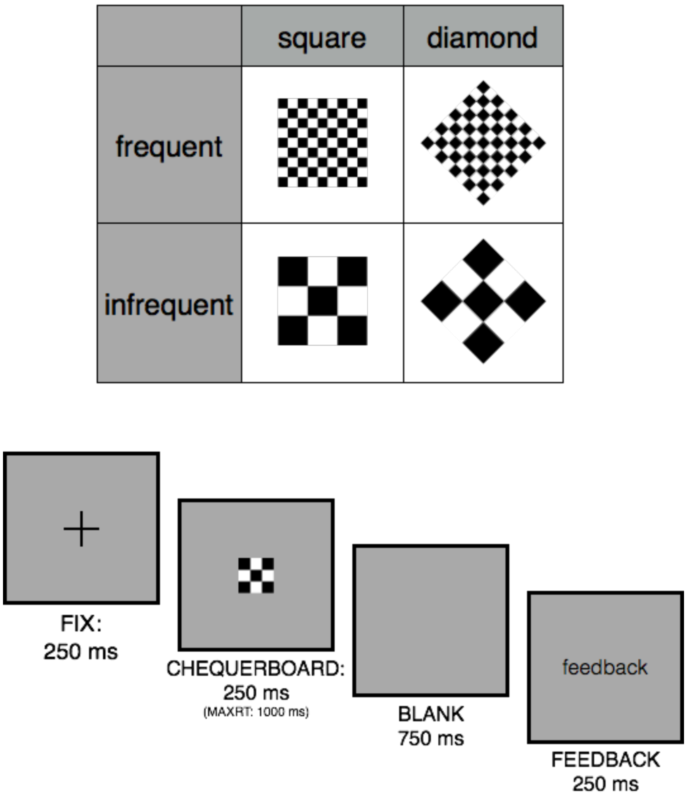

Figure 5.1. Stimuli and trial structure. Each trial began with the presentation of a fixation cross in the centre of the screen for 250 ms. Following the removal of the fixation cross, a chequerboard was presented in the centre of the screen. After 250 ms, the chequerboard was replaced with a blank screen. The maximum reaction time (MAXRT) was 1000 ms. Feedback (‘correct’, ‘incorrect’ or ‘too slow’ in case participants did not response before the RT threshold in Experiment 9; ‘too slow’ or a blank screen in Experiment 10) was presented for 250 ms. Note: frequent/infrequent refers to the spatial frequency of the chequerboards not stimulus presentation frequency.

Experiment 9

Method

Participants. Sixty-four students from the University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 19.11 years, *SD* = 1.25, 48 females, 58 right-handed). Six participants were removed and replaced due to technical problems with the stimulus presentation computers; no participants were removed for performance-related reasons. All experiments of the present chapter were approved by the local research ethics committee at the School of Psychology, University of Exeter. Written informed consent was obtained after the nature and possible consequences of the study were explained. The target sample size and exclusion criteria were decided in advance of data collection to ensure that I had enough power to detect medium-sized effects in my within-subject comparisons and within-between-subject interactions.

Apparatus, stimuli, and procedure. The stimuli were presented on 19-inch monitors (screen size: 1280 × 1024 pixels) using Psychtoolbox (Brainard, 1997) in a multiple testing environment. The stimuli consisted of four chequerboards (size: 140 × 140 pixels; Figure 5.1), which varied along two dimensions: frequency (3 × 3 or 9 × 9) and orientation (0° or 45°). All stimuli were presented on a grey background (RGB: 169 169 169).

Participants were randomly assigned to the go/no-go task condition or to the two-choice task condition (between-subjects). All participants started with a training phase, which consisted of twelve blocks, followed by a test phase, which consisted of two blocks. Each block consisted of 80 trials and each stimulus was presented 20 times per block. Half the participants decided whether the chequerboards were high or low frequency in the training phase and decided whether the chequerboards were orientated as diamond or square in the test phase. For the other half, the order was reversed. Participants received instructions on the screen for the test phase after they finished the training phase and were required to press any key to confirm that they understood the new judgment task or to alert the experimenter if they had any questions. In the test phase, half of the stimuli were congruent (the same response was required in the training and test phases) and half of the stimuli were incongruent (e.g. no-go stimulus in the training phase and go stimulus in the test phase in the go/no-go task; respond left in the training phase and respond right in the test phase in the two-choice task).

On each trial, a fixation cross (30 × 30 pixels) was presented in the middle of the screen for 250 ms. The chequerboard appeared in the centre of the screen (replacing the fixation cross). After 250 ms, the chequerboard was replaced with a blank screen. In the go/no-go task condition, half of the participants had to press the ‘c’ key (with their left index finger) when the go stimulus was presented; the other half had to press the ‘m’ key (with their right index finger) when the go stimulus was presented. All participants in this condition had to refrain from pressing any key when the no-go stimulus was presented. In the two-choice task condition, participants had to press the ‘c’ key (with the left index finger) for one set of chequerboards (e.g high spatial frequency) and the ‘m’ key (with the right index finger) for the other set of chequerboards (e.g. low spatial frequency; Figure 5.1). All mappings were counterbalanced. The response deadline was 1000 ms post-chequerboard presentation.

Fast responding was encouraged via a (covert) response threshold that was adaptively tracked throughout task performance. The response threshold was initially set at 500 ms and was continually adjusted using a four-up/one-down tracking procedure: the response threshold decreased by 50 ms following four consecutive correct responses executed within the response threshold but increased by 50 ms for all other responses (regardless of whether they were correct or incorrect). The minimum response threshold value was 50 ms and the maximum response threshold value was 950 ms (MAXRT - 50 ms). In the go/no-go task, the response threshold tracking procedure was based on responding on go trials only. Separate tracking procedures were implemented for the left and right responses in the two-choice task. All response threshold values were reset to 500 ms at the beginning of the test phase. After the response deadline elapsed (1000 ms), feedback was presented ('correct', 'incorrect', or 'too slow' in case the participant did not respond before the response threshold), which remained on the screen for 250 ms. The feedback and threshold procedure were designed to encourage fast and accurate responding. The next trial started immediately after the feedback.

There was a break after every block. During the break, feedback on mean RTs (on go trials in the go/no-go task condition; on left and right response trials in the two-choice task condition), the number of errors (two-choice task condition only), the number of missed responses (on go trials in the go/no-go task and on all trials in the two-choice task condition), and the percentage of responses on no-go trials (go/no-go task condition only). After 15 s had elapsed, participants had to press a key to start the next block.

Analyses. All data processing and analyses were completed using R (R Development Core Team, 2014). All data files and R scripts are deposited in Dropbox (https://www.dropbox.com/sh/1bj5jb9xfcabu0j/AADw7XgnC4X-Su1JkMe-U_zRa?dl=0). I performed ANOVAs on correct RTs[39] on go trials and the probability of responding on no-go trials [*p*(respond|no-go)] in the go/no-go task; and on correct RTs and the probability of correct responses on all trials in the two-choice task. Note, choice accuracy refers to the probability that participants selected the correct response between two alternatives (left key press/right key press); consequently it was not possible to perform an analysis with 'task' as a between-subjects factor as participants responded using only one response key on go trials in the go/no-go task. Similarly, it was not possible to include 'task' as a between-subjects factor in the *p*(respond|no-go) analyses because participants responded on all trials in the two-choice task (i.e. there were zero no-go trials). Missed responses were rare (2% in Experiment 9; 1% in Experiment 10) so were not analysed further. RTs ≤ 1 ms were removed prior to analyses (0.003% of all trials) as these indicate that participants held down the response key between trials. Where appropriate, I applied the Huyhn-Feldt correction for violations of sphericity. For pairwise comparisons, Hedge's g_{av} is the reported effect size measure (Lakens, 2013). For the test phase analyses, I also calculated Bayes factors for all main effects and interaction contrasts in the ANOVA designs (Rouder, Morey, et al., 2012). Bayes factors compare the likelihood of the data under the null hypothesis of no difference between groups against the alternative hypothesis of a difference greater than 0. Bayes factors were calculated with the BayesFactor package in R, using the default prior (0.707; Morey, Rouder, et al., 2015).

---

[39] Note, correct RTs were analysed regardless of whether the response was executed before or after the go response threshold (as adjusted via the tracking procedure), so long as the response was executed before the response deadline (1000 ms following chequerboard presentation).

Results and Discussion

Training phase. Training phase. All stimuli were consistently associated with only one response during the training phase so I did not expect any differences between the congruent and incongruent items. Initial analyses (not reported) confirmed that there was no main effect of congruency in the training phase (all *p*s ≥ 0.290) so I did not analyse this further.

Analyses of correct go RTs revealed that the main effect of block was significant ($p < 0.001$; Table 5.1) indicating that mean RTs decreased as a function of task practice. Follow-up analyses showed that the main effect of block was reliable in the go/no-go task ($p < 0.001$; Table 5.1) and in the two-choice task ($p = 0.017$; Table 5.1)[40]. As expected, responding was reliably slower in the two-choice task (386 ms) than in the go/no-go task (349 ms; $p < 0.001$; Table 5.1). The two-way interaction between task and block was also reliable ($p = 0.009$; Table 5.1) reflecting the steeper decrease in RTs with practice in the go/no-go task than in the two-choice task (Figure 5.2).

In addition to the analyses of RTs, I also conducted two additional analyses on choice accuracy in the two-choice task and on the probability of responding on no-go trials in the go/no-go task. Analyses of the probability of correct responses in the two-choice task revealed a main effect of block ($p < 0.001$; Table 5.1) indicating that choice accuracy increased as a function of practice. Analyses of the probability of responding on no-go trials in the go/no-go task revealed that the main effect of block was not reliable ($p = 0.179$; Table 5.1). The absence of a learning effect in the probability of responding is probably due to a floor effect as values were already low in the first block and reached asymptote from block two onwards (Figure 5.4).

Taken together, the training phase data provides strong evidence for learning in both tasks (as expected): as can be seen in Figures 5.2 and 5.3, the pattern of responding in the RT and choice accuracy data shows that task performance benefited from practice.

---

[40] Planned comparisons confirmed that responding at the beginning of training (blocks 1-2) were reliably slower than responding at the end of training (blocks 11-12) in the two-choice task, $t(31) = 1.83$, $p = 0.038$, $g_{av} = 0.24$ (one-tailed), and in the go/no-go task, $t(31) = 4.35$, $p < 0.001$, $g_{av} = 0.67$ (one-tailed).

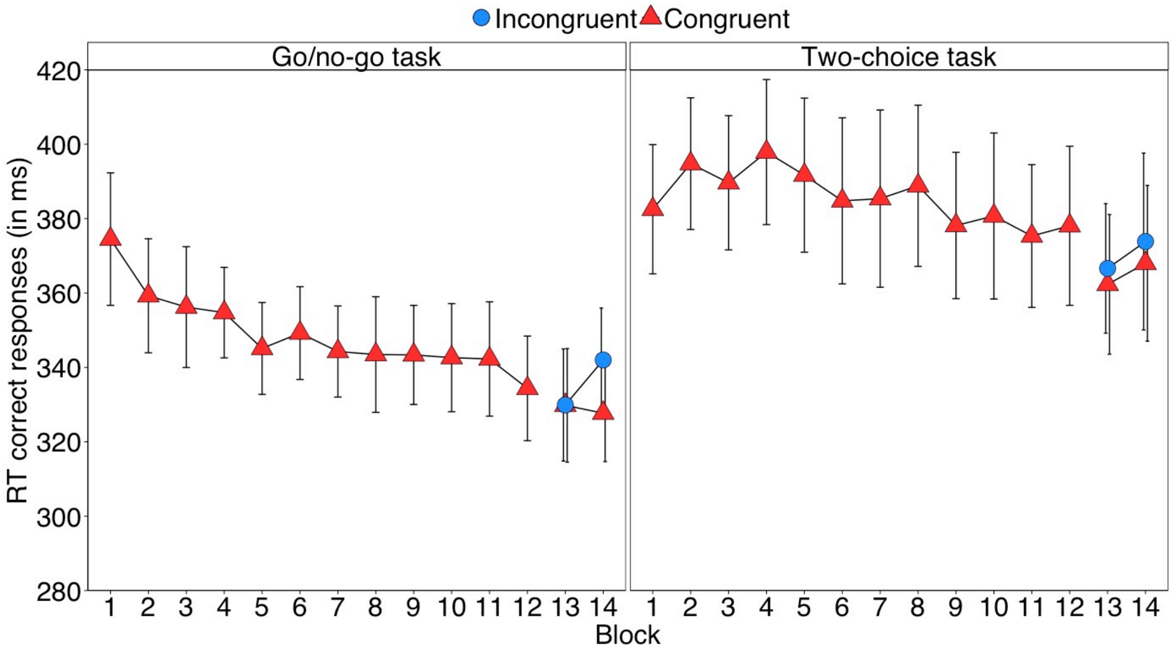

Figure 5.2. Reaction times of correct responses on go trials in the go/no-go task (left panel) and on all trials in the two-choice task (right panel) in Experiment 9. Blocks 1-12 were the training phase; blocks 13-14 were the test phase. Error bars are 95% confidence intervals.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Combined task analysis							
Reaction times							
Task	1	62	256275.53	1434674.6	11.08	**0.001**	0.126
Block	11	682	39128.03	349635.1	6.94	**< 0.001**	0.021
Task by block	11	682	18204.83	349635.1	3.24	**0.009**	0.010
Go/no-go task only							
Reaction times: go trials							
Block	11	341	39598.56	151812.80	8.09	**< 0.001**	0.063
p(respond|no-go): no-go trials							
Block	11	341	0.06	1.19	1.45	0.180	0.025
Two-choice task only							
Reaction times: all trials							
Block	11	341	17550.56	197049.60	2.76	**0.018**	0.014
Choice accuracy: all trials							
Block	11	341	0.12	0.76	4.85	**< 0.001**	0.056

Table 5.1. Overview of the Analyses of Variance on the *training phase* data (blocks 1-12) in Experiment 9 with block as a within-subjects factor and task as a between-subjects factor. The separate analyses for the go/no-go task and the two-choice task are also presented. *p*s < 0.05 are highlighted in bold.

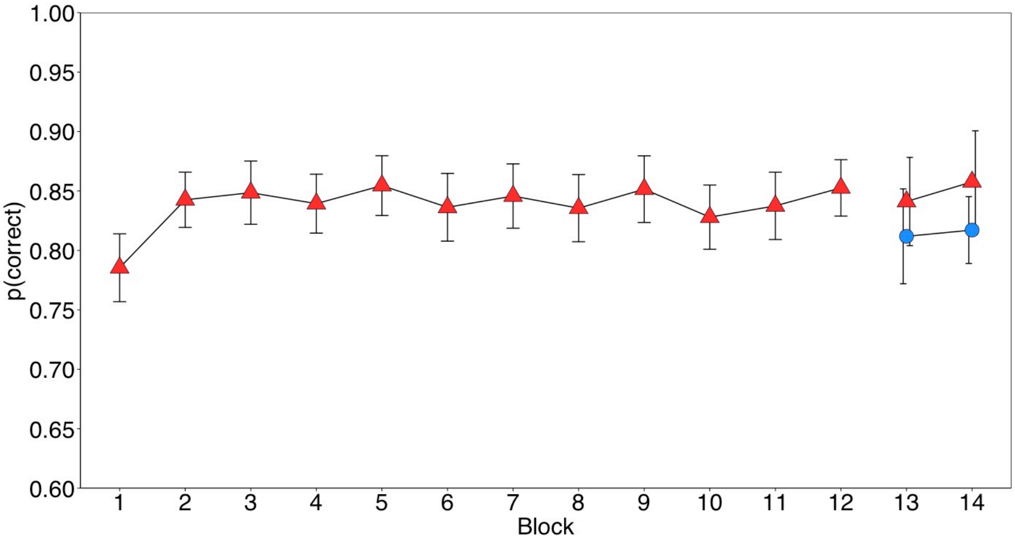

Figure 5.3. The probability of correct responses in the two-choice task in Experiment 9. Blocks 1-12 were the training phase; blocks 13-14 were the test phase. Error bars are 95% confidence intervals.

Test phase. As can be seen in Figure 5.2, RTs were slightly longer for the incongruent stimuli than for the congruent stimuli in the two-choice task (incongruent: 370 ms, congruent: 365 ms) and in the go/no-go task (incongruent: 336 ms, congruent: 329 ms). Analyses of the RTs revealed that the main effect of congruency was reliable ($p = 0.007$; Table 5.2). The main effect of task was reliable ($p < 0.001$; Table 5.2) but the two-way interaction between task and congruency was not reliable ($p = 0.632$; Table 5.2). Table 5.3 shows the outcome of the Bayesian analysis. As can be seen, the results are largely consistent with the ANOVAs reported (Table 5.2). Table 5.3 shows that dropping task, congruency and block had a deleterious effect on the model. All other factors or interactions could be dropped (Table 5.3 also shows the separate models for the two-choice and go/no-go tasks). However, the separate model for the two-choice task showed that the congruency factor could be dropped without impairing the model fit so I remain cautious about interpreting the numerical difference observed in the two-choice task as (entirely) due to the retrieval of the acquired stimulus-response associations in this task.

Analyses of choice accuracy in the two-choice task revealed a reliable main effect of congruency ($p = 0.008$; Table 5.2; see also Figure 5.3) reflecting lower accuracy for the incongruent stimuli than for the congruent stimuli. The decreased accuracy for the incongruent stimuli than for the congruent stimuli is consistent with the Instance theory (Logan, 1988; Logan, 2002): following the presentation of the incongruent stimuli, the previously encoded instances are automatically retrieved and when these instances win the race against the algorithmic processes, the incorrect response is executed. Consistent with the absence of a block effect in the training phase, analyses of the p(respond|no-go) data revealed that the main effect of congruency was not reliable ($p = 0.694$; Table

5.2; Figure 5.4) and Bayesian analyses showed that the congruency factor could be dropped without impairing the model fit (Table 5.3).

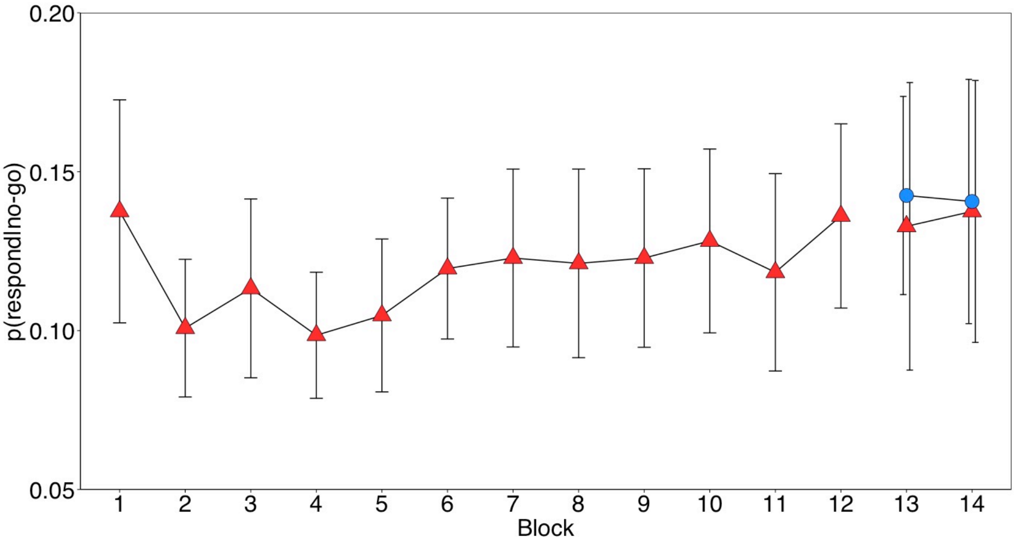

Figure 5.4. The probability of responding on no-go trials in the go/no-go task in Experiment 9. Blocks 1-12 were the training phase; blocks 13-14 were the test phase. Error bars are 95% confidence intervals.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Combined task analysis							
Reaction times							
Task	1	62	79949.22	527580.92	9.40	**0.003**	0.119
Congruency	1	62	2409.17	19476.63	7.67	**0.007**	0.004
Block	1	62	2094.19	31320.01	4.15	**0.046**	0.004
Task by congruency	1	62	72.96	19476.63	0.23	0.632	< 0.001
Task by block	1	62	33.56	31320.01	0.07	0.797	< 0.001
Congruency by block	1	62	994.47	14086.41	4.38	**0.041**	0.002
Task by congruency by block	1	62	638.58	14086.41	2.81	0.099	0.001
Go/no-go task only							
Reaction times: go trials							
Congruency	1	31	1660.31	12355.65	4.17	0.050	0.008
Block	1	31	798.78	8781.62	2.82	0.103	0.004
Congruency by block	1	31	1613.42	6436.08	7.77	**0.009**	0.008
p(respond|no-go): no-go trials							
Congruency	1	31	0.00	0.18	0.01	0.694	0.001
Block	1	31	0.00	0.26	0.16	0.919	< 0.001
Congruency by block	1	31	0.00	0.20	0.05	0.820	< 0.001
Two-choice task only							
Reaction times: all trials							
Congruency	1	31	821.82	7120.98	3.58	0.068	0.002
Block	1	31	1328.96	22538.39	1.83	0.186	0.003
Congruency by block	1	31	19.63	7650.33	0.08	0.780	< 0.001
Choice accuracy: all trials							
Congruency	1	31	0.04	0.15	7.92	**0.008**	0.028
Block	1	31	0.00	0.21	0.56	0.458	0.003
Congruency by block	1	31	0.00	0.14	0.22	0.642	0.001

Table 5.2. Overview of the Analyses of Variance on the *test phase* data (blocks 13-14) in Experiment 9 with congruency and block as within-subjects factors and task as a between-subjects factor. The separate analyses for the go/no-go task and the two-choice task are also presented. *p*s < 0.05 are highlighted in bold.

Omitted Factor(s)	Bayes Factor	Confidence Interval
Combined task analysis		
Reaction times		
Task	0.10	±12.74%
Congruency	0.27	±12.28%
Block	0.39	±11.20%
Block by congruency	1.47	±12.20%
Block by task:by congruency	1.74	±11.70%
Task by congruency	4.58	±11.40%
Block by task	5.37	±11.64%
Go/no-go task only		
Reaction times: go trials		
Block by congruency	0.40	±4.31%
Congruency	0.41	±4.27%
Block	1.55	±4.16%
p(respond|no-go): no-go trials		
Block by congruency	3.95	±3.94%
Congruency	5.03	±3.91%
Block	5.49	±3.93%
Two-choice task only		
Reaction times: all trials		
Block	1.13	±5.37%
Congruency	1.98	±5.42%
Block by congruency	3.76	±5.31%
Choice accuracy: all trials		
Congruency	0.20	±3.68%
Block by congruency	3.59	±3.58%
Block	3.79	±3.57%

Table 5.3. Bayesian analysis for the *test phase* data (blocks 13-14) in Experiment 9. Bayes factors < 1 indicate that the removal of the factor or interaction had a deleterious effect on the model whereas Bayes factors > 1 indicate that the factor or interaction could be removed without substantially impairing the model fit. Note that 'participant' was included as a factor for all models, but this factor is not added to the model descriptions in the tables to reduce the amount of text.

Training vs. test phase. In addition to the main analyses, I also performed two extra analyses to compare performance between the training phase and the test phase. First, I performed an analysis to explore whether the observed RT difference between the incongruent and congruent stimuli in the go/no-go task was due to the interference between the acquired stimulus-stop (no-go) associations and go processing (rather than the absence of go learning in the training phase). To this end, I compared RTs in the go/no-go task in the first two blocks in the training phase (blocks 1-2) with RTs in the two blocks of the test phase (blocks 13-14). If the difference in RT between the incongruent and congruent stimuli reflects the automatic retrieval of the no-go response associated to the incongruent stimuli during training, responding should speed up more in the training phase (blocks 1-2) than in the test phase (blocks 13-14) for the incongruent stimuli than for the congruent stimuli[41]. Analyses showed that the three-way interaction between congruency, phase and block was marginally significant, $F(1, 31) = 3.84$, $p = 0.059$, *gen.* $\eta^2 = 0.007$. As this interaction was marginally significant, I conducted some follow-up analyses. These follow-up comparisons revealed that the two-way interaction between block and phase was reliable for the incongruent stimuli, $F(1, 31) = 13.66$, $p < 0.001$, *gen.* $\eta^2 = 0.077$: responding became faster between the first and second blocks in the training phase ($M_{\text{diff}} = 17$ ms) but responding became slower between the first and second blocks of the test phase ($M_{\text{diff}} = -12$ ms). This interaction did not reach significance for the congruent stimuli, $F(1, 31) = 3.24$, $p = 0.081$, *gen.* $\eta^2 = 0.011$, but the numerical difference between the first and second blocks in the training phase ($M_{\text{diff}} = 13$ ms) was greater than in the first and second blocks in the test phase ($M_{\text{diff}} = 3$ ms). This suggests that the response slowing induced by the retrieval of trained stimulus-stop (no-go) associations in addition to a more general slowing induced by the retrieval of trained associations in a novel task context.

Second, I compared overall RTs at the end of the training phase (blocks 11-12) with RTs in the test phase (blocks 13-14) to investigate whether the introduction of the novel judgment task induced a general decrement in performance. These analyses revealed that there were no costs induced by the introduction of the novel judgment task. If anything, responding was faster in the test phase ($M = 350$ ms, $SD = 3$ ms) than in the final blocks of the training phase ($M = 358$ ms, $SD = 5$ ms): this difference was marginally significant, $F(1, 62) = 3.99$, $p = 0.050$, *gen.* $\eta^2 = 0.006$. The two-way interaction between phase (end of training vs. test phase) and task was not reliable, $F(1, 62) = 0.19$, $p = 0.667$, *gen.* $\eta^2 < 0.001$. Furthermore, there was no reliable difference in performance between the end of training and the test phase in choice accuracy in the two-choice task, $F(1, 31) = 1.60$, $p = 0.215$, *gen.* $\eta^2 = 0.008$, nor in the p(respond|no-go) in the go/no-go task, $F(1, 31) = 0.57$, $p = 0.457$, *gen.* $\eta^2 = 0.005$. Table 5.4 shows the outcome of the Bayesian analyses. As can be seen, the results are broadly consistent with the ANOVAs reported. Thus, it is possible that the effects of reversing the acquired mappings were short-lived and participants rapidly acquired the new mappings. Taken together, this suggests that any prior stimulus-category learning during the training phase did not substantially impair performance in the test phase when the previous categories were no longer relevant. The role of stimulus-category learning is discussed in detail in Chapter 7 (see *Mediation via a Category Representation,* Chapter 7).

---

[41] Note that this analysis was originally conducted by Verbruggen and Logan (2008a).

Omitted Factor(s)	**Bayes Factor**	**Confidence interval**
Combined task analysis		
Reaction times		
Task	0.05	±5.91%
Phase	0.08	±5.78%
Task by phase	4.88	±6.68%
Factor	**Bayes Factor**	**Confidence interval**
Go/no-go task only		
Reaction times: go trials		
Phase	0.95	±2.41%
p(respond|no-go): no-go trials		
Phase	0.27	±1.94%
Two-choice task only		
Reaction times: all trials		
Phase	2.76	±2.59%
Choice accuracy: all trials		
Phase	0.35	±2.07%

Table 5.4. Bayesian analyses for the difference between performance at the end of training vs. in the test phase in Experiment 9. For details, see Table 5.3. For the separate analyses of the go/no-go task and the two-choice task, 'phase' was the only factor. Therefore, for these analyses, the Bayes factors of the model including the phase factor was computed against the intercept-only model. As can be seen, all Bayes Factors (except for RTs in the two-choice only model) are < 1 indicating that the model including the phase factor provided a worse fit for the data than the intercept-only model. Note, the Bayes Factor for the RTs in the two-choice only model reflects the speeding (not slowing) of responding in the test phase than in the final blocks of training, providing further evidence against the idea that the introduction of the novel judgment task in the test phase impaired performance.

Experiment 10

In Experiment 9, it was found that performance improved with task practice in both the go/no-go and two-choice tasks during the training phase. In the test phase, I found some reliable differences between the congruent and incongruent stimuli in go response latencies in the go/no-go task and in choice accuracy in the two-choice task. However, these differences were numerically smaller than expected based on the Verbruggen and Logan (2008a, Experiment 2) findings. Therefore, in Experiment 10, three main changes were made to the task procedure. First, the response tracking procedure was modified to push participants to respond as quickly as possible. It was important to encourage rapid responding, as sluggish responding could reduce the probability that any acquired stimulus-response or stimulus-stop associations would influence task performance (for RT percentiles, see Figure B.4, Chapter 2). In Experiment 9, I observed that the implementation of the tracking procedure resulted in a higher response threshold in the two-choice task (*M* = 695 ms, *SD* = 129) than in the go/no-go task (*M* = 435 ms, *SD* = 66 ms), $F(1, 62) = 103.27$, $p < 0.001$, *gen.* $\eta^2 = 0.625$.

Consequently, participants in the two-choice task condition received fewer ‘too slow’ feedback messages (relative to ‘incorrect’ feedback messages) than participants in the go/no-go task condition. In Experiment 10, the tracking procedure was adjusted based only on responding on correct trials to ensure that the tracking procedure did not take into account response accuracy (only response latencies). I expected that this modified tracking procedure would push participants to respond quickly more than in Experiment 9. Furthermore, ‘too slow’ feedback messages were only presented when the RT exceeded the threshold and did not present feedback on other trials (cf. ‘correct’, ‘incorrect’, ‘too slow’ feedback messages in Experiment 9).

Second, an instruction screen with the stimulus-response and stimulus-stop (no-go) mappings was presented at the beginning of each block to better match the procedures in the training phase and in the test phase (cf. in Experiment 9, participants were instructed only the beginning of each phase). This modification ensured that the improvements during the training phase could not be due to participants forgetting the task instructions (and therefore learning by feedback).

Third, the number of blocks in the training (blocks 1-18; cf. blocks 1-12 in Experiment 9) and test (blocks 19-21; cf. blocks 12-14 in Experiment 9) phases was increased. The addition of extra blocks in the training and test phases provided useful pilot experiment for the experiments reported in Chapter 6 where additional training (and test; see Appendix K) phase blocks were needed to increase the signal-to-noise ratio in the event-related potential analyses (i.e. I wanted to make sure that the effects of training/test were still present when the number of blocks increased).

Method

Participants. Sixty-four students from University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 19.59 years, *SD* = 1.51, 54 females, 60 right-handed). None of the participants took part in Experiment 9. Eight participants were removed and replaced; three due to technical problems with the stimulus presentation computers, four in the two-choice task because the percentage of correct trials was < 60%[42] and one in the go/no-go task because 73 (out of 80 trials) in the final block had RTs of 1 ms. Similar exclusion criteria were applied in Experiment 9 (see *Stimuli, apparatus, procedure, & analyses* above) suggesting that this participant held down the response key throughout the block.

Stimuli, apparatus, procedure, and analyses. The stimuli, apparatus, procedure and analyses were identical to Experiment 9 except for the following: there were eighteen blocks in the training phase and three blocks in the test phase. After the response deadline elapsed (1000 ms), I presented the words ‘too slow’ in the centre of the screen when participants did not respond before the

---

[42] Note that it is a common finding that accuracy is lower in the two-choice task than in the go/no-go task (e.g. Gomez, Ratcliff, et al., 2007). Therefore, the exclusion criteria was set as < 60% correct trials in the two-choice task rather than < 70% correct trials as used in the other go/stop tasks presented in this thesis (see, e.g. Chapters 3 & 4).

response threshold. The feedback message on ‘too slow’ trials remained on the screen for 250 ms; on all other trials, a blank screen was presented for 250 ms. The response threshold tracking procedure was modified such that the response threshold decreased by 50 ms following four consecutive correct responses executed within the threshold but increased by 50 ms for correct responses executed after the response threshold (in Experiment 9, the threshold increased by 50 ms for all other responses, regardless of whether they were correct/incorrect). To ensure that participants paid attention to the new judgment task and to encourage them to prepare in advance, I presented a new instruction screen at the beginning of the test phase blocks with the words ‘ATTENTION: NEW RULE!’ at the bottom of the screen in red (RGB: 255 0 0). Participants were required to press the ‘s’ key confirm that had understood the new rule and to start the next block. At the beginning of all other blocks, I presented a rule instruction screen with the word ‘REMINDER!’ at the bottom of the screen to encourage participants to attend to the rules. Participants were required to press the ‘n’ key to start the next block.

Results and Discussion

Training phase. Analyses of correct RTs revealed that the main effect of block was significant ($p < 0.001$; Table 5.5) indicating that mean RTs decreased as a function of practice (Figure 5.5). Follow-up analyses indicated that the main effect of block was reliable in the go/no-go task ($p < 0.001$; Table 5.5) and in the two-choice task ($p < 0.001$; Table 5.5). The two-way interaction between task and block was reliable ($p = 0.020$; Table 5.5) reflecting the increased the speeding of RTs as a function of block in the two-choice task than in the go/no-go task (i.e. the reverse of the interaction found in Experiment 9; Figure 5.2) suggesting that my modified tracking procedure was effective in encouraging fast responding, particularly in the two-choice task[43]. Furthermore, a post-hoc comparison of response latencies in Experiment 10 and in Experiment 9 confirmed that responding was reliably faster in Experiment 10 ($M = 339$ ms, $SD = 42$ ms) than in Experiment 9 ($M = 367$ ms, $SD = 47$ ms), $F(1, 124) = 14.00$, $p < 0.001$, *gen.* $\eta^2 = 0.101$[44]. Surprisingly, there was no reliable main effect of task in the overall RTs ($p = 0.547$; Table 5.5) so I plotted RT distributions to explore the absence of this effect. These distributions showed that responding was faster in the two-choice task than in the go/no-go task in the fastest RTs but that this difference reversed for the slowest RTs (as reflected by a reliable two-way interaction between task and percentile; see Appendix J).

---

[43] Consistent with these findings, there was also no reliable difference between response threshold values in the go/no-go task ($M = 395$ ms, $SD = 45$ ms) and in the two-choice task ($M = 412$ ms, $SD = 56$ ms), $F(1, 62) = 1.84$, $p = 0.180$, *gen.* $\eta^2 = 0.028$.

[44] Note that I only examined blocks 1-12 in this comparison to ensure that any between-experiment differences were not due to the increased number of training blocks in Experiment 10 compared with Experiment 9.

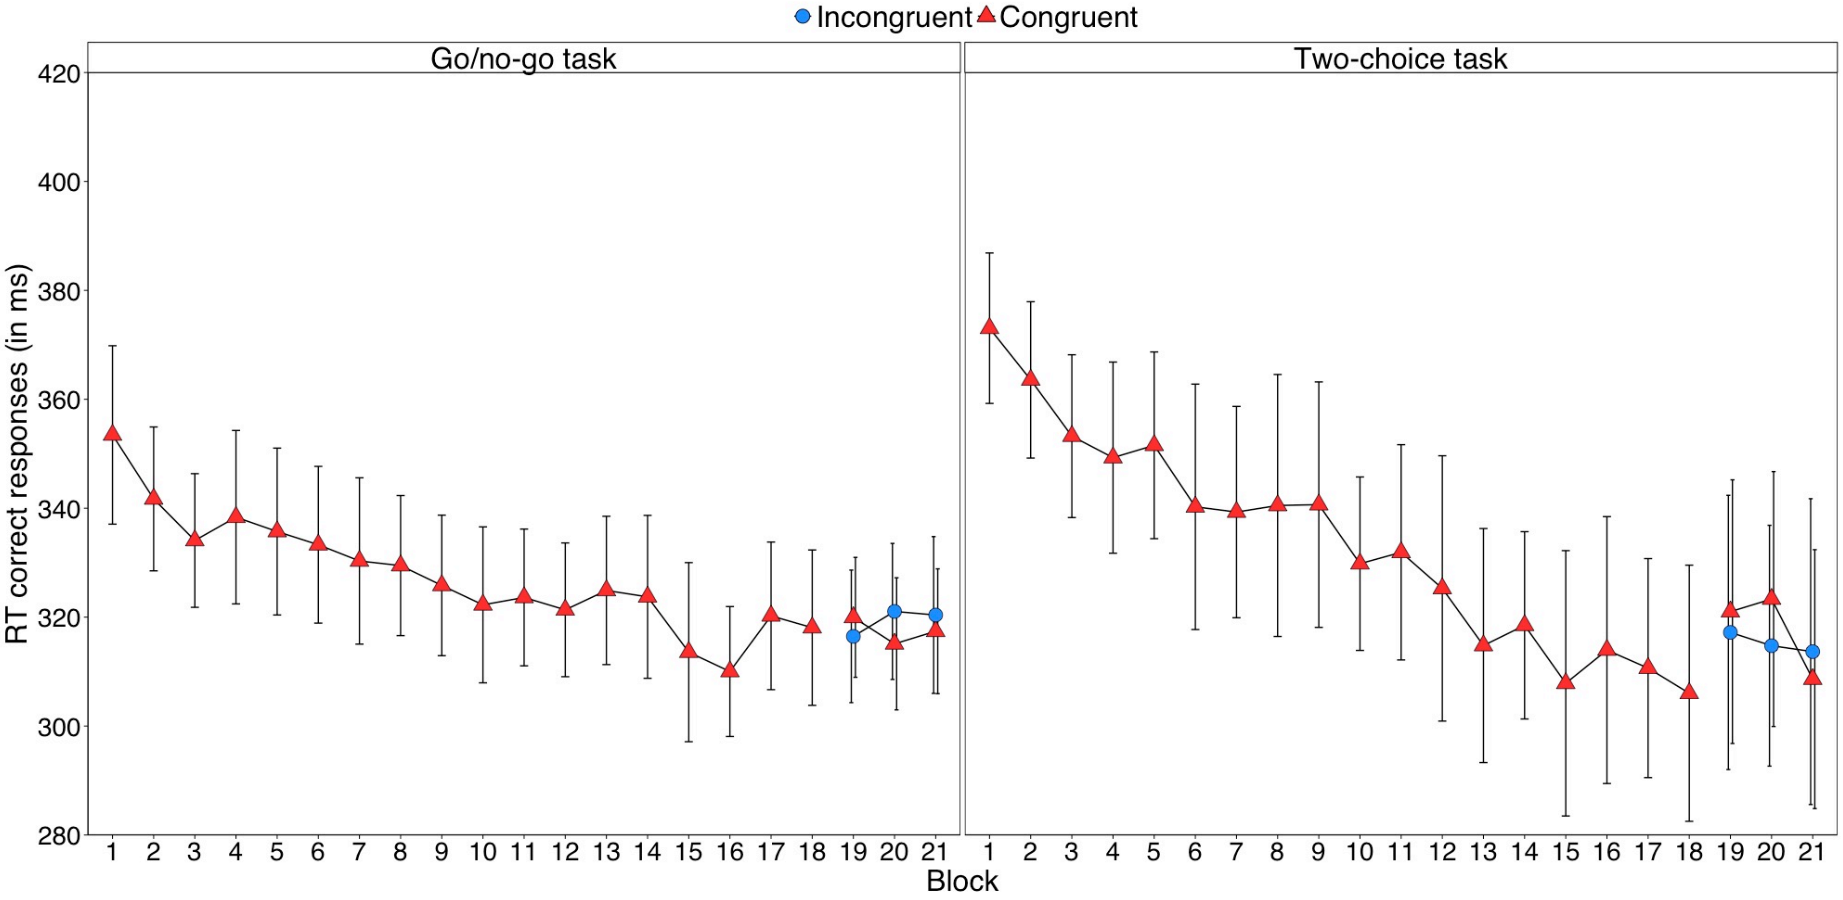

Figure 5.5. Reaction times of correct responses on go trials in the go/no-go task (left panel) and on all trials in the two-choice task (right panel) in Experiment 10. Blocks 1-18 were the training phase; blocks 19-21 were the test phase. Error bars are 95% confidence intervals.

Analyses of the probability of correct responses in the two-choice task revealed a main effect of block ($p < 0.001$; Table 5.5). However, as can be seen in Figure 5.6, this reflects decreasing choice accuracy with task practice. Analyses of the probability of responding on no-go trials in the go/no-go task also revealed a similar pattern: the main effect of block was reliable ($p < 0.001$; Table 5.5) but this reflects the increasing *p*(respond|no-go) with training (Figure 5.7). This suggests that whilst the modified tracking procedure and feedback was successful in encouraging fast responding, this was at the expense of more error-prone performance. In support of this idea, post-hoc tests revealed that choice accuracy in the two-choice task was reliably lower in Experiment 10 (0.79) than in Experiment 9 (0.84)[3], $F(1, 62) = 9.47$, $p = 0.003$, *gen.* $\eta^2 = 0.132$, and the *p*(respond|no-go) was reliably greater in Experiment 10 ($M = 0.15$, $SD = 0.08$) than in Experiment 9 ($M = 0.12$, $SD = 0.05$), $F(1, 62) = 4.40$, $p = 0.039$, *gen.* $\eta^2 = 0.066$. Note that the relatively high error rates in these tasks was most likely due to the adaptive response threshold tracking procedure (see *Method* above) that encouraged participants to prioritise rapid responding.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen.*** η^2
Combined task analysis							
Reaction times							
Task	1	62	10792.83	1826833.80	0.37	0.547	0.004
Block	17	1054	240122.17	776055.70	19.18	**< 0.001**	0.084
Task by block	17	1054	31222.00	776055.70	2.49	**0.020**	0.012
Go/no-go task only							
Reaction times: go trials							
Block	17	527	60604.22	163305.30	11.50	**< 0.001**	0.066
p(respond|no-go): no-go trials							
block	17	527	0.53	3.53	4.68	**< 0.001**	0.065
Two-choice task only							
Reaction times: all trials							
Block	17	527	210739.90	612750.40	10.66	**< 0.001**	0.107
Choice accuracy: all trials							
Block	17	527	0.68	2.98	7.04	**< 0.001**	0.091

Table 5.5. Overview of the Analyses of Variance on the *training phase* data (blocks 1-18) in Experiment 10 with block as a within-subjects factor and task as a between-subjects factor. The separate analyses for the go/no-go task and for the two-choice task are also presented. *p*s < 0.05 are highlighted in bold.

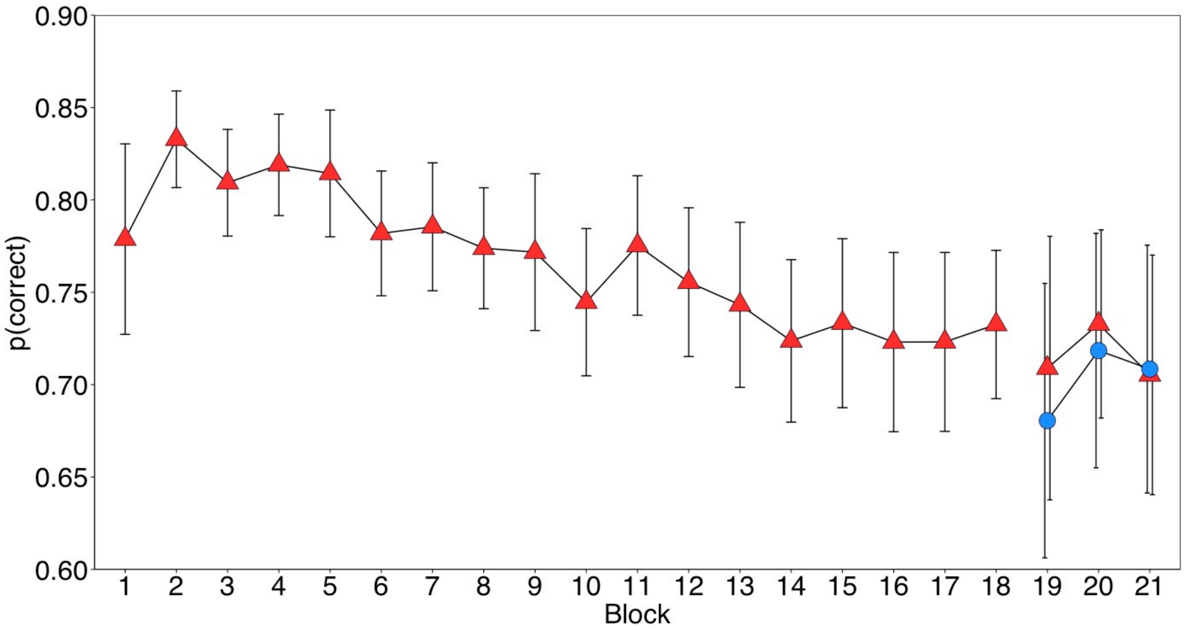

Figure 5.6. The probability of correct responses in the two-choice task in Experiment 10. Blocks 1-18 were the training phase; blocks 19-21 were the test phase. Error bars are 95% confidence intervals.

Test phase. Analyses of the RTs revealed no reliable main effect of congruency ($p = 0.875$; Table 5.6; see Figure 5.3). The main effect of task was not reliable ($p = 0.876$; see Appendix J for RT distribution analyses) and the two-way interaction between task and congruency was also not reliable ($p = 0.295$; Table 5.6); Table 5.7 shows the outcome of the Bayesian analysis. As can be seen, the results are consistent with the ANOVAs reported (Table 5.6). Table 5.7 shows that all factors or interactions could be dropped (Table 5.7 also shows the separate models for the two-choice and go/no-go tasks). Taken together, I conclude that there was no difference between the incongruent and congruent stimuli in the go/no-go and two-choice tasks.

Analyses of the probability of correct responses in the two-choice task revealed no reliable main effects of congruency ($p = 0.546$; Table 5.6; Figure 5.6). Bayesian analyses confirmed that the congruency factor could be dropped from the model (Table 5.7). As can be seen in Figure 5.7, the probability of responding on no-go trials in the go/no-go task was slightly higher for the incongruent stimuli than for the congruent stimuli. Analyses revealed that this difference was reliable ($p = 0.026$; Table 5.6). This suggests that when pushed to respond quickly during training, stopping performance is influenced by the retrieval of old stimulus-go associations.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen.*** η^2
Combined task analysis							
Reaction times							
Task	1	62	372.51	951905.26	0.02	0.877	< 0.001
Congruency	1	62	9.59	24145.97	0.02	0.876	< 0.001
Block	2	124	1097.40	61588.26	1.10	0.332	0.001
Task by congruency	1	62	434.76	24145.97	1.12	0.295	< 0.001
Task by block	2	124	1608.95	61588.26	1.62	0.204	0.001
Congruency by block	2	124	988.02	40487.96	1.51	0.226	0.001
Task by congruency by block	2	124	1286.74	40487.96	1.97	0.150	0.001
Go/no-go task only							
Reaction times: go trials							
Congruency	1	31	157.60	12389.39	0.39	0.535	0.001
Block	2	62	25.13	20866.71	0.04	0.963	< 0.001
Congruency by block	2	62	746.25	11713.66	1.97	0.147	0.003
p(respond|no-go): no-go trials							
Congruency	1	31	0.05	0.26	5.45	**0.026**	0.014
Block	2	62	0.02	0.53	1.42	0.249	0.008
Congruency by block	2	62	0.01	0.41	0.92	0.404	0.004
Two-choice task only							
Reaction times: all trials							
Congruency	1	31	286.75	11756.57	0.76	0.391	< 0.001
Block	2	62	2681.22	40721.55	2.04	0.145	0.003
Congruency by block	2	62	1528.51	28774.30	1.65	0.207	0.002
Choice accuracy: all trials							
Congruency	1	31	0.01	0.70	0.37	0.546	0.001
Block	2	62	0.03	1.81	0.53	0.506	0.005
Congruency by block	2	62	0.01	0.31	0.79	0.457	0.001

Table 5.6. Overview of the Analyses of Variance on the *test phase* data (blocks 19-21) in Experiment 10 with congruency and block as within-subjects factors and task as a between-subjects factor. The separate analyses for the go/no-go task and the two-choice task are also presented. *p*s < 0.05 are highlighted in bold.

Omitted Factor(s)	Bayes Factor	Confidence interval
Combined task analysis		
Reaction times		
Task	2.90	±22.69%
Block by task	4.29	±29.29%
Block by task by congruency	4.64	±24.93%
Task by congruency	6.14	±20.27%
Block by congruency	6.96	±17.34%
Congruency	9.93	±18.59%
Block	10.35	±17.2%
Go/no-go task only		
Reaction times: go trials		
Block by congruency	3.54	±6.92%
Congruency	4.53	±7.26%
Block	15.87	±6.58%
p(respond|no-go): no-go trials		
Congruency	0.42	±5.91%
Block	4.61	±5.32%
Block by congruency	5.65	±5.34%
Two-choice task only		
Reaction times: all trials		
Block	2.12	±7.61%
Block by congruency	3.35	±7.58%
Congruency	5.22	±7.96%
Choice accuracy: all trials		
Congruency	5.18	±6.41%
Block by congruency	9.31	±5.93%
Block	9.36	±5.9%

Table 5.7. Bayesian analysis for the *test phase* data in Experiment 10. For details, see Table 5.3.

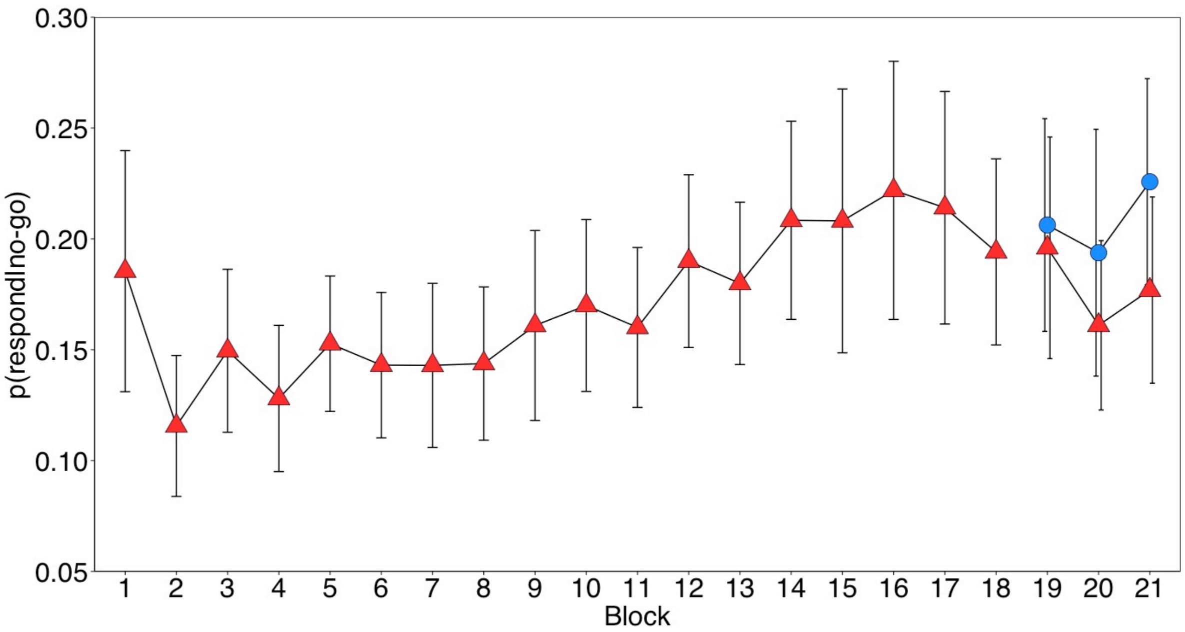

Figure 5.7. The probability of responding on no-go trials in the go/no-go task in Experiment 10. Blocks 1-18 were the training phase; blocks 19-21 were the test phase. Error bars are 95% confidence intervals.

Figure 5.7. The probability of responding on no-go trials in the go/no-go task in Experiment 10. Blocks 1-18 were the training phase; blocks 19-21 were the test phase. Error bars are 95% confidence intervals.

Training vs. test phase. As in Experiment 9, I performed two additional analyses to investigate whether there were any overall differences between the training and test phases. First, I compared RTs in the go/no-go task for the first three blocks in the training phase (blocks 1-3) with RTs in the three blocks in the test phase (blocks 19-21). Contrary to my predictions, this revealed that the three-way interaction between congruency, phase, and block was not reliable, $F(1, 31) = 0.131$, $p = 0.720$, *gen.* $\eta^2 < 0.001$. Follow-up comparisons revealed that the two-way interaction between block (the first three blocks in the training phase vs. the three blocks in the test phase) and phase (training vs. test) was reliable for the incongruent stimuli, $F(1, 31) = 6.52$, $p = 0.015$, *gen.* $\eta^2 = 0.037$, and for the congruent stimuli, $F(1, 31) = 6.29$, $p = 0.018$, *gen.* $\eta^2 = 0.047$. Consistent with the pattern of results in Experiment 9, responding for the incongruent stimuli became faster during the first three blocks of the training phase (block 1: 353 ms; block 2: 341 ms; block 3: 338 ms) but responding for these items became slower during the test phase blocks (block 19: 316 ms; block 20: 321 ms; block 21: 320 ms); responding for the congruent stimuli became faster in the first three blocks in the training phase (block 1: 354 ms; block 2: 342 ms; block 3: 330 ms) which was greater in magnitude that the speed-up across the test phase blocks (block 19: 320 ms; block 20: 315 ms; block 21: 317 ms). Combined, this analysis provides support for the idea that performance in the test phase was influenced partially by the retrieval of old stimulus-stop (no-go) associations.

Second, I compared performance between the end of the training phase (blocks 16-18) and the test phase (blocks 19-21). As in Experiment 9, introducing the new judgment task in the test phase did not result in a difference in performance compared with at the end of training in RTs, $F(1, 62) = 1.16$, $p = 0.284$, $gen.\ \eta^2 = 0.001$, choice accuracy, $F(1, 31) = 1.13$, $p = 0.295$, $gen.\ \eta^2 = 0.005$, or in the p(respond|no-go), $F(1, 31) = 1.06$, $p = 0.310$, $gen.\ \eta^2 = 0.006$. The two-way interaction between phase and task on RTs was also not reliable, $F(1, 62) = 1.85$, $p = 0.668$, $gen.\ \eta^2 < 0.001$. Bayesian analyses were also broadly consistent with ANOVAs reported (Table 5.8) . Taken together, these findings suggest that the stimulus-stop (no-go) and stimulus-response associations were not mediated via strong stimulus-category associations in either the go/no-go task or in the two-choice task because (strong) mediation via a category representation should result in a decrement in performance between the end of training and the test phase (for a discussion of the factors that could influence the extent to which the stimulus-stop associations are mediated via a category representation, see *Mediation via a Category Representation*, Chapter 7).

Omitted factor(s)	**Bayes Factor**	**Confidence Interval**
Combined task analysis		
Reaction times		
Task	2.26	±6.56%
Phase	2.43	±6.49%
Task by phase	5.58	±6.42%
Factor	**Bayes Factor**	**Confidence Interval**
Go/no-go task only		
Reaction times: go trials		
Phase	0.20	±2.47%
p(respond|no-go): no-go trials		
Phase	0.34	±2.12%
Two-choice task only		
Reaction times: all trials		
Phase	0.46	±3.05%
Choice accuracy: all trials		
Phase	0.24	±2.20%

Table 5.8. Bayesian analyses for the difference between performance at the end of training vs. in the test phase in Experiment 10. For details, see Table 5.4.

Chapter Discussion

Across two experiments presented in this chapter, I show that performance improved with task practice; response latencies decreased in the go/no-go and two-choice tasks (Experiments 1-2) and choice accuracy increased in the two-choice task (Experiment 9) during the training phase. In the test phase, I examined whether memory-retrieval influenced performance via the introduction of a novel judgment task. For half of the stimuli the response changed between the training and test phases (incongruent stimuli); for the other half, the response stayed the same (congruent stimuli). Across both experiments, I found some evidence of differences between the incongruent and congruent stimuli in go RTs in the go/no-go task (Experiment 9), in the *p*(respond|no-go) in the go/no-go task (Experiment 10), and in choice accuracy in the two-choice task (Experiment 9). However, these differences were much smaller than expected. In what follows, I propose two possible non-associative influences on performance during the training phase of inhibition and non-inhibition tasks.

The Role of Non-Associative Processes

The pattern of responding observed during the training phase in Experiments 9-10 is consistent with the idea that, as a function of task practice, performance shifted from an entirely algorithmic action selection process *towards* the retrieval of the appropriate response from memory (i.e. Instance theory Logan, 1988, 1990). However, as mentioned above, it is notable that the differences between the congruent and incongruent items in the test phase were much smaller than expected. This indicates that performance was not *fully* automatised by the end of training (see Verbruggen & Logan, 2008a).

In previous stop learning studies, improvements observed during training have been attributed to the acquisition of associations between specific stimuli and the correct response (e.g. ‘go’ or ‘no-go’). According to this idea, when stimuli are repeated, the acquired associations are retrieved from memory, and the repetition priming effects observed in the training phase are the first step towards automatisation (see e.g. Logan, 1988, 1990). However, the negligible cost in performance observed between the congruent and incongruent items in both the go/no-go and two-choice tasks in Experiments 9-10 suggests that the effects of training on behavioural performance in these tasks is unlikely to reflect entirely associative (memory-retrieval) processes.

One (non-associative) explanation is that participants require some time to find an appropriate balance between speed and accuracy, particularly when the task instructions emphasise that participants should prioritise speed over accuracy (as in Experiment 10) or vice versa. Consistent with this suggestion, Chein and Schneider (2012) suggested that there are three hierarchically organised learning systems: a metacognitive system (involved in the generation of new routines, task sequencing, and monitoring), a cognitive control network (involved in attention control and selection, action sequencing, reinforcement, and arousal), and a representational system (involved in associative learning). According this framework, when people begin performing a new task there is an initial formation stage in which the metacognitive system is involved in establishing the task strategies and behavioural routines that are required for the execution of the task. Once these routines are established, the person moves into a controlled-execution stage, in which the ‘cognitive control network’ is involved in guiding action selection via attentional control, arousal, and reinforcement signals. During these earlier learning stages, associative mechanisms within the representation system are involved in slowly strengthening the key stimulus-response relationships that underlie task execution, gradually allowing cognitive-control resources to disengage as task performance becomes practiced and, eventually automatic. Thus, the early stages of learning involve monitoring for information consistent with the task goals. It is only after sufficient practice that performance is supported entirely by the representational system. For the purposes of the present chapter, the key feature of this model is that it highlights that the effects of training on behaviour do not necessarily reflect entirely associative processes. Thus, it seems that the effects of task practice in inhibition and non-inhibition tasks depends on an interplay of bottom-up learning processes and top-down control processes, alongside non-specific practice-effects (e.g. rule learning, monitoring etc.) arising from adjustments to response priorities.

The findings in the present chapter also provide some evidence that practice-effects may be sensitive to task demands. In Experiment 10, I modified the go response thresholds and the trial feedback to push participants to respond faster than in Experiment 9. Analyses of the training phase data suggests task performance was sensitive to these modifications: responding was faster, choice accuracy was lower, and the *p*(respond|no-go) was higher in Experiment 10 than in Experiment 9. Furthermore, the practice-effects on response latencies in the training phase (i.e. shorter RTs at the end of the training phase compared with the beginning of the training phase) were larger in Experiment 10 than in Experiment 9, although the practice-effects previously observed in choice accuracy in the two-choice task were abolished. Analyses of the test phase in Experiment 10 showed no reliable differences in go/no-go and two-choice task performance between the incongruent and congruent stimuli in the test phase in response latencies or in choice accuracy (cf. Experiment 9 where I observed some small but reliable differences) but I did find a small but reliable cost in the probability of responding on no-go trials in the go/no-go task. These findings were supported by Bayesian analyses. The absence of a cost in go performance (despite stronger practice-effects on go performance during training) indicates participants engaged top-down control processes to mitigate any response delay incurred by the retrieval of an incorrect response (Logan, 1988, 1990). Logan argued that there is a race between algorithmic processes and memory-retrieval processes. This could explain why the effects of memory-retrieval on task performance appear to be sensitive to response priorities; if the task does not push participants enough, algorithmic processing is slower and there is a higher probability that the memory-retrieval process will win the race (hence a small congruency effect in Experiment 9). By contrast, if the task pushes participants a lot, they may exert more top-down control to bias action selection in favour of fast algorithmic processing, reducing the probability that the memory-retrieval process will win the race[45] (see Figure 5.8). Pushing participants might also encourage them more to shield attention from distracting features of the stimuli that are not part of the specified task set (i.e. those associated with the old trained response; cf. Dreisbach & Haider, 2009).

---

[45] To be clear, I refer here to fast responding as a result of strategic adjustments rather than fast responding as a result of impulsive tendencies.

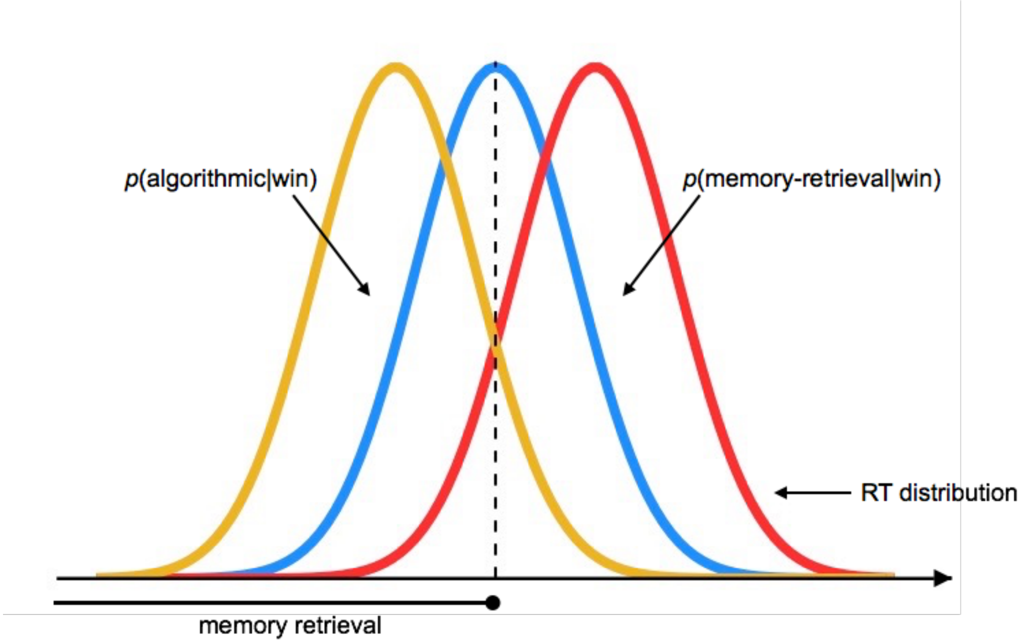

Figure 5.8. Visual depiction of the probability of the algorithmic process ‘winning’ the race based on a Logan’s (1988, 1990) race model. The probability of the algorithmic process winning the race is represented by the area under the curve to the left of the dashed line. As the go RT distribution is shifted to the right, the probability of the algorithmic process winning decreases; as the distribution is shifted to the left, the probability of the algorithmic process winning increases. Note: for simplicity, I represent the duration of the memory retrieval process as constant in this figure, but in reality the duration of the memory retrieval would most likely vary depending on the number of instances stored in memory.

A key outstanding issue addressed in this chapter concerned the specificity of learning in response inhibition tasks. The comparison of behavioural performance between the go/no-go and two-choice tasks provides some insight into this issue. Across both tasks, I find strong practice-effects in response latencies during training. These findings are consistent with previous research that has shown that action selection in inhibition and non-inhibition tasks are strongly related (e.g. Mostofsky & Simmonds, 2008). However, the conclusion that learning has similar effects on action selection in inhibition and non-inhibition tasks is entirely based on the observation that the learning curves were similar in both experiments. It is possible that different learning mechanisms produce similar learning effects. Therefore, in Chapter 6, I follow-up on these findings using event-related potentials to index the processes modulated by practice in go/no-go and two-choice tasks. I also primarily use a within-subjects design in Chapter 6 to reduce the possibility that between-task differences could be accounted for by indirect consequences of task performance (such as fatigue) and to increase the power to detect small to medium-sized differences between-tasks (the present chapter was sufficiently powered to detect within-task differences and within-between interactions, but not between-task differences).

An outstanding issue raised the present chapter is why the observed differences between the incongruent and congruent stimuli in the go/no-go task are smaller than in previous work (Verbruggen and Logan, 2008a). There are at least three explanations that could account for the small differences observed in the present chapter. First, it is possible that the behavioural effects of interference between the incongruent and congruent stimuli in the present chapter were simply short-lived and were consequently masked by rapid re-learning of the new associations. However, note that a visual inspection of the data argues against this suggestion; if anything, the differences between the incongruent and congruent stimuli seem larger in the second test phase block than in the first (this conclusion is further supported by a reliable two-way interaction between congruency and block in go RTs in Experiment 9; see Table 5.1). Second, it is possible that the use of a smaller stimulus set than in the Verbruggen and Logan, 2008a) task, may have enabled participants to adjust their response settings in preparation for, or rapidly during, the test phase. Third, it is possible that the Verbruggen and Logan (2008a) go/no-go task was more difficult. Consistent with this idea, mean reaction times on go trials in the Verbruggen and Logan (2008a) task were noticeably slower than in my task. Slower reaction times could have increased the probability that the memory retrieval process won the race (see above for a discussion of this idea). For an extended discussion of these ideas, see Chapter 7.

Conclusion

Taken together, I conclude that during practice of inhibition and non-inhibition tasks participants can acquire various associations between a stimulus and the response-related information contained within an ‘instance’ (Logan, 1988, 1990). However, in addition to memory-based effects, these findings highlight that non-associative practice-effects also play a role in both inhibition and non-inhibition tasks. Furthermore, the findings presented in this chapter also suggest that what is learned and retrieved in go/no-go and choice tasks depends on task design and response priorities. Thus, these findings open the door to further investigations into the specificity of response inhibition training (such as the experiments reported in Chapter 6).

Appendix J

RT Distributions

To investigate the absence of difference in response latencies between the go/no-go and two-choice tasks in Experiment 10, I plotted RT percentiles for the training and test phases. For comparison, I also plotted the response latencies for Experiment 9 where I found, as expected, that responding was faster in the go/no-go task than in the two-choice task (Tables 6.1 & 6.2). Visual inspection of the percentile plots suggests that in Experiment 9 responding was faster in the the go/no-go task than in the two-choice task across the whole RT distribution, but the biggest difference in response latencies between the go/no-go and two-choice tasks emerged in the slow end of the RT distribution. This conclusion is supported by a reliable interaction between task and percentile in the training phase, $F(8, 496) = 9.77$, $p < 0.001$, *gen.* $\eta^2 = 0.013$, and in the test phase, $F(8, 496) = 12.33$, $p < 0.001$, *gen.* $\eta^2 = 0.015$. However, in Experiment 10 responding was faster in the two-choice task than in the go/no-go task in the fast end of the RT distribution but was faster in the go/no-go task than in the two-choice task in the slow end of the RT distribution (i.e the direction of the difference between tasks reversed across the RT distribution). This conclusion is also supported by a reliable two-way interaction between task and percentile in the training phase, $F(8, 496) = 7.80$, $p < 0.001$, *gen.* $\eta^2 = 0.016$, and in the test phase, $F(8, 496) = 7.46$, $p < 0.001$, *gen.* $\eta^2 = 0.015$. This suggests that the absence of a main effect of task on RTs in Experiment 10 is because the distribution of RTs was different between the conditions: in the two-choice task, the distribution was more spread out than in the go/no-go task and the RT at the onset of the distribution (10th percentile) was faster in the two-choice task than in the go/no-go task.

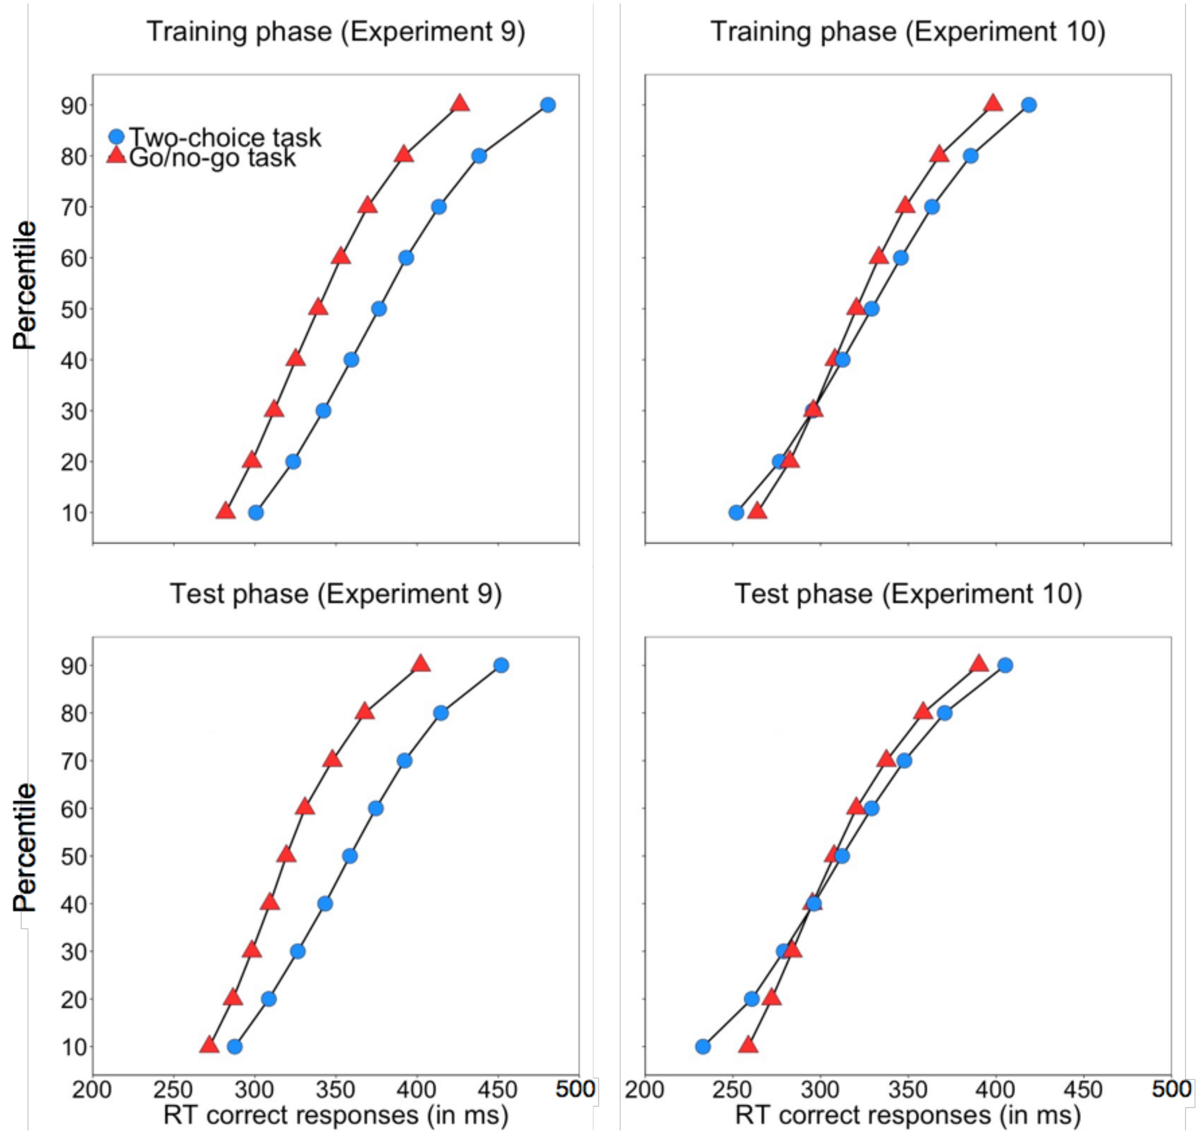

Figure J.1. The RT distributions for Experiment 9 (left panels) and Experiment 10 (right panels) as a function of phase (training phase: upper panels; test phase: lower panels) and task (go/no-go task; two-choice task).

CHAPTER 6

THE ELECTROPHYSIOLOGY OF PRACTICE-EFFECTS IN INHIBITION AND NON-INHIBITION TASKS

In Chapter 5, I found qualitatively similar effects of training on behavioural performance in inhibition and non-inhibition tasks. This could indicate that learning to inhibit a response is achieved through the same or similar mechanisms as learning to select and execute a motor response. This would provide strong support for the idea that response inhibition and response execution depend on similar basic processes (McLaren & Verbruggen, 2016; Verbruggen, McLaren, et al., 2014). However, evidence of similar behavioural outcomes of training does not necessarily constitute evidence that similar (neural) mechanisms were involved. Therefore, in the present chapter, I follow-up on the behavioural findings reported in Chapter 5 by comparing the effects of training in go/no-go and two-choice tasks on two event-related potential components thought to index processes involved in action selection (the N2 and the P3 components).

Event-related potentials (ERPs) provide an online measure of brain activity with up to millisecond precision[46]. Performance of response inhibition tasks is associated with two ERP components: the N2, a negative component peaking approximately 200-300 ms following stimulus presentation, and the P3, a positive component peaking approximately 300-500 ms following stimulus presentation (e.g. Eimer, 1993; Falkenstein, Koshlykova, Kiroj, Hoorman, & Hohnsbein, 1995; Falkenstein, Hoorman, & Hohnsbein, 1999; Jodo & Kayama, 1992; Kiefer, Marzinzik, Weisbrod, Scherg, & Spitzer, 1998; Pfefferbaum, Ford, Weller, & Kopell,1985). The established finding is that the N2 and P3 components are larger (i.e. more negative and more positive, respectively) on stop (or no-go) trials than on go trials (e.g. Bekker, Kenemans, Hoeksma, Talsma, & Verbaten, 2005; Bruin & Wijers, 2001; Jodo & Kayama, 1992; Smith, Johnstone, & Barry, 2008).

Whilst some studies have attributed the N2 and P3 components solely to processes associated with response inhibition (e.g. Bruin & Wijers, 2001; Kok, Ramautar, De Ruiter, Band, & Ridderinkhof, 2004; Lavric, Pizzagalli, & Forstmeier, 2004; Jodo & Kayama, 1992), other studies have suggested that these components index more general processes that are also involved in performance of non-inhibition tasks. For example, the N2 component has been linked to processes involved in the resolution of conflict between response options and the P3 component has been linked to evaluative and attentional processes involved in memory-updating processes (Gajewski &

---

[46] ERP amplitudes are thought to reflect the level of engagement in cognitive processes whereas peak latencies reflect the timing of cognitive processes (e.g. Luck, 2005).

Falkenstein, 2012; Huster, Enriquez-Geppert, et al., 2012) and, more generally, the neural response to the outcome of internal decision-making processes (Nieuwenhuis, Aston-Jones, & Cohen, 2005). In support of conflict interpretation of the N2 component, it has been shown that the N2 is more negative on relatively infrequent trials regardless of whether going or stopping is required (Enriquez-Geppert, Konrad, Pantev, & Huster, 2010; Wessel, Danielmeier, Morton, & Ullsperger, 2012) and more negative when the similarity between the go and no-go stimuli is high (Nieuwenhuis, Yeung, & Cohen, 2004; Smith & Douglas, 2011)[47]. In support of task-general interpretations of the P3 component, research suggests that the P3 component consists of at least two subcomponents with different scalp distributions and corresponding underlying processes; a frontally-maximal 'P3a' and a parietally-maximal 'P3b' (e.g. Berti, 2008; Donchin, 1981; Duncan-Johnson & Donchin, 1997; Friedman, Cycowicz, & Gaeta, 2001; Polich, 2007; Verleger, Jaskowski, & Wascher, 2005). The P3a and P3b components are thought to occur during memory updating (Polich & Kok, 1995) with the P3a reflecting an attentional process involved in the selection of relevant information from memory and the P3b reflecting context-updating processes and/or processes that mediates the link between perception and motor response selection (e.g. Donchin, 1981; Friedman, Cycowicz, et al., 2001; Polich, 2007; Verleger, Jaskowski, et al., 2005). Note that the aforementioned 'no-go P3'effect (i.e. larger P3 amplitudes for no-go stimuli than for go stimuli) corresponds to the P3a in this model, whereas more parietal P3 components (such as the P3b) can display the opposite effect (i.e. more positive amplitudes for go stimuli than for no-go; Pfefferbaum, Ford, et al., 1985).

To date, a small number of studies have utilised ERPs to index the neural signatures of training in the go/no-go task. Note that whilst some studies have reported effects of training on early sensory and attentional components in the go/no-go task (Benikos, Johnstone, et al., 2013; see also Manuel, Grivel, et al., 2010) these effects are most likely to reflex associative processes influencing the detection of the no-go stimulus rather than direct activation of an inhibition or no-go network (i.e. as suggested by the automatic inhibition hypothesis; Verbruggen, Best, et al., 2014; see also *Stimulus Detection*, Chapter 7). Therefore, as I am interested in the effects of training on action selection processes, I focus on the N2 and P3 components in this chapter.

Although research in the stop learning literature generally converges on the finding that the N2 and P3 components peak earlier at the end of training than at the beginning of training, the effects of training on the amplitudes of these component are less clear. As can be seen in Table 6.1, Benikos and colleagues showed that N2 amplitudes became *less* negative after go/no-go training (Benikos, Johnstone, et al., 2013; Benikos, unpublished PhD thesis), and they proposed that this could reflect a reduction in the level of response conflict triggered by the go and no-go stimuli with training

---

[47] Thus, according to the conflict interpretation, the larger (more negative) N2 amplitudes typically observed for no-go stimuli could reflect a bias towards go responding in response inhibition tasks, such that no-go or stop-signal trials require the reconfiguration of responses and/or the reconfiguration of attentional settings from the default 'go' setting(s) (see Huster, Enriquez-Geppert, et al., 2013).

(Benikos, Johnstone, et al., 2013). However, in another experiment, Schapkin and colleagues found that amplitudes in the N2 range became *more* negative with training for the no-go stimuli but not for the go stimuli (Schapkin, Falkenstein, Marks, & Griefahn, 2007). Schapkin, Johnstone, et al. attributed this difference to an 'improvement in the cognitive inhibition process in terms of a more coherent activation of the related brain units' (Schapkin, Johnstone, et al., 2007, p. 1286). A similar picture emerges for the P3 amplitudes; whilst some research reported that P3 amplitudes became more positive after training (Benikos, Johnstone, et al., 2013; Benikos, unpublished PhD thesis), other research reported no reliable change (Jodo & Inoue, 1990).

Inconsistencies between these studies are most likely due, at least in part, to (1) the varying number of training trials and sessions; (2) variations in go response deadlines (see Benikos, Johnstone, et al., 2013); and (3) low statistical power, with the majority of previous studies testing eleven or fewer participants (i.e. Jodo & Inoue, 1990; Schapkin, Falkenstein, et al., 2007).

Thus, although the N2 and P3 components appear to be sensitive to the effects of go/no-go training, it remains an outstanding issue exactly how go/no-go training modulates the specific processes underlying task performance. Furthermore, it is unclear the extent to which the reported effects on the N2 and P3 components are specific to training in inhibitory control tasks or whether they reflect more general training effects also found in non-inhibition tasks.

Study	Task	# Training Trials	# Sessions	# Participants	Behavioural outcomes	ERP amplitudes	ERP latencies	Specificity
Benikos (unpublished PhD thesis, Study 3)	vGNG	800 (30% no-go condition; 15% no-go condition; 70% no-go condition)	1	54 (18 per condition)	↓ Go RT no change: *p*(respond|no-go)	↓ N1 (in Fz) ↑ P2 (in Pz) ↓ N2 (in Fz) ↑ P3 (in Pz)	↓ N1 (in Fz) ↓ P2 (in Pz) ↓ N2 (in Fz)	Go & no-go
Benikos (unpublished PhD thesis, Study 4)	vGNG	800 (70% no-go)	1	18 (go/no-go between-subjects condition)	↓ Go RT	↓ N1 (in Fz) ↑ P2 (in Pz) ↓ N2 (in Fz) ↑ P3 (in Pz)	↓ P2 (in Pz) ↑ N2 (in Fz)	Go & no-go
Benikos, Johnstone, et al. (2013)	vGNG	800 (30% no-go)	1	60 (20 per low/medium/high difficulty condition)	↓ Go RT ↓ or ↑ *p*(respond|no-go)*	↓ N1 (in Fz) ↑ P2 (in Pz) ↓ N2 (in Fz) ↑ P3 (in Pz)	↓ N1 (in Fz) ↓ P2 (in Pz) ↓ N2 (in Fz) ↓ P3 (in Pz)	N2 amplitude: go only
Jodo & Inoue (1990)	vGNG	1200 (50% no-go)	6 (1 per day)	10	↓ Go RT *p*(respond|no-go): *NS*	no change: P3 (in Fz, Cz & Pz)	↓ P3 (in Fz, Cz & Pz)	No-go only
Schapkin, Johnstone, et al. (2007)	vGNG	~ 600 out of 6000 (50 % no-go)	6 (2 x per day for 3 days)	8	↓ Go RT ↓ *p*(respond|no-go)	↑ N2 (in Fz) **	*NS*	No-go only

Table 6.1. Overview of go/no-go training studies with behavioural and N2 and/or P3 outcome measures. Note*:* vGNG: visual go/no-go task; ↓ amplitudes = more negative, ↑ amplitudes = more positive; ↓ latencies = earlier, ↑ latencies = later; *NS* = not studied. For an extended overview of previous inhibitory training studies and outcomes, see Spierer, Chavan, et al., 2013). * The direction of this result depended on task difficulty (as manipulated by the go reaction time deadline). ** Note that these effects were observed after three days of a three week training protocol (total trials = 6000); however, there were no further reliable differences after three days.

The Present Chapter

In the present chapter I used ERPs to extend the behavioural findings reported in Chapter 5 and to examine how the neural signatures of action selection are modulated by training in the go/no-go and two-choice tasks. Although some research has compared the effects of go/no-go and two-choice task performance on the N2 and P3 components, no research to date compared the effects of training[48]. Evidence for similar training effects on the N2 and P3 components in these tasks would not only have important implications for our understanding of the brain mechanisms linked to inhibitory control training, but would also contribute to the wider on-going debate in the ERP literature concerning whether the N2 and P3 components index inhibition-related processes or more task-general processes. For the action control literature, evidence of similar effects of training on the N2 and/or P3 components in the go/no-go and two-choice tasks would add weight to the claim that response execution and response inhibition share similar overlapping processes (see Bunge, 2004; Duncan & Owen, 2000; Buch, Mars, Boorman, & Rushworth, 2010; Mostofsky & Simmons, 2008; Verbruggen, McLaren, et al., 2014). For the ERP literature, evidence of similar effects of training on the N2 and/or P3 components in the go/no-go and two-choice tasks would add weight to claims that these components index task-general processes rather than those exclusively linked to response inhibition.

As previous studies have found inconsistent effects of training on the N2 and P3 components (see above), I first conducted an exploratory pilot ERP study using a between-subjects task similar to the tasks used in Experiment 9-10 (Chapter 5). There was a training (acquisition) phase in which the stimulus category defined how the participant should respond, followed by a test phase in which the acquired associations were reversed. Consistent with my predictions, the behavioural task performance data confirmed that reaction times decreased during the training phase for the go and two-choice stimuli. Analyses of the ERPs showed that mean amplitudes in the N2 range (200-300 ms post-stimulus[49]) were (numerically) less negative at the end of training than at the beginning of training for the go, no-go and two-choice stimuli. For the two-choice stimuli, mean amplitudes in the P3a range (310-410 ms post-stimulus) and in the P3b range (300-400 ms post-stimulus) were more positive at the end of training than at the beginning of training for the two-choice task. However, for the go (and no-go) stimuli there were no reliable differences between the beginning and the end of training. Whilst this could indicate that there are some differences between the effects of training between the go/no-go and two-choice tasks, it is also possible that these differences reflect indirect consequences of task performance (such as fatigue; see Kato, Endo & Kizuka, 2009). Furthermore, the power in this pilot experiment was too low to allow for firm conclusions about the absence of a reliable effect in the go/no-go task. Nevertheless, these findings confirm prior findings that training influences N2 and P3

---

[48] For example, Smith and colleagues found that the N2 amplitude was more negative for any unexpected stimulus, regardless of whether it required response execution on go and two-choice trials or response inhibition on no-go trials (Smith, Smith, et al., 2009).

[49] Details regarding how these time-windows were selected are reported in *EEG data acquisition* below.

amplitudes, even if the specificity of this effect requires further investigation. In addition, the direction of the observed amplitude N2 differences are also consistent with some prior research discussed above (see Benikos, Johnstone, et al., 2013; Benikos, unpublished PhD thesis; Table 6.1).

To advance on the pilot experiment, in Experiments 11-12 (presented in the main body of this chapter) I used a within-subjects design in which participants alternated between performance of the go/no-go and two-choice tasks on a block-by-block basis. This eliminated any indirect consequences of task performance and increased the power to detect small to medium-sized differences between-tasks, however the within-subjects nature of the design meant that it was not possible to include a test (reversal) phase. Note that in the pilot study, I found there was a small but reliable performance cost in the test phase (relative to the end of training) and amplitudes in the P3a range (for the two-choice stimuli only) and in the P3b range were less positive in the test phase than at the end of training. However, I did not find any statistically reliable evidence that this reversal influenced the ERP amplitudes in the go/no-go task or that this reversal influence ERP peak latencies. Furthermore, the work presented in Chapter 5 also indicated that the reversal effects in the test phase were small. Thus, I focussed only on the effects of training in Experiments 11-12.

Experiment 11

In Experiment 11, participants alternated between performance of the go/no-go and two-choice tasks on a block-by-block basis. There were two stimuli per task with a 1:1 stimulus-to-response mapping (e.g. diamond = go; square= no-go; blue circle = left keypress; yellow circle = right keypress). Different stimulus sets were used in each task to avoid between-task interference (e.g. chequerboards = go/no-go task; coloured circles = two-choice task, or vice versa). Performance at the beginning of training (the first three blocks of each task) was compared with performance at the end of training (the final three blocks of each task). Training-effects were indexed via three behavioural indices of performance: reaction times, the probability of responding on no-go trials (in the go/no-go task) and choice accuracy (in the two-choice task). Alongside the behavioural indices, mean amplitudes and peak latencies were also measured in time-windows and electrodes corresponding to the N2, P3a, and P3b components.

To examine the effects of task practice on motor-related processes, I also computed response-locked readiness potentials for the go and two-choice stimuli. The readiness potential is thought to be generated in the primary motor cortex and supplementary motor areas, with activity larger in the hemisphere that is contralateral to the response hand (for an overview, see e.g. Eimer, 1998). For the experiments reported in this chapter, the response-locked components were particularly important for the interpretation of the stimulus-locked ERP training differences in the P3a and P3b ranges as differences in these components could reflect temporally overlapping differences in motor-related processes rather than the processes linked to the P3a and/or P3b components.

Method

Participants. Thirty-two right-handed volunteers from the University of Exeter participated for monetary compensation (£10) or partial course credit (*M* = 20.03 years, *SD* = 4.05, 27 females). Eleven participants were removed and replaced due to low trial numbers in the ERP averages (see the *ERP analyses* section for the exclusion criteria; for similar exclusion criteria see Elchlepp, Lavric, et al., 2016). All experiments of the present chapter were approved by the local research ethics committee at the School of Psychology, University of Exeter. Written informed consent was obtained after the nature and possible consequences of the study were explained. The target sample size and exclusion criteria were decided in advance of data collection to ensure that I had enough power (0.80) to detect medium-sized effects.

Apparatus, stimuli, and procedure. The experiment was run on a 21-inch iMac (screen size: 1920 × 1080 pixels) using Psychtoolbox (Brainard, 1997). The stimuli consisted of two chequerboards, which varied along the orientation dimension (0° or 45°) and two circles, which varied along the colour dimension (yellow [RGB: 255 204 102] or blue [RGB: 102 102 255]). All stimuli were 140 x 140 pixels and were presented on a grey background (RGB: 169 169 169). Participants alternated between the go/no-go task and the two-choice task on a block-by-block basis. The order of the first task block was counterbalanced across participants (e.g. participant 1: GN-TC-GN-TC…; participant 2: TC-GN-TC-GN…). In total there were 18 blocks; 9 blocks of the go/no-go task and 9 blocks of the two-choice task. Each block consisted of 80 trials and each stimulus within the relevant stimulus set (e.g. chequerboards or circles; Figure 6.1) was presented 40 times per block. Half of the participants decided whether the chequerboards were orientated as diamond or square in the go/no-go task and whether the circles were coloured blue or yellow in the two-choice task. For the other half of participants, the judgments were reversed. Participants received instructions on the screen at the beginning of each block to remind them of the stimulus-response and/or stimulus-stop (no-go) mappings in the upcoming block. If the average RT of correct responses in the previous block (of the relevant task) was greater than the running average RT of all preceding blocks (of the relevant task) participants received a warning message presented in red (RGB: 255 0 0) at the bottom of the instruction screen (‘*Please try to respond faster in this block!*’). Similarly if the average number of correct responses (in the two-choice task) was less than the running average of all the preceding blocks or if the number of responses on no-go trials (in the go/no-go task) was greater than the running average of all the preceding blocks of the go/no-go task, participants also received a warning message presented in red (RGB: 255 0 0) at the bottom of the instruction screen (‘*Please try to be more accurate in this block!*’). Participants were required to press any key to start the next block.

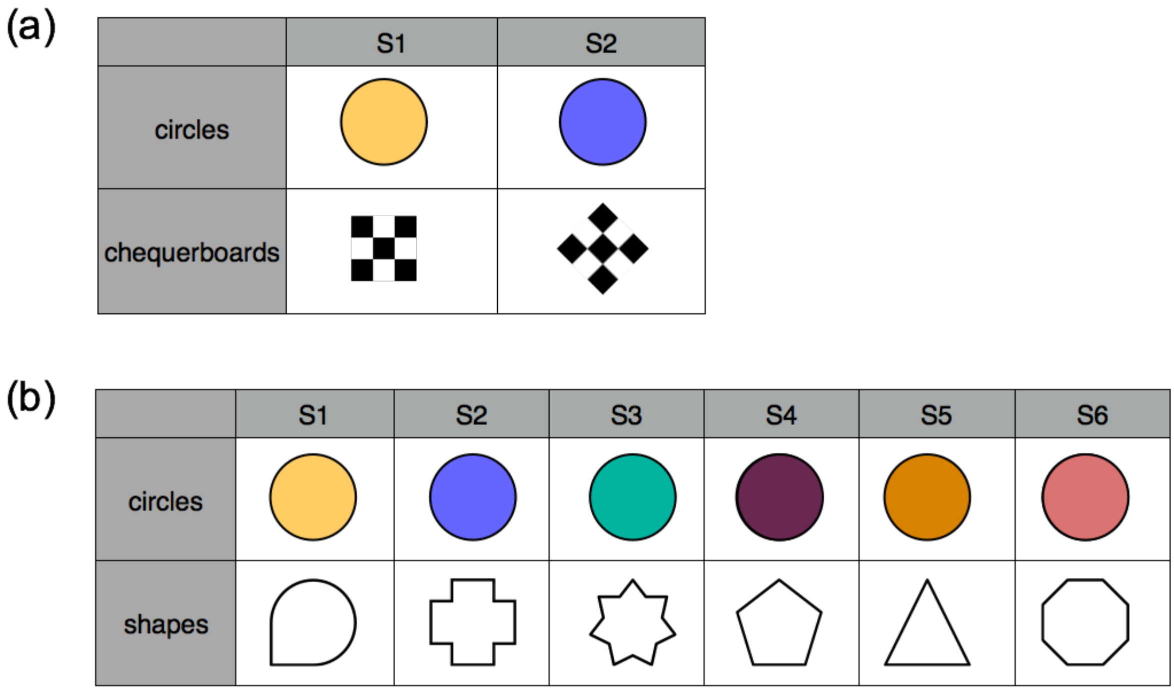

Figure 6.1. Stimulus sets used in (a) Experiment 11 and (b) Experiment 12.

On each trial, a fixation cross (30 × 30 pixels) was presented in the middle of the screen for 500 ms. The stimulus (a chequerboard or a circle) appeared in the centre of the screen (replacing the fixation cross). After 250 ms, the stimulus was replaced with a blank screen. In the go/no-go task, half of the participants had to press the 'c' key (with their left index finger) when the go stimulus was presented; the other half had to press the 'm' key (with their right index finger) when the go stimulus was presented. Participants had to refrain from pressing any key when the no-go stimulus was presented. In the two-choice task, participants had to press the 'c' key (with their left index finger) for one stimulus (e.g. blue circle) and the 'm' key (with their right index finger) for the other (e.g. yellow circle; Figure 6.1). All stimulus-response mappings were counterbalanced. After 1000 ms, the response deadline elapsed and a blank screen was presented for 250 ms.

There was a break after every block. During the break, feedback on mean correct RTs (on go trials in the go/no-go task; or on left keypress trials and right keypress trials in the two-choice task), the number of choice errors (two-choice task only), the number of missed responses (on go trials in the go/no-go task and on all trials in the two-choice task), and the percentage of responses on no-go trials (go/no-go task only) was presented on the screen. After 15 s elapsed, participants had to press any key to start the next block.

EEG data acquisition. The electroencephalogram (EEG) was acquired using 64 Ag/AgCl active electrodes (ActiCap, Brain Products, Munich, Germany) connected to BrainAmp amplifiers (Brain Products, Munich, Germany). The EEG was sampled continuously at 500 Hz with a bandpass of 0.016-100 Hz, the reference was at CZ and the ground was at AFz. There were 62 electrodes on the scalp and one on each earlobe (for the configuration, see Figure 6.2). Electrode impedances were kept below 10 k**Ω.**

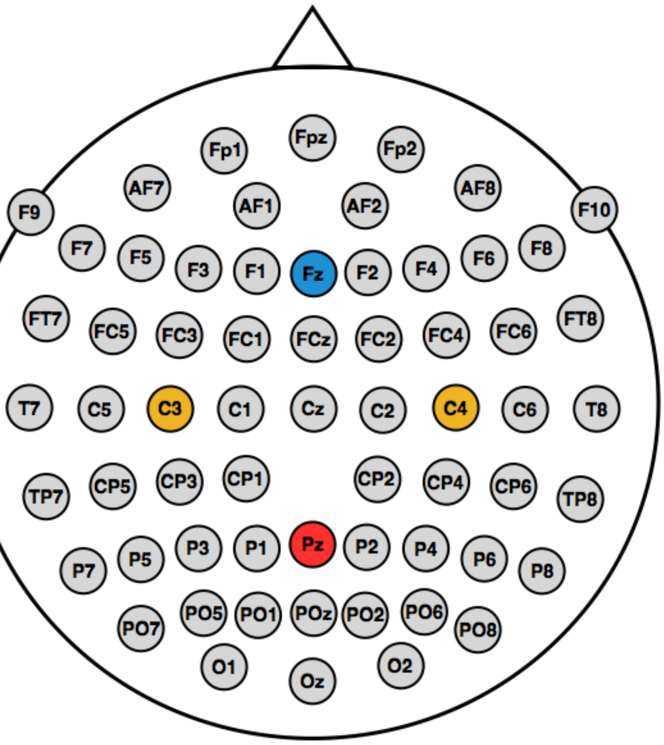

Figure 6.2. Electrode arrangement. The amplitudes and latencies of the highlighted electrodes were subjected to statistical analyses (Fz: stimulus-locked N2 and P3a components; Pz: stimulus-locked P3b component; C3 & C4: response-locked readiness potentials).

Analyses. All raw and processed behavioural and EEG data are deposited in Dropbox (https://www.dropbox.com/sh/emswlmmu941080s/AADj-qKRDQ2oSCffOeO7zVkda?dl=0). Prior to all analyses I subsetted the data; for participants who responded with a left keypress on go trials in the go/no-go task, I analysed only left keypress responses in the two-choice task; for participants who responded with a right keypress response on go trials, I analysed only right keypress responses in the two-choice task. Thus, I could ensure equal trial numbers across each stimulus type (go, no-go, two-choice). This was especially important for the ERP analyses (see *ERP analyses*). For planned and follow-up comparisons, Hedge's g_{av} is the reported effect size measure (Lakens, 2013). Where appropriate, I applied the Huyhn-Feldt correction for violations of sphericity (unadjusted degrees of freedom are reported). For all planned comparisons, Bayes factors were calculated with the BayesFactor package in R, using the default prior of 0.707 (Morey, Rouder, et al., 2015). A Bayes factor of greater than 3 is often taken as substantial evidence for the alternative hypothesis, whereas a Bayes factor of less than 0.33 is taken as substantial evidence for the null hypothesis (Dienes, 2011).

Behavioural analyses. All data processing and analyses were completed using R (R Development Core Team, 2014). I performed analyses of variance on correct RTs on go trials and the probability of responding on no-go trials [*p*(respond|no-go)] in the go/no-go task; and on correct RTs and the probability of correct responses [*p*(correct)] in the two-choice task. The probability of missed responses were very low (≤ 0.01) so were not analysed further. Reaction times were analysed as a function of 'part' (there were 3 blocks per part: 'beginning' = first three training blocks of each task; 'end' = final three training blocks of each task) and stimulus (go, two-choice). Note, *p*(correct) refers here to the probability that participants executed the correct response from two alternatives (i.e. left keypress vs. right keypress), consequently I could not include the 'stimulus' factor here. Similarly, I could not include the 'stimulus' factor in the *p*(respond|no-go) analyses because, by their very nature, participants responded on all go trials.

ERP analyses. The EEG was filtered off-line with a 20 Hz low-pass (24 dB/oct) and a 0.1 Hz high pass (12 dB/oct) filter. The raw EEG data was visually inspected for non eye-movement artifacts: muscle activity, drifts, and large amplitudes in the alpha frequency band. To correct eye blink artifacts, I ran an Independent Component Analysis on the whole data (Infomax ICA, Bell & Sejnovski, 1995), implemented in Vision Analyzer (BrainProducts, Munich, Germany). Sixty-three ICA components were obtained from every participant's EEG (the same as the number of electrodes submitted to the ICA). An average of 5 components with characteristic eye-blink and eye-movement topographies and time-courses were excluded. Following this ICA-based artifact subtraction, the EEG was re-referenced to the linked ears. For the stimulus-locked ERPs, segments were cut from -100 ms preceding stimulus presentation to 600 ms following stimulus presentation, time-locked to the presentation of the stimulus and baseline corrected[50] relative to the average amplitude of the 100 ms preceding the stimulus presentation. For the response-locked ERPs, segments were cut from -1000 ms preceding the response to 200 ms following the response, and the first 100 ms of that segment (from -1000 ms to -900 ms) was used as a baseline.

Consistent with the behavioural analyses, I excluded from all ERP analyses segments associated with errors and all segments in blocks 7 to 12 as these were not in the 'beginning' or in the 'end' parts. The resulting segments were visually inspected for residual ocular, muscle, movement, and other artifacts. Segments containing such artifacts were removed. The remaining EEG segments were averaged for every participant as a function of part (beginning, end)[51] and stimulus (go, no-go, two-

---

[50] Baseline correction is the subtraction of the mean amplitude of the baseline interval from each time-sample following the baseline. The baseline intervals used in the present study are consistent with previous research (see, e.g. Elchlepp, Lavric, et al., 2016).

[51] The data was analysed as 'beginning' and 'end' parts to ensure that there were sufficient trial numbers for the ERP analyses. Note that I did not analyse the middle three blocks of each task as the data in this middle interval was expected to fall somewhere between the beginning and end phases. The beginning/end approach is consistent with previous ERP studies that have investigated the effects of go/no-go training (e.g. Manuel, Grivel, et al., 2010).

choice). The segments were not separated until after the artifact rejection to ensure that the experimenter was blind to the segment condition during the artifact rejection process. As can be seen in Table 6.2, all stimulus-locked averages in Experiment 11 contained ≥ 88 segments meaning that there were sufficient trials to examine the N2 and P3 components (for guidelines, see Rietdijk, Franken & Thurik, 2014)[52]. All descriptive and inferential statistics were computed using R (R Development Core Team, 2014). Analyses were restricted to the sites Fz (N2 and frontal P3a) and Pz (parietal P3b), although the amplitudes at the other electrodes are represented in the topographic maps represented in Figure 6.4. The selection of these electrodes for the N2, P3a, and P3b components was on the basis of previous research (see, e.g. Smith, Smith, et al., 2009).

	Stimulus-locked				**Response-locked**			
	Beginning		***End***		***Beginning***		***End***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Experiment 11								
Go	101	12	94	14	102	12	95	14
No-go	97	11	89	12	-	-	-	-
Two-choice	96	11	88	15	95	15	85	16
Experiment 12								
Go	97	21	96	22	98	20	100	20
No-go	86	23	87	22	-	-	-	-
Two-choice	85	20	84	21	83	21	83	21

Table 6.2. Overview of the average number of artifact-free ERP segments contributing to the stimulus-locked and response-locked waveform grand averages as a function of experiment (Experiment 11, Experiment 12), stimulus (go, no-go, two-choice) and part (beginning, end). *M* = mean; *SD* = standard deviation.

Some studies in the go/no-go literature have used the peak minima or peak maxima to measure the amplitudes of the N2 and P3 components, respectively (e.g. Bokura, Yamaguchi, Kobayashi, 2001; Falkenstein, Hoorman, et al., 1999; Niewenhuis, Yeung, et al., 2003; Smith, Johnstone, & Barry, 2008) whereas others have used the mean amplitude in specified N2 and P3 time-windows (e.g. Bruin, Wijers, van Staveren, 2001; Eimer, 1993). I adopt the latter approach as peak amplitude measures have the following shortcomings (for an extended discussion of these ideas see Luck, 2005): (1) one must use relatively large time windows in order to capture the peak minima/maxima. Consequently, when the minimum or maximum voltage is measured, the rising or falling edge of an overlapping component could be measured rather than the desired component; (2) peak amplitude uses a single time point to

---

52. To avoid excessive fatigue in my participants, I aimed to keep the duration of the experiment at around 1 hr, excluding the EEG preparation procedures. Furthermore, I did not provide any instructions about restricting eye movements or eye-blinks as research has demonstrated that spontaneous eye blinking is related to inhibitory control function (e.g. Colzato, van den Wildenberg, van Wouwe, Pannebakker, & Hommel, 2009). This meant that I had 120 trials per cell in Experiment 11 and 126 trials per cell in Experiment 12 prior to EEG preprocessing. I had previously established that each participants' ERP average should contain at least 25 trials (for a similar criterion, see Elchlepp, Lavric, et al., 2016). After applying the exclusion criteria, there were 11 participants in Experiment 11 and 8 participants in Experiment 12 that had less than 25 segments remaining and were therefore removed and replaced.

measure a component that lasts several hundred milliseconds making it sensitive to noise; (3) peak amplitude will be artificially increased when the noise level is higher (for example, when there are fewer trials contributing to the average); and (4) peak amplitude is a non-linear measure meaning that the peak amplitudes may not correspond to the grand average waveforms. In contrast, the mean average amplitude has several (relative) advantages: (1) a narrower time window can be used; (2) mean amplitude is less sensitive to high-frequency noise as a range of time points are used; (3) mean amplitude measures do not become biased when noise levels increase (consequently, one can compare mean average amplitudes from waveforms with differing numbers of trials and/or differing lengths of time window); (4) one can measure amplitudes in the N2 and P3 ranges, even when there is no clear peak in the participant's data; and (5) mean amplitude is a linear measure meaning that the analysed amplitude is represented in the grand average waveform. However, note that analyses of the peak amplitudes (not presented) revealed a broadly similar pattern of results.

Peak latencies were first computed within the following time-windows: 180-350 ms (N2), 300-500 ms (P3a), and 250-450 ms (P3b). These values were based on visual inspection of the data and are broadly consistent with the time ranges identified in the ERP literature (for an overview, see e.g. Rietdijk, Franken, et al., 2014). Although several of the aforementioned shortcomings of peak amplitude measures also apply to peak latencies, there are few good alternatives for latency analyses. However, I did take some steps to mitigate these shortcomings. First, by subsetting the two-choice trials I ensured that the trial numbers were similar across the go, no-go, and two-choice stimuli (for ERP segment numbers, see Table 6.2). Thus, I could be sure that noise in the peaks for these stimuli were similar and could therefore draw comparisons accordingly. Second, as participant's waveforms are typically too noisy to identify each peak I used the 'jackknifing' technique (Miller, Patterson, & Ulrich, 1998). Thus, I computed the grand average of all but one of the participants' waveforms, omitting each participant's data in turn. Next, I selected the greatest positive or negative amplitudes within the three manually specified time-windows. Peak detections were performed separately for each sub-average, stimulus, and part. Mean amplitude time-windows were defined on the basis of the peak latencies values and visual inspection of the grand average waveforms[53]. The mean average amplitude and peak latencies were computed within the following time windows: 180-280 ms (N2), 300-400 ms (P3a), and 270-370 ms (P3b).

For each mean average amplitude and peak latency, I computed an ANOVA to compare the go and no-go stimuli in the go/no-go task with part (beginning; end) and stimulus (go; no-go) as within-subjects factors; and an ANOVA to compare the go and two-choice stimuli with part and stimulus (go;

---

[53] The peak latencies were first obtained and this peak latency value was rounded to the nearest 10 ms. The mean amplitude windows were defined by subtracting 50 ms from this value and adding 50 ms to this value to produce a 100 ms window that captured the peak. 100 ms windows for these components are common; however, to ensure that these windows corresponded to and sufficiently captured the component of interest each participants raw data was visually inspected within these windows. However, no further amendments were made on the basis of this final inspection stage.

two-choice) as within-subject factors. The former analysis examines the effects of task practice on processes relating to response execution and response inhibition, whilst the latter analysis compares the effects of task practice on response execution in the go/no-go and two-choice tasks. Before testing for significance, the F values for the jack-knifed peak latencies were adjusted according to $F_c = F/(n - 1)^2$ and follow-up T values were adjusted according to $t_c = t(n-1)$, where n is the number of observations per cell (for a proof of these adjustments, see Ulrich & Miller, 2001). I followed up on any null findings between the beginning and end of training using Bayesian analyses. Note, however, that it was not possible to compute Bayes factors for the peak latencies because no Bayesian equivalent has been developed yet for the jack-knifing procedure.

Finally, response-locked waveforms were analysed for the go and two-choice stimuli. These analyses were particularly useful for the interpretation of stimulus-locked effects: any differences in the amplitudes of the P3a or P3b sub-components between the beginning and the end of training for the go and two-choice stimuli could otherwise be attributed to temporally overlapping differences in motor-related activity between these stimuli. To isolate response-related activity from non-lateralised ERP defections arising from components occurring at approximately the same time as response preparation and/or execution (e.g. the P3a and P3b), I used the averaging procedure outlined by Coles (1989). I subtracted the amplitude in an electrode positioned over the motor cortex ipsilateral to the response hand (C4 for right-hand responses and C3 for left-hand responses; see Figure 6.2) from the amplitude in an electrode contralateral to the response hand (C3 for right-hand responses and C4 for left-hand responses). I averaged these difference waveforms for the left-hand and right-hand responses. As I subsetted the two-choice data (see *Analyses* above), these averages were computed between-subjects for both the go and two-choice stimuli. This two-step procedure should remove all sensory, cognitive, and non-lateralised motor potentials as the corresponding stimuli were not lateralised in the experiments presented in this chapter. The mean amplitude in the 200 ms preceding response execution was submitted to an ANOVA with stimulus (go, two-choice) and part (beginning, end) as within-subjects factors.

Results

Behavioural Results

As expected, analyses of the RTs revealed that responding was reliably slower in the two-choice task (359 ms) than in the go/no-go task (329 ms; $p = 0.006$; Tables 6.3 & 6.4). However, there was no reliable difference between RTs at the beginning and the end of training ($p = 0.526$; Table 6.4). The two-way interaction between task and part was also not reliable ($p = 0.317$; Table 6.5), and Bayesian analyses supported the null hypothesis of no difference between the beginning and the end of training ($BF = 0.24$).

	RT				***p*(correct)**				***p*(respond|no-go)**			
	Beginning		***End***		***Beginning***		***End***		***Beginning***		***End***	
	M	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***	***M***	***SD***
Experiment 11												
Go	331	47	326	42	-	-	-	-	-	-	-	-
No-go	-	-	-	-	-	-	-	-	0.05	0.05	0.07	0.06
Two-choice	356	43	355	50	0.91	0.07	0.91	0.07	-	-	-	-
Experiment 12												
Go	417	57	382	45	-	-	-	-	-	-	-	-
No-go	-	-	-	-	-	-	-	-	0.10	0.07	0.07	0.07
Two-choice	473	65	430	58	0.85	0.13	0.87	0.10	-	-	-	-

Table 6.3. Overview of the behavioural data in Experiment 11 and Experiment 12. Average reaction times (RT) for go and two-choice stimuli, the probability of an accurate response to the two-choice task [*p*(correct)] and the probability of responding to the no-go stimuli in the go/no-go task [*p*(respond|no-go)] as a function of stimulus (go, no-go, two-choice) and part (beginning, end). Accuracy is the ratio of correct two-choice trials to the number of correct and incorrect two-choice trials (missed trials are excluded). *M* = mean; *SD* = standard deviation.

In addition to RTs, I also analysed the *p*(correct) in the two-choice task and the *p*(respond|no-go) in the go/no-go task (Table 6.3). These analyses revealed no reliable difference in choice accuracy between the beginning of training and the end of training, $t(31) = 0.96$, $p = 0.347$, $g_{av} = 0.18$, $BF = 0.29$, and that the *p*(respond|no-go) reliably increased between the beginning of training and the end of training, $t(31) = -2.70$, $p = 0.011$, $g_{av} = -0.31$, $BF = 4.03$. Taken together, these results provide no evidence that task practice improved behavioural task performance in the go/no-go or two-choice tasks. If anything, no-go performance became worse (not better) with training.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
Experiment 11							
Stimulus	1	31	22800.14	109985.81	6.43	**0.017**	0.081
Part	1	31	297.15	16941.38	0.54	0.466	0.001
Stimulus by part	1	31	95.54	5063.86	0.58	0.450	< 0.001
Experiment 12							
Stimulus	1	39	108167.70	117473.19	35.91	**< 0.001**	0.177
Part	1	39	58748.40	28201.23	81.24	**< 0.001**	0.105
Stimulus by part	1	39	618.61	14279.88	1.69	0.201	0.001

Table 6.4. Overview of the Analyses of Variance of the reaction time data for the go and two-choice stimuli in Experiment 11 and Experiment 12 with stimulus and part as within-subjects factors. ps < 0.05 are highlighted in bold.

ERP Results

Figures 6.3 present the stimulus-locked grand average waveforms, Figure 6.4 presents the topographies, and Figure 6.5 presents the mean average amplitudes and the peak latencies in the analysed time windows. Tables 6.5 and 6.6 present an overview of the inferential statistics.

N2 range.

Mean amplitudes. Consistent with previous research, the mean amplitude in the N2 range was more negative for the no-go stimuli than for the go stimuli in the go/no-go task ($p < 0.001$; Table 6.5). Consistent with the behavioural data, there was no reliable change in the mean amplitude between the beginning and the end of training ($p = 0.744$; Table 6.5) and Bayesian analyses supported the null hypothesis of no difference ($BF = 0.19$). The two-way interaction between stimulus (go, no-go) and part reached significance ($p = 0.043$; Table 6.5); as can be seen in Figure 6.5, amplitudes in the N2 range were slightly *more* negative at the end of training than at the beginning of training for the go stimuli, but were slightly *less* negative at the end of training than at the beginning of training for the no-go stimuli. However, follow-up analyses revealed no reliable main effect of part for the go stimuli, $t(31) = 0.59$, $p = 0.562$, $g_{av} = 0.07$, $BF = 0.22$, or for the no-go stimuli, $t(31) = -1.37$ $p = 0.180$, $g_{av} = -0.16$, $BF = 0.44$.

A comparison of the go and two-choice stimuli revealed no reliable difference in the mean amplitude between these stimuli ($p = 0.083$; Table 6.5). Whilst there was no reliable main effect of part ($p = 0.137$, $BF = 0.54$; Table 6.5), the two-way interaction between stimulus (go, two-choice) and part was significant ($p = 0.005$; Table 6.5). Follow-up analyses confirmed that, unlike for the go stimuli (see above), the mean amplitude for the two-choice stimuli was reliably *less* negative at the end of training than at the beginning of training, $t(31) = -3.85$, $p = 0.001$, $g_{av} = -0.38$, $BF = 55.39$ (see Figures 6.3 & 6.5).

Peak latencies. Analyses of the peak latencies in the N2 range revealed that amplitudes peaked 11 ms earlier for the no-go stimuli than for the go stimuli ($p_c = 0.019$; Table 6.6; Figure 6.5). However, analyses in the go/no-go task showed no reliable difference between the peak latencies in the beginning and the end of training ($p_c = 0.442$; Table 6.6) and no reliable two-way interaction between stimulus and part ($p_c = 0.685$; Table 6.6). The comparison of the go and two-choice stimuli revealed no reliable main effects or interactions (p_cs ≥ 0.717; Table 6.6).

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen. η^2***
N2 range: Mean amplitude							
Go vs. No-go							
Stimulus	1	31	90.43	119.44	23.47	**< 0.001**	0.071
Part	1	31	0.35	99.88	0.11	0.744	< 0.001
Stimulus by part	1	31	3.51	24.33	4.48	**0.043**	0.003
Go vs. Two-choice							
Stimulus	1	31	21.49	207.21	3.21	0.083	0.018
Part	1	31	4.97	66.27	2.33	0.137	0.004
Stimulus by part	1	31	12.34	42.23	9.06	**0.005**	0.011
P3a range: Mean amplitude							
Go vs. No-go							
Stimulus	1	31	202.11	328.26	19.09	**< 0.001**	0.053
Part	1	31	21.59	219.51	3.05	0.091	0.006
Stimulus by part	1	31	1.76	63.52	0.86	0.361	< 0.001
Go vs. Two-choice							
Stimulus	1	31	28.18	333.65	2.62	0.116	0.010
Part	1	31	1.11	135.27	0.25	0.618	< 0.001
Stimulus by part	1	31	5.13	57.89	2.75	0.107	0.002
P3b range: Mean amplitude							
Go vs. No-go							
Stimulus	1	31	530.06	424.59	38.70	**< 0.001**	0.105
Part	1	31	38.12	284.25	4.16	0.050	0.008
Stimulus by part	1	31	3.51	45.46	2.39	0.132	0.001
Go vs. Two-choice							
Stimulus	1	31	93.28	386.13	7.49	**0.010**	0.022
Part	1	31	27.15	251.69	3.34	0.077	0.006
Stimulus by part	1	31	0.83	35.49	0.72	0.402	< 0.001

Table 6.5. Overview of the Analyses of Variance of the stimulus-locked mean amplitudes in Experiment 11 for the go vs. no-go comparison with part (beginning, end) and stimulus (go, no-go) as a within-subjects factors; and for the go vs. two-choice comparison with part (beginning, end) and stimulus (go, two-choice) as within-subjects factors. *p*s < 0.05 are highlighted in bold.

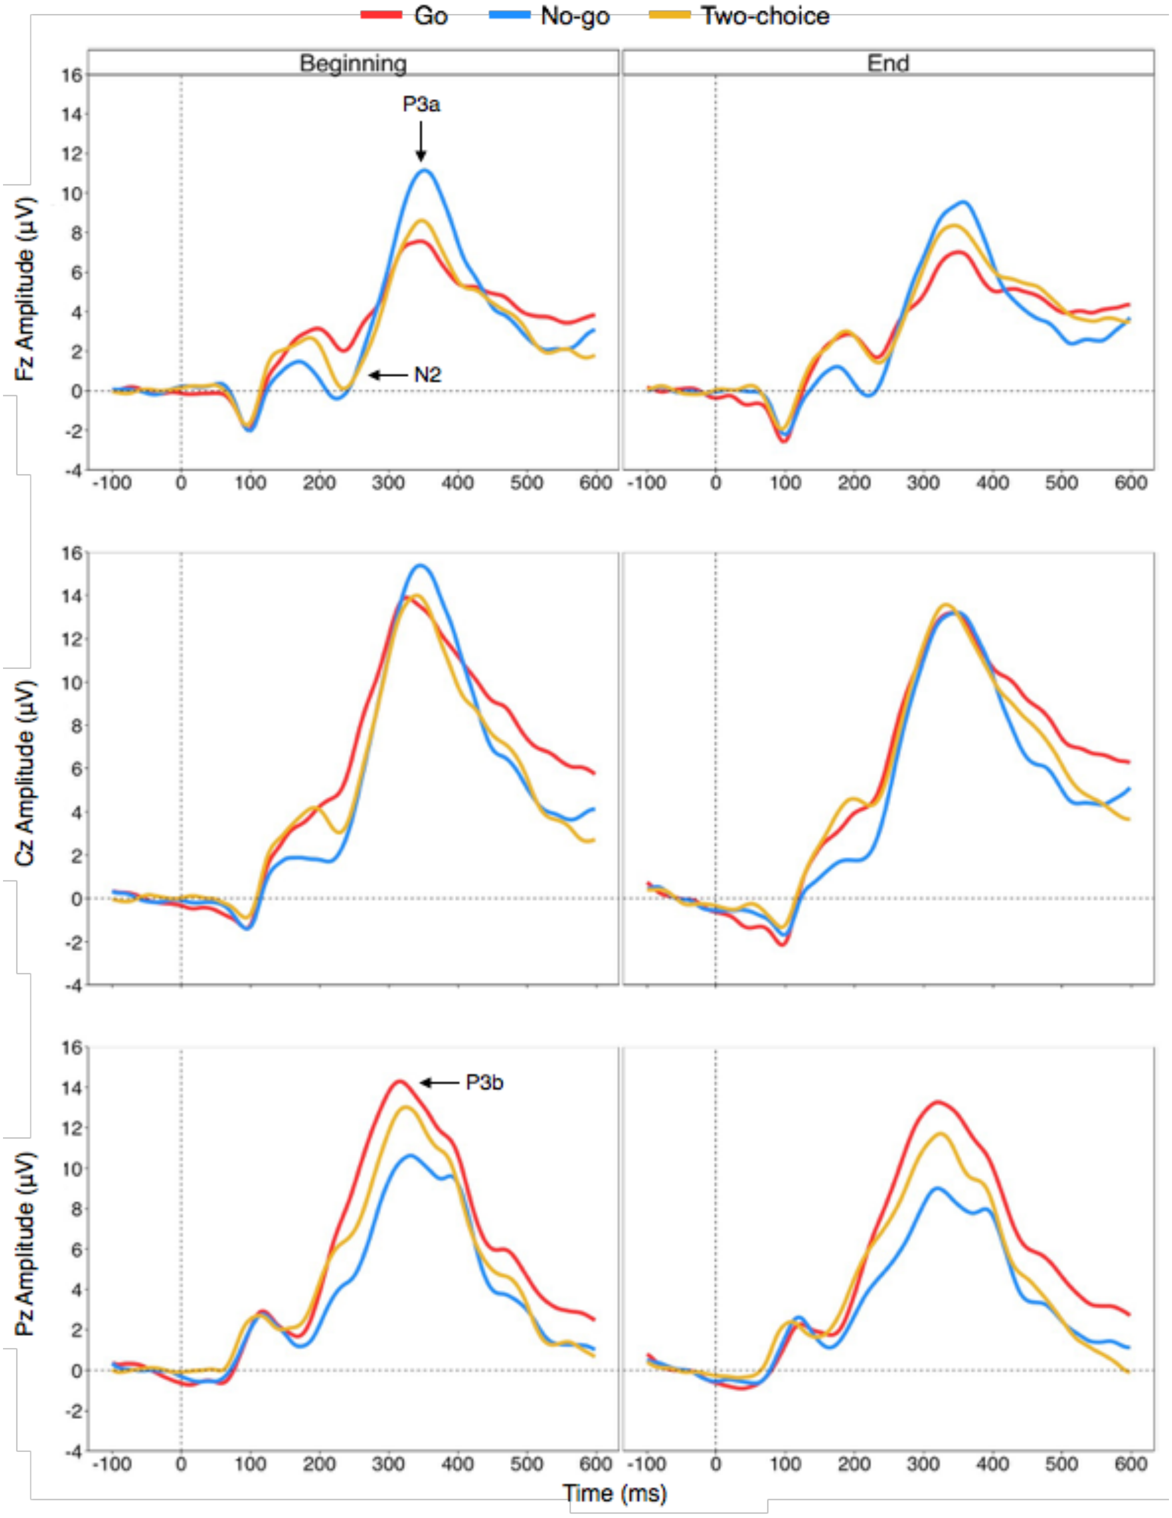

Figure 6.3. Stimulus-locked waveforms for the go, no-go, and two-choice stimuli in Experiment 11 as a function of part (beginning, end) in electrodes Fz, Cz, and Pz. The stimulus was presented at 0 ms.

P3a range.

Mean amplitudes. Consistent with the go/no-go ERP literature, the mean amplitude in the P3a range was more positive for the no-go stimuli than the go stimuli ($p < 0.001$; Table 6.5; Figures 6.3 & 6.5). However, there was no reliable main effect of part ($p = 0.091$; Table 6.6) and no reliable two-way interaction between stimulus and part ($p = 0.361$; Table 6.6). However, the Bayesian analyses only revealed anecdotal support for the null hypothesis of no difference between the beginning and the end of training ($BF = 0.73$).

A comparison of the go and two-choice stimuli revealed no reliable main effect of stimulus ($p = 0.116$; Table 6.5) or part ($p = 0.618$; Table 6.5). The two-way interaction between stimulus and part was also not reliable ($p = 0.107$; Table 6.5). The Bayesian analyses supported the null hypothesis of no difference between the beginning and the end of training ($BF = 0.21$).

Peak latencies. There were no reliable main effects or interactions in the peak latencies in the P3a range in the comparison of the go and no-go stimuli (p_cs ≥ 0.300; Table 6.6) or in the comparison of the go and two-choice stimuli (p_cs ≥ 0.491; Table 6.6).

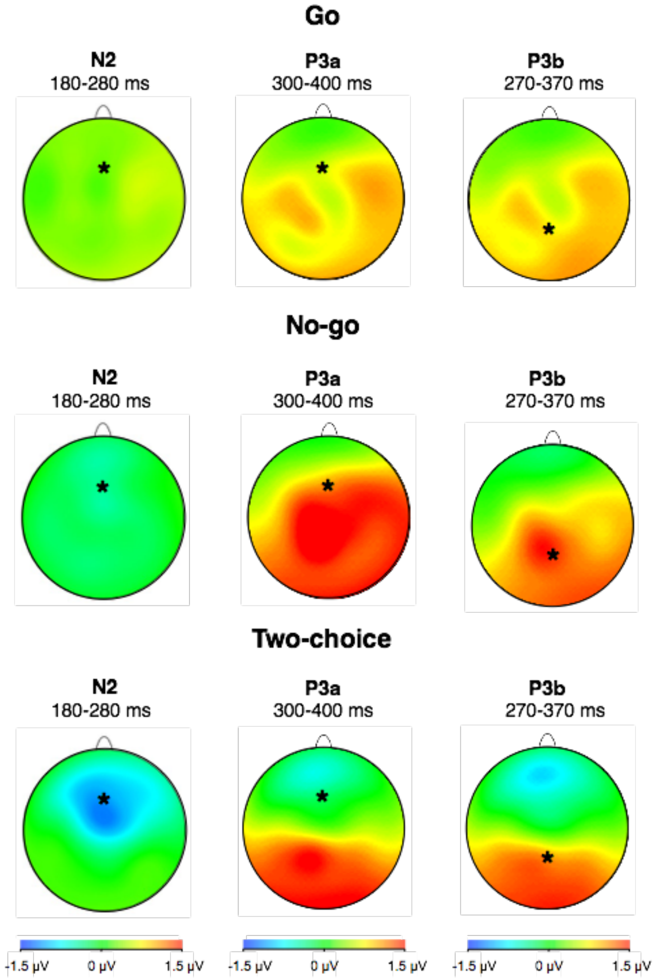

Figure 6.4. Topographies for the beginning minus end difference as a function of stimulus (go, no-go, two-choice) and analysed time-window (N2: 180-280 ms; P3a: 300-400 ms; P3b: 270-370 ms) for Experiment 11. * denotes the location of the analysed electrode in each time-window (N2 & P3a: Fz; P3b: Pz).

P3b range.

Mean amplitudes. The mean amplitude in the P3b range was more positive for the go stimuli than for the no-go stimuli ($p < 0.001$; Table 6.5; see also Figures 6.3 & 6.5). Thus, the decision to analyse the P3a and P3b separately appears justified given the differential pattern of amplitudes between the go and no-go stimuli across these sub-components (cf. amplitudes were more positive for the no-go stimuli than for the go stimuli in the P3a range). As can be seen in Figure 6.3, the mean amplitudes in the P3b range were numerically less positive at the end of training than at the beginning of training. However, this difference was only marginally significant ($p = 0.050$; Table 6.5). The corresponding Bayes factor was 1.16[54]. The two-way interaction between stimulus and part was not reliable (p = 0.132; Table 6.5).

A comparison of the go and two-choice stimuli revealed that the mean amplitude was more positive for the go stimuli than for the two-choice stimuli (p = 0.010; Table 6.5; this difference is discussed further in the *Response-locked analyses* section). However, there was no reliable main effect of part (p = 0.077; Table 6.5) and no reliable two-way interaction between stimulus and part (p = 0.402; Table 6.5). The Bayes factor for the difference between the beginning and the end of training was 0.83.

Peak latencies. Amplitudes in the P3b range peaked 8 ms later for the no-go stimuli than for the go stimuli (p_c = 0.011; Table 6.6; Figure 6.5). However, there were no other reliable main effects or interactions in the comparison of the go and no-go stimuli (p_cs ≥ 0.087; Table 6.6) or in the comparison of the go and two-choice stimuli (p_cs ≥ 0.319; Table 6.6).

---

54. Note that this difference is in the opposite direction to the predicted amplitude difference based on the ERP pilot study (Appendix K) and also the findings previously reported in the literature (Benikos, Johnstone, et al., 2013; Benikos, unpublished PhD thesis). Given the absence of reliable behavioural training-effects, it is possible that the observed effects in the present study were due to fatigue rather than learning. Consistent with this idea, Kato and colleagues showed that amplitudes of the P3 component became less positive with the time spent on the task (Kato, Endo, et al., 2009). Kato, Endo, et al. attributed this to the effects of fatigue which impaired the intensity of response-related processes. Another, potentially compatible, explanation is that the observed P3b difference reflects habituation to the visual stimuli. Consistent with this idea, it has been shown that habituation results in less positive P3 amplitudes (Ravden & Polich, 1998). This could account for differences between the results in Experiment 11 and previous research as there were only 2 stimuli used per task in Experiment 11 (cf. 4 shapes per task in the ERP pilot study and 8 shapes were used by Benikos, Johnstone, et al., 2013).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	F_c	p_c
N2: Peak latency							
Go vs. No-go							
Stimulus	1	31	3938.28	20.72	5892.57	6.13	**0.019**
Part	1	31	371.28	19.72	583.69	0.61	0.442
Stimulus by part	1	31	42.78	8.22	161.37	0.17	0.685
Go vs. Two-choice							
Stimulus	1	31	81.28	29.72	84.79	0.09	0.768
Part	1	31	34.03	20.97	50.31	0.05	0.821
Stimulus by part	1	31	47.53	11.47	128.48	0.13	0.717
P3a: Peak latency							
Go vs. No-go							
Stimulus	1	31	1313.28	63.72	638.93	0.66	0.421
Part	1	31	790.03	22.97	1066.27	1.11	0.300
Stimulus by part	1	31	7.03	25.97	8.39	0.01	0.926
Go vs. Two-choice							
Stimulus	1	31	153.13	152.88	31.05	0.03	0.859
Part	1	31	8.00	34.00	7.29	0.01	0.931
Stimulus by part	1	31	512.00	34.00	466.82	0.49	0.491
P3b: Peak latency							
Go vs. No-go							
Stimulus	1	31	1785.03	7.97	6944.12	7.23	**0.011**
Part	1	31	371.28	25.72	447.52	0.47	0.500
Stimulus by part	1	31	2032.03	20.97	3004.14	3.13	0.087
Go vs. Two-choice							
Stimulus	1	31	1200.50	37.50	992.41	1.03	0.317
Part	1	31	136.13	29.88	141.25	0.15	0.704
Stimulus by part	1	31	200.00	14.00	442.86	0.46	0.502

Table 6.6. Overview of the Analyses of Variance of the stimulus-locked peak latencies in Experiment 11 with part (beginning, end) and stimulus (go, no-go, two-choice) as a within-subjects factors. As these analyses were performed on jack-knifed averages, the *F* values were adjusted according to $F_c = F/(n - 1)^2$, where *n* is the number of observations per cell (see *Analyses* section for a detailed explanation of this procedure). p_cs < 0.05 are highlighted in bold.

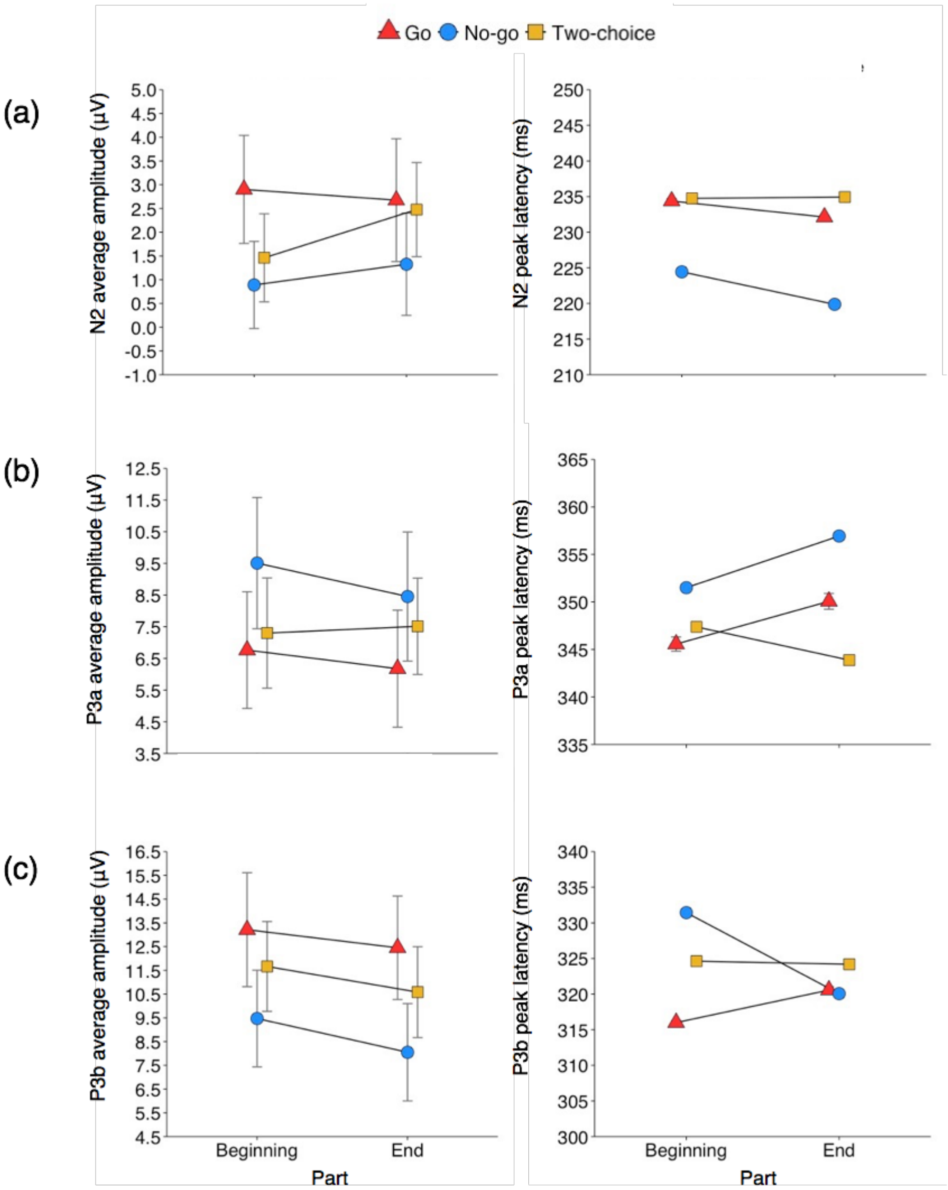

Figure 6.5. (a) Mean average amplitude in the N2 range (180-280 ms) and N2 peak latency (b) mean average amplitude in the P3a range (300-400 ms) and peak P3a latency; and (c) mean average amplitude in the P3b range (270-370 ms) and peak P3b latency, in Experiment 11 as a function of stimulus (go, no-go, two-choice) and part (beginning, end). Error bars are 95% confidence intervals. Note that the confidence intervals are small for the peak latencies due to the jack-knifing procedure (see *ERP analyses*).

Response-locked analyses. Analyses of the mean amplitude in the 200 ms preceding response execution revealed that mean amplitudes were less negative for the go stimuli than for the two-choice stimuli ($p = < 0.001$; Table 6.7; for grand average waveforms, see Figure 6.6). This finding is consistent with previous work suggesting that (lateralised) readiness potentials are sensitive to the level of competition between responses (i.e. the ‘Gratton dip’; Gratton, Coles, Sirevaag, Eriksen, & Donchin, 1988). There was no reliable difference between the beginning and the end of training ($p =$

0.207; Table 6.7) and no reliable two-way interaction between stimulus and part (*p* = 0.214; Table 6.7). However, the Bayesian analyses provided only anecdotal support for the null hypothesis of no difference between the beginning and the end of training (*BF* = 0.40).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Stimulus	1	31	38.11	60.16	19.63	**< 0.001**	0.066
Part	1	31	2.91	54.29	1.66	0.207	0.005
Stimulus by part	1	31	1.09	20.98	1.61	0.214	0.002

Table 6.7. Overview of the Analyses of Variance of the response-locked average amplitudes (-200 to 0 ms) in Experiment 11 with part (beginning, end) and stimulus (go, two-choice) as a within-subjects factors. *p*s < 0.05 are highlighted in bold.

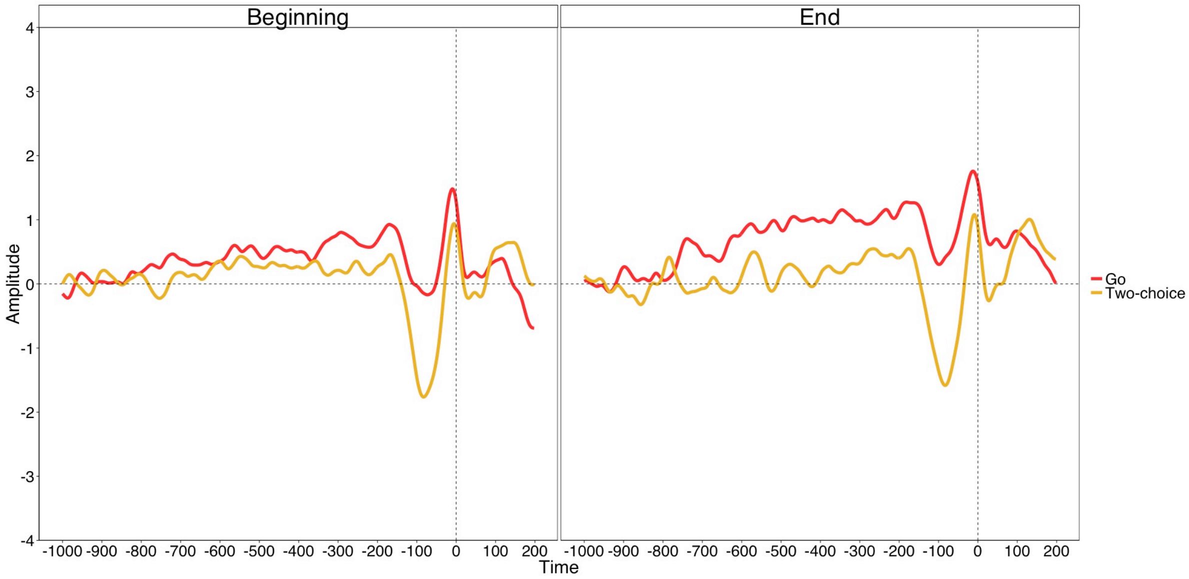

Figure 6.6. Lateralised readiness potentials for the go and two-choice stimuli in Experiment 11 as a function of part (beginning, end). The response was executed at 0 ms.

Interim Conclusions

In Experiment 11, there was no evidence that behavioural performance improved with task practice in the go/no-go and two-choice tasks (for similar results, see Manuel, Grivel, et al., 2010). Furthermore, I found no significant overall main effects of part on mean amplitudes or peak latencies in the N2, P3a, and P3b time-windows. Interestingly, however, further interactions highlighted an effect of training on amplitudes in the N2 range for the two-choice stimuli but not for the go stimuli. Follow-up comparisons confirmed that the mean amplitude in the N2 range was significantly less negative at the end of training than at the beginning of training for the two-choice stimuli.

Although the exact functional correlates of the N2 component are debated (see Folstein & Van Petten, 2008), research broadly converges on the finding that amplitudes in the N2 range are less negative in conditions with low response conflict than in conditions with high response conflict (e.g. Enriquez-Geppert, Konrad, et al., 2010; Nieuwenhuis, Yeung, et al., 2003; Ramautar, Kok, et al., 2004). These findings are consistent with the interpretation that the N2 component reflects control processes over response selection (see, e.g. Huster, Enriquez-Geppert, et al., 2013). This interpretation could is account for the present results; as can be seen in the RT analyses, the present study replicates the common finding that response latencies are shorter in the go/no-go task than in the two-choice task. Several researchers have attributed this difference to greater competition between response alternatives in the two-choice task relative to the go/no-go task (for an overview, see Chapter 5). Thus, if amplitudes in the N2 range are assumed to reflect the intensity of conflict resolution processes, it is perhaps unsurprising that short-term training would result in a larger effect on the N2 component for two-choice stimuli than for go stimuli. It is, however, puzzling that changes in N2 amplitudes were not accompanied by any training-effects on behavioural task performance. For example, responding is usually faster and accuracy higher in conditions where response conflict is low compared with conditions in which response conflict is high (Donkers & Van Boxtel, 2004) so training-induced changes in the level of response conflict should manifest in decreasing RTs and/or increasing accuracy between the beginning and the end of training.

One possibility is that ERP amplitudes afford a more direct measure of cognitive processes and, as such, they may be more sensitive to the effects of practice than behavioural measures. Consistent with this suggestion, previous research has shown that changes in ERPs can precede the emergence of training-effects on behaviour (Atienza, Cantero, & Dominguez-Marin, 2002). If this were the case in Experiment 11, one would expect the corresponding effects of training on behaviour to emerge in longer-term training designs. Another possibility is that the observed change in N2 amplitudes reflects an effect independent of training, for example due to habituation of N2 amplitudes over time. However, it is unclear why habituation would differentially affect two-choice stimuli over go or no-go stimuli. Furthermore, unlike amplitudes in the P3 range, there is no evidence in the literature to suggest that the N2 amplitudes become less negative due to repeated exposure to visual stimuli independent of response execution (Ravden & Polich, 1998). Similarly, although fatigue has been shown to result in less positive P3 amplitudes as a function of time spent on the task, no reliable fatigue-induced change on N2 amplitudes was found (Kato, Endo, et al., 2009). Importantly, participants alternated between performance of the go/no-go and two-choice tasks on a block-by-block basis in Experiment 11 meaning that the effects of fatigue should similarly influence amplitudes in the go/no-go and the two-choice tasks.

Before proceeding, it is important to investigate further why I did not observe any behavioural practice-effects in this experiment. There are at least two possibilities. First, it is possible that the use of a 1:1 stimulus-response mapping made the task too simple (cf. a 2:1 mapping was used in Experiments 9-10, Chapter 5, and in the EEG pilot experiment, Appendix K). If this were the case, the majority of the training effect may have occurred rapidly, within the 'beginning' part of the training phase[55]. Second, it is possible that the requirement to alternate between performance of the go/no-go task and the two-choice task on a block-by-block basis interfered with the practice-effects on performance. After all, research in the task-switching literature indicates that there is a performance decrement when switching between tasks relative to repeatedly performing the same task (see e.g. Monsell, 2003). To investigate these possibilities, I conducted a behavioural control experiment. A detailed description of the procedure and the results of that experiment can be found in Appendix L. To briefly summarise, there were two between-subjects groups: an alternating task order group and a serial task order group. In the alternating task order group, participants alternated between the go/no-go task and the two-choice task on a block-by-block basis (as in Experiment 11). In the serial task order group, participants performed all the training blocks of the go/no-go task followed by all the training blocks of the two-choice task (task order was counterbalanced). Furthermore, four stimuli were used in each task (versus two stimuli per task in Experiment 11) to increase the task difficulty. The results showed that RTs generally became faster between the beginning and the end of training. Crucially, there was no difference in the magnitude of the training effect between the alternating and serial conditions. The presence of a behavioural improvement for go and two-choice stimuli in both the alternating and serial conditions argues against the idea that switching between performance of the go/no-go and two-choice tasks interferes with the effects of training on behaviour. Thus, it seems most likely that the absence of a practice-effect in Experiment 11 was due to task simplicity.

Experiment 12

In Experiment 12, I used the same procedure as in Experiment 11 except that I used twelve stimuli across the experiment (six per task). Based on the RT results in the behavioural control experiment (Appendix L), it was predicted that the increased stimulus-set size would result in a reliable performance improvement between the beginning and the end of training. As the duration of the task was the same as in Experiment 11, any additional differences between the 'beginning' and the 'end' phases of training on the amplitude and/or the latency of the N2 and P3 components can be more confidently attributed to effects of task practice rather than general non-specific effects of stimulus exposure or repeated task practice, such as habituation and/or fatigue.

Method

---

[55] Consistent with this suggestion, there was a reliable main effect of block (1-3) within the 'beginning' part of training was reliable in RTs, $F(2, 62) = 5.63$, $p = 0.006$, *gen.* $\eta^2 = 0.010$. The two-way interaction between block and stimulus (go, two-choice) was not reliable, $F(2, 62) = 0.16$, $p = 0.814$, *gen.* $\eta^2 < 0.001$.

Participants. Forty right-handed volunteers[56] from the University of Exeter participated for monetary compensation (£10) or partial course credit (*M* = 20.15 years, *SD* = 1.78, 31 females). Two participants were removed and replaced as their percentage of correct go trials in the go/no-go task was < 70% (i.e. they missed too many responses and/or pressed an incorrect key on these trials) and seven participants were removed and replaced due to low trial numbers in the ERP averages (see *ERP analyses* above, for the exclusion criteria).

Stimuli, apparatus, and procedures. The stimuli, apparatus, and procedure were identical to Experiment 11, except for the following: the stimuli consisted of six shapes and six coloured circles (see Figure 6.1). Each block consisted of 84 trials and each stimulus within the relevant set (e.g. shapes or circles) was presented 16 times per block.

Analyses. The analyses were identical to Experiment 11, except for the following: I excluded an average of 5 ICA components with characteristic eye-blink and eye-movement topographies and time-courses. The peak latency time-windows were identical to Experiment 11. However, unlike in Experiment 11, amplitudes in the P3a range had a broad plateau meaning that estimating the latency of this component was not appropriate as small amplitude variations on the plateau could result in large latency variations. As it was not possible to obtain an accurate peak latency measure in the P3a window, I instead extracted the latencies for the amplitude rising 50% of the peak maximum (cf. Kiesel, Miller, Jolicoeur, & Brisson, 2008). The time point at which the criterion was reached was determined using linear interpolation (Kiesel, Miller, et al., 2008). Note, I could not examine the falling 50% of the peak maximum as this occurred beyond the end of the segment for some stimuli.

As I could not use the peak latency as a guide to identify the mean amplitude window for the P3a component, I adopted the approach used in Elchlepp, Lavric, et al. (2016) and extended the mean amplitude time-window for the P3a component to 300-500 ms in order to sufficiently capture the broad plateau. Mean amplitudes were computed within the following time windows: 190-290 ms (N2), 300-500 ms (P3a), and 350-450 ms (P3b)[57].

Results

Behavioural Results

Analyses of the RTs revealed that responding was reliably slower to the two-choice stimuli (445 ms) than to the go stimuli (406 ms; $p < 0.001$; Tables 6.3-6.4). The main effect of part was reliable ($p < 0.001$; Table 6.4), reflecting the decrease in RTs between the beginning of training and the end of training. There was no reliable two-way interaction between stimulus and part ($p = 0.201$; Table 6.4).

---

[56] Due to the increase in the number of stimuli, I required more participants in Experiment 12 than in Experiment 11 to ensure that the task, stimulus, and response mappings were fully counterbalanced.

[57] Note that although I make no direct comparisons between components or indeed experiments in these analyses, mean average amplitude measures do not become biased when noise levels increase. Thus, the use of a shorter window for this component in Experiment 11 should not differentially bias the measurement compared with the present experiment.

Analyses of no-go performance revealed that the probability of responding reliably decreased between the beginning of training and the end of training, $t(39) = 2.67$, $p = 0.011$, $g_{av} = 0.33$, $BF = 3.80$. However, analyses of choice accuracy in the two-choice task revealed no reliable main effect of part reliable, $t(39) = - 0.72$, $p = 0.478$, $g_{av} = - 0.10$, $BF = 0.22$. Thus, unlike in Experiment 11, go and stop performance benefitted from task practice. This suggests that the increased stimulus-set used in this experiment was effective in increasing the magnitude of the training-effects on behavioural performance.

ERP Results

Figure 6.7 presents the stimulus-locked grand average waveforms, Figure 6.8 presents the topographies, and Figure 6.9 presents the mean average amplitudes and the peak latencies for the N2, P3a, and P3b components; Tables 6.8 and 6.9 present an overview of the inferential statistics for the mean average amplitudes and the peak latencies, respectively.

N2 range.

Mean amplitudes. As expected, the mean amplitude in the N2 range was reliably more negative for the no-go stimuli than for the go stimuli in the go/no-go task ($p < 0.001$; Table 6.8). Importantly, the main effect of part was reliable ($p < 0.001$; Table 6.8); the N2 amplitude was less negative at the end of training than at the beginning of training (Figures 6.7 and 6.9). The two-way interaction between part and stimulus was not reliable ($p = 0.226$; Table 6.8).

A comparison of the go and two-choice stimuli showed that the mean amplitude in the N2 range was more negative for the two-choice stimuli than for the go stimuli ($p = 0.036$; Table 6.8). Consistent with the go/no-go analyses, the mean amplitude was reliably less negative at the end of training than at the beginning of training ($p = 0.001$; Table 6.8). However, the two-way interaction between stimulus and part was also reliable ($p = 0.005$; Table 6.8). As can be seen in Figure 6.7, the mean amplitude difference between the beginning and the end of training was numerically larger for the go stimuli than for the two-choice stimuli. Follow-up analyses showed that the main effect of part was reliable for the go stimuli, $t(39) = -4.51$, $p < 0.001$, $g_{av} = -0.38$, $BF = 395.12$, and was marginally significant for the two-choice stimuli, $t(39) = -1.90$, $p = 0.064$, $g_{av} = -0.16$, $BF = 0.87$.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
N2 range: Mean amplitude							
Go vs. No-go							
Stimulus	1	39	29.12	59.47	19.09	**< 0.001**	0.015
Part	1	39	57.21	124.71	17.89	**< 0.001**	0.030
Stimulus by part	1	39	1.58	40.79	1.51	0.226	0.001
Go vs. Two-choice							
Stimulus	1	39	17.96	147.66	4.74	**0.036**	0.009
Part	1	39	37.84	109.49	13.48	**0.001**	0.019
Stimulus by part	1	39	7.12	30.61	9.07	**0.005**	0.004
P3a range: Mean amplitude							
Go vs. No-go							
Stimulus	1	39	16.55	128.08	5.04	**0.031**	0.007
Part	1	39	58.83	294.38	7.79	**0.008**	0.026
Stimulus by part	1	39	1.35	31.02	1.70	0.120	0.001
Go vs. Two-choice							
Stimulus	1	39	6.19	277.98	0.87	0.357	0.003
Part	1	39	99.21	223.93	17.28	**0.002**	0.043
Stimulus by part	1	39	1.27	71.13	0.70	0.409	0.001
P3b range: Mean amplitude							
Go vs. No-go							
Stimulus	1	39	316.87	366.81	33.69	**< 0.001**	0.067
Part	1	39	8.47	404.73	0.82	0.372	0.002
Stimulus by part	1	39	14.13	77.89	7.08	**0.012**	0.003
Go vs. Two-choice							
Stimulus	1	39	9.96	395.16	0.98	0.328	0.002
Part	1	39	73.21	249.83	11.43	**0.002**	0.017
Stimulus by part	1	39	3.56	119.13	1.16	0.287	0.001

Table 6.8. Overview of the Analyses of Variance of the stimulus-locked mean amplitudes in Experiment 12 with part (beginning, end) and stimulus (go, no-go, two-choice) as within-subjects factors. *p*s < 0.05 are highlighted in bold.

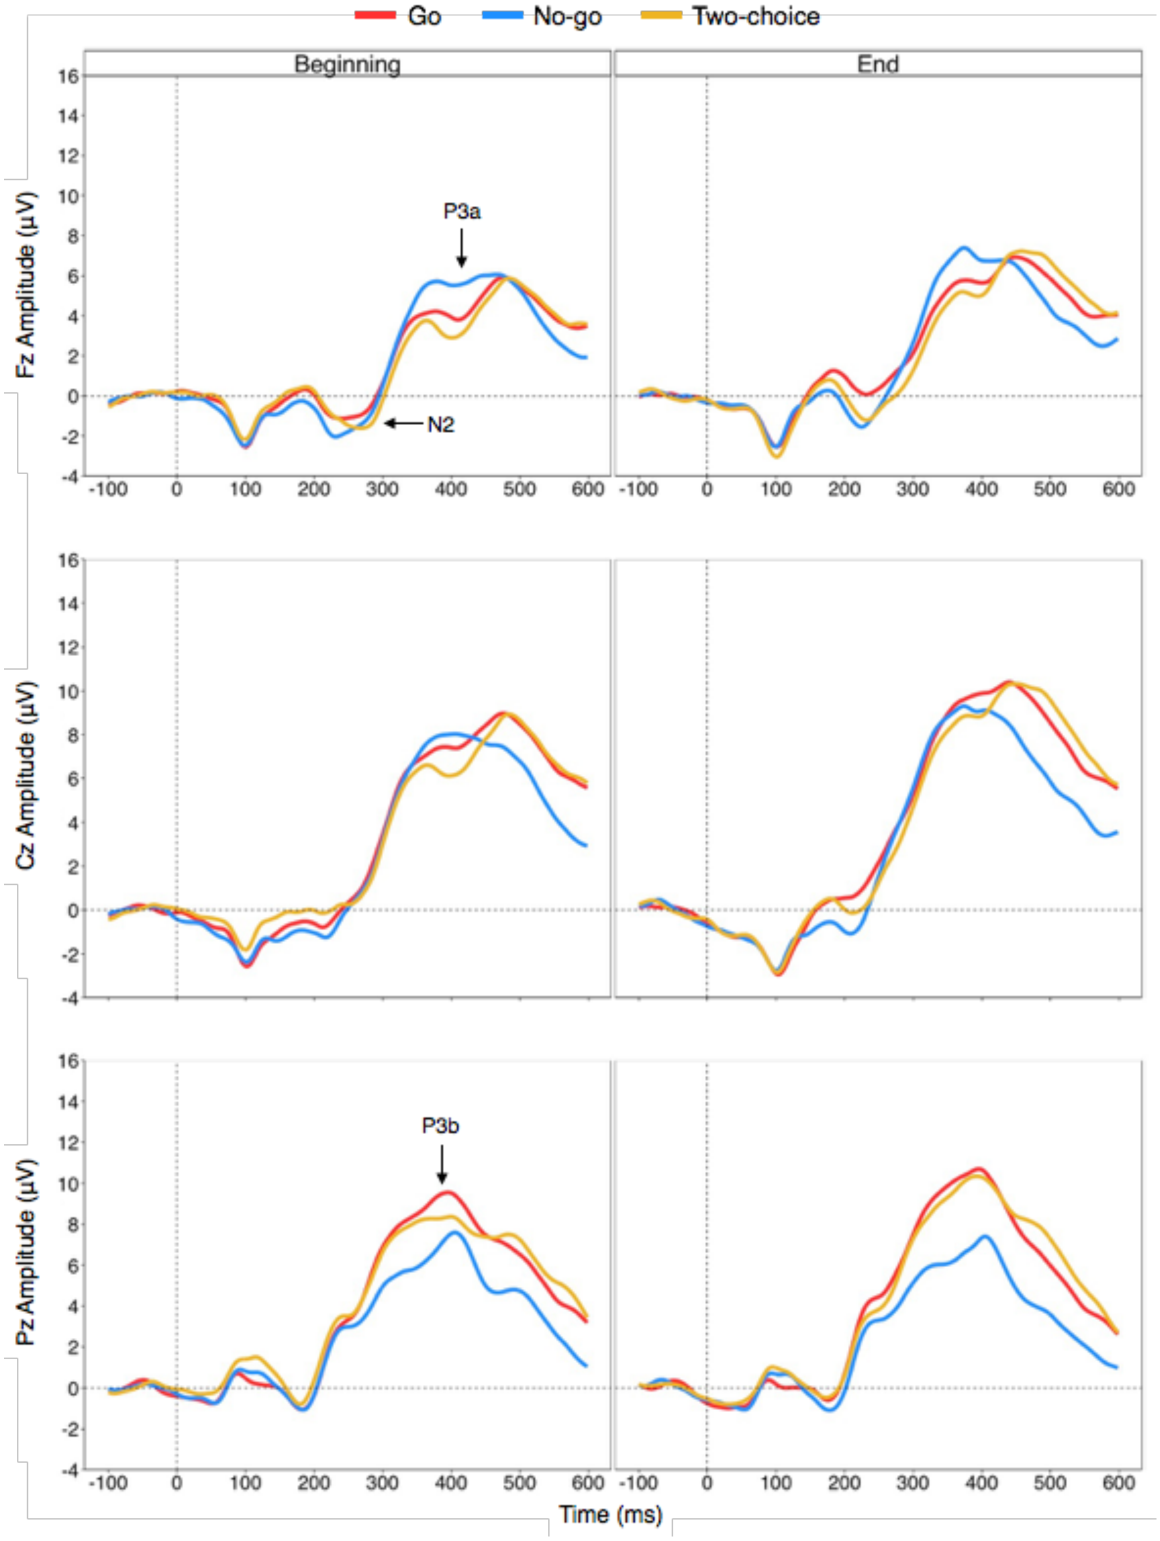

Figure 6.7. Stimulus-locked waveforms for the go, no-go, and two-choice stimuli in Experiment 12 as a function of part (beginning, end) in electrodes Fz, Cz, and Pz. The stimulus was presented at 0 ms.

Peak latencies. A comparison of the go and no-go stimuli revealed no main effect of stimulus ($p_c = 0.348$; Table 6.9). Although amplitudes peaked numerically earlier at the end of training than at the beginning of training (Figure 6.9), this difference was not reliable ($p_c = 0.414$; Table 6.9). There was also no reliable two-way interaction between stimulus and part ($p_c = 0.921$).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*F_c*	*p_c*
N2: Peak latency							
Go vs. No-go							
Stimulus	1	39	2739.03	77.98	1369.95	0.90	0.348
Part	1	39	2418.03	90.98	1036.58	0.68	0.414
Stimulus by part	1	39	27.23	69.78	15.22	0.01	0.921
Go vs. Two-choice							
Stimulus	1	39	7507.60	134.40	2178.55	1.43	0.239
Part	1	39	19624.90	117.10	6536.05	4.30	**0.045**
Stimulus by part	1	39	7344.10	109.90	2606.19	1.71	0.198
P3a: Peak latency (50%)							
Go vs. No-go							
Stimulus	1	39	777.32	27.18	1115.16	0.73	0.397
Part	1	39	5219.94	26.14	7788.76	5.12	**0.029**
Stimulus by part	1	39	232.92	8.39	1082.15	0.71	0.404
Go vs. Two-choice							
Stimulus	1	39	5926.80	76.99	3002.46	1.97	0.168
Part	1	39	2750.64	21.90	4897.52	3.22	0.081
Stimulus by part	1	39	20.62	16.38	49.09	0.03	0.858
P3b: Peak latency							
Go vs. No-go							
Stimulus	1	39	3572.10	27.90	4993.26	3.28	0.078
Part	1	39	0.40	27.60	0.57	0.00	0.985
Stimulus by part	1	39	25.60	18.40	54.26	0.04	0.851
Go vs. Two-choice							
Stimulus	1	39	13.23	43.78	11.78	0.01	0.930
Part	1	39	429.03	35.98	465.10	0.31	0.583
Stimulus by part	1	39	632.03	32.98	747.50	0.49	0.487

Table 6.9. Overview of the Analyses of Variance of the stimulus-locked peak latencies in Experiment 12 with part (beginning, end) and stimulus (go, no-go, two-choice) as a within-subjects factors. As these analyses were performed on jack-knifed averages, the *F* values were adjusted according to $F_c = F/(n - 1)^2$, where *n* is the number of observations per cell (see *ERP analyses* section, for a detailed explanation of this procedure). p_cs < 0.05 are highlighted in bold.

A comparison of the go and two-choice stimuli revealed no reliable difference in the peak amplitudes between these stimuli ($p_c = 0.239$; Table 6.9). Amplitudes peaked 22 ms earlier at the end of training than at the beginning of training; this difference reached significance ($p_c = 0.045$; Table 6.9). Visual inspection of Figure 6.9 indicates that the difference between the beginning and the end of training was greater for the two-choice stimuli ($M_{\mathrm{diff}} = 36$ ms) than for the go stimuli ($M_{diff} = 8$ ms). However, there was no reliable two-way interaction between stimulus and part ($p_c = 0.198$; Table 6.9).

P3a range.

Mean amplitudes. Consistent with the ERP literature (for an overview, see the introduction of this Chapter), the mean amplitude in the P3a range was more positive for the no-go stimuli than for the go stimuli ($p = 0.031$; Table 6.8). The mean amplitude was reliably more positive at the end of training than at the beginning of training ($p = 0.008$; Table 6.8). There was no reliable two-way interaction between stimulus (go, no-go) and part was not reliable ($p = 0.120$; Table 6.8).

A comparison of the go and two-choice stimuli revealed no reliable main effect of stimulus ($p = 0.357$; Table 6.8). However, the mean amplitude in the P3a range was more positive at the end of training than at the beginning of training ($p = 0.002$; Table 6.8; Figure 6.7). The two-way interaction between stimulus and part was not reliable ($p = 0.409$; Table 6.8).

Peak latencies. Amplitudes for the go and no-go stimuli rose to 50% of peak amplitude at 320 ms (see *Method*). There was no reliable difference between the criterion latencies for the go and no-go stimuli ($p_c = 0.397$; Table 6.9). However, amplitudes rose to criterion 11 ms earlier in the go/no-go task at the end of training than at the beginning of training; this difference was reliable ($p_c = 0.029$; Table 6.9; Figure 6.9). The two-way interaction between stimulus and part ($p_c = 0.404$; Table 6.9).

The P3a for the go and two-choice stimuli rose to 50% of peak amplitude at 328 ms. There was no reliable difference between the go and two-choice stimuli ($p_c = 0.168$; Table 6.9). Although the P3a rose to criterion 8 ms earlier at the end of training than at the beginning of training, this difference did not reach significance ($p_c = 0.081$; Table 6.9). The two-way interaction between stimulus and part was not reliable ($p_c = 0.858$; Table 6.9).

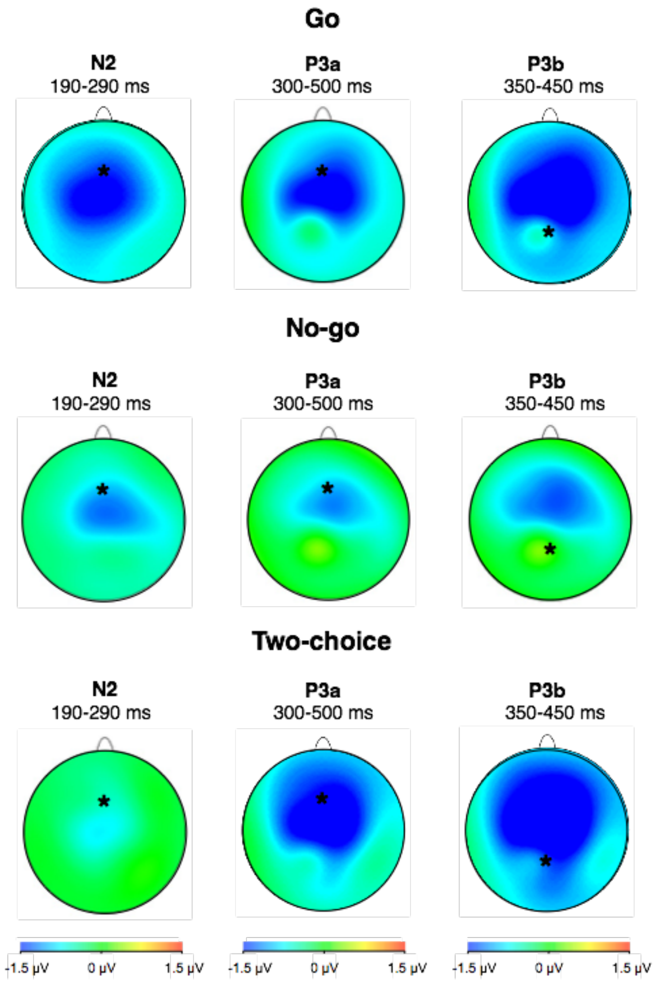

Figure 6.8. Topographies for the beginning minus end difference as a function of stimulus (go, no-go, two-choice) and analysed time-window (N2: 190-290 ms; P3a: 300-500 ms; P3b: 350-450 ms) in Experiment 12. * denotes the location of the analysed electrode in each time-window (N2 & P3a: Fz; P3b: Pz).

P3b range.

Mean amplitudes. Consistent with the results in Experiment 11, the mean amplitude in the P3b range was more positive for the go stimuli than for the no-go stimuli ($p < 0.001$; Table 6.8). There was no reliable main effect of part ($p = 0.372$; Table 6.8), but the two-way interaction between stimulus (go, no-go) and part was reliable ($p = 0.012$; Table 6.8); as can been seen in Figure 6.9, the mean amplitude in the P3b range was numerically more positive at the end of training than at the beginning of training for the go stimuli, but there was no similar change for the no-go stimuli. Follow-up comparisons revealed a reliable main effect of part for the go stimuli, $t(39) = -2.14$, $p = 0.039$, $g_{av} = -0.19$, $BF = 1.31$, but no reliable main effect of part for the no-go stimuli, $t(39) = 0.22$, $p = 0.828$, $g_{av} = 0.03$, $BF = 0.17$. Note that the Bayesian analyses provide support for the null hypothesis of no difference between the beginning and the end of training for the no-go stimuli.

There was no reliable difference in the mean amplitude in the P3b range between the go and two-choice stimuli ($p = 0.328$; Table 6.8). Crucially, the mean amplitude was more positive at the end of training than at the beginning of training ($p = 0.002$; Table 6.8). The two-way interaction between stimulus and part was not reliable ($p = 0.287$; Table 6.8).

Peak latencies. There were no reliable main effects or interactions in the P3b peak latencies (p_cs ≥ 0.078; Table 6.9).

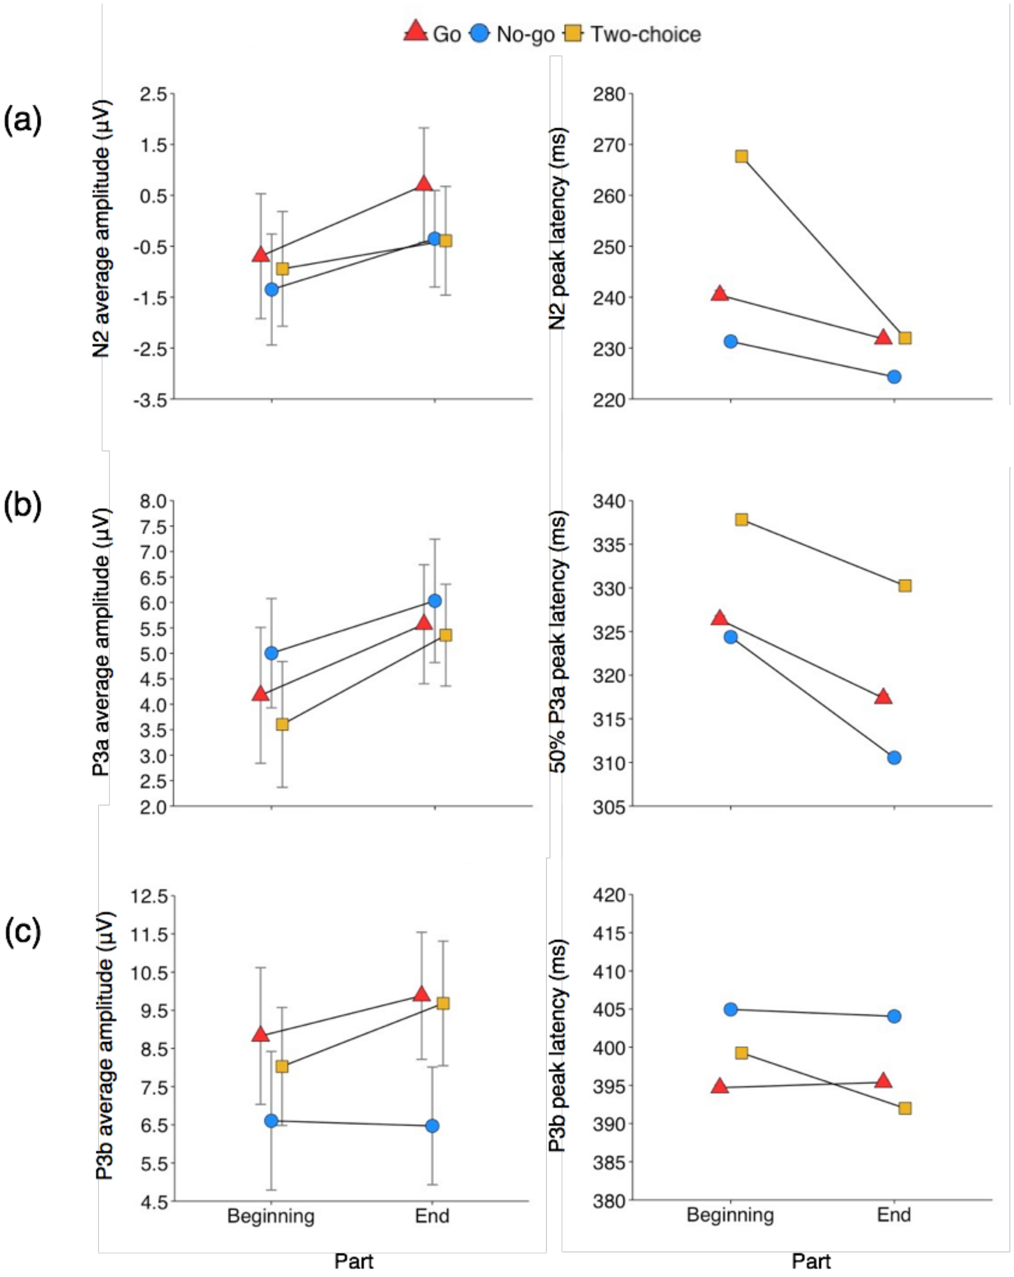

Figure 6.9. (a) Mean average amplitude in the N2 range (190-290 ms) and N2 peak latency (b) mean average amplitude in the P3a range (300-500 ms) and peak P3a latency; and (c) mean average amplitude in the P3b range (350-450 ms) and peak P3b latency, in Experiment 12 as a function of stimulus (go, no-go, two-choice) and part (beginning, end). Error bars are 95% confidence intervals. Due to the broad peak of the P3a component, I could not obtain an accurate peak latency (particularly for the no-go stimuli). Therefore, I used the latency at 50% of peak amplitude for this component (see Kiesel, Miller, et al., 2007). Note that the confidence intervals are small for the peak latencies due to the jack-knifing procedure (see *ERP analyses* above).

Response-locked analyses. Consistent with the findings in Experiment 11, the mean amplitude in the 200 ms preceding response execution was less negative for the go stimuli than for the two-choice stimuli (*p* = 0.013; Table 6.10; Figure 6.10). Similarly, there was no reliable difference between the amplitude at the beginning of training and the amplitude at the end of training (*p* = 0.364; Table 6.10). The corresponding Bayes factor was 0.25. There was also no reliable two-way interaction between stimulus and part (*p* = 0.994; Table 6.10). Thus, the training-effects observed in the P3a and P3b ranges are unlikely to be entirely due to motor response-related effects.

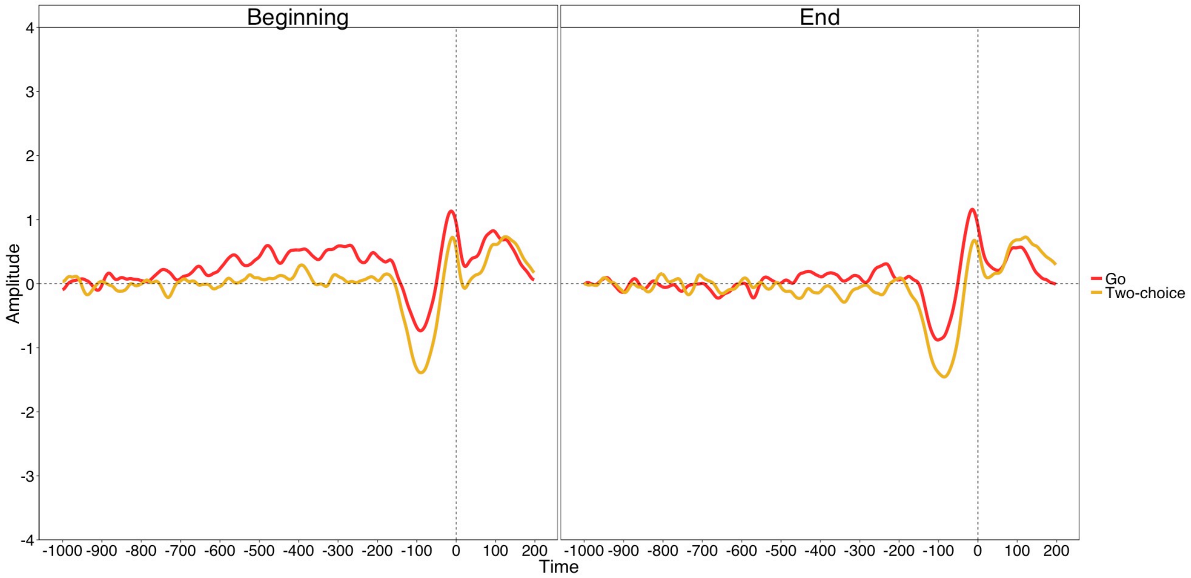

Figure 6.10. Lateralised readiness potentials for the go and two-choice stimuli in Experiment 12 as a function of part (beginning, end). The response was executed at 0 ms.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Stimulus	1	39	9.43	53.77	6.84	**0.013**	0.013
Part	1	39	0.81	37.24	0.84	0.364	0.001
Stimulus by part	1	39	0.00	17.08	0.00	0.994	< 0.001

Table 6.10. Overview of the Analyses of Variance of the response-locked average amplitudes (-200 to 0 ms) in Experiment 12 with part (beginning, end) and stimulus (go, two-choice) as a within-subjects factors.

Chapter Discussion

The aim of this chapter was two-fold: (1) to investigate how training in inhibition tasks influences the ERP signatures associated with action selection processes and (2) to compare the effects of training in inhibition and non-inhibition tasks on ERPs in a within-subjects training design. To this end, I compared the mean amplitudes and peak latencies of the N2, P3a, and P3b components at the beginning and the end of task performance.

Based on an exploratory ERP pilot experiment (Appendix K), I predicted that the N2 would become less negative (i.e. smaller) at the end of training than at the beginning of training and that the P3a and P3b sub-components would become more positive (i.e. larger) at the end of training than at the beginning of training[58]. I also predicted that behavioural performance would improve during training.

Contrary to my predictions, Experiment 11 showed no reliable improvement in behavioural performance between the beginning and the end of training in the go/no-go and two-choice tasks. If anything, performance on no-go trials declined across the training phase. Consistent with these behavioural results, I also found no reliable overall differences in the peak latencies or mean amplitudes of the N2, P3a, or P3b components. Nevertheless, despite no reliable overall differences, further interactions revealed that, unlike for the go stimuli in the go/nogo task, mean amplitudes in the N2 range for the two-choice stimuli became less negative with training. A common interpretation is that the N2 component reflects control processes involved in the resolution of conflict between competing responses (e.g. Enriquez-Geppert, Konrad, et al., 2010; Nieuwenhuis, Yeung, et al., 2003; Ramautar, Kok, et al., 2004). Thus, it is plausible that the observed amplitude effect is linked to training-induced improvements in the resolution of conflict, especially as the competition between responses in the two-choice task is considered to be greater than the competition between the 'go' and 'no-go' responses the go/no-go task (see, e.g. Gomez, Ratcliff & Perea, 2007) . However, the absence of a behavioural training effect fundamentally complicates the interpretation of this effect; without a reliable training-effect on behaviour it is impossible to conclusively attribute the N2 amplitude difference observed in the two-choice task to the effects of training.

Based on a behavioural follow-up experiment (Appendix L), I concluded that the absence of a behavioural training-effect in Experiment 11 was most likely due to the use of a 1:1 stimulus-response mapping which limited the magnitude of training-effects on performance. Therefore, in Experiment 12 I introduced eight additional stimuli in order to increase the task difficulty. Based on the results of the behavioural follow-up experiment, it was anticipated that this would result in reliable behavioural task improvements between the beginning and the end of training. All other aspects of the task remained the same as in Experiment 11. Analyses of the behavioural data showed that increasing the stimulus set-size was effective; the RTs on correct go and two-choice trials and the probability of responding on no-go trials decreased between the beginning and the end of training. Thus, any effects of training on the ERPs in Experiment 12 can be more confidently attributed to the effects of training than those in Experiment 11 (where there was no reliable behavioural improvement). Therefore, in the remainder of this Chapter Discussion, I focus on the differences between the beginning and the end of training on the N2, P3a, and P3b components in Experiment 12.

---

[58] It is important to note that the pilot study also indicated that these differences could be larger in the two-choice task than in the go/no-go task. However, as a between-subjects design was used, it is possible that these differences reflect between-task differences rather than qualitative differences in the effects of training. The within-subjects design used in Experiments 11-12 addressed this limitation.

The N2 Component

In the N2 range, I found that mean amplitudes became less negative with training for the go, no-go, and two-choice stimuli. This finding is consistent with my predictions based on the exploratory pilot study presented in Appendix K. Evidence that training similarly influenced all three stimulus types provides support for the interpretation of the N2 as reflecting task-general conflict monitoring processes rather than processes linked only with response inhibition.

The idea training influences the resolution of conflict between responses is entirely consistent with memory-retrieval accounts (Hommel, 1998, 2004; Logan, 1988; for a full overview of memory-retrieval accounts, see *Modelling Performance in Response Inhibition Tasks*, Chapter 1). As discussed in Chapter 1, memory-retrieval accounts state that, with consistent training, stimuli become associated or ‘bound’ with the correct response and these associations are stored as instances or event-files in memory. When participants begin performing a task they rely on general algorithms and require top-down control to resolve conflict between the various available responses in the task. However, with training, participants begin to rely on the retrieval of past responses from memory, reducing the need for conflict monitoring. In other words, memory-retrieval accounts suggest that less effort is needed to monitor for conflict between responses when the correct response can be automatically retrieved from memory. Importantly, the finding that the N2 component became less negative across the go, no-go and two-choice stimuli supports the idea that the resolution of conflict plays an important role in training in both inhibition and non-inhibition tasks.

However, whilst the observed reduction in the N2 amplitude in the go/no-go task is consistent with some previous findings (Benikos, unpublished PhD thesis, Study 3 and Study 4), it is at odds with other ERP findings in the go/no-go training literature. In particular, Schapkin and colleagues report that following go/no-go training, the amplitude of the N2 component were was more negative at the end of training than at the beginning of training (Schapkin, Johnstone, et al., 2007). There are several plausible explanations that could account for the inconsistency between the findings in Experiment 12 and those of Schapkin, Johnstone, et al.. Firstly, it is possible that the Schapkin, Johnstone, et al. task induced additional conflict between the go and no-go stimuli than in Experiment 12. For example, Schapkin, Johnstone, et al. used a go/no-go task in which there were both response compatible and response incompatible trials. On go trials, participants in the Schapkin, Johnstone, et al. study had to respond to the uppercase German word for ‘press’ (‘DRÜCK’) and the lowercase word for ‘stop’ (‘stopp’); on no-go trials, participants had to withhold their response to the uppercase word for ‘stop’ (‘STOPP’) and the lowercase word for ‘press’ (‘drück’). Thus, as features of the stimuli became associated with the go or stop goal with training, it is plausible that the presentation of response incompatible trials would begin to partially activate the incorrect response further increasing the need for conflict monitoring processes with training. This could have manifested in the more negative N2 amplitudes observed at the end of training in their study. Secondly, it is possible that inconsistencies between the findings reported in Experiment 12 and those of Schapkin, Johnstone, et al. could reflect the effects of long-term versus short-term training. In Schapkin, Johnstone, et al.'s experiment the N2 amplitude became more positive after three days of training (comprising two 20 minute go/no-go training sessions per day). Furthermore, prior to commencement of the training phase, participants in their study completed a practice phase until the number of omission errors on go trials did not exceed two out of 200 trials. Thus, whilst Experiment 12 primarily indexed the early stages of learning (as the task was entirely novel to participants), it is likely that the Schapkin, Johnstone, et al. study indexed the later stages of learning. Future research is required to investigate whether the effects of short-term and long-term training have different (and possibly opposing) effects on the processes underlying performance. Thirdly, only eight participants were tested in the Schapkin, Johnstone, et al. study meaning that the findings reported in their experiment are likely to be influenced by individual differences. On this basis, it is important to determine whether the Schapkin, Johnstone, et al. findings could be replicated in a sufficiently powered experimental design. In Experiment 12, I tested 40 participants and I ensured that I had sufficient trial numbers contributing to the ERP averages to obtain a reliable measure. Thus, Experiment 12 was sufficiently powered to detect medium-sized effects in the ERP analyses.

The P3 Component(s)

It is commonly accepted that the P3 component is best considered a non-unitary component comprising a frontal P3a sub-component and a temporo-parietal P3b sub-component that reflect separable processes. Whilst the amplitude of the frontal P3 is commonly more positive for the no-go stimuli than for the go stimuli, the parietal P3 component has been reported to be more positive for the go stimuli than for the no-go stimuli (Pfefferbaum, Ford, et al., 1985). To ensure that I captured these sub-components, I analysed and here report activity in frontal and parietal electrodes during the P3 time-window separately. Although it remains currently unclear *exactly* what the P3a and P3b sub-component reflects, an increasing body of literature has set about investigating what factors influence the amplitude and latency of the P3 component(s)[59]. The present chapter adds to this body of research in considering the effects of training in inhibition and non-inhibition tasks.

The frontal P3a component. In the P3a range, amplitudes were reliably more positive at the end of training than at the beginning of training for the go, no-go, and two-choice stimuli (Experiment 12). This finding is also consistent with the pilot ERP experiment (Appendix K) and previous ERP results obtained for go and no-go stimuli following go/no-go training (Benikos, Johnstone, et al., 2013; Benikos, unpublished PhD thesis; see also Song, Ding, Fan, & Chen, 2002).

---

[59] Note that the lack of complete agreement on the interpretation of these components is not a particular concern for the purposes of the present chapter. Importantly, I was interested in how similar the effects of training for the go, no-go, and two-choice stimuli on these components were.

A common interpretation is that the frontal P3a sub-component is thought to reflect frontal lobe activation associated with attentional processes required for the selection of relevant information from working memory (e.g. Berti, 2008; Duncan-Johnson & Donchin, 1977; for a theoretical overview, see Polich, 2003, 2007). The idea that attentional processes involved in memory updating may be influenced by training is also consistent with previous findings showing that attention plays a key role in the encoding and the retrieval of instances from memory (e.g. Boronat & Logan, 1997; Logan & Etherton, 1994). After all, instances stored in memory can contain various information about a stimulus (including a representation of its visual features). Thus, it is plausible that as performance begins to depend more on the retrieval of instances stored in memory and less on algorithmic processes, the engagement of attentional processes associated with memory updating would increase. Importantly, the presence of similar increased P3a amplitudes across the go, no-go and two-choice stimuli argues against the interpretation of this component as solely reflecting (top-down) inhibitory control processes, as has been previously assumed (for an overview, see e.g. Huster, Enriquez-Geppert, et al., 2013).

However, whilst the memory-retrieval interpretation seems plausible in the context of the observed training results, it is important to highlight that this account is purely speculative. Therefore, future research is required to directly test this explanation. For example, it is firstly important to establish whether the observed increase in amplitudes of the frontal P3a component and the less negative amplitudes of the frontal N2 component are additive, as it is possible that the observed more positive frontal P3a amplitude difference is not functionally separable to the observed N2 effect. For example, it is possible that the P3a effect simply reflects the ‘down-stream’ effects of the observed N2 reduction (i.e. less negative amplitudes). An inspection of Figure 6.7 suggests that this is unlikely as the absolute amplitude difference observed in the frontal P3 range is larger than the absolute amplitude difference observed in the N2 range. However, it is not possible to conclusively rule this out.

The parietal P3b component. Analyses in the parietal P3b range revealed that amplitudes became more positive at the end of training than at the beginning of training. However, unlike for the amplitudes in the frontal P3a range, this difference was only reliable for the go and two-choice stimuli, and was not reliable for the no-go stimuli (Figure 6.9). Bayesian analyses supported the null hypothesis of no difference for the no-go stimuli.

Consistent with these findings, several studies have shown a linear increase in the amplitude of the parietal P3b component during learning (Donchin & Coles, 1988; Pineda, Westerfield, Kronenberg, & Kubrin, 1997; Luft, Takase, & Bhattacharya, 2014). Furthermore, the P3b component has been shown to increase during performance of the Wisconsin card sorting task but to decrease when the task required that participants acquire a new rule (Barcelo, Munoz-Cespedes, Pozo, & Rubia, 2000; see also Morgan, Luu, & Tucker, 2016). However, despite being the subject of extensive research, the exact processes indexed by the parietal P3b component remains unclear.

Whilst some accounts suggest that the parietal P3b component reflects processes relating the updating of context information in working memory (Polich, 2007), other accounts suggest that the P3b component reflects the reactivation of acquired stimulus-response associations (Verleger, Jaskowski, et al., 2005). According to the latter account, the P3b component indexes the 'bridge' from stimulus to response (see also Verleger, Schroll, & Hamker, 2013) with larger P3b amplitudes when the stimulus-response associations can be readily retrieved from memory. The observed increase in the P3b amplitudes at the end of training is consistent with both interpretations. Firstly, it is possible that the increase in the P3b amplitude reflects the formation of representations of the task context with training (i.e. consistent with the context-updating account; Polich, 2007). The retrieval of this contextual information could serve to facilitate subsequent retrieval of the contextually-relevant response (see, e.g;. Morgan, Luu, & Tucker, 2016). Secondly, it is also possible that the increase in the P3b amplitudes reflects the formation of associations between the go and two-choice stimuli and the corresponding motor response (consistent with the 'stimulus-response associations account'; Verleger, Jaskowski, et al., 2005). According this this account, the P3b component could index processes that are specifically linked with the selection of a motor response. Consistent with this idea, across Experiments 11 and 12 I found that overall P3b amplitudes (i.e. independent of training) were more positive for the go stimuli than the no-go stimuli.

Importantly, the finding that the P3b amplitudes only reliably increased for the go and two-choice stimuli and not for the no-go stimuli could indicate that there are additional processes influenced during the acquisition of associations between specific stimuli and motor responding, that are not influenced during the acquisition of associations between specific stimuli and response inhibition. In other words, this could indicate that, there might be differences in the effects of training on response execution and response inhibition trials as well as similarities (Bender, Filmer, et al., 2016; Rae, Hughes, et al., 2014). To test this idea, future research should aim to distinguish between the effects of training on the frontal P3a and the parietal P3b components further. In particular, the use of principle component analyses or independent component analyses would help to further differentiate the effects of training on these P3 sub-components (see Enriquez-Geppert, et al., 2010 for a practical example).

Conclusion

The aim of the present chapter was to investigate how training in inhibition and non-inhibition tasks modulates the neural signatures of action selection and response inhibition processes. Using a within-subjects design in which participants alternated between performance of the go/no-go and two-choice tasks on a block-by-block basis, I demonstrate that the N2, P3a, and P3b components are sensitive to the effects of training in these tasks (Experiment 12) and that these ERP effects are mostly eliminated when training does not influence behavioural task performance (Experiment 11). Importantly, I find similar effects of training on the N2 and frontal P3a components across the go, no-go, and two-choice stimuli. Thus, these results extend the work presented in Chapter 5, address some of the methodological limitations highlighted in Chapter 5, and open the door to future research using ERPs to investigate the specificity of response inhibition training. On the basis of the results presented in this chapter, I conclude that there is most likely strong overlap between the effects of training on the action-selection processes that are involved in performing inhibition and non-inhibition tasks.

Appendix K

Between-Subjects Pilot ERP Experiment

This experiment was designed as a pilot experiment to compare the effects of training in the go/no-go and two-choice tasks on the N2, P3a, and P3b components. In this experiment, participants were randomly allocated (between-subjects) to the go/no-go task condition or to the two-choice task condition (for a similar design, see Experiments 9 & 10, Chapter 5). Participants completed a *training* phase (blocks 1-18) whereby the stimulus-response (and stimulus-stop (no-go) in the go/no-go task) mappings were consistent throughout task performance (e.g. diamond = go, square = no-go *or* diamond = left key press, square = right key press); followed by a *test* phase (blocks 19-21) in which the acquired mappings were reversed (e.g. square = go, diamond = no-go *or* square = left key press, diamond = right key press). As in Experiments 11-12, training phase performance was analysed as a function of 'part': the first three blocks in the training phase were the 'beginning' part, the final three blocks in the training phase were the 'end' part; the three blocks in the test phase were the 'test' part.

Method

Participants. Forty right-handed volunteers from University of Exeter participated for monetary compensation (£10) or partial course credit (*M* = 21.85 years, *SD* = 2.92, 24 females). Seven participants were removed and replaced due to low trial numbers in the ERP averages (see the *ERP analyses* section in the main text, for the exclusion criteria). Participants were randomly allocated to the go/no-go task condition or to the two-choice task condition (between-subjects).

Apparatus, stimuli, procedure, and behavioural analyses. The apparatus, stimuli, and procedure were identical to Experiment 11, except for the following: the stimuli consisted of four chequerboards, which varied along the orientation dimension (0° or 45°) (the stimuli were identical to those used in Chapter 5, see Figure 5.1). All participants started with a training phase, which consisted of eighteen blocks followed by a test phase, which consisted of three blocks. Each block consisted of 80 trials and each stimulus was presented 20 times per block. In the training phase, the stimuli were consistently mapped onto the same response; in the test phase, the stimulus-response mapping was reversed. All mappings were fully counterbalanced.

Separate ANOVAs were performed for the training and test phases with 'part' (beginning/end in the training phase analyses, end/test in the test phase analyses) as a within-subjects factor. In the analyses of the correct RTs, I also included stimulus (go, two-choice) as a between-subjects factor.

EEG/ERP analyses. The EEG set-up and pre-processing steps were the same as in Experiment 11 & 12. For each participant, I excluded an average of 5 components following the ICA with characteristic eye-blink and eye-movement topographies and time-courses. The average number of segments in each average are reported in Table K.1. There was no peak minima in the N2 range (180-350 ms), especially for the go stimuli (see Figure K.1), so it was not possible to estimate the peak latencies within this window. Otherwise, the peak latency time-windows for the P3a and P3b components were the same as in Experiments 11-12. Furthermore, the P3b had a broad plateau, particularly for the no-go stimuli (see Figure K.1), so I instead extracted the latencies for the amplitude rising 50% of the peak maximum (see Experiment 12 for details). Mean amplitudes were analysed in the following windows: N2 (200-300 ms), P3a (310-410 ms), and P3b (300-400 ms). For the P3a and P3b components, the mean amplitude windows were based on the peak latencies and visual inspection. It was not possible to use the peak latency to identify a suitable mean amplitude window for the N2 range. Therefore, a 200-300 ms window was used on the basis of previous research (for an overview see e.g., Huster, Enriquez-Geppert, et al., 2013). Visual inspection of the grand average waveforms confirmed that this window was appropriate. Note, for the analyses of the go and two-choice stimuli 'stimulus' was included as a between-subjects factor. All other ANOVAs were as reported for Experiment 11.

	Stimulus-locked						**Response-locked**					
	Beginning		**End**		**Test**		**Beginning**		**End**		**Test**	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Go	101	16	97	17	100	14	101	15	96	15	99	16
No-go	89	17	88	13	87	13	-	-	-	-	-	-
Two-choice	89	14	89	11	86	11	86	14	84	12	82	12

Table K.1. Overview of the average number of artifact-free ERP segments contributing to the stimulus-locked and response-locked waveform grand averages as a function of stimulus (go, no-go, two-choice) and part (beginning, end, test) for the pilot experiment. *M* = mean; *SD* = standard deviation.

Behavioural results

Descriptive statistics are reported in Table K.2 and inferential statistics are reported in Table K.3.

Training phase. Analyses of correct RTs showed that there was no reliable difference between the overall RTs for go and two-choice stimuli ($p = 0.271$; Table K.3). However, the main effect of part was significant ($p = 0.002$; Table K.3) indicating that mean RTs decreased between the beginning and the end of training. There was no reliable two-way interaction between stimulus and phase ($p = 0.110$; Table K.3). Analyses of the probability of responding on no-go trials in the go/no-go revealed no reliable main effect of part, $t(19) = 0.67$, $p = 0.510$, $g_{av} = 0.10$, $BF = 0.28$. Similarly, analyses of the probability of correct responses in the two-choice task revealed that there was no reliable difference between the beginning and the end of training, $t(19) = -0.37$, $p = 0.715$, $g_{av} = -0.09$ $BF = 0.25$.

	RT						***p*(correct)**						***p*(respond|no-go)**					
	Beginning		***End***		***Test***		***Beginning***		***End***		***Test***		***Beginning***		***End***		***Test***	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Go	349	42	339	36	338	38	-	-	-	-	-	-	-	-	-	-	-	-
No-go	-	-	-	-	-	-	-	-	-	-	-	-	0.13	0.02	0.12	0.01	0.15	0.02
Two-choice	374	60	344	50	368	60	0.83	0.09	0.85	0.06	0.82	0.08	-	-	-	-	-	-

Table K.2. Overview of the behavioural data in the between-subjects EEG pilot experiment. Average reaction time (RT) for go and two-choice stimuli, probability of an accurate response in the two-choice task [*p*(correct)] and the probability of responding on no-go trials in the go/no-go task [*p*(respond|no-go)] as a function of task (go/no-go, two-choice) and part (beginning, end, test). *p*(correct) is the ratio of correct two-choice trials to the number of correct and incorrect two-choice trials (missed trials are excluded). *M* = mean; *SD* = standard deviation.

Test phase. RTs were numerically longer in the test phase than at the end of training (see Table K.2); this difference was marginally significant ($p = 0.052$; Table K.3). There was no reliable difference in RTs between the go and two-choice stimuli ($p = 0.204$; Table K.3) but the two-way interaction between stimulus and part did reach significance ($p = 0.048$; Table K.3). Follow-up analyses revealed that the cost in reaction times between the end of training and the test phase was significant for the two-choice stimuli, $t(19) = -2.17$, $p = 0.043$, $g_{av} = -0.42$, $BF = 1.57$, but was not reliable in the go/no-go task, $t(19) = 0.051$, $p = 0.960$, $g_{av} = 0.01$, $BF = 0.23$. Note, however, that although the Bayesian analyses revealed substantial support for the null hypothesis of no difference between the end of training and the test phase for the go stimuli, there was only anecdotal support for the alternative hypothesis for the two-choice stimuli.

Analyses of choice accuracy in the two-choice task revealed no reliable difference between the end of training and the test phase, $t(19) = 1.13$, $p = 0.272$, $g_{av} = 0.21$, $BF = 0.41$. Similarly, the probability of responding on no-go trials in the go/no-go task was higher in the test phase than at the end of training, $t(19) = -2.40$, $p = 0.027$, $g_{av} = -0.44$, $BF = 2.29$. Thus, reversing the acquire stimulus-go and stimulus-response associations resulted in a small, but reliable cost in the two-choice and go/no-go tasks.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
Training phase							
Stimulus	1	38	4856.19	147794.39	1.25	0.271	0.027
Part	1	38	8296.48	26933.67	11.71	**0.002**	0.045
Task by part	1	38	1898.52	26933.67	2.68	0.110	0.011
Test phase							
Stimulus	1	38	6221.90	141710.69	1,67	0.204	0.036
Part	1	38	2685.59	25248.85	4.04	0.052	0.016
Stimulus by part	1	38	2784.12	25248.85	4.19	**0.048**	0.016

Table K.3. Overview of the Analyses of Variance of the reaction time data during the *training phase* and the *test phase* with part (end, test) as a within-subjects factor and stimulus (go, two-choice) as a between-subjects factor for the between-subjects EEG pilot experiment. *p*s < 0.05 are highlighted in bold.

ERP results

Figures K.1 presents the stimulus-locked grand average waveforms, Figure K.2 presents the mean average amplitudes and the peak latencies, and Figure K.3 presents the response-locked (lateralised) readiness potentials. Tables K.4- K.8 present overviews of the inferential statistics.

Training phase. As expected, the mean amplitude in the N2 range were less negative for the go stimuli than for the no-go stimuli (Figure K.1). Furthermore, consistent with the findings in Experiment 12, the amplitudes in the N2 range became less negative at the end of training than at the beginning of training (Table K.4). Analyses of the go and no-go stimuli showed that this difference was marginally significant in the comparison of the go and no-go stimuli ($p = 0.075$; Table K.4) and reached significance in the comparison of the go and two-choice stimuli ($p < 0.001$; Table K.4).

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen. η^2*
N2 range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	73.18	73.19	19.00	**< 0.001**	0.091
Part	1	19	18.01	96.51	3.54	0.075	0.024
Stimulus by part	1	19	2.22	29.22	1.44	0.244	0.003
Go vs. Two-choice							
Stimulus	1	38	0.69	877.33	0.03	0.864	0.001
part	1	38	95.44	170.34	21.29	**< 0.001**	0.083
stimulus by part	1	38	16.28	170.34	3.63	0.064	0.015
P3a range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	36.64	177.47	3.92	0.062	0.018
Part	1	19	48.55	225.34	4.09	0.057	0.024
Stimulus by part	1	19	0.53	36.18	0.28	0.604	< 0.001
Go vs. Two-choice							
Stimulus	1	38	31.40	1573.43	0.76	0.389	0.016
Part	1	38	7.32	260.46	1.07	0.308	0.004
Stimulus by part	1	38	108.17	260.46	15.78	**< 0.001**	0.056
P3b range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	175.46	230.65	14.45	**0.001**	0.099
Part	1	19	28.14	192.72	2.77	0.112	0.017
Stimulus by part	1	19	0.20	34.66	0.11	0.742	< 0.001
Go vs. Two-choice							
Stimulus	1	38	35.77	1809.94	0.75	0.392	0.017
Part	1	38	0.048	310.44	0.01	0.939	< 0.001
Stimulus by part	1	38	30.65	310.44	3.75	0.060	0.014

Table K.4. Overview of the Analyses of Variance of the mean average amplitude in the N2 range (200-300 ms post-stimulus), in the P3a range (310-410 ms post-stimulus), and in the P3b range (300-400 ms) in the *training phase* in the EEG pilot experiment. In the comparison between the go and no-go stimuli, part (beginning, end) and stimulus (go, no-go) were within-subjects factors. In the comparison between the go and two-choice stimuli, part (beginning, end) is a within-subjects factor and stimulus (go, two-choice) is a between-subjects factor. *p*s < 0.05 are highlighted in bold.

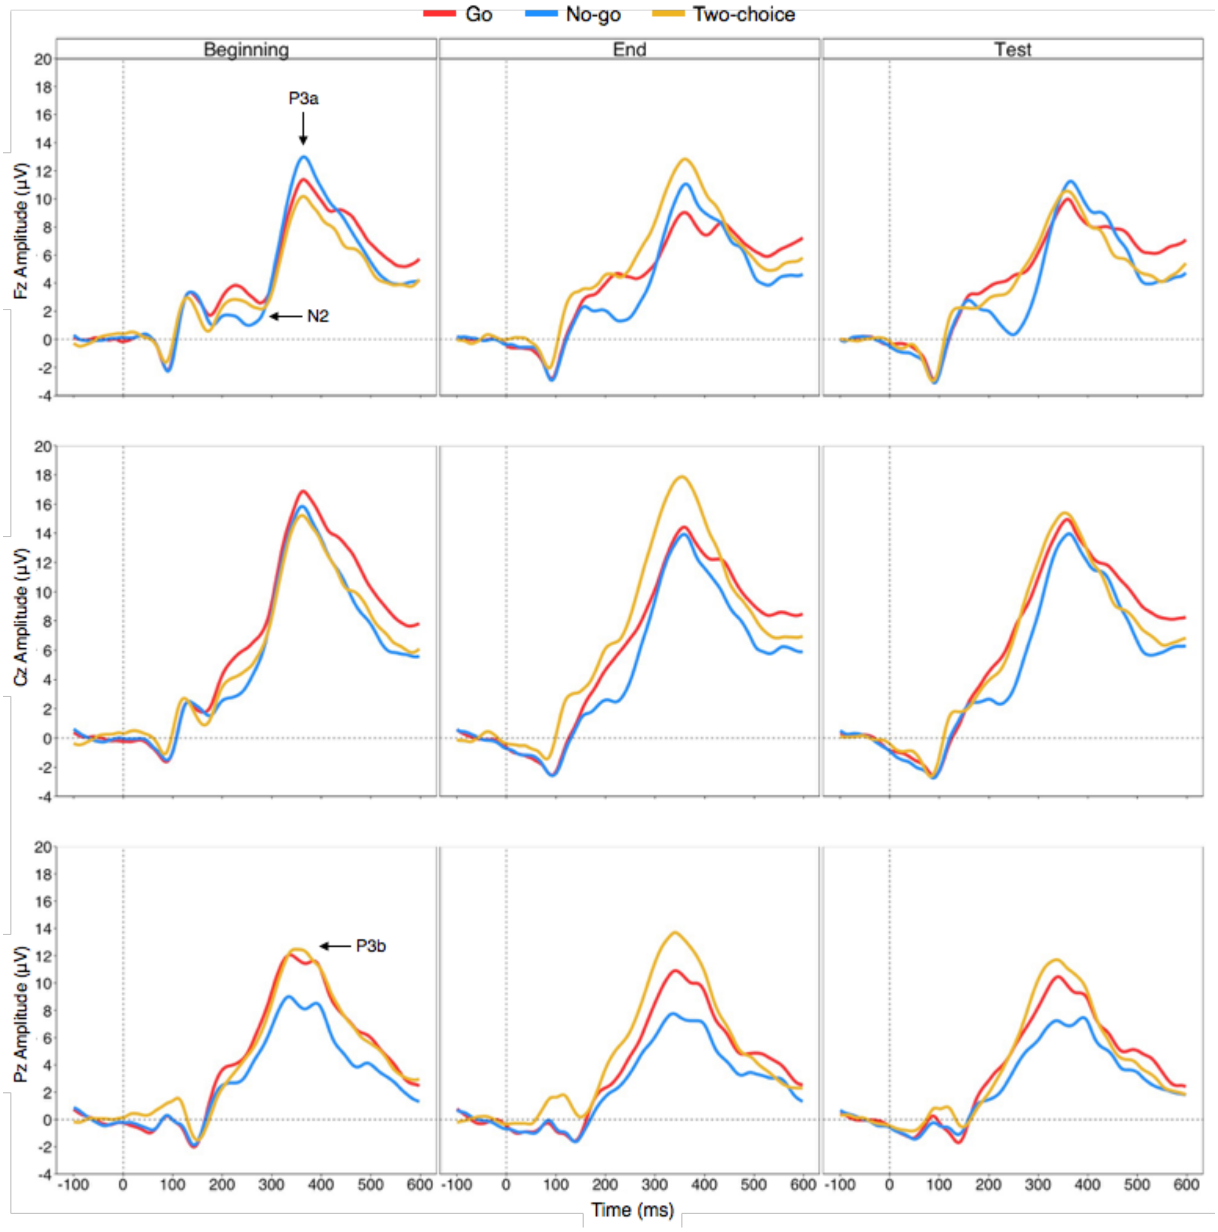

Figure K.1. (a) Stimulus-locked waveforms for the go, no-go, and two-choice stimuli in the ERP pilot experiment as a function of part (beginning, end, test) in electrodes Fz, Cz, and Pz. The stimulus was presented at 0 ms.

Consistent with the pattern of results in Experiment 11, analyses of mean amplitudes in the P3a range revealed that amplitudes were numerically less positive at the end of training than at the beginning of training in the go/no-go task; this difference was marginally significant ($p = 0.057$; Table K.4). The main effect of part in the comparison of the go and two-choice stimuli was not significant ($p = 0.308$; Table K.4), but the two-way interaction between stimulus and part was reliable ($p < 0.001$). As can be seen in Figure K.2, amplitudes became more positive for the two-choice stimuli at the end of training, but became less positive at the end of training for the go stimuli. There were no overall reliable differences between the beginning and the end of training on the amplitudes P3b range (*ps* ≥ 0.112; Table K.4). In addition to training-effects on amplitudes, a comparison of the go and two-choice stimuli also showed an effect of training on peak latencies; amplitudes in the P3b range reached 50% of peak amplitudes 12 ms earlier at the end of training than at the beginning of training ($p = 0.038$; Table K.5). There were no main effects of part on the P3a peak latencies ($ps \geq 0.237$; Table K.5). Note that I could not examine N2 peak latencies in this experiment because there was no identifiable peak minima, particularly for the go stimuli (see *Analyses*).

Consistent with Experiments 11 and 12, the analyses of the response-locked waveforms revealed no reliable main effect of part ($p = 0.250$; Table K.8) and Bayesian analyses supported the null hypothesis of no difference between the beginning and the end of training ($BF = 0.32$). There was also no reliable two-way interaction between stimulus and part ($p = 0.412$; Table K.8). Thus, it is unlikely that the differences observed for the two-choice stimuli are due to training-induced differences in processes related to motor preparation and execution.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***F_c***	***p_c***
P3a: Peak latency							
Go vs. No-go							
Stimulus	1	19	45.00	7.00	122.14	0.34	0.568
Part	1	19	245.00	15.00	310.33	0.86	0.365
Stimulus by part	1	19	24.20	11.80	38.97	0.11	0.746
Go vs. Two-choice							
Stimulus	1	38	0.45	29.50	0.58	0.00	0.968
Part	1	38	328.05	23.90	521.59	1.44	0.237
Stimulus by part	1	38	6.05	23.90	9.62	0.03	0.871
P3b: 50% peak latency							
Go vs. No-go							
Stimulus	1	19	77.24	42.29	34.70	0.10	0.760
Part	1	19	668.80	45.66	278.31	0.77	0.391
Stimulus by part	1	19	283.60	16.47	327.26	0.91	0.353
Go vs. Two-choice							
Stimulus	1	38	22.86	324.26	2.68	0.01	0.932
Part	1	38	2936.58	67.16	1661.58	4.60	**0.038**
Stimulus by part	1	38	2040.30	67.16	1154.45	3.20	0.082

Table K.5. Overview of the Analyses of Variance of the stimulus-locked peak latencies in the P3a range (310-410 ms post-stimulus), and in the P3b range (300-400 ms) in the *training phase* in the EEG pilot experiment. In the comparison between the go and no-go stimuli, part (beginning, end) and stimulus (go, no-go) are within-subjects factors. In the comparison between the go and two-choice stimuli, part (beginning, end) is a within-subjects factor and stimulus (go, two-choice) is a between-subjects factor. Note that unlike Experiments 11-12 I could not analyse peak latencies in the N2 range as there was no clear peak minima, particularly for the go stimuli. As these analyses were performed on jack-knifed averages, the *F* values were adjusted according to $F_c = F/(n - 1)^2$, where *n* is the number of observations per cell (see *Analyses* section for a detailed explanation of this procedure). p_cs < 0.05 are highlighted in bold.

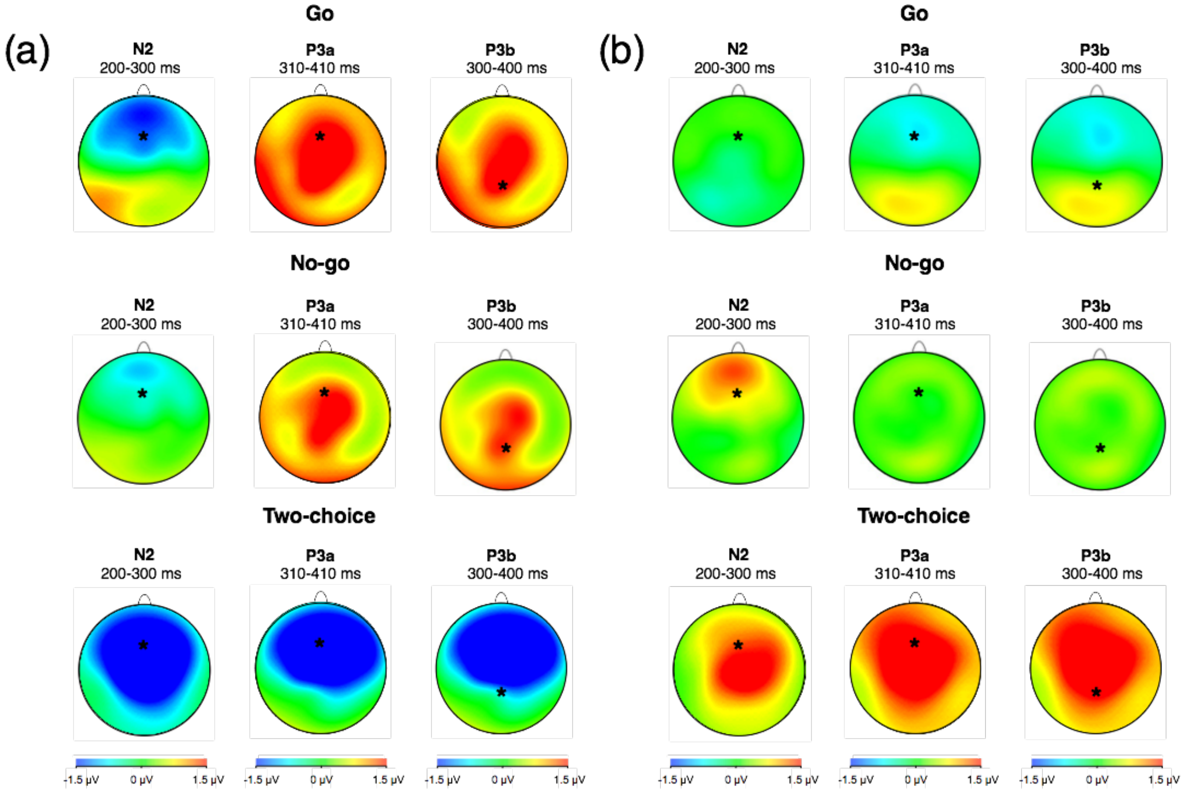

Figure K.2. Topographies for (a) the beginning minus end difference and (b) the end minus test difference as a function of stimulus (go, no-go, two-choice) and analysed time-window (N2: 200-300 ms; P3a: 310-410 ms; P3b: 300-400 ms) for the ERP pilot experiment.

Test phase. Consistent with Experiments 11 and 12, the mean amplitude in the N2 range was less negative for the go stimuli than for the no-go stimuli ($p < 0.001$; Table K.6; Figure K.1), the mean amplitude in the P3a range was more positive for the no-go stimuli than for the go stimuli ($p = 0.021$; Table K.6), and the mean amplitude in the P3b range was more positive for the go stimuli than for the no-go stimuli ($p < 0.001$; Table K.6). In the comparison of the go and two-choice stimuli, amplitudes in the P3b range were reliably less positive in the test phase than at the end of training ($p = 0.009$; Table K.6). Furthermore, there was a reliable two-way interaction between stimulus and part on amplitudes in the P3a range ($p = 0.001$; Table K.6). Follow-up analyses revealed that amplitudes became reliably less positive for the two-choice stimuli, $t(19) = 3.38$, $p = 0.003$, $g_{av} = 0.39$, $BF = 13.71$, but there was no reliable difference for the go stimuli, $t(19) = -1.53$, $p = 0.142$, $g_{av} = -0.20$, $BF = 0.63$. There were also no reliable differences in the mean amplitudes in the N2, P3a, or P3b ranges between the end of training and the test phase in the comparisons of the go and no-go stimuli ($ps \geq 0.213$; Table K.6) or in the mean amplitude in the N2 range comparison of the go and two-choice stimuli ($p = 0.103$; Table K.6). There were also no reliable main effects of part in the latency analyses ($ps \geq 0.120$; Table K.2); and no reliable main effects of part in the response-locked waveforms ($p = 0.614$; Table K.8; Figure

K.5) and no reliable two-way interaction between stimulus and part (*p* = 0.299; Table K.5). Bayesian analyses supported the null hypothesis of no difference between the beginning and the end of training in the response-locked waveforms (*BF* = 0.19).

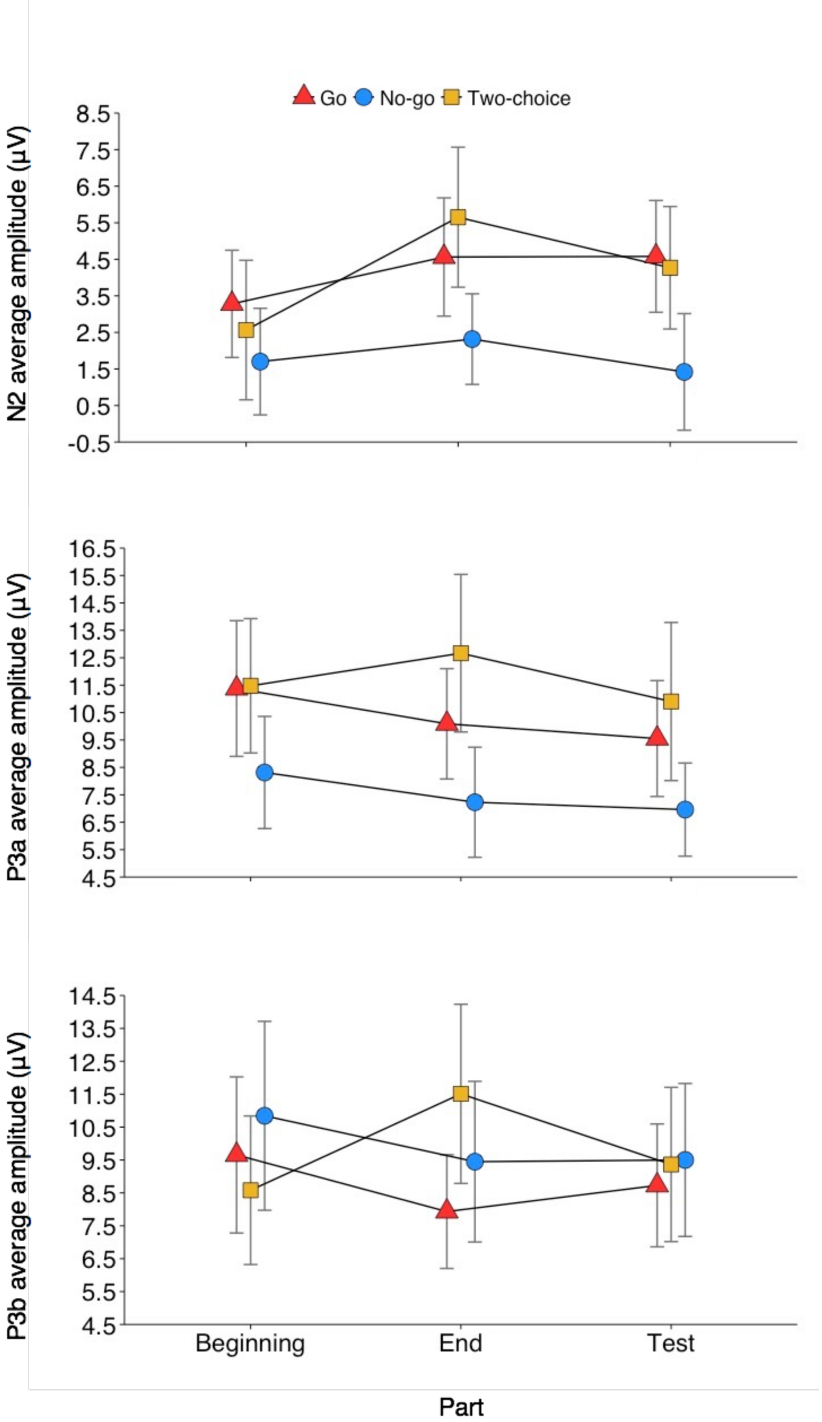

Figure K.3. Mean average amplitudes in the N2 range (200-300 ms), in the P3a range (310-410 ms), and in the P3b range (300-400 ms) in the ERP pilot experiment as a function of stimulus (go, no-go, two-choice) and part (beginning, end, test). Error bars are 95% confidence intervals.

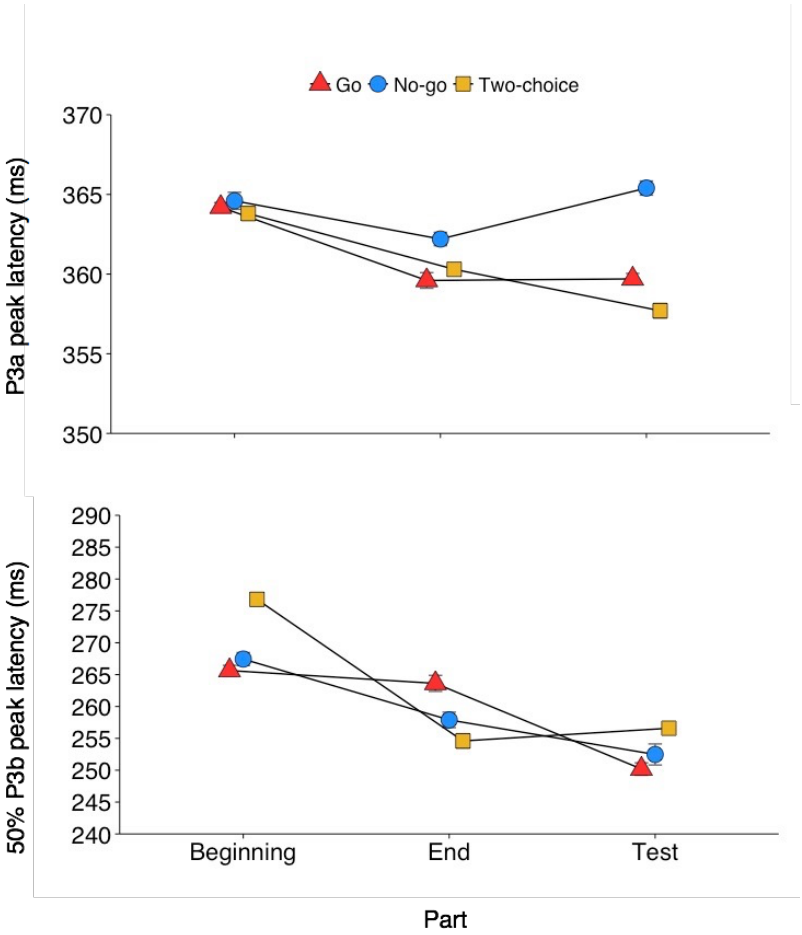

Figure K.4. Peak latencies in the P3a range and in the P3b range. Error bars are 95% confidence intervals. Note, due to the broad peak of the P3b component, I could not obtain an accurate peak latency, particularly for the no-go stimuli. Therefore, I used the latency at 50% of peak amplitude (see Kiesel, Miller, et al., 2007).

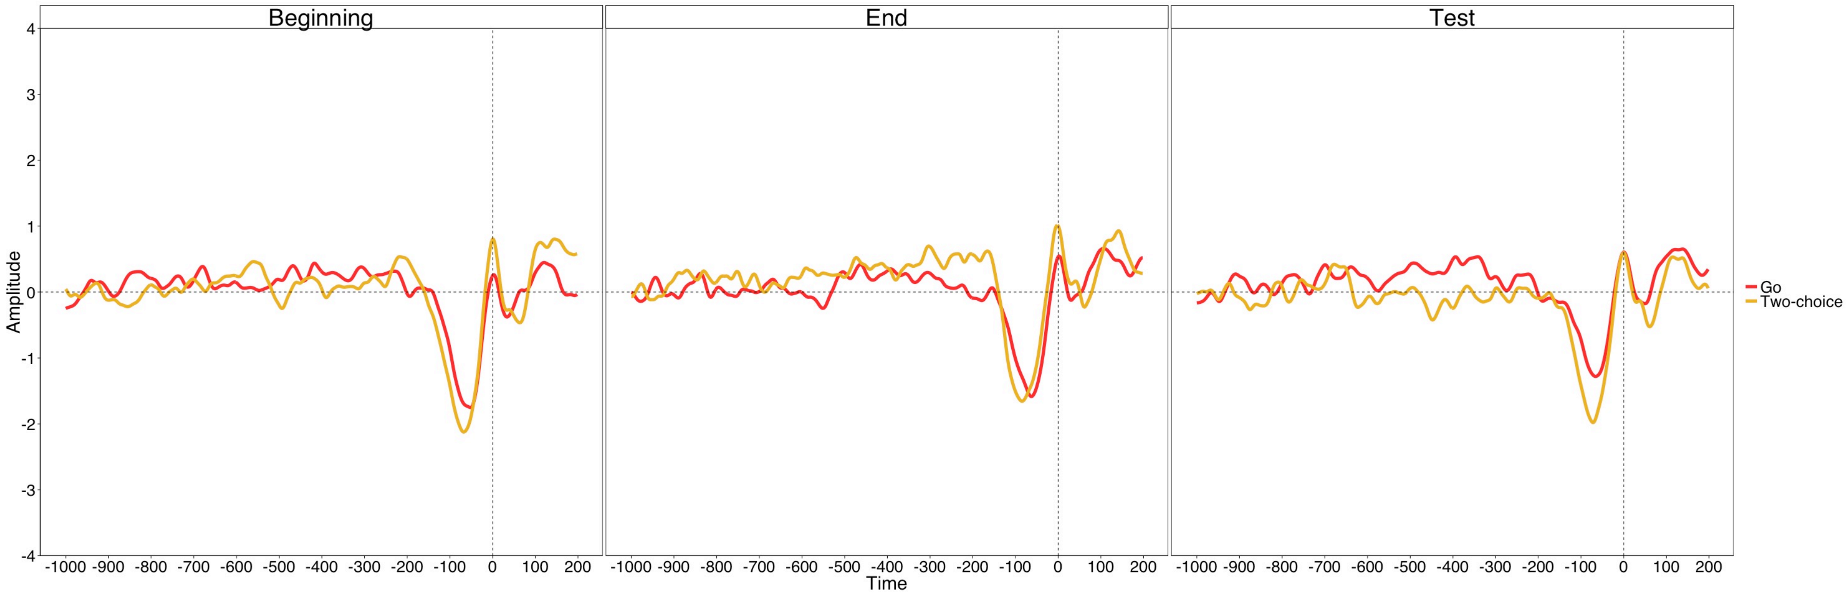

Figure K.5. Lateralised readiness potentials for the go and two-choice stimuli in the ERP pilot experiment as a function of part (beginning, end, test). The response was executed at 0 ms.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
N2 range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	146.15	87.82	31.62	**< 0.001**	0.157
Part	1	19	3.88	44.43	1.66	0.213	0.005
Stimulus by part	1	19	4.18	53.31	1.49	0.237	0.005
Go vs. Two-choice							
Stimulus	1	38	3.02	864.72	0.13	0.718	0.003
Part	1	38	9.32	126.74	2.80	0.103	0.009
Stimulus by part	1	38	9.78	126.74	2.93	0.095	0.010
P3a range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	26.21	79.02	6.30	**0.021**	0.017
Part	1	19	3.60	103.14	0.66	0.426	0.002
Stimulus by part	1	19	2.76	60.08	0.87	0.362	0.002
Go vs. Two-choice							
Stimulus	1	38	88.79	1554.00	2.17	0.149	0.050
Part	1	38	9.14	127.78	2.72	0.107	0.005
Stimulus by part	1	38	43.31	127.78	12.88	**0.001**	0.025
P3b range: Mean amplitude							
Go vs. No-go							
Stimulus	1	19	148.62	98.35	28.71	**< 0.001**	0.100
Part	1	19	3.25	57.82	1.07	0.315	0.002
Stimulus by part	1	19	0.36	47.81	0.14	0.708	< 0.001
Go vs. Two-choice							
Stimulus	1	38	77.09	2045.70	1.43	0.239	0.034
Part	1	38	26.44	130.88	7.68	**0.009**	0.012
Stimulus by part	1	38	7.49	130.88	2.17	0.149	0.003

Table K.6. Overview of the Analyses of Variance of the mean average amplitude in the N2 range (200-300 ms post-stimulus), in the P3a range (310-410 ms post-stimulus), and in the P3b range (300-400 ms) in the *test phase* in the EEG pilot experiment. In the comparison between the go and no-go stimuli, part (beginning, end) and stimulus (go, no-go) as within-subjects factors. In the comparison between the go and two-choice stimuli, part (beginning, end) is a within-subjects factor and stimulus (go, two-choice) is a between-subjects factor. *p*s < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	F_c	p_c
P3a: Peak latency							
Go vs. No-go							
Stimulus	1	19	344.45	8.55	765.44	2.12	0.162
Part	1	19	54.45	10.55	98.06	0.27	0.608
Stimulus by part	1	19	48.05	4.95	184.43	0.51	0.483
Go vs. Two-choice							
Stimulus	1	38	8.45	41.10	7.81	0.02	0.884
Part	1	38	31.25	18.30	64.89	0.18	0.674
Stimulus by part	1	38	36.45	18.30	75.69	0.21	0.650
P3b: 50% peak latency							
Go vs. No-go							
Stimulus	1	19	59.43	65.58	17.22	0.05	0.829
Part	1	19	1778.90	35.41	954.52	2.64	0.120
Stimulus by part	1	19	321.12	57.16	106.74	0.30	0.593
Go vs. Two-choice							
Stimulus	1	38	34.57	364.99	3.60	0.01	0.921
Part	1	38	654.79	45.89	542.25	1.50	0.228
Stimulus by part	1	38	1190.80	45.89	986.14	2.73	0.107

Table K.7. Overview of the Analyses of Variance of the stimulus-locked peak latencies in the P3a range (310-410 ms post-stimulus), and in the P3b range (300-400 ms) in the *test phase* in the EEG pilot experiment. In the comparison between the go and no-go stimuli, part (beginning, end) and stimulus (go, no-go) as within-subjects factors. In the comparison between the go and two-choice stimuli, part (beginning, end) is a within-subjects factor and stimulus (go, two-choice) is a between-subjects factor. As these analyses were performed on jack-knifed averages, the *F* values were adjusted according to $F_c = F/(n - 1)^2$, where *n* is the number of observations per cell (see *Analyses* section for a detailed explanation of this procedure). p_cs < 0.05 are highlighted in bold.

	Df1	*Df2*	*Sum of squares effect*	*Sum of squares error*	*F*	*p*	*gen.* η^2
Training phase							
Stimulus	1	38	0.04	350.81	0.00	0.950	< 0.001
Part	1	38	0.92	25.60	1.36	0.250	0.003
Stimulus by part	1	38	0.46	25.60	0.69	0.412	0.002
Test phase							
Stimulus	1	38	0.04	301.42	0.00	0.944	< 0.001
Part	1	38	0.27	39.41	0.26	0.614	0.001
Stimulus by part	1	38	1.15	39.41	1.11	0.299	0.003

Table K.8. Overview of the Analyses of Variance on the response-locked average amplitudes (-200, 0) in the ERP pilot experiment in the *training phase* and the *test phase* with part (beginning, end phases; end, test phases) as within-subjects factors and stimulus (go, two-choice) as a between-subjects factor. *p*s < 0.05 are highlighted in bold.

Appendix L

Behavioural Control Experiment

This experiment served two purposes: (1) to investigate whether alternating between performance of the go/no-go task and the two-choice task on a block-by-block basis interferes with practice-effects on behavioural performance; and (2) to investigate whether increasing the number of stimuli in each task (and consequently the stimulus:response mappings) influenced the magnitude of the practice-effect on behavioural task performance (relative to Experiment 11).

Method

Participants. Thirty-two right-handed volunteers from University of Exeter participated for monetary compensation (£5) or partial course credit (*M* = 20.59 years, *SD* = 3.60, 25 females). Five participants were removed and replaced; one due to technical problems with the stimulus presentation computer; four as their choice accuracy on two-choice trials was ≤ 65%.

Apparatus, stimuli, procedure, and analyses. The stimuli, apparatus, procedure, and analyses were identical to Experiment 11 except for the following: Participants were randomly allocated to the alternating condition or to the serial condition (between-subjects). In the alternating condition, participants alternated between performing the go/no-go task and the two-choice task on a block-by-block basis (as in Experiment 11, the order of the first block was counterbalanced); in the serial condition, half of the participants performed all the training blocks of the go/no-go task followed by all the training blocks of the two-choice task (and vice versa for the other half of participants). The stimuli consisted of four shapes and four coloured circles (stimuli S1-S4 in Figure 6.1). Each stimulus was presented 20 times per block; there was a two-to-one stimulus-to-response mapping in each task. For half of the participants, the coloured circles were used in the go/no-go task and the shapes in the two-choice task; for the other half, the shapes were used in the go/no-go task and the coloured circles in the two-choice task. The analyses were identical to Experiment 11, except that condition (alternating, serial) was included as a between-subjects factor.

Results and Discussion

Descriptive statistics are presented in Table L.1 and inferential statistics are presented in Table L.2.

Analyses of RTs revealed that the main effect of part was significant ($p < 0.001$, Table L.2) indicating that mean RTs decreased with task practice (Table L.1). As expected, participants were faster to respond to go stimuli than to two-choice stimuli ($p < 0.001$; Table L.2), but there was no reliable two-way interaction between stimulus and part ($p = 0.971$; Table L.2). There was no reliable overall difference in mean RTs between the alternating and serial conditions ($p = 0.061$; Table L.2). Crucially, the two-way interaction between condition (alternating vs. serial) and part was not significant ($p = 0.659$, Table L.2). Similarly, there was no reliable three-way interaction between condition, part, and task ($p = 0.793$; Table L.2). All other main effects and interactions were also not reliable (Table L.2).

	RT				***p*(correct)**				***p*(respond|no-go)**			
	Beginning		***End***		***Beginning***		***End***		***Beginning***		***End***	
	M	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Alternating condition												
Go	368	29	350	24	-	-	-	-	-	-	-	-
No-go	-	-	-	-	-	-	-	-	0.97	0.03	0.96	0.02
Two-choice	418	59	398	57	0.87	0.07	0.90	0.07	-	-	-	-
Serial condition												
Go	393	46	369	39	-	-	-	-	-	-	-	-
No-go	-	-	-	-	-	-	-	-	0.97	0.02	0.96	0.03
Two-choice	449	54	427	60	0.88	0.06	0.85	0.07	-	-	-	-

Table L.1. Overview of the behavioural data of the control experiment. Average go reaction time (RT) for go and two-choice stimuli, probability of an accurate response in the two-choice task [*p*(correct)] and the probability of responding on no-go trials in the go/no-go task [*p*(respond|no-go)] as a function of condition (alternating task order, serial task order), task (go/no-go, two-choice) and part (beginning, end). Accuracy is the ratio of correct two-choice trials to the number of correct and incorrect two-choice trials (missed trials are excluded). *M* = mean; *SD* = standard deviation.

	Df1	***Df2***	***Sum of squares effect***	***Sum of squares error***	***F***	***p***	***gen.*** η^2
Reaction times							
Condition	1	30	22297.55	177005.08	3.77	0.061	0.075
Part	1	30	13848.65	17377.98	23.91	**< 0.001**	0.048
Stimulus	1	30	90252.25	70035.67	38.66	**< 0.001**	0.246
Condition by part	1	30	115.08	17377.98	0.20	0.659	< 0.001
Condition by stimulus	1	30	618.86	70035.67	0.27	0.610	0.002
Part by stimulus	1	30	0.54	11888.07	0.00	0.971	< 0.001
Condition by part by stimulus	1	30	26.67	11888.07	0.07	0.793	< 0.001
***p*(respond|no-go)**							
Condition	1	30	0.00	0.04	1.46	0.236	0.041
Part	1	30	0.00	0.00	2.09	0.158	0.008
Condition by part	1	30	0.00	0.00	0.59	0.450	0.002
Two-choice accuracy							
Condition	1	30	0.01	0.25	1.22	0.278	0.033
Part	1	30	0.00	0.05	0.16	0.900	0.000
Condition by part	1	30	0.00	0.05	3.98	**0.024**	0.029

Table L.2. Overview of the Analyses of Variance with task and part as within-subjects factors and condition (alternating task order, serial task order) as a between-subjects factor for the control experiment. *p*s < 0.05 are highlighted in bold.

Table L.3 shows the outcome of the Bayesian ANOVA. As can be seen, the results are largely consistent with the ANOVA reported in Table L.2. Table L.3 shows that dropping the stimulus and part factors had a deleterious effect on the model. All other factors or interactions could be dropped. These analyses provide further support for the conclusion that alternating between performance of the go/no-go and two-choice tasks does not influence the magnitude of practice-effects on RTs much.

In addition to analysing RTs, I also analysed choice accuracy in the two-choice task and the probability of responding on no-go trials in the go/no-go task. There was no reliable main effect of part in choice accuracy in the two-choice task ($p = 0.900$; Table L.2) and no reliable difference between overall choice accuracy in the alternating and serial task order conditions ($p = 0.278$; Table L.2). The two-way interaction between condition and part did reach significance ($p = 0.024$; Table L.2). As can be seen in Table L.1, choice accuracy numerically increased between the beginning and the end of training in the alternating condition, but numerically decreased between the beginning and the end of training in the serial condition. However, follow-up comparisons revealed that the main effect of part was not reliable in the alternating condition, $t(15) = -1.62$, $p = 0.126$, $g_{av} = -0.35$, $BF = 0.75$, or in the serial condition, $t(15) = 1.78$, $p = 0.096$, $g_{av} = 0.31$, $BF = 0.92$. Consistent with previous results (see Chapter 5), there was no reliable difference in the p(respond|no-go) between the beginning of training and the end of training ($p = 0.158$; Table L.2). Furthermore, there was no reliable main effect of condition ($p = 0.236$; Table L.2) and no reliable two-way interaction between part and condition ($p = 0.450$; Table L.2). Thus, even when task difficulty was increased, choice accuracy and the p(respond|no-go) did not change with task practice.

Omitted factor(s)	***Bayes Factor***	***Confidence interval***
Stimulus	0.00	±88.02%
Part	0.00	±57.20%
Condition	0.48	±45.34%
Condition:part	1.78	±42.18%
Condition:stimulus	2.46	±44.98%
Condition:part:stimulus	2.65	±44.26%
Part:stimulus	4.35	±42.56%

Table L.3. Bayesian ANOVA for the reaction time data (for the go and two-choice stimuli).

CHAPTER 7

GENERAL DISCUSSION

The aim of this thesis was to investigate how stimulus-stop associations are acquired and to explore the conditions under which stop learning influences behaviour. Three main outstanding issues were addressed: (1) what is learned in response inhibition tasks; (2) the role of expectancies during stop learning; and (3) the specificity of stop learning. In this chapter, I summarise how the research presented in this thesis contributes to our understanding of each of these issues, highlight some limitations, and provide some suggestions for future research.

What is Learned?

A recent review proposed that there are at least two pathways through which learning could arise in response inhibition tasks (Verbruggen, Best, et al., 2014; for a visual depiction of these pathways, see Figure 1.4, Chapter 1); the *direct* pathway in which an unmediated link is formed between a specific stimulus and the stop goal, and the *indirect* pathway in which a link is formed between a stimulus and the stop goal that is mediated via a representation of whatever stop-signal or no-go category is used in the experimental task. According to this framework, the acquisition of direct associations would result in the automatic retrieval of the stop goal from memory following the presentation of stop-associated stimuli (i.e. the original ‘automatic inhibition’ hypothesis; Verbruggen & Logan, 2008a) whereas the acquisition of indirect associations would result in priming of stop-signal detection or some (weak) activation of the stop goal if the link is strong enough. Thus, these pathways have important implications for understanding what is learned in response inhibition tasks. One of the main contributions of this thesis was to investigate whether there is any evidence to support the existence of these proposed associative pathways.

Can Participants Acquire Direct (Stimulus-Stop) Associations?

Most previous research in the stop learning literature could not distinguish between the *direct* and the *indirect* pathways because the stop signal or no-go category remained the same throughout task performance (e.g. Lenartowicz, Verbruggen, et al., 2011; Verbruggen & Logan, 2008a). Therefore, in Chapter 2 (Experiments 1-4) I developed a task that combined features of the go/no-go task and the stop-signal task in which the stop signal rule changed at the beginning of each block.

Across four experiments, I showed that responding on go trials was slower and that the probability of responding on stop trials was lower for stop-associated items than for go-associated items and control items (that were not particularly associated with either go or stop; for the specific contingencies, see Table 2.1). As the representation of the stop signal was constantly changing, I conclude that the most parsimonious explanation for these findings is that participants acquired a direct association between the stop-associated stimuli and the stop goal, rather than an association between the stop-associated stimuli and the representation of a single signal representation. Together, these findings provide the strongest evidence to date for the original ‘automatic inhibition’ hypothesis that participants can learn direct stimulus-stop associations that are not mediated via a single stop signal representation. More generally, the idea that response inhibition can become a ‘learned reflex’ contributes to a growing body of research suggesting that control functions, such as response inhibition, interference control, and task switching, can rely on *both* bottom-up control and top-down control[60] (for an overview, see Chapter 5).

Can Participants Acquire Indirect (Stimulus-Signal) Associations?

The potential for stimulus-signal associations could suggest that previously observed RT slowing and neural activations (e.g. Lenartowicz, Verbruggen, et al., 2011) for stop-associated stimuli could be mediated by a link between the stimulus, the stop signal, and the stop goal. Thus, the issue of whether participants can acquire stimulus-signal associations plays an important role in our understanding of what is learned in response inhibition tasks.

In Experiments 7-8 (Chapter 4), I present a direct investigation of whether participants could acquire associations between specific stimuli and the spatial location of the stop signal. In addition to manipulating the contingencies between specific stimuli and the stop goal (as in typical stop learning tasks; e.g. Verbruggen & Logan, 2008a, Experiment 5), I also manipulated the contingencies between the stop- and go-associated stimuli and the spatial location of the stop/go signal. If participants could acquire stimulus-signal associations in this task, I predicted that (1) performance would improve during training for items consistent paired with the same signal location compared with items paired with two signal locations with equal probability (Experiment 7) and (2) performance would be better on trials where the signal appeared in the trained signal location than on trials where the signal appeared in the untrained location (Experiment 8). However, I found no evidence to support these predictions and Bayesian analyses provided support for the null hypotheses of no difference. Importantly, despite finding no evidence to support the idea that participants acquired stimulus-signal associations, there was reliable evidence in both experiments that participants acquired the stimulus-stop/go associations[61].

---

[60] For example, in the task-switching domain Logan and colleagues argued that, in addition to top-down reconfiguration, switching between tasks could also be achieved via the retrieval of acquired associations between cues, stimuli, and responses (Logan & Bundesen, 2003; Schneider & Logan, 2005).

[61] To maximise the number of stimulus-signal training trials, there were no control items (that were not particularly associated with stopping or going) in Experiments 7-8. Consequently, I could not distinguish between stimulus-stop and stimulus-go learning. However, as discussed in Chapter 2, it is likely that learning in response inhibition tasks reflects some combination of both stimulus-stop and stimulus-go learning.

Taken together, the findings reported in Chapter 4 clearly demonstrate that participants did not acquire associations between specific stimuli and the spatial location of the stop signal.

Nevertheless, it could be possible that the paradigm used in Experiments 7-8 (Chapter 4) may have underestimated the contribution of stimulus-signal learning. After all, the pattern of behavioural results obtained in Experiment 8 is consistent with the predictions of the indirect pathway (i.e. an effect of learning on stop performance but no effect of learning on go performance). Thus, one could perhaps argue that the paradigm used in Experiments 7-8 was not sensitive enough to detect attentional shifts on the basis of the stimulus-signal associations or that the stop/go task used in Experiments 7-8 does not well generalise to the kind of tasks that are used in the majority of stop learning studies. However, I suggest that this explanation is unlikely on the basis of the findings reported in Chapter 3.

In Experiments 5-6 (Chapter 3), I adapted a paradigm developed to investigate the role of signal detection processes in top-down response inhibition (Verbruggen, Stevens, et al., 2014). In this study, the stop signal was always presented in the periphery of the screen. The introduction of a perceptual distractor manipulation meant that it was possible to identify whether participants adjusted their attentional settings to detect the stop signal during task performance (as reflected by an increased distractor effect when participants expected a stop signal to appear in the periphery; see Verbruggen, Stevens, et al., 2014). The use of the perceptual distractor manipulation was a particularly good way to to test the main prediction of the indirect (stimulus-signal) pathway that the presentation of stop-associated should prime the detection of the stop signal; if participants acquired stimulus-signal associations, the distractor effect on no-signal trials should have been larger following the presentation of stop-associated (80%-stop) items than following the presentation of go-associated (20%-stop) items. In Chapter 3, I found clear evidence that item-specific stop learning influenced task performance and expectancy ratings obtained following task completion. However, there was no evidence that the distractors influenced performance more for the go-associated stimuli than for the stop-associated stimuli. Thus even if we assume that participants acquired stimulus-signal associations in this task, the effects of mediation on performance were negligible.

Combined, the research presented in this thesis provides no clear support for the stimulus-signal pathway (for similar conclusions, see Bowditch, unpublished PhD thesis). It might be possible to salvage evidence for the stimulus-signal route if we assume that stimulus-signal learning prioritises certain features of the stop signal (e.g. colour) more than others. In support of this suggestion, research in the visual attention domain has suggested that attentional selection can occur on the basis of independent features, with colour and shape gaining priority over spatial location (Andersen, Müller, & Hillyard, 2009; Kasten, & Navon, 2008; Saenz, Buracas, & Boynton, 2002). To investigate this suggestion in the context of stimulus-signal learning, it could be informative to cue attention to various signal attributes (e.g. colour, shape, location, orientation) on trial-by-trial basis (for an example, see Muller, Reimann, & Krummenacher, 2003). Cueing attention to various signal attributes would allow future research to directly investigate the possibility that some features of the signal representation may mediate the stimulus-stop association more strongly than other features. Nevertheless, it is important to highlight that even if it does turn out that the effects of mediation are stronger for some stimulus attributes than others, this would suggest that the stimulus-signal pathway operates under very restricted conditions. Thus, on the basis of the evidence presented in this thesis, the idea that response inhibition will *always* operate through the indirect (stimulus-signal) pathway when the stimulus-stop associations are mediated via a single representation of the stop signal throughout training seems highly unlikely.

Given that evidence for stimulus-signal learning is weak, an important issue is raised: why does stimulus-stop learning sometimes influence stop performance but not go performance (e.g. as in Verbruggen, Best, et al., 2014, Experiment 2)? In Chapter 4, I suggested an alternative explanation that could account for an effect of stop learning on stop performance, but not on go performance. According to the race model (Logan & Cowan, 1984; see *Modelling Performance in Response Inhibition Tasks*, Chapter 1), performance of response inhibition tasks depends on a balance between going and stopping; increasing the go response threshold will increase the finishing time of the go process and decrease the probability of responding on stop trials; decreasing the go response threshold will decrease the finishing time of the go process and increase the probability of responding on stop trials (e.g. Verbruggen & Logan, 2009a; see also Ratcliff, 1978). In other words, the probability of responding on stop trials is sensitive to the go RT distribution. Thus, it is possible that stop learning could primarily influence the fastest RTs because when responses are slowed as a consequence of strategic waiting the retrieval of stimulus-stop associations would no longer have an effect. In other words, strategic waiting and automatic slowing could be non-additive factors on performance. It is therefore possible that if stop learning primarily influences the fastest reaction times in the distribution, the effects of stop learning could manifest in the probability of responding on stop trials, but have no substantial effect on the go RT (which represents the mean of the whole go RT distribution). If this were the case, differences between the effects of learning on go and stop performance (e.g. as observed in Verbruggen, Best, et al., 2014, Experiment 2; see also Experiment 4, Chapter 2) would not reflect the acquisition qualitatively different kinds of learning than in other stop learning experiments, but rather differences in the sensitives of go RTs and the probability of responding on stop trials to stimulus-stop learning. To provide an initial test of this idea, I plotted RT distributions for the item-specific stop learning experiments presented in this thesis. As can be seen in Appendix M, there is some evidence to suggest that when the stimulus-stop contingencies are acquired via task practice (i.e. under uninstructed conditions) the slowing for the stop-associated items emerges (numerically) at the fast end of the RT distribution (see also Experiment 4, Figure B.4, Appendix B). However, it is important to note that (consistent with the RT distribution under instructed conditions; see Figure M.1), the RT distributions for Experiments 1-3 (Figure B.1-B.3, Appendix B) showed the reverse pattern: the slowing for the stop-associated stimuli emerged at the slow end of the distribution. Therefore, to test this idea more directly in future research, it could be useful to use a paradigm in which there would be training phase whereby the stimulus-stop mappings are consistent. In a subsequent test phase there would be one condition in which successful stopping would be rewarded over fast go responding (stop-reward) and other condition in which fast go response would be rewarded over successful stopping (go-reward). If the latency of the go response plays a role, it follows that the slowing for the old stop-associated items would be greater in the go-reward condition than in the stop-reward condition.

Is Stop Learning Special?

It was recently suggested that one of the biggest problems in the action control literature is that researchers often fall into the trap of referring to general constructs, such as ‘inhibition’ without specifying what the underlying mechanisms are (Verbruggen, McLaren, et al., 2014). Consistent with this idea, the effects of training are often explained in terms of strengthening or improving general constructs, such as ‘inhibitory control’, rather than specifying how the processes underlying action control are influenced by training. In this thesis, I aimed to advance on this position by considering whether the effects of learning are specific to inhibition tasks or whether they represent task-general (non-inhibition) effects. To this end, Experiments 9-12 compared the effects of training in the go/no-go and two-choice tasks. Several studies have suggested that response selection and response inhibition in these tasks are ‘two sides of the same coin’ (e.g. Mostofsky & Simmonds, 2008) and similar training/test phase effects to those found in the go/no-go task have also been found in non-inhibition (go only) tasks (e.g. Horner & Henson, 2009). However, the effects of learning in the go/no-go and two-choice tasks had not previously been compared in the same experimental design. Thus, it was previously impossible to draw any conclusions about the specificity of learning in response inhibition tasks as similarities (or indeed differences) between inhibition and non-inhibition tasks could be due to between-experiment effects. Furthermore, it was not possible to determine whether the similar behavioural effects observed for the inhibition and non-inhibition tasks could be attributed to the same underlying (neural) processes.

Therefore, in Experiments 9-10 (Chapter 5), I compared performance (between-subjects) in the go/no-go (inhibition) and two-choice (non-inhibition) tasks during a training phase in which the stimulus category determined how the participants should respond and in a test phase where the acquired associations were reversed (for further details of the experimental design, see *Mediation via a Category Representation,* Chapter 7). Importantly, the results showed similar improvements in performance during the training phase in both the go/no-go and two-choice tasks. Furthermore, the difference between the congruent and the incongruent test items[62] was small (but reliable) in both tasks

---

[62] For the ‘congruent items’ the response remained the same between the training and test phases; for the ‘incongruent items’ the response changed between the training and test phases.

(Experiment 9) or absent in both tasks (Experiment 10). The similarity between effects of training in the go/no-go and two-choice tasks is consistent with memory-retrieval accounts which suggest that, as a function of task practice, performance shifts from entirely rule-based, algorithmic processing towards performance on the basis of retrieved instances from memory (Logan, 1988, 1990). However, as discussed in Chapter 5, non-specific training effects are likely to also contribute to behavioural improvements during the training phase of both tasks.

However, similarities in the behavioural outcomes of training in these tasks do not necessarily mean that the same (neural) mechanisms were in play across inhibition and non-inhibition tasks. Furthermore, Experiments 9-10 used a between-subjects design in which participants were randomly assigned to the go/no-go task group or to the two-choice task group. Although this allowed for the inclusion of the test phase[63], it left open the possibility that the observed behavioural similarities were the result of indirect consequences of task performance (such as fatigue) rather than due to the effects of learning.

Therefore, in Experiments 11-12 (Chapter 6), I used event-related potentials to investigate whether the neural signatures of go/no-go and two-choice task performance were similarly influenced by training. I used a within-subjects design in these experiments in which participants alternated between performance of the go/no-go task and the two-choice task on a block-by-block basis to reduce any between-task confounds. In the analyses of the ERP data, I focussed on the N2 and P3 components as these components have been linked to processes required for performance of both the go/no-go task and the two-choice task (for an overview, see Chapter 6). Analyses of the ERP data showed similar effects of training on the N2 and P3a components across the go, no-go, and two-choice stimuli; the N2 component was less negative at the end of training than at the beginning of training and the P3a component was more positive at the end of training than at the beginning of training. The observed pattern ERP of results fits with the common interpretation that the N2 component reflects conflict monitoring processes (Nieuwenhuis, Yeung, et al., 2003; Yeung, Botvinick, & Cohen, 2004),and the P3a reflects the attentional processes involved in memory updating (Polich, 2007). Furthermore, these findings were consistent with other ERP research that previously investigated the effects of training in the go/no-go task (Benikos, unpublished PhD thesis; Benikos, Johnstone, et al., 2013). However, it is also important to note that I also found some differences between the go, two-choice stimuli, and no-go stimuli; whilst the P3b component was more positive at the end of training than at the beginning of training for the go and two-choice stimuli, there was no change between the beginning and the end of training for the no-go stimuli. One interpretation of the P3b component is that it reflects the formation of links between stimuli and a specific (motor) response (Verleger, Jaskowski, & Wascher, 2005). Thus, it plausible that this component would index additional effects of training for go and two-choice stimuli

---

[63] It would not have been possible to include a test phase in a within-subjects version as the results would have most likely been confounded by order-effects of task performance.

(where a motor response was required) than for no-go stimuli. Importantly, these results suggest that similar general processes are involved in the acquisition of associations in inhibition and non-inhibition tasks. However, learning to respond (on go trials in the go/no-go task and on all trials in the two-choice task) may additionally influence processes relating to motor responding.

From Associations to Behaviour

In addition to considering the mechanisms involved in the acquisition of stimulus-stop associations, it is also important to consider the mechanisms through which acquired stimulus-stop associations could influence behaviour. Therefore, in this section, I discuss what the research presented in this thesis tells us about how stimulus-stop learning could influence task performance. Firstly, I discuss the contribution of this thesis to theoretical models of stop learning; secondly, I discuss the role of expectancies in mediating the link between associations and behaviour.

How do Stimulus-Stop Associations Influence Behaviour?

Verbruggen and Logan (2008a) suggested that the effects of stop learning on behaviour can be explained by the interactive race model[64] (Boucher, Palmeri, et al., 2007; for an extended discussion of this model, see *Modelling Performance in Response Inhibition Tasks,* Chapter 1): following training, the go unit can be activated via the retrieval of an association between the go stimulus and going, and the stop association can be activated via the retrieval of an association between the go stimulus and stopping. Following the retrieval of the stimulus-stop association, activation in the stop unit slows the rate of accumulation in the go unit, increasing go RTs relative to control stimuli that were not particularly associated with either stopping or going. Note that this interpretation assumes that the activation of the stop unit is weaker when it is incorrectly activated on go trials than when it is correctly activated on no-go or stop-signal trials, which explains why go responding is slowed down following the retrieval of a stimulus-stop association but is not (or is only occasionally) completely inhibited. This suggests that there is an interplay between top-down control (activated on the basis of the go stimulus) and bottom-up control (activated via the retrieval of the stimulus-stop association). The findings reported in Experiments 1-6 (Chapters 2 & 3) are consistent with the predictions of this model; the presentation of stop-associated items under uninstructed conditions slowed responding on go trials but did not result in outright stopping (i.e. there was no increased probability of missed responses for the stop-associated items).

---

[64] To briefly recap, response inhibition can be modelled as a race between a ‘go’ process triggered by the presentation of the go stimulus and a ‘stop’ process triggered by the presentation of the no-go stimulus or the stop signal (Logan & Cowan, 1984). The key premise is that the process that finishes first determines whether a response is executed or not: when the stop process finishes before the go process, response inhibition is successful; when the go process finishes before the stop process, response inhibition is unsuccessful. According to the interactive race model, the go and stop process remain independent for the majority of their latencies before they interact towards the end.

However, whilst the interactive race model can account for the effects of stop learning under uninstructed (implicit) conditions, the findings of Experiments 5-6 (Chapter 3) suggest that the slowing for stop-associated items could operate through a slightly different mechanism when the stimulus-stop contingencies are acquired via explicit instruction. In Chapter 3, analyses of task performance on no-signal trials revealed that, alongside the expected slowing for the stop-associated (80%-stop) items, accuracy was also lower for the stop-associated items than for the other item types when the stimulus-stop associations were acquired through explicit task instructions. There was no evidence of this accuracy effect in the uninstructed condition. In Chapter 3, I proposed that these results can be explained by the blocked-input model of response inhibition (Logan, Yamaguchi, et al., 2015); consistent with the interactive race model, the blocked-input account predicts decreased go activity following the presentation of stop-associated stimuli. However, the blocked-input model suggests that decreased go activity is achieved by removing input to the go unit rather than by slowing the rate of accumulation in the go unit. Thus, when input to the relevant go unit is blocked, the probability of the incorrect response (and indeed the probability of a missed response) may increase. Future research is required to test this blocked-input idea and to examine whether similar effects on go accuracy that are observed under instructed conditions would also emerge under uninstructed conditions with sufficient training.

More generally, the accuracy data obtained in Chapter 3 highlights an unanswered question: do stimulus-stop associations have a specific or a global effect on go performance? The aforementioned Verbruggen and Logan (2008a) theoretical interpretation suggests that stimulus-stop associations that are acquired through task practice may have a global effect on responding (i.e. inhibit all go responses). Consistent with this idea, research in the sequential after-effects literature shows that the retrieval of (implicit) stimulus-stop associations acquired on a 'prime' trial influenced performance on a subsequent 'probe' trial regardless of the specific keypress response that was inhibited on the prime trial (Giesen & Rothermund, 2014; for an extended discussion of this experiment, see *Does Stimulus-Stop Learning Have a Global or Specific Effect on Responding?*, Chapter 1). However, the research presented in Chapter 3 suggests that when participants have explicit instructed awareness of the stimulus-stop contingencies, the contingencies could have more selective effects on go responding (i.e. stop left keypress). Consistent with this idea, it has been shown that stopping becomes more selective when participants have foreknowledge of which response to stop, but employ a global stopping mechanism when they do not have such foreknowledge (Aron & Verbruggen, 2008).

Although very speculative, this could indicate that stimulus-stop associations acquired through task practice could operate through different neural pathways to stimulus-stop associations acquired through task instructions. For example, Aron and Verbruggen (2008) suggest that that global stopping engages the 'hyperdirect' neural pathway (between the prefrontal cortex and the basal ganglia) whereas selective stopping engages the 'indirect' neural pathway (between the thalamus, external and internal globus pallidus, and subthalmic nucleus; for further details of these pathways, see *The Neural Mechanisms of Response Inhibition*, Chapter 1). This stimulus-stop associations acquired via experience could operate through the ‘hyperdirect’ pathway whereas stimulus-stop associations acquired via task instructions could operate through the ‘indirect’ pathway. Future research could utilise neuroimaging techniques to test this idea.

What Role do Expectancies Play in Stop Learning?

Evidence that participants generate expectancies that are consistent with the stimulus-stop contingencies could indicate that the response slowing for old stop items is due to strategic top-down effects rather than the development of bottom-up control (as originally assumed by Verbruggen and Logan (2008a); see Chapter 1). Yet, prior to the research presented in this thesis, the role of expectancies in stop learning tasks had not been investigated. One of the main contributions of this thesis was to investigate the role of expectancies during stop learning.

Across eight experiments presented in this thesis (Experiments 1-8), I found that participants generated expectancies that were consistent with the stimulus-stop contingencies in play following task completion (Experiments 1-3, 5-8) or during task completion (Experiment 4). However, in Chapter 3 I found that evidence of expectancies do not necessarily indicate that the slowing for stop-associated items is the result of controlled, explicit processes rather than automatic processes. To examine the contributions of explicit expectancies and learning, I introduced a between-subjects expectancy manipulation in which one group of participants were informed about the stimulus-stop contingencies (instructed group) whereas another group did not receive any information about the stimulus-stop contingencies (uninstructed group). The results showed that the stimulus-stop contingencies influenced task performance under both instructed and uninstructed conditions and participants in both conditions generated expectancies that were consistent with the stimulus-stop contingencies. However, these experiments also highlighted some differences: in the instructed condition, the expectancy ratings correlated with task performance whereas in the uninstructed condition, the expectancy ratings did not correlate with performance. These results suggest that the effect of stimulus-stop contingencies on behaviour can arise from explicit, top-down expectancies or from learning through experience.

However, it is important to note that although these findings suggest different contributions of explicit and implicit processes to stop learning under instructed and uninstructed conditions, they do not necessarily provide conclusive evidence that explicit and implicit processes have a separable influence on learning under these conditions. In future research, it would therefore be beneficial to further distinguish between the relative contributions of explicit and implicit processes during uninstructed item-specific stop learning. One feature of explicit processing is that it is resource-intensive (e.g. McLaren, Green, & Mackintosh, 1994). For example, Waldron and Ashby (2001) found that the execution of a concurrent task (a numerical Stroop task) resulted in a larger decrement on performance in a rule based task than in an information-integration task. The inclusion of a concurrent load manipulation during the acquisition of uninstructed stimulus-stop associations could be an effective way to examine the contribution of implicit (automatic) processes and rule-based processes to learning.

Another outstanding issue that arises from the expectancy data presented in this thesis concerns the extent to which the expectancy ratings obtained under uninstructed conditions reflect awareness of the stimulus-stop contingencies in play. It is notable that whilst the expectancy ratings obtained under uninstructed conditions were consistent with the stimulus-stop associations in play, the differences between the stimulus-types were numerically small especially relative to the differences observed under instructed conditions (Chapter 3). Consequently, I assumed that the effects of the stimulus-stop contingencies on expectancy ratings under uninstructed conditions were not necessarily indicative of conscious, verbalisable awareness of the stimulus-stop contingencies. However, I did not directly test this assumption. Furthermore, even if the expectancy ratings did not capture awareness in these tasks, this does not rule out the possibility that the expectancy ratings would begin to reflect explicit awareness in longer-term training designs. Indeed, it has been suggested that associative learning processes can give rise to conscious representations as learning develops (Perruchet, 2015). Note, however, that follow-up tests reported in Chapter 3 showed no reliable correlation between the expectancy ratings and task performance in final part of training in the uninstructed condition suggesting that this explanation cannot entirely account for the differences observed between the instructed and the uninstructed conditions. Nevertheless, future research should examine the extent to which the expectancy ratings capture conscious awareness of stimulus-stop associations. Dienes (2008) suggests that in order to show evidence for unconscious knowledge one must be able to show that a person has knowledge but does not know they have it (the ‘guessing criterion’). Thus it would be beneficial to obtain confidence ratings for each stimulus-stop expectancy rating; if the effects of the stimulus-stop contingencies on expectancy ratings are reliable but the person believes that they are guessing (i.e. the confidence ratings are not above baseline), this would provide much stronger evidence for the conclusion that participants are not explicitly aware of the contingencies under uninstructed conditions.

Furthermore, if it does turn out that these expectancies reflect awareness, it is important to consider whether awareness of the contingencies means that inhibition driven by the stimulus-stop contingencies is no longer automatic. For example, Tzelgov (1997) argued that awareness does not necessarily mean that performance is not automatic. Instead, Tzelgov argued that the defining feature of non-automatic processing is ‘monitoring’ and not awareness. In this context, monitoring refers to ‘the intentional setting of the goal of behaviour and to intentional evaluation of the outcome of the process’ (Tzelgov, 1997, p. 444). Furthermore, research in the associative learning domain suggests that main difference between the stimulus-stop learning under aware and unaware conditions could be the kind of representation that is linked with the stimulus: an abstract, rule-like representation (X–'if x then stop'), or more concrete stimulus–response associations (X–stop; McLaren, et al., 1994; see also Verbruggen, McLaren, et al., 2014). In other words, even if subjects are aware of the contingencies, this does not necessarily imply that an entirely different form of learning has taken place compared with situations in which subjects were not aware of the contingencies.

Top-Down and Bottom-Up Influences on Response Inhibition

One of the major contributions of the research presented in this thesis is to highlight that learning in response inhibition tasks depends on an interplay of top-down and bottom-up influences. In this section, I integrate some of these findings and propose how top-down and bottom-up control could interact during stimulus-detection and action selection stages to influence what is learned in response inhibition and when stop learning influences in behaviour. A visual depiction of these ideas is presented in Figure 7.1.

Stimulus Detection

In Chapter 2, I proposed that attention may be necessary for stimulus-stop associations to be retrieved from memory. In Experiments 1-3, I found no evidence that the acquired stimulus-stop associations influenced task performance in the test phase (when the mappings were reversed), despite effects of learning in the training phase and on expectancy ratings obtained following task completion. The presence of an effect of stimulus-stop learning in the expectancy ratings suggests that participants had not forgotten the stimulus-stop associations in the test phase. However, unlike in other stop learning experiments in the literature, Experiments 1-3 did not require participants to attend to the stop-associated stimuli meaning that participants could have begun to ignore the stop-associated stimuli without impairing overall task performance. Consistent with this possibility, there was some evidence to suggest that the influence of stimulus-stop learning began to diminish towards the end of the training phase (Experiment 1 & Experiment 3). In Experiment 4, I made some changes to the task design to encourage participants to attend to the stop-associated stimuli; the stop-associated stimuli were presented before the stop or go signals and participants were required to provide an expectancy rating on each trial. The results showed a clear effect of stimulus-stop learning in both the training phase (on the *p*(respond|stop) and expectancy ratings) and test phase (on go RTs, the *p*(respond|stop), and on expectancy ratings) in this experiment. The idea that top-down attentional settings might play a role in mediating the effects of stimulus-stop learning on behaviour is consistent with theoretical evidence.

For example, the Instance theory suggests that processing episodes will only be retrieved from memory if participants attend to each stimulus presentation (Logan, 1988; Logan & Etherton, 1994)[65].

It is plausible that when the stop-associated stimuli are task-irrelevant (Experiments 1-3) or the task demands rapid responding (Experiments 9-10) participants can apply a 'selection bias' towards a specific stimulus attribute. In the visual attention literature, it is an established idea that stimulus representations held in working memory can enable top-down control of attentional settings via the activation of cells selective for the features of the stimulus held within the representation (Desimone & Duncan, 1995). These settings can be flexibly adjusted; for example, behavioural experiments show that during task performance participants can dynamically change the attentional weights for specific stimulus attributes in accordance with task demands (Olivers & Eimer, 2011).

Consistent with this 'selection bias' idea, research in the stop learning literature suggests that not only is the task-relevance of the stimuli important for stop learning (see above) but so is the relative importance of the go and stop processes. Benikos and colleagues demonstrated that in conditions where the go response deadline was easy (1000 ms) or medium difficulty (500 ms) the amplitudes of the N1 component[66] for the go and no-go stimuli were the same at the end of training (Benikos, Johnstone, et al., 2013). However, when the go response deadline was strict (300 ms) the amplitude of the N1 component was reduced (i.e. less negative) for the no-go stimuli relative to the go stimuli at the end of training. Note that there was no difference between the N1 amplitudes between the go and no-go stimuli at the beginning of training. Benikos, Johnstone, et al. suggest that participants in the strict deadline condition directed attention away from the no-go stimuli and towards the the go stimuli in order to respond within the strict deadline. In other words, participants adjusted the attentional weighting of features corresponding to the go and no-go stimuli to prioritise attentional processing of the go stimuli (for similar ERP effects, see Manuel, Grivel, et al., 2010).

However, attention to the stop-associated stimuli is unlikely to be influenced by only top-down factors, it is highly plausible that attention is influenced by bottom-up factors (driven by the acquisition of the stimulus-stop associations) as well. Consistent with this suggestion, it has been shown that attentional selection can be modulated by past experiences, including learning (Awh, Belopolsky, & Theeuwes, 2012). For example, research in the associative learning literature shows that stimuli that are predictive of a given response outcome capture attention faster than stimuli that do not consistently predict the same outcome (Le Pelley, Vadillo, & Luque, 2013). Thus, stimuli that are consistently associated with stopping could capture attention faster than stimuli that are paired with both stopping

---

65. As discussed in Chapter 1, Giesen and Rothermund (2014) found that the identity of a task-irrelevant stimulus could become associated with stopping and influence task performance. However, the argument here is that participants can strategically bias attention away from stimulus attributes in accordance with top-down goals. Thus, although the task-irrelevance of the stop-associated stimuli is sufficient to induce a strategic bias, it is not an inevitable consequence.

66. The N1 component is sensitive to perceptual features of a stimuli and can be modulated by attention (Näätänen and Picton, 1987).

and going with equal probability due to the 'learned predictiveness' of the stop-associated stimuli. Another way stimulus-stop learning could influence stimulus detection is via a link between the appetitive/aversive systems (for an overview, see *Alternative routes*, below). For example, it has been shown that the acquisition of stimulus-stop associations could devalue and decrease the motivational value of stop-associated stimuli (Ferrey, Frischen, & Fenske, 2012; Veling, Aarts, et al., 2013a; Veling, Holland, & van Knippenberg, 2008. This could indirectly influence attention to stop-associated stimuli; for example, it has been shown that the predicted value of a stimulus influences how much attention is paid to it, with attention preferentially allocated to motivationally-relevant stimuli (e.g. Fenske & Eastwood, 2003). Thus, if (and when) devaluation does occur, this could indicate that attention to the stop-associated stimuli decreases with learning. These accounts suggest that attention to the stop-associated stimuli could be both increased and decreased by learning. Future research should examine the effects of training on the attentional capture of stop-associated stimuli and how these various influences on attention could interact.

In summary, it seems that attention can influence learning in response inhibition tasks, but learning could also influence attention. Alongside top-down influences on attention during learning, I suggest that the acquisition of stimulus-stop associations could influence attentional capture. This could even facilitate or impair further learning; according to Awh and colleagues, past attentional selection episodes are retrieved in subsequent trials when the relevant context is encountered again (Awh, Belopolsky, & Theeuwes, 2012). In future research, it could be useful to examine under what conditions task-irrelevant stimuli that co-occur with the presentation of task-relevant stop-associated stimuli also become associated with the stop goal. Furthermore, it is possible that the use of neutral images increased the extent to which participants began to 'tune out' their attention in Experiments 1-3. Motivationally-salient images capture attention even if they are task-irrelevant (e.g., Anderson, Laurent & Yantis, 2011). Consequently, if task-irrelevant, but motivationally-salient images are used as the stop-associated stimuli, the attentional capture to the images should be increased and the 'tuning out' of attention could be slowed. Thus, the salience of task-irrelevant stop-associated stimuli could be a key consideration for future research.

Action Selection

Response inhibition depends on the outcome of a race between a 'go' process and a 'stop' process; inhibition is successful when the stop process finishes before the go process whereas response inhibition is unsuccessful when the go response reaches the response threshold first (Logan & Cowan, 1984). Activation in the stop unit can occur via the presentation of a stop-signal (or no-go stimulus) or following the retrieval of a stimulus-stop association from memory (as demonstrated in this thesis). I propose that the activation of the stop unit via the retrieval of stimulus-stop associations slows the rate of accumulation in the go unit (see *How do Stimulus-Stop Associations Influence Behaviour?* above).

As shown in this thesis, this will result in longer reaction times. However, training in response inhibition tasks does not necessarily mean that learning will always influence go responding.

For example, in Experiments 9-10, I found that performance benefited from task practice: reaction times on go and two-choice trials decreased (Experiments 9-10) and the accuracy on two-choice trials became lower (Experiment 9) throughout the training phase . However, in the test phase of both experiments, there was little or no difference between the congruent items (where the response remained the same between the training and test phases) and incongruent items (where the response changed between training and test phases). Importantly, the absence of a difference was most obvious when the task instructions emphasised rapid responding (Experiment 10). To account for this pattern of results, I suggest that when the task demands the rapid reconfiguration of responding (in order to maximise accuracy or response speed) participants engage top-down control processes (in accordance with these higher-order task goals) in order to reduce any potential interference resulting from the retrieval of acquired stimulus-stop or stimulus-response associations. For example, when the instructed (explicit) task goal is to maintain rapid or accurate responding, participants may deploy a top-down 'task-set' to overcome any potential interference from prior stimulus-response or stimulus-stop learning. Dreisbach and colleagues (2008, 2009, 2011) suggest that participants can form task-sets to shield task performance from potential interference, such that when participants are informed of a task rule, performance is less susceptible to interference from competing response options than when performance is driven by stimulus-response associations (Dreisbach & Haider, 2009). According to this idea, one of the main functions of top-down control is to shield processing from potential interference between response tendencies. Shielding could occur by biasing attention to specific stimulus features (see *Stimulus Detection*, Chapter 7) or by activating top-down bias signals to enable the selection of the required response (e.g. go; Verbruggen, McLaren, & Chambers, 2014; see also Chun, Golomb, & Turk-Browne, 2011).

The presentation of new task rules at the beginning of the test phase, in combination with the instruction to respond as quickly as possible, may have encouraged task shielding in Experiments 9-10. Following the reversal of acquired stimulus-stop associations, the need for top-down 'shielding' could also be identified following the detection of conflict at the response level between the required response (i.e. 'go') and the response retrieved from memory (i.e. 'stop'; e.g. Botvinick, Braver, et al., 2001). This could be an effective way to activate top-down control. For example, research using the Stroop task has shown that the detection of conflict enables the prioritisation of colour name responses over more automatic responses associated with the word meaning (for an overview, see Botvinick, Cohen, & Carter, 2004). The conflict monitoring account states that when conflict is detected, top-down control adjustments are made to enhance task-relevant processing and eliminate or, at the very least, diminish the effects of task-irrelevant responses. Consistent with this idea, in Experiment 2, I found that participants were slower to respond to the stop-associated images than to the go-associated images in the first half of trials in block 13 (the first block of the test phase), but this was in the opposite direction in the second half of block 13. To account for this pattern, I suggested that this reversal could be due to an increased error signal in the first half of the test phase. In other words, the detection of conflict between the go and stop responses could rapidly activate top-down control to reduce the interference from the acquired stimulus-stop associations.

In future research, it would be useful to test whether the interference from the retrieval of old stimulus-stop associations is sensitive to top-down control by comparing the magnitude of the interference between the previously acquired stimulus-stop associations and the new task instructions in a condition where participants are explicitly instructed to avoid the effects of retrieving the old stimulus-stop associations, with a condition in which participants are not provided with any instructions to minimise interference (e.g. as in the original Verbruggen & Logan 2008a study).

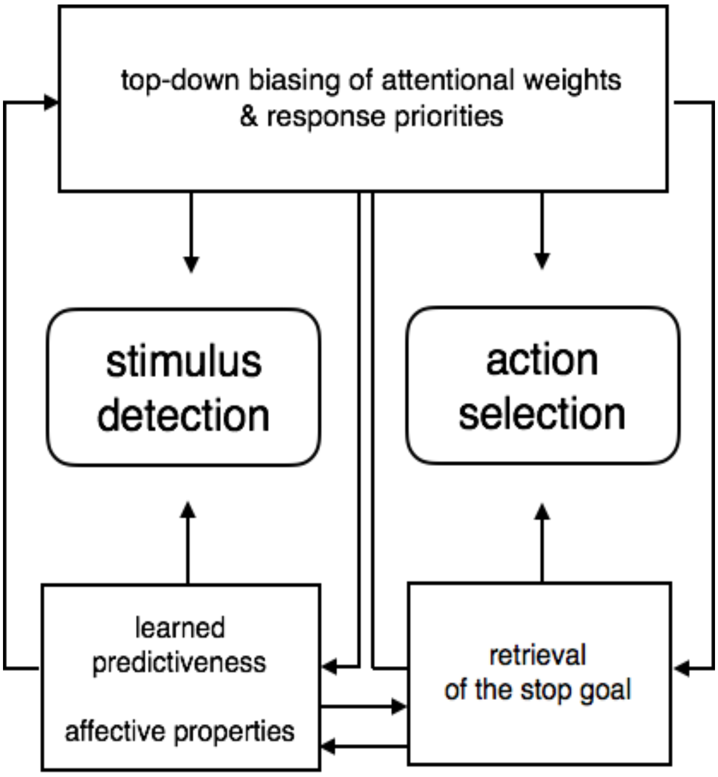

Figure 7.1. An visual depiction of how top-down control and bottom-up learning could influence stimulus detection and action selection processes during performance of response inhibition tasks.

Is 'automatic inhibition' truly automatic? Before proceeding, it is important to briefly address the extent to which 'automatic inhibition' is driven by 'automatic' processes. After all, a defining feature of automatic processes is the idea that they occur independent of top-down influences (e.g. Posner & Snyder, 1975). This is in direct contrast to the account provided above in which the effects of stimulus-stop learning on behaviour can be modulated by top-down control settings (see also

Verbruggen & Logan, 2009c; for a discussion of this research, see *Response inhibition as a 'prepared reflex',* Chapter 1). In this way, 'automatic inhibition' clearly does not fit with the traditional description of 'automatic' processing. Nevertheless, I highlight that this does not rule out the idea that stimulus-stop learning is automatic in the sense that it relies on memory-retrieval processes (consistent with theories of automaticity; Logan, 1988). Indeed, the research presented in this thesis provides strong support for the idea that inhibition can become 'automatic' in that it can be triggered by the retrieval of direct stimulus-stop associations from memory. In other words, stop learning can be 'automatic' in the Logan sense because the acquisition of stimulus-stop associations can facilitate the automatisation of response inhibition which may, on some trials, rely exclusively on the retrieval of stimulus-stop associations. However, stop learning is not 'automatic' in the Posner and Snyder sense because it is can still be influenced by top-down attentional and executive control settings. Future research is required to test exactly how top-down and bottom-up control interact in response inhibition tasks and to examine the factors that influence the relative contributions of top-down and bottom-up factors.

Implications for Applied Research

In addition to contributing to our theoretical understanding of stimulus-stop learning, this thesis also has implications for more applied research. In this section, I highlight three of the main implications.

First, the results presented in Chapter 2 indicate that attentional settings influence learning in response inhibition tasks. Even when salient images are used as stimuli (e.g. as in the food studies; see *Practical Applications*, Chapter 1), participants may still adjust their attentional settings, and ignore the images to a certain degree. Currently, the task-relevance of the stop-associated images used in applied stop-training studies varies. Whilst the task-relevance of the images may not influence engagement in impulsive behaviours (e.g. impulsive eating can be prompted by incidental processing of food cues in the environment; see, e.g. Lawrence, Hinton, Parkinson, & Lawrence, 2012), the results presented in this thesis suggest that designs in which participants must attend to the images should produce 'stronger' stimulus-stop associations that will have a more pronounced influence on stop-learning.

Second, in order to maximise the inhibitory control training effects, this thesis suggests that it is important to consider how other features of the stop learning task design could influence what is learned. For example, in Experiments 1-4, I devised a novel task that combined features of the go/no-go task and the stop-signal task. In Experiments 1-3, there was no delay between the presentation of the images and the go/stop signal (i.e. zero ms); in Experiment 4, there was a variable delay between the images and the go/stop signal but, unlike in typical stop-signal tasks, the go and stop signal occurred at the same moment (i.e. there was no delay between the go and the stop signals). But to avoid the possibility that subjects would simply wait on all trials, I used a low overall proportion of stop trials (0.25), imposed a relatively strict response deadline (750 ms) and provided feedback if the participant did not respond in time. Combined, this provided an effective design in which which to investigate stop learning it allowed me to manipulate the go/stop signal representation whilst maximising the number of correct stop trials. After all, previous work indicates that stimulus-stop associations are less likely to be learned when inhibition is unsuccessful (Verbruggen & Logan, 2008a, 2008b; Verbruggen, Logan, et al., 2008). The idea that the stop outcome is important is further supported by studies in the applied domain. Stop learning effects on task performance and on food- and alcohol consumption have been observed after both go/no-go and stop-signal training. However, a recent meta-analysis indicates that go/no-go training has stronger effects on appetitive behaviour than stop training (Jones, Di Lemma, et al., 2016). This could be due to generally higher success rates in the go/no-go task (Jones, Di Lemma, et al., 2016; Verbruggen & Logan, 2008a). Thus, I suggest that a consideration of features of applied stop training tasks, such as the timing, the proportion of stop trials, and the go response deadline, could be key to understanding the efficacy of training on food and alcohol consumption. Therefore, future research should manipulate these task features to discern how they influence outcome measures, such as food or alcohol consumption.

Third, a key concern in the applied domain is whether qualitatively different learning processes occur when the stimulus-stop associations are implicit (i.e. uninstructed) and when stimulus-stop associations are explicit (i.e. instructed). Although Chapter 3 (Experiments 5-6) indicates that expectancies do not necessarily correlate with performance, I find consistent evidence that participants generate expectancies that are consistent with the stimulus-stop contingencies in play (Experiments 1-8). This raises the question; to what extent are training effects in the inhibitory training domain due to expectancies and task demands (Boot, Simons, et al. 2013 raised a similar issue)? Furthermore, the extent to which the effects of stimulus-stop learning on expectancy ratings equates to other dependent variables used in the stop-training studies, such as food intake, is also unclear. Most studies in the inhibitory control training domain do not include a dependent variable that is sensitive enough to measure expectancies (see Newell & Shanks, 2014), such as expectancy ratings (e.g. Stothart, Simons, Boot & Kramer, 2014). However, the research presented in this thesis suggests that this would provide a useful addition to applied inhibitory control training studies.

Avenues for Future Research

Alongside the outstanding issues outlined above, there are also some more general issues relating to nature of associations acquired in response inhibition tasks that should be addressed in future research.

Alternative Routes

Alongside the direct and indirect pathways (see *What is Learned?,* Chapter 7), Verbruggen, Best, et al. (2014) also proposed a third possibility that stimulus-stop associations could influence behaviour through a link with the Pavlovian aversive and appetitive centres (Dickinson & Balleine, 2002;

Dickinson & Dearing, 1979; Konorski, 1967). Verbruggen, Best, et al. suggest that learning in response inhibition tasks could operate through links between the 'stop system' and the aversive system, and between the 'go system' and the appetitive system (for an overview, see McLaren & Verbruggen, 2014).

Consistent with this suggestion, several studies in the applied inhibitory control training literature have shown that images associated with stopping show a reduction in rated valence (Veling, Aarts, et al., 2013a; Veling, Holland, & van Knippenberg, 2008) and more negative implicit affective reactions (Houben, Nederkoorn, & Jansen, 2012b; Veling & Aarts, 2009). Furthermore, Ferrey, Frischen, and Fenske (2012) showed that stop associations not only influence the hedonic value of stop-associated stimuli but also the motivational value. In a first experiment, Ferrey, Frischen, et al. paired sexually attractive images with either stopping or going. Following a training phase, participants were required to rate the attractiveness of the images. These ratings showed that participants rated images that were paired with stopping as less attractive than the images associated with going. In a second experiment, Ferrey, Frischen, et al. showed that heterosexual males made fewer key presses to view the sexually-appealing images that were previously paired with stopping (no-go) than to un-appealing images or control (scrambled) images. Thus, not only does stop learning devalue the stop-associated stimuli but it could also influence the motivational value of stimuli, discouraging participants from responding to previously appealing 'approach' stimuli.

In the experiments presented in this thesis, I always used affectively neutral images or words. However, manipulating the appetitive/aversive properties of the stop-associated stimuli could be an effective way to enhance the effects of stimulus-stop learning on behaviour. For example, in addition to measuring the perceived value of the stop-associated stimuli at the end of training, it could be useful to compare the effects of stop learning for aversive stimuli (e.g. stimuli associated with monetary loss in a prior training task) compared with the effects of stop learning for appetitive stimuli (e.g. stimuli associated with monetary reward in a prior training task). If the appetitive/aversive systems play a role, I predict that the effects of stop learning should be enhanced for the aversive stimuli compared with the appetitive stimuli. This would have theoretical implications for our understanding of the links between the appetitive/aversive systems and action control, and could also provide a potential avenue to increase the effects of stimulus-stop learning in the applied inhibition training domain.

Mediation via a Category Representation

The majority of experiments in the stop learning literature can be placed into two groups: those in which stopping is trained by manipulating the pairings of specific stimuli and the presentation of a stop signal (e.g. as in Experiments 1-8 in this thesis) and those in which the stimulus-stop (no-go) associations are instructed by a stimulus-category rule (e.g. as in Experiments 9-10 in this thesis). As discussed above, this thesis found no evidence to support the idea that the stimulus-stop association is strongly mediated by a representation of the stop signal. However, this does not rule out the possibility that the stimulus-stop associations acquired in the go/no-go task are mediated via the go/no-go category (e.g., ‘desk = non-living -> non-living = no-go’, instead of ‘desk = no-go’; Verbruggen, Best, et al., 2014). Indeed, mediation of the stimulus-stop association via a category representation could even present an adaptive advantage in that it reduces the memory load of having to store all the individual stimuli and would instead allow for optimal performance by storing only the category rule (e.g. natural vs. human-made; for a similar idea in the context of task-switching, see Arrington & Logan, 2004).

Initial research by Verbruggen and Logan (2008a) supports the idea that participants acquire stimulus-category associations during performance of the go/no-go task; despite the reversal of the go/no-go associations in a test phase, Verbruggen and Logan found that performance was initially slightly better for the old no-go items than for novel items that were not presented during the training phase. This suggests that performance for the old items benefitted from the retrieval of stimulus-category associations acquired during training whereas the novel items could not benefit from prior stimulus-category learning (Verbruggen & Logan, 2008a, Experiment 1). Furthermore, in another experiment, Verbruggen and Logan (2008a, Experiment 2) showed that when a new categorisation task was introduced in the test phase, there was general (non-specific) cost between the end of the training phase and the test phase. This suggests that training in this task probably benefited from the consistent stimulus-category associations, such that when the new judgment task was introduced participants could no longer rely on the old category labels.

However, the research presented in this thesis suggests that mediation via the go/no-go category on performance is not an inevitable consequence of go/no-go task performance, especially when the stimulus set is small. In Experiments 9-10 (Chapter 5), I adapted the aforementioned Verbruggen and Logan (2008a, Experiment 2) task in order to examine the specificity of learning in the go/no-go task[67]. The paradigm comprised a training phase in which the stimulus-category and stimulus-response (go or no-go) mappings remained the same (e.g. diamond = go; square = no-go), followed by a test phase in which a new judgment task was introduced (e.g. ‘infrequent’ (3 x 3) chequerboard = go; ‘frequent’ (9 x 9) chequerboard = no-go). Thus, just as in the Verbruggen and Logan (2008a) paradigm, performance could benefit from stimulus-category associations during the training phase but the introduction of a new judgment task in the test phase meant that performance could no longer benefit from the retrieval of the acquired stimulus-category associations. The main differences between the task used in Experiments 9-10 and the Verbruggen and Logan (2008a, Experiment 2) version were that I used a smaller stimulus-set (4 items vs. 60 items) and that I used visual discrimination tasks

---

[67] Note that although there was a two-choice task condition in Experiments 9-10, I focus only on the go/no-go version here for comparison with the Verbruggen and Logan (2008a) experiment. However, there is no clear reason why the effects of category mediation would differentially influence performance in the go/no-go task more than in the two-choice task, or vice versa.

(diamond/square; frequent/infrequent chequerboards) rather than semantic judgment tasks (living/non-living; smaller/bigger). The results showed that, unlike in the Verbruggen and Logan (2008a, Experiment 2) experiment, there was no reliable difference in performance between the final blocks in the training phase and blocks in the test phase.

Therefore, I propose that the extent of mediation via a category could depend, at least in part, on the stimulus-set size; when the stimulus-set is large, participants are more likely to rely on abstract category representations (thus acquiring mediated stimulus-category-stop associations) and when the stimulus-set is small(er) participants could acquire more direct stimulus-stop associations. The suggestion that stimulus set-size may influence the acquisition of stimulus-stop versus more abstract category-stop associations is consistent with research in the task-switching literature which suggests that people tend to use 'task-sets'[68] when the stimulus-set is large but rely on associations when the stimulus-set is small (Forrest, Monsell, & McLaren, 2014). Future research should test this idea by varying the stimulus set-size in go/no-go training tasks and measuring the magnitude of the cost between the end of training and the test phase. It could also be useful to sufficiently debrief participants in that experiment in order to discern whether they based performance on a category representation (e.g. living/non-living) or whether they relied on remembering individual stimulus-stop (and stimulus-go) associations.

Transfer to a New Stimulus Set

A related outstanding issue concerns the extent to which stop learning transfers to novel settings. Whilst the potential effects of stimulus-signal learning on performance were construed by Verbruggen, Best, et al. (2014) as being highly specific to trials in which the trained signal appeared, mediation of the stimulus-stop association via a category representation may actually be beneficial as it could indicate a higher level of abstraction held within the associations. For example, it has been suggested that rules have 'compositional hierarchy' in which more abstract representations can modulate more specific stimulus-action representations (e.g. rather than apple -> stop, apple = fruit = stop; Cole, Laurent, & Stocco, 2013). Thus, if stimulus-stop associations become mediated by a category representation this could encourage transfer to a novel stimuli. Consistent with this idea, it has recently been found that participants are able to generalise category-response associations to novel stimuli whereas stimulus-response associations do not improve performance above baseline for novel stimuli (Longman, Milton, Wills, & Verbruggen, under review). This suggests that the acquisition of category-stop associations could even facilitate learning in novel contexts. However, category-stop learning and stimulus-specific stop learning have not been directly compared and it has not been investigated how stimulus-specific stop learning could transition into abstract category-stop learning.

---

[68] A 'task set' refers to the configuration of the cognitive system that is required to perform a given task (e.g. Rogers & Monsell, 1995).

A consideration of how stop learning could transfer to novel stimulus sets and the situations under which transfer could occur was beyond the scope of the research presented in this thesis. However, alongside being of theoretical interest, transfer is a central concern to research in the applied inhibitory control domain (for an overview of this research, see *Practical Applications,* Chapter 1). Indeed, some effects of applied inhibitory control training on food and alcohol consumption have been reported outside of the lab suggesting that some level of abstraction must occur in these studies (as people never see the exact training images in real life; for an example, see Houben, Havermans, et al. 2012). As a first step, future research should investigate the acquisition of stimulus-category stop rules by using stimuli in which a category-stop rule can be abstracted during training. It seems likely that the acquisition of an abstract category-stop association would transfer to novel situations more than the acquisition of a stimulus-specific stop association. To test this idea in future research, one could insert a 'transfer' phase with a new stimulus set (e.g. a new set of living/non-living words) following training. If participants acquired specific stimulus-stop associations, the effects of training should not transfer to the novel stimuli, whereas if participants abstracted a category-stop rule during the course of learning some transfer to to the novel items should be observed.

Concluding Remarks

In conclusion, this thesis has demonstrated that response inhibition can become a 'bottom-up' act of control driven by the retrieval of stimulus-stop associations from memory. However, a major contribution of this thesis is to highlight that the effects of stop learning on behaviour most likely depends on an interplay of top-down and bottom-up factors. To account for these findings, I proposed how top-down and bottom-up factors could interact during stimulus detection and action selection stages of processing. Considering how top-down and bottom-up control interacts, and establishing the factors that influence the relative contributions of top-down and bottom-up factors are key questions for future research. In this way, 'banishing the control homunculus' from theories of action control lies not in distinguishing between top-down and bottom-up factors, but in understanding how top-down and bottom-up factors interact in order to give rise to a flexible, adaptive, efficient system capable of outsourcing control where possible and overriding bottom-up influences whenever the situation requires. Therefore, this thesis calls for researchers in the action control domain to consider top-down and bottom-up control as dependent, rather than independent, influences on behaviour.

Appendix M

RT Percentiles

To investigate the possibility that the effects of stop learning were sensitive to response speed, I plotted RT percentiles for the item-specific training data in Experiments 5-6 (combined; Chapter 3) and Experiment 8 (Chapter 4). Under conditions where the stimulus-stop contingencies were not instructed (Figure M.1(b) and Figure M.2) there was a trend towards the slowing for the stop-associated items emerging at the faster end of the distribution. However, the interactions between stimulus type and percentile did not reach statistical significance in Experiments 5-6 ($p = 0.255$) or in Experiment 8 ($p = 0.390$). When the stimulus-stop contingencies were acquired via explicit instructions, there was no similar trend (Figure M.1(a); Experiments 5-6); the slowing was largest in the middle of the RT distribution (the stimulus-type by percentile interaction was reliable, $p < 0.001$).

RT percentiles for Experiments 1-4 can be found in Appendix B. To briefly recap, these percentile analyses revealed that for Experiments 1-3 the slowing for the stop-associated stimuli emerged at the slow end of the distribution (which I attributed to the processing of the images) whereas for Experiment 4 the slowing for the stop-associated stimuli emerged at the fast end of the RT distribution (which I attributed to the presentation of the images prior to the presentation of the stop/go signal).

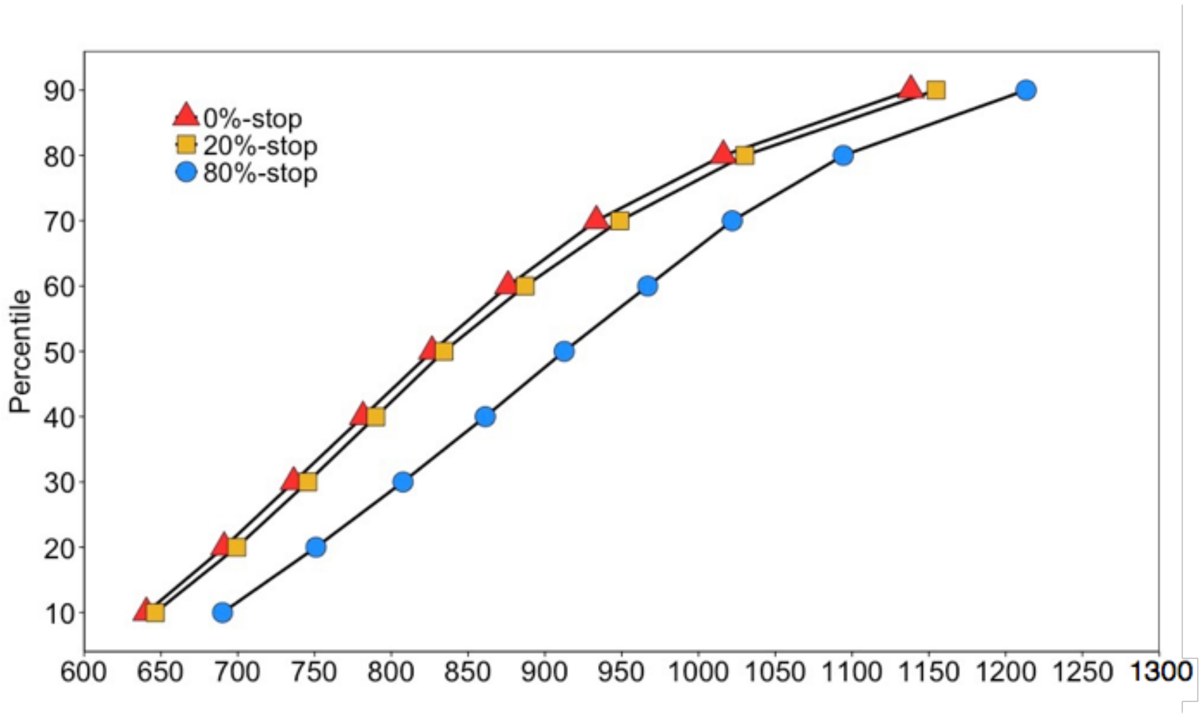

Figure M.1. Go RTs (in ms) in Experiments 5-6 (combined) in the instructed condition for the three stimulus-types (0%-stop, 20%-stop, 80%-stop) as a function of percentile.

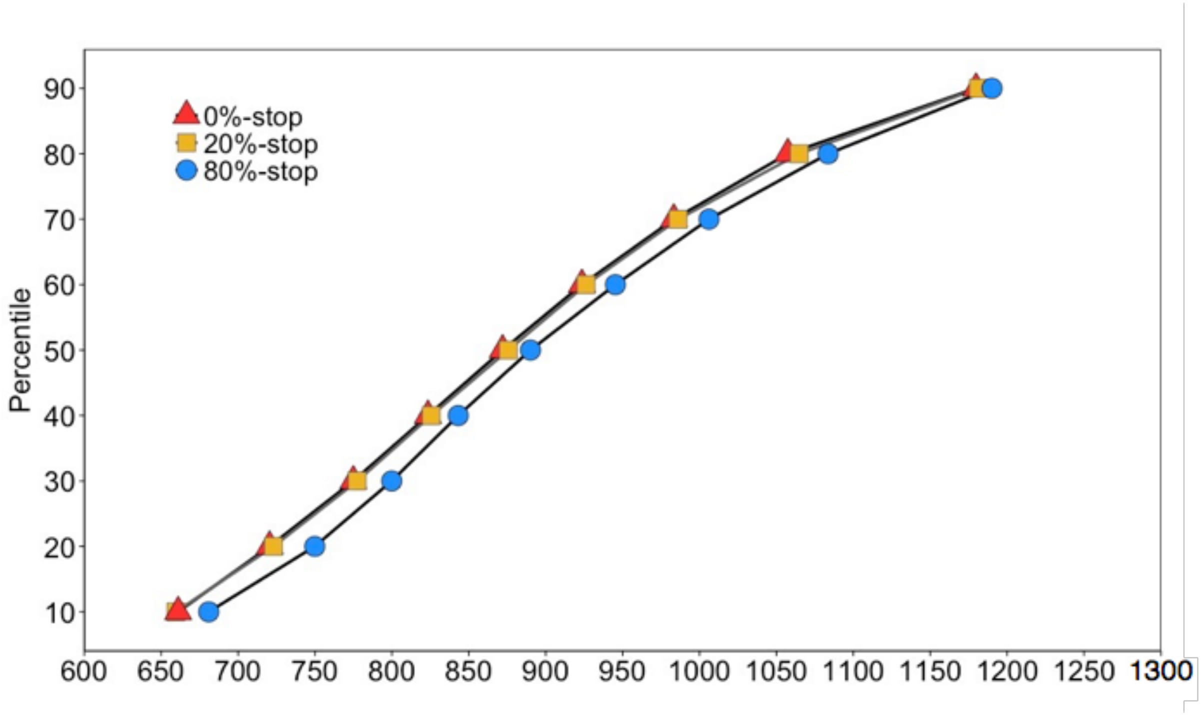

Figure M.2. Go RTs (in ms) in Experiments 5-6 (combined) in the uninstructed condition for the three stimulus-types (0%-stop, 20%-stop, 80%-stop) as a function of percentile.

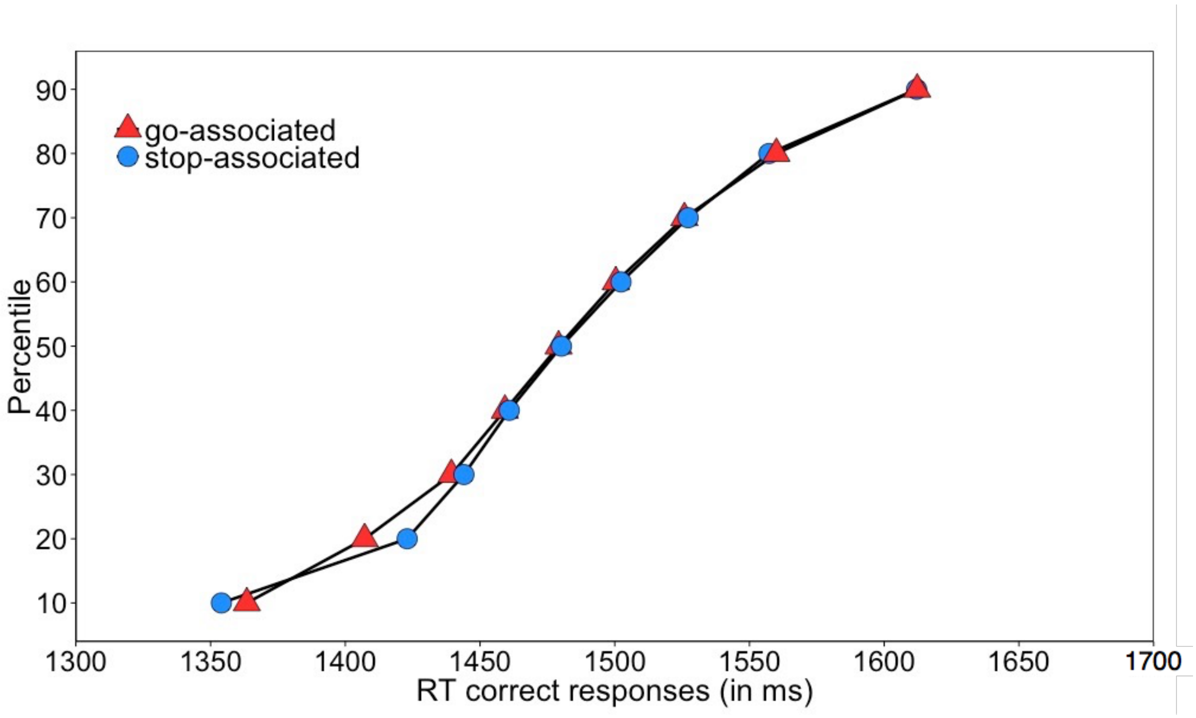

Figure M.3. Go RTs (in ms) in Experiment 8 for the two stimulus-types (go-associated, stop-associated) as a function of percentile.

REFERENCES

Adams, R. C., & Chambers, C. D. (2012). Mapping the timecourse of goal-directed attention to location and colour in human vision. *Acta Psychologica*, *139*(3), 515–523. doi:10.1016/j.actpsy.2012.01.014

Allom, V., & Mullan, B. (2015). Two inhibitory control training interventions designed to improve eating behaviour and determine mechanisms of change. *Appetite*, *89*, 282–290. doi:10.1016/j.appet.2015.02.022

Allom, V., Mullan, B., & Hagger, M. (2015). Does inhibitory control training improve health behaviour? A meta-analysis. *Health Psychology Review*, *7199*, 1–38. doi:10.1080/17437199.2015.1051078

Andersen, S. K., Müller, M. M., & Hillyard, S. A. (2009). Color-selective attention need not be mediated by spatial attention. *Journal of Vision*, *9*(6), 1–7. doi:10.1167/9.6.2.

Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences USA*, *108*(25), 10367–10371. doi:10.1073/pnas.1104047108

Anderson, B. A., & Folk, C. L. (2012). Contingent involuntary motoric inhibition: the involuntary inhibition of a motor response contingent on top-down goals. *Journal of Experimental Psychology. Human Perception and Performance*, *38*(6), 1348–1352. doi:10.1037/a0030514

Anderson, B. A., & Folk, C. L. (2014). Conditional automaticity in response selection: contingent involuntary response inhibition with varied stimulus-response mapping. *Psychological Science*, *25*(2), 547–554. doi:10.1177/0956797613511086

Anselme, P., Robinson, M. J. F., & Berridge, K. C. (2013). Reward uncertainty enhances incentive salience attribution as sign-tracking. *Behavioural Brain Research*, *238*, 53–61. doi:10.1016/j.bbr.2012.10.006

Aron, A. R. (2007). The neural basis of inhibition in cognitive control. *The Neuroscientist*, *13*(3), 214–228. doi:10.1177/1073858407299288

Aron, A. R., Durston, S., Eagle, D. M., Logan, G. D., Stinear, C. M., & Stuphorn, V. (2007). Converging evidence for a fronto-basal-ganglia network for inhibitory control of action and cognition. *The Journal of Neuroscience*, *27*(44), 11860–11864. doi:10.1523/JNEUROSCI.3644-07.2007

Aron, A. R., & Verbruggen, F. (2008). Stop the presses! Dissociating a selective from a global mechanism for stopping. *Psychological Science*, *19*(11), 1146–1153. doi:10.1111/j.1467-9280.2008.02216.x

Aron, A. R. (2011). From reactive to proactive and selective control: Developing a richer model for stopping inappropriate responses. *Biological Psychology*, *69*, e55–e68. doi:10.1016/j.biopsych.2010.07.024

Aron, A. R., & Poldrack, R. A. (2006). Cortical and subcortical contributions to stop signal response inhibition: Role of the subthalamic nucleus. *The Journal of Neuroscience*, *26*(9), 2424–2433. doi:10.1523/JNEUROSCI.4682-05.2006

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2014). Inhibition and the right inferior frontal cortex: One decade on. *Trends in Cognitive Sciences*, *18*(4), 177–185. doi:10.1016/j.tics.2013.12.003

Aron, A. R., Robbins, T. W., & Poldrack, R. A. (2004). Inhibition and the right inferior frontal cortex. *Trends in Cognitive Sciences*, *8*(4), 170–177. doi:10.1016/j.tics.2004.02.010

Arrington, C. M., & Logan, G. D. (2004). The cost of a voluntary task switch. *Psychological Science*, *15*(9), 610–615. doi:10.1111/j.0956-7976.2004.00728.x

Ashby, F. G., Ell, S. W., & Waldron, E. M. (2003). Procedural learning in perceptual categorization. *Memory & Cognition*, *31*(7), 1114–1125. doi:10.3758/BF03196132

Atienza, M., Cantero, J. L., & Dominguez-Marin, E. (2002). The time course of neural changes underlying auditory perceptual learning. *Learning & Memory*, *9*(3), 138–150. doi:10.1101/lm.46502

Awh, E., Belopolsky, A. V., & Theeuwes, J. (2012). Top-down versus bottom-up attentional control: A failed theoretical dichotomy. *Trends in Cognitive Sciences*, *16*(8), 437–443. doi:10.1016/j.tics.2012.06.010

Band, G. P. H., Ridderinkhof, K. R., & van der Molen, M. W. (2003). Speed-accuracy modulation in case of conflict: the roles of activation and inhibition. *Psychological Research*, *67*(4), 266–279. doi:10.1007/s00426-002-0127-0

Barcelo, F., Munoz-Cespedes, Juan, M., Pozo, M. A., & Francisco, J. R. (2000). Attentional set shifting modulates the target P3b response in the Wisconsin card sorting test. *Neuropsychologia*, *38*(10), 1342–1355. doi:10.1016/S0028-3932(00)00046-4

Bari, A., & Robbins, T. W. (2013). Inhibition and impulsivity: Behavioral and neural basis of response control. *Progress in Neurobiology*, *108*, 44–79. doi:10.1016/j.pneurobio.2013.06.005

Bechara, A., Noel, X., & Crone, E. (2006). Loss of willpower: Abnormal neural mechanisms of impulse control and decision making in addiction. In R. W. Wiers & A. W. Stacy (Eds.), *Handbook of implicit cognition and addiction* (pp. 225–232). Thousand Oaks, CA: Sage Publications.

Bekker, E. M., Kenemans, J. L., Hoeksma, M. R., Talsma, D., & Verbaten, M. N. (2005). The pure electrophysiology of stopping. *International Journal of Psychophysiology*, *55*(2), 191–198. doi:10.1016/j.ijpsycho.2004.07.005

Bell, A. J., & Sejnovski, T. J. (1995). An information-maximisation approach to blind separation and blind deconvolution. *Neural Computation*, 7, 1004–1034.

Bender, A. D., Filmer, H. L., Garner, K. G., Naughtin, C. K., & Dux, P. E. (2016). On the relationship between response selection and response inhibition: An individual differences approach. *Attention, Perception, & Psychophysics*. doi:10.3758/s13414-016-1158-8

Benikos, N. (2014). *Optimising the training-induced changes of inhibitory control* (Unpublished doctoral dissertation). University of Wollongong, New South Wales.

Benikos, N., Johnstone, S. J., & Roodenrys, S. J. (2013). Short-term training in the Go/Nogo task: Behavioural and neural changes depend on task demands. *International Journal of Psychophysiology*, *87*(3), 301–312. doi:10.1016/j.ijpsycho.2012.12.001

Berkman, E. T., Kahn, L. E., & Merchant, J. S. (2014). Training-induced changes in inhibitory control network activity. *The Journal of Neuroscience*, *34*(1), 149–157. doi:10.1523/JNEUROSCI.3564-13.2014

Berti, S. (2008). Object switching within working memory is reflected in the human event-related brain potential. *Neuroscience Letters*, *434*(2), 200–205. doi:10.1016/j.neulet.2008.01.055

Best, M., Lawrence, N. S., Logan, G. D., McLaren, I. P. L., & Verbruggen, F. (2015). 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. doi:10.1037/xhp0000116

Bestmann, S. (2012). Functional modulation of primary motor cortex during action selection. In R. Chen & J. C. Rothwell (Eds.), *Cortical connectivity* (pp. 183–205). Berlin, Germany: Springer.

Bissett, P. G., & Logan, G. D. (2011). Balancing cognitive demands: control adjustments in the stop-signal paradigm. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *37*(2), 392–404. doi:10.1037/a0021800

Bokura, H., Yamaguchi, S., & Kobayashi, S. (2001). Electrophysiological correlates for response inhibition in a Go/NoGo task. *Clinical Neurophysiology*, *112*(12), 2224–2232. doi:10.1016/S1388-2457(01)00691-5

Boot, W. R., Simons, D. J., Stothart, C. ., & Stutts, C. (2013). The pervasive problem with placebos in psychology: Why active control groups are not sufficient to rule out placebo effects. *Perspectives on Psychological Science*, *8*(4), 445–454. doi:10.1177/1745691613491271

Boronat, C. B., & Logan, G. D. (1997). The role of attention in automatization: does attention operate at encoding, or retrieval, or both? *Memory & Cognition*, *25*(1), 36–46. doi:10.3758/BF03197283

Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Evaluating the demand for control: Anterior cingulate cortex and conflict monitoring. *Psychological Review*, *108*(3), 624–652. doi:10.1038/46035

Botvinick, M. M., Cohen, J. D., & Carter, C. S. (2004). Conflict monitoring and anterior cingulate cortex: An update. *Trends in Cognitive Sciences*, *8*(12), 539–546. doi:10.1016/j.tics.2004.10.003

Boucher, L., Palmeri, T. J., Logan, G. D., & Schall, J. D. (2007). Inhibitory control in mind and brain: an interactive race model of countermanding saccades. *Psychological Review*, *114*(2), 376–397. doi:10.1037/0033-295X.114.2.376

Bowditch, W.A. (2016). Response inhibition and associative learning: Training stimulus specific response inhibition (Unpublished doctoral dissertation). University of Exeter, UK.

Bowditch, W. A., Verbruggen, F., & McLaren, I. P. L. (2015). Associatively mediated stopping: Training stimulus-specific inhibitory control. *Learning & Behavior*, 1–13. doi:10.3758/s13420-015-0196-8

Bowley, C., Faricy, C., Johnstone, S. J., & Smith, J. L. (2013). The effects of inhibitory control training on alcohol consumption, implicit alcohol-related cognitions and brain electrical activity. *International Journal of Psychophysiology*, *89*, 342–348.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, *10*(4), 433–436.

Bruin, K. J., Wijers, A. A, & van Staveren, A. S. J. (2001). Response priming in a go/nogo task: do we have to explain the go/nogo N2 effect in terms of response activation instead of inhibition? *Clinical Neurophysiology : Official Journal of the International Federation of Clinical Neurophysiology*, *112*(9), 1660–1671. doi:10.1016/S1388-2457(01)00601-0

Bruin, K. J., & Wijers, A. A. (2002). Inhibition, response mode, and stimulus probability: A comparative event-related potential study. *Clinical Neurophysiology*, *113*(7), 1172–1182. doi: 10.1016/S1388-2457(02)00141-4

Buch, E. R., Mars, R. B., Boorman, E. D., & Rushworth, M. F. S. (2010). A network centered on ventral premotor cortex exerts both facilitatory and inhibitory control over primary motor cortex during action reprogramming. *Journal of Neuroscience*, *30*(4), 1395–1401. doi:10.1523/JNEUROSCI.4882-09.2010

Bunge, S. A. (2004). How we use rules to select actions: A review of evidence from cognitive neuroscience. *Cognitive, Affective, & Behavioral Neuroscience*, *4*(4), 564–579. doi:10.3758/CABN.4.4.564

Chambers, C. D., Bellgrove, M. A., Gould, I. C., English, T., Garavan, H., McNaught, E., Kamke, M., & Mattingley, J. B. (2007). Dissociable mechanisms of cognitive control in prefrontal and premotor cortex. *Journal of Neurophysiology*, *98*(6), 3638–3647. doi:10.1152/jn.00685.2007

Chambers, C. D., Bellgrove, M. A., Stokes, M. G., Henderson, T. R., Garavan, H., Robertson, I. H., Morris, A. P., & Mattingley, J. B. (2006). Executive "brake failure" following deactivation of human frontal lobe. *Journal of Cognitive Neuroscience*, *18*(3), 444–455. doi: 10.1162/089892906775990606

Chambers, C. D., Garavan, H., & Bellgrove, M. A. (2009). Insights into the neural basis of response inhibition from cognitive and clinical neuroscience. *Neuroscience and Biobehavioral Reviews*, *33*(5), 631-646. doi:10.1016/j.neubiorev.2008.08.016

Chavan, C. F., Mouthon, M., Draganski, B., van der Zwaag, W., & Spierer, L. (2015). Differential patterns of functional and structural plasticity within and between inferior frontal gyri support training-induced improvements in inhibitory control proficiency. *Human Brain Mapping*, *36*(7), 2527–2543. doi:10.1002/hbm.22789

Chein, J. M., & Schneider, W. (2012). The brain's learning and control architecture. *Current Directions in Psychological Science*, *21*, 78–84. doi:10.1177/0963721411434977

Chevrier, A., Noseworthy, M. D., & Schachar, R. (2007). Dissociation of response inhibition and performance monitoring in the stop signal task using event-related fMRI. *Human Brain Mapping*, *28*, 1347-1358. doi:10.1002/hbm.20355

Chikazoe, J., Jimura, K., Hirose, S., Yamashita, K., Miyashita, Y., & Konishi, S. (2009). Preparation to inhibit a response complements response inhibition during performance of a stop-signal task. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, *29*(50), 15870–15877. doi:10.1523/JNEUROSCI.3645-09.2009

Chiu, Y.C., & Aron, A. R. (2014). Unconsciously Triggered Response Inhibition Requires An Executive Setting. *Journal of Experimental Psychology: General*, *143*(1), 56–61. doi:10.1037/a0031497

Chiu, Y.C., Aron, A. R., & Verbruggen, F. (2012). Response suppression by automatic retrieval of stimulus-stop association: evidence from transcranial magnetic stimulation. *Journal of Cognitive Neuroscience*, *24*(9), 1908–18. doi:10.1162/jocn_a_00247

Cohen-Kdoshay, O., & Meiran, N. (2009). The representation of instructions operates like a prepared reflex: flanker compatibility effects found in first trial following S-R instructions. *Experimental Psychology*, *56*(2), 128–133. doi:10.1027/1618-3169.56.2.128

Cohen, J. R., & Poldrack, R. A. (2008). Automaticity in motor sequence learning does not impair response inhibition. *Psychonomic Bulletin & Review*, *15*(1), 108–115. doi:10.3758/PBR.15.1.108

Cohen, M. X., & Frank, M. J. (2009). Neurocomputational models of basal ganglia function in learning, memory and choice. *Behavioural Brain Research*, *199*(1), 141–156. doi:10.1016/j.bbr.2008.09.029

Cole, M. W., Laurent, P., & Stocco, A. (2013). Rapid instructed task learning: a new window into the human brain's unique capacity for flexible cognitive control. *Cognitive, Affective & Behavioral Neuroscience*, *13*(1), 1–22. doi:10.3758/s13415-012-0125-7

Coles, M. G. H. (1989). Modern mind-brain reading: Psychophysiology, physiology, and Ccgnition. *Psychophysiology*, *26*(3), 251–269. doi:10.1111/j.1469-8986.1989.tb01916.x

Colzato, L. S., Van Den Wildenberg, W. P. M., Van Wouwe, N. C., Pannebakker, M. M., & Hommel, B. (2009). Dopamine and inhibitory action control: Evidence from spontaneous eye blink rates. *Experimental Brain Research*, *196*(3), 467–474. doi:10.1007/s00221-009-1862-x

Connor, C. E., Egeth, H. E., & Yantis, S. (2004). Visual attention: Bottom-up versus top-down. *Current Biology, 14,* R850-R852. doi:10.1016/j.cub.2004.09.041

Cornelissen, F. W., Peters, E. M., & Palmer, J. (2002). The eyelink toolbox: Eye tracking with MATLAB and the psychophysics Toolbox. *Behavior Research Methods, Instruments, & Computers*, *34*(4), 613–617. doi:10.3758/BF03195489

Cramer, A. O. J., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P. P. P., Waldorp, L., & Wagenmakers, E.J. (2016). Hidden multiplicity in exploratory multiway ANOVA: Prevalence and remedies. *Psychonomic Bulletin & Review*, *23*(2), 640–647. doi: 10.3758/s13423-015-0913-5

Crews, F. T., & Boettiger, C. A. (2009). Impulsivity, frontal lobes and risk for addiction. *Pharmacology, Biochemistry, and Behavior*, *93*(3), 237–47. doi:10.1016/j.pbb.2009.04.018

de Wit, H. (2009). Impulsivity as a determinant and consequence of drug use: a review of underlying processes. *Addiction Biology*, *14*, 22–31. doi:10.1111/j.1369-1600.2008.00129.x

Della Libera, C., & Chelazzi, L. (2009). Learning to attend and to ignore is a matter of gains and losses. *Psychological Science*, *20*(6), 778–784. doi:10.1111/j.1467-9280.2009.02360.x

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18*, 193-222. doi:10.1146/annurev.neuro.18.1.193

Dickinson, A., & Balleine, B. (2002). The role of learning in the operation of motivational systems. In C. R. Gallistel (Ed.), *Steven's handbook of experimental psychology, Vol. 3. Learning, motivation, & emotion* (pp. 497–533). New York: Wiley.

Dickinson, A., & Dearing, M. F. (1979). Appetitive-aversive interactions and inhibitory processes. In A. Dickinson & R. A. Boakes (Eds.), *Mechanism of learning and motivation* (pp. 203–231). Hilldale, N.J: Erlbaum.

Dienes, Z. (2011). Bayesian versus orthodox statistics: Which side are you on? *Perspectives on Psychological Science*, *6*(3), 274–290. doi:10.1177/1745691611406920

Dienes, Z. (2007). Subjective measures of unconscious knowledge. *Progress in Brain Research*, *168*, 1–45. doi:10.1016/S0079-6123(07)68005-4

Dienes, Z., Altmann, G. T. M., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *2*(5), 1322–1338. doi:10.1037/0278-7393.21.5.1322

Ditye, T., Jacobson, L., Walsh, V., & Lavidor, M. (2012). Modulating behavioral inhibition by tDCS combined with cognitive training. *Experimental Brain Research*, *219*(3), 363–368. doi: 10.1007/s00221-012-3098-4

Dodds, C. M., Morein-Zamir, S., & Robbins, T. W. (2011). Dissociating inhibition, attention, and response control in the frontoparietal network using functional magnetic resonance imaging. *Cerebral Cortex*, *21*(5), 1155–1165. doi:10.1093/cercor/bhq187

Donchin, E. (1981). Surprise!... Surprise? *Psychophysiology*, *18*(5), 493–513. doi:10.1111/j.1469-8986.1981.tb01815.x

Donchin, E., & Coles, M. G. H. (1988). Is the P300 component a manifestation of context updating? *Behavioral and Brain Sciences*, *11*(03), 357-427. doi:10.1017/S0140525X00058027

Donders, F. C. (1868/1969). On the speed of mental processes. *Acta Psychologica*, *30*, 412–431. doi: 10.1016/0001-6918(69)90065-1

Donkers, F. C. L., & Van Boxtel, G. J. M. (2004). The N2 in go/no-go tasks reflects conflict monitoring not response inhibition. *Brain and Cognition*, *56*, 165–176. doi:10.1016/j.bandc.2004.04.005

Dreisbach, G., & Haider, H. (2009). How task representations guide attention: further evidence for the shielding function of task sets. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*(2), 477–486. doi:10.1037/a0014647

Dreisbach, G., & Haider, H. (2008). That's what task sets are for: Shielding against irrelevant information. *Psychological Research*, *72*(4), 355–361. doi:10.1007/s00426-007-0131-5

Dreisbach, G., & Wenke, D. (2011). The shielding function of task sets and its relaxation during task switching. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(6), 1540–1546. doi: 10.1037/a0024077

Duncan-Johnson, C. C., & Donchin, E. (1977). On quantifying surprise: The variation of event-related potentials with subjective probability. *Psychophysiology*, *14*(5), 456–467. doi:10.1111/j.1469-8986.1977.tb01312.x

Duncan, J. (2006). EPS Mid-Career Award 2004: Brain mechanisms of attention. *Quarterly Journal of Experimental Psychology (2006)*, *59*(1), 2–27. doi:10.1080/17470210500260674

Duncan, J., & Owen, A. M. (2000). Common regions of the human frontal lobe recruited by diverse cognitive demands. *Trends in Neurosciences*, *23*(10), 475–483. doi:10.1016/S0166-2236(00)01633-7

Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature Neuroscience*, *8*(12), 1784–1790. doi: 10.1038/nn1594

Eimer, M., & Coles, M. G. H. (2003). The lateralized readiness potential as an on-line measure of central response activation processes. *Behavior Research Methods, Instruments, & Computers*, *30*(1), 146–156. doi:10.3758/BF03209424

Eimer, M. (1993). Effects of attention and stimulus probability on ERPs in a Go/Nogo task. *Biological Psychology*, *35*(2), 123–138. doi:10.1016/0301-0511(93)90009-W

Elchlepp, H., Lavric, A., Chambers, C. D., & Verbruggen, F. (2016). Proactive inhibitory control: A general biasing account. *Cognitive Psychology*, *86*, 27–61. doi:10.1016/j.cogpsych.2016.01.004

Enriquez-Geppert, S., Konrad, C., Pantev, C., & Huster, R. J. (2010). Conflict and inhibition differentially affect the N200/P300 complex in a combined go/nogo and stop-signal task. *NeuroImage*, *51*(2), 877–887. doi:10.1016/j.neuroimage.2010.02.043

Enticott, P. G., Bradshaw, J. L., Bellgrove, M. a, Upton, D. J., & Ogloff, J. R. P. (2009). Stop task after-effects: The extent of slowing during the preparation and execution of movement. *Experimental Psychology*, *56*(4), 247–251. doi:10.1027/1618-3169.56.4.247

Falkenstein, M., Hoormann, J., & Hohnsbein, J. (1999). ERP components in Go/Nogo tasks and their relation to inhibition. *Acta Psychologica*, *101*(2-3), 267–291. doi:10.1016/S0001-6918(99)00008-6

Falkenstein, M., Koshlykova, N. A., Kiroj, V. N., Hoormann, J., & Hohnsbein, J. (1995). Late ERP components in visual and auditory Go/Nogo tasks. *Electroencephalography and Clinical Neurophysiology/ Evoked Potentials*, *96*(1), 36–43. doi:10.1016/0013-4694(94)00182-K

Fenske, M. J., & Eastwood, J. D. (2003). Modulation of focused attention by faces expressing emotion: evidence from flanker tasks. *Emotion*, *3*(4), 327–343. doi:10.1037/1528-3542.3.4.327

Fernie, G., Peeters, M., Gullo, M. J., Christiansen, P., Cole, J. C., Sumnall, H., & Field, M. (2013). Multiple behavioural impulsivity tasks predict prospective alcohol involvement in adolescents. *Addiction*, *108*(11), 1916–1923. doi:10.1111/add.12283

Ferrey, A. E., Frischen, A., & Fenske, M. J. (2012). Hot or not: Response inhibition reduces the hedonic value and motivational incentive of sexual stimuli. *Frontiers in Psychology*, *3*, 575. doi:10.3389/fpsyg.2012.00575

Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, *18*(4), 1030–1044. doi:10.1037/0096-1523.18.4.1030

Folstein, J. R., & Van Petten, C. (2008). Influence of cognitive control and mismatch on the N2 component of the ERP: A review. *Psychophysiology*, *45*(1), 152–170. doi:10.1111/j.1469-8986.2007.00602.x

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. doi:10.1037/a0035981

Frank, M. J. (2006). Hold your horses: A dynamic computational role for the subthalamic nucleus in decision making. *Neural Networks*, *19*(8), 1120–1136. doi:10.1016/j.neunet.2006.03.006

Friedman, D., Cycowicz, Y. M., & Gaeta, H. (2001). The novelty P3: An event-related brain potential (ERP) sign of the brain's evaluation of novelty. *Neuroscience and Biobehavioral Reviews*, *25*(4), 355–373. doi:10.1016/S0149-7634(01)00019-7

Frings, C., & Moeller, B. (2012). The horserace between distractors and targets: Retrieval-based probe responding depends on distractor–target asynchrony. *Journal of Cognitive Psychology*, *24*(5), 582–590. doi:10.1080/20445911.2012.666852

Gajewski, P. D., & Falkenstein, M. (2013). Effects of task complexity on ERP components in Go/Nogo tasks. *International Journal of Psychophysiology : Official Journal of the International Organization of Psychophysiology*, *87*(3), 273–278. doi:10.1016/j.ijpsycho.2012.08.007

Garavan, H., & Stout, J. C. (2005). Neurocognitive insights into substance abuse. *Trends in Cognitive Sciences*, *9*(4), 195–201. doi:10.1016/j.tics.2005.02.008

Gaschler, R., & Nattkemper, D. (2012). Instructed task demands and utilization of action effect anticipation. *Frontiers in Psychology*, *3*, 578. doi:10.3389/fpsyg.2012.00578

Gebauer, G. F., & Mackintosh, N. J. (2007). Psychometric intelligence dissociates implicit and explicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *33*(1), 34–54. doi:10.1037/0278-7393.33.1.34

Giesen, C., & Rothermund, K. (2014). You better stop! Binding "stop" tags to irrelevant stimulus features. *Quarterly Journal of Experimental Psychology (2006)*, *67*(4), 809–832. doi: 10.1080/17470218.2013.834372

Gollwitzer, P. M. (1999). Implementation intentions. *American Psychologist*, *54*(7), 493–503. doi: 10.1177/0146167207311201

Golomb, J., & Turk-Browne, N. B. (2010). A taxonomy of external and internal attention. *Annual Review of Psychology*, *62*(1), 73–101. doi:10.1146/annurev.psych.093008.100427

Gomez, P., Ratcliff, R., & Perea, M. (2007). A model of the go/no-go task. *Journal of Experimental Psychology: General*, *136*(3), 389–413. doi:10.1037/0096-3445.136.3.389

Gratton, G., Coles, M.G.H., Sirevaag, E.J., Eriksen, C.W., Donchin, E., 1988. Pre- and poststimulus activation of response channels: a psychophysiological analysis. Journal of Experimental Psychology: *Human Perception & Performance, 14*(3)*,* 331–344.

Graybiel, A. M. (2005). The basal ganglia: Learning new tricks and loving it. *Current Opinion in Neurobiology*, *15*(6), 638–644. doi:10.1016/j.conb.2005.10.006

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. doi:10.1016/j.appet.2012.04.006

Hall, G. (2002). Associative structures in Pavlovian and instrumental conditioning. In C. R. Gallistel (Ed.), *Stevens' Handbook of Experimental Psychology* (3rd ed., pp. 1–45). New York: John Wiley & Sons.

Hampshire, A., Chamberlain, S. R., Monti, M. M., Duncan, J., & Owen, A. M. (2010). The role of the right inferior frontal gyrus: inhibition and attentional control. *NeuroImage*, *50*(3), 1313–1319. doi:10.1016/j.neuroimage.2009.12.109

Hanes, D. P., Patterson, W. F., & Schall, J. D. (1998). Role of frontal eye fields in countermanding saccades: visual, movement, and fixation activity. *Journal of Neurophysiology*, *79*(2), 817–834. doi:10.1007/s00221-008-1455-0

Hartmann, L., Sallard, E., & Spierer, L. (2015). Enhancing frontal top-down inhibitory control with Go/NoGo training. *Brain Structure and Function*, 1131–1137. doi:10.1007/s00429-015-1131-7

Henson, R. N., Eckstein, D., Waszak, F., Frings, C., & Horner, A. J. (2014). Stimulus-response bindings in priming. *Trends in Cognitive Sciences*, *18*(7), 376–384. doi:10.1016/j.tics.2014.03.004

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: the role of uncertainty. *Quarterly Journal of Experimental Psychology (2006)*, *61*(11), 1658–1668. doi:10.1080/17470210701643439

Hommel, B. (1998). Event Files: Evidence for automatic integration of stimulus-response episodes. *Visual Cognition*, *5*(1-2), 183–216. doi:10.1080/713756773

Hommel, B. (2004). Event files: feature binding in and across perception and action. *Trends in Cognitive Sciences*, *8*(11), 494–500. doi:10.1016/j.tics.2004.08.007

Horner, A. J., & Henson, R. N. (2011). Stimulus-response bindings code both abstract and specific representations of stimuli: evidence from a classification priming design that reverses multiple levels of response representation. *Memory & Cognition*, *39*(8), 1457–1471. doi:10.3758/s13421-011-0118-8

Horner, A. J., & Henson, R. N. (2009). Bindings between stimuli and multiple response codes dominate long-lag repetition priming in speeded classification tasks. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *35*(3), 757–779. doi:10.1037/a0015262

Houben, K. (2011). Overcoming the urge to splurge: influencing eating behavior by manipulating inhibitory control. *Journal of Behavior Therapy and Experimental Psychiatry*, *42*(3), 384–388. doi:10.1016/j.jbtep.2011.02.008

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 (Abingdon, England)*, *107*(7), 1280–1287. doi: 10.1111/j.1360-0443.2012.03827.x

Houben, K., & Jansen, A. (2011). Training inhibitory control. A recipe for resisting sweet temptations. *Appetite*, *56*(2), 345–349. doi:10.1016/j.appet.2010.12.017

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. doi:10.1016/j.drugalcdep.2010.12.011

Hughes, G., Velmans, M., & De Fockert, J. (2009). Unconscious priming of a no-go response. *Psychophysiology*, *46*(6), 1258–1269. doi:10.1111/j.1469-8986.2009.00873.x

Huster, R. J., Enriquez-Geppert, S., Lavallee, C. F., Falkenstein, M., & Herrmann, C. S. (2013). Electroencephalography of response inhibition tasks: Functional networks and cognitive contributions. *International Journal of Psychophysiology*, *87*(3), 217–233. doi:10.1016/j.ijpsycho.2012.08.001

Jahanshahi, M., Obeso, I., Rothwell, J. C., & Obeso, J. A. (2015). A fronto-striato-subthalamic-pallidal network for goal-directed and habitual inhibition. *Nature Reviews. Neuroscience*, *16*(12), 719–732. doi:10.1038/nrn4038

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. doi:10.1162/jocn.2009.21307

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. doi:10.1523/JNEUROSCI.0902-12.2012

Jasinska, A. J. (2013). Automatic inhibition and habitual control: Alternative views in neuroscience research on response inhibition and inhibitory control. *Frontiers in Behavioral Neuroscience*, *7*, 1–4. doi:10.3389/fnbeh.2013.00025

Jodo, E., & Inoue, K. (1990). Effects of practice on the P300 in a Go/NoGo task. *Electroencephalography and Clinical Neurophysiology*, *76*(3), 249–257. doi: 10.1016/0013-4694(90)90019-G

Jodo, E., & Kayama, Y. (1992). Relation of a negative ERP component to response inhibition in a Go/No-go task. *Electroencephalography and Clinical Neurophysiology*, *82*(6), 477–482. doi: 10.1016/0013-4694(92)90054-L

Jones, A., Di Lemma, L. C. G., 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. doi:10.1016/j.appet.2015.11.013

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. doi: 10.1037/a0030683

Kasten, R., & Navon, D. (2008). Is location cueing inherently superior to color cueing? Not if color is presented early enough. *Acta Psychologica*, *127*(1), 89–102. doi:10.1016/j.actpsy.2007.02.002

Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex, *Annual Review of Neuroscience, 23,* 315–341. doi:10.1146/annurev.neuro.23.1.315

Kato, Y., Endo, H., & Kizuka, T. (2009). Mental fatigue and impaired response processes: Event-related brain potentials in a Go/NoGo task. *International Journal of Psychophysiology*, *72*(2), 204–211. doi:10.1016/j.ijpsycho.2008.12.008

Kiefer, M., Marzinzik, F., Weisbrod, M., Scherg, M., & Spitzer, M. (1998). The time course of brain activations during response inhibition: evidence from event-related potentials in a go/no go task. *Neuroreport*, *9*(4), 765–770.

Kiesel, A., Miller, J., Jolicoeur, P., & Brisson, B. (2008). Measurement of ERP latency differences: A comparison of single-participant and jackknife-based scoring methods. *Psychophysiology*, *45*(2), 250–274. doi:10.1111/j.1469-8986.2007.00618.x

Koch, I., & Allport, A. (2006). Cue-based preparation and stimulus-based priming of tasks in task switching. *Memory & Cognition*, *34*(2), 433–444. doi:10.3758/BF03193420

Kok, A., Ramautar, J. R., De Ruiter, M. B., Band, G. P. H., & Ridderinkhof, K. R. (2004). ERP components associated with successful and unsuccessful stopping in a stop-signal task. *Psychophysiology*, *41*(1), 9–20. doi:10.1046/j.1469-8986.2003.00127.x

Konorski, J. (1967). *Integrative activity of the brain: An interdisciplinary approach*. Chicago: University of Chicago Press.

Kruschke, J. K. (2003). Attention in learning. *Current Directions in Psychological Science*, *12*(5), 171–175. doi:10.1111/1467-8721.01254

Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: A practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, *4*(NOV), 1–12. doi:10.3389/fpsyg.2013.00863

Langford, Z. D., Krebs, R. M., Talsma, D., Woldorff, M., & Boehler, C. N. (2016). Strategic down-regulation of attentional resources as a mechanism of proactive response inhibition. *European Journal of Neuroscience, 44*(4), 2095-2103. doi:10.1111/ejn.13303

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. doi:10.1111/j.1460-9568.2004.03683.x

Lawrence, N. S., O’Sullivan, J., Parslow, D., Javaid, M., Adams, R. C., Chambers, C. D., Kos, K., & Verbruggen, F. (2015a). Training response inhibition to food is associated with weight loss and reduced energy intake. *Appetite*, *95*, 17–28. doi:10.1016/j.appet.2015.06.009

Lawrence, N. S., Verbruggen, F., Morrison, S., Adams, R. C., & Chambers, C. D. (2015b). Stopping to food can reduce intake. Effects of stimulus-specificity and individual differences in dietary restraint. *Appetite*, *85C*, 91–103. doi:10.1016/j.appet.2014.11.006

Lawrence, N. S., Hinton, E. C., Parkinson, J. A., & Lawrence, A. D. (2012). Nucleus accumbens response to food cues predicts subsequent snack consumption in women and increased body mass index in those with reduced self-control. *NeuroImage*, *63*(1), 415–422. doi:10.1016/j.neuroimage.2012.06.070

Le Pelley, M. E., & McLaren, I. P. L. (2003). Learned associability and associative change in human causal learning. *The Quarterly Journal of Experimental Psychology. B, Comparative and Physiological Psychology*, *56*(1), 68–79. doi:10.1080/02724990244000179

Le Pelley, M. E., Suret, M. B., & Beesley, T. (2009). Learned predictiveness effects in humans: a function of learning, performance, or both? *Journal of Experimental Psychology. Animal Behavior Processes*, *35*(3), 312–327. doi:10.1037/a0014315

Le Pelley, M. E., Mitchell, C. J., & Johnson, A. M. (2013). Outcome value influences attentional biases in human associative learning: Dissociable effects of training and instruction. *Journal of Experimental Psychology: Animal Behavior Processes*, *39*(1), 39–55. doi:10.1037/a0031230

Le Pelley, M. E., Pearson, D., Griffiths, O., & Beesley, T. (2014). When goals conflict with values: Counterproductive attentional and oculomotor capture by reward-related stimuli. *Journal of Experimental Psychology: General*, *144*(1), 158–171. doi:10.1037/xge0000037

Le Pelley, M. E., Vadillo, M., & Luque, D. (2013). Learned predictiveness influences rapid attentional capture: evidence from the dot probe task. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *39*(6), 1888–1900. doi:10.1037/a0033700

Lenartowicz, A., & Verbruggen, F. (2011). Inhibition-related activation in the right inferior frontal gyrus in the absence of inhibitory cues. *Journal of Cognitive Neuroscience*, *23*(11), 3388–3399. doi:10.1162/jocn_a_00031

Leocani, L., Cohen, L., & Wassermann, E. (2000). Human corticospinal excitability evaluated with transcranial magnetic stimulation during different reaction time paradigms. *Brain*, *123*, 1161–1173. doi:10.1093/brain/ 123.6.1161

Levy, B. J., & Wagner, A. D. (2011). Cognitive control and right ventrolateral prefrontal cortex: Reflexive reorienting, motor inhibition, and action updating. *Annals of the New York Academy of Sciences*, *1224*(1), 40–62. doi:10.1111/j.1749-6632.2011.05958.x

Liefooghe, B., Degryse, J., & Theeuwes, M. (2016). Automatic effects of No-Go instructions. *Canadian Journal of Experimental Psychology*. doi:10.1037/cep0000080

Lin, Z., & Murray, S. O. (2015). Automaticity of unconscious response inhibition: comment on Chiu and Aron (2014). *Journal of Experimental Psychology. General*, *144*(1), 244–54. doi:10.1037/xge0000042

Livesey, E. J., & McLaren, I. P. L. (2007). Elemental associability changes in human discrimination learning. *Journal of Experimental Psychology: Animal Behavior Processes*, *33*(2), 148–159. doi:10.1037/0097-7403.33.2.148

Logan, G. D. (1981). Attention, automaticity, and the ability to stop a speeded choice response. In J. Long, & A. D. Baddeley (Eds.), *Attention and performance IX*. Hillsdale, NJ: Erlbaum.

Logan, G. (1983). On the ability to inhibit simple thoughts and actions: I. Stop-signal studies of decision and memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *9*(4), 585–606. doi:10.1037/0278-7393.9.4.585

Logan, G. D. (1988). Toward an instance theory of automatization. *Psychological Review*, *95*(4), 492–527. doi:10.1037/0033-295X.95.4 .492

Logan, G. (1990). Repetition Priming and Automaticity: Common underlying mechanisms? *Cognitive Psychology*, *35*, 1–35. doi:10.1016/0010-0285(90)90002-L

Logan, G. D. (1992). Shapes of reaction-time distributions and shapes of learning curves: a test of the instance theory of automaticity. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *18*(5), 883–914. doi:10.1037//0278-7393.18.5.883

Logan, G. D., & Bundesen, C. (2003). Clever homunculus: Is there an endogenous act of control in the explicit task-cuing procedure? *Journal of Experimental Psychology. Human Perception and Performance*, *29*(3), 575–599. doi:10.1037/0096-1523.29.3.575

Logan, G. D., & Etherton, J. L. (1994). What is learned during automatization ? The Role of Attention in Constructing an Instance. *Journal of Experimental Psychology Learning Memory and Cognition*, *20*(5), 1022–1050.

Logan, G. D., & Gordon, R. D. (2001). Executive control of attention in dual-task situations. *Psychological Review, 108*(2), 393-434. doi:10.1037//0033-295X.108.2.393

Logan, G. D., Yamaguchi, M., Schall, J. D., & Palmeri, T. J. (2015). Inhibitory control in mind and brain 2.0: Blocked-input models of saccadic countermanding. *Psychological Review*, *122*(2), 115–147. doi:10.1037/a0038893

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. doi: 10.1037/0096-1523.12.4.549

Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, *91*(3), 295–327. doi:10.1037//0033-295X.91.3.295

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. doi:10.1037/a0035230

Longman, C. S., Lavric, A., & Monsell, S. (2013). More attention to attention? An eye-tracking investigation of selection of perceptual attributes during a task switch. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *39*(4), 1142–1151. doi:10.1037/a0030409

Longman, C. S., Lavric, A., Munteanu, C., & Monsell, S. (2014). Attentional inertia and delayed orienting of spatial attention in task-switching. *Journal of Experimental Psychology. Human Perception and Performance*, *40*(4), 1580–1602. doi:10.1037/a0036552

Longman, C. S., Milton, F., Wills, A. J., & Verbruggen, F. (under review). Category-response associations enhance the early stages of instructed learning and rapidly generalize to novel stimuli.

Luck, S. J. (2005). *An introduction to the event-related potential technique*. Cambridge, MA: MIT Press.

Luck, S. J. (2009). The spatiotemporal dynamics of visual-spatial attention. In F. Aboitiz, & D. Cosmelli (Eds.), *From attention to goal-directed behavior.* Springer-Verlag: Berlin.

Luft, C. D. B., Takase, E., & Bhattacharya, J. (2014). Processing Graded Feedback: Electrophysiological correlated of learning from small and large errors. *Journal of Cognitive Neuroscience*, *26*(5), 1180–1193. doi:10.1162/jocn_a_00543

Mackintosh, N. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review*, *82*(4), 276–298. doi:10.1037/h0076778

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* (Vol. 43, pp. 163–168). San Diego, CA: Academic Press.

MacLeod, C. (1991). Half a Century of Research on the Stroop Effect: An Integrative Review. *Psychological Bulletin*, *109*(2), 163–203. doi:10.1037/0033-2909.109.2.163

Maizey (2016). (Unpublished doctoral dissertation). Cardiff University, UK.

Manuel, A. L., Grivel, J., Bernasconi, F., Murray, M. M., & Spierer, L. (2010). Brain Dynamics Underlying Training-Induced Improvement in Suppressing Inappropriate Action. *Journal of Neuroscience*, *30*(41), 13670–13678. doi:10.1523/JNEUROSCI.2064-10.2010

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. doi:10.1016/j.cortex.2012.12.009

Mars, R. B., Piekema, C., Coles, M. G. H., Hulstijn, W., & Toni, I. (2007). On the programming and reprogramming of actions. *Cerebral Cortex*, *17*(12), 2972–2979. doi:10.1093/cercor/bhm022

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. doi:10.1037/a0026718

McAndrew, A., Yeates, F., Verbruggen, F., & McLaren, I. P. L. (2013). Modeling a reaction time variant of the Perruchet effect in humans. In M. Knauff, M. Pauen, N. Sebanz, & I. Wachsmuth (Eds.), Proceedings of the 35th Annual Conference of the Cognitive Science Society (pp. 3014– 3019). Austin, TX: Cognitive Science Society.

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. doi:10.1016/j.nlm.2013.09.014

McLaren, I. P. L., Wills, A. J., & Graham, S. (2010). Attention and perceptual learning. In C. Mitchell & M. Le Pelley (Eds.), *Attention and associative learning: From brain to behaviour* (pp. 131–158). Oxford: Oxford: Oxford University Press.

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*. Academic Press.

McLaren, I. P. L., & Verbruggen, F. (2015). Associations, inhibition, and action. In R. A. Murphy & R. C. Honey (Eds.), *The Wiley Handbook on the Cognitive Neuroscience of Learning*. Chichester, West Sussex: John Wiley & Sons.

Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 2. Accounts of psychological refractory-period phenomena. *Psychological Review*, *104*(4), 749–791. doi:10.1037/0033-295X.104.4.749

Miller, E. K., & Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annual Review of Neuroscience*, *24*, 167–202. doi:10.1146/annurev.neuro.24.1.167

Miller, J., Patterson, T., & Ulrich, R. (1998). Jackknife-based method for measuring LRP onset latency differences. *Psychophysiology*, *35*(1), 99–115. doi:10.1111/1469-8986.3510099

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. doi:10.1017/S0140525X09000855

Mitchell, C. J., & Le Pelley, M. E. (2010). *Attention and associative learning: From brain to behaviour.* Oxford: Oxford University Press.

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. doi:10.1006/cogp.1999.0734

Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, *7*(3), 134–140. doi:10.1016/S1364-6613(03)00028-7

Monsell, S., & Driver, J. (2000). Banishing the control homunculus. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and performance XVIII* (pp. 3–32). Cambridge, MA: 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. doi:10.1037/0096-1523.32.3.493

Moreau, D., Kirk, I. J., & Waldie, K. E. (2016). Seven pervasive statistical flaws in cognitive training interventions. *Frontiers in Human Neuroscience*, *10*(153), 1–17. doi:10.3389/fnhum.2016.00153

Morey, R. D., Rouder, J. N., & Jamil, T. (2015). BayesFactor: Computation of Bayes factors for common designs (Version 0.9.11-1). Retrieved from https://cran.r-project.org/web/packages/BayesFactor/

Morgan, K. K., Luu, P., & Tucker, D. M. (2016). Changes in P3b latency and amplitude reflect expertise acquisition in a football visuomotor learning task. *PLoS ONE*, *11*(4), e0154021. doi:10.1371/journal.pone.0154021

Mostofsky, S. H., & Simmonds, D. J. (2008). Response inhibition and response selection: Two sides of the same coin. *Journal of Cognitive Neuroscience*, *20*(5), 751–761. doi:10.1162/jocn.2008.20500

Muller, H. J., Reimann, B., & Krummenacher, J. (2003). Visual search for singleton feature targets across dimensions: Stimulus- and expectancy-driven effects in dimensional weighting. *Journal of Experimental Psychology. Human Perception and Performance*, *29*(5), 1021–1035. doi:10.1037/0096-1523.29.5.1021

Näätänen, R., & Picton, T. (1987). The N1 wave of the human electric and magnetic response to sound: a review and an analysis of the component structure. *Psychophysiology*, *24*(4), 375–425. doi:10.1111/j.1469-8986.1987.tb00311.x

Nambu, A., Tokuno, H., & Takada, M. (2002). Functional significance of the cortico-subthalamo-pallidal “hyperdirect” pathway. *Neuroscience Research*, *43*(2), 111–117. doi:10.1016/S0168-0102(02)00027-5

Neill, W., & Valdes, L. (1992). Persistence of negative priming: Steady state or decay? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *18*(3), 565–576. doi: 10.1037/0278-7393.18.3.565

Neill, W., Valdes, L., Terry, K., & Gorfein, D. (1992). Persistence of negative priming: II. Evidence for episodic trace retrieval. *Journal of Experimental Psychology Learning Memory and Cognition*, *18*(5), 993–1000. doi:10.1037/0278-7393.18.5.993

Newell, B. R., & Shanks, D. R. (2014). Unconscious influences on decision making: A critical review. *The Behavioral and Brain Sciences*, *37*(1), 1–19. doi:10.1017/S0140525X12003214

Nieuwenhuis, S., Aston-Jones, G., & Cohen, J. D. (2005). Decision making, the P3, and the locus coeruleus-norepinephrine system. *Psychological Bulletin*, *131*(4), 510–532. doi: 10.1037/0033-2909.131.4.510

Nieuwenhuis, S., Yeung, N., & Cohen, J. D. (2004). Stimulus modality, perceptual overlap, and the go/no-go N2. *Psychophysiology*, *41*(1), 157–160. doi:10.1046/j.1469-8986.2003.00128.x

Nieuwenhuis, S., Yeung, N., van den Wildenberg, W., & Ridderinkhof, K. R. (2003). Electrophysiological correlates of anterior cingulate function in a go/no-go task: Effects of response conflict and trial type frequency. *Cognitive, Affective & Behavioral Neuroscience*, *3*(1), 17–26. doi:10.3758/CABN.3.1.17

Nigg, J. (2001). Is ADHD a disinhibitory disorder? *Psychological Bulletin*, *127*(5), 571–598. doi: 10.1037/0033-2909.127.5.571

Noël, X., Brevers, D., & Bechara, A. (2013). A neurocognitive approach to understanding the neurobiology of addiction. *Current Opinion in Neurobiology*, *23*, 1–7. doi:10.1016/j.conb.2013.01.018

Noël, 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. doi:10.1016/j.jbtep.2016.01.003

Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. *Advances in Research: Vol. IV. Consciousness and Self-Regulation*, *99*(99), 1–18. doi: 10.1007/978-1-4757-0629-1

O’Shea, J., & Walsh, V. (2007). Transcranial magnetic stimulation. *Current Biology*, *17*(6), R196–9. doi:S0960-9822(07)00868-8 [pii]\r10.1016/j.cub.2007.01.030

Oldenburg, J. F. E., Roger, C., Assecondi, S., Verbruggen, F., & Fias, W. (2012). Repetition priming in the stop signal task: The electrophysiology of sequential effects of stopping. *Neuropsychologia*, *50*(12), 2860–2868. doi:10.1016/j.neuropsychologia.2012.08.014

Olivers, C. N. L., & Eimer, M. (2011). On the difference between working memory and attentional set. *Neuropsychologia*, *49*(6), 1553–1558. doi:10.1016/j.neuropsychologia.2010.11.033

Papies, E. K., & Aarts, H. (2016). Automatic self-regulation: From habit to goal pursuit. In Handbook of self regulation: Research, theory, and applications (3rd ed). New York, NY: Guilford Press.

Passingham, R. E., & Toni, I. (2001). Contrasting the dorsal and ventral visual systems: Guidance of movement versus decision making. *NeuroImage*, *14*, S125–S131. doi:10.1006/nimg.2001.0836

Pearce, J., & Hall, G. (1980). A model for Pavlovian learning: variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*, *87*(6), 532–552. doi: 10.1037/0033-295X.87.6.532

Pearce, J., & Mackintosh, N. (2010). Two theories of attention: a review and a possible integration. In C. Mitchell & M. E. Le Pelley (Eds.), *Attention and Associative Learning: From Brain to Behaviour* (pp. 11–39). Oxford: Oxford University Press.

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. doi:10.1037/xan0000060

Perruchet, P., Cleeremans, A., & Destrebecqz, A. (2006). Dissociating the effects of automatic activation and explicit expectancy on reaction times in a simple associative learning task. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *32*(5), 955–965. doi: 10.1037/0278-7393.32.5.955

Pfefferbaum, A., Ford, J. M., Weller, B. J., & Kopell, B. S. (1985). ERPs to response production and inhibition. *Electroencephalography and Clinical Neurophysiology*, *60*(5), 423–434. doi: 10.1016/0013-4694(85)91017-X

Pineda, J. A., Westerfield, M., Kronenberg, B. M., & Kubrin, J. (1997). Human and monkey P3-like responses in a mixed modality paradigm: Effects of context and context-dependent noradrenergic influences. *International Journal of Psychophysiology*, *27*, 223–240. doi: 10.1016/S0167-8760(97)00061-5

Polich, J. (2003). Overview of the P3a and P3b. In J. Polich (Ed.), *Detection of change: event-related potential and fMRI findings.* (pp. 83–98). Boston, MA: Kluwer.

Polich, J. (2007). Updating P300: An integrative theory of P3a and P3b. *Clinical Neurophysiology*, *118*(10), 2128–2148. doi:10.1016/j.clinph.2007.04.019

Polich, J., & Kok, A. (1995). Cognitive and biological determinants of P300: an integrative review. *Biological Psychology*, *41*(2), 103–146. doi:10.1016/0301-0511(95)05130-9

Posner, M. A. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. doi:10.1080/00335558008248231

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. In H. Bouma & D. G. Bouwhuis (Eds.), *Attention and performance, vol. 10* (pp. 531–556). Hilldale, N.J: Erlbaum.

Posner, M. I., & Snyder, C. R. (1975). Attention and cognitive control. In R. L. Solso (Ed.), *Information processing and cognition* (pp. 55–85). Hilldale, N.J: Erlbaum.

R Development Core Team. (2014). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from http://www.R-project.org/

Race, E. A., Shanker, S., & Wagner, A. D. (2008). Neural Priming in Human Frontal Cortex : Multiple Forms of Learning Reduce Demands on the Prefrontal Executive System, 1766–1781. doi: 10.1162/jocn.2009.21132

Rae, C. L., Hughes, L. E., Weaver, C., Anderson, M. C., & Rowe, J. B. (2014). Selection and stopping in voluntary action: A meta-analysis and combined fMRI study. *NeuroImage*, *86*, 381–391. doi: 10.1016/j.neuroimage.2013.10.012

Ramautar, J. R., Kok, A., & Ridderinkhof, K. R. (2004). Effects of stop-signal probability in the stop-signal paradigm: The N2/P3 complex further validated. *Brain and Cognition*, *56*(2 SPEC. ISS.), 234–252. doi:10.1016/j.bandc.2004.07.002

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108. doi: 10.1037/0033-295X.85.2.59

Ratcliff, R., & Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time, *111*(2), 333–367. doi:10.1037/0033-295X.111.2.333

Ratcliff, R., Smith, P. L., Brown, S. D., & McKoon, G. (2016). Diffusion decision model: Current issues and history. *Trends in Cognitive Sciences*, *20*(4), 260–281.

Ravden, D., & Polich, J. (1998). Habituation of P300 from visual stimuli. *International Journal of Psychophysiology*, *30*(3), 359–365. doi:10.1016/S0167-8760(98)00039-7

Rehder, B., & Hoffman, A. B. (2005). Eyetracking and selective attention in category learning. *Cognitive Psychology*, *51*(1), 1–41. doi:10.1016/j.cogpsych.2004.11.001

Ridderinkhof, K. R., Ullsperger, M., & Crone, E. A. (2004). The role of the medial frontal cortex in cognitive control. *Science*, *306*(5695), 443–447. doi:10.1126/science.1100301

Rieger, M., & Gauggel, S. (1999). Inhibitory after-effects in the stop signal paradigm. *British Journal of Psychology*, *90*, 509–518.

Rietdijk, W. J. R., Franken, I. H. A., & Thurik, A. R. (2014). Internal consistency of event-related potentials associated with cognitive control: N2/P3 and ERN/Pe. *PLoS ONE*, *9*(7), 3–9. doi: 10.1371/journal.pone.0102672

Rogers, R. D., & Monsell, S. (1995). Costs of a predictable switch between simple cognitive tasks. *Journal of Experimental Psychology: General*, *124*(2), 207–231.

Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, *56*(5), 356–374. doi:10.1016/j.jmp.2012.08.001

Rubia, K., Taylor, E., Smith, A. B., Oksannen, H., Overmeyer, S., & Newman, S. (2001). Neuropsychological analyses of impulsiveness in childhood hyperactivity. *British Journal of Psychiatry*, *179*, 138–143.

Rushworth, M. F. S., Passingham, R. E., & Nobre, A. C. (2002). Components of switching intentional set. *Journal of Cognitive Neuroscience*, *14*(8), 1139–1150. doi:10.1162/089892902760807159

Saenz, M., Buracas, G. T., & Boynton, G. M. (2002). Global effects of feature-based attention in human visual cortex. *Nature Neuroscience*, *5*(7), 631–2. doi:10.1038/nn876

Salinas, E., & Stanford, T. R. (2013). The countermanding task revisited: Fast stimulus detection is a key determinant of psychophysical performance. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *33*(13), 5668–5685. doi:10.1523/JNEUROSCI.3977-12.2013

Schall, J. D. (2001). Neural basis of deciding, choosing and acting. *Nature Reviews Neuroscience*, *2*(1), 33–42. doi:10.1038/35049054

Schapkin, S., Falkenstein, M., Marks, A., & Griefahn, B. (2007). Practice-related effects in a go-nogo task. *Perceptual and Motor Skills*, *105*, 1275–1288. doi:10.2466/PMS.105.4.1275-1288

Schneider, D. W. (2015). Attentional Control of Response Selection in Task Switching. *Journal of Experimental Psychology: Human Perception and Performance*, *41*(5), 1315–1324. doi: 10.1037/xhp0000091

Schneider, D. W., & Logan, G. D. (2005). Modeling task switching without switching tasks: a short-term priming account of explicitly cued performance. *Journal of Experimental Psychology. General*, *134*(3), 343–367. doi:10.1037/0096-3445.134.3.343

Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, *84*(1), 1–66. doi: 10.1037/0033-295X.84.1.1

Shanks, D. R., & John, M. F. S. (1994). Characteristics of dissociable learning systems. *Behavioral and Brain Sciences*, *17*, 367–395. doi:10.1017/S0140525X00035032

Shanks, D. R. (2010). Learning: From association to cognition. *Annual Review of Psychology*, *61*, 273–301. doi:10.1146/annurev.psych.093008.100519

Shiffrin, R., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, *84*(2), 127–190. doi:10.1037/0033-295X.84.2.127

Simmonds, D. J., Pekar, J. J., & Mostofsky, S. H. (2008). Meta-analysis of Go/No-go tasks demonstrating that fMRI activation associated with response inhibition is task-dependent. *Neuropsychologia*, *46*(1), 224–232. doi: 10.1016/j.neuropsychologia.2007.07.015

Simon, J. R., & Rudell, A. P. (1967). Auditory S-R compatibility: The effect of an irrelevant cue on information processing. *Journal of Applied Psychology*, *51*(3), 300–304. doi:10.1037/h0020586

Smith, J. L., Johnstone, S. J., & Barry, R. J. (2008). Movement-related potentials in the Go/NoGo task: The P3 reflects both cognitive and motor inhibition. *Clinical Neurophysiology*, *119*(3), 704–714. doi:10.1016/j.clinph.2007.11.042

Smith, J. L., Smith, E. A., Provost, A. L., & Heathcote, A. (2010). Sequence effects support the conflict theory of N2 and P3 in the Go/NoGo task. *International Journal of Psychophysiology*, *75*(3), 217–226. doi:10.1016/j.ijpsycho.2009.11.002

Smith, J. L., & Douglas, K. M. (2011). On the use of event-related potentials to auditory stimuli in the Go/NoGo task. *Psychiatry Research*, *193*(3), 177–81. doi:10.1016/j.pscychresns.2011.03.002

Smith, P. L., & Ratcliff, R. (2004). Psychology and neurobiology of simple decisions. *Trends in Neurosciences*, *27*(3), 161–168. doi:10.1016/j.tins.2004.01.006

Snyder, K. M., Ashitaka, Y., Shimada, H., Ulrich, J. E., & Logan, G. D. (2014). What skilled typists don’t know about the QWERTY keyboard. *Attention, Perception & Psychophysics*, *76*(1), 162–71. doi:10.3758/s13414-013-0548-4

Song, Y., Ding, Y. L., Fan, S. L., & Chen, L. (2002). An event-related potential study on visual perceptual learning under short-term and long-term training conditions. *Neuroreport*, *13*(16), 2053–2057. doi:10.1097/00001756-200211150-00013

Spierer, L., Chavan, C. F., & Manuel, A. L. (2013). Training-induced behavioral and brain plasticity in inhibitory control. *Frontiers in Human Neuroscience*, *7*, 427. doi:10.3389/fnhum.2013.00427

Stothart, C. R., Simons, D. J., Boot, W. R., & Kramer, A. F. (2014). Is the effect of aerobic exercise on cognition a placebo effect? *PloS One*, *9*(10), e109557. doi:10.1371/journal.pone.0109557

Strayer, D. L., & Kramer, A. F. (1994). Strategies and automaticity: II. Dynamic aspects of strategy adjustment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *20*(2), 342–365. doi:10.1037/0278-7393.20.2.342

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. doi: 10.1037/0096-1523.13.1.89

Suret, M., & McLaren, I. (2005). Elemental representation and associability: An integrated model. In A. J. Wills (Ed.), *New directions in human associative learning* (pp. 155–188). Mahwah, NJ: Lawrence Erlbaum.

Sutherland, N., & Mackintosh, N. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.

Theeuwes, J., & Van der Burg, E. (2011). On the limits of top-down control of visual selection. *Attention, Perception & Psychophysics*, *73*(7), 2092–103. doi:10.3758/s13414-011-0176-9

Tipper, S. P. (1985). The negative priming effect. Inhibitory priming by ignored objects. *The Quarterly Journal of Experimental Psychology, 37*(A), 571-590. doi:10.1080/14640748508400920

Tukey, J. W. (1977). *Exploratory Data Analysis*. Reading, MA: Addison-Wesley.

Tzelgov, J. (1997). Specifying the relations between automaticity and consciousness: A theoretical note. *Consciousness and Cognition, 6*, 441–451.

Ulrich, R., Mattes, S., & Miller, J. (1999). Donders's assumption of pure insertion: an evaluation on the basis of response dynamics. *Acta Psychologica*, *102*(1), 43–76. doi:10.1016/S0001-6918(99)00019-0

Ulrich, R., & Miller, J. (2001). Using the jackknife-based scoring method for measuring LRP onset effects in factorial designs. *Psychophysiology*, *38*(5), 816–827. doi:10.1017/S0048577201000610

Valls-Solé, J., Pascual-Leone, A., Wassermann, E. M., & Hallett, M. (1992). Human motor evoked responses to paired transcranial magnetic stimuli. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *85*(6), 355–364. doi: 10.1016/0168-5597(92)90048-G

Van Boxtel, G. J. M., Van der Molen, M. W., Jennings, J. R., & Brunia, C. H. M. (2001). A psychophysiological analysis of inhibitory motor control in the stop-signal paradigm. *Biological Psychology*, *58*(3), 229–262. doi:10.1016/S0301-0511(01)00117-X

van Gaal, S., Ridderinkhof, K. R., Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). Frontal cortex mediates unconsciously triggered inhibitory control. *Journal of Neuroscience*, *28*(32), 8053–8062. doi:10.1523/JNEUROSCI.1278-08.2008

van Gaal, S. Van, Lamme, V., & Fahrenfort, J. (2011). Dissociable brain mechanisms underlying the conscious and unconscious control of behavior. *Journal of Cognitive*, 1–15. doi:10.1162/jocn.2010.21431

van Gaal, S., Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. a F. (2010). Unconscious activation of the prefrontal no-go network. *The Journal of Neuroscience*, *30*(11), 4143–4150. doi:10.1523/JNEUROSCI.2992-09.2010

van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P. M., & Lamme, V. a F. (2009). Dissociating consciousness from inhibitory control: Evidence for unconsciously triggered response inhibition in the stop-signal task. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(4), 1129–1139. doi:10.1037/a0013551

Van Koningsbruggen, G. M., Veling, H., Stroebe, W., & Aarts, H. (2014). Comparing two psychological interventions in reducing impulsive processes of eating behaviour: Effects on self-selected portion size. *British Journal of Health Psychology*, *19*(4), 767–782. doi:10.1111/bjhp.12075

Veling, H., & Aarts, H. (2009). Putting behavior on hold decreases reward value of need-instrumental objects outside of awareness. *Journal of Experimental Social Psychology*, *45*(4), 1020–1023. doi:10.1016/j.jesp.2009.04.020

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. doi:10.1016/j.brat.2011.08.005

Veling, H., Aarts, H., & Stroebe, W. (2013a). Using stop signals to reduce impulsive choices for palatable unhealthy foods. *British Journal of Health Psychology*, *18*(2), 354–368. doi:10.1111/j.2044-8287.2012.02092.x

Veling, H., Aarts, H., & Stroebe, W. (2013b). Stop signals decrease choices for palatable foods through decreased food evaluation. *Frontiers in Psychology*, *4*, 875. doi:10.3389/fpsyg.2013.00875

Veling, H., Holland, R. W., & van Knippenberg, A. (2008). When approach motivation and behavioral inhibition collide: Behavior regulation through stimulus devaluation. *Journal of Experimental Social Psychology*, *44*(4), 1013–1019. doi:10.1016/j.jesp.2008.03.004

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. doi:10.1016/j.appet.2014.03.014

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. doi:10.1371/journal.pone.0070155

Verbruggen, F., Adams, R., & Chambers, C. D. (2012). Proactive motor control reduces monetary risk taking in gambling. *Psychological Science*, *23*(7), 805–815. doi:10.1177/0956797611434538

Verbruggen, F., Aron, A. R., Stevens, M. A., & Chambers, C. D. (2010). Theta burst stimulation dissociates attention and action updating in human inferior frontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, *107*(31), 13966–13971. doi: 10.1073/pnas.1001957107

Verbruggen, F., Best, M., Bowditch, W., Stevens, T., & McLaren, I. P. L. (2014). The inhibitory control reflex. *Neuropsychologia*, *65*, 263–278. doi: 10.1016/j.neuropsychologia.2014.08.014.

Verbruggen, F., Chambers, C. D., & Logan, G. D. (2013). Fictitious inhibitory differences: How skewness and slowing distort the estimation of stopping latencies. *Psychological Science*, *24*, 352–362. doi:10.1177/0956797612457390

Verbruggen, F., & Logan, G. D. (2008a). 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. doi:10.1037/a0013170

Verbruggen, F., & Logan, G. D. (2008b). Response inhibition in the stop-signal paradigm. *Trends in Cognitive Sciences*, *12*(11), 418–424. doi:10.1016/j.tics.2008.07.005

Verbruggen, F., & Logan, G. D. (2008c). Long-term aftereffects of response inhibition: Memory retrieval, task goals, and cognitive control. *Journal of Experimental Psychology. Human Perception and Performance*, *34*(5), 1229–1235. doi:10.1037/0096-1523.34.5.1229

Verbruggen, F., & Logan, G. D. (2009a). Proactive adjustments of response strategies in the stop-signal paradigm. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(3), 835–854. doi:10.1037/a0012726

Verbruggen, F., & Logan, G. D. (2009b). Models of response inhibition in the stop-signal and stop-change paradigms. *Neuroscience and Biobehavioral Reviews*, *33*(5), 647–661. doi:10.1016/j.neubiorev.2008.08.014

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. doi:10.1037/a0016645.

Verbruggen, F., & Logan, G. D. (2015). Evidence for capacity sharing when stopping. *Cognition*, *142*, 81–95. doi:10.1016/j.cognition.2015.05.014

Verbruggen, F., Logan, G. D., Liefooghe, B., & Vandierendonck, A. (2008). Short-term aftereffects of response inhibition: repetition priming or between-trial control adjustments? *Journal of Experimental Psychology. Human Perception and Performance*, *34*(2), 413–426. doi: 10.1037/0096-1523.34.2.413

Verbruggen, F., McAndrew, A., Weidemann, G., Stevens, T., & McLaren, I. P. L. (2016). Limits of executive control: Sequential effects in predictable environments. *Psychological Science*, *27*(5), 748–757. doi:10.1177/0956797616631990

Verbruggen, F., McLaren, I. P. L., & Chambers, C. D. (2014). Banishing the control homunculi in studies of action control and behaviour change. *Perspectives on Psychological Science.*, *9*(5), 497–524. doi:10.1177/1745691614526414

Verbruggen, F., Stevens, T., & Chambers, C. (2014). Proactive and reactive stopping when distracted : An attentional account. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(4), 1295–1300. doi:10.1037/ a003654.

Verleger, R., Jaskowski, P., & Wascher, E. (2005). Evidence for an integrative role of P3b in linking reaction to perception. *Journal of Psychophysiology*, *19*(3), 165-181. doi: 10.1027/0269-8803.19.3.165

Verleger, R., Schroll, H., & Hamker, F. H. (2013). The unstable bridge from stimulus processing to correct responding in Parkinson's disease. *Neuropsychologia*, *51*(13), 2512–2525. doi:10.1016/ j.neuropsychologia.2013.09.017

Vidal, F., Burle, B., Grapperon, J., & Hasbroucq, T. (2011). An ERP study of cognitive architecture and the insertion of mental processes: Donders revisited. *Psychophysiology*, *48*(9), 1242–1251. doi:10.1111/j.1469-8986.2011.01186.x

Waldron, E. M., & Ashby, F. G. (2001). The effects of concurrent task interference on category learning: evidence for multiple category learning systems. *Psychonomic Bulletin & Review*, *8*(1), 168–176. doi:10.3758/BF03196154

Wallis, J. D., Dias, R., Robbins, T. W., & Roberts, A. C. (2001). Dissociable contributions of the orbitofrontal and lateral prefrontal cortex of the marmoset to performance on a detour reaching task. *European Journal of Neuroscience*, *13*(9), 1797–1808. doi:10.1046/ j.0953-816x.2001.01546.x

Waszak, F., Hommel, B., & Allport, A. (2004). Semantic generalization of stimulus-task bindings. *Psychonomic Bulletin & Review*, *11*(6), 1027–1033. doi:10.3758/BF03196732

Waszak, F., Hommel, B., & Allport, A. (2005). Interaction of task readiness and automatic retrieval in task switching: Negative priming and competitor priming. *Memory & Cognition*, *33*(4), 595–610. doi:10.3758/BF03195327

Waszak, F., Hommel, B., & Allport, A. (2003). Task-switching and long-term priming: Role of episodic stimulus–task bindings in task-shift costs. *Cognitive Psychology*, *46*(4), 361–413. doi:10.1016/S0010-0285(02)00520-0

Waszak, F., Pfister, R., & Kiesel, A. (2013). Top-down versus bottom-up: When instructions overcome automatic retrieval. *Psychological Research*, *77*(5), 611–617. doi:10.1007/s00426-012-0459-3

Webb, T. L., & Sheeran, P. (2006). Does changing behavioral intentions engender behavior change? A meta-analysis of the experimental evidence. *Psychological Bulletin*, *132*(2), 249–268. doi:10.1037/0033-2909.132.2.249

Wendt, M., & Kiesel, A. (2008). The impact of stimulus-specific practice and task instructions on response congruency effects between tasks. *Psychological Research*, *72*(4), 425–432. doi:10.1007/s00426-007-0117-3

Wenke, D., De Houwer, J., De Winne, J., & Liefooghe, B. (2014). Learning through instructions vs. learning through practice: flanker congruency effects from instructed and applied S-R mappings. *Psychological Research*, *79*(6), 899–912. doi:10.1007/s00426-014-0621-1

Wessel, J. R., Danielmeier, C., Morton, J. B., & Ullsperger, M. (2012). Surprise and Error: Common Neuronal Architecture for the Processing of Errors and Novelty. *Journal of Neuroscience*, *32*(22), 7528–7537. doi:10.1523/JNEUROSCI.6352-11.2012

Wiecki, T. V., & Frank, M. J. (2013). A computational model of inhibitory control in frontal cortex and basal ganglia. *Psychol Rev*, *120*(2), 329–355. doi:10.1037/a0031542

Woodman, G. F. (2013). Viewing the dynamics and control of visual attention through the lens of electrophysiology. *Vision Research*, *80*, 7–18. doi:10.1016/j.visres.2013.01.003

Woodman, G. F. (2010). A brief introduction to the use of event-related potentials in studies of perception and attention. *Attention, Perception & Psychophysics*, *72*(8), 2031–2046. doi:10.3758/APP.72.8.2031

Woodman, G. F., & Luck, S. J. (2003). Serial deployment of attention during visual search. *Journal of Experimental Psychology. Human Perception and Performance*, *29*(1), 121–138. doi:10.1167/1.3.103

Woolard, A. A., Kose, S., Woodward, N. D., Verbruggen, F., Logan, G. D., & Heckers, S. (2010). Intact Associative Learning in Patients with Schizophrenia: Evidence from a Go/NoGo Paradigm. *Schizophrenia Research*, *122*(1-3), 131–135. doi:10.1016/j.schres.2010.02.1057

Wylie, G., & Allport, A. (2000). Task switching and the measurement of “switch costs”. *Psychological Research*, *63*, 212–233. doi:10.1007/s004269900003

Yang, J., & Li, P. (2012). Brain networks of explicit and implicit learning. *PLoS ONE*, *7*(8), e42993. doi:10.1371/journal.pone.0042993

Yeung, N., Botvinick, M. M., & Cohen, J. D. (2004). The neural basis of error detection: conflict monitoring and the error-related negativity. *Psychological Review*, *111*(4), 931–959. doi: 10.1037/0033-295X.111.4.939

Zandbelt, B. B., Bloemendaal, M., Neggers, S. F. W., Kahn, R. S., & Vink, M. (2012). Expectations and violations: Delineating the neural network of proactive inhibitory control. *Human Brain Mapping*, *34*(9), 2015–2024. doi:10.1002/hbm.22047

Zandbelt, B. B., & Vink, M. (2010). On the role of the striatum in response inhibition. *PLoS ONE*, *5*(11), e13848. doi:10.1371/journal.pone.0013848