Investigating the Temporal Dynamics of Covert Visual Spatial Attention: Exploitative and Explorative Attentional Mechanisms

Salvatore Simone Russo

School of Psychology

The University of Adelaide

A thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy

February 2020

Table of Contents

Abstract

This thesis explores how trial and error learning affects attentional processes. Previous research has shown that selective attention tends to be biased towards cues that accurately and consistently predict future events. Paying attention to predictive cues is adaptive because an animal can use the information that these cues convey to predict future events and change their behaviour accordingly. However, some research has shown that selective attention can be biased towards non-predictive, or uncertain, stimuli. Paying attention to non-predictive cues could also be adaptive because it could help establish the true nature of these currently uncertain stimuli. Although selective attention can be driven by both high and low levels of uncertainty, the factors that determine which driving force prevails are not fully understood. This thesis investigates whether time is one such factor. The experiments presented here involved training participants on a categorisation task where some stimuli were predictive (P) of the categorisation response while others were non-predictive (NP). These stimuli were then used as uninformative spatial cues to a target stimulus in a dot probe task. The time course of attention to the cues was investigated by manipulating the stimulus onset asynchrony (SOA) between the cues and the target. Behavioural and electrophysiological (EEG) data were collected. It was hypothesised that P cues would be preferentially processed early in a trial. However, after these cues were processed, we predicted that they would be inhibited, and that this inhibition would bias attention towards the currently NP cue.

Experiments 1-3 (Chapter 2) explored the dot probe paradigm by using different stimuli and different dot probe tasks. Using two SOAs (250 and 1200 ms) and an intermixed dot probe and categorisation task, the reaction time (RT) and N2 posterior contralateral (N2pc) results showed that targets that appeared over P cues after short SOAs were easier to process compared to targets that appeared over NP cues. Therefore, P cues were preferentially processed early in a trial. However, no evidence of inhibition of the P cue was

found at the longer SOA. Experiments 4-6 (Chapter 3) tested a wider and earlier range of SOAs and included additional behavioural measures (e.g., dot probe errors and premature responding). In these experiments, RTs were faster to targets that appeared over P cues compared to NP cues, and this advantage increased proportionally with SOA. This novel RT interaction suggested that the P cues were being strategically processed. The N2pc results also showed an interaction between predictiveness and SOA, but one that suggested a shift in attention from P to NP cues. The error data indicated that RTs in the localisation version of the dot probe task could be contaminated by a non-attentional response bias towards selecting a response that was congruent with the location of the P cues. Experiments 7-9 (Chapter 4) tested whether participants were strategically or automatically attending to the P and NP cues, and also investigated whether the attentional effects observed in the previous chapters could influence subsequent cue-outcome learning. Blocking the dot probe and categorisation task resulted in the loss of the interaction between predictiveness and SOA on RTs. Instead, the blocked design resulted in a small, but consistent, RT advantage towards the P cues across all SOAs. This suggested that the current task relevance of the cues is an important factor that determines whether they are strategically processed. In the final experiment, different target stimuli were used after short and long SOAs to investigate whether the changes in attention measured via the N2pc could impact subsequent cue-target learning. However, no evidence of biased learning was found. Chapter 5 presents a summary of the results, including meta-analyses of the behavioural data. The idea that time may be an important factor which moderates exploitative and explorative behaviour is discussed further. This discussion pivots around a real-time model that was modified to take attention into account. Ideas for future research are also presented.

Declaration

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

I acknowledge the support I have received for my research through the provision of an Australian Government Research Training Program Scholarship.

Salvatore Russo

Date: 21/02/2020

Acknowledgements

Firstly, I would like to thank my supervisors, Irina and Nick, for encouraging me to explore my own ideas (for better or worse). Thank you both for your help and advice. Special thanks to my primary supervisor, Irina, for teaching me how to turn my ideas into real experiments, and giving me a bucket load of skills in the process. You are a superhuman.

My sincere thanks to Carolyn Semmler for letting me use her EEG system. Thank you also to Oren Griffiths and Scott Gwinn for teaching me the frequency tagging technique. I had a lot of fun working with you.

To my offices mates, Brittany and Lauren, thank you for putting up with my constant worrying. Your friendship has made this journey extra special for me.

To my wife, Louise, what can I say? This achievement is yours as much as it is mine (maybe more!). I know it’s your job to put up with me, but thanks for signing up for it. To my son, Sam, thank you for helping me put all of this in perspective. I love you both.

List of Publications

Experiments 4 and 5 (Chapter 3) of this thesis are published as:

Russo, S., Burns, N., & Baetu, I. (2019). Mackintosh, Pearce-Hall and time: An EEG study on Inhibition of return. *Biological Psychology, 146*.

List of Conference Presentations

Russo, S., Baetu, I., & Burns, N. (2017, November). The Interaction Between Attention and Associative Learning: Predictive vs. Non-Predictive Cues. Poster session presented at 7th Annual Meeting of the Australasian Cognitive Neuroscience Society. Adelaide, Australia.

Russo, S., Baetu, I., & Burns, N. (2017, December). The Interaction Between Attention and Associative Learning: Predictive vs. Non-Predictive Cues. Paper presented at the meeting of the Australian Learning Group Christmas Workshop. Sydney, Australia.

Russo, S., Baetu, I., & Burns, N. (2018, April). The speed of target processing predicts attentional facilitation in an intermixed associative learning and dot probe task. Paper presented at the meeting of the Associative Learning Symposium. Newtown, Wales.

Russo, S., Baetu, I., & Burns, N. (2018, September). The Interaction Between Attention and Associative Learning. Poster session presented at the meeting of the 12th Annual Florey Postgraduate Conference. Adelaide, Australia.

Russo, S., Baetu, I., & Burns, N. (2018, November). Reconciling Mackintosh and Pearce-Hall: Evidence from human electrophysiology. Paper presented at the meeting of the 8th Annual Biological Psychiatry Australia Conference. Adelaide, Australia.

Russo, S., Baetu, I., & Burns, N. (2018, November). Reconciling Mackintosh and Pearce-Hall: An EEG study on inhibition of return. Paper presented at the meeting of the 8th Annual Meeting of the Australasian Cognitive Neuroscience Society. Melbourne, Australia.

Russo, S., Baetu, I., & Burns, N. (2018, December). Reconciling Mackintosh and Pearce-Hall: An EEG study on inhibition of return. Paper presented at the meeting of the Australian Learning Group Christmas Workshop. Sydney, Australia

CHAPTER 1: General Introduction

The environment is a complex place, with events occurring sometimes at random but also sometimes in predictable ways (Nasser & Delamater, 2016). As information-processing systems, humans, and other non-human animals, must detect, learn about and respond appropriately to stimuli that reliably signal events which are critical for survival, such as food, sex, or danger (Nasser & McNally, 2013). An animal's ability to detect such contingencies is highly adaptive because it can use this information to predict future events and adjust its behaviour accordingly. However, the environment contains more information than we can process at one time (Awh, Belopolsky, & Theeuwes, 2012; Theeuwes, Olivers, & Belopolsky, 2010). Therefore, stimuli, or cues, that are presented simultaneously will compete with one another for internal representation and further analysis (Desimone & Duncon, 1995; Reynolds, Chelazzi, & Desimone, 1999).

This thesis investigates the relationship between two central phenomena in psychology, attention and associative learning. Visual selective attention refers to a set of mechanisms that allow us to preferentially process a subset of stimuli in the environment, which enhances perception of the attended stimuli at the expense of other stimuli that are ignored (Yeshurun & Carrasco, 1998). Associative learning, on the other hand, is the formation of associations among stimuli, actions, and outcomes (Niv et al., 2015). Intuitively, it seems obvious that for learning to occur attention must be paid to what is being learned about. However, the complex nature between attention and learning is not fully understood and the nature of the relationship between the two is an important topic facing associative learning researchers.

1.1 Attention

In his famous chapter on attention, William James wrote:

> Everyone knows what attention is. It is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneous possible objects or trains of thought. Focalization, concentration, of consciousness are of its essence. It implies withdrawal from some things in order to deal effectively with others. (1890/1950, p.403-404)

This quote alludes to the idea that the brain's ability to process information is limited. A consequence of this limited capacity is that increased processing of one stimulus will necessarily be accompanied by a corresponding decrease in processing of other stimuli (Le Pelley, 2010). This idea has also been referred to as the *inverse hypothesis* (Thomas, 1970). Lennie (2003) provided an interesting physiological perspective on this view of attention. After estimating the energy cost of individual spikes (i.e., action potentials) in the human cortex, he showed that there was a limit to the number of neurons that could be substantially active concurrently. He proposed that this limitation necessitates flexible energy management. In other words, information processing costs energy, which is a limited resource. Therefore, the brain must prioritise what information deserves to be processed and allocate its resources accordingly. Selective attention refers to the set of mechanisms by which the brain differentially allocates information processing resources (Klein & Lawrence, 2012).

In the visual modality, selective attention can be divided into three pairs of contrasting types of attention (Moore & Zirnsak, 2017). The first dichotomy is between space-based and object-based attention. Object-based attention involves the observer selecting objects that are behaviourally relevant, whereas spatial attention involves the observer attending to a region of space, which results in facilitated processing of stimuli presented in that location. In addition, visual attention can be deployed either overtly or covertly. Overt attention occurs in the presence of motor commands (e.g., eye movements), whereas covert attention involves the selection of stimuli in the absence of orienting movements (Moore & Zirnsak, 2017; Carrasco, 2011). Therefore, overt attention can be seen by an observer, whereas covert attention must be inferred from changes in performance (Klein & Lawrence, 2012).

The final important distinction is between exogenous and endogenous attention. Exogenous ('bottom-up') factors are external to the observer, based solely on the physical salience of the stimulus (e.g., brightness or colour), and control attention automatically. For example, stimuli that contain features that differ strongly to surrounding items (e.g., a red item surrounded by green ones) are said to 'pop-out' and capture attention despite the intentions of the observer (Awh et al., 2012). Endogenous ('top-down') factors are based on the current internal behavioural goals of the observer (e.g., motivational state or strategy), and direct attention in a controlled manner (Awh et al., 2012; Moore & Zirnsak, 2017). For example, when looking for a friend in a crowd we can shift our attention at will, from person to person, until what we see matches our expectations about the physical appearance of our friend (Macaluso & Doricchi, 2013).

It is important to note, however, that the term "top-down" is not equivalent to the term "voluntary" because some forms of top-down influence occur outside of the participant's control. For example, the phonemic restoration effect occurs when listeners hear a sentence as intact even though some words are masked by an extraneous sound (e.g., a cough). Clearly this effect depends on top-down influences (e.g., the listener's expectation) and occurs regardless of the listeners intentions (Gaspelin & Luck, 2018).

It has been suggested that endogenous and exogenous mechanisms represent two attention systems that can influence information processing in distinct ways (Carrasco, 2011). However, attentional selection likely depends on the dynamic interplay between top-down and bottom-up mechanisms (Theeuwes, 2010). For example, Theeuwes argued that initially visual selection is completely bottom-up, and that only later, through recurrent feedback loops, do top-down processes bias selection.

1.2 Attention and Learning

The exogenous-endogenous dichotomy has proven to be a useful heuristic. However, in recent years this dichotomy has been challenged (Awh et al., 2012) by findings which show that prior experience can result in selection of stimuli that are neither task relevant nor intrinsically salient (Failing & Theeuwes, 2018). In other words, previous attentional deployments can cause persistent selection biases that are unrelated to top-down and bottom-up factors (Theeuwes, 2019). Consequently, it has been proposed that the lingering biases of selection history should constitute a third factor that drive changes in attention (Awh et al., 2012).

In the field of associative learning, a distinction is often made between instrumental and classical conditioning. Instrumental conditioning involves animals learning associations between behavioural responses and their consequences. In his Law of Effect, Thorndike (1911) proposed that behaviours that are followed by reward are reinforced, which increases the likelihood that they will be repeated in similar contexts. Conversely, behaviours that are followed by negative consequences are less likely to be repeated in similar situations. Pavlovian learning (i.e., classical conditioning), on the other hand, occurs when an initially neutral conditioned stimulus (CS), such as a bell, is repeatedly paired with a biologically relevant unconditioned stimulus (US), such as food. After repeated CS-US pairings, the response normally evoked by the US (e.g., salivation) will be evoked by the CS (i.e., the conditioned response; CR). Recent studies on attention have suggested that attentional orienting behaviour might be governed by similar principles (Failing & Theeuwes, 2018). In other words, attention can be conditioned just as other behaviour can be conditioned (Le Pelley et al., 2016; Luque et al., 2017).

Anderson, Laurent, and Yantis (2011a) developed a procedure that demonstrates the effect of prior learning on attention. Participants first completed a training phase in which they learned associations between colour targets and reward. For example, some participants learned that red was associated with high reward (10¢) while green was associated with low reward (2¢). These same colours then acted as task-irrelevant distractors during a subsequent test phase in which no rewards were given and participants had to search for a target defined by shape. On half the trials, one of the distractors was the same colour as a formerly rewarded stimulus (i.e., red or green), whereas on the remaining trials the distractors were rendered in colours not previously associated with reward. Participants were slower to report the orientation of a line within the shape singleton when a reward associated colour distractor was present. This result was striking because these colours were non-salient, task-irrelevant, and shared no features with the shape-singleton target. A control experiment showed that when reward feedback is removed in the training phase the attentional capture by former targets is abolished. This suggested that the slowing of reaction times (RTs) to targets by previously rewarded distractors was not simply because the distractors were previous targets (Anderson, 2015). In a follow up study, Anderson, Laurent & Yantis (2011b) showed that high-value distractors capture attention more so than low value-distractors, confirming that attentional capture can be driven by prior learning about the reward outcome associated with a stimulus. This effect has been termed value-modulated attentional capture (VMAC) (Le Pelley et al., 2016).

Following the findings of Anderson et al. (2011a, 2011b) a number of studies have been reported that both replicate and extend the VMAC effect (Anderson, 2015). Using a similar design, Theeuwes & Belopolsky (2012) had participants make saccades to targets (a vertical or horizontal bar among distractors) and manipulated the reward that participant's received when they made correct responses. For example, some participants received a high reward after making a saccade to a red horizontal bar and a low reward after making a saccade to a red vertical bar. In a subsequent test phase participants had to make saccades to a grey circle among red circle distractors. However, on two-thirds of trials a red horizontal or red vertical bar was presented as a distractor. They found that participants made more erroneous saccades to the high value distractor compared to the low value distractor, demonstrating that stimuli previously associated with high reward can also capture the eyes (i.e., oculomotor capture). In addition, Wang, Yu, & Zhou (2013) showed that attentional capture can occur for loss-associated or pain-associated distractors in visual search, suggesting that motivational significance, rather than valence, of the predicted outcome is the crucial factor for eliciting the VMAC effect (Le Pelley et al., 2016).

Le Pelley, Pearson, Griffiths, & Beesley (2015) extended this research by showing that a stimulus paired with reward can capture attention even when that stimulus has never been task-relevant. Participants had to stare at a fixation cross that disappeared just before the onset of a visual search array that contained a grey shape singleton target and circle distractors. One of the circle distractors was coloured (e.g., red or blue) and the remaining circles were grey. Participants were asked to make saccades to the shape singleton as fast as possible. Importantly, the colour distractor presented on two thirds of trials predicted that a correct response would result in high or low reward (e.g., red and blue predicted high and low reward, respectively). If participants made saccades to the reward-signalling colour distractor the reward for that trial was omitted. Even though looking at the distractor was counterproductive, participants made more erroneous saccades to the high value distractor compared to the low value distractor. Therefore, this study demonstrated that Pavlovian learning (i.e., classical conditioning) of the distractor-outcome relationship was the crucial driving force of attention because orienting responses towards the distractor were never rewarded. In fact, orienting towards the distractor was punished by omission of reward. In addition, the stimuli were never task relevant because there was no separate training session in which they learnt distractor-outcome contingencies. Therefore, attentional capture was dependent on the signal-value of the stimulus, rather than its response-value (Le Pelley et al., 2015).

As the studies reviewed above demonstrate, the idea that learning can drive changes in attention has become increasingly popular. However, this idea is not new. William James wrote about the idea of *derived attention* as attention to a stimulus that "owes its interest to association with some other immediately interesting thing" (1890/1950, p 416). The idea that attention to stimuli can change as a function of reinforcement history is at the heart of attentional learning theories, which try and predict how selective attention and associative learning interact.

1.3 Attentional Learning Theories

Over the last 40 years, formal models of associative learning have attempted to account for how attention and learning interact. These attentional learning theories propose that the amount of attention allocated to a stimulus depends on the certainty or uncertainty of the predictions that it makes about future events (Luque et al., 2016). Some stimuli help us make perfect predictions, indicating that a future event definitely will or will not occur, while other stimuli signal uncertain outcomes (e.g., a bell is followed by food 50% of the time; Gottlieb, 2012). In addition, these theories propose a reciprocal relationship between attention and learning, such that learning influences attention and attention influences subsequent learning.

Two attentional learning theories have dominated the associative learning literature. The first was proposed by Mackintosh (1975), who argued that the amount of attention a stimulus receives depends on how reliably it predicts important outcomes. This theory has been said to embody the *predictiveness principle* (Griffiths, Johnson, & Mitchell, 2011; Le Pelley et al., 2016) because cues that most accurately and consistently predict future events will receive the most attentional processing (Luque et al., 2016). When animals interact with their environment they often aim to obtain rewards and avoid punishments. Indeed, both appetitive and aversive stimuli elicit sustained orienting responses, which prolongs perceptual processing and facilitates action selection (Lang & Bradley, 2013). Mackintosh's theory exemplifies an attention exploitation process whereby cues that reliably predict outcomes are given attentional priority in order to exploit the information that they convey (Le Pelley, Beesley, & Griffiths, 2011). When framed in this way, there seems no reason to waste limited cognitive resources processing cues that do not reliably predict outcomes.

The second theory, proposed by Pearce and Hall (1980), takes a different view. According to this theory, stimuli that reliably predict the events that follow them should receive little attention. Instead, stimuli whose consequences are not well predicted should be paid more attention. This idea has been termed the *uncertainty principle* (Griffiths, Johnson, & Mitchell, 2011; Le Pelley et al., 2016). Paying attention to uncertain stimuli may be adaptive because it could allow for the true nature of these cues to be established. The Pearce-Hall model exemplifies an attention exploration process, whereby processing resources are allocated towards cues whose consequences are currently unknown in order to reduce the uncertainty associated with them (Beesley et al., 2015). When framed in this way, there seems no reason to continue processing stimuli whose meanings are already known. Indeed, once uncertainty disappears there is no longer any need for further information processing (Bernstein, 1979). Intuitively, both the Mackintosh (1975) and Pearce-Hall (1980) models appear plausible (Le Pelley, Haselgrove, & Esber, 2012) and evidence can be found in favour of both.

1.3.1 Mackintosh.(1975) Model

Evidence in favour of the idea that more attention is paid to predictive cues than non-predictive cues comes from studies that compare intradimensional (ID) and extradimensional (ED) shifts. For example, Mackintosh & Little (1969) presented pigeons with coloured lines of different orientation (e.g., red and yellow lines that were either horizontally or vertically oriented). In stage one, half of the pigeons had to learn to peck the correct colour (e.g., red) to gain access to a food reward, while the other half had to learn to peck the correct orientation (e.g., vertical). After learning these discriminations all pigeons were then trained on a new discrimination involving different instances of the previously trained dimensions (e.g., blue and green lines that were tilted 45° to the left or right). Half of the pigeons experienced an ID shift such that the dimension that was relevant during phase one was also relevant during phase two. For example, those that learned to peck a colour in phase one had to learn to peck a different colour in phase two. The remaining pigeons experienced an ED shift such that the dimension that was relevant during phase one was irrelevant during phase two. For example, those that learned to peck lines based on orientation in phase one had to learn to peck lines of a certain colour in phase two. The results showed that pigeons who experienced an ID shift made fewer errors in phase two compared to those that experienced an ED shift. The retardation in learning experienced by the ED group can easily be explained by Mackintosh's (1975) model: pigeons in the ID group learned to pay attention to the relevant dimension during phase one, which would have aided their phase two learning. Conversely, pigeons in the ED group learned to pay attention to a dimension in phase one that was irrelevant in phase two, which would have hindered their phase two learning.[1] This effect has also been found in other species, including rats (Shepp & Eimas, 1964; Trobalon, Miguelez, McLaren, & Mackintosh, 2003), monkeys (Baxter & Gaffan, 2007) and humans (Eimas, 1966; Owen, Roberts, Polkey, Sahakian, Robbins, 1991).

---

[1] Note, however, that Mackintosh's (1975) model has attention operating at the level of individual features, such as red, and not stimulus dimensions, such as colour (Le Pelley et al., 2016). Mackintosh attempts to account for this by suggesting that attention can generalise between features based on their similarity, and that features from one dimension are more similar to each other compared to features from different dimensions (Le Pelley et al., 2016; Pearce & Mackintosh, 2010).

Further evidence in favour of the predictiveness principle comes from the learned predictiveness effect, first reported by Le Pellley and McLaren (2003). Similarly to the ID-ED shift effect, experiments using a learned predictiveness design involve an initial learning phase that establishes the predictive status of stimuli. However, in phase two, instead of learning about new cues, participants learn that the same cues used in phase one predict new outcomes. Consequently, participants in learned predictiveness experiments learn about specific features, and not feature dimensions as in an ID-ED shift design. In their design, Le Pelley and McLaren gave participants a causal learning task, an analogue of classical conditioning (Griffiths & Mitchell, 2008), in which they had to examine the allergic effects of different foods on fictitious patients. In phase one, participants were presented with pictures of a meal made up of two food cues and had to predict which of two allergic reaction outcomes would occur if their patient, Mr. X., had eaten those foods. Half of the meals caused one allergic reaction (e.g., itchiness), while the remaining half caused the other (e.g., nausea). Importantly, within each meal one food cue was predictive of the allergic reaction outcome while the other was non-predictive. After enough trials with corrective feedback participants could predict the correct outcome for each meal. In phase two, participants were informed that they were seeing a new patient, Mr. Y, who ate the same foods as Mr. X but had different allergic reactions (e.g., sweating and dizziness). These new outcomes ensured that participants had to learn new associations between foods and allergies in phase two. Crucially, all foods in phase two were perfect predictors of these new outcomes. Despite the perfect correspondence between foods and allergies, a final test revealed that participants rated foods that were previously predictive in phase one as more causal of phase two outcomes compared to foods that were non-predictive.

The learned predictiveness effect has been replicated many times (for a recent review see Le Pelley et al., 2016). However, a problem with these learned predictiveness studies is that learning measures provide only an indirect measure of attention, and the results of these experiments can be explained by other non-attentional processes (Le Pelley, Beesley, and Griffiths, 2016). For example, Mitchell, Griffiths, Seetoo & Lovibond (2012) have suggested that participants may directly infer that cues that were causal in the first phase of learning are likely to be causal in the second phase, such that the resulting bias in subsequent learning is not mediated by attention, but rather by rational inference. Consequently, recent research has started to examine the relationship between learning and attention by using more direct measures of attention borrowed from perceptual and cognitive psychology.

Le Pelley, Beesley & Griffiths (2011) used a similar design to that of Le Pelley and Mclaren (2003) except that participants were presented with pairs of nonsense words that predicted distinct sounds. In phase one, pairings were made up of predictive and non-predictive words, but in phase two all words were perfect predictors. The aim of this study was to test whether any biased learning found in favour of previously predictive cues was accompanied by changes in attention to those cues as measured by eye tracking. The learning results of this experiment were the same as those found by Le Pelley and McLaren. Participants were better at learning the relationships between words and sounds in phase two for words that were good predictors in phase one, compared to words that were poor predictors. The phase one eye tracking data showed that participants spent longer looking at words that were predictive compared to the non-predictive words. Importantly, this bias in eye gaze continued into phase two, even though all cues were equally predictive in this subsequent phase. This result gave direct evidence in support of the idea that predictive cues demand more attention than non-predictive cues and showed that biased learning towards predictive cues coincides with a change in overt attention towards them. The result of greater eye gaze to predictive cues compared to non-predictive cues has also been replicated many times (see Le Pelley et al., 2016).

1.3.2 Pearce-Hall (1980) Model

One example that lends support to the uncertainty principle comes from Kaye and Pearce (1984) who measured the orienting response of rats towards a light stimulus. The orienting response (i.e., rearing in front of the light or touching the light) was used as an index of attention. Half of the rats were put on a continuous reinforcement schedule such that every time the light came on it was followed by a food reward. The remaining rats experienced a partial reinforcement schedule such that a random half of the light presentations were followed by food reward. The results showed that rats that experienced partial reinforcement oriented towards the light more so than rats that experienced continuous reinforcement. Therefore, rats that learned that the light was a poor predictor of the food outcome paid more attention to the light than rats that learned that the light was a good predictor, a result consistent with the Pearce-Hall model.

Hogarth, Dickinson, Austin, Brown, and Duka (2008) showed that an orienting response in humans might also be governed by the uncertainty principle. In this study, participants were presented with stimuli that signalled an aversive noise outcome with different degrees of predictive validity. One stimulus signalled the noise with a probability of 1 (i.e., perfectly predict the onset of noise), and another stimulus signalled the noise with a probability of 0 (i.e., perfectly predict the absence of noise). A third stimulus was paired with the noise on a random 50% of trials (i.e., non-predictive of noise). Each of these stimuli were paired with a common fourth stimulus. Eye fixations were measured between each unique cue and the common cue. The researchers found that the ratio of time spent looking at the unique cues compared to the common cue was greater for the non-predictive cue compared to the cues the perfectly predicted the presence or absence of noise. In other words, participants spent longer looking at the non-predictive cues compared to the predictive cues.

Further evidence in support of the uncertainty principle comes from studies that show release from negative-transfer (Hall & Pearce, 1982). To set up negative transfer, a CS is first learned to predict a moderate outcome (e.g., weak shock). After this learning, the CS is then paired with a stronger outcome (e.g., strong shock). The negative transfer-effect occurs when initial training of the CS with the moderate outcome interferes with the subsequent conditioning of the CS with the stronger outcome. In their design, Hall and Pearce presented rats with a tone CS followed by a mild shock US. Half of these rats then received a few presentations of the tone by itself (Change group), while the remaining rats did not receive tone-alone trials (Negative Transfer group). All rats were then given trials in which the tone was paired with a stronger shock. There was also a third group of rats (Novel group) that did not experience training with the tone before tone/strong-shock conditioning. The results showed that rats in the Negative Transfer group learned the tone/strong-shock relationship slower than rats in the Novel group (i.e., negative transfer occurred). Importantly, the Change group learned the tone/strong-shock relationship faster than the Negative Transfer group (i.e., this group was released from negative transfer). The implication was that the tone-alone trials introduced uncertainty into the tone-shock relationship, which, according to the Pearce-Hall model, would restore attention to the tone, thus resulting in rapid conditioning of the tone with the larger shock.

Using an allergist task, Griffiths, Johnson, & Mitchell (2011) found evidence of release from negative transfer in humans. Participants examined the allergic effects of different foods on a fictitious patient, Mr. X. In phase one, participants in a Negative Transfer and Change group learned that a certain food produced a mild allergic reaction. In phase two, both groups learned that the same food produced a larger, more serious, allergic reaction. However, before phase two learning, the Change group received two unreinforced trials with the food that produced the allergic reaction (i.e., food followed by no reaction). The results showed that the Change group rapidly learned the food/strong-allergy relationship compared to the Negative Transfer group whose learning was slowed. Therefore, the surprising unreinforced trials experienced by the Change group increased the uncertainty of the food-allergy relationship, which increased attention to the food cue and protected the Change group from negative transfer. This result is uniquely predicted by the uncertainty principle embodied in the Pearce-Hall (1980) model. It should be noted, however, that attempts to replicate the negative transfer effect have been mixed (Le Pelley et al., 2016).

Recent studies have also shown that attention and learning in humans can be moderated by uncertainty. For example, in their Experiment 1, Beesley, Nguyen, Pearson, and Le Pelley (2015) participants were presented with pairs of cues, with one predictive and one non-predictive of a subsequent outcome. Importantly, these pairs also differed in their level of overall uncertainty. That is, some pairs, which contained a perfectly predictive cue, were always followed by one outcome, while other pairs, which contained a partially predictive cue, were paired with their outcomes in a probabilistic fashion. The researchers found that participants spent more time looking at cues with uncertain compounds compared to certain compounds. Therefore, cues that belonged to uncertain compounds had higher attentional priority, consistent with the Pearce-Hall (1980) model.

Easdale, Le Pelley & Beesley (2019) had participants learn about cue compounds containing one P and one NP cue. In phase one, participants were trained with either certain or uncertain contingencies such that the P cues predicted the correct outcome 100% or 80% of the time, respectively. In phase two, the P cues from phase one predicted the correct outcome with a probability of 80% for all participants, whereas the NP cues predicted the correct outcome with a probability of 100%. Therefore, participants in the certain group experienced an increase in uncertainty, whereas participants in the uncertain group experienced no change in overall uncertainty. The results showed that participants in the certain group looked longer at the previously NP cues in phase two, and learned phase two associations faster than those in the uncertain group. Therefore, those who experienced a sudden increase in uncertainty showed greater amounts of exploration compared to those who experienced sustained, or expected, uncertainty. Importantly, this increase in exploration also resulted in faster learning of new cue-outcome contingencies.

As the studies summarised above indicate, there is substantial evidence in favour of both the Mackintosh (1975) and Pearce-Hall (1980) models. However, it should be noted that in human research there is more evidence if favour of the predictiveness principle compared to the uncertainty principle (Le Pelley et al., 2016). It should also be noted that the Mackintosh and Pearce-Hall models are not exclusive. That is, both exploitation and exploration are possible. Indeed, it has been suggested that neither the Mackintosh nor Pearce-Hall model alone can account for how prior experience influences cue processing (Le Pelley, 2004). Consequently, hybrid models that try and incorporate both the predictiveness and uncertainty principle have been developed (e.g., Le Pelley, 2004; Pearce & Mackintosh, 2010; Le Pelley, Haselgrove, & Esber, 2012). However, it could be argued that the Mackintosh and Pearce-Hall models are *mutually* exclusive. In other words, one cannot exploit and explore at the same time, and the preference to exploit or explore might change over time. Therefore, a learning model that attempts to reconcile these two mutually exclusive types of attention might also have to account for changes in attention with time.

1.4 Time

In associative learning research time has often been studied in the context of contiguity, whereby the temporal interval between the presentation of the CS and the presentation of the US is varied (for a recent review see Boakes, & Costa, 2014). One typical finding is that as the interval between the CS and the US is lengthened (i.e., delay of reinforcement), more trials are required to elicit a CR and the strength of the CR is reduced

(Baslam, Drew, & Gallistel, 2010). For example, Shneiderman and Gormenzano (1964) used air puffs to condition a nictitating membrane response in albino rabbits. They found that rates of conditioning were inverse to the CS-US interval. Rabbits that experienced a CS-US interval of 250 ms reliably acquired CRs, whereas rabbits that experienced an interval of 4 s showed reduced CRs. However, another important finding in temporal research was that the effect of reinforcement delay (T) on learning depends on the intertrial duration (I). Specifically, if I is fixed then the number of reinforcements needed to acquire a CR will increase as T increases. In other words, if the I/T ratio decreases then more reinforcements will be required to acquire a CR. However, if the I/T ratio is held constant (e.g., by proportionally increasing I as T increases) then the number of reinforcements to acquire a CS response remains constant. Therefore, the delay of reinforcement (T) only affects acquisition if the I/T ratio is altered (Gallistel & Gibbon, 2000), with high I/T ratios (e.g., short reinforcement delays and long intertrial intervals) resulting in faster (or more efficient) learning compared to low I/T ratios.

Temporal research has also been conducted in the context of the behavioural systems framework developed by William Timberlake (Timberlake and Lucas, 1989). In this framework, a motivational *system* (e.g., feeding) is made up of different *subsystems* (e.g., predation or foraging). The subsystems are broken down further into different *modes*, which are differentially expressed depending on the spatio-temporal distance to the goal object (e.g., food). For example, when food is spatially or temporally distant an animal is thought to engage in a *general search* mode (i.e., attention to novelty and search for food-related cues). When cues predictive of immediate food are present the animal enters a *focal search* mode (i.e., more focused action patterns related to the immediate procurement of food). The animal finally enters a *handling/consuming* mode (i.e., action patterns focused on dealing with the food). Once in a mode, the animal engages in a repertoire of pre-organised (i.e., selected over generations and modified by previous experience) perceptual-motor *modules* that make the animal more sensitive to certain stimuli and increase the likelihood of certain motor actions (Timberlake, 2001).

In the behavioural system approach, the idea of modes is particularly related to time because animals are thought to transition between these different behavioural repertoires depending on the spatio-temporal proximity of the outcome (e.g., food). Therefore, the behavioural systems approach predicts that different CS-US intervals will differentially condition different modes (i.e., general search, focal search, handling/consuming) and hence manifest different behavioural repertoires (e.g., scanning, sniffing, pouncing, chewing) and stimulus sensitivities (Timberlake, 2001). For example, when the CS is always followed by the US after a short interval, the animal is conditioned to respond to the CS with behaviour related to focal search, such as nosing the food tray (if the US is food). However, when the CS-US interval is lengthened, animals will engage in responses related to general search, such as locomotion and scanning (Timberlake, 2001). To investigate different responses during a CS-US interval researchers have used interfood clocks in fixed-time schedules. An interfood clock involves a sequence of different stimuli that typically divide the inter food interval into equal segments. For example, Silva and Timberlake (1998) investigated the behaviour of rats during the presentation of four equal light durations, with the termination of the final light coinciding with food delivery (i.e., S1-S2-S3-S4-food). On average, the rats responded with rearing near the feeder during S1, locomotion away from the feeder during S2 and S3, and nosing in the feeder during S3 and S4.

Matthews and Lerer (1987) used an interfood clock to investigate the behaviour of pigeons during an interfood interval. They gave hungry pigeons grain after a 30 second presentation of a light that went from dim (10s) to medium intensity (10s) to very bright (10s). During the first 10 second interval the pigeons showed circling away from the food hopper. During the second 10 second interval they showed hopper directed behaviour. Finally, during the last interval they showed increased pecking of the light. This behaviour is consistent with a general search mode being followed by focal and a handling/consuming mode. Matthews and Lerer also ran another condition in which the change in illumination of the light was reversed (i.e., bright to dim) and the behaviour of the pigeons also reversed. That is, they pecked the bright light during the first 10 seconds, even though this was now temporally further away from the food reward, and circled during the dim light presentation in the last 10 seconds, even though this was now temporally closer to food. Therefore, it seemed the pigeons were conditioned to express different behavioural responses depending on the temporal proximity between the CS and the US.

Researchers have also used unconditioned probe stimuli to investigate the extent to which temporal delay elicits different behaviours. For example, Silva and Timberlake (1997) presented rats with a tone followed by the delivery of food. Different rats received a long (e.g., 16 sec) or short (e.g., 4 sec) tone duration, after which food was immediately delivered. After this conditioning, the rats were exposed to unconditioned presentations of a rolling ball bearing and their baseline level of interacting with the ball was assessed (i.e., touching the ball with nose, mouth or paw). The ball bearing was then used as an unconditioned probe stimulus in a subsequent test phase in which the CS was presented two seconds before the bearing and no food was given. The results showed that contact with the bearing increased above baseline for rats conditioned with a long, but not short, CS-US interval. Following the CS, the rats in the short group increased nosing in the feeder more so than rats in long group. These results were consistent with a general search mode conditioned to the CS under the longer interval, and a focal search mode conditioned under the short interval.

Interestingly, uncertainty in the CS-US relationship can be related to the temporal proximity between the CS and the US. Taking into account both reinforced and unreinforced trials, a particular time point within a trial is, on average, further away from the US during partial reinforcement compared to continual reinforcement (Silva & Timberlake, 2005). For example, when a continuous reinforcement schedule is used there is a 1:1 relationship between CS exposure and reinforcement. Therefore, the expected amount of CS exposure per reinforcement is T, the trial duration. However, if a partial reinforcement schedule of 2:1 is used (i.e., an average of 2 presentations of the CS per reinforcement), then the expected amount of CS exposure per reinforcement is 2T (Gallistel & Gibbon, 2000). In this way, uncertainty (i.e., partial reinforcement) can be related to time. Silva and Timberlake (2005) investigated the effect of partial or continuous reinforcement on the distribution of focal and general search behaviour. During a 48 second trial, rats were presented with four 12 second CSs (i.e., an interfood clock), which were followed by food delivery (i.e., S1-S2-S3-S4-food). Rats that received partial reinforcement (i.e., S4 was followed by food 50% of the time) showed increased general search compared to the continuously reinforced group. In addition, peak locomotion (a behaviour associated with general search) occurred during S3 for the partially reinforced group and during S2 for the continuously reinforced group. The implication of this result was that rats in the partial reinforcement group perceived each clock segment as further away from the US. In other words, peak locomotion occurred closer to S4 for partially reinforced rats because this segment was, on average, closer to the middle of the interfood interval (i.e., S2), which was the time of peak locomotion for the continuously reinforced rats (Silva & Timberlake, 2005).

As the studies reviewed above indicate, a short temporal distance between the CS and the US seems to evoke a mode of behaviour (focal search and handling/consuming) analogous to exploitation, whereas a long CS-US interval appears to evoke a mode (general search) analogous to exploration. It is possible that animals engaged in general search would be more likely to form new associations between stimuli in their environment compared to animals engaged in focal search or handling/consuming. Therefore, time may be an important variable to consider when trying to reconcile the Mackintosh (1975) and Pearce-Hall (1980) theories.

Another way time can be related to exploitation and exploration comes from the phenomenon of inhibition of return (Klein, 2000, Wang & Klein, 2010). Posner and Cohen (1984) first demonstrated inhibition of return using a spatial cueing paradigm. In these experiments participants are presented with a central fixation point, which is flanked by peripheral placeholders to the left and right. One of the two peripheral locations is cued (e.g., by briefly brightening it) and a short time after the cue a target appears in either the cued or un-cued location. The observer has to make a speeded response as soon as they detect the target. The results typically show that when the time between the onset of the cues and the onset of the target (i.e., stimulus onset asynchrony; SOA) is short observers respond faster to the target when it is in the cued location rather than the un-cued location. This facilitation effect is thought to reflect better processing of the cued location due to the observer automatically orienting attention toward the cue. However, at long SOAs observers respond faster to the target when it is presented in the un-cued location. A plausible interpretation of this latter finding is that attention is initially captured by the peripheral cue, but during the long SOA attention reorients back to fixation before the target appears. Attention is then thought to be inhibited from returning to the previously attended cued location and encouraged toward the previously un-attended un-cued location. This interpretation led Posner, Rafal, Choate, and Vaughan (1985) to coin the term ‘inhibition of return’ (IOR). Temporally, IOR usually occurs when the cue-target SOAs are greater than approximately 300 ms (Satel, Wang, Hilchey, & Klein, 2012) and has been shown to occur with SOAs as long as 3000 ms (Samuel & Kat, 2003).

IOR has been dubbed a “foraging facilitator” in visual search (Klein & MacInnes, 1999), and has been invoked as a novelty seeking-mechanism that encourages foraging behaviour (Klein, 2000; Wang & Klein, 2010). The similarities between the foraging account of IOR and the attentional exploration exemplified by the Pearce and Hall (1980) model suggests that time (i.e., SOA) and inhibition could help explain how animals explore their environment in order to reduce uncertainty. For example, it may be that P cues capture attention early so that animals can make an immediate response based on the information that these cues convey. However, when no immediate response is required, these cues may be inhibited so that the animal can scan the environment, giving it an opportunity to reduce the uncertainty associated with other stimuli present.

Using a modified version of the dot probe task, Le Pelley, Vadillo, and Luque (2013) investigated the time course of covert attentional capture by stimuli that had different predictive validity. They first gave participants a pre-training phase in which they were presented a green square and some oblique lines that appeared to the left and right of a central fixation cross. Participants were required to make a categorisation response (e.g., up or down) based on the stimuli that were presented. For half of the participants, the shade of the square was predictive of category membership (e.g., light or dark green indicated that the correct response was up or down, respectively), while the oblique lines were non-predictive. For the remaining participants, the thickness of the oblique lines (e.g., thin or thick) was predictive of the categorisation response and the shade of green was non-predictive. After learning the predictive status of these stimuli with the aid of corrective feedback, the participants then completed a dot probe task in which the same stimuli were used as uninformative spatial cues to a target stimulus (a small white square). The SOA between the onset of the cues and the onset of the target was either 250 or 1000 ms. Participants were faster to respond to the location of the target when it appeared over a P cue after the short SOA, while at the long

SOA no difference in RT was found. The results were taken as evidence that the P cues automatically captured covert attention because the short SOA gave participants little time to consciously process the stimuli. In addition, if participants were strategically shifting their attention towards the P cue then the long SOA, which provided more time for top-down processing, should have produced a larger, or at least equivalent, RT facilitation effect (Le Pelley et al., 2016).

Le Pelley et al (2013) did not observe IOR, but instead found no RT differences to targets that appeared over the P or NP cue at their long (1000 ms) SOA. However, Le Pelley et al (2013) manipulated SOA in order to investigate the automaticity of attention, not to investigate IOR. Therefore, they may have used SOAs that were not optimal to find the effect. Another reason may be due to the stimuli that were used. One could argue that the physical properties, and hence the intrinsic saliences, of the stimuli used in Le Pelley et al.'s dot probe task were very different. For example, perhaps the green square was more inherently salient than the oblique lines. In accordance with stimulus-driven attentional capture, observers will initially select the most physically salient object in the environment (Theeuwes, Olivers & Belopolsky, 2010). If the inherent salience of an object interacts with its derived salience (i.e., the salience acquired due to the object's associative history), the fact that the two stimuli started with very different physical properties may have confounded the RT results. It has also been argued that behavioural measures, such as RT, are inherently ambiguous and may not be the most sensitive measure to observer IOR. Consequently, the use of converging techniques may be needed to investigate the issue

1.5 Electroencephalography

EEG is useful technique to investigate stimulus processing because its high temporal resolution allows for the study of immediate brain responses to stimuli. This makes it suitable for studying changes in attention. A number of studies have attempted to identify electrophysiological markers for facilitation and inhibition, but the results have been mixed. Nevertheless, a recent review of these studies has highlighted the N2-posteriror contralateral component (N2pc) as reliable marker for the mechanisms underlying cueing effects (Martin-Arevalo, Chica, & Lupianez, 2016). The N2pc is a lateralised event related potential (ERP) that is known to reflect the location of covert visual spatial attention. The N2pc is defined as a greater negativity recorded over posterior electrodes that are contralateral to an attended item relative to ipsilateral sites. The amplitude of the N2pc is thought to reflect the amount of attention deployed to a stimulus and its latency to the point in time of attentional deployment (Luck et al., 2006)

Feldmann-Wustefeld, Uengoer, & Shubo (2015) used N2pc and RT measures to investigate changes in attention after learning. During a learning task, participants viewed displays containing eight shapes arranged in a circle around a central fixation point. The learning display contained six grey circles, one colour singleton (e.g., a green or blue circle) and one shape singleton (e.g., a grey pentagon or triangle). Participants were assigned to either a colour or shape predictive group, and learnt that colour or shape, respectively, predicted the correct categorisation response. After learning, participants completed a visual search task. The search display contained a grey shape singleton among grey distractors. On half of the search trials a colour singleton was added as a distractor. Participants had to report the orientation of a line segment within the shape singleton. The results showed that participants in both groups were slower to report the orientation of the line when the colour distractor was present. However, the slowing effect of the colour distractor was greater for those in the colour-predictive group. In addition, participants in the colour predictive group showed an early N2pc effect in the direction of the colour distractor, followed by a N2pc effect in the direction of the shape target. Participants in the shape predictive group showed only a N2pc effect in the direction of the shape target. The authors concluded that participants in the colour, but not the shape, predictive group had their attention captured by the colour distractor before attending to the shape target.

1.6 General Aims and Chapter Summary

Trial-level models, such as those proposed by Mackintosh (1975) and Pearce and Hall (1980), do not include a temporal component to account for changes in attention within a learning trial. Such real-time models exist (e.g., Harris, 2006; Harris & Livesey, 2010; McLaren & Mackintosh, 2000; Kokkola, Mondragon, & Alonso, 2019) and have been shown to produce more complex emergent behaviour than their trial-level counterparts (Kokkola et al., 2019). However, the experimental attention devoted to investigating the temporal dynamics of attention within a learning trial has been lacking. The current thesis aims to fill this gap by investigating whether the temporal dynamics of attention can be used to reconcile exploitative and explorative attentional learning theories.

Overall, this thesis aims to measure learning induced changes in attention using behavioural (e.g., RTs, ratings, errors) and EEG (e.g., N2pc) measures The measurement of within-trial changes in attention might provide novel insights into exploitation and exploration mechanisms that can inform real-time models of associative learning. It is also possible that measuring changes in attention with time could help inform the debate of whether the learned predictiveness effect is generally under top-down or bottom-up control (Le Pelley et al., 2013).

In all experiments presented here, participants first learned to categorise cue compounds, with one cue in each pairing predictive of the categorisations response and the other non-predictive. The same stimuli were then used as uninformative spatial cues in a dot probe task in which participants had to respond to a target stimulus that appeared randomly over one of the stimuli. All experiments employed the dot probe task first used by Le Pelley et al. and EEG (e.g., probe-locked and cue locked N2pc) and behavioural data (e.g., RTs, ratings, dot probe errors, premature responses) were collected in an attempt to understand how attention to the cues changed with time. Table 1.1 summarises the experiments presented in each chapter.

Chapter 2 presents a series of experiments in which the physical attributes of the cues and the dot probe target were manipulated to explore the temporal dynamics of attention. All experiments in this chapter used SOAs between the cues and the dot probe target of 250 and 1200 ms. Experiments 1 and 2 used a dot probe localisation task (i.e., left or right response to target that appeared over the P and NP cues). However, in Experiment 1 the P and NP cues were black and white gratings that needed to be categorised by attending to orientation, whereas in Experiment 2 the cues were coloured gratings that could be discriminated by colour and/or orientation. In Experiment 3, the cues were coloured gratings but the dot probe task was changed from a simple localisation task to a target discrimination task (i.e., report the location of a vertical line within a sideways T target). In addition, Experiment 3 introduced a rating measure that involved asking the participants how often they noticed the target appear over the P and NP cues.

Chapter 3 presents a series of experiments in which additional SOAs were tested to further explore the timing of attentional facilitation and inhibition to the cues. Experiments 4 and 5 introduce error and premature response data. Experiment 6 employs the frequency tagging technique to measure attention to the cues in an attempt to better understand the N2pc measure of attention used in all experiments.

Chapter 4 contains experiments that attempt to address whether the learned predictiveness effect is generally under top down or bottom up control and to test the reciprocal relationship between attention and learning. Experiment 7 used an additional longer SOA to test the time course of strategic processing. Experiment 8 blocked the categorisation and dot probe tasks to test whether P cues that have no task relevance are strategically or automatically processed. Experiment 9 manipulated the appearance of the dot probe target to test whether attention towards the cues could result in biased learning of cue-target associations.

Chapter 5 contains the General Discussion, which provides a summary of the results (including meta-analyses) and a discussion of how cue-outcome learning could be influenced by changes in attention over time. This discussion pivots around a real-time model (a modified version of McClelland & Rumelhart's (1988) auto-associator). Ideas for future experiments are also suggested.

Table 1.1

Summary of Experiments

Chapter	Experiment	P & NP Cues	Cue Discrimination	Categorisation & Dot Probe Task	SOAs (ms)
2	1	Black and white	Feature (orientation)	Intermixed	250 and 1200
	2	Colour	Feature (colour and/or orientation)	Intermixed	250 and 1200
	3	Colour	Feature (colour and/or orientation)	Intermixed	250 and 1200
3	4	Colour	Feature (colour and/or orientation)	Intermixed	0,100,200,300, 400 and 500
	5	Colour	Feature (colour and/or orientation)	Intermixed	0,100,200,300, 400 and 500
	6	Colour	Conjunction (colour and orientation)	Intermixed	0,100,200,300, 400, 500 and 600

Table 1.1 (continued)

Summary of Experiments

Chapter	Experiment	P & NP Cues	Cue Discrimination	Categorisation & Dot Probe Task	SOAs (ms)
4	7	Colour	Feature (colour and/or orientation)	Intermixed	0,100,200,300, 400, 500 and 1000
	8	Colour	Feature (colour and/or orientation)	Blocked	0,100,200,300, 400 and 500
	9	Colour	Feature (colour and/or orientation)	Intermixed	0,100,200,300, 400 and 500

CHAPTER 2: Exploring the Dot Probe Task and the N2pc

The aim of the current experiments was to investigate facilitation and inhibition effects caused by derived attention. Human subjects learned the predictive status of four cues, which were then used in a dot probe task similar to that described by Le Pelley et al. (2013). To investigate the time course of covert visual-spatial attention the SOA between the cues and the dot probe target was manipulated and behavioural (RTs) and ERP (N2pc) measures thought to reflect attention were obtained. With short SOAs a Mackintosh-like facilitation effect was expected. Conversely, a Pearce-Hall-like IOR effect was expected after longer SOAs. This would provide evidence that the two attention mechanisms operate in sequential manner, with predictive cues capturing attention initially, followed by a shift of attention towards non-predictive cues.

2.1. Experiment 1

Experiment 1 was based on the design of Le Pelley et al. (2013; also see Luque et al., 2017 and Cobos, Vadillo, Luque, & Le Pelley, 2018). Participants first completed a categorisation task, the design of which can be seen in Table 2.1. Participants were presented with two equally salient black and white sinusoidal gratings. One of the gratings in the stimulus pair was oriented near-vertical or near-horizontal (hereinafter referred to as vertical and horizontal, respectively), while the other grating was oriented 45° to the left or right of vertical (hereinafter referred to as left and right, respectively). Each stimulus pair belonged to one of two categories and participants learned how to categorise each pair with the aid of corrective feedback. Only two of the gratings were predictive (P) of the correct categorisation response, while the other two were non-predictive (NP). A P grating was always paired with an NP grating.

After learning the predictive nature of these stimuli, participants completed an intermixed categorisation and dot probe task. The same gratings used in the categorisation task were presented to the left and right of a central fixation cross. A small white square target then appeared randomly over the position of either the P or the NP grating, and participants were required to make a speeded response to the location of the target. The SOA was also manipulated in order to investigate the time-course of attention. Although we used the same short SOA duration of 250ms used by Le Pelley et al. (2013), we used a long SOA of 1200ms. This was done on the assumption that the 1000ms SOA used by Le Pelley et al. may not have been long enough to capture IOR in RTs. EEG was used to monitor each participant's brain activity throughout the experiment in order to examine ERPs (N2pc) elicited by the lateralised target when it appeared over the P and NP cues at each SOA.

Table 2.1

Design of the Categorisation Task

	Correct Response	
Stimulus Pair	Horizontal-Vertical Predictive Condition	Left-Right Predictive Condition
Horizontal & Right	R1	R1
Horizontal & Left	R1	R2
Vertical & Right	R2	R1
Vertical & Left	R2	R2

Note. Horizontal and Vertical refer to gratings with orientations of 4° and 94°, respectively. Left and Right refer to gratings with orientations of 45° and 135°, respectively. R1 and R2 refer to the correct categorisation response (the up and down buttons on a button box, counterbalanced). Participants in the "Horizontal-Vertical Predictive" condition learned that gratings with near-vertical or near-horizontal orientation were predictive of the categorisation response, while gratings oriented 45° to the left or right of vertical were non-predictive.

Participants in the “Left-Right Predictive” condition learned that gratings orientated 45° to the left or right of vertical were predictive of the categorisation response, while gratings with near-vertical or near-horizontal orientation were non-predictive.

2.1.1. Method

Participants

Twenty-four volunteers (15 male; mean age 24 years, age range 18 to 29 years) with normal or corrected to normal vision participated in this experiment. Due to age-related effects on reaction time, only individuals aged between 18 and 40 years were eligible to participate. Other eligibility criteria included not suffering from migraines, not suffering from a drug or alcohol dependency (either a current or previous condition), not smoking more than five cigarettes per day, and not using medication that affects neurological function (e.g. antidepressants, sedatives, antipsychotics). All participants gave written informed consent and were told they could withdraw from the experiment at any time.

Apparatus

The experiment was run using E-Prime (version 2.0; Psychology Software Tools, Pittsburgh, PA, USA) to control the presentation of visual stimuli on 23.5-in computer monitor (Eizo FORIS FG2421, 1920 × 1080 pixels, vertical refresh rate = 100 Hz). Participants used the four rectangular buttons on a Cedrus response box (model RB-540) to respond throughout the experiment.

Stimuli

The four cues were sinusoidal luminance gratings (spatial frequency = 0.05, phase = 0) that were spatially modulated in contrast by a Gaussian envelope (standard deviation = 30) and had orientations of 4° (near-horizontal), 45° (tilted to the left), 94° (near-vertical), or 135° (tilted to the right). Near-horizontal and near-vertical gratings were used to avoid preferential selection that has been shown to occur for non-oblique orientations (Meigen, Lagreze, & Bach, 1994). The grating stimuli were presented in white square frames with sides subtending 3.5° of visual angle at a viewing distance of 80 cm. Each square was presented on the horizontal midline of the screen and centred 5.5° to the left and right of a central fixation cross. The target in the dot probe task was a white square with sides subtending 0.8° that appeared in the centre of one of the white square frames, superimposed over one of the grating cues. The screen background was mid-gray (R128, G128, B128).

Procedure

Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n=12) and “Left-Right Predictive” (n=12). Experiment 1 contained two phases. Phase 1 was a categorisation task that established the cue-outcome contingencies. Phase 2 combined the categorisation task with a dot probe task.

Participants received written and verbal instructions for the categorisation task and the dot probe task. For the categorisation task, participants were asked to fixate the centre of the screen throughout the task. They were informed that on each trial they would be shown two images, one to the left and one to the right of a central cross. Participants were also informed that the images would look like lines rotated at different angles and that for each pair presented they were required to make a categorisation response by pressing the up or down buttons on a button box. Participants were asked to learn how to categorise these stimuli through trial and error via corrective feedback that would be provided after every response. Table 2.1 shows the four cue compounds and their corresponding correct responses. On every trial a horizontal or vertical grating was presented alongside a left or right grating. Participants assigned to the “Horizontal-Vertical Predictive” condition learned that the horizontal and vertical gratings were predictive of the categorisation response, while the left and right gratings were non-predictive. Participants assigned to the “Left-Right Predictive” condition learned that left and right gratings were predictive of the categorisation response, while the horizontal and vertical gratings were non-predictive (note that participants were not informed that some stimuli were predictive of the correct response whereas others were non-predictive). Participants were asked to minimise errors and were informed that after an initial learning phase, 100% accuracy was attainable.

The Phase 1 categorisation task is summarised in Figure 2.1A. Each trial began with the presentation of a pre-stimulus screen that consisted of a central cross flanked by two empty placeholder boxes. This pre-stimulus screen was presented for a randomly jittered interval between 900 and 1125ms, after which the two cue stimuli appeared centred in the placeholder boxes and the text “Up or Down?” appeared in the centre of the screen replacing the cross. After participants made a categorisation response using the up or down buttons on the button box, corrective feedback appeared in the centre of screen for 1500ms. After the feedback was provided, the pre-stimulus screen for the next trial was presented. The categorisation task was split into blocks of eight trials. This consisted of each stimulus pair shown in Table 2.1 presented once with the predictive cue appearing on the left and once with the predictive cue appearing on the right, in a random order. A performance criterion was used to determine the end point of Phase 1. Each participant completed a minimum number of 20 blocks (160 trials). If the participant achieved at least 7/8 trials correct in the 20th block, or any subsequent block, they proceeded to Phase 2. If participants did not achieve this learning criterion after 40 blocks, the experiment was terminated.

After completing the categorisation task participants received instructions for the Phase 2 intermixed categorisation and dot probe task. Phase 2 is summarised in Figure 2.1B. Participants were instructed that they would still have to fixate the centre of the screen and that the same images as before would appear to the left and right. On every trial, soon after these images appeared a small white square appeared randomly either over the left or the right image. The participant's task was to respond as fast as they could to the location of the target square by pressing the left or right buttons on the button box. After responding to the target, it disappeared from the screen and participants were required to make a categorisation response based on the grating stimuli that remained on the screen by pressing the up or down buttons on the button box. This requirement for participants to continue categorising the stimuli throughout Phase 2 ensured that any attentional bias produced by the categorisation task would not extinguish over time (Le Pelley et al., 2013). Participants were explicitly told that they should ignore the grating stimuli until after they had responded to the dot probe target because the target would appear randomly over the gratings. Participants were also informed that their categorisation response (i.e., up or down) did not have to be rushed. Feedback regarding their categorisation response was provided after their up or down response.

The dot probe task was split into blocks of 32 trials. Each block consisted of all possible combinations of four stimulus pairs (Table 2.1), two cue locations (predictive or non-predictive cue on the left or right), two SOAs (250ms or 1,200ms), and two target locations (left or right). The 32 trials within a block were presented in random order. Participants completed 10 blocks of the dot probe task (320 trials). There were four within-subjects conditions. These included when the target appeared over the predictive cue at the short SOA (P250) and over the non-predictive cue at the short SOA (NP250), and when the target appeared over the predictive and non-predictive cues at the long SOA (P1200 and NP1200, respectively).

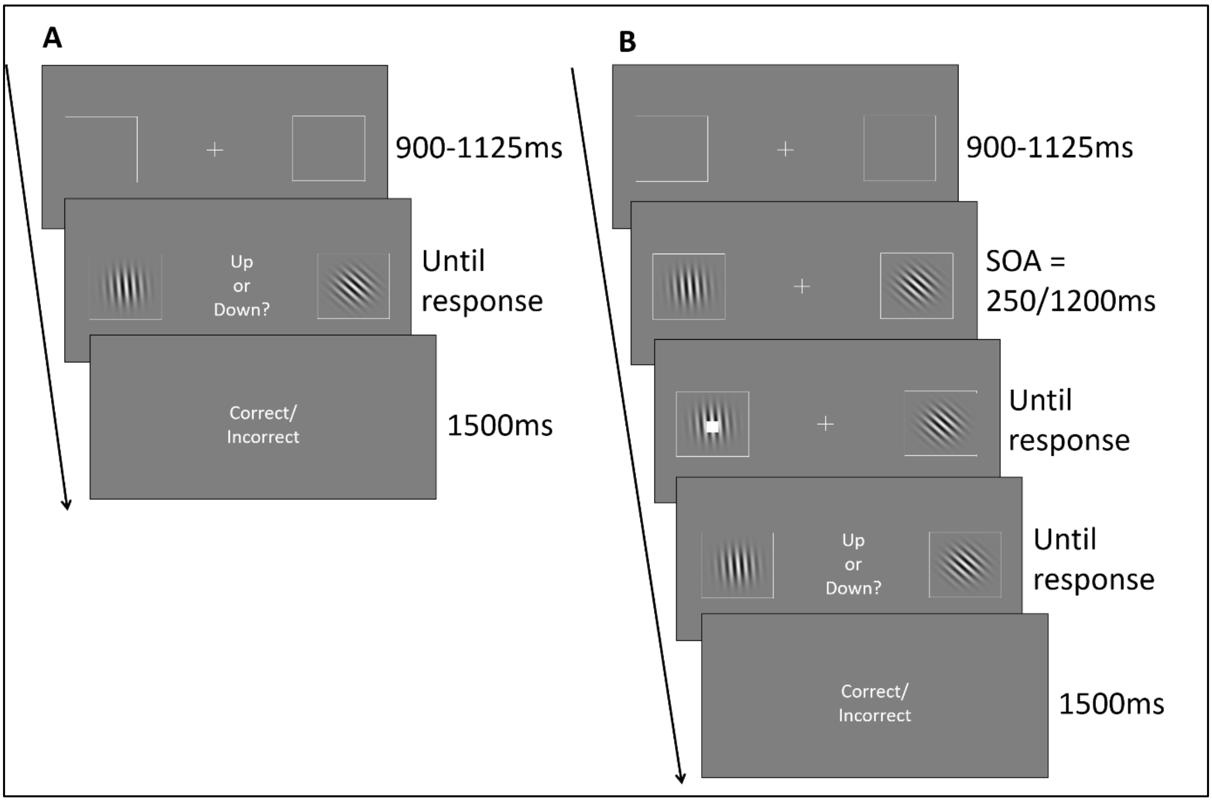

Figure 2.1: Examples of trial types for Phases 1 and 2 of Experiment 1. (A) Phase 1 was a categorisation task. Each trial began with a pre-stimulus screen, which consisted of a central fixation cross that was flanked by two placeholder boxes. After a jittered delay, two grating stimuli of different orientation appeared within each placeholder box. Participants categorised each pair presented to them by pressing either the up or down buttons on a button box. (B) Phase 2 was an intermixed categorisation and dot probe task. It was identical to Phase 1 except that either 250 or 1200ms after the onset of the gratings, a white square target appeared randomly over the grating on the left or right. Participants responded as fast as they could to the location of the square on the screen by pressing either the left or right buttons on the button box. After responding to the target it disappeared from the screen and participants categorised the gratings that remained in the same way as in Phase 1.

EEG Recording and Analysis

Throughout the experiment continuous EEG was recorded from nine tin electrodes embedded in a cap (Electrode-Cap International, Ohio) that was fitted onto the participant's head. Active electrodes were located at sites Fz, F3, F4, Cz, Pz, P3, P4, O1 and O2 according to the International 10-20 system. An additional active electrode was placed on the right earlobe. All electrodes were referenced to the left earlobe. A ground electrode was located at AFz. Impedances were generally kept below 5kΩ and never exceeded 10kΩ. A vertical electrooculogram (EOG) was recorded from electrodes placed above and below the left eye. A horizontal EOG was recorded from electrodes placed to the left and right outer canthi. EEG and EOG were recorded at a sampling rate of 1000 Hz and amplified using a BioNomadix wireless system (Biopac Systems Inc., Goleta, CA, USA). EEG and EOG data were filtered online with a 0.1-100Hz and 0.005-35Hz bandpass filter, respectively.

The data were analysed offline using EEGLAB (Delorme & Makeig, 2004) and ERPLAB (Lopez-Calderon & Luck, 2014). The EEG data was re-referenced to the average of the two earlobes and filtered using a 50Hz notch filter and a 30Hz low-pass filter (12 dB/octave). The continuous EEG was segmented into epochs ranging from 100ms prior to the onset of the dot probe target to 250ms post target onset, and baseline corrected using the 100ms pre-stimulus interval. Events that were contaminated by eye movement and blink artefacts were excluded from ERP averaging. These events were detected by a function in ERPLAB that detects step-like artefacts in the horizontal and vertical EOG channels, as recommended by Luck (2014). Trials in which the participant responded to the dot probe faster than 250ms were also excluded from the EEG analysis to ensure motor responding did not contaminate the epoch of interest. Participant's data were included in the analysis if they had more than 30 valid (artefact-free and with correct response times longer than 250ms) trials in each of the four within-subject conditions (P250, NP250, P1200, NP1200). The data from all participants met this inclusion criterion.

ERP waveforms were time-locked to the onset of the target in the dot probe task. To compute lateralised ERP waveforms we collapsed events over left and right target locations and left and right recording hemispheres. ERPs ipsilateral to the target were computed by averaging the ERPs recorded over the right scalp when the target was on the right with ERPs recorded over the left scalp when the target was on the left. ERPs contralateral to the target were computed by averaging the ERPs recorded over the right scalp when the target was on the left with ERPs recorded over the left scalp when the target was on the right. Lateralised ERPs were computed for each of the four within-subjects conditions (i.e., when the target appeared over the predictive and non-predictive cues at each SOA). The N2pc component was measured from contralateral minus ipsilateral difference waveforms at the P3 and P4 electrode sites. The 50% fractional area latency of the N2pc component was calculated using a time window between 100 and 250ms. Four separate latencies (i.e., one for each within-subjects condition) were calculated for each participant and then averaged. The 50% fractional area latency was chosen because, compared to other measures, it is well suited for making comparisons with reaction time and is less sensitive to noise (Luck, 2014).

As significant differences in N2pc latencies were found between SOA conditions (see Results), the mean amplitude of the N2pc component elicited by the target was measured during different 50ms time windows centred on the median latencies for each SOA (see McDonald et al., 2009 who also used a 50 ms time window post target onset). The median N2pc 50% fractional area latency elicited by targets that appeared after a 250 ms and 1200 ms SOA was 204 ms and 188 ms, respectively. Therefore, the mean amplitude of the N2pc component was calculated as the mean voltage of the contralateral minus ipsilateral difference waveforms at the P3 and P4 electrode sites between 179 and 229ms for the 250ms

SOA and between 163 and 213ms for the 1200ms SOA. Mean amplitude was chosen because, unlike other measures, it is less sensitive to differences in the number of trials between conditions (Luck, 2014).

Statistical Analysis

In order for predictiveness to have an influence on spatial cueing during the dot probe task, adequate associative learning of the Phase 1 contingencies was required (Le Pelley et al., 2013; Luque et al., 2017). Although we had a Phase 1 learning criterion in place, which participants had to pass in order to proceed to Phase 2, it was quite lenient. As such, it was possible that participants could have reached the learning criterion in Phase 1 by chance. Therefore, we also imposed a selection criterion of 60% correct categorisation responses during Phase 2 (see Le Pelley et al., 2013 and Luque et al., 2017 who also excluded participants with less than 60% categorisation accuracy). Participants who did not meet this criterion were removed from further analysis. RTs faster than 150ms or slower than 1500ms were removed from each participant's data. For each participant, a median RT was calculated for each within subject condition (i.e., P250, NP250, P1200 and NP1200). We also calculated a NP-P difference score for each participant for each SOA (i.e., NP250-P250 and NP1200-P1200). As a guard against outliers, any participant who had a difference score lying three or more standard deviations away from the mean of the P-NP difference scores in either SOA was removed from the RT analysis. The effects of predictiveness (target over the P or NP cues) and SOA (250ms or 1200ms) on median RTs were compared using a two-way repeated-measures ANOVA. The effects of predictiveness and SOA on the N2pc latencies and N2pc mean amplitudes were also compared using two separate two-way repeated-measures ANOVAs. Unbiased effect sizes (d_{unb}), sometimes referred to as Hedges' g (Hedges, 1981), are reported (Bornstein, Hedges, Higgins, & Rothstein, 2009).

2.1.2. Results

One participant failed to reach the Phase1 learning criterion and did not proceed to Phase 2. Of the 23 remaining participants, 21 completed the Phase 1 categorisation task in the minimum number of blocks (20), with two participants requiring 21 and 24 blocks to reach the learning criterion. Categorisation accuracy remained high during Phase 2 (mean accuracy = 97%, range = 87-100%) and the number of errors made during the dot probe task was very low across all participants (mean errors= 1, range = 0 -4). One participant was removed from the RT analysis because they had RT difference scores that lay more than three standard deviations away from the average.

Figure 2.2 shows median RTs to the target in the dot probe task as a function of predictiveness and SOA. There was a significant main effect of SOA on median RT such that RTs to the target were slower when the target appeared after a short SOA compared to a long SOA ($F_{(1,21)} = 50.26$, $p < .001$, $\eta^2 = 0.657$, MSE = 712.30, $d_{unb} = 0.56$). The main effect of predictiveness was not significant ($F_{(1,21)} = 1.45$, $p = .243$, $\eta^2 = 0.003$, MSE = 96.89, $d_{unb} = 0.045$), nor was the predictiveness×SOA interaction ($F_{(1,21)} = 0.87$, $p = .362$, $\eta^2 = 0.001$, MSE = 69.60).

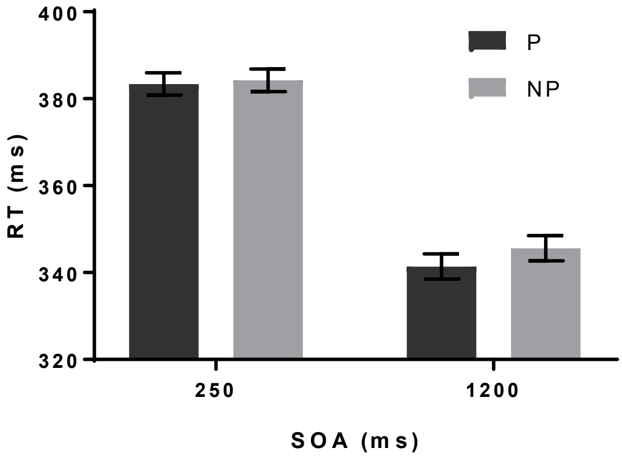

Figure 2.2: Median reaction times to the location of the white square target during the dot probe task in Experiment 1 as a function of SOA and predictiveness. Errors bars represent standard error of the mean.

Figure 2.3 shows the ERP waveforms recorded from lateral occipital electrode sites (P3 and P4) time-locked to the onset of the white square target in each of the four within-subjects conditions (P250, NP250, P1200 and NP1200), averaged across participants. [2] Separate waveforms are shown for electrodes that were contralateral and ipsilateral to the target. Figure 2.3 also shows contralateral minus ipsilateral difference waveforms. The N2pc component was present as a negative deflection beginning at approximately 150ms after the onset of the target. All latency and amplitude analyses were conducted on these difference waves.

---

[2] Note the non-zero baseline period in the 250ms SOA condition. Presumably, this was due to residual activity from the cues, which only shortly preceded the target in this condition. In the 1200ms SOA condition a flat baseline period can be seen. We were only interested in comparing ERPs within each SOA (i.e. P250 vs. NP250 and P1200 vs. NP1200). Therefore, the non-zero baseline in the short SOA condition does not impact our comparisons, as the residual activity from the cues was identical between the P250 and NP250 conditions.

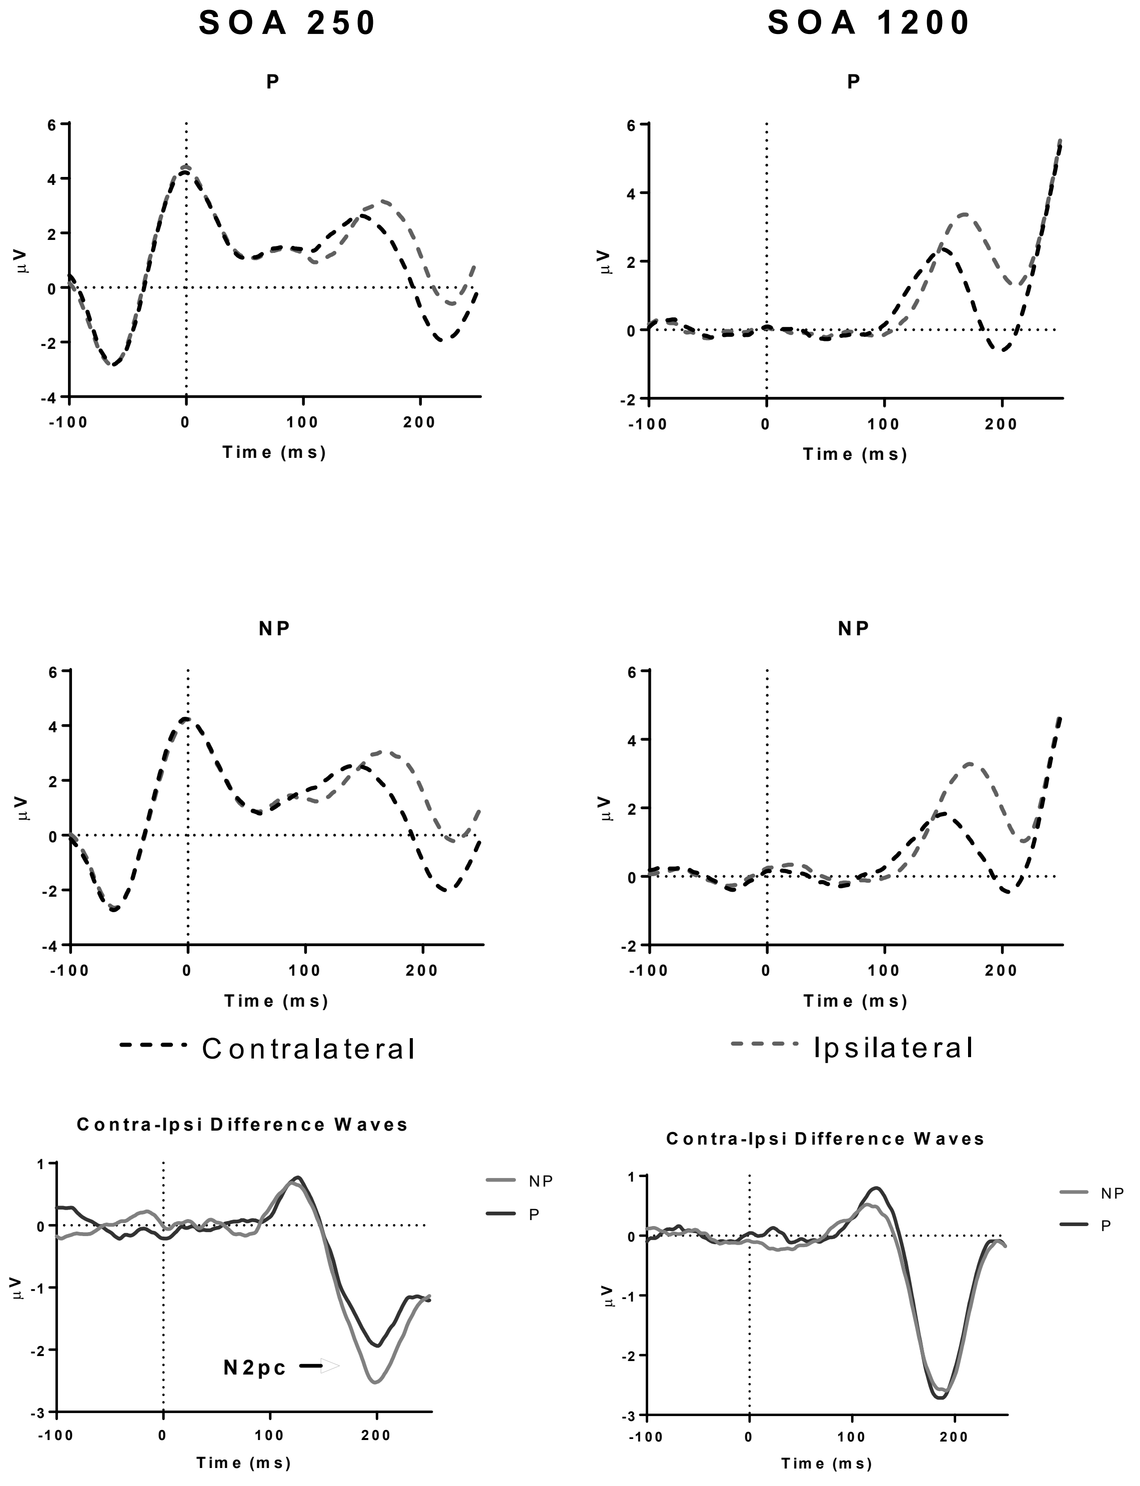

Figure 2.3: Grand average ERP waveforms from Experiment 1 for each within-subjects condition at contralateral versus ipsilateral P3/P4 electrode sites, along with contralateral minus ipsilateral ERP differences waves. All waveforms were time-locked to the onset of the white square target.

Figure 2.4 shows the median 50% fractional area latencies of the target elicited N2pc as a function of predictiveness and SOA. There was a significant main effect of SOA on latency such that latencies to the target were longer when the target appeared after a short SOA compared to a long SOA ($F_{(1,22)} = 44.30$, $p < .001$, $\eta^2 = 0.493$, MSE = 83.68, $d_{unb} = 2.34$). The main effect of predictiveness was not significant ($F_{(1,22)} = 0.08$, $p = .781$, $\eta^2 = 0.001$, MSE = 54.67, $d_{unb} = 0.039$), nor was the predictiveness×SOA interaction ($F_{(1,22)} = 0.15$, $p = .699$, $\eta^2 = 0.001$, MSE = 34.31).

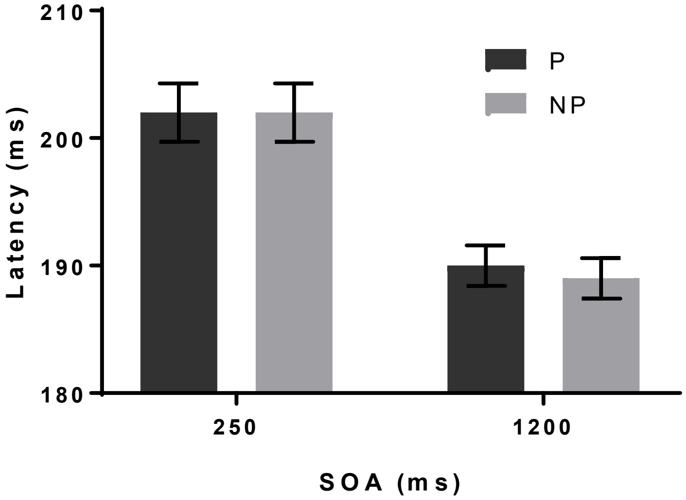

Figure 2.4: Latencies of the N2pc at the P3/P4 electrode sites elicited by the white square target during the dot probe task in Experiment 1 as a function of SOA and predictiveness. Errors bars represent standard error of the mean.

Figure 2.5 shows the average of the N2pc mean amplitudes evoked by the target as a function of predictiveness and SOA. The main effect of SOA was significant because targets that appeared after a short SOA elicited smaller amplitudes than targets that appeared after a long SOA ($F_{(1,22)} = 4.58$, $p = .044$, $\eta^2 = 0.083$, MSE = 0.47, $d_{unb} = 0.34$). There was a significant main effect of predictiveness ($F_{(1,22)} = 7.81$, $p = 0.011$, $\eta^2 = 0.059$, MSE = .196, $d_{unb} = 0.30$), such that targets that appeared over NP cues had a greater negative mean amplitude (M = -2.19μV) than those that appeared over P cues (M = -1.93 μV). However, there was also a significant predictiveness×SOA interaction ($F_{(1,22)} = 6.63$, $p = .0.17$, $\eta^2 = 0.067$, MSE = 0.26). This interaction was investigated using paired sample t-tests to compare the N2pc mean amplitudes elicited by the target when it appeared over the P and NP cues at each SOA. As shown in Figure 2.5, the mean amplitude of the N2pc was significantly larger (more negative) when the target appeared over a NP cue compared to a P cue at the short SOA ($t_{(22)} = 3.29$, $p = 0.003$, $d_{unb} = 59$). At the long SOA there was no difference in the N2pc mean amplitude to the target over the P and NP cues ($t_{(22)} = 0.15$, $p = .88$, $d_{unb} = .017$).

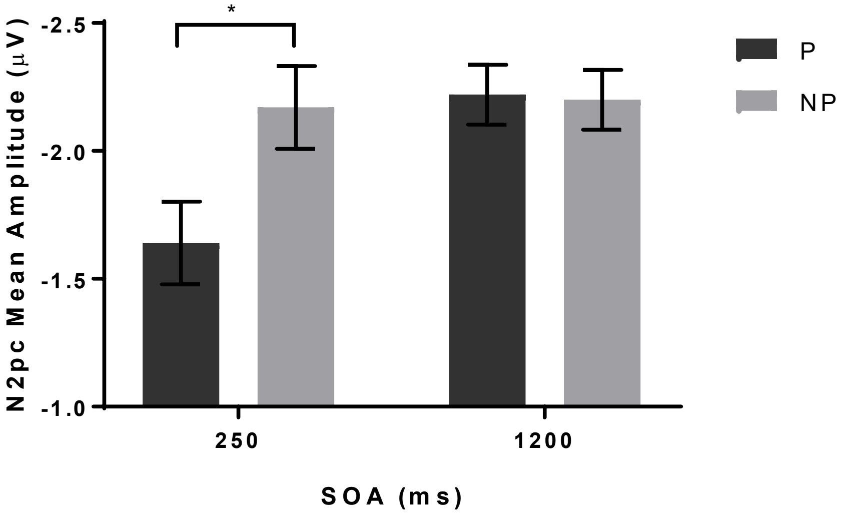

Figure 2.5: Mean amplitudes of the N2pc at the P3/P4 electrode sites elicited by the white square target during the dot probe task in Experiment 1 as a function of SOA and predictiveness. Errors bars represent standard error of the mean (* $p < .05$).

2.1.3. Discussion

It is clear from the RT results (Figure 2.2) that we were unable to replicate the RT facilitation effect at the 250ms SOA reported by previous studies (Le Pelley et al., 2013; Luque et al., 2017). We were also unable to find behavioural evidence for IOR. However, the N2pc results were interesting. The main effect of SOA found in RTs (Figure 2.2) was nicely complemented by the main effect of SOA found in the N2pc latencies (Figure 2.4). This could be taken as evidence that the N2pc is measuring attentional processes that affect responding (i.e., the aging-foreperiod effect; Drazin, 1961). Although no evidence of IOR was observed at the 1200ms SOA, there was an intriguing N2pc amplitude effect at the 250ms SOA. Targets that appeared 250ms after cue onset elicited a significantly smaller N2pc mean amplitude when they appeared over the P cues compared to the NP cues (Figure 2.5). We believe this was because participants were already attending to the P cues before the target appeared. Therefore, processing the white square over a P cue did not require a large shift in attention. Conversely, processing the white square over a (presumably unattended) NP cue required a greater shift in attention.

Admittedly, this interpretation is post hoc and is in the opposite direction that one might expect from the N2pc literature. That is, stimuli that elicit a greater negative voltage at contralateral electrodes are usually interpreted as receiving more attention. However, it seems unlikely that this interpretation holds in the current case for two reasons. Firstly, the attentional bias set up by the categorisation task in Phase 1, and the fact that participants needed to attend to P cues to make a categorisation in Phase 2, predicts an initial shift of covert attention in the direction of the P cues. Secondly, the RT facilitation effect at short SOAs reported by previous studies (Le Pelley et al., 2013; Luque, et al., 2017) supports the idea that attention initially shifts towards the P cue. However, as already noted, no RT facilitation effect was found in Experiment 1. Therefore, the aim of Experiment 2 was to find supporting evidence of our N2pc mean amplitude interpretation by replicating the RT facilitation effect at the short SOA.

2.2. Experiment 2

Experiment 2 was designed to replicate the previously reported RT facilitation effect in favour of the P cues at a 250ms SOA (Le Pelley et al., 2013; Luque et al., 2017). One obvious difference between Experiment 1 and the studies that have reported the RT effect is the nature of the cues. Specifically, our cues were black and white and needed to be discriminated by orientation, whereas previous studies have used an easier colour discrimination for their cues. It is possible that the discrimination difficulty of the cues preceding the target affects the processing of the target itself. This altered target processing could have influenced behavioural responses, resulting in an elimination of the RT effect at the 250ms SOA. In Experiment 2, we added colour to our gratings so that participants could use colour and/or orientation to discriminate between the P and NP cues.

2.2.1. Method

Participants

Twenty-seven volunteers (12 male; mean age 20 years, age range 18 to 25 years) with normal or corrected to normal vision participated in the experiment. The same eligibility criteria as Experiment 1 applied to Experiment 2.

Apparatus

All apparatus were identical to those described in Experiment 1.

Stimuli

The four cues were sinusoidal luminance gratings (spatial frequency = 0.05, phase = 0) that were spatially modulated in contrast by a Gaussian envelope (standard deviation = 30) and had orientations of 4° (near-vertical), 45° (tilted to the left), 94° (near-horizontal), or 135° (tilted to the right). Near-horizontal and near-vertical gratings were used to avoid preferential selection that has been shown to occur for non-oblique orientations (Meigen, Lagreze, & Bach, 1994). Blue (34.2 cd/m^2), green (34.0 cd/m^2), orange (34.5 cd/m^2) and pink (34.8 cd/m^2) gratings with approximately equal luminance were chosen (one for each orientation).

Procedure

Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n =13) and “Left-Right Predictive” (n=14). Seven participants in the “Vertical-Horizontal Predictive” and eight participants in the “Left-Right Predictive” condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings, and green left gratings.

All tasks and instructions were identical to those used in Experiment 1.

EEG Recording and Analysis

Similarly to Experiment 1, significant differences in N2pc latencies were found between SOA conditions. The median N2pc 50% fractional area latency elicited by targets that appeared after a 250ms and 1200ms SOA was 197ms and 188ms, respectively. Therefore, the mean amplitude of the N2pc component was calculated as the mean voltage of the contralateral minus ipsilateral difference waveforms at the P3 and P4 electrode sites between 172 and 222ms for the 250ms SOA and between 163 and 213ms for the 1200ms SOA. All other EEG recording and analyses were identical to Experiment 1.

Statistical Analysis

Statistical analyses and exclusion criteria were identical to Experiment 1.

2.2.2. Results

One participant failed to reach the Phase1 learning criterion and did not proceed to Phase 2. Of the 26 remaining participants, 24 completed the Phase 1 categorisation task in the minimum number of blocks (20), with two participants requiring 21 and 31 blocks to reach the learning criterion. Categorisation accuracy remained high during Phase 2 (mean accuracy = 96%, range = 74%-100%) and the number of errors made during the dot probe task was very low across all participants (mean errors = 1, range = 0 -5). One participant was removed from the RT analysis because they had RT difference scores that lay more than three standard deviations away from the average.

Figure 2.6 shows the average of median RTs to the target in the dot probe task as a function of predictiveness and SOA. There was a significant main effect of SOA on median RTs ($F_{(1,24)} = 38.57$, $p < .001$., $\eta^2 = 0.51$, MSE = 562.90, $d_{unb} = 0.45$) such that RTs to the target were slower when the target appeared after a short SOA (M = 379 ms) compared to a long SOA (M = 350 ms). There was also a main effect of predictiveness ($F_{(1,24)} = 10.52$, p = .0.003, $\eta^2 = 0.032$, MSE = 131.25, $d_{unb} = 0.12$) because participants responded faster to targets that appeared over P cues (M = 361 ms) compared to NP cues (368 ms). However, there was a significant predictiveness×SOA interaction ($F_{(1,24)} = 4.76$, $p = .039$, $\eta^2 = 0.011$, MSE = 94.89). This interaction was investigated using paired sample t-tests to compare the RTs to the target when it appeared over the P and NP cues at each SOA. As shown in Figure 2.6, RTs were significantly faster when the target appeared over the P cues compared to the NP cues at the short SOA ($t_{(24)} = 3.95$, $p = .001$, $d_{unb} = .17$). At the long SOA there was no difference in RTs to the target over the P and NP cues ($t_{(24)} = 1.04$, $p = .31$, $d_{unb} = .047$).

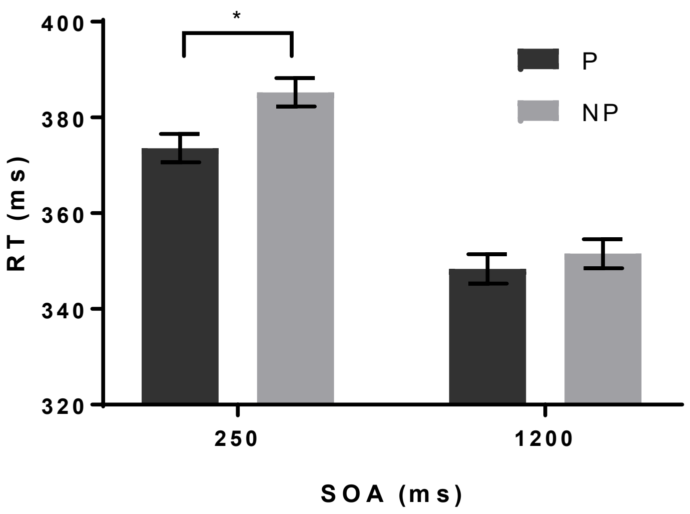

Figure 2.6: Median reaction times to the location of the white square target during the dot probe task in Experiment 2 as a function of SOA and predictiveness. Errors bars represent standard error of the mean (* p < .05).

Figure 2.7 shows the ERP waveforms recorded from lateral occipital electrode sites (P3 and P4) time-locked to the onset of the white square target in each of the four within-subjects conditions (P250, NP250, P1200 and NP1200), averaged across participants.

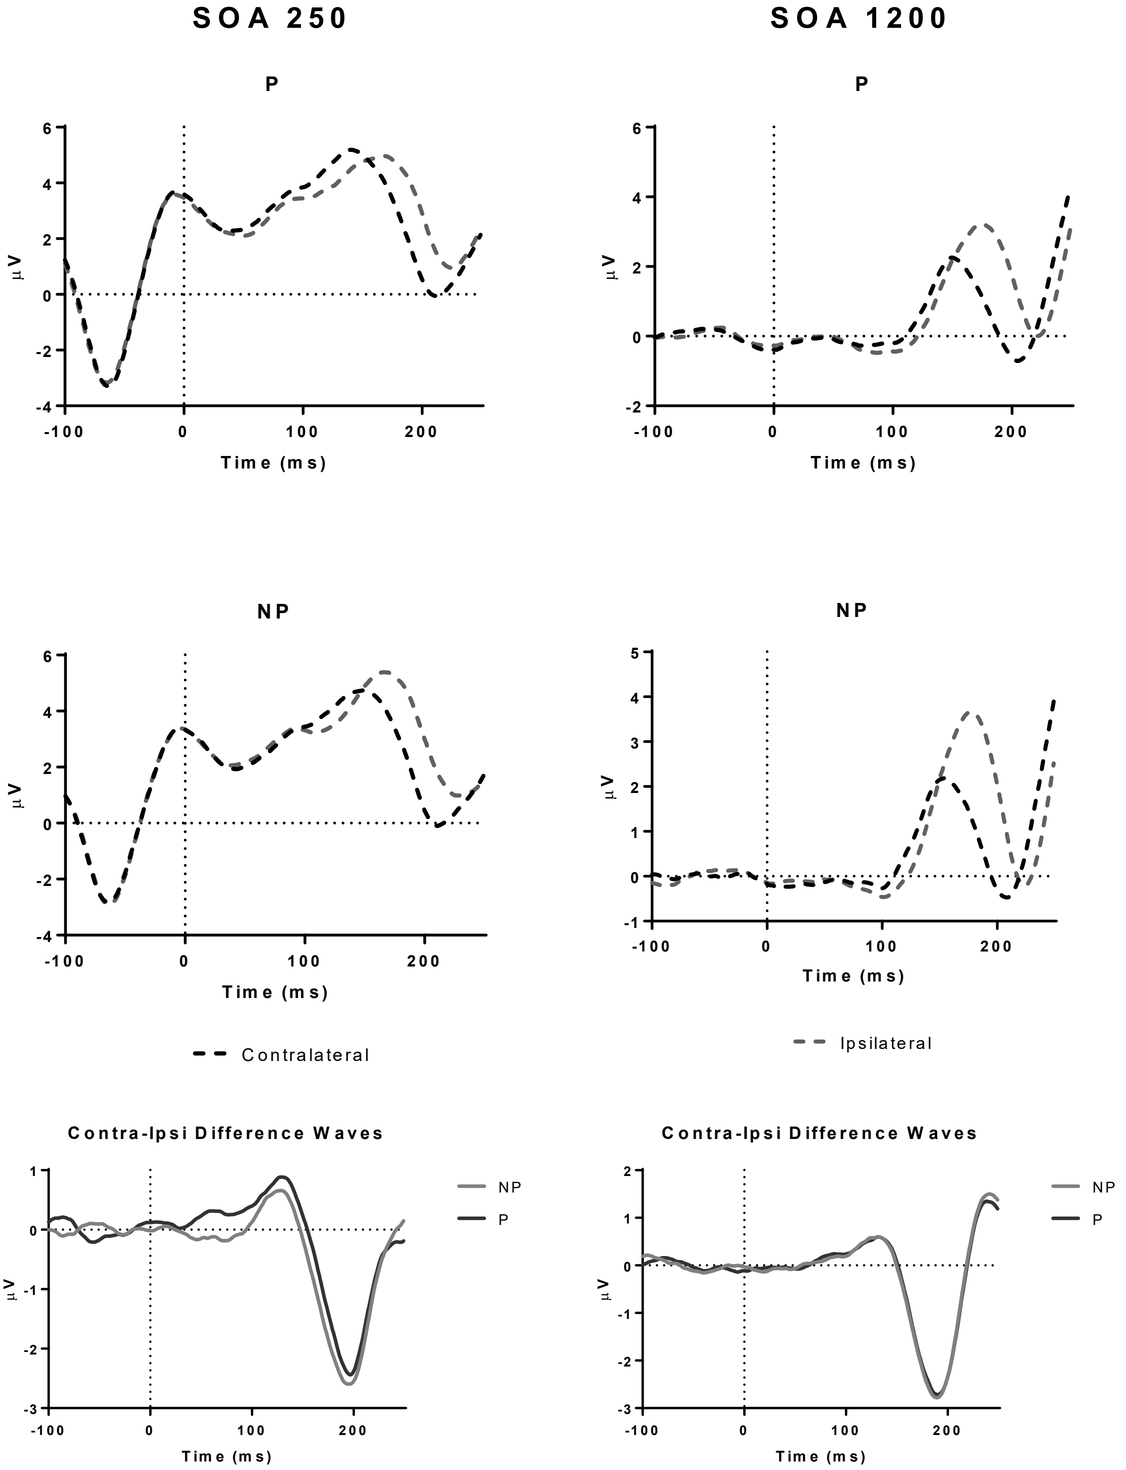

Figure 2.7: Grand average ERP waveforms from Experiment 1 for each within-subjects condition at contralateral versus ipsilateral P3/P4 electrode sites, along with contralateral minus ipsilateral ERP differences waves. All waveforms were time-locked to the onset of the white square target.

Figure 2.8 shows the median 50% fractional area latencies of the target elicited N2pc as a function of predictiveness and SOA. There was a significant main effect of SOA on latency such that latencies to the target were longer when the target appeared after a short SOA compared to a long SOA ($F_{(1,24)} = 15.01$, $p = .001$, $\eta^2 = 0.230$, MSE = 103.96, $d_{unb} = 0.66$)[3]. The main effect of predictiveness was not significant ($F_{(1,24)} = 0.25$, $p = .625$, $\eta^2 = 0.002$, MSE = 62.17, $d_{unb} = 0.11$), nor was the predictiveness×SOA interaction ($F_{(1,24)} = 0.19$, $p = .666$, $\eta^2 = 0.001$, MSE = 50.19).

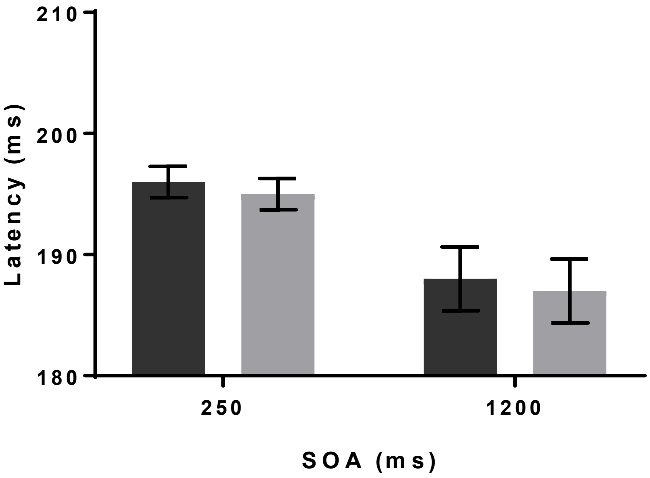

Figure 2.8: Latencies of the N2pc at the P3/P4 electrode sites elicited by the white square target during the dot probe task in Experiment 2 as a function of SOA and predictiveness. Errors bars represent standard error of the mean.

Figure 2.9 shows the average of the N2pc mean amplitudes evoked by the target as a function of predictiveness and SOA. The main effect of SOA was not significant ($F_{(1,25)} = 0.67$, $p = .422$, $\eta^2 = 0.011$, MSE = .722, $d_{unb} = 0.093$), nor was the main effect of

---

[3] The degrees of freedom is 24 because one participant did not elicit an N2pc, and hence no N2pc latency, to the target when it appeared over a P cue at the 250ms SOA.

predictiveness ($F_{(1,25)} = .96$, $p = .335$, $\eta^2 = 0.012$, MSE = .567, $d_{unb} = 0.095$) or the predictiveness×SOA interaction ($F_{(1,25)} = 0.65$, $p = .43$, $\eta^2 = 0.007$, MSE = .464).

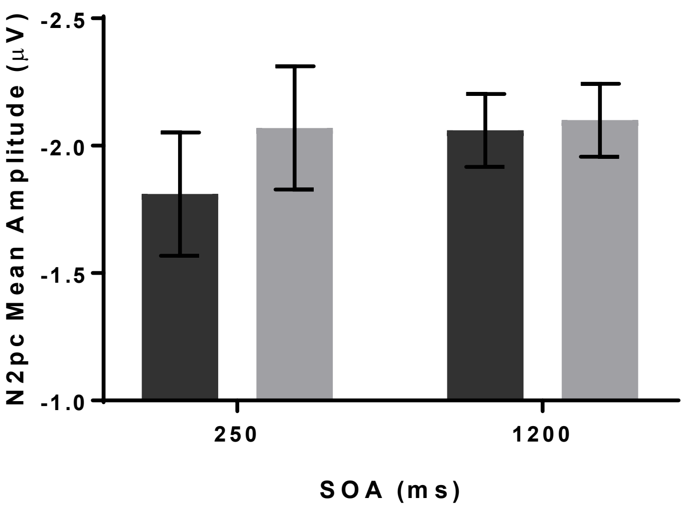

Figure 2.9: Mean amplitudes of the N2pc at the P3/P4 electrode sites elicited by the white square target during the dot probe task in Experiment 2 as a function of SOA and predictiveness. Errors bars represent standard error of the mean.

2.2.3. Discussion

The RT results showed that participants were faster to respond to the target when it appeared over the P cues compared to the NP cues after a 250ms SOA (Figure 2.6). Therefore, Experiment 2 successfully replicated the RT facilitation effect found by Le Pelley et al. (2013) and Luque, et al. (2017). Similarly to Experiment 1, there was a main effect of SOA in the N2pc latency results (Figure 2.8), which mirrored the main effect of SOA found in RTs (Figure 2.6). In contrast to Experiment 1, the N2pc mean amplitude effect at the 250ms SOA was not significant (Figure 2.9), although the effect was trending in the same direction. Once again, the direction of the effect was opposite to what would be expected given the definition of the N2pc and appears to indicate that attention was biased towards the NP cue. However, the RT facilitation effect towards the P cue indicates that attention was allocated to the P cue. Therefore, it seems likely that the N2pc effect reported here actually reflects attention towards the P cue.

Similarly to Experiment 1, there was no evidence of behavioural IOR at the long SOA, or any evidence of IOR in the mean amplitudes of the N2pc. A possible reason for this could be floor effects at the long SOA in both the RT and N2pc latency data. By looking at Figures 2.6 (RTs) and 2.8 (target elicited N2pc latencies), it is plausible that processing of the white square target at the 1200 ms SOA occurred too fast for differences in processing as a function of predictiveness to be observed.

The aim of Experiment 3 was to test whether floor effects were attributing to the lack of RT effects at the long SOA. We attempted to slow down the participants' responding by changing the dot probe task from a localisation task to a discrimination task. IOR has been shown to occur in discrimination tasks, albeit at longer SOAs compared to simple detection tasks (Lupianez, Milan, Tornay, Madrid, & Tudela, 1999). An additional aim of Experiment 3 was to investigate whether the N2pc effect reflects processing of the P or NP cue. As of yet, a significant RT and N2pc effect have not been observed within the same experiment (or the same SOA). In Experiment 1, there was an N2pc effect after the short SOA, but no RT effect. In Experiment 2, there was a RT facilitation effect after the short SOA, but the N2pc effect was not significant. At present, it is difficult to see why these two measures of attention do not coincide. In Experiment 3, another behavioural measure was included (ratings) in attempt to gain further behavioural evidence that might help interpret the N2pc effect. That is, if participants are paying attention to the P cue during the dot probe task then they may be more likely to rate that they notice the target appear over the P cue compared to the NP cue (i.e., a learned predictiveness effect).

2.3. Experiment 3

In Experiment 3 we changed the target from a square to a sideways T, and used a target-discrimination task instead of a target-localisation task. Participants were required to respond to the location of the vertical line of the sideways T (i.e., left or right side of the T), irrespective of the location of the T on the screen. This would allow us to test whether floor effects were responsible for the lack of IOR observed in Experiments 1 and 2.

In neither Experiment 1 nor Experiment 2 were we able to obtain both a N2pc effect and a RT effect at the same SOA. This suggests that the two measures are decoupled for some reason. Therefore, we have not been able to confirm that our N2pc effect reflects attentional facilitation towards the P cue with RT data alone. At the end of Experiment 3 participants were asked to rate how often they noticed the sideways T target appear over each one of the cues. This was done to obtain another behavioural measure to help interpret the N2pc effect obtained in Experiment 1. It also provided an opportunity to replicate the learned predictiveness effect using a dot probe task. A learned predictiveness effect occurs when participants associate novel outcomes more readily with cues that have been previously learned to be predictive compared to cues that have been previously learned to be non-predictive (Le Pelley et al., 2016). If the sideways T target in the dot probe task is considered to be a novel outcome occurring over the previously learned gratings, then participants should report that the target appeared over the P gratings more often compared to the NP gratings.

2.3.1. Method

Participants

Twenty-seven volunteers (9 male; mean age 21 years, age range 18 to 29 years) with normal or corrected to normal vision participated in the experiment. The same eligibility criteria as Experiments 1 and 2 applied to Experiment 3.

Apparatus

All apparatus were identical to those described in Experiment 1.

Stimuli

All grating stimuli were identical to those used in Experiment 2. The target in the dot probe task was now a white sideways T. This was made up of a horizontal line with a length of 0.8° and height of 0.18° connected to the centre of vertical line with a height of 0.8° and length of 0.18°. The sideways T target appeared in the centre of one of the white square frames, superimposed over one of the grating cues. The screen background and viewing distance were identical to those used in Experiments 1 and 2.

Procedure

Participants were alternately assigned to the conditions "Vertical-Horizontal Predictive" (n =12) and "Left-Right Predictive" (n=15)[4]. Six participants in the "Vertical-Horizontal Predictive" and six participants in the "Left Right Predictive" condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings, and green left gratings.

The instructions for the Phase 1 categorisation task were identical to Experiments 1 and 2. However, the instructions for the Phase 2 dot probe task were modified to suit the sideways T discrimination task. Participants were asked to respond to the location of the vertical line on either side of the sideways T by pressing the corresponding button on the button box as fast as possible. The importance of both speed and accuracy was emphasised. The dot probe task was split into blocks of 64 trials. Each block consisted of all possible

---

[4] The three participants excluded from this experiment (see Results) were all assigned to the "Left-Right Predictive" condition, and were replaced by three additional participants.

combinations of four stimulus pairs (Table 1.1), two cue locations (predictive or non-predictive cue on the left or right), two SOAs (250ms or 1200ms), two target locations on the screen (left or right), and two locations of the vertical line within the sideways T (left or right). The 64 trials within a block were presented in random order. Participants completed five blocks of the dot probe task (320 trials). The same four within-subjects conditions (P250, NP250, P1200, and NP1200) as in Experiments 1 and 2 applied to Experiment 3. So as not to influence the rating data obtained at the end of Phase 2, the participants were not informed that the target would appear randomly over the gratings.

At the end of Phase 2, participants were asked to rate how often they noticed the sideways T target appear over each of the cues. Each cue was shown individually in a random order and the participant entered a number on the keyboard from 1 (Rarely) to 6 (Often). This resulted in a rating for each cue.

EEG Recording and Analysis

Similarly to Experiments 1 and 2, significant differences in N2pc latencies were found between SOA conditions. The 50% fractional area latency of the N2pc component was calculated using a time window between 100 and 350ms. The median N2pc 50% fractional area latency elicited by targets that appeared after a 250ms and 1200ms SOA was 206ms and 196ms, respectively. Therefore, the mean amplitude of the N2pc component was calculated as the mean voltage of the contralateral minus ipsilateral difference waveforms at the P3 and P4 electrode sites between 181 and 231ms for the 250ms SOA and between 171 and 221ms for the 1200ms SOA. All other EEG recording and analyses were identical to Experiments 1 and 2.

Statistical Analysis

For each participant, we averaged their ratings for the P cues and for the NP cues, which resulted in one rating for each cue type. The P and NP ratings were compared using a paired sample t-test. All other exclusion criteria and statistical analyses were identical to Experiments 1 and 2.

2.3.2. Results

One participant's data was removed due to excessive eye movement artefacts during Phase 2. Two additional participants had their data removed due to low Phase 2 categorisation accuracies of 50% and 54%. These participants passed the learning criterion in Phase 1 (in 29 and 34 blocks, respectively), but may have done so by chance. Three additional participants were recruited and assigned to the same conditions as the excluded participants. All included participants completed the Phase 1 categorisation task in the minimum number of blocks (20). Categorisation accuracy remained high during Phase 2 (mean accuracy = 95%, range = 64%-99.7%). Not surprisingly, the number of errors made during the dot probe task was higher compared to the previous experiments due to the increased difficulty associated with the discrimination task. Nevertheless, the number of errors made during the dot probe task was still very low across all participants (mean errors= 9, range = 1 - 39). One participant was removed from the RT analysis because they had a RT difference score that lay more than three standard deviations away from the average.

Figure 2.10 shows the average of median RTs to the target in the dot probe task as a function of predictiveness and SOA. There was a significant main effect of SOA on median RT such that RTs to the target were slower when the target appeared after a short SOA (M = 531 ms) compared to a long SOA (M = 520 ms; $F_{(1,22)} = 4.64$, p = .042, η^2 = 0..073, MSE = 559.90, d_{unb} = 0.14). There was also a significant main effect of predictiveness such that participants responded faster to the target when it appeared over the P cues (M = 520 ms) compared to the NP cues (M = 530 ms; $F_{(1,22)} = 8.51$, $p = .008$, $\eta^2 = 0.069$, MSE = 288.10, $d_{unb} = 0.13$). Although the predictiveness×SOA interaction was not significant ($F_{(1,22)} = 2.03$, $p = .168$, $\eta^2 = 0.028$, MSE = 495.10), *a priori* comparisons of interest were investigated using paired sample t-tests on RTs to the target when it appeared over the P and NP cues at each SOA. As shown in Figure 2.10, at the 250ms SOA there was no difference in RTs when the target appeared over the P or NP cues ($t_{(22)} = 0.76$, $p = 0.456$, $d_{unb} = 0.05$). However, at the 1200ms SOA, participants were significantly faster to respond when the target appeared over the P cues compared to the NP cues ($t_{(22)} = 2.55$, $p = 0.018$, $d_{unb} = 0.19$).

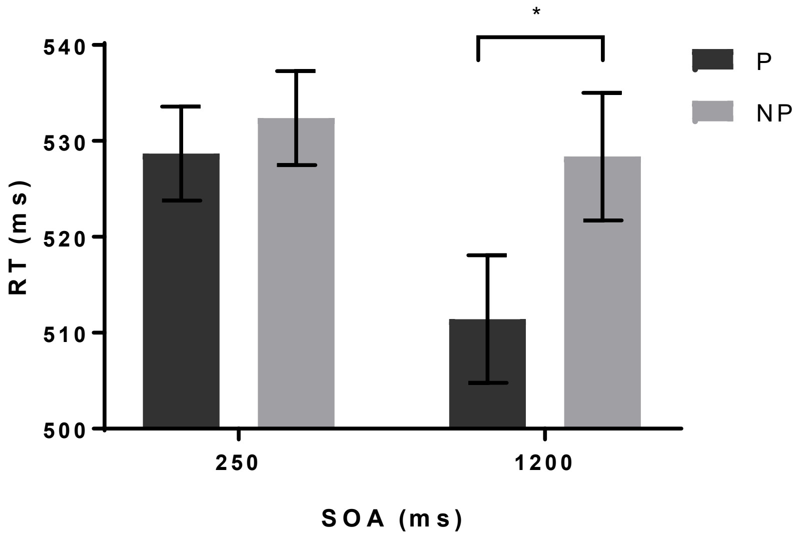

Figure 2.10: Median reaction times to the location of the vertical line within the sideways T target during the dot probe task in Experiment 3 as a function of SOA and predictiveness. Errors bars represent standard error of the mean (* $p < .05$).

Figure 2.11 shows the ERP waveforms recorded from lateral occipital electrode sites (P3 and P4) time-locked to the onset of the sideways T target in each of the four within-subjects conditions (P250, NP250, P1200 and NP1200), averaged across participants.

Figure 2.12 shows the median 50% fractional area latencies of the target elicited N2pc as a function of predictiveness and SOA. There was a significant main effect of SOA on latency such that latencies to the target were longer when the target appeared after a short SOA compared to a long SOA ($F_{(1,23)} = 17.23$, $p < .001$, $\eta^2 = 0.218$, MSE = 141.05, $d_{unb} = 0.69$). The main effect of predictiveness was not significant ($F_{(1,23)} = 0.05$, $p = .825$, $\eta^2 = 0.0005$, MSE = 110.77, $d_{unb} = 0.05$), nor was the predictiveness×SOA interaction ($F_{(1,23)} = 0.61$, $p = .444$, $\eta^2 = 0.007$, MSE = 123.91).

Figure 2.13 shows the average of the N2pc mean amplitudes evoked by the target as a function of predictiveness and SOA. The main effect of SOA was not significant ($F_{(1,23)} = 0.44$, $p = .52$, $\eta^2 = .006$, MSE = 0.61, $d_{unb} = .069$). The main effect of predictiveness was significant ($F_{(1,23)} = 5.40$, $p = .029$, $\eta^2 = 0.08$, MSE = 0.67, $d_{unb} = .28$), such that targets that appeared over NP cues had a greater negative mean amplitude (M = -2.69μV) than those that appeared over P cues (M = -2.30 μV). The predictiveness×SOA interaction was not significant ($F_{(1,23)} = 3.26$, $p = .084$, $\eta^2 = 0.033$, MSE = 0.47). Although the predictiveness×SOA interaction did not reach statistical significance, *a priori* comparisons of interest were investigated using paired sample t-tests on the N2pc mean amplitudes elicited by the target when it appeared over the P and NP cues at each SOA. As shown in Figure 2.13, the mean amplitude of the N2pc was significantly larger (more negative) when the target appeared over the NP cues compared to the P cues at the short SOA ($t_{(23)} = 2.58$, $p = .017$, $d_{unb} = .47$). At the long SOA there was no difference in the N2pc mean amplitude to the target over the P and NP cues ($t_{(23)} = 0.76$, $p = .46$, $d_{unb} = .07$).

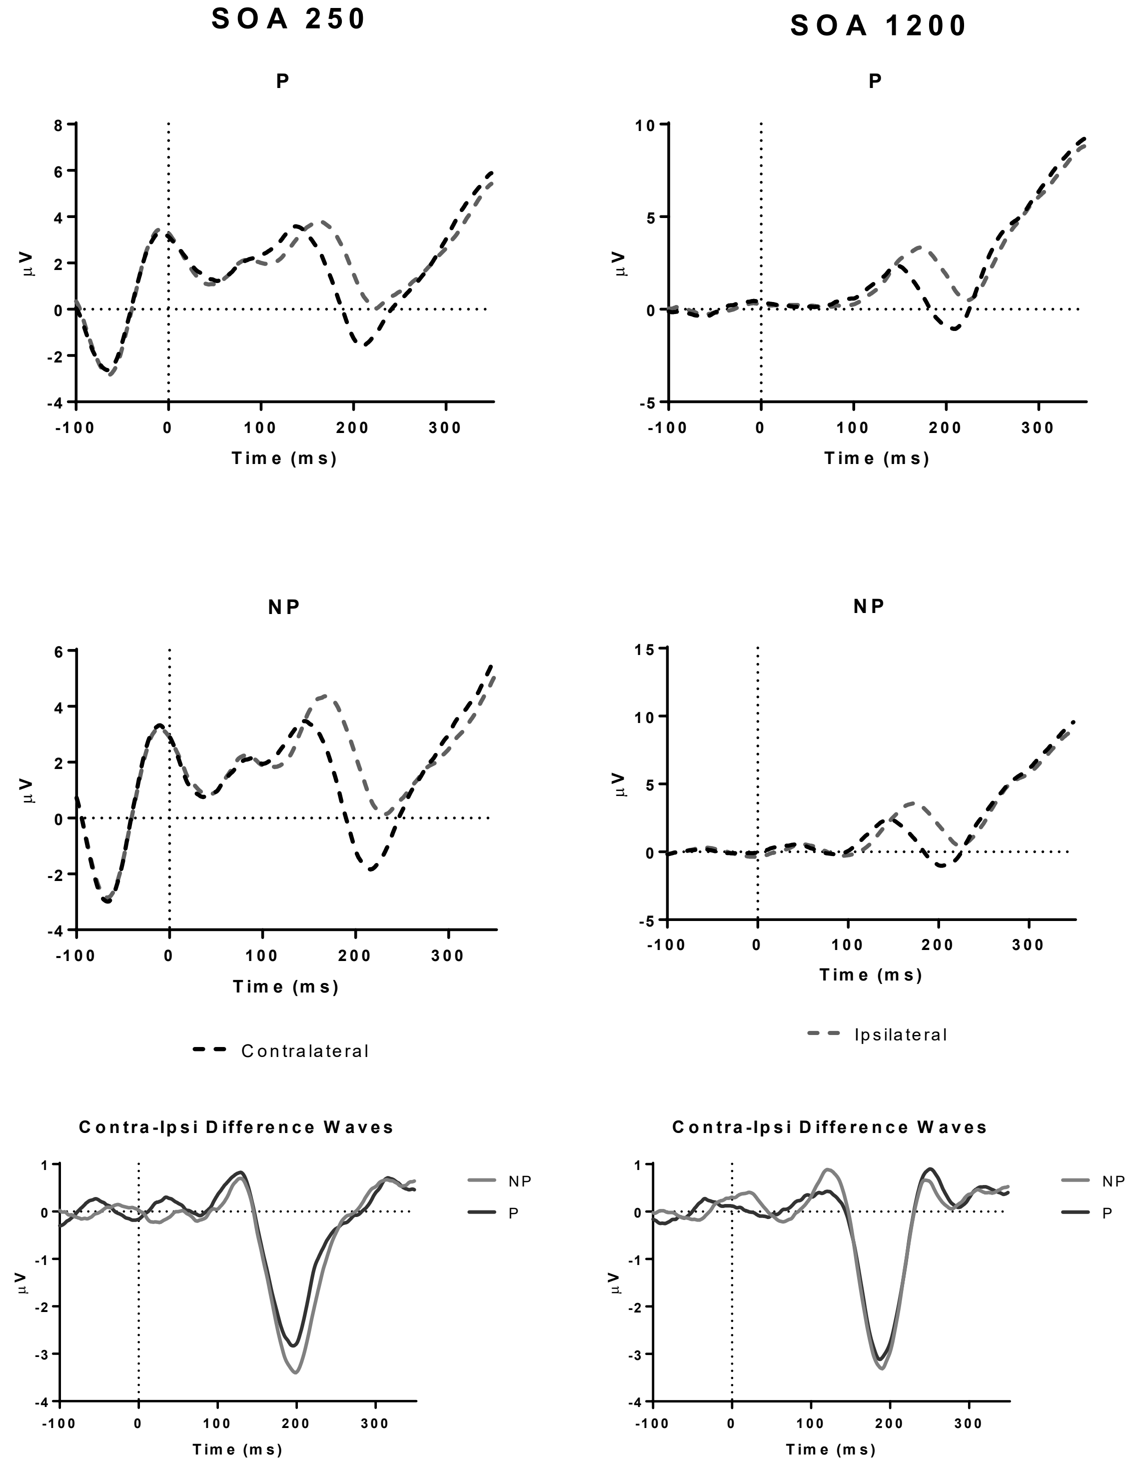

Figure 2.11: Grand average ERP waveforms from Experiment 3 for each within-subjects condition at contralateral versus ipsilateral P3/P4 electrode sites, along with contralateral minus ipsilateral ERP differences waves. All waveforms were time-locked to the onset of the white square target.

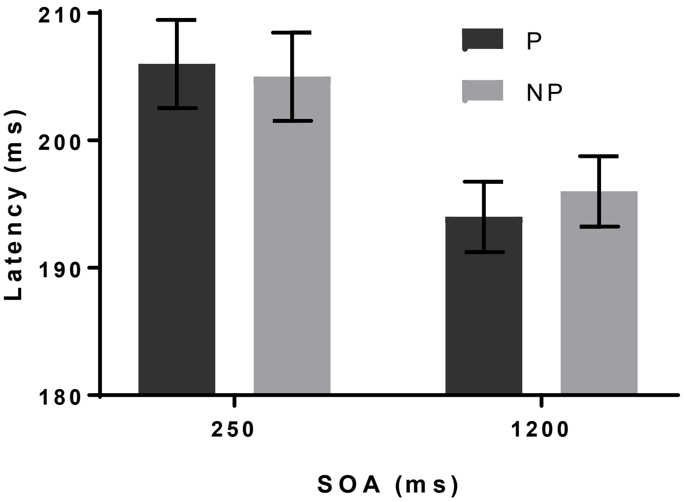

Figure 2.12: Latencies of the N2pc at the P3/P4 electrode sites elicited by the white sideways T target during the dot probe task in Experiment 3 as a function of SOA and predictiveness. Errors bars represent standard error of the mean.

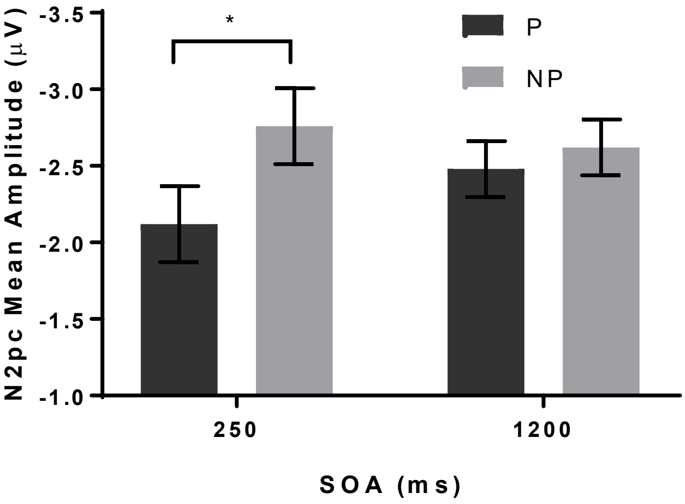

Figure 2.13: Mean amplitudes of the N2pc at the P3/P4 electrode sites elicited by the white square target during the dot probe task in Experiment 3 as a function of SOA and predictiveness. Errors bars represent standard error of the mean (* $p < .05$).

A paired sample t-test revealed that participants rated that they noticed the T appear more often over the P cues (M = 4.4) compared to NP cues (M=3.5; $t_{(23)} = 2.78$, p = .011, $d_{unb} = 0.77$)[5].

2.3.3. Discussion

The stimuli and SOAs used in Experiment 3 were the same as those used in Experiment 2. However, the dot probe task in Experiment 3 was a discrimination task with a sideways T target, instead of the localisation task and white square target used in Experiment 2. Not surprisingly, the discrimination task resulted in slower overall RTs (Figure 2.10). Similarly to Experiments 1 and 2, there was a main effect of SOA in the N2pc latencies (Figure 2.12), which mirrored the main effect of SOA found in RTs. Also in agreement with the previous experiments, targets that appeared after a 250 ms SOA elicited a larger N2pc amplitude when they appeared over a NP cue compared to a P cue (Figure 2.10).

The average N2pc latency at the 1200ms SOA in Experiment 3 (M = 195ms) was significantly longer than that in Experiment 2 (M = 188ms; $t_{(48)} = 2.38$, p = .021, $d_{unb} = 0.66$). In addition, the average RT at the 1200ms SOA in Experiment 3 (M = 520ms) was significantly longer than that in Experiment 2 (M =350ms; $t_{(46)} = 8.87$, $p < .001$, $d_{unb} = 2.5$). Despite this increase in N2pc latency and RT, no evidence of IOR was found in the RT or N2pc data. In fact, there was a RT facilitation effect in favour of the P cues at the 1200ms SOA. One interpretation of this finding is that participants were consciously shifting their attention to the P cues before the target appeared, perhaps as a strategy to get through the task faster. Le Pelley et al. (2013) originally used a 1000ms SOA in their dot probe task to rule out this interpretation being applied to the facilitation effect they observed at the 250ms SOA. Similarly to Experiments 1 and 2, Le Pelley et al. found no differences in RTs to the target as

---

[5] Please note that in subsequent experiments rating data was collected but not always reported in text. For a meta-analysis of the rating data from each experiment, please see the General Discussion in Chapter 5.

a function of predictiveness at the long SOA. They interpreted this null result as evidence that the earlier RT facilitation effect was a result of automatic, not strategic, shifts of attention towards the P cues. Therefore, the RT facilitation effect observed at the 1200 ms SOA in the current experiment could be taken as evidence that participants were strategically shifting their attention towards the P cue during the SOA.

At the end of Experiment 3 participants were asked to rate how often they noticed the sideways T target appear over each one of the cues. The rating data showed a learned predictiveness-like effect whereby participants rated that they noticed the target appear more often over the P cues than the NP cues (for a similar result see Luque, Vadillo, Gutierrez-Cobo & Le Pelley, 2016).

The N2pc effect obtained at the short SOA in Experiment 3 was in the same direction as observed in Experiments 1 and 2 (i.e., greater negativity contralateral to a target that appeared over a NP cue compared to a P cue). Similarly to the first two experiments, we were unable to obtain both the N2pc effect and the RT facilitation effect at the same SOA. However, all of the behavioural data obtained so far (RTs and ratings) suggest attentional facilitation towards the P cue. Therefore, it seems parsimonious to interpret the N2pc effect as attention towards the P cue, despite it being counterintuitive given the definition of the N2pc.

2.5. General Discussion

The experiments presented in this chapter were designed to investigate potential facilitation and inhibition effects caused by derived attention. Participants learned to categorise pairs of grating stimuli, with one grating in each pair predictive of the categorisation response and the other non-predictive. The same gratings were then used as uninformative spatial cues in a dot probe task, from which behavioural (RT and ratings) and electrophysiological (N2pc) measures of attention were obtained. Each experiment used a 250ms and 1200ms SOA to investigate the time course of attention. It was predicted that participants would be faster to respond to the target when it appeared over P cues compared to NP cues at the short SOA. Such a facilitation effect would be a replication of previous work (Le Pelley et al., 2013; Luque et al., 2017, Cobos, Vadillo, Luque, & Le Pelley, 2018), and in accordance with the Mackintosh model. At the long SOA, however, it was predicted that participants would be faster to respond to targets that appeared over NP cues compared to P cues, in accordance with the Pearce-Hall model and IOR.

In Experiment 1, the P and NP stimuli were black and white gratings that differed only in orientation. After learning which gratings were predictive of a categorisation response, participants completed an intermixed dot probe localisation and categorisation task in which they responded as fast as possible to the location of a white square that appeared randomly over either the P or NP gratings. The RT results in Experiment 1 did not replicate the previously reported RT facilitation effect at the 250 ms SOA (Le Pelley et al., 2013; Luque et al., 2017, Cobos et al., 2018). However, the EEG results showed a novel N2pc effect at the 250ms SOA. The mean amplitude of the N2pc elicited by the target was significantly smaller when the target appeared over P cues compared to NP cues (Figure 2.5). This N2pc effect was in the opposite direction to what would be expected given the definition of the N2pc. That is, the N2pc results appeared to suggest that attention was allocated towards the NP cue. However this interpretation seems problematic given that participants had to attend to the P cue in order to make the categorisation response in the intermixed categorisation and dot probe task. Therefore, it seems parsimonious to attribute the N2pc effect to an attentional shift towards the P cue.

The aim of Experiment 2 was to replicate the RT facilitation effect to support our interpretation of the N2pc effect. The previous studies that have reported a RT facilitation effect used coloured stimuli, so colour was added to the gratings to allow participants to discriminate between P and NP cues based on orientation and/or colour information. This change was sufficient to replicate the RT effect. Participants in Experiment 2 responded significantly faster to the target when it appeared over the P cues compared to the NP cues at the 250ms SOA (Figure 2.6). In contrast to Experiment 1, the N2pc mean amplitude effect at the 250ms SOA in Experiment 2 was not significant (Figure 2.9), although the effect was in the same direction.

In both Experiments 1 and 2 no RT effects were found at the 1200 ms SOA, and hence no evidence of inhibition was obtained. The aim of Experiment 3 was to test whether floor effects could be contributing to the lack of IOR at the 1200 ms SOA. To slow down responding the dot probe task was changed from a localisation task to a discrimination task in which participants had to respond to the location of the vertical line (left or right) within a sideways T target, irrespective of where the target appeared on the screen. This manipulation successfully slowed down target processing (as indexed by the target elicited N2pc latencies) and RTs. However, no evidence of inhibition was found at the 1200 ms SOA in the RT or N2pc data. In fact, an RT facilitation effect in the direction of the P cue was found at the long SOA in Experiment 3. One interpretation of this RT facilitation effect is that participants were strategically processing the cues during the cue-target interval (Le Pelley et al, 2013; Le Pelley et al, 2016). Nevertheless, it is clear from Experiment 3 that the lack of behavioural IOR was not due to floor effects in RTs.

In Experiment 3, targets that appeared after a 250 ms SOA elicited larger N2pc amplitudes when they appeared over the NP cues compared to the P cues. This result is in agreement with Experiments 1 and 2, which showed a similar N2pc effect at the short SOA. In all experiments, the N2pc effect was never found to coincide with a RT effect suggesting that these two measures of attention are decoupled for some reason. However, all the behavioural data suggest an attentional shift towards the P cues. This includes the rating data obtained at the end of Experiment 3 in which participants rated that they noticed the sideways T target appear more often over the P cues compared to the NP cues. However, it is possible that this rating result was not due to an attention bias towards the P cue, but instead reflected post-hoc rational inference (Mitchell et al., 2012). In other words, participants may have attributed higher ratings to the P cues because they made the assumption that cues that were important during the categorisation task would also be important during the dot probe task, not because they actually noticed the target appear more often over the P cue. Nevertheless, the behavioural data suggest an overall attention bias towards the P cues.

As already mentioned, it seems most likely that the novel N2pc effect found in these experiments reflects attention towards the P cue. However, the definition of the N2pc (i.e., greater negativity contralateral to an attended item) indicates that this effect reflects attention to the NP cue. As IOR can occur quite early (e.g., 200 ms; Cao, Wu, & Whang, 2017) it is possible that the targets appearing after a 250 ms SOA were already in an inhibition time window. Indeed, a reduction in N2pc amplitude similar to that found in these experiments has been reported as a marker of IOR (McDonald et al., 2009; Pierce, Crouse, & Green, 2017). Therefore, it is possible that attention was already biased towards the NP cue after the 250 ms SOA.

The next chapter presents a series of experiments that survey a broader range of SOAs with the aim of clarifying the N2pc effect reported in the current chapter. By testing multiple (and earlier) SOAs it might be possible to obtain a fine grain analysis of the time course of attention and elucidate whether the N2pc effect reflects facilitation or inhibition of the P cue.

CHAPTER 3: Multiple SOAs

Statement of Authorship (Experiments 4 and 5)

Title of Paper: Mackintosh, Pearce-Hall and time: An EEG study on Inhibition of Return

Publication Status: Published

Publication Details: Russo, S., Burns, N., & Baetu, I. (2019). Mackintosh, Pearce-Hall and time: An EEG study on Inhibition of return. *Biological Psychology, 146*.

Principal Author

Name of Principal Author (Candidate)	Salvatore Russo		
Contribution to the Paper	Design, data collection, data analysis and write up		
Overall percentage (%)	70		
Certification:	This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.		
Signature		Date	21/02/2020

Co-Author Contributions

By signing the Statement of Authorship, each author certifies that:

i. the candidate's stated contribution to the publication is accurate (as detailed above);
ii. permission is granted for the candidate in include the publication in the thesis; and
iii. the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Name of Co-Author	Nicholas Burns		
Contribution to the Paper	Statistical analyses and editing		
Signature		Date	21/02/2020

Name of Co-Author	Irina Baetu		
Contribution to the Paper	Design, data analyses and editing		
Signature		Date	21/02/2020

The time course of facilitation and IOR effects may be more complicated than what was assumed in the previous chapter. Experiments 1 through 3 used a 1200 ms SOA on the assumption that IOR occurs late. Although, IOR has been shown to be quite stable for SOAs between 300-1600 ms (Samuel & Kat, 2003), it has also been shown to occur as fast as 250 (Castel, Chasteen, Scialfa, & Pratt, 2003) and 200 ms (Cao, Wu, & Wang, 2017). Therefore, the 250 ms SOA used in the first three experiments might already be in an inhibition time window, or at a cross over point from facilitation to inhibition. To address this possibility the experiments presented in this chapter survey a wider and earlier range of SOAs. A fine-grained analysis of the time course of attention might help clarify whether the N2pc effect reported in the previous chapter was due to facilitation or inhibition of the P cue.

All experiments in the previous chapter time locked the N2pc to the onset of the target in the dot probe task. This approach is novel because N2pc studies typically lock the ERP to the onset of a visual search array that contains a lateralised target surrounded by distractors (e.g., Feldmann-Wustefeld et al., 2015). It is possible that this difference can account for the unusual direction of the N2pc effect observed in the previous experiments. That is, a probe that appears after a delay on top of an attended item might elicit a smaller N2pc because attention was already allocated to that area. Conversely, a probe that appears over an unattended item might elicit a larger N2pc because it is more difficult to process the probe in that location. This interpretation is consistent with N2pc literature that suggests that when more attention is required to process a target it produces a larger N2pc. For example, Luck, Girelli, McDermott, & Forde (1997) found that conjunction targets (i.e., targets defined by a combination of features) elicited larger N2pc amplitudes compared to single feature targets. They also found that targets crowded by a distractor elicited larger N2pc amplitudes compared to isolated targets.

In the following experiments the N2pc will be time locked to the onset of the cues in the categorisation task, as well as to the probe in the dot probe task. By locking ERPs to the onset of the cues it might be possible to gain more insight into the probe locked N2pc.[6]

3.1. Experiments 4 and 5

Experiments 4 and 5 were designed to obtain a fine grained analysis of the time course of attention across the cue display period. Attention to the cues in the dot probe task could be shifting endogenously (i.e., voluntarily), exogenously (i.e., automatically), or by a combination of both processes. In either case the shift of attention would be rapid, with endogenous shifts taking between 200-300 ms, and automatic shifts taking around 75-175 ms (Carlson, Hogendoorn, & Verstraten, 2006). Experiments 4 and 5 used six SOAs (0, 100, 200, 300, 400, and 500 ms) in the dot probe task. These SOAs should be sufficient to capture the early facilitation effect and the later inhibition effect, if it is present.

In addition to the probe locked N2pc, experiments 4 and 5 also time locked ERPs to the onset of the cues in the categorisation task. This might further clarify the nature of the probe locked N2pc (i.e., does it reflect attention to the P or NP cue). These experiments also included additional behavioural measures that could help explain what participants are attending to during the dot probe task. In Experiment 4, error data during the dot probe task was recorded. A dot probe error was classified as any trial in which the participant pressed a button (left or right) that was opposite to the location of the probe. Experiment 5 also recorded premature responses (i.e., left or right responses in the dot probe task that occurred before the probe appeared). These data could also be used to assess what participants are attending to during the dot probe task. For example, if participants are attending to the P cue,

---

[6] Using cue-locked data in the dot probe task did not seem ideal because the probe appears after different SOAs and so would interrupt different ERP components.

then they might also be inclined to respond towards the P cue prematurely, or respond towards the P cue even though the target appears over the NP cue.

For the RT results, it was predicted that participants would respond faster to targets when they appeared over the P cue compared to the NP cue after the short SOAs, and that this effect might reverse after a longer SOA. It was also expected that the probe locked N2pc results would show a similar effect. If the N2pc results from the previous chapter replicate, it is expected that after short SOAs targets will elicit a larger N2pc amplitudes when they appear over a NP cue compared to a P cue. At longer SOAs this pattern should reverse if attention shifts from the P cue to the NP cue such that targets that appear over P cues elicit larger amplitudes compared to targets over NP cues.

3.1.1. Method

Participants

Fourteen participants (8 male; mean age 22 years, age range 19 to 33 years) took part in Experiment 4 and sixteen participants (5 male; mean 22 age years, 18 age range to 29 years) took part in Experiment 5. All participants reported normal or corrected-to-normal vision and participated in two 1.5 hour sessions (on separate days). Participants gave written informed consent and were compensated with $100 (Experiment 1) and $80 (Experiment 2) gift cards.

Apparatus

All apparatus were identical to those previously described except that the electroencephalogram and an electrooculogram (EOG) were recorded using a BioSemi ActiveTwo system (BioSemi, Amsterdam, Netherlands).

Stimuli

All stimuli were identical to those used in Experiment 2 (i.e., coloured gratings and dot probe localisation task with white square target).

Procedure

Experiments 4 and 5 were identical to Experiment 2 except multiple SOAs were tested (0, 100, 200, 300, 400 and 500 ms). Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n = 7 for Experiment 4 and n = 8 for Experiment 5) and “Left-Right Predictive” (n = 7 for Experiment 4 and n = 8 for Experiment 5). In each experiment, four participants in the “Vertical-Horizontal Predictive” and four participants in the “Left-Right Predictive” condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings and green left gratings.

The experiments consisted of two testing sessions run on different days. Each session contained two phases. Phase 1 was a categorisation task and Phase 2 combined the categorisation task with a dot probe task. There were a total of 12 within-subject conditions (i.e., the target appearing over the P and NP cues at each of the six SOAs), and a total of 144 trials in each condition.

At the end of Phase 2 in the second session, participants were asked to rate how often they noticed the white square target appear over each of the cues. Each cue was shown individually in a random order and the participant entered a number on the keyboard from 1 (rarely) to 6 (often), resulting in one rating for each cue.

EEG Recording and Analysis

EEG was recorded from 32 electrodes (Fp1, Fp2, AF3, AF4, F7, F8, F3, F4, Fz, FC5, FC6, FC1, FC2, C3, C4, Cz, CP5, CP6, CP1, CP2, P7, P8, P3, P4, Pz, PO7, PO8, PO3, PO4, O1, O2, and O2) at a sampling rate of 2048 Hz. The continuous EEG was segmented into epochs ranging from 100 ms prior to the onset of the dot probe target to 250 ms (for the 200, 300, 400 and 500 ms SOAs) or 300 ms (for the 0 and 100 ms SOAs) post target onset, and baseline corrected using the 100 ms pre-stimulus interval. Trials in which the participant responded to the dot probe faster than 250 ms (for 250 ms epochs) or faster than 300 ms (for 300 ms epochs) were excluded from the EEG analysis to ensure motor responding did not contaminate the epoch of interest. A longer epoch length was used for the shorter SOAs because participants were slower to respond to the probe, allowing us to extend the epoch.

ERP waveforms were time-locked to the onset of the white square target.[7] ERPs ipsilateral to the target were computed by averaging the ERPs recorded over the right scalp when the target was on the right with ERPs recorded over the left scalp when the target was on the left. ERPs contralateral to the target were computed by averaging the ERPs recorded over the right scalp when the target was on the left with ERPs recorded over the left scalp when the target was on the right. These lateralised ERPs were computed for each of the 12 within-subjects conditions (i.e., when the target appeared over the P and NP cues at each SOA). The 50% fractional area latency of the N2pc component was calculated using a time window of 150-300 ms for the 0 and 100 ms SOAs and a time window of 100-250 ms for the 200, 300, 400 and 500 ms SOAs. Twelve separate latencies (i.e., one for each within-subjects condition) were calculated for each participant.

---

[7] We also time-locked ERPs to the onset of the grating cues during the categorisation task. For a detailed discussion of the cue onset results please see Appendix C.

Similarly to the previous experiments, significant differences in N2pc latencies were found between SOA conditions (see row three of Figure3.2). There was a significant main effect of SOA on N2pc latencies at the PO7/PO8 (Experiment 4: $F_{(2.17, 26.04)} = 47.75$, $p < .001$, $\eta^2 = .58$, MSE = 569.09; Experiment 5: $F_{(2.28, 22.83)} = 40.42$, $p < .001$, $\eta^2 = .62$, MSE = 284.14) and O1/O2 (Experiment 4: $F_{(1.73, 20.70)} = 82.03$, $p < .001$, $\eta^2 = .72$, MSE = 432.30; Experiment 5: $F_{(2.841, 29.41)} = 41.07$, $p < .001$, $\eta^2 = .62$, MSE = 211.92) electrode sites. Therefore, the mean amplitude of the N2pc component elicited by the target was measured during different 50 ms time windows centred on the median latencies for each SOA. The median latencies for each SOA measured at the PO7/PO8 and O1/O2 electrode sites can be found in Appendix A.

For each participant we calculated the mean amplitude of the N2pc component from their contralateral minus ipsilateral difference waveforms recorded at PO7/PO8 and O1/O2 electrode sites. The mean amplitudes recorded at PO7/PO8 and O1/O2 were then averaged together.

Statistical Analysis

The effects of predictiveness and SOA on the mean number of dot probe errors was compared using a two-way repeated-measures ANOVA. In Experiment 5 we also recorded premature responses to the dot probe (i.e., responses after the cues appeared but before the probe appeared)[8]. The effects of predictiveness and SOA on the mean number of premature dot probe responses was compared using a two-way repeated measures ANOVA. We also used a two-way repeated measures ANOVA to analyse the effects of predictiveness and SOA on the mean number of categorisation errors during the intermixed dot probe and categorisation task. Greenhouse-Geisser corrections were applied when Mauchly's test

---

[8] Please note that dot probe errors and premature responses were collected for multiple experiments but not always reported in text. For a meta-analysis of the dot probe errors and premature responses from each experiment, please see the General Discussion in Chapter 5.

indicated that the assumption of sphericity was violated. Where this assumption was violated, the corrected p-values are reported.

3.1.2. Results

Experiment 4

Figure 3.1 (top) shows the average of the median RTs to the target in the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on RTs was significant ($F_{(1, 13)} = 59.0$, $p < .001$, $\eta^2 = .10$, MSE = 489.65, $d_{unb} = .38$) as participants responded faster to the target when it appeared over a P cue (M = 372 ms) compared to a NP cue (M = 398 ms). There was main effect of SOA as participants were slower to respond to the target when it appeared after short SOAs compared to long SOAs ($F_{(1.68, 21.85)} = 85.98$, $p < .001$, $\eta^2 = .73$, MSE = 1425.15). There was also a significant Predictiveness×SOA interaction ($F_{(2.03, 26.45)} = 5.80$, $p = .008$, $\eta^2 = .01$, MSE = 246.01). The Predictiveness×SOA interaction was best captured by a linear trend ($F_{(1, 13)} = 8.26$, $p = .013$, $\eta^2 = .87$, MSE = 307.17) such that the difference in RT to the target over the P and NP cues increased proportionally with SOA. Paired sample t-tests revealed that participants were significantly faster to respond to the target when it appeared over a P cue compared to a NP cue at each SOA (largest $p = .004$ at the 0 ms SOA).

Figure 3.1 (middle) shows the mean number of dot probe errors. There was a main effect of Predictiveness ($F_{(1, 13)} = 6.76$, $p = .022$, $\eta^2 = .14$, MSE = 3.95, $d_{unb} = .69$) as participants made more errors when the target appeared over the NP cue (M = 1.3) compared to the P cue (M = 0.5). This meant that participants were more likely to indicate that the target appeared over the P cue when in fact it had appeared over the NP cue, rather than vice versa. There was a main effect of SOA ($F_{(2.36, 30.72)} = 5.55$, $p = .006$, $\eta^2 = .09$, MSE = 1.35) because participants made more errors when targets appeared after long SOAs compared to short SOAs. The Predictiveness×SOA interaction was not significant ($F_{(5, 65)} = .76$, $p = .58$).

Figure 3.1 (bottom) shows the mean number of categorisation errors. There was a main effect of Predictiveness ($F_{(1, 13)} = 14.09$, $p = .002$, $\eta^2 = .193$, MSE = 12.79, $d_{unb} = .76$) as participants made more categorisation errors when the target appeared over the NP cue (M = 4.6) compared to the P cue (M = 2.52). The main effect of SOA was not signficant ($F_{(5, 65)} = 1.114$, $p = .36$) nor was the Predictiveness×SOA interaction ($F_{(5, 65)} = 2.16$, $p = .07$).

A paired sample t-test revealed that participants rated that they noticed the target appear more often over the P cues (M = 4.4) than the NP cues (M = 3.5; $t_{(13)} = 2.20$, $p = .047$, $d_{unb} = .74$).

Figure 3.2 shows contralateral minus ipsilateral difference waveforms for each SOA when the target appeared over a P cue (row one) or a NP cue (row two). The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix A. Figure 3.2 (row four) shows the average of the N2pc mean amplitudes evoked by the target as a function of Predictiveness and SOA. The main effect of SOA on N2pc mean amplitudes was significant because targets elicited larger (i.e., more negative) voltages when they appeared after long SOAs compared to short SOAs ($F_{(2.21, 28.76)} = 6.37$, $p = .004$, $\eta^2 = 0.16$, MSE = 1.76). The main effect of Predictiveness was not significant ($F_{(1, 13)} = 3.82$, $p = .073$, $\eta^2 = 0.02$, MSE = 0.62, $d_{unb} =.19$). However, there was a significant Predictiveness×SOA interaction ($F_{(2.32, 30.18)} = 4.12$, $p = .021$, $\eta^2 = .10$, MSE = 1.65). The Predictiveness×SOA interaction was best captured by a cubic trend ($F_{(1, 13)} = 5.85$, $p = .031$, $\eta^2 = .88$, MSE = 1.54) such that at short SOAs targets that appeared over a NP cue elicited a larger (i.e., more negative) voltage, but at longer SOAs this pattern reversed. Paired sample t-tests revealed that targets elicited greater negative voltages when they appeared over NP cues compared to

P cues at the 100 ms ($M_{P\text{-}NP}$ = 1.20, $t_{(13)}$ = 2.533, p = .025, d_{unb} = 0.81) and 200 ms ($M_{P\text{-}NP}$ = .52, $t_{(13)}$ = 2.21, p = .046, d_{unb} = .35) SOAs. However, targets elicited greater negative voltages when they appeared over P cues compared to NP cues at the 400 ms SOA ($M_{P\text{-}NP}$ = -.87, $t_{(13)}$ = 2.48, p = .028, d_{unb} = 0.54).

Experiment 5

The analyses of N2pc data at different SOAs were not corrected for multiple comparisons. Therefore, in order to confirm these results a direct replication was run. The replication experiment yielded a very similar pattern of RT results. Figure 3.1 (top) shows the average of the median RTs to the target in the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on RTs was significant ($F_{(1, 15)}$ = 26.36, p < .001, η^2 = .072, MSE = 1288.69, d_{unb} = .23) as participants responded faster to the target when it appeared over a P cue (M = 438 ms) compared to a NP cue (M = 465 ms). There was a main effect of SOA as participants were slower to respond to the target when it appeared after short SOAs compared to long SOAs ($F_{(1.30, 19.54)}$ = 80.07, p < .001, η^2 = .71, MSE = 3248.02). There was also a significant Predictiveness×SOA interaction ($F_{(2.60, 39.02)}$ = 5.77, p = .003, η^2 = .011, MSE = 351.35). The Predictiveness×SOA interaction was best captured by a linear trend ($F_{(1, 15)}$ = 10.23, p = .006, η^2 = .71, MSE = 367.43) such that the difference in RT to the target over the P and NP cues increased proportionally with SOA. Paired sample t-tests revealed that participants were significantly faster to respond to the target when it appeared over a P cue compared to a NP cue at each SOA (largest p = .018 at the 0 ms SOA).

Figure 3.1 (middle) shows the mean number of dot probe errors. Unlike Experiment 4, the main effect of Predictiveness was not significant ($F_{(1, 15)}$ = 3.19, p = .094, d_{unb} = .35). However, the direction of the effect found in Experiment 4 was the same in Experiment 5, with participants making more errors when the target appeared over the NP cue (M = 1.6) compared to the P cue (M = 0.8). There was a main effect of SOA ($F_{(3.082, 46.23)} = 4.26$, p = .009, $\eta^2 = .09$, MSE = 2.93) because participants made more errors when targets appeared after long SOAs compared to short SOAs. The Predictiveness×SOA interaction was not significant ($F_{(5, 75)} = 1.11$, p = .36).

In Experiment 5 we also recorded premature responses to the dot probe target (i.e., left or right responses when the cues were on the screen but the target had not yet appeared). Very few premature responses were made ($M_{100} = 0$, $M_{200} = 0$, $M_{300} = .13$, $M_{400} = .53$, and $M_{500} = 1.25$). As no premature responses were possible at the 0 ms SOA, we conducted a repeated measures ANOVA on the 100 to 500 ms SOAs. This revealed a tendency for participants to respond prematurely in the direction of the P cues ($M_P = 0.55$ and $M_{NP} = 0.08$; $F_{(1, 15)} = 4.22$, p = .058, $\eta^2 = .036$, MSE = 3.0, $d_{unb} = .49$) and to make more premature responses at longer SOAs ($F_{(1.20, 18.06)} = 3.91$, p = .057, $\eta^2 = .10$, MSE = 7.71). There was also a tendency for the bias in premature responses towards the P cues to increase with SOA ($F_{(1.38, 20.71)} = 3.02$, p = .086, $\eta^2 = .06$, MSE = 14.3).

Figure 3.1 (bottom) shows the mean number of categorisation errors. Unlike Experiment 5, the main effect of Predictiveness was not significant ($F_{(1, 15)} = 1.02$, p = .329, $d_{unb} = .14$). However, the direction of the effect found in Experiment 4 was the same in Experiment 5, with participants making more categorisation errors when the target appeared over the NP cue (M = 5.1) compared to the P cue (M = 4.6). The main effect of SOA was not significant ($F_{(3.14, 47.07)} = 1.04$, p = .38), nor was the Predictiveness×SOA interaction ($F_{(5, 75)} = 1.20$, p = .32).

Unlike Experiment 4, a paired sample t-test on the rating data showed no significant difference ($t_{(15)} = 1.45$, p = .17, $d_{unb} = .46$) between the P (M = 3.4) and NP cues (M = 2.8), although the effect was in the same direction.

The N2pc effects found in Experiment 5 were very similar to those found in Experiment 4. Figure 3.2 shows contralateral minus ipsilateral difference waveforms for each SOA when the target appeared over a P cue (row one) or a NP cue (row two). The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix A. Figure 3.2 (row four) shows the average of the N2pc mean amplitudes evoked by the target as a function of Predictiveness and SOA. The main effect of SOA on N2pc mean amplitudes did not reach statistical significance ($F_{(1.459, 21.89)} = 2.39$, $p = .13$, $\eta^2 = 0.08$, MSE = 5.49). The main effect of Predictiveness was not significant ($F_{(1, 15)} = .13$, $p = .73$, $\eta^2 = 0.001$, MSE = 1.17, $d_{unb} = .03$). However, there was a significant Predictiveness×SOA interaction ($F_{(2.95, 44.30)} = 5.13$, $p = .004$, $\eta^2 = .096$, MSE = 1.59). The Predictiveness×SOA interaction was best captured by a cubic trend ($F_{(1, 15)} = 7.55$, $p = .015$, $\eta^2 = .90$, MSE = 1.23) such that at short SOAs targets that appeared over a NP cue elicited a larger (i.e., more negative) voltage, but at longer SOAs this pattern reversed. Paired sample t-tests revealed that targets elicited greater negative voltages when they appeared over NP cues compared to P cues at the 100 ms ($M_{P\text{-}NP} = 0.78$, $t_{(15)} = 2.82$, $p = .013$, $d_{unb} = 0.53$) and 200 ms ($M_{P\text{-}NP} = 1.02$, $t_{(15)} = 2.65$, $p = .018$, $d_{unb} = .44$) SOAs. However, targets elicited greater negative voltages when they appeared over P cues compared to NP cues at the 400 ms SOA ($M_{P\text{-}NP} = -.92$, $t_{(15)} = 2.35$, $p = .033$, $d_{unb} = 0.36$).

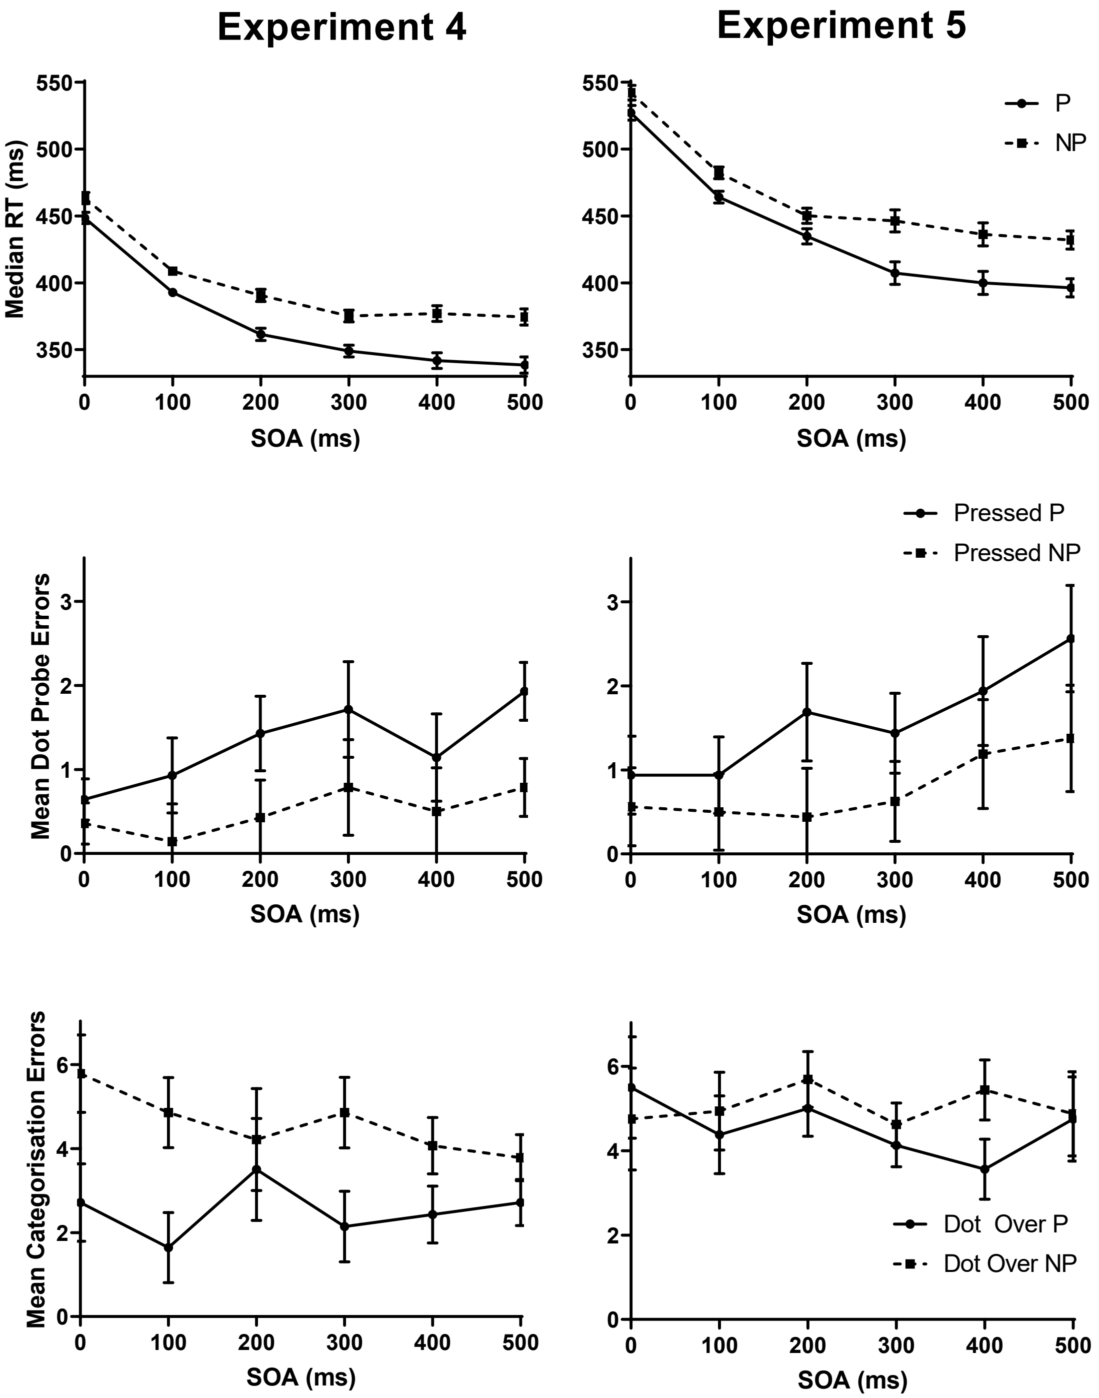

Figure 3.1: Median reaction time during the dot probe task (top), mean dot probe errors (middle), and mean categorisation errors (bottom) as a function of predictiveness and SOA in Experiment 4 (left) and Experiment 5 (right). For dot probe errors, Pressed P errors occurred when the target appeared over a NP cue and Pressed NP errors occurred when the target appeared over a P cue. Error bars represent standard error of the paired difference scores.

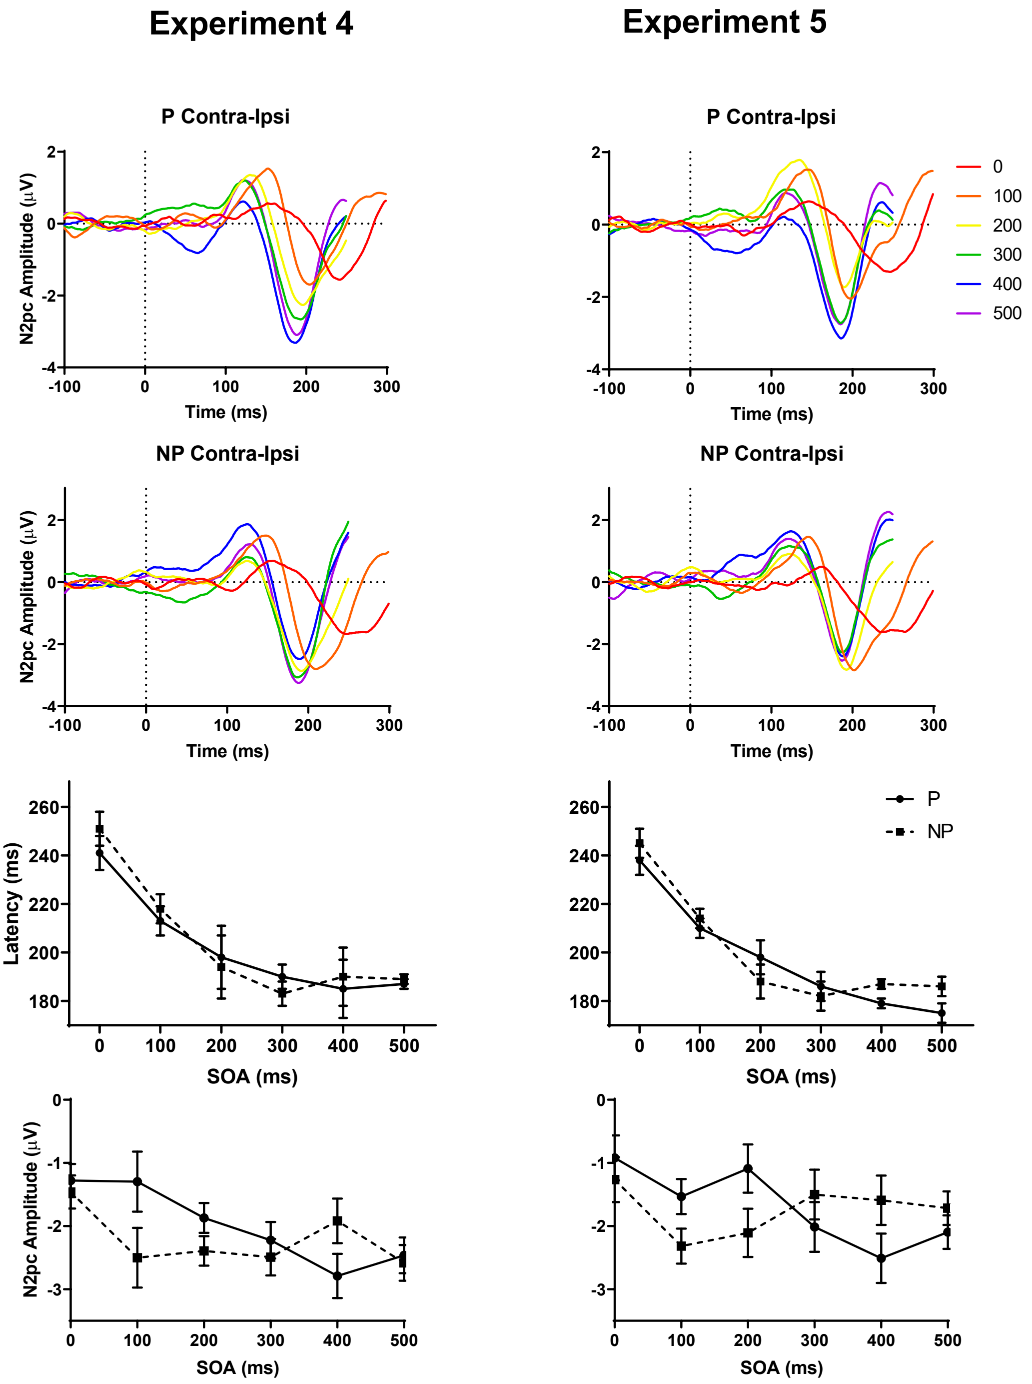

Figure 3.2: Contralateral minus ipsilateral difference waves for each SOA when targets appeared over a P cue (row one) or a NP cue (row two) for each SOA. N2pc latency (row three) and N2pc mean amplitudes (row four) as function of predictiveness and SOA are also shown. All data were averaged over PO7/PO8 and O1/O2 electrodes. Experiment 4 is on the left and Experiment 5 is on the right. Error bars represent standard error of the paired difference scores.

3.1.3. Discussion

The critical findings of Experiment 4 were replicated in Experiment 5. However, in both experiments the RT and N2pc results showed different patterns. The RT data suggested only a facilitation effect in favour of the P cues. In contrast, the N2pc amplitudes showed a pattern consistent with a shift from facilitation to inhibition of the P cues. Before we discuss the N2pc results in more detail, we present a thorough discussion of our behavioural data, in particular the extent to which this type of data reflects attentional processes.

Contrary to the predictions, RTs were significantly faster when the target appeared over a P cue compared to a NP cue at each SOA, including the longer SOAs. Moreover, the magnitude of the RT facilitation effect increased proportionally with SOA. This suggests that participants were strategically processing the cues during the SOA (Le Pelley et al., 2013; Le Pelley et al., 2016). Using a similar intermixed task, Le Pelley et al. (2013; Experiment 3) did not find evidence of strategic processing in their RT data. Instead, they found a RT facilitation effect after a 250 ms SOA and no effect after a 1000 ms SOA. One explanation for the lack of effect after the 1000 ms SOA is that strategic processing of the cues was completed by this point. Our data suggest that people engage in strategic processing much sooner than 1000 ms. In addition, the discrimination of the cues in both our study and Le Pelley et al.'s was quite easy, and our data suggest that participants could discriminate between the cues very rapidly.

Our RT data suggest that participants were strategically processing the cues (to prepare a categorisation response) before the dot probe target appeared. However, it is also possible that the strategic processing of the cues continued after the SOA had elapsed and the target was on the screen. In other words, participants may not have been responding as fast as possible to the target, but instead let it linger briefly on the screen while they continued to process the cues and consider their categorisation response. If true, this would contaminate the RT data due to the extra time taken (after target onset but before target response) for cue processing and categorisation response preparation.

To test whether participants were delaying their dot probe response in order to process the cues and prepare a categorisation response, we compared Phase 1 categorisation RTs with Phase 2 categorisation RTs (i.e., categorisation responses that followed the dot probe response). We limited our analysis to Session 2 as all participants were experts at the categorisation response by this stage. In Experiment 4 the average categorisation response accuracy in Session 2 Phase 2 was 98%, with a minimum of 95%. In Experiment 5 it was 97%, with a minimum of 93%. RTs from inaccurate categorisation responses were excluded. For Phase 2 data we collapsed across P and NP conditions and analysed the 0 ms SOA only. In this condition the probe and the cues appeared simultaneously on the screen. Therefore, participants had no delay in which to process the cues before the target appeared. Participants were not instructed to respond as fast as possible when making their categorisation response in Phase 1 or Phase 2. In Phase 2 they were explicitly told that only the dot probe response needed to be fast. If participants were responding to the dot probe as fast as possible, then cue processing should be uninitiated prior to the dot probe response in the 0 ms SOA condition. Therefore, the categorisation response made after the dot probe response should take approximately the same time as the categorisation response in Phase 1. The data can be seen in Figure 3.3.

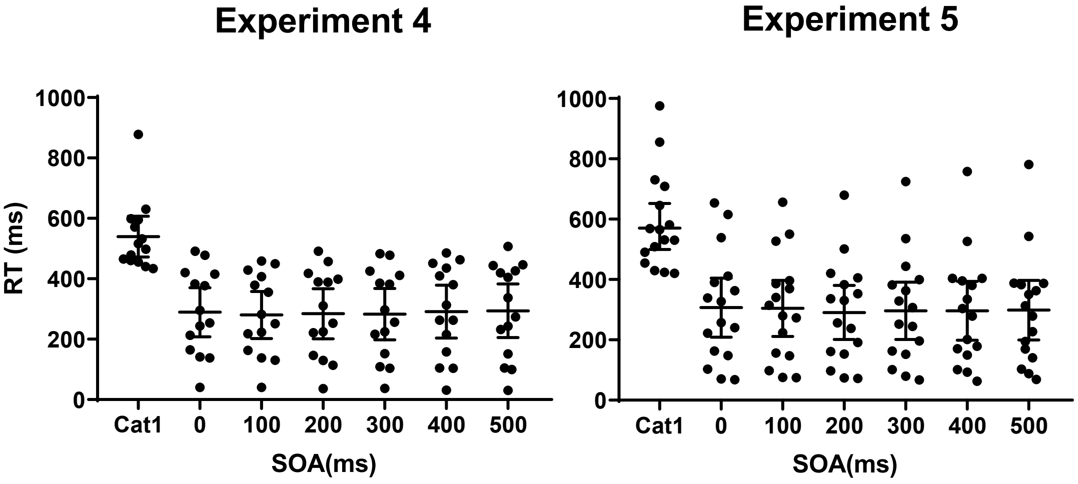

Figure 3.3: RT data from Session 2. Phase 1 (Cat1) and Phase 2 median categorisation RTs at each SOA. Categorisation RT in Phase 2 was the time between the dot probe response and the subsequent categorisation response in the dot probe task. Experiment 4 is on the left and Experiment 5 is on the right. Error bars represent 95% confidence intervals of the mean.

From Figure 3.3, it is clear that the Phase 2 categorisation response in the 0 ms SOA condition (Experiment 4: M = 300 ms; Experiment 5: M = 307 ms) was faster than the Phase 1 response (Experiment 4: M = 540 ms; Experiment 5: M = 589 ms). These differences were large (Experiment 4: $M_{diff} = 250$ ms; $t_{(13)} = 7.72$, $p < .001$, $d_{unb} = 1.8$; Experiment 5: $M_{diff} = 281$ ms; $t_{(15)} = 8.53$, $p < .001$, $d_{unb} = 1.5$). It is also clear that the RT floor of the Phase 2 categorisation responses drops towards 0 ms, with *accurate* median categorisation RTs getting as fast as 40 ms in Experiment 4 and 68 ms in Experiment 5. The delay between the dot probe response and the categorisation response is so brief that it seems likely that participants were preparing their categorisation response *before* responding to the dot probe target. That is, they did not commence responding to the target until after they had decided on a categorisation response. We will refer to this pattern of fast categorisation responding following the dot probe response as *strategic responding.*

The analysis at the 0 ms SOA challenges the assumption that dot probe RTs in the intermixed version of the task are measuring the locus of attention at each SOA. Although the Phase 2 strategic responding was not unique to the 0 ms SOA condition (Figure 3.3), it is not possible to know the extent to which dot probe RTs were affected after each SOA. It could be argued that the extra cue processing and response preparation that occurred in the 0 ms SOA condition would occur during the non-zero SOAs. If so, perhaps strategic responding confounds dot probe RTs more at short SOAs compared to long SOAs. In other words, as more time becomes available to process the cues and prepare a categorisation response before probe onset, the amount of time spent doing this after probe onset should decrease. However, it is also possible that the onset of the probe disrupts cue processing and/or categorisation response preparation. This would limit (to some extent) any processing advantage afforded to the cues during the SOAs. The Phase 2 categorisation errors (Figure 3.1, bottom) suggest that the onset of the probe did disrupt the subsequent categorisation response because participants made more categorisation errors when the dot appeared over the NP cue (the effect was in the same direction in both experiments, but only significant in Experiment 4).

Non-attentional response biases could also be affecting the dot probe RT data in the localisation version of the dot probe task. [9] That is, the left or right response to the dot probe could have been contaminated by a non-attentional response bias towards selecting a response that was congruent with the position of the P cue. The errors made in response to the dot probe target (Figure 3.1, middle) could indicate such an automatic response bias. Participants made more errors when the target appeared over a NP cue. Therefore, they were more likely to select a response that was congruent with the position of the P cue, even when the target had already appeared over the NP cue. The premature responses measured in Experiment 5 also showed a tendency for participants to prematurely select a left or right

---

[9] We would like to thank an anonymous reviewer for highlighting this possibility.

response that was congruent with the direction of the P cue. Therefore, the RT differences to the dot probe found in this task may be the result of an unknown combination of attentional processes as well as non-attentional response biases.

In summary, the strategic responding and automatic response biases found in the current intermixed categorisation and dot probe localisation task suggest that it may be problematic to infer the locus of attention from RT facilitation. The benefit of intermixing the two tasks was that we could be confident the participants were paying attention to the cues. However, the dual nature of the task meant that the P cues were task relevant on every dot probe trial. This could have encouraged participants to strategically process the cues and to strategically respond during the dot probe task. However, we would like to note that the validity of RT data as a measure of attention may depend strongly on specific task parameters. For example, strategic attention and strategic responding might be mitigated if the tasks were blocked (i.e., categorisation task followed by dot probe task). In a blocked design the cues have no task relevance during the dot probe phase so participants might be less motivated to engage in strategic processing. Furthermore, it is unlikely that participants would strategically respond in a blocked design because only one response (to the probe) is required on every dot probe trial. Therefore, RT measures in a blocked design may better reflect automatic attentional processes. In addition, a non-selective dot probe response (e.g., a single key press as used in a detection task) would mitigate the effect of non-attentional response biases.

The RT results of the current experiments may relate to the question of whether the processing of the P cues in Phase 2 of the learned predictiveness design is automatic or strategic. Consider the intermixed categorisation and dot probe task as analogous to Phase 2 of the learned predictiveness design, where participants have to attend to previously P and NP stimuli, but all stimuli now equally predict new outcomes. By keeping the previously P cues task relevant using an intermixed design we found that participants strategically processed the P cues. Similarly, in Phase 2 of the learned predictiveness design the P cues are still task relevant because they can be used to predict Phase 2 outcomes. Therefore, perhaps participants strategically process the P cues in Phase 2. It would be difficult to show that attention to cues is purely automatic when the cues are task relevant. However, it is possible that the initial shift of attention towards the P cue is automatic. Evidence for an automatic shift of attention towards P cues can be found in Le Pelley et al.'s (2013) Experiment 1, which used a non-intermixed detection dot probe task. The authors found a RT benefit in favour of the P cue after a 350 ms SOA that was perhaps less contaminated by strategic attention and non-attentional response biases. After learning to attend to P cues in Phase 1 it may be somewhat automatic for participants to continue this behaviour in Phase 2. This might bias participants to continue to strategically process the P cues during Phase 2. As using the P cues to complete Phase 2 is reinforced (because P cues predict Phase 2 outcomes) participants might continue to ignore the now task relevant, but also redundant, NP cues.

Our rating data showed a bias towards P cues (for similar results see Luque, Vadillo, Gutierrez-Cobo, & Le Pelley, 2016) as participants rated that they noticed the target appear more often over P cues compared to NP cues (the effect was in the same direction in both experiments, but only significant in Experiment 4). However, it is unknown whether this occurred due to biased learning caused by biased attention, or for some other non-attentional reason, such as rational inference. For example, it is possible that participants inferred that cues that were important in the first phase of learning are likely to be important in the second phase. This kind of inference would explain the observed results without assuming the contribution of attentional processes.

We have argued that the behavioural data collected in our experiments may not be the most sensitive measure of attention. In contrast, the immediate brain responses measured via the N2pc might provide a measure of attention (automatic and/or controlled) that is less contaminated by non-attentional processes. The target-locked N2pc data showed an interaction between predictiveness and SOA that implied changes in attention towards the cues with time (Figure 3.2). Targets that appeared after a 100 ms and 200 ms SOA elicited a more negative N2pc when they appeared over a NP cue compared to a P cue. We interpret this larger component as more focused attention being required to process the target over a NP cue. That is, at these early SOAs it was more difficult for participants to process the target when it appeared over a NP cue because attention was oriented towards the P cue. Conversely, targets that appeared after a 400 ms SOA elicited a more negative N2pc when they appeared over a P cue compared to a NP cue. Therefore, it seems that at this later SOA it was more difficult for participants to process the target when it appeared over a P cue compared to a NP cue.

The interaction between predictiveness and SOA on N2pc amplitudes supports our prediction that attention between the P and NP cues changes with time in a manner consistent with IOR (Klein, 2000). The initial direction of attention towards the P cue seems logical. Participants were conditioned to attend to the P cue during the categorisation task and a categorisation response was required after every dot probe response. After processing the P cue, the N2pc data suggest that it was inhibited. The early N2pc facilitation effect could be measuring a combination of automatic and strategic attention. The subsequent inhibition might simply be a consequence of attention moving with time, such that there is a processing cost for returning attention back to areas that have just been attended (Rayner, Juhasz, Ashby, & Clifton Jr., 2003).

We would like to note that it is possible to interpret our N2pc data in the opposite direction (i.e., greater negativity elicited by the dot probe over the NP cue means more attention to the NP cue). The N2pc is defined as a greater negativity contralateral to an attended item, so at first glance it may appear that our interpretation is incorrect. However, we believe that the definition of the N2pc is consistent with our interpretation when it is considered that our N2pc was time-locked to the onset of a target that appeared over attended cues. If attention was first paid to the P cue, then a dot probe appearing over the NP cue at short SOAs would be more difficult to process, which could cause a larger N2pc. There is evidence in the N2pc literature that larger amplitudes correspond with greater processing difficulty. For example, Luck, Girelli, McDermott, and Ford (1997) found that conjunction targets (i.e., targets that are defined by a combination of features) elicited larger N2pc amplitudes than single feature targets. Targets crowded by a distractor also elicited larger N2pc amplitudes compared to isolated targets. Therefore, when more focused attention was required to process a target it elicited a larger N2pc (Luck et al., 1997). This is in line with our interpretation of more focused attention being required to process a target that appeared over a NP or P cue at short and long SOAs, respectively.

It would also be important to investigate whether the facilitation and inhibition time windows can be shifted and/or enlarged. At present these time points are very rapid, challenging the relevance they would have for learning. It might be possible to change these time points by changing the nature of the discrimination between the P and NP cues. It seems logical that an easy feature-based discrimination would occur sooner, and be completed faster, compared to a more difficult discrimination based on a conjunction of features. If so, perhaps the facilitation and inhibition time windows measured during a conjunction discrimination may be delayed and/or extended compared to those found using a feature discrimination.

3.3. Experiment 6

Experiments 4 and 5 showed that probes that appeared after short SOAs (100 and 200 ms) elicited larger N2pc amplitudes when they appeared over NP cues compared to probes that appeared over P cues. In addition, probes that appeared after a longer SOA (400 ms) elicited a larger N2pc amplitude when they appeared over a P cue compared to a NP cue. The results also showed that the effect at the early SOAs was larger in magnitude than the later reversal. Once again, there was ambiguity regarding how to interpret the direction of the N2pc effect. However, all of the behavioural data indicated that attention was directed towards the P cue. RTs to the probe were faster when it appeared over a P cue than over the NP cue. Participants also made more dot probe errors and premature responses in the direction that would be expected if they were attending to the P cue, and they rated noticing the target more often over the P cues. Therefore, the interpretation that best suits the probe locked N2pc is that more attention (hence a larger N2pc) was required to process the probe over the NP cue after short SOAs because attention was initially deployed to the P cue. According to this interpretation, the smaller N2pc amplitude elicited by the probe when it appeared over P cues reflects facilitated detection of the probe. At the longer SOA, however, more attention was needed to process the probe over the P cue because the P cue was inhibited after initial processing.

Figure 3.4 summarises the counter-intuitive interpretation offered above. Figure 3.4A shows how attention might change from facilitation to inhibition over time. It also shows hypothetical N2pc waveforms time-locked to a probe that appears over an attended item after different SOAs. Figure 3.4B shows hypothetical N2pc waveforms time-locked to a probe that appears after different SOAs over an attended (left) or unattended (right) item. Probes that appear over an unattended item after short (100 and 200 ms) SOAs elicit larger N2pc amplitudes compared to probes that appear over an attended item (yellow and orange solid lines). Probes that appear over an attended item after a longer (400 ms) SOA elicit larger N2pc amplitudes compared to probes that appear over an unattended item (blue solid line).

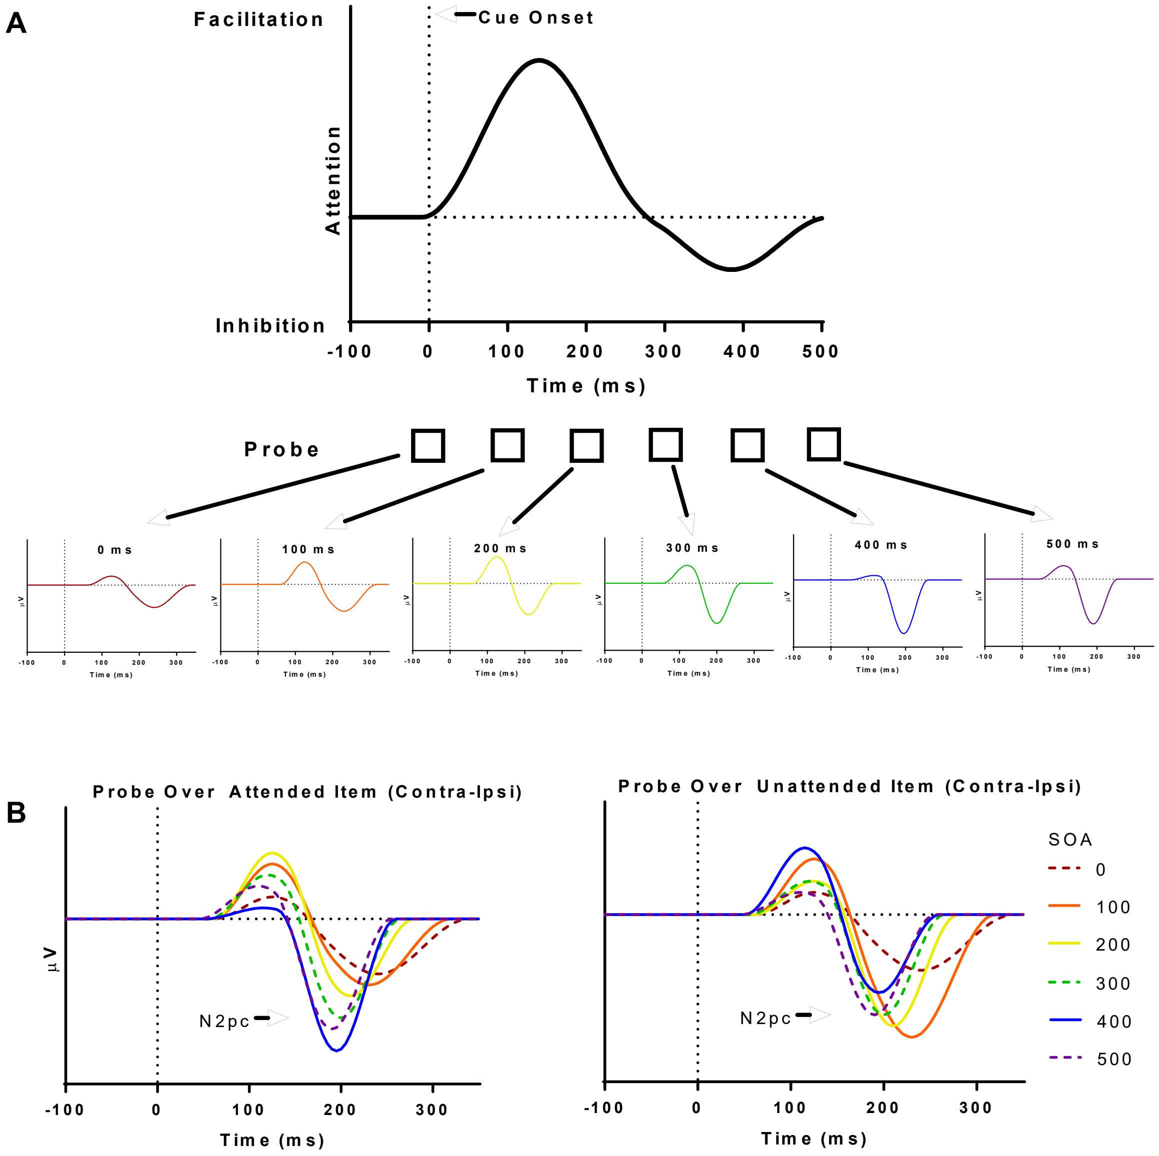

Figure 3.4: (A) Hypothetical change in attention as a function of time and hypothetical contra-ipsi N2pc waveforms time-locked to a probe that appears over an initially attended item after various SOAs. (B) Hypothetical contra-ipsi N2pc waveforms to a probe that appears over an attended (left) or unattended (right) item after different SOAs. The solid lines are N2pc waveforms that have different mean amplitudes when the probe appears over an

attended or unattended item. There is no difference in mean amplitudes between the dotted waveforms.

The aim of Experiment 6 was twofold. One aim was to continue to investigate the counterintuitive idea that probes that appear over attended items will elicit smaller N2pc amplitudes compared to probes that appear over unattended items. The second aim was to test whether the facilitation and inhibition time windows reported in Experiments 4 and 5 could be extended. It could be argued that the very rapid changes in attention observed in the N2pc data of the previous experiments would have little relevance to learning and to the Mackintosh (1975) and Pearce-Hall (1980) models. Demonstrating that facilitation and inhibition can be shifted in time depending on the task parameters might be important to make the case that these attention shifts are relevant for learning.

Participants in Experiment 6 completed the same categorisation and dot probe tasks as Experiments 4 and 5, but with two important differences. Firstly, the grating stimuli flickered at different frequencies (15 or 17 Hz) so that we could utilise the frequency tagging technique (for a reviews, see Norcia, Appelbaum, Ales, Cottereau & Rossion, 2015; Wieser, Miskovic, & Keil, 2016) to obtain another electrophysiological measure of attention. In the previous experiments we failed to obtain cue-locked N2pc results (i.e., locked to the onset of the coloured gratings), which could have aided in interpreting our probe-locked N2pc. For example, if during the categorisation task we found a cue-locked N2pc that was in the direction of the P cue (i.e., greater negativity contralateral to the P cue) then it stands to reason that attention during the dot probe task was also biased towards the P cue. One reason for the null cue-onset result could be that N2pc studies typically present stimuli for a brief duration (e.g., < 200ms; Feldmann-Wustefeld et al., 2015; McDonald et al., 2009; Mazza, Pagano, & Caramazza, 2013; Brisson. Robitaille, & Jolicoeur, 2007; Robitaille & Jolicoeur, 2006), whereas our stimuli were on the screen until a categorisation response was made.

Therefore, we decided to employ the frequency tagging technique, which is perhaps better suited to measuring attention to stimuli that are presented continuously (Anderson, Muller, & Hillyard, 2012).

When presented with a visual stimulus that flickers at a certain rate the brain will produce a steady-state visual evoked potential (SSVEP) with the same fundamental frequency as the oscillating stimulus. Importantly, when presented with two visual stimuli that flicker at different rates, the amplitude of the SSVEP corresponding to the attended stimulus will be enhanced (Müller et al., 1998). Recently, we used this technique to investigate learning driven changes in attention for the first time (Gwinn, Russo, Baetu, Nicholls, Griffiths, in preparation). To help break the deadlock of the probe-locked N2pc interpretation, we measured SSVEPs during the categorisation task in Phase 1. If participants were attending to the P cues then we expected amplitudes in the frequency spectrum to be larger at the P frequency compared to the NP frequency.

The second change introduced in the current study was that participants had to complete a conjunction discrimination to make the correct categorisation response. That is, participants were required to attend to the orientation of a specific coloured grating while ignoring another grating of a different colour. For example, some participants had to learn to attend to blue gratings and to ignore green ones, and that vertical blue was associated with an ‘up’ response and horizontal blue with a ‘down’ response (see Table 3.1). In Experiments 4 and 5, which used a feature-based discrimination, we found that probes elicited a larger N2pc when they appeared over NP cues compared to P cues at SOAs of 100 and 200 ms (we refer to this as the *early time window*), and a smaller reverse effect at a 400 ms SOA (we refer to this as the *late time window*). In the current experiment we expected to find a similar pattern of results. However, we predicted that the conjunction discrimination would influence the time windows measured. This is because the conjunction discrimination would require attention to remain on the P cue for longer (to process the grating orientation), compared to the previous feature-based discrimination (which did not require processing of the grating orientation). An extension of the early N2pc time window would show that the attentional effects measured by the N2pc are flexible and depend on how much scrutiny the cues require to be processed. In addition, an extension of the early time window would provide evidence that the direction of the N2pc effect during this window (i.e., larger N2pc amplitude for probes over NP cues compared to P cues) reflects attention being allocated to the P cue. We also expected that the N2pc reversal in the late time window would be delayed as a consequence of the extended early time window.

Table 3.1

Design of the Categorisation Task

Stimulus Pair		Correct Response
Blue Predictive	Green Predictive	
Blue-Horizontal & Green-Right	Green-Horizontal & Blue-Right	R1
Blue-Horizontal & Green-Left	Green-Horizontal & Blue-Left	R1
Blue-Vertical & Green-Right	Green-Vertical & Blue-Right	R2
Blue-Vertical & Green-Left	Green-Vertical & Blue-Left	R2

Note. Horizontal and Vertical refer to gratings with orientations of 4° and 94°, respectively. Left and Right refer to gratings with orientations of 45° and 135°, respectively. R1 and R2 refer to the correct categorisation response (the up and down buttons on a button box, counterbalanced). Horizontal and vertical gratings were always predictive and left and right gratings were always non predictive.

For the RT results, we expected to replicate Experiments 4 and 5. That is, we expected to find faster responding to the probe when it appeared over a P cue compared to an

NP cue at each SOA, and for the magnitude of the effect to increase proportionally with SOA. Similarly to experiments 4 and 5, we expected participants to delay their dot probe responses so that they could process the cues and prepare a categorisation response.

3.3.1. Method

Participants

Twenty University of Adelaide students (5 male; mean age 21 years, age range 18 to 29 years) participated in the experiment. [10] One participant did not meet the Phase 1 learning criterion and was excluded. All participants reported normal or corrected-to-normal vision. All participants reported no history of epilepsy. Participants gave written informed consent and were compensated with credit for an undergraduate psychology course.

Apparatus

The apparatus were identical to Experiments 4 and 5 except that this experiment was run on MATLAB 2017b (Mathworks, Natick, MA, USA) using the Psychophysics Toolbox extensions (Brainard, 1997; Pelli, 1997; Kleiner, Brainard, & Pelli, 2007) to control the presentation of visual stimuli on a 23.5-in computer monitor (Eizo FORIS FG2421, 1920 × 1080 pixels, vertical refresh rate = 120 Hz).

Stimuli

The cues were blue or green sinusoidal gratings of equal luminance (spatial frequency = 0.05, phase = 0) that were spatially modulated in contrast by a Gaussian envelope (standard deviation = 30) and had orientations of 4° (near-vertical), 45° (tilted to the left), 94° (near-

---

[10] In Experiments 4 and 5 we had fewer participants complete more trials. Specifically, in Experiment 4 we had 14 participants complete a total of 144 trials in each of the 12 within subject conditions (i.e., total of 2016 trials for each condition). We estimated that we would need to average a similar number of trials for each condition to replicate the N2pc effects. In the current study, we aimed to test 20 participants, but could only analyse the data from 19, who completed 96 trials in each of the 14 within subject conditions (i.e., total of 1824 trials in each condition, and this number is fairly close to the number of trials in our previous study).

horizontal), or 135° (tilted to the right). Two gratings were presented on the horizontal midline of the screen and centred 5.5° (viewing distance 80cm) to the left and right of a central fixation cross. One grating flickered at 15.004 Hz and the other at 17.147 Hz (henceforth 15 Hz and 17 Hz). The probe in the dot probe task was a white square with sides subtending 0.8° that appeared superimposed over the centre of one of the gratings. The screen background was mid-grey.

Procedure

Table 3.1 shows the cue compounds and their corresponding correct categorisation responses. Participants were assigned to the conditions “Blue Predictive” (n=10) or “Green Predictive” (n=10). Half of the participants learnt that the horizontal gratings predicted an “up” response and the vertical gratings predicted a “down” response. The other half of the participants learnt the reverse contingencies. Six participants in the “Blue Predictive” condition and six participants in the “Green Predictive” condition were shown P gratings that flickered at 15 Hz and NP gratings that flickered at 17 Hz. The remaining participants were shown P gratings that flickered at 17 Hz and NP gratings that flickered at 15 Hz.

Phase 1 and 2 in the current experiment were the same as in Experiments 4 and 5 with the following exceptions. During Phase 1, participants had only 2.4 seconds to make a categorisation response. If participants did not respond within 2.4 seconds the gratings disappeared from the screen and the words “No Response” replaced the fixation cross for 1000 ms. In Phase 2, The SOA between the onset of the images and of the white square varied between 0 and 600 ms. The extra 600 ms SOA was added in anticipation of an extended facilitation time window (and hence a delayed inhibition time window) due to the conjunction discrimination. Therefore, the dot probe task was split into blocks of 112 trials. Each block consisted of the four possible combinations of stimulus pairs (Table 3.1), two cue locations (P or NP cue on the left or right), seven SOAs for the appearance of the dot probe (0, 100, 200, 300, 400, 500 and 600 ms), and two probe locations (left or right). The 112 trials within a block were presented in random order. Participants completed a total of twelve blocks, with a self-paced break between each block. There were fourteen within-subjects conditions (the probe appearing over the P and NP cues at each of the seven SOAs), and a total of 96 trials in each condition.

EEG Recording and Analysis

The data were analysed offline using EEGLAB (Delorme & Makeig, 2004), ERPLAB (Lopez-Calderon & Luck, 2014) and Letswave 6 (letswave.org). For the frequency tagging analyses the data were filtered using a fourth order zero-phase Butterworth band-pass filter, with cutoff values of .1-120 Hz. A Fast Fourier Transform (FFT) multi-notch filter with a width of .5 Hz was also applied to remove electrical noise at three harmonics of 50 Hz. The data were segmented to include 200 ms before and after the beginning of stimulation before referencing to the average of all 32 scalp electrodes. Trials were then re-segmented to begin 467 ms after the onset of the gratings, a length that fits within the screen refresh rate of 120 Hz and the cycles of 15. Hz and 17 Hz. This was done to reduce onset responses in the EEG signal at the beginning of each trial. The segment duration was 933 ms, a length which also fits with the screen refresh rate and stimulus presentation rates. Trials were then concatenated within each condition[11]. The EEG data were then transformed into separate frequency-domain amplitude spectra by computing an FFT for each participant, condition and channel. Recordings were analysed using channels O1/O2, Oz and PO3/PO4. To account for differences in baseline noise across participants as well as across the frequency spectrum, we

---

[11] Participants completed 160 categorisation trials each. However, for one participant we did not record EEG data for their first 48 trials. As each participant was required to have the same number of trials for concatenation, we deleted the first 48 categorisation trials for each remaining participant. This left 112 categorisation trials to be concatenated for each participant.

performed a baseline subtraction in which the average of the twenty surrounding bins, excluding the immediately adjacent bins and the local maximum and minimum amplitude bins, was subtracted from the bin of interest (Gwinn & Jiang, 2019; Gwinn, Matera, O'Neil, & Webster, 2018).

The ERP analyses were the same as in Experiments 4 and 5, with the EEG segmented into epochs ranging from 100 ms prior to the onset dot probe to 300 ms (for the 200, 300, 400, 500 and 600 ms SOAs) or 350 ms (for the 0 and 100 ms SOAs) post probe onset.

The N2pc component was measured from the difference waveforms (contralateral minus ipsilateral) at the PO7/PO8, P7/P8 and O1/O2 electrode sites (see Brisson & Jolicoeur, 2007, and Brisson, Robitaille & Jolicoeur, 2007, who also averaged across these sites)[12]. For the categorisation task, the 50% fractional area latency of the N2pc component was calculated using a time window of 300-550 ms. The latency was searched for in this range because visual inspection of the grand averaged ERP waveforms at contralateral and ipsilateral electrodes showed that the N1 component was completed by 300 ms (see Figure 3.6). The mean amplitude of the N2pc component was measured during a 50 ms time window centred on the median latency of each electrode pair. The median latencies measured during the categorisation task at electrodes PO7/PO8, P7/P8, and O1/O2 can be found in Appendix B.

For the dot probe task, the 50% fractional area latency of the N2pc component was calculated using a time window of 150-300 ms for the 0 and 100 ms SOAs and a time window of 100-250 ms for the 200, 300, 400 and 500 ms SOAs. Fourteen separate latencies (i.e., one for each within-subjects condition) were calculated for each participant. The median

---

[12] We averaged over an additional pair of electrodes to help reduce noise that was present in the N2pc component. The N2pc component may have been noisier due to the fact that this experiment had more participants completing fewer trials compared to Experiments 4 and 5. However, the general pattern of results was unaffected by the inclusion of the additional electrode pair.

N2pc latency for each within-subject condition was calculated. As there was a significant main effect of SOA on N2pc latencies (see bottom row of Figure 3.7) at the PO7/PO8 ($F_{(1.36, 17.71)} = 22.33$, $p < .001$, $\eta^2 = .52$, MSE = 2658.14), P7/P8 ($F_{(1.34, 16.07)} = 16.57$, $p < .001$, $\eta^2 = .47$, MSE = 3185.45) and O1/O2 ($F_{(1.64, 22.91)} = 20.81$, $p < .001$, $\eta^2 = .44$, MSE = 2119.19) electrode sites, the mean amplitude of the N2pc component elicited by the probe was measured during different 50 ms time windows centred on the median latencies for each SOA. The median latencies for each SOA measured at the PO7/PO8, P7/P8 and O1/O2 electrode sites can be found in Appendix B.

For each participant we calculated the mean amplitude of the N2pc component from their contralateral minus ipsilateral difference waveforms recorded at PO7/PO8, P7/P8 and O1/O2 electrode sites. These mean amplitudes were then averaged together.

Statistical Analysis

The statistical analyses were the same as in Experiments 4 and 5 except that we also analysed SSVEP amplitudes. The amplitudes were entered into a mixed ANOVA with a within subjects factor of Predictiveness and a between subjects factor of Frequency Condition.

3.3.2. Results

Behavioural Results

Figure 3.5 (top) shows the average of the median RTs to the probe in the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on RTs was significant ($F_{(1, 18)} = 27.0$, $p < .001$, $\eta^2 = .13$, MSE = 5925.69, $d_{unb} = .42$) as participants responded faster to the probe when it appeared over a P cue (M = 516 ms) compared to a NP cue (M = 565 ms). There was a main effect of SOA as participants were slower to respond to the probe when it appeared after short SOAs compared to long SOAs ($F_{(2.10, 37.86)} = 132.88$, p

$< .001$, $\eta^2 = .64$, MSE = 2938.97). There was also a significant Predictiveness×SOA interaction ($F_{(2.94, 52.88)} = 3.66$, $p = .019$, $\eta^2 = .01$, MSE = 1123.46). The Predictiveness×SOA interaction was captured by a linear trend ($F_{(1, 18)} = 7.12$, $p = .016$, $\eta^2 = .69$, MSE = 1175.36) such that the difference in RT to the probe over the P and NP cues increased proportionally with SOA. Paired sample t-tests revealed that participants were significantly faster to respond to the probe when it appeared over a P cue compared to a NP cue at each SOA (largest $p = .013$ at the 100 SOA).

Figure 3.5 (bottom) shows the participants; categorisation RTs in Phase 1 and at each SOA during phase 2 (i.e., the time between the dot probe response and the subsequent categorisation response). We compared participants' categorisation RTs in Phase 1 (M = 860 ms) with those in the 0 ms SOA condition in Phase 2 (M = 388 ms) using a paired samples t-test. Phase 1 categorisation responses were significantly slower than those measured in the 0 ms SOA condition in Phase 2 ($M_{diff} = 473$ ms, $t_{(18)} = 10.9$, $p < .001$, $d_{unb} = 3.2$).

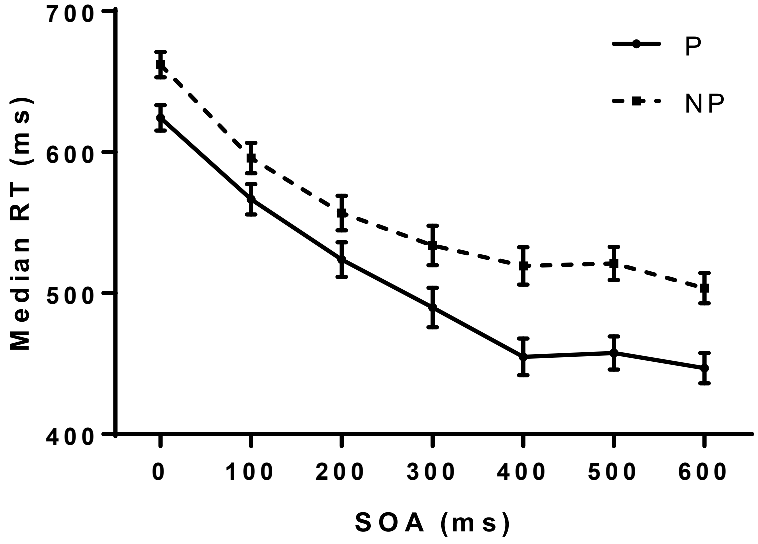

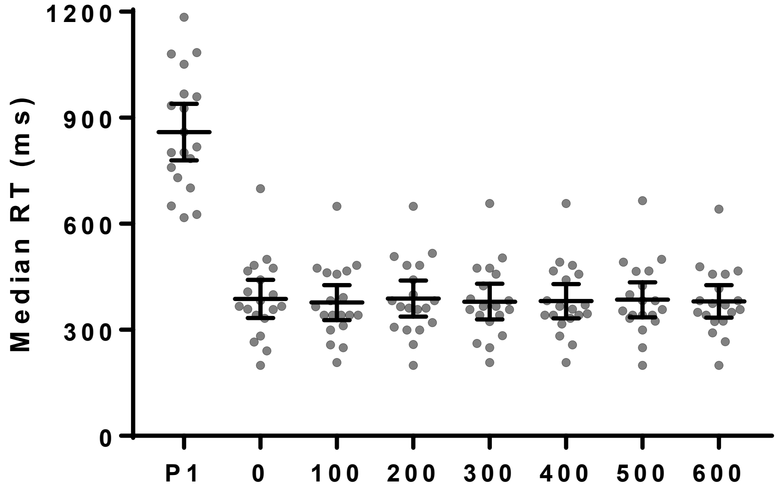

Figure 3.5: Median reaction time to the dot probe (top) as a function of Predictiveness and SOA. Error bars represent standard error of the paired difference scores. Median categorisation reaction time (bottom) during Phase 1 (P1) and after each SOA (0-600 ms) in Phase 2. Categorisation RT in Phase 2 was the time between the dot probe response and the subsequent categorisation response in the dot probe task. Error bars represent 95% confidence intervals of the mean.

N2pc Results

Figure 3.6 shows the grand average ERP waveforms at contralateral and ipsilateral (to the P cue) electrodes time-locked to the onset of the cues during the categorisation task. We compared participants' N2pc difference wave amplitudes (averaged over the PO7/PO8, P7/P8, and O1/O2 electrodes) to zero using a one-sample t test. The results showed that participants amplitudes were significantly below 0 (i.e., more negative contralateral to the P cue; M = -.84 μV, $t_{(18)} = 2.75$, p = .013, $d_{unb} = 0.60$).

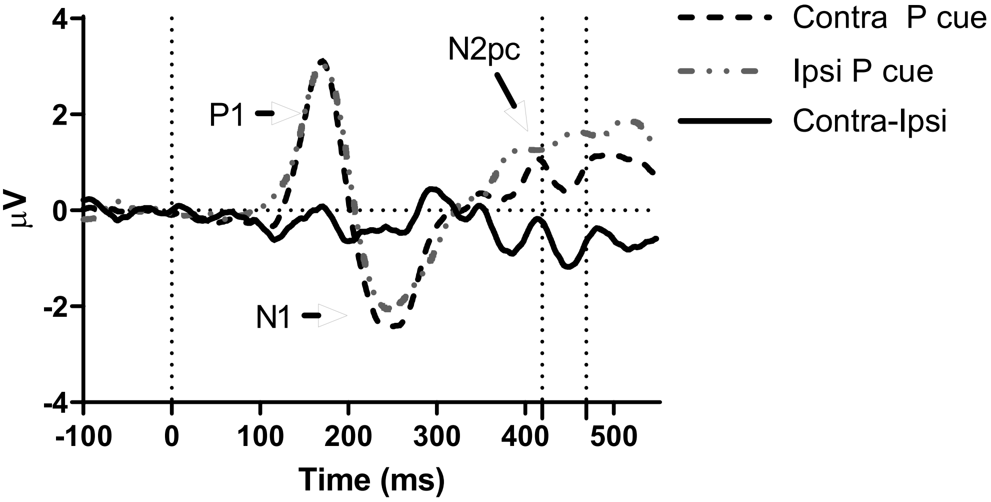

Figure 3.6: Grand average ERP waveforms at contralateral versus ipsilateral electrode sites (averaged over PO7/PO8, P7/P8, and O1/O2) time-locked to the onset of the cues during the categorisation task. The vertical dotted lines represent the 50 ms time window used to measure the N2pc amplitude. The contralateral minus ipsilateral waveform is also shown.

Figure 3.7 (top row) shows contralateral minus ipsilateral difference waveforms for each SOA when the probe appeared over a P cue (left) or a NP cue (right) during the dot probe task. The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix B. Figure 3.7 (bottom row) shows the average of the N2pc mean amplitudes evoked by the probe as a function of Predictiveness and SOA. The main effect of SOA on

N2pc mean amplitudes was significant because probes elicited larger (i.e., more negative) voltages when they appeared after long SOAs compared to short SOAs ($F_{(2.18, 39.24)} = 13.14$, $p < .001$, $\eta^2 = 0.27$, $MSE = 5.29$). The main effect of Predictiveness was not significant ($F_{(1, 18)} = .428$, $p = .52$, $\eta^2 = 0.002$, $MSE = 2.37$, $d_{unb} = .22$). However, there was a significant Predictiveness×SOA interaction ($F_{(6,108)} = 4.37$, $p = .001$, $\eta^2 = 0.054$, $MSE = 1.16$). The Predictiveness×SOA interaction was captured by a cubic trend ($F_{(1, 18)} = 5.27$, $p = .034$, $\eta^2 = 0.83$, $MSE = 1.54$) such that at short SOAs (100-300 ms) probes that appeared over a NP cue elicited a larger (i.e., more negative) voltage, but at longer SOAs (500-600 ms) this pattern reversed. Paired sample t-tests revealed that probes elicited greater negative voltages when they appeared over NP cues compared to P cues at the 100 ms ($M_{P\text{-}NP} = 0.81$ µV, $t_{(18)} = 3.56$, $p = .002$, $d_{unb} = 0.61$) and 300 ms ($M_{P\text{-}NP} = 1.2$ µV, $t_{(18)} = 2.77$, $p = .013$, $d_{unb} = .67$) SOAs. The effect at the 200 ms SOA was in the same direction but did not reach statistical significance ($M_{P\text{-}NP} = .57$ µV, $t_{(18)} = 1.43$, $p = .17$, $d_{unb} = .33$). Probes elicited greater negative voltages when they appeared over P cues compared to NP cues at the 500 ms ($M_{P\text{-}NP} = -.66$ µV, $t_{(18)} = 2.01$, $p = .06$, $d_{unb} = 0.43$) and 600 ms ($M_{P\text{-}NP} = -.51$ µV, $t_{(18)} = 1.13$, $p = .27$, $d_{unb} = 0.31$) SOAs, however the simple effects were not significant.

SSVEP Results

Figure 3.8 shows participants amplitudes in the frequency domain. For eleven participants the P stimuli flicker at 15 Hz and the NP stimuli flicker at 17 Hz. The remaining eight participants had the reverse contingency. The amplitudes were entered into a mixed ANOVA with a within-subjects factor of Predictiveness and a between subjects factor of Frequency Condition. There was a main effect of Predictiveness ($F_{(1, 17)} = 4.7$, $p = .044$) as amplitudes at the P frequency were higher than those at the NP frequency ($M_{P\text{-}NP} = .12$ µV, $d_{unb} = 0.35$). The Predictiveness×Frequency interaction was not significant ($F_{(1, 17)} = .01$, $p = .52$), nor was the between subjects effect of Frequency Condition ($F_{(1, 17)} = 2.2$, $p = .15$). We also collapsed across the two frequency conditions and compared the amplitudes of the P and NP stimuli using a Wilcoxon Signed-Ranks Test. A non-parametric test was used because the P-NP difference scores were not normally distributed. The Wilcoxon Signed-Ranks Test revealed that amplitudes at the P frequency were significantly higher than those at the NP frequency (Z = -2.1, p = .04, d_{unb} = 0.35).

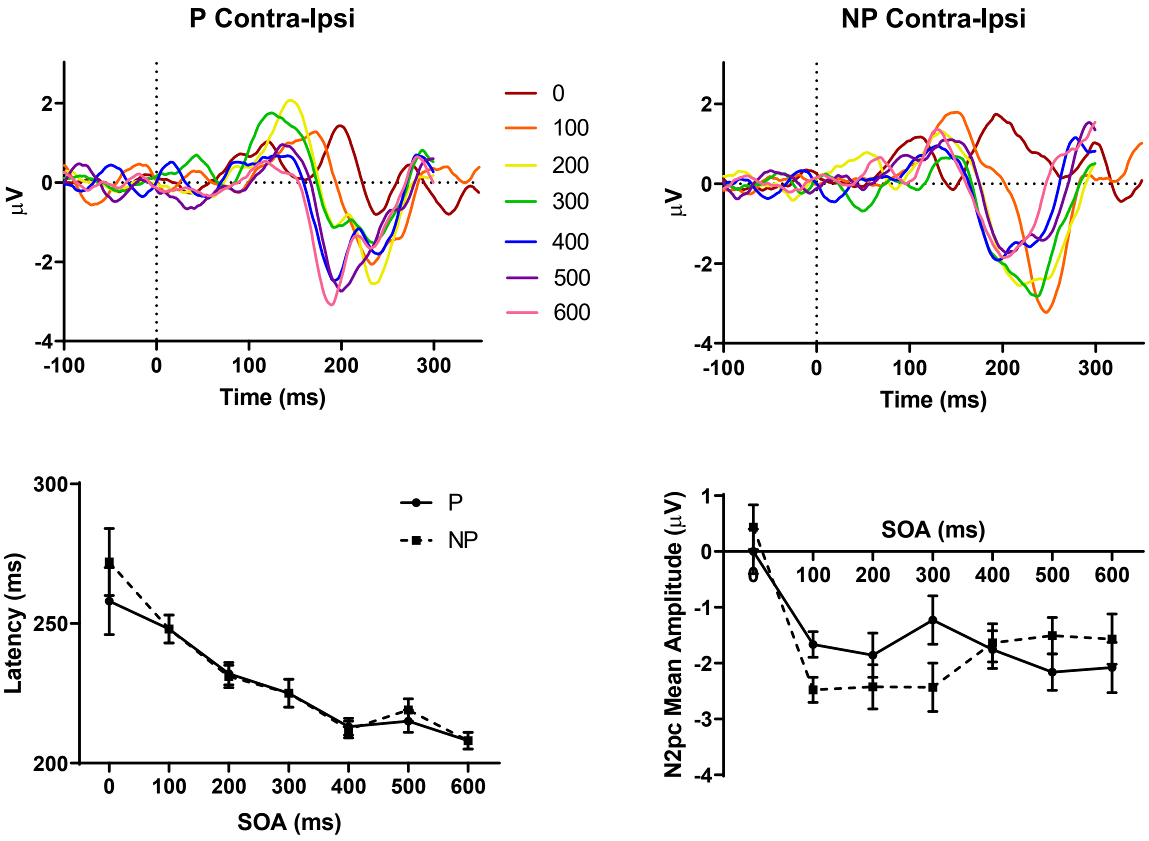

Figure 3.7: The top row shows contralateral minus ipsilateral difference waves for each SOA when the probe appeared over a P cue (left) or a NP cue (right). For individual P vs. NP waveforms at each SOA please Appendix B. The bottom row shows median N2pc latencies (left) and N2pc mean amplitudes (right) as function of predictiveness and SOA. All data were averaged over PO7/PO8, P7/P8, and O1/O2 electrodes. Error bars represent standard error of the paired difference scores.

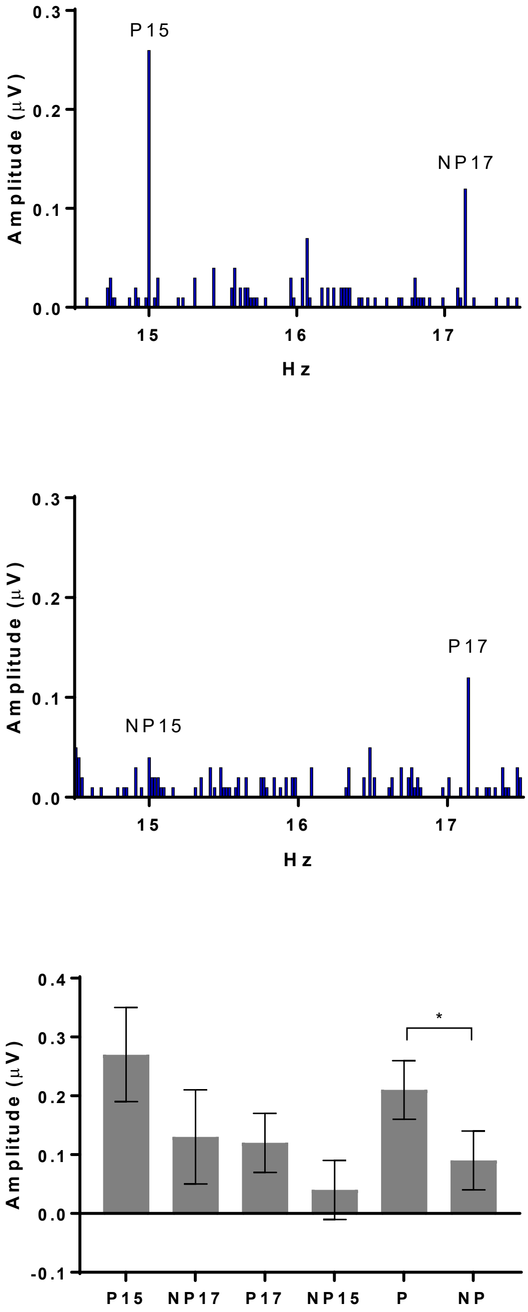

Figure 3.8: Amplitudes in the frequency domain (averaged over electrodes PO3/PO4, O1/O2 and Oz) measured during the Phase 1 categorisation task. Eleven participants viewed the P cue at 15.004 Hz (P15) and the NP cue at 17.147 Hz (NP17). Eight participants viewed the P cue at 17.147 Hz (P17) and the NP cue at 15.004 Hz (NP15). The last two columns represent the grand average collapsed across the two groups. Error bars represent standard error of the paired difference scores (* $p < .05$).

3.3.3. Discussion

One aim of Experiment 6 was to test the idea that probes that appear over attended items will elicit smaller N2pc amplitudes compared to probes that appear over unattended items. The second aim was to test whether the facilitation and inhibition time windows reported in Experiments 4 and 5 could be extended. Participants first learned to categorise pairs of gratings, with one grating in each pair predictive of the categorisation response and the other non-predictive. Participants had to make use of both colour and orientation information to make the correct categorisation response. During the categorisation task we obtained two electrophysiological measures of attention (N2pc, SSVEPs) and we expected them to show that attention was biased towards the P cues. The results of our frequency tagging analysis during the categorisation task confirmed that participants were attending to the P cue. SSVEP amplitudes at the P frequency were larger than those at the NP frequency (Figure 3.8). Further, the cue-locked N2pc also confirmed that participants were attending to the P cue. Posterior electrodes contralateral to the P cue showed a more negative voltage compared to electrodes ipsilateral to the P cue (Figure 3.6). Coupled with our SSVEP data, this result suggests that participants were indeed attending to the P cues during the categorisation task.

The same gratings were then used in an intermixed categorisation and dot probe task from which we obtained a behavioural (RT) and an electrophysiological measure (N2pc) of attention. Behaviourally, we expected to replicate RT facilitation effect in favour of the P cue at each SOA, and for the magnitude of the effect to increase with SOA. We also expected participants to delay their dot probe responses so that they could prepare their categorisation response. For the EEG data, we expected probes that appeared after short SOAs to elicit larger N2pc amplitudes when they appeared over NP cues compared to P cues (early time window). Conversely, probes that appeared after longer SOAs were expected to elicit larger amplitudes when they appeared over P cues compared to NP cues (late time window). Compared to the feature discrimination task used in Experiments 4 and 5, it was expected that the more difficult conjunction discrimination would result in an extension of the early N2pc time window and a delay of the late time window.

RTs were significantly faster when the probe appeared over a P cue compared to an NP cue at each SOA and the magnitude of the RT facilitation effect increased proportionally with SOA (Figure 3.5, top). This suggests that participants were strategically processing the cues (to prepare a categorisation response) during the SOA, consistent with Experiments 4 and 5. Also in line with Experiments 4 and 5, it was found that participants were strategically responding during the dot probe task. Figure 3.5 (bottom) shows that Phase 1 categorisation responses (i.e., responses not preceded by a dot probe response) were significantly longer than Phase 2 categorisation responses (i.e., time between dot probe response and subsequent categorisation response). This indicates that participants were delaying their dot probe response in order to inspect the cues and prepare a categorisation response.

The probe-locked N2pc results showed a significant interaction that implied changes in attention with time (Figure 3.7). Probes that appeared after short SOAs (100-300 ms) elicited a larger (i.e., more negative) N2pc amplitude when they appeared over a NP cue compared to a P cue. Conversely, probes that appeared after longer SOAs (500-600 ms) elicited a larger N2pc amplitude when they appeared over a P cue compared to a NP cue. Note, however, that although the direction of these latter effects was as expected, the simple effects were not significant (see Results). These results are in line with the prediction that the early N2pc time window would be extended because the conjunction discrimination required attention to remain with the P cue for longer compared to the feature based discrimination used in Experiments 4 and 5. The extension of the early time window also supports the idea that the early N2pc effect (i.e., more negative N2pc to probe over NP cue compared to P cue) actually reflects attention being allocated to the P cue. That is, when the probe appeared over an unattended NP cue it elicited a large N2pc amplitude because it was more difficult for the probe to pull attention away from the P cue, and this effect persisted even to an SOA of 300 ms given that attention to the P cue was likely still required to make the more difficult categorisation response. Stated the other way, when the probe appeared over a currently attended P cue it elicited a small N2pc amplitude because participants were attending to a spatial location proximate to the location of the probe when it appeared. The EEG results from the categorisation task also support this idea. It seems unlikely that participants would attend to the P cue during the categorisation task, and then switch to paying attention to the NP cue in the dot probe task. Therefore our interpretation of the larger N2pc to NP cues during the dot probe task – that it reflects greater attention to the P cue – seems most likely.

3.4. General Discussion

The experiments presented in this chapter surveyed a broader and earlier range of SOAs compared to the experiments presented in Chapter 2. Experiments 4 and 5 tested SOAs of 0, 100, 200, 300, 400, and 500 ms. Both experiments showed that targets that appeared over NP cues after short SOAs (100 and 200 ms) elicited larger N2pc amplitudes compared to targets that appeared over P cues (Figure 3.2, bottom). These early N2pc effects are in the same direction as those observed in Chapter 2. In addition, both experiments showed that the direction of the N2pc effect reversed when targets appeared after a later SOA (400 ms). Considering the behavioural and EEG data together, it seems likely that the earlier (and larger) N2pc effect reflects attention being directed towards the P cue. Participants were faster to respond to the target when it appeared over a P cue at each SOA (Figure 3.1, top), indicating attentional facilitation towards the P cue. Participants also made more dot probe errors (Figure 3.1, middle) and premature responses in the direction of the P cues. The error results indicate that dot probe RTs were contaminated by a non-attentional response bias towards selecting a response that was congruent with the position of the P cue. However, in order to select a response that was congruent with a P cue it seems logical that participants would have had to flag, and therefore attend to, the location of the P cue. If this is true for both premature responses and dot probe errors, then these data suggest that participants were indeed attending to the P cue during the dot probe task (for meta-analyses of dot probe errors and premature responses from all experiments, see the General Discussion in Chapter 5).

The RT data collected in Experiments 4 and 5 also showed a significant interaction (Figure 3.1, top). The RT advantage in favour of the P cues increased proportionally with SOA. This novel result indicates that participants were strategically processing the P cues during the SOA to prepare a categorisation response. The results also showed that categorisation RTs were significantly faster during Phase 2 (i.e., the intermixed dot probe and categorisation task) compared to Phase 1 (i.e., the categorisation task alone; Figure 3.3). This result suggests that participants continued to process the cues to prepare a categorisation response after the SOA had elapsed and may help explain why the pattern of RT results did not match the N2pc response pattern. In addition, the strategic responding observed in the intermixed categorisation and dot probe task may inform theories related to dual task interference and the psychological refractory period (Davis, 1959).

In Experiments 4 and 5, the earlier range of SOAs was chosen because automatic and endogenous shifts of covert attention have been shown to occur quite rapidly (Carlson, Hogendoorn, & Verstraten, 2006) and IOR has been shown to occur as early as 200-250 ms (Cao et al., 2017; Castel et al., 2003). Therefore, the N2pc effect found in Chapter 2 at the 250 ms SOA could reflect inhibition, not facilitation, of the P cue. If it is assumed that inhibition follows facilitation (Klein, 2000), then the fact that the N2pc effect reversed after 300 ms suggests that the early N2pc effect reflects facilitation. However, one could argue that it is the processing of the NP, not P, cue that is facilitated during the early time window.

Although this seems unlikely given the aforementioned behavioural data, one aim of Experiment 6 was to further investigate whether the early N2pc effect reflects facilitation of the P or NP cue. A second, and related, aim of Experiment 6 was to test whether the time windows of the N2pc effects could be shifted and/or enlarged depending on the difficulty of the categorisation task (and hence the amount of time spent processing the P cue in order to make a correct categorisation response).

Experiment 6 required participants to discriminate between the P and NP cues using both colour and orientation information (i.e., a more difficult conjunction discrimination). In addition, Experiment 6 employed the frequency tagging technique to obtain another electrophysiological measure of attention. The behavioural results of Experiment 6 replicated those observed in Experiments 4 and 5. The RT results showed that participants were faster to respond to targets that appeared over P cues compared to NP cues at each SOA and the magnitude of this facilitation effect increased proportionally with SOA (Figure 3.5, top). Participants in Experiment 6 also delayed their dot probe response so that they could continue to process the cues and prepare a categorisation response (Figure 3.5, bottom).

The ERP and SSVEP results from the categorisation task indicated that participants were attending to the P cue. Posterior electrodes contralateral to the P cue showed a greater negative voltage compared to ipsilateral electrodes (Figure 3.6) and SSVEP amplitudes were larger at the P frequency compared to the NP frequency (Figure 3.8). In the subsequent dot probe task, the N2pc response pattern was very similar to that observed in Experiments 4 and 5. The early N2pc effect in Experiment 6 was in the same direction as the previous experiments (i.e., larger N2pc amplitudes elicited by targets that appeared over the NP cues). However, the early N2pc time window in Experiment 6 was extended from the 200 ms SOA observed in Experiments 4 and 5 to 300 ms. This result agrees with the idea that the conjunction discrimination used in Experiment 6 required participants to attend to the P cue for longer (to process orientation information) compared to Experiments 4 and 5 (which required only colour processing). This result also agrees with the idea that the early N2pc effect reflects facilitation of the P cue because it seems unlikely that participants would pay attention to the NP cue for longer in a conjunction discrimination. Experiment 6 also showed a reversal of the N2pc effect after a 500 ms SOA. This latter reversal supports the prediction that the inhibition effect observed in Experiments 4 and 5 at the 400 ms SOA would be delayed. Note, however, that although this effect was in the expected direction it did not reach statistical significance (see Results of Experiment 6).

The N2pc results of the experiments presented in this chapter suggest that attention might switch between the P and NP cues over the time course of a trial. This possibility has been largely overlooked by computational modellers. In the final chapter (Chapter 5) we present a mathematical model that attempts to capture this movement of attention, and show how it might influence subsequent cue-outcome learning. In the interim, Chapter 4 presents a series of experiments that aim to investigate whether participants strategically or automatically attend to the P cues, and test whether the attentional effects found in the previous experiments can affect subsequent learning.

CHAPTER 4: Automaticity and Learning

An important question that remains to be resolved is whether learning driven changes in attention are generally under top-down or bottom-up control. In other words, to what extent are participants deliberately or automatically attending to previously predictive cues? This question has been approached in different ways. For example, Mitchell et al. (2012) tested whether the learned predictiveness effect is sensitive to top-down attentional processes by manipulating the instructions given to participants between phase one and phase two. Participants in a continuity group were instructed that the same stimuli relevant in phase one would be relevant in phase two, whereas those in a change group were instructed that cues that were relevant in phase one would be irrelevant in phase two. They found that those in the change group learnt more about the relationship between previously non-predictive cues and phase two outcomes, and looked longer at these cues compared to the previously predictive cues. The fact that instruction could reverse the learned predictiveness effect suggests that the effect can, to some extent, be explained by top-down factors.

In their Experiment 3, Le Pelley et al. (2010) had participants complete a learned predictiveness experiment but provided them with written summaries of the cue-outcome contingencies. That is, in Phases 1 and 2 participants had constant access to information regarding which cues were followed by which outcomes, and the frequencies of these co-occurrences. The authors reasoned that such information would encourage the participants to engage in high level reasoning by reducing cognitive load (e.g., memory). If the learned predictivness effect is the product of rational inference, then the summary information should, if anything, strengthen the effect. Contrary to this, the results did not elicit a learned predictiveness effect, suggesting that the learned predictiveness effect is not solely due to rational inference.

Another approach used to investigate the automaticity of the learned predictiveness effect involves looking at the time course of attention. In their dot probe task, Le Pelley et al. (2013) manipulated SOA to investigate the automaticity of attention. After a short SOA (250 ms) participants responded faster to the probe when it appeared over a predictive cue compared to a non-predictive cue. However, after a longer SOA (1000 ms) no difference in reaction time was found (also see Cobos et al., 2018). The results were taken as evidence that the previously predictive cues automatically captured attention because the short SOA gave participants little time to consciously process the stimuli. In addition, if the participants were strategically processing the cues then the long SOA, which provided more time for top-down processing, should have produced a larger, or at least equivalent, RT facilitation effect compared to the short SOA. Le Pelley et al. (2013) also argued that during the long SOA participants could have used controlled processes to reorient their attention away from the P cue and back to fixation. This is because the participants knew that the probe was equally likely to appear over either the P or NP cue and that the best strategy in the dot probe task was to ignore the cues and attend centrally throughout.

In contrast to previous dot probe research (Le Pelley et al., 2013; Luque et al., 2017; Cobos et al., 2018), the experiments presented in Chapter 3 used a wider and early range of SOAs (e.g., 0-500 ms) and the results showed that the RT advantage in favour of the P cues increased proportionally with the SOAs tested. This could suggest that the long SOA used in previous studies (i.e., 1000 ms) was too late to capture the strategic attention observed in our experiments. In other words, perhaps strategic processing of the cues is completed before 1000 ms. This would mean that the RT effect found after intermediate SOAs in the previous chapter (e.g., 500 ms) would reduce if a longer SOA was tested. However, there are also methodological differences that could explain the discrepant results. For example, it is possible that the response set-up in our experiments encourages a stronger strategic bias towards the P cue[13]. The results of our experiments have shown that participants were not responding to the probe as fast as possible, but instead let it linger briefly on the screen while they continued to process the cues and prepare a categorisation response. After they had prepared a categorisation response, participants made both the dot probe response and the categorisation response in quick succession (see Figures 3.3 and 3.5). If the response set up encourages this strategic behaviour, then perhaps the RT bias will not reduce after a longer SOA. On the other hand, if strategic processing of the cues is completed by 1000 ms, then the RT advantage afforded to P cues should be eliminated, or significantly reduced, after a 1000

---

[13] We would like to thank an anonymous reviewer for this suggestion.

ms SOA. Experiment 7 aimed to investigate this by using a 1000 ms SOA in addition to the SOAs used in Experiments 4 and 5.

4.1. Experiment 7

Using an intermixed design, previous dot probe research (Le Pelley et al., 2013; Cobos et al., 2018) has shown a RT facilitation effect for P cues after a short SOA (250 ms) and no RT effect after a long SOA (1000 ms). In other words, previous research has shown that increasing SOA significantly weakens the influence of predictiveness on dot probe biases. This pattern of results has been taken as evidence that the early RT facilitation effect is due to automatic, rather than controlled, attentional biases (Le Pelley et al., 2016). However, the intermixed tasks presented in the previous chapter (Experiments 4, 5, and 6) used a broader and earlier range of SOAs and showed that the RT advantage afforded to the P cues increases proportionally with SOA. This pattern of results suggests that the P cues are being processed strategically. It is possible that by using only a short and long SOA previous dot probe research may have missed the RT effects that occur at intermediate SOAs. That would mean that RT effects found after an intermediate (e.g., 500 ms) SOA in the previous experiments would significantly reduce if a longer (e.g., 1000 ms) SOA was tested. However, it is also possible the response set-up in our experiments might have encouraged a strategic bias towards P cues because it was easy for participants to execute the dot probe and categorisation responses together.

Experiment 7 aimed to investigate this idea by using an additional 1000 ms SOA. If strategic processing of the cues is complete before 1000 ms, then the 1000 ms SOA should show a reduced or no RT bias. If, on the other hand, the response set up in our experiments encourages participants to maintain a strategic bias towards the P cues, then the RT effect at the 1000 ms SOA should increase or remain at a high level. This experiment also provided another opportunity to replicate the behavioural (RT) and EEG (N2pc) effects observed in our previous experiments.

4.2.1. Method

Participants

Nineteen participants (9 male; mean age 22 years, age range 19 to 26 years) took part in the Experiment. One participant did not meet the Phase 1 learning criterion and was excluded from further analysis. All participants reported normal or corrected-to-normal vision and participated in two 1.5 hour sessions. Participants gave written informed consent and were compensated with $80 gift cards.

Apparatus

All apparatus were identical to Experiments 4 and 5.

Stimuli

All stimuli were identical to those used in Experiments 4 and 5.

Procedure

The procedure was identical to Experiments 4 and 5 except that an additional SOA was tested (0, 100, 200, 300, 400, 500 and 1000 ms). Each participants completed a total of 144 trials for each of the 14 within-subject conditions. Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n=9) and “Left-Right Predictive” (n=9). Five participants in the “Vertical-Horizontal Predictive” and five participants in the “Left-Right Predictive” condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings and green left gratings.

EEG Recording and Analysis

The continuous EEG was segmented into epochs ranging from 100 ms prior to the onset of the dot probe target to 250 ms (for the 200, 300, 400, 500, and 1000 ms SOAs) or 300 ms (for the 0 and 100 ms SOAs) post target onset, and baseline corrected using the 100 ms pre-stimulus interval.

ERP waveforms were time-locked to the onset of the white square target. Lateralised ERPs were computed for each of the 14 within-subjects conditions (i.e., when the target appeared over the P and NP cues at each SOA). The 50% fractional area latency of the N2pc component was calculated using a time window of 150-300 ms for the 0 and 100 ms SOAs and a time window of 100-250 ms for the 200, 300, 400, 500 and 1000 ms SOAs. Fourteen separate latencies (i.e., one for each within-subjects condition) were calculated for each participant. The median N2pc latency for each within-subject condition was calculated.

Similarly to the previous experiments, significant differences in N2pc latencies were found between SOA conditions (see Figure 4.3). There was a significant main effect of SOA on N2pc latencies at PO7/PO8 ($F_{(2.64, 44.93)} = 121.39$, $p < .001$, $\eta^2 = ..795$, MSE = 329.06), P7/P8 ($F_{(2.80, 47.56)} = 157.6$, $p < .001$, $\eta^2 = .84$, MSE = 287.2) and O1/O2 ($F_{(3.09, 52.45)} = 109.03$, $p < .001$, $\eta^2 = .77$, MSE = 300.02) electrode sites. Therefore, the mean amplitude of the N2pc component elicited by the target was measured during different 50 ms time windows centred on the median latencies for each SOA. The median latencies for each SOA measured at the PO7/PO8, P7/P8 and O1/O2 electrode sites can be found in Appendix D. All other EEG recording and ERP analyses were identical to Experiments 4 and 5.

Statistical Analysis

Statistical analyses were identical to Experiments 4 and 5.

4.2.2. Results

Figure 4.1 (top) shows the average of the median RTs to the target in the dot probe task as a function of Predictiveness and SOA. There was a main effect of Predictiveness ($F_{(1, 17)} = 37.75$, $p < .001$, $\eta^2 = .0.091$, MSE = 991.41, $d_{unb} = .35$) because participants responded faster to the target when it appeared over a P cue (M = 387 ms) compared to a NP cue (M = 411 ms). There was a main effect of SOA because participants were slower to respond to the target when it appeared after short SOAs compared to long SOAs ($F_{(1.78, 30.27)} = 129.14$, $p < .001$, $\eta^2 = .74$, MSE = 1311.83). The Predictiveness×SOA interaction was also significant ($F_{(6,102)} = .10.13$, $p < .001$, $\eta^2 = .012$, MSE = 83.52). Paired sample t-tests revealed that participants were significantly faster to respond to the target when it appeared over a P cue compared to a NP cue at each SOA (largest p = .007 at the 0 ms SOA). Figure 4.1 (bottom) shows the effect size at each SOA.

Figure 4.2 shows participants' categorisation RTs in Phase 1 and at each SOA during Phase 2 (i.e., the time between the dot probe response and the subsequent categorisation response). We compared participants categorisation RTs in Phase 1 (M = 524 ms) with those in the 0 ms SOA condition in Phase 2 (M = 219 ms) using a paired samples t-test. Phase 1 categorisation responses were significantly slower than those measured in the 0 ms SOA condition in Phase 2 ($M_{diff} = 305$ ms, $t_{(17)} = 12.69$, $p < .001$, $d_{unb} = 2.51$).

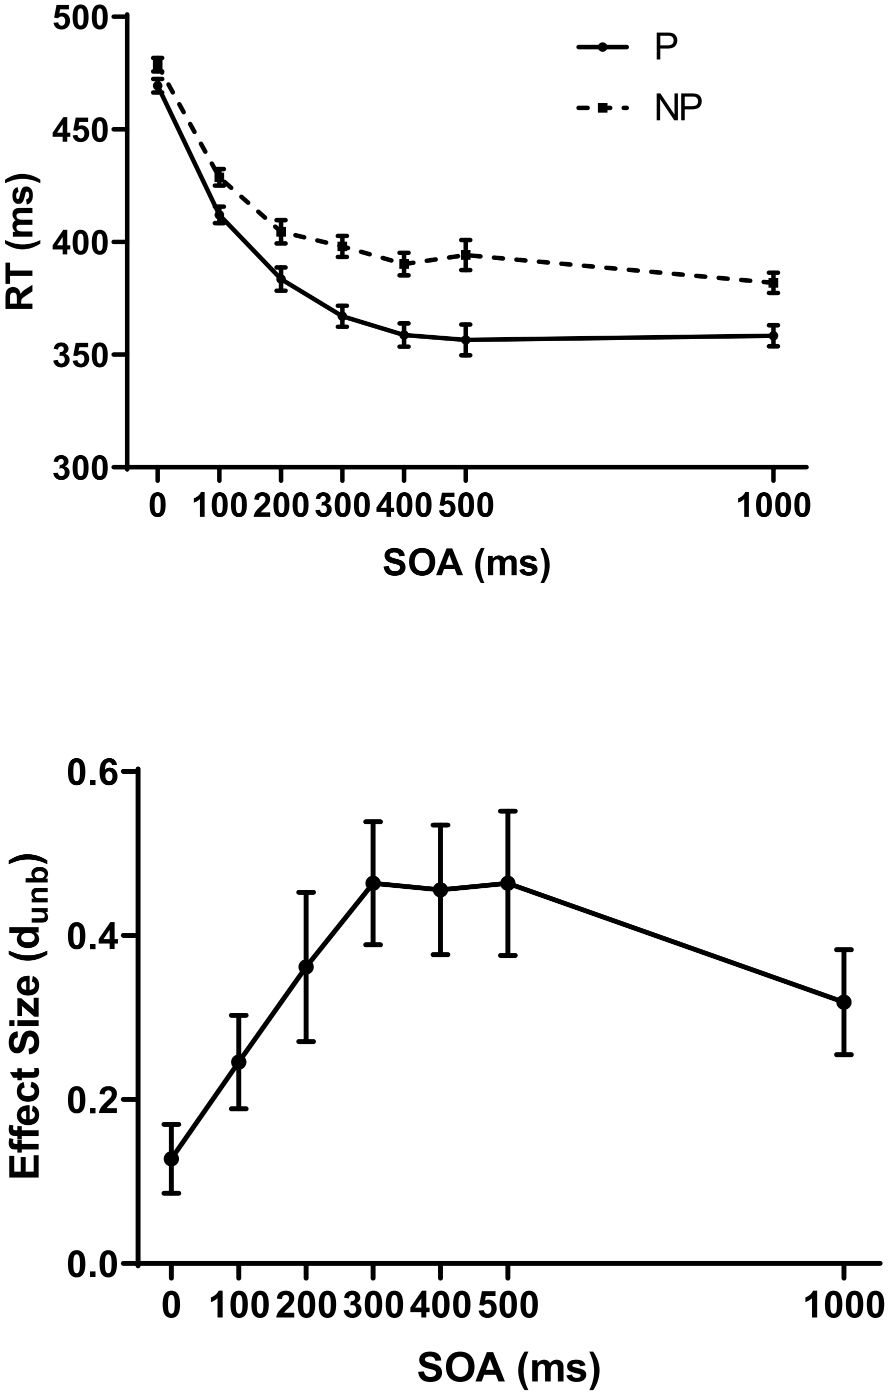

Figure 4.1: Median reaction time to the dot probe target as a function of Predictiveness and SOA (top) and effect size as a function of SOA (bottom). Error bars represent standard error of the paired difference scores (top) and standard errors for dunb (bottom).

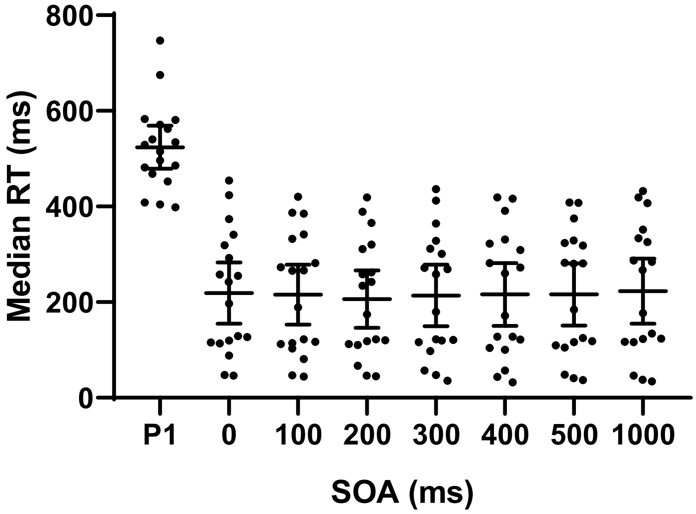

Figure 4.2: RT data from Session 2. Phase 1 (P1) and Phase 2 median categorisation RTs at each SOA. Categorisation RT in Phase 2 was the time between the dot probe response and the subsequent categorisation response in the dot probe task. Error bars represent 95% confidence intervals of the mean.

Figure 4.3 (top row) shows contralateral minus ipsilateral difference waveforms for each SOA when the target appeared over a P cue or a NP cue. The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix D. Figure 4.3 (bottom right) shows the average of the N2pc mean amplitudes evoked by the target as a function of Predictiveness and SOA. The main effect of SOA on N2pc mean amplitudes was significant because targets elicited larger (i.e., more negative) voltages when they appeared after long SOAs compared to short SOAs ($F_{(3.22, 54.79)} = 6.21$, $p = .001$, $\eta^2 = 0.15$, MSE = 2.02). The main effect of Predictiveness was not significant ($F_{(1, 17)} = 1.62$, $p = .22$, $\eta^2 = 0.008$, MSE = 1.31, $d_{unb} = .14$), nor was the Predictiveness×SOA interaction ($F_{(3.57, 60.76)} = 1.06$, $p = .38$, $\eta^2 = 0.022$, MSE = 1.60).

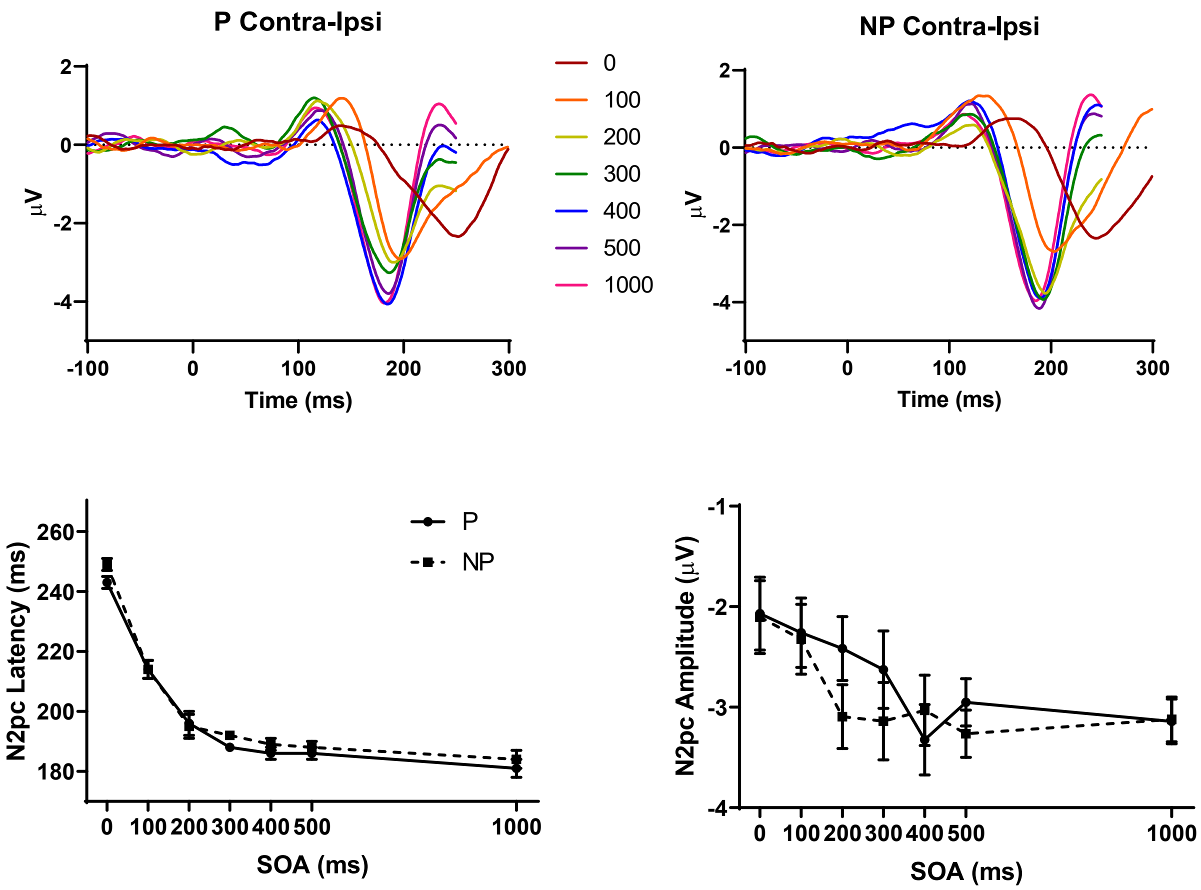

Figure 4.3: The top row shows contralateral minus ipsilateral difference waves for each SOA when the probe appeared over a P cue (left) or a NP cue (right). The bottom row shows median N2pc latencies (left) and N2pc mean amplitudes (right) as function of predictiveness and SOA. All data were averaged over PO7/PO8, P7/P8, and O1/O2 electrodes. Error bars represent standard error of the paired difference scores.

4.2.3. Discussion

Experiment 7 was designed to investigate the time course of strategic attention by using an additional long SOA (1000 ms). This was done because previous dot probe research (Le Pelley et al., 2013, Cobos et al., 2018) has found a reduced dot probe bias at longer SOAs. Using an intermixed dot probe and categorisation task, both Cobos et al., and Experiment 3 in Le Pelley et al. found no difference in RTs to a probe that appeared over P and NP cues after a 1000 ms SOA. However, in the experiments presented in Chapter 3 we found that the RT advantage afforded to the P cues increased proportionally with SOA (0-600 ms). One way to explain why previous dot probe research has found no dot probe bias after a 1000 ms SOA is that strategic processing of the cues is completed before 1000 ms. However, the RT results of the current experiment suggest that this is not the case. From Figure 4.1 it is clear that there is still a RT advantage in favour of the P cues after the 1000 ms SOA, although the effect does appear to reduce slightly after this longer SOA. Therefore, it is difficult to see how the time course of strategic attention can explain the null result usually found at the 1000 ms SOA (Le Pelley et al., 2013; Cobos et al., 2018).

An alternative explanation relates to methodological factors. For example, it is possible that the response set up in our experiments encourages participants to engage in a stronger strategic bias compared to previous dot probe research. Indeed, Figure 4.2 shows that participants were once again responding strategically during the dot probe task in Experiment 7. However, Experiments 1 and 2 (Chapter 2) used an identical response set up to the current experiment and the RT results of those experiments did not find evidence of strategic processing. To test whether participants were responding strategically during Experiments 1 and 2 we analysed the categorisation RT data from Phases 1 and 2. The data can be seen below in Figure 4.4.

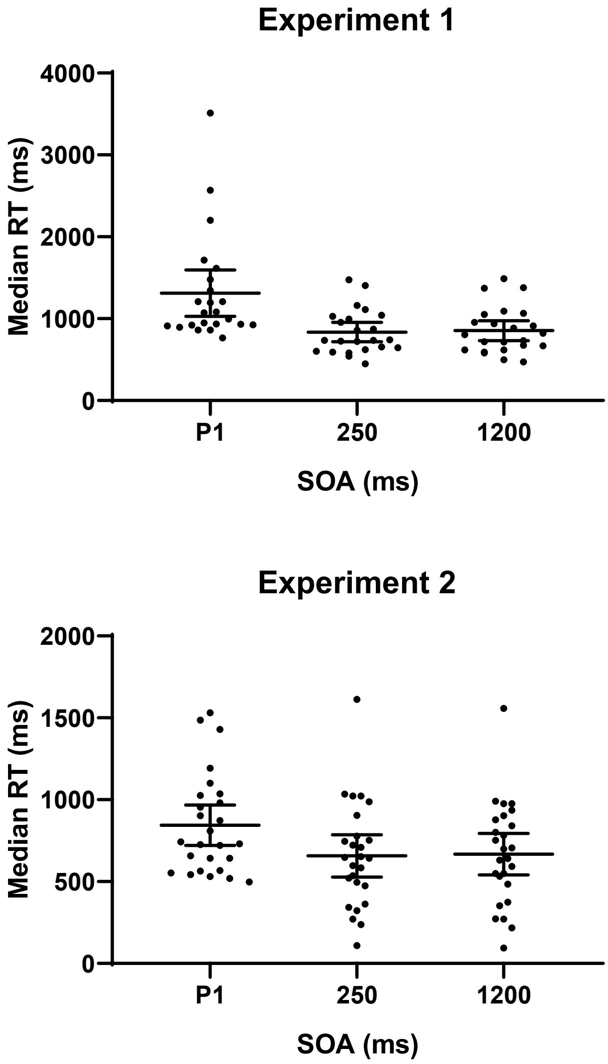

Figure 4.4 Median categorisation RT during Phase 1 (P1) and after each SOA (250 and 1200 ms) in Phase 2 for Experiments 1 (top) and 2 (bottom). Categorisation RT in Phase 2 was the time between the dot probe response and the subsequent categorisation response in the dot probe task. Error bars represent 95% confidence intervals of the mean.

In Experiments 1 and 2 there was no 0 ms SOA in Phase 2. Therefore, we compared the Phase 1 categorisation RTs to the categorisation RTs in the 250 ms SOA condition. Paired sample t-tests showed that Phase 1 categorisation RTs were slower than Phase 2 RTs in both Experiment 1 (M_{diff}= 474 ms, $t_{(22)}$ = 3.98, p = .001, d_{unb} = .80) and Experiment 2 (M_{diff}= 187 ms, $t_{(25)}$ = , p = .032, d_{unb} = .58). Therefore, participants were still responding strategically during those experiments. However, the effect sizes of the strategic response (i.e., the difference between Phase 1 categorisation RTs and Phase 2 categorisation RTs) obtained in Experiments 1 and 2 were smaller than those obtained in Experiments 4, 5, 6, and 7. [14] Therefore, perhaps participants in Experiments 1 and 2 were better able to ignore the cues before responding to the dot probe target. Experiments 4-7 tested a wider range of SOAs compared to Experiments 1 and 2, so there were more within subject conditions that each participant had to complete. This resulted in a longer experiment. In addition, testing multiple SOAs (e.g., six) could have increased interference between each SOA condition because participants were more uncertain as to when the dot probe target would appear. A longer experiment with multiple randomly interleaved SOAs could be the driving force that encourages participants to respond strategically.

The N2pc results of the current experiment only partially replicate those seen in Experiments 4 and 5. The early N2pc effects were in the expected direction (i.e., greater N2pc amplitudes to targets that appear over NP cues compared to P cues) but there was no significant interaction between cue type and SOA. In addition, the early N2pc effects observed in Chapter 3 were largest after the 100 ms SOA, whereas in the current experiment the 200 ms SOA showed the largest effect. Although the direction of the N2pc effect reversed after the 400 ms SOA, the difference between the P and NP cues was not statistically significant. The only difference between the current experiment and Experiments 4 and 5 was the addition of the 1000 ms SOA, but it is not clear how the addition of this SOA could affect the N2pc results. To speculate, one possibility is that the 1000 ms SOA is somewhat of an outlier compared to the other SOAs tested and this SOA could have altered the temporal expectation of the target within each trial. For example, using SOAs of 0, 100, 200, 300, 400,

---

[14] Experiment 1: dunb = .80, S.E. = 0.23; Experiment 2: dunb = .58, S.E. = 0.28; Experiment 4: dunb = 1.8, S.E. = 0.39; Experiment 5: dunb = 1.5, S.E. = 0.27; Experiment 6: dunb = 3.2, S.E. = 0.74; Experiment 7: dunb = 2.5, S.E. = 0.42

and 500 ms the participant might implicitly develop an average expectation of target arrival to be 250 ms (i.e., average of the SOAs). By including the 1000 ms SOA, this average expectation would increase to 357 ms. It is possible that this change in average expectancy could have altered attentional orienting, response preparation and the time windows of the N2pc results. Temporal expectancy has been shown to modulate the shift from facilitation to IOR (Gabay & Henik, 2010), however this was shown in the context of the aging-foreperiod effect (i.e., the expectancy of target occurrence increases as cue target interval elapses) and not a simple average wait time. Despite this failure to replicate the pattern of N2pc results observed in Chapter 3, the direction of the effects observed in the current study agree with our previous results.

The intermixed within-subject design might also encourage strategic processing of the cues because the cues are task relevant on every dot probe trial. Using a blocked design could mitigate this type of strategic attention because the cues have no task relevance during the dot probe phase. In addition, only one response (to the probe) is required in the blocked version of the dot probe task, which should reduce strategic responding. Therefore, it might be possible to use a blocked design to investigate whether P cues that have no task relevance are automatically or strategically attended to. Experiment 2 of Le Pelley et al. (2013) used a blocked design and found a RT facilitation effect in favour of the P cue after a 250 ms SOA, and no effect after 1000 ms. This pattern of results suggested that the early facilitation effect was automatic because top down processes should have increased the effect after 1000 ms. However, the RT results observed in Experiments 4-7 of the current thesis suggests that it might be informative to test earlier SOAs (0-500 ms) in a blocked design. For example, if participants do not strategically process the cues in a blocked design then the RT interaction observed in Experiments 4 to 7 should be reversed. That is, the initial RT facilitation effect should reduce after long SOAs. In contrast, if participants are still strategically processing the cues for some reason (e.g., they think the probe will appear over the P cues more often), then the longer SOAs should produce a stronger RT facilitation effect.

4.2. Experiment 8

The aim of Experiment 8 was to investigate what effect blocking, rather than intermixing, the associative learning and dot probe tasks would have on the attentional allocation towards the P and NP cues during the dot probe phase. Specifically, we were interested to see whether a blocked design would result in automatic or strategic processing of the P cues during the dot probe task. Unlike the intermixed version of the task, the cues are not task relevant during the dot probe phase in the blocked design. Therefore, participants might be less motivated to engage in strategic processing of the cues. In accordance with previous research that has blocked the categorisation and dot probe tasks (e.g., Experiment 2 of Le Pelley et al., 2013), we expected participants to respond faster to the probe over P cues compared to NP cues after short SOAs. If no strategic processing of the cues takes place during the blocked design, then the RT interaction observed in Experiments 4 to 7 should reverse. In other words, the RT facilitation effect in favour of the P cues should reduce after longer SOAs. However, if some strategic processing remains then the RT advantage in favour of the P cues should increase with SOA. In a blocked design, each dot probe trial also acts as an extinction trial for any learned attentional bias. Consequently, the RT effects in the blocked design are likely to be smaller in magnitude compared to those found using an intermixed design (Le Pelley et al.). In Experiment 8, we also analysed the RT data to test whether the RT effects extinguished over the course of a dot probe block (see Statistical Analysis).

4.1.1. Method

Participants

Twenty -participants (15 male; mean age 20 years, age range 18 to 37 years) took part in the experiment. One participant did not meet the learning criterion for the categorisation task and was excluded. All participants reported normal or corrected-to-normal vision. Participants gave written informed consent and were compensated with credit for an undergraduate psychology course.

Apparatus

All apparatus were identical to Experiments 4 and 5 (i.e., the experiment was again run on E-Prime).

Stimuli

All stimuli were identical to Experiments 4 and 5.

Procedure

Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n=14) and “Left-Right Predictive” (n=14). In each experiment, seven participants in the “Vertical-Horizontal Predictive” and seven participants in the “Left-Right Predictive” condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings and green left gratings.

The experiment was very similar to experiments 4 and 5 except that participants alternated between short blocks consisting of either the categorisation task or the dot probe task. Overall participants completed 10 categorisation tasks and 10 dot probe tasks. The categorisation task is summarised in Figure 4.5 (left). The first categorisation task consisted of thirty blocks (240 trials). Participants completed a further nine categorisation tasks. These were identical to the first, except they consisted of five blocks (40 trials) instead of thirty.

The dot probe task is summarised in Figure 4.5 (right). The dot probe task was split into blocks of 96 trials. Each block consisted of all possible combinations of four stimulus pairs (Table 1), two cue locations (P or NP cue on the left or right), six SOAs (0, 100, 200, 300, 400 and 500 ms), and two target locations (left or right). The 96 trials within a block were presented in random order. Participants completed one block of the dot probe task after each categorisation task (i.e., ten blocks total). The tasks were alternated in this way so that learning about the predictive status of the cues was "topped up" prior to each dot probe task (Le Pelley et al, 2013). There were twelve within-subjects conditions (the target appearing over the P and NP cues at each of the six SOAs). Each participant completed a total of 80 trials in each condition.

EEG Recording and Analysis

Similarly to the previous experiments, significant differences in N2pc latencies were found between SOA conditions (see bottom left of Figure 4.4). There was a significant main effect of SOA on N2pc latencies at PO7/PO8 ($F_{(2.97, 59.40)} = 117.72$, $p < .001$, $\eta^2 = .75$, MSE = 289.02), P7/P8 ($F_{(3.27, 68.62)} = 127.4$, $p < .001$, $\eta^2 = .73$, MSE = 321.42) and O1/O2 ($F_{(3.01, 60.17)} = 73.52$, $p < .001$, $\eta^2 = .64$, MSE = 371.16) electrode sites. Therefore, the mean amplitude of the N2pc component elicited by the target was measured during different 50 ms time windows centred on the median latencies for each SOA. The median latencies for each SOA measured at the PO7/PO8, P7/P8 and O1/O2 electrode sites can be found in Appendix E. All other EEG recording and ERP analyses were identical to Experiments 4 and 5.

Statistical Analysis

All statistical analyses were identical to Experiments 4 and 5. However, in this experiment, we also analysed each participants RT data to test whether the RT effects extinguished over the course of the dot probe trials. Each dot probe trial was sorted into one of twelve arrays based on the within subject factors of SOA (0, 100, 200, 300, 400, 500 ms) and Predictiveness (P or NP) condition (i.e., each array corresponded to P0, P100, P200, P300, P400, P500, NP0, NP100, NP200, NP300, NP400, NP500). Each array ranked the order in which each trial was presented. Therefore, the first eight rows of each array corresponded to first block of dot probe trials (i.e., 8 × 12 = 96 trials in a block). Each array contained eighty cells, which accounted for all ten blocks. Within each array, we averaged corresponding trials from each block (i.e., ten values of trial one, ten values for trial two, etc.), which resulted in eight RT values in each array. We then averaged across SOAs to get eight P RTs and eight NP RTs for each participant. These values were entered into a repeated measures ANOVA with two levels of Predictiveness (P and NP) and eight levels of Trial within a block.

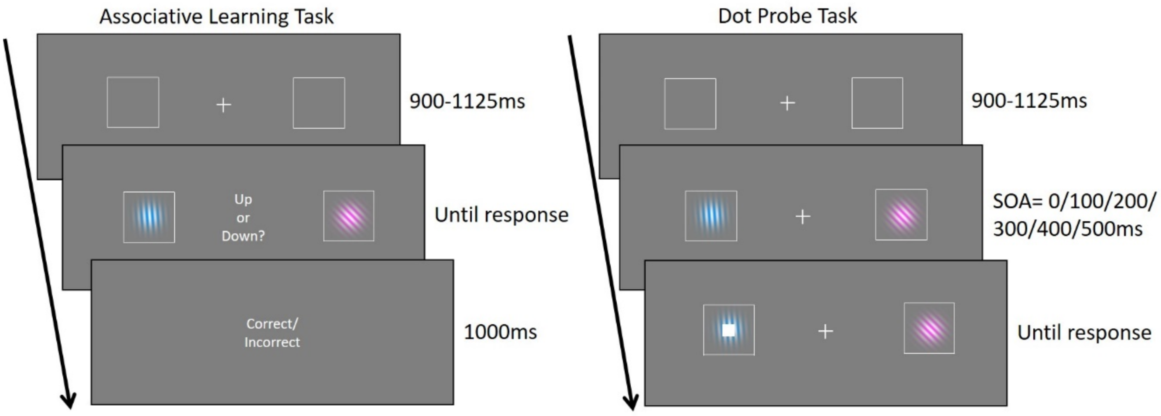

Figure 4.5: Examples of trials from the associative learning and dot probe tasks. Participants were instructed to maintain central fixation for both tasks. Each trial began with a pre-stimulus screen, which consisted of a central cross that was flanked by two placeholder boxes. After a jittered delay, two grating stimuli of different colour and orientation appeared within each placeholder box. During the associative learning task, participants categorised each pair presented to them by pressing either the up or down buttons on a button box. After responding to the gratings a feedback screen was presented. The same gratings were presented during the dot probe task. After a variable SOA a white square probe appeared randomly over the grating on the left or right. Participants responded as fast as possible to the location of the square on the screen by pressing either the left or right button on the button box.

4.1.2. Results

Figure 4.6 (top) shows the average of the median RTs to the target in the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on RTs was significant ($F_{(1, 26)} = 31.0$, $p < .001$, $\eta^2 = .0.01$, MSE = 78.18, $d_{unb} = .11$) as participants responded faster to the target when it appeared over a P cue (M = 346 ms) compared to a NP cue (M = 351 ms). There was a main effect of SOA as participants were slower to respond to the target when it appeared after short SOAs compared to long SOAs ($F_{(1.25, 32.51)} = 118.44$, $p < .001$, $\eta^2 = .80$, MSE = 2423.47). The Predictiveness×SOA interaction was not significant

($F_{(3.19, 83.0)} = .32$, $p = .82$, $\eta^2 < .001$, MSE = 84.91). Figure 4.6 (bottom) shows the effect size at each SOA.

Figure 4.7 (top) shows the average of median RTs to the target in the dot probe task as a function of Predictiveness and Trial position within a block. Trial 1 corresponds to the earliest part of a block and Trial 8 to the latest. The main effect of Predictiveness on RTs was significant ($F_{(1, 26)} = 15.22$, $p < .01$, $\eta^2 = .0.036$, MSE = 186.62, $d_{unb} = .090$) because participants were faster to respond to targets that appeared over P cues (M = 361 ms) compared to NP cues (M = 366 ms). There was also a main effect of Trial ($F_{(3.99, 103.71)} = 8.20$, $p < .001$, $\eta^2 = .15$, MSE = 353.89) because participants' RTs at the start of a block were faster than their RTs at the end of a block. However, there was no significant Predictiveness×Trial interaction ($F_{(5.55, 144.16)} = .82$, $p = .55$, $\eta^2 = .009$, MSE = 150.92), suggesting that the RT advantage in favour of P cues was constant from the start of a block to the end.

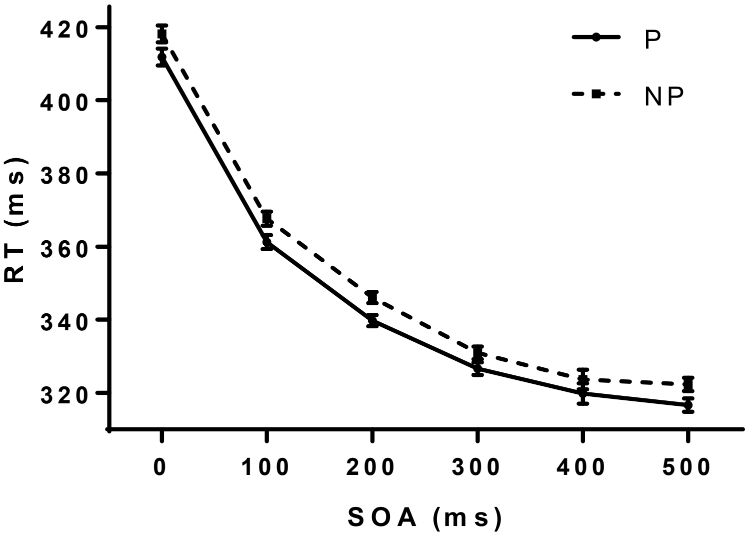

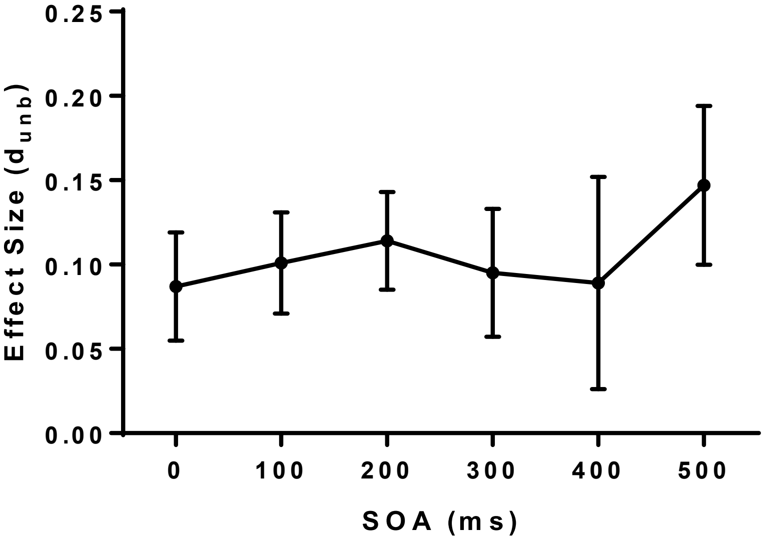

Figure 4.6: Median reaction time to the dot probe target as a function of Predictiveness and SOA (top) and effect size as a function of SOA (bottom). Error bars represent standard error of the paired difference scores (top) and standard errors for d_{unb} (bottom).

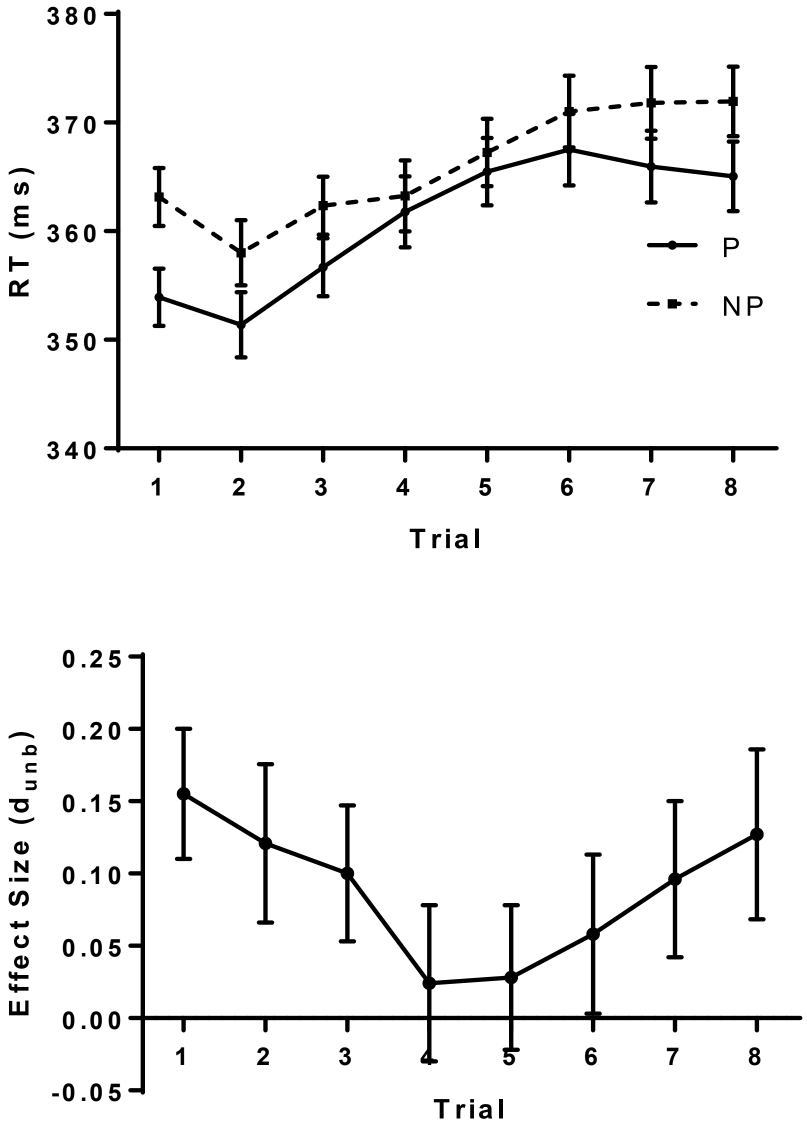

Figure 4.7: Median reaction time to the dot probe target as a function of Predictiveness and Trial position within a block (top) and effect size as a function of Trial (bottom). Error bars represent standard error of the paired difference scores (top) and standard errors for d_{unb} (bottom).

Figure 4.8 (top row) shows contralateral minus ipsilateral difference waveforms for each SOA when the target appeared over a P cue or a NP cue. The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix E. Figure 4.8 (bottom right) shows the average of the N2pc mean amplitudes evoked by the target as a function of

Predictiveness and SOA. The main effect of SOA on N2pc mean amplitudes was significant because targets elicited larger (i.e., more negative) voltages when they appeared after long SOAs compared to short SOAs ($F_{(2.91, 75.71)} = 6.85$, $p < .001$, $\eta^2 = 0.11$, MSE = 1.79). The main effect of Predictiveness was not significant ($F_{(1, 26)} = .51$, $p = .48$, $\eta^2 = 0.003$, MSE = 1.02, $d_{unb} = .09$), nor was the Predictiveness×SOA interaction ($F_{(3.45,89.68)} = .52$, $p = .69$, $\eta^2 = 0.006$, MSE = 1.09).

Figure 4.9 shows the mean number of dot probe errors measured during the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on Errors was significant ($F_{(1, 26)} = 5.47$, $p = .027$, $\eta^2 = .0.21$, MSE = 1.77, $d_{unb} = .29$) as participants made more errors when the target appeared over the NP cue (M = 1.4) compared to the P cue (M = 1.0). This meant that participants were more likely to indicate that the target appeared over the P cue when in fact it had appeared over the NP cue, rather than vice versa. There was a main effect of SOA because participants made more errors when the target appeared after long SOAs compared to short SOAs ($F_{(3.14, 81.70)} = 6.13$, $p = .001$, $\eta^2 = .10$, MSE = 2.31). The Predictiveness×SOA interaction was not significant ($F_{(5, 130)} = .021$, $p = .99$, $\eta^2 < .001$, MSE = 1.29).

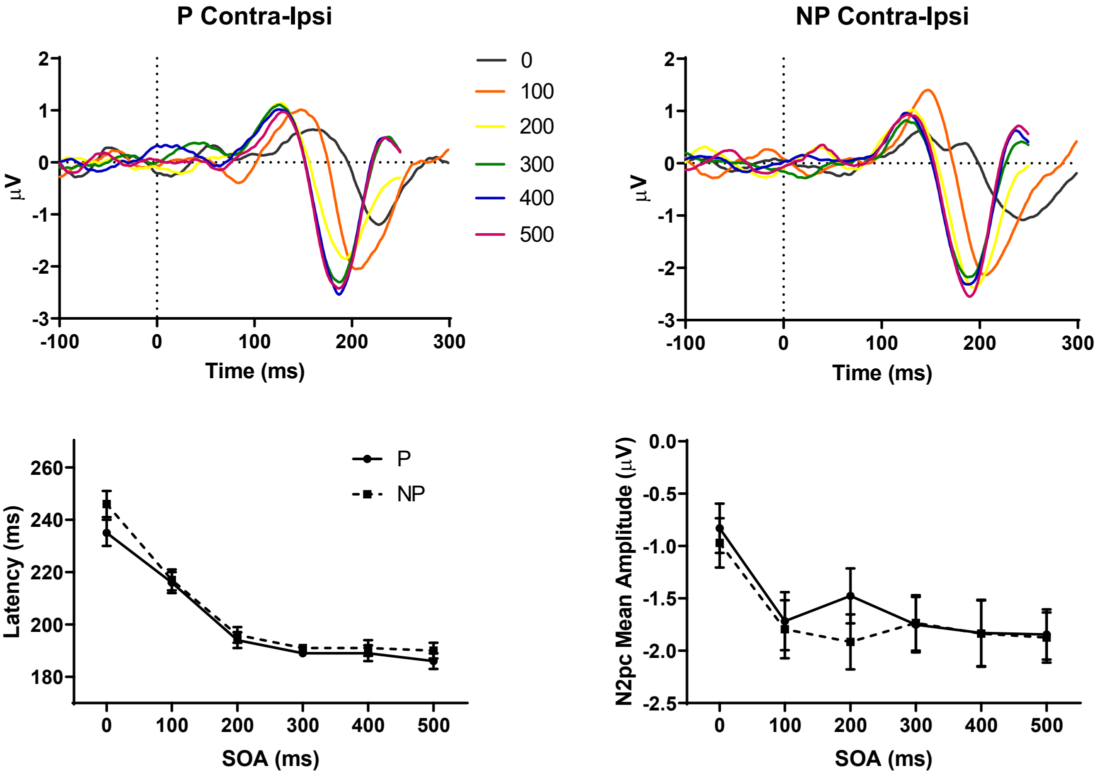

Figure 4.8: The top row shows contralateral minus ipsilateral difference waves for each SOA when the probe appeared over a P cue (left) or a NP cue (right). The bottom row shows median N2pc latencies (left) and N2pc mean amplitudes (right) as function of predictiveness and SOA. All data were averaged over PO7/PO8, P7/P8, and O1/O2 electrodes. Error bars represent standard error of the paired difference scores.

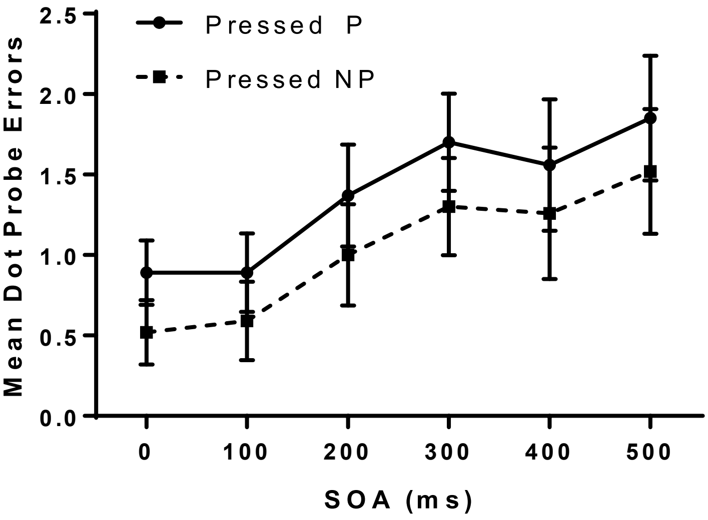

Figure 4.9: Mean dot probe errors as a function of Predictiveness and SOA. Pressed P errors occurred when the target appeared over a NP cue and Pressed NP errors occurred when the target appeared over a P cue. Errors bars represent standard error of the paired difference scores.

4.1.3. Discussion

The aim of Experiment 8 was to investigate what effect blocking, rather than intermixing, the associative learning and dot probe tasks would have on the attentional allocation towards the P and NP cues during the dot probe phase. Participants learned to categorise pairs of gratings, with one grating in each pair predictive of the categorisation response and the other non-predictive. The same gratings were then used as uninformative spatial cues to a probe stimulus in a dot probe task, from which we obtained behavioural (RTs and errors) and electrophysiological measures (N2pc) thought to reflect attention. At short SOAs we expected participants to respond faster to targets that appeared over P cues compared to NP cues. RTs after the longer SOAs could help determine whether the early attentional effects were automatic or strategic. If participants were strategically shifting attention towards the P cues during the dot probe task, then we expected the RT bias to increase with SOA. However, if participants were automatically shifting their attention towards the P cues, then the RT bias should diminish at the longer SOAs (Le Pelley et al., 2013 Le Pelley et al., 2016). Over the course of a dot probe block we also expected the RT facilitation effect in favour of the P cue to reduce because each dot probe trial acted as an extinction trial for the learnt attention towards the P cue.

RTs were significantly faster when the target appeared over a P cue compared to an NP cue. However, the RT results did not show an interaction between Predictiveness and SOA, suggesting that the RT bias in favour of the P cue did not vary as a function SOA (Figure 4.6). Therefore, it appears that the strategic processing of the P cues that occurs when the tasks are intermixed (Experiments 4-7) is reduced when the tasks are blocked. More generally, this result suggests that intermixing the two tasks strongly encourages strategic processing of the P cues. In an intermixed design the cues are task relevant on every dot probe trial because the participant has to use the cues to subsequently make a categorisation response. Therefore, the significant Predictiveness×SOA interaction on RTs observed when the tasks are intermixed could indicate that participants are using one of the cues to solve the categorisation task. In the blocked design, no categorisation response is required during the dot probe task and the Predictiveness×SOA interaction is eliminated. In addition, the blocked design resulted in a smaller main effect of Predictiveness on RTs (d_{unb} = .11, S.E. = 0.019) compared to experiments that intermixed the tasks (Experiment 4: d_{unb} = 0.38, S.E. = 0.044; Experiment 5: d_{unb} = 0.23, S.E. = 0.042; Experiment 6: d_{unb} = 0.42, S.E. = 0.083; Experiment 7: d_{unb} = 0.35, S.E. = 0.059).[15] This suggests that when the tasks are blocked there is an overall reduction of attention towards the P cues.

It should be noted that the RT data obtained at different SOAs does not by itself provide enough information to label the attentional processes as automatic or strategic.

---

[15] For a meta-analysis of RT data across all experiments, please see General Discussion in chapter 5.

Although automatic processes tend to be rapid, it is possible for a process to be controlled and relatively fast or automatic and relatively slow. For example, fast eye movements that are top-down in origin have been reported (Hollingworth, Matsukura & Luck, 2013; Gaspelin, Leonard, & Luck, 2013). A hallmark of automatic behaviour is that it is shown to go against one's own top-down goals (Moors & De Houwer, 2006). Therefore, a stronger test of the automaticity of a behaviour is provided by tasks in which performing the behaviour is counterproductive to the task being conducted (Luque, Molinero, Jevtovic & Beesley, 2020).

One reason why attention to the cues might be lower when the tasks are blocked is that each dot probe trial acts as an extinction trial for any learned bias in attention towards the cues. Therefore, as the dot probe block progresses perhaps participants pay less attention to the now task irrelevant cues. However, our extinction RT analysis did not show evidence of extinction across the dot probe trials. Although the analysis showed that participants slowed down as the block progressed, the small RT facilitation effect in favour of the P cue remained fairly constant from the start of a block to the end (Figure 4.7). Overall then, the RT data suggest that the facilitation effect was small during the dot probe block, and this small effect did not extinguish over the course of a block.

It is not immediately clear how to interpret the small residual RT facilitation effects in relation to the automaticity of attention. On the one hand, the RT facilitation effect in favour of the P cue did not increase with SOA, which could suggest that these small RT effects are due to automatic rather than strategic attentional processes. On the other hand, there was no evidence of a decrease in RT facilitation across SOAs. Similarly to the current experiment, Le Pelley et al. (2013; Experiment 2) blocked their dot probe localisation task and found that participants were faster to respond to probes that appeared over P cues after a 250 ms SOA. The numerical advantage in favour of the P cue was small (5 ms) and similar to the main effect of Predictiveness found in the current experiment. However, after a 1000 ms SOA, Le

Pelley et al. found no difference between RTs to the probe over the P or NP cues.[16] The authors argued that the long SOA gave participants time to use controlled processes to reorient attention back to central fixation, thus correcting for the early automatic attentional orienting towards the P cue. In contrast, the current experiment found that the small RT advantage in favour of the P cue was consistent across the SOA range tested (0-500 ms). Similarly to Le Pelley et al., our participants were told to ignore the cues during the dot probe task in order to respond as fast as possible to the probe. If the small RT effects found in the current experiment are due to automatic attention, it is possible that more time is required for participants to override the automatic tendency to attend to the P cue and return attention back to fixation. However, it is difficult to say whether the residual RT facilitation effect observed in this experiment is due to purely automatic attentional processes.

The probe locked N2pc amplitudes did not differ as a function of Predictiveness (Figure 4.8). Experiments 4, 5, 6 and 7 showed that when the dot probe task is intermixed with the categorisation task the probe elicits different N2pc amplitudes depending on whether it appeared over a P or NP cue and that these differences are moderated by SOA (although the interaction was not replicated in Experiment 7). Therefore, the current N2pc results might suggest that attention towards the P and NP cues was evenly distributed across the cue display period. If this is the case, then perhaps the small RT facilitation effect in favour of the P cue was not due to differential attention between the cues, but for some other non-attentional reason. For example, the localisation version of the dot probe task requires participants to make a left or right response to the location of the probe. This response could be contaminated by a non-attentional response bias that involves selecting a response that is congruent with the position of the P cue. The dot probe error data suggest that such a

---

[16] Please note that similarly to the current experiment, Experiment 2 of Le Pelley et al. (2013) manipulated SOA within subjects.

response bias was present because participants made more dot probe errors when the probe appeared over the NP cue compared to the P cue (Figure 4.9). In other words, participants were more likely to select a response (left or right) that was congruent with the P cue even when the target had already appeared over the NP cue, rather than vice versa. This result is similar to the error data reported in Experiments 4 and 5 when the categorisation and dot probe tasks were intermixed (see General Discussion in Chapter 5 for a meta-analysis of the error data). Therefore, it appears that the localisation version of the dot probe task produces a response bias in the direction of the P cue irrespective of whether an intermixed or blocked design is used.

So far the experiments presented in this thesis have investigated how learning influences attention. However, attentional theories of associative learning propose a reciprocal relationship between attention and learning. That is, learning influences attention and the resulting changes in attention can influence subsequent learning (Easdale et al., 2019). Therefore, the aim of Experiment 9 was to investigate whether the attention effects observed in the previous experiments can influence the learning of new cue-outcome relationships.

4.3. Experiment 9

Attentional learning theories propose a reciprocal relationship between attention and learning, such that prior learning influences attention *and* attention influences subsequent learning (Easdale, et al., 2019). However, the attentional effects shown in our N2pc data have not been related to subsequent learning. At present, our N2pc data suggest that P cues are preferentially processed (Experiments 1-7), and after that processing they are inhibited (Experiments 4-7). Experiment 9 aims to test whether the early and late N2pc effects observed in the previous experiments can influence subsequent cue-outcome learning. This was done by modifying the dot probe task to include two different dot probe targets (i.e., outcomes). The targets were a white square and a white diamond (i.e., square rotated 45°). After short SOAs (i.e., the early facilitation time window) only one target (e.g., the white square) was presented over the P and NP cues. After the longer SOAs (i.e., the late inhibition time window) the other target (e.g., the white diamond) was presented. Both targets appeared equally often over both of the cues. If these N2pc measures reflect a processing advantage in favour of the P and NP cues then one would expect participants to better notice the squares over the P cues and the diamonds over the NP cues. This may lead participants to associate the squares with the P cues and the diamonds with the NP cues. That is, we expected learning of the P cue-target associations to proceed more readily for targets presented during the early facilitation window, and conversely, we expected learning of the NP cue-target associations to be processed more readily for targets presented during the late inhibition window.

4.3.1. Method

Participants

Twenty-three participants (5 male; mean age 20 years, age range 18 to 25 years) took part in the experiment. All participants reported normal or corrected-to-normal vision. Participants gave written informed consent and were compensated with credit for an undergraduate psychology course.

Apparatus

All apparatus were identical to Experiments 4 and 5.

Stimuli

All stimuli were identical to those used in Experiments 4 and 5 with one exception. On half of the dot probe trials the white square target was rotated 45° to appear as a white diamond.

Procedure

Participants were alternately assigned to the conditions “Vertical-Horizontal Predictive” (n=12) and “Left-Right Predictive” (n=11). In each experiment, six participants in the “Vertical-Horizontal Predictive” and six participants in the “Left-Right Predictive” condition were shown blue vertical gratings, green horizontal gratings, orange right gratings and pink left gratings. The remaining participants were shown orange vertical gratings, pink horizontal gratings, blue right gratings and green left gratings.

The procedure was the same as Experiments 4 and 5 except that we had more participants complete fewer trials. Specifically, each participant completed 10 blocks of the dot probe task (i.e., 960 dot probe trials). Therefore, there were 80 trials per subject for each of the 12 within-subject conditions (collapsing over target type). In addition, after short SOAs (0, 100 and 200 ms) the dot probe was a white square, whereas after long SOAs (300, 400, and 500 ms) the dot probe was a white diamond (counterbalanced). Participants were told that the probe would be both a white square and white diamond, but were asked to respond the same way (i.e., left or right) to both. At the end of the experiment, participants were asked to rate how often they thought the white square and the white diamond appeared over each of the cues. Participants were not told that the target was equally likely to appear on the left or right, regardless of cue type.

EEG Recording and Analysis

Similarly to the previous experiments, significant differences in N2pc latencies were found between SOA conditions (see Figure 4.11, bottom left). There was a significant main effect of SOA on N2pc latencies at PO7/PO8 ($F_{(1.94, 29.16)} = 73.53$, $p < .001$, $\eta^2 = .703$, MSE = 604.64), P7/P8 ($F_{(1.75, 26.33)} = 56.84$, $p < .001$, $\eta^2 = .65$, MSE = 815.60) and O1/O2 ($F_{(1.69, 23.62)} = 99.09$, $p < .001$, $\eta^2 = .71$, MSE = 441.88) electrode sites. Therefore, the mean amplitude of the N2pc component elicited by the target was measured during different 50 ms time windows centred on the median latencies for each SOA. The median latencies for each SOA measured at the PO7/PO8, P7/P8 and O1/O2 electrode sites can be found in Appendix F. All other EEG recording and ERP analyses were identical to Experiments 4 and 5.

Statistical Analysis

The effects of predictiveness (P and NP) and SOA (short and long) on rating data was compared using a two-way repeated measures ANOVA. All other statistical analyses were identical to Experiment 4 and 5.

4.3.2. Results

Figure 4.10 shows the average of the median RTs to the target in the dot probe task as a function of Predictiveness and SOA. The main effect of Predictiveness on RTs was significant ($F_{(1, 22)} = 20.43$, $p < .001$, $\eta^2 = .057$, MSE = 492.35, $d_{unb} = .23$) as participants responded faster to the target when it appeared over a P cue (M = 432 ms) compared to a NP cue (M = 455 ms). There was a main effect of SOA as participants were slower to respond to the target when it appeared after short SOAs compared to long SOAs ($F_{(1.74, 38.23)} = 205.5$, $p < .001$, $\eta^2 = .75$, MSE = 1416.68). There was also a significant Predictiveness×SOA interaction ($F_{(5, 110)} = 8.94$, $p < .001$, $\eta^2 = .015$, MSE = 222.23). The Predictiveness×SOA interaction was best by a linear trend ($F_{(1, 22)} = 233.21$, $p < .001$, $\eta^2 = .84$, MSE = 414.16) such that the difference in RT to the target over the P and NP cues increased proportionally with SOA. Paired sample t-tests revealed that participants were significantly faster to respond to the target when it appeared over a P cue compared to a NP cue at each SOA (largest p = .047 at the 0 ms SOA), except for the 100 ms SOA (($t_{(22)} = 1.44$, $p = .17$).

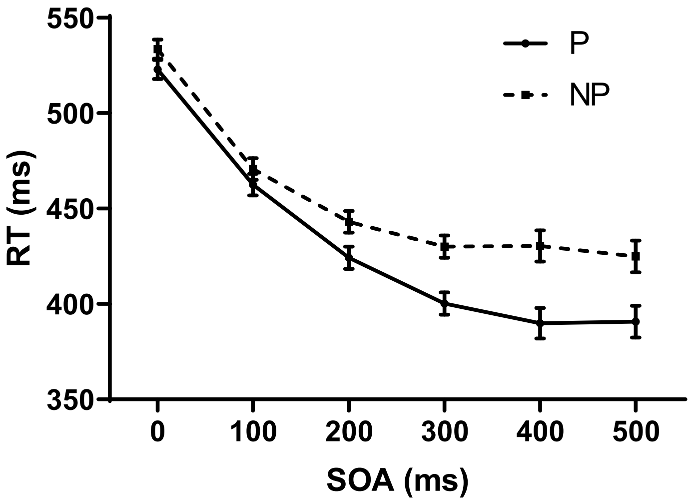

Figure 4.10: Median reaction time to the dot probe target as a function of Predictiveness and SOA. Error bars represent standard error of the paired difference scores.

Figure 4.11 (top row) shows contralateral minus ipsilateral difference waveforms for each SOA when the target appeared over a P cue (left) or a NP cue (right). The waveforms recorded at contralateral and ipsilateral electrodes can be found in Appendix F. Figure 4.11 (bottom right) shows the average of the N2pc mean amplitudes evoked by the target as a function of Predictiveness and SOA. There was a main effect of SOA ($F_{(3.25, 71.41)} = 11.73$, $p < .001$, $\eta^2 = 0.191$, MSE = 2.34) because targets that appeared after long SOAs elicited larger (i.e., more negative) amplitudes compared to targets that appeared after short SOAs. The main effect of Predictiveness was not significant ($F_{(1, 22)} = 3.08$, $p = .093$, $\eta^2 = 0.005$, MSE = .76, $d_{unb} = .099$), nor was the Predictiveness×SOA interaction ($F_{(2.60, 57.10)} = 1.36$, $p = .27$, $\eta^2 = .024$, MSE = 3.22).

The ANOVA on the rating data did not show a significant main effect of SOA ($F_{(1, 22)} = .212$, $p = .65$, η^2 = ., MSE = 1.28, $d_{unb} = .13$) or Predictiveness ($F_{(1, 22)} = 1.78$, $p = .20$, η^2 = ., MSE = .61, $d_{unb} = .30$), nor a Predictiveness×SOA interaction ($F_{(1, 22)} < .001$, $p = .99$, $\eta^2 < .001$, MSE = .478). Paired sample t-tests revealed no difference in ratings to the targets over the P cues (M = 3.6) and the NP cues (M = 3.4) at the short SOAs ($t_{(22)} = 1.02$, $p = .318$, $d_{unb} = .23$) or to targets over the P cues (M = 3.5) and the NP cues (M = 3.3) at the long SOAs ($t_{(22)} = .98$, $p = .338$, $d_{unb} = .21$).

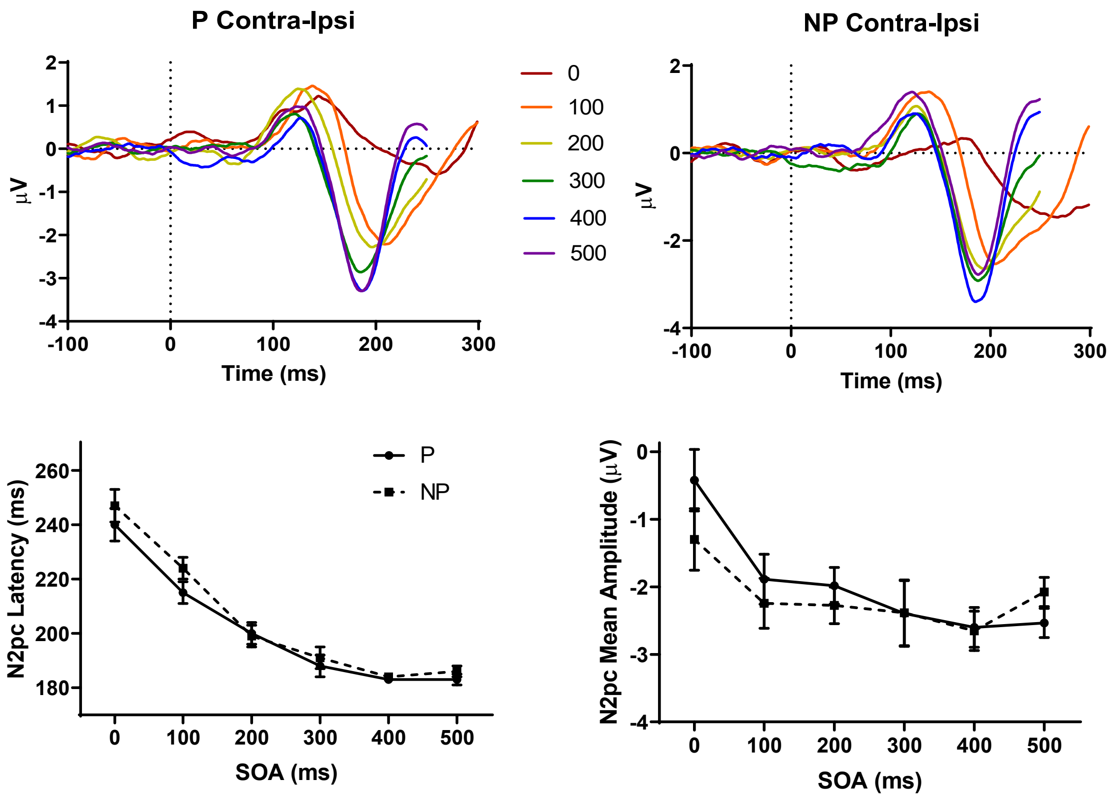

Figure 4.11: The top row shows contralateral minus ipsilateral difference waves for each SOA when the probe appeared over a P cue (left) or a NP cue (right). The bottom row shows median N2pc latencies (left) and N2pc mean amplitudes (right) as function of predictiveness and SOA. All data were averaged over PO7/PO8, P7/P8, and O1/O2 electrodes. Error bars represent standard error of the paired difference scores.

4.3.3. Discussion

In Experiment 9, we attempted to test whether the N2pc effects observed in the previous experiments could bias learning towards the P and NP cues. We presented two different targets (a square and a diamond) over the P and NP cues such that one target appeared after short delays and the other appeared after longer delays. We expected the target that appeared after short delays to be associated with P cues and for the target that appeared after long delays to be associated with NP cues. We measured these associations by having participants rate how often they noticed each target appear over each of the cues. However, the rating data for each cue did not show any differences with respect to target type. If anything, the ratings were in the direction of the classic learned predictiveness effect. That is, for both targets there was a small numerical advantage in favour of the P cues compared to the NP cues. Therefore, we did not obtain any evidence to demonstrate that the N2pc measures observed in our previous experiments can influence subsequent learning.

It should be noted, however, that the rating data is a fairly course measure of learning that is vulnerable to top-down rational biases (Mitchell et al., 2012). Methodological factors could also be contributing to the null rating results. For example, the dot probe targets were on the screen until a left or right response was made. Therefore, even if attention was not congruent with target location after the SOA had elapsed, participants had time to attend to the target and to the cue it appeared over before responding. This could have reduced any potential for biased learning that may have occurred if the targets were briefly flashed. In addition, participants were not explicitly told to try and learn the cue-target relationships, and so only implicit learning (i.e., learning without conscious awareness; Martin-Pichora, Mankovsky-Arnold, & Katz, 2011) may have taken place. In addition, the P and NP cues were both uninformative with respect to dot probe location. Therefore, the null rating results may simply reflect the fact that participants learned the uncertain relationship between cues and outcomes. Once it was evident to participants that there was nothing to learn (i.e., expected uncertainty; Easdale et al., 2019) any attempt to learn cue-outcome relationships would be further reduced. In a typical learned predictiveness design, there is an explicit requirement for participants to learn the cue-outcome relationships and both the previously P and NP cues are informative of their Phase 2 outcomes. Therefore, participants are required to attend to both the cues and the outcomes in order to solve Phase 2.

For the dot probe task, one way to address these issues would be to instruct participants to try and learn cue-target relationships and to use fewer trials so that participants do not learn the uncertain relationship between cues and outcomes. Additional ideas on how to test whether the time course of attention can influence subsequent learning will be discussed in the section on future experiments presented in the next chapter.

There was a RT advantage in favour of the P cues at each SOA, which increased proportionally with SOA (Figure, 4.10). Therefore, the RT results found in Experiments 4-7 were replicated once again. However, the N2pc results observed in the previous experiments were not replicated because there was no interaction between cue type and SOA (Figure 4.11, bottom). To speculate, one possible reason for the null N2pc results could be related to the number trials each participant completed per condition. Particularly relevant to ERP research is the balance between statistical power and experiment length. A disadvantage of the ERP technique is that it requires a large number of trials (Luck, 2014) and it is recommended to collect as many trials from each participant as possible (Woodman, 2010). In Experiment 9, we collected data from twenty-three participants and each one completed 80 trials per within subject condition. In Experiments 4 (N = 14) and 5 (N = 16) each participant completed 144 trials per condition. Therefore, it is possible that the reduction in power that occurs by reducing the number of trials each participant completes cannot be fully recovered by simply increasing the number of participants. However, this is speculation and may not fully explain the failure to replicate. Experiment 7, for example, only partially replicated the N2pc results seen in Experiments 4 and 5, and used an almost identical design. Nevertheless, the number of trials each participant completes per condition may be an important factor to consider when trying to obtain significant ERP results (Boudewn, Luck, Farrens, & Kappenman, 2018).

4.4. General Discussion

Experiment 7 of the current chapter aimed to investigate why previous dot probe research has not found evidence of strategic processing. In Experiments 4-6 (Chapter 3) the RTs from the dot probe task showed a clear interaction between Predictiveness and SOA, such that the RT facilitation effect in favour of the P cues increased proportionally with SOA. Using SOAs of 250 and 1000 ms, previous dot probe research (Le Pelley et al., 2013 and Cobos et al., 2018) has found that the early facilitation effect is significantly reduced after the longer SOA (Le Pelley et al., 2016). The experiments presented in Chapter 3 tested a broader and earlier range of SOAs (e.g., 0, 100, 200, 300, 400, and 500 ms) compared to previous dot probe research. It was possible, therefore, that the 1000 ms SOA used by previous studies was too late to observe the strategic RT effects reported in Chapter 3. To test this, Experiment 7 included an additional SOA of 1000 ms in addition to the earlier SOAs (0-500 ms) tested. The results showed a RT facilitation effect in the direction of the P cue at each SOA. Furthermore, the RT effect increased with SOA. Although, the RT effect at the 1000 ms SOA was slightly reduced compared to the 300-500 ms SOAs, it was clearly present and statistically significant. Therefore, the choice of SOA cannot be the reason why previous dot probe studies have not found evidence of strategic processing in RTs.

Another way to explain why our RT results have differed from previous dot probe studies relates to methodological differences. For example, it was possible that the response set up in our experiments encouraged a strategic bias towards the P cues because it was easy for participants to strategically respond during the intermixed categorisation and dot probe task. Indeed, the RT results of Experiment 7 showed, once again, that participants were preparing their categorisation and dot probe responses together (Figure, 4.2). One problem with this explanation was that participants in Experiments 1 and 2 (Chapter 2) used an identical response set up and were also responding strategically (Figure 4.4), but no evidence of strategic processing in RTs was found. However, the effect size of our strategic response metric (i.e., the difference between Phase 1 categorisation RTs and Phase 2 categorisation RTs) was smaller during Experiments 1 and 2 compared to Experiments 4-7. Therefore, participants in Experiments 1 and 2 seemed better able to follow instruction and to ignore the cues before responding to the dot probe target. One reason why strategic responding was reduced in Experiments 1 and 2 could be related to the number of within subject conditions they had to complete. In Experiments 4-7 participants completed at least six SOAs, whereas those in Experiments 1 and 2 completed only two. A longer experiment with multiple randomly interleaved SOAs could have strongly encouraged participants in Experiments 4-7 to strategically process the cues because they felt this was the fastest way to get through the experiment.

Interestingly, in their intermixed tasks, Le Pelley et al. (2013) and Cobos et al. (2018) manipulated SOA in a between-subjects design. This would have kept the experiment length short and may have reduced strategic responding because participants knew the amount of time they had before dot probe onset. For example, participants in their 1000 ms SOA condition might quickly learn that they have a reasonable amount of time after cue onset before the target appears. Participants could then use this time to prepare a categorisation response and/or correct for the attentional bias towards the P cue by returning attention back to fixation. Our participants were thus exposed to multiple SOAs and were thus more uncertain about the time of probe onset and so their best strategy may have been to wait for the probe before deciding on their categorisation response. Overall then, a combination of experiment length, and within- and between-subject factors might be needed to explain when strategic processing of the cues takes place.

The intermixed dot probe and categorisation task could also encourage strategic processing of the P cues because these cues are task relevant on every dot probe trial. Using an SOA range of 0-500 ms our intermixed tasks have shown evidence of strategic processing in the RT data. Therefore, in Experiment 8, we thought it might be informative to use the same SOA range in a blocked design to test whether P cues that have no task relevance are strategically or automatically processed. The RT results showed a small facilitation effect in the direction of the P cues. However, this RT facilitation effect remained constant across all SOAs (Figure 4.6). This result was difficult to interpret in relation to the automaticity of attention. The fact that the Predictiveness×SOA interaction observed in the intermixed tasks was abolished when the tasks were blocked could suggest that the attention deployed to the P cues was relatively automatic. However, it could be argued that automatic attention should have resulted in a significant reduction in the dot probe bias after the longer SOAs. Indeed, previous dot probe research has argued that attention to the P cue is automatic because RT facilitation has been found at early but not long SOAs (Le Pelley et al., 2013; Le Pelley et al., 2016).

In Experiment 8, two pieces of evidence suggested that the small RT facilitation effect was not due to attentional processes. Firstly, the N2pc measure of attention showed no differences between the P and NP cues (Figure 4.8), suggesting that attention was evenly distributed between the cues during the dot probe task. Secondly, the dot probe error data showed that participants made more errors when the target appeared over a NP cue (Figure 4.9). This could indicate the presence of a non-attentional response bias towards selecting a response that was congruent with the location of the P cue. Note, however, that for participants to select a response that was congruent with P cue, it follows that they would have had to flag, and hence attend to, the location of the P cue. Nevertheless, any attention directed towards the P cues during the dot probe task was clearly small. Importantly, these results suggest that future dot probe studies that rely on RTs to infer the location of covert spatial attention should use a version of the task that controls for response biases. For example, a detection task could be used that requires participants to press only a single button (e.g., spacebar) once they detect the probe (see Experiment 1 of Le Pelley et al., 2013). After response biases are controlled for in this way, any RT effects that are found could more confidently be interpreted as being due to attention.

The main effect of Predictiveness on RTs found in Experiment 8 was smaller compared to Experiments 4-7, which intermixed the dot probe and categorisation tasks (see General Discussion in Chapter 5 for a meta-analysis of RTs). One reason for this could be that attention to the cues diminishes over the course of a dot probe block. However, our extinction analysis showed that the RT facilitation effect in favour of the P cue remained fairly constant across a dot probe block (Figure 4.7). Therefore, any attention that was deployed to the cues was low from the outset of the dot probe block. This may simply be due to participants following the instruction to ignore the cues during the dot probe task. Contextual modulation of attentional deployment could also explain this reduced attention (Uengoer, Pearce, Lachnit, & Koenig, 2018). For example, the change from categorisation to dot probe task, and the corresponding drop in cue relevance, could have acted as a salient context change that may have encouraged a sudden drop in attention to the cues at the outset of a dot probe block.

Experiment 9 aimed to test whether the early facilitation and late inhibition time windows illuminated by the N2pc can influence the learning of new cue-outcome relationships. This is an important next step because attentional learning theories propose a reciprocal relationship between attention and learning. We presented two different targets (square and diamond) over the P and NP cues such that one target appeared after short delays (corresponding to early facilitation) and the other appeared after long delays (corresponding to late inhibition). We expected participants to associate the P and NP cues with the target that appeared after the short and long delays, respectively. However, our rating data did not show any differences with respect to target type. Therefore, we did not obtain evidence to support the idea that our N2pc measures can influence subsequent learning. This null result is more problematic for our inhibition time window because previous learned predictiveness research has consistently found a bias in attention and subsequent learning for previously P cues (Le Pelley et al., 2016).

CHAPTER 5: General Discussion

Summary and Meta-Analyses

The experiments presented in this thesis attempted to measure real-time changes in covert attention using behavioural (RTs, ratings and errors) and EEG (N2pc, SSVEPs) measures. It was hoped that these measurements would provide novel insights into exploitative and explorative attentional mechanisms. We hypothesised that P cues would capture attention early within a trial, consistent with the predictiveness principle embodied in the Mackintosh model (1975). However, later in a trial, after the P cues had been processed, we suspected that they might be inhibited (Klein, 2000). This inhibition could bias attention towards the NP cues, consistent with the uncertainty principle embodied in the Pearce-Hall model (1980). In Chapter 1, we explored different aspects of the intermixed dot probe and categorisation task that was first used by Le Pelley et al. (2013). Experiments 1-3 used SOAs of 250 and 1200 ms to investigate the time course of attention. Using black and white P and NP gratings (Experiment 1) we were unable to elicit any RT effects. Therefore, we were unable to replicate the RT facilitation effect usually seen for targets that appear over P cues after a 250 ms SOA (Le Pelley et al., Luque et al., 2017; Cobos et al., 2018). However, adding colour to the stimuli (Experiment 2) so that participants could use colour information to discriminate between the P and NP cues was sufficient to replicate the effect. The general finding that participants were faster to respond to targets that appeared over P cues was replicated (at least directionally) in all experiments. This can be seen in Figure 5.1, which is a forest plot representing effect size coefficients (d_{unb}) for the main effect of Predictiveness found in the RT data from all Experiments($d = .21$, $z = 5.56$, $p < .001$).[17]

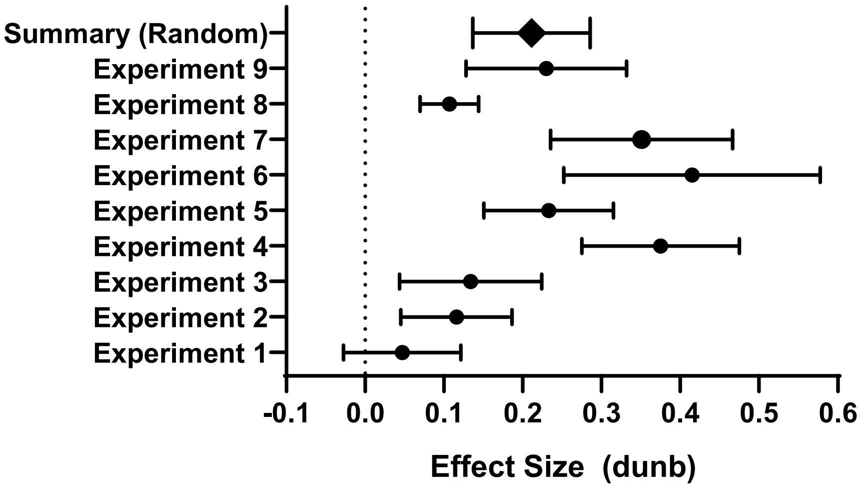

Figure 5.1: Results of a meta-analysis on the main effect of Predictiveness observed in RT data. Positive values indicate a facilitation effect in favour of the P cues.

The N2pc results of Experiments 1-3 showed a novel effect after the 250 ms SOA (although not significant in Experiment 2), such that targets that appeared over NP cues elicited larger N2pc amplitudes compared to targets that appeared over P cues. At first glance, this result appeared to suggest that attention was initially allocated towards the NP cue. However, given our approach involved time locking the N2pc to a probe that appeared over attended stimuli, we reasoned that this effect might actually reflect attention being directed towards the P cue. That is, if attention was directed towards the P cue first, then it would be more difficult to process a dot probe that appeared over the NP cue after a short SOA, which could have resulted in a larger N2pc amplitude (Luck et al., 1997). In addition, Experiment 3 introduced a rating measure. Participants were asked how often they thought the target appeared over the P and NP cues. The ratings indicated that they noticed the target

---

[17] Please note that all meta-analyses were conducted using the metagen package in R. With the exception of the RT analysis (Figure 5.1), all tests of heterogeneity showed no systematic differences across studies. For the RT analyses, a random effects summary is reported. All other meta-analyses report fixed effect summaries.

appear more often over the P cues, a result that was fairly consistent across subsequent experiments. This can be seen in Figure 5.2, which shows effect sizes included in a meta-analysis of the rating data across all experiments (d =0.41, z = 3.68, p < .001).

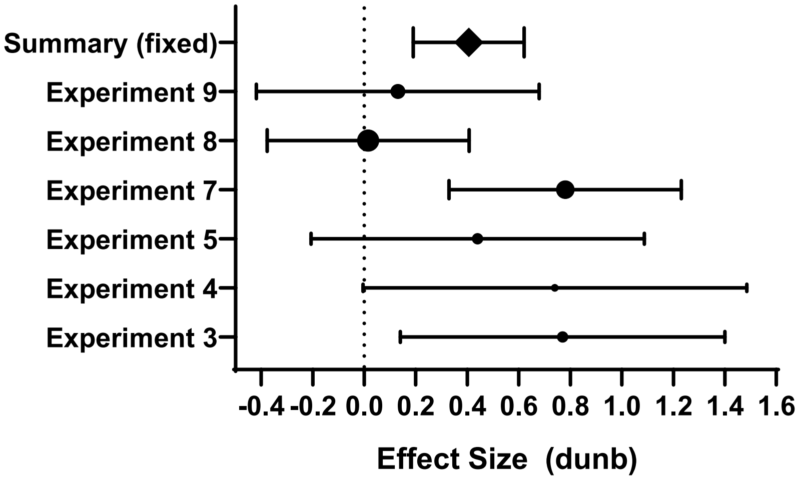

Figure 5.2: Results of a meta-analysis on the main effect of Predictiveness observed in rating data. Positive values indicate higher ratings for P cues. Note that for Experiment 9, the ratings were collapsed across target type.

Using six equally spaced SOAs that ranged from 0-500 ms, Experiments 4 and 5 (Chapter 3) found a novel interaction in RTs. The RT facilitation effect in favour of the P cues increased proportionally with SOA, which suggested that the P cues were being strategically processed (Le Pelley et al., 2013). This result, which was replicated in Experiments 6, 8 and 9, was in contrast to previous dot probe research that has argued that P cues are automatically processed because the RT advantage to P cues at short SOAs is usually reduced at longer SOAs (Le Pelley et al., 2013; Cobos et al., 2018). However, it is important to note that in the intermixed dot probe and categorisation task, the P, but not NP, cues are task relevant on every dot probe trial. This is because participants have to make a categorisation response after every dot probe response. Therefore, it is not surprising that participants were using the P cues strategically during the dot probe phase. A more interesting test to determine whether the learned predictiveness effect is under top-down or bottom up control would be to employ a dot probe task during a Phase 2 categorisation, when both the P and NP cues are task relevant (see Future Experiments). In addition, the RT analysis of these experiments showed that participants were responding strategically during the dot probe phase. Instead of responding as fast as possible to the dot probe, participants let it linger briefly on the screen while they continued to inspect the cues and prepare a categorisation response. This finding helped explain why our RT and N2pc measures of attention showed different patterns. Furthermore, there appeared to be a response bias towards selecting a response that was congruent with the P cues. Participants made more dot probe errors, and premature responses, in the direction of the P cues. Forest plots for the coefficients included in the meta-analysis of dot probe errors ($d = .34$, $z = 4.29$, $p < .001$) and premature responses ($d = .28$, $z = 2.86$, $p < .01$) can be seen in Figure 5.3. The dot probe errors and premature responses measured in the localisation version of the dot probe task suggest that future dot probe studies should control for response biases (e.g., by using a detection task; see Experiment 1 in Le Pelley et al., 2013) to more confidently interpret RT effects as being due to changes in attention.

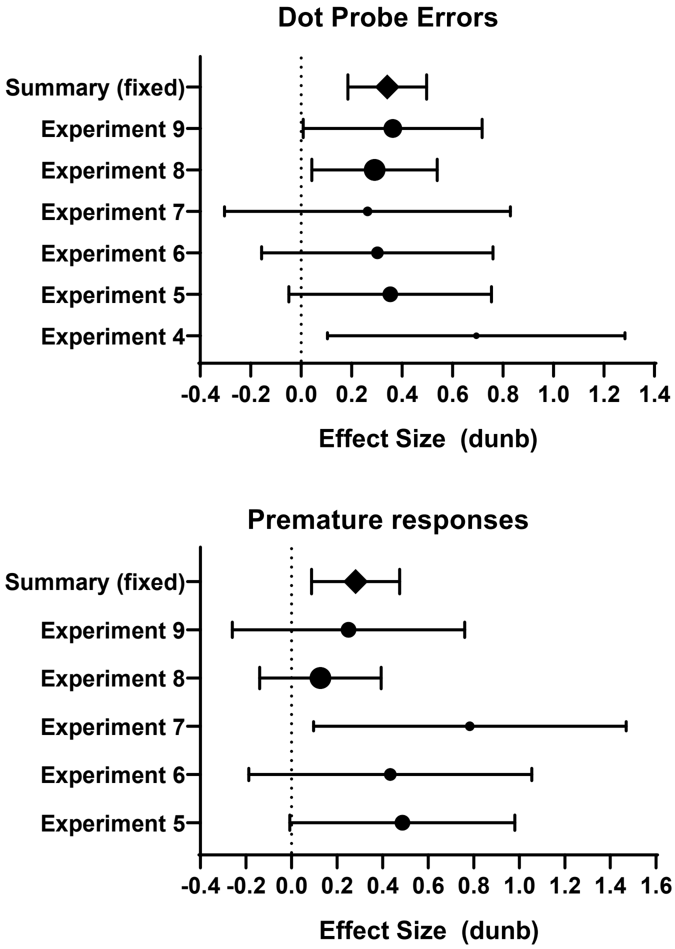

Figure 5.3: Results of a meta-analysis on the main effect of Predictiveness observed in the dot probe error (top) and premature response (bottom) data. Positive values indicate responses in the direction of the P cues.

The N2pc results from Experiments 4 and 5 also showed a novel interaction, but one that suggested an attention shift from P to NP cues. Probes that appeared after short SOAs (100 and 200 ms) elicited larger N2pc amplitudes when they appeared over NP cues compared to P cues, in agreement with the experiments presented in Chapter 2. However, after a longer SOA (400 ms) this pattern reversed. This interaction, which was replicated again in Experiment 6, agreed with our predictions that attention shifts between the P and NP cues in a sequential fashion. However, there was still ambiguity regarding which cue was attended to first. Nevertheless, all of the behavioural data suggested that the larger, and longer lasting, early N2pc effects should be interpreted as attention towards the P cues.

Participants rated that they noticed the targets appear more often over the P cues (Figure 5.2) and they made more dot probe errors and anticipatory responses in the direction of the P cue (Figure 5.3).

Experiment 6 aimed to resolve the N2pc ambiguity by using a conjunction discrimination in the categorisation task. Participants had to use both colour and orientation information of the P gratings to determine the correct categorisation response. We reasoned that this change would require attention to remain with the P cue for longer, which should extend the early N2pc time window. Consistent with this prediction, the probe locked ERP analysis showed that the early N2pc time window was extended because targets that appeared 300 ms after cue onset elicited larger amplitudes over NP cues compared to P cues. In addition, SSVEP and N2pc amplitudes measured during the categorisation task indicated that attention was allocated towards the P cues. Taken together, this pattern of results supports the counterintuitive idea that probes that appear over unattended stimuli elicit larger N2pc amplitudes compared to probes that appear over attended stimuli.

There are some important caveats regarding the interpretation of our EEG data. Firstly, it is possible that motor responses may have contaminated the N2pc data. Originally, we attempted to control for this by removing trials from the EEG analysis if the participant responded to the dot probe faster than 250 ms (i.e., the length of the epoch of interest). However, this approach does not control for motor related ERP components that occur in the lead up to a response. For example, when responding with the right hand a negative potential occurs over the left motor cortex (and vice-versa). Importantly, this lateralised readiness potential (LRP) begins before the overt response (Smulders & Miller, 2013). In the current experiments, the N2pc was measured at parietal/occipital electrode sites, and not at central sites that are closer to motor cortex. Nevertheless, the possibility remains that lateralised motor components (e.g., the LRP) overlapped with our lateralised measure of attention (i.e., the N2pc). Consequently, independent analysis of motor-related ERP components would be required to ensure the validity of our N2pc results. Secondly, the baseline used in our ERP analysis was the 100 ms prior to the dot probe onset. Therefore, each SOA used a different 100 ms baseline period. Although we were only interested in comparing ERPs elicited by probes *within* each SOA condition, it might be worth reanalysing the data using the same 100 ms baseline period (e.g., the 100 ms before the onset of the cues).

Experiments 7-9 (Chapter 4) were designed to investigate the automaticity of attention and whether the early and late N2pc effects could influence subsequent learning. Experiment 7 investigated why previous dot probe research has not been able to find RT effects at long SOAs, which has led to the conclusion that attention shifts towards P cues automatically (Le Pelley et al., 2013). It was possible that the long SOA used in previous studies (e.g., 1000 ms; Le Pelley et al., 2013; Cobos et al., 2018) was too late to capture the RT effects observed when SOAs of 0-500 ms are used. Indeed, Experiments 1 and 2 (Chapter 2) found no RT effects after a 1200 ms SOA. Therefore, in Experiment 7 we included a 1000 ms SOA, in addition to our earlier SOAs, to see if the RT facilitation effect decreased to zero. Contrary to this, we found a significant RT effect in favour of the P cue after 1000 ms, although it was slightly reduced compared to the 300-500 ms SOAs. It appeared, therefore, that the choice of SOA was not the reason why previous studies have not found RT effects at their longer SOAs. We argue that a mixture of between and within-subject factors, including the number of within-subject conditions, are probably contributing to the discrepant results.

Experiment 8 demonstrated that one important factor that future dot probe studies must consider is whether, or when, to block or intermix the dot probe and categorisation tasks in the dot probe phase. Experiment 8 blocked the two tasks, which caused markedly different results in the dot probe phase. Firstly, the interaction between Predictiveness and SOA observed on RTs in Experiments 4-6 was abolished. However, there was a small, but statistically significant, RT advantage in favour of the P cue that was constant across all SOAs. In addition, the RT advantage was consistent across the length of a dot probe block, which suggested that the small RT facilitation effect did not extinguish. These results were difficult to interpret in regard to the automaticity of attention. On the one hand, strategic processing of the P cues was reduced because the RT effect no longer increased with SOA. On the other hand, if the RT effects were due purely to automatic attention, then one would expect the RT effect to diminish at longer SOAs (Le Pelley et al., 2013; Le Pelley et al, 2016). In addition, there was still a response bias in the direction of the P cues, which could suggest that the small RT effects found when the task is blocked are not due to attention. Taken together, the results of Experiment 8 suggested that the context change from categorisation task to dot probe task resulted in a sudden drop in attention towards the cues and that any attention deployed during the dot probe phase was small but directed towards the P cues consistently over the course of a dot probe trial and a dot probe block.

Finally, in Experiment 9 we attempted to test whether the early and late N2pc effects observed in Experiments 4-6 could influence subsequent learning. This was important because attentional learning theories (Mackintosh, 1975; Pearce & Hall, 1980) propose a reciprocal relationship between attention and learning. In other words, learning influences attention and subsequent attention deployments influences new learning (Easdale et al., 2019). However, the results of this experiment did not provide any evidence that the early and late N2pc time windows could result in biased cue-target learning. Note, however, that the time course of attention within a trial might still be an important factor to consider when trying to reconcile exploitative and explorative attentional mechanisms and there might be a better, and more general, way of testing whether this is the case (see Future Experiments).

Over the last 40 years, learning models have attempted to formally capture the rules that govern changes in attention and learning. However, these models have generally overlooked the importance of time and how attention shifts within trials. We attempted to formally model changes in attention during cue presentation to demonstrate that these changes could be captured mathematically. The next section will describe a real-time model to demonstrate how within trial changes in attention could influence subsequent learning. In addition, this model makes a novel prediction regarding the learned predictiveness effect.

Real-time model: The auto-associator

Given that in real life events unfold over time several associative models have tried to account for “real-time” stimulus representations. One such model is the auto-associator (McClelland & Rumelhart, 1988). Unlike other models (e.g., Harris, 2006; McLaren & Mackintosh, 2000; Schmajuk, Lam, & Gray, 1996) the auto-associator does not modulate attention to cues based on prior experience (Baetu & Baker, 2016). However, in this section, we describe a modified version of the model which does take attention into account to demonstrate how real-time changes in attention could influence learning.

The auto-associator models a network of units that represent different stimuli (e.g., cues and outcomes). In the network, the units are connected via unidirectional connections and the strength of each connection is updated as learning takes place. When a stimulus is presented, the activation level of the unit that represents it increases gradually from zero into the positive range. This change in activation level is the real-time component of this model. For example, when a stimulus is turned on (i.e., it is physically present) its activation gradually and continuously changes with time. The model simulates this by dividing time into discrete units called cycles, and the activation level of the stimulus is changed on every cycle by a small amount. If the stimulus is on for an extended period of time then its activation will reach an asymptote. When the stimulus is no longer physically present its activity level gradually decays back to zero. Therefore, although the occurrence of the stimulus is a binary event (i.e., it either does or does not occur), the activity level of its mental representation is graded and changes with time (Baetu & Baker, 2016).

For two stimuli to be associated (i.e., for their connection to be strengthened), the activation levels of the two stimuli need to be non-zero at the same time. For example, if an outcome stimulus occurs long after the activity of a preceding cue has decayed to zero, then the connection between the cue and the outcome will not be strengthened. Connection strength is updated as a function of prediction error and activity level. If the activity level of a cue and an outcome is high, and the prediction error of the outcome is also high (i.e., the outcome is unexpected) then the change in association between the cue and the outcome will be large (i.e., there will be a large change in connection strength). However, if the activity level of the cue and the outcome is high, but the prediction error of the outcome is low (i.e., the outcome is expected), then the change in association between the cue and the outcome will be small.

In addition to sensory input, the activation level of a unit can be increased by other active units in the network that share excitatory associations with it. That is, once a connection has formed between a cue and an outcome the activity level of the outcome will increase when that cue is presented, even if the outcome is not physically present. Therefore, an outcome unit can receive external activation (i.e., when it is physically present) and internal activation from an active cue unit that has been associated with it. When the external activation of an outcome is high (e.g., because it is physically present) and the internal activation of the outcome is low (e.g., because active cues have not been associated with it) then a large prediction error will be generated. Conversely, when an outcome has both high external and internal activation (e.g., because active cues have been associated with it) then prediction error will be small.

In the original auto-associator, when two cues are presented simultaneously their external activation levels increase at the same rate. We adapted the model so that attention to the cues would modulate their external activity. For example, if a P and NP cue are presented together, their external activation levels will rise at different rates depending on which cue is being attended. Given the interaction between predictiveness and SOA on N2pc amplitudes observed in Experiments 4-6, we assumed attention would initially shift towards the P cue (i.e., soon after the cues appeared) but then shift away from the P cue and towards the NP cue. This assumption is in agreement with the behavioural systems framework which suggests that short CS-US intervals result in focal search while long CS-US intervals result in general search (Timberlake, 2001). Therefore, the external activation of a P cue will be boosted by attention early in a trial. However, as the trial progresses attention to the P cue reduces and attention to the NP cue increases. This means that late in a trial the external activity of the NP cue increases at a faster rate compared to the external activity of the P cue. Note that the model has to learn which cues are P and NP, and the attention that both cues receive at the start of training is identical.

Similarly to other hybrid models (Le Pelley, 2004; Pearce & Mackintosh, 2010) our modified version of the auto-associator assigned two attention values to each unit. One represented a Mackintosh-like mechanism (alphaM) and the other a Pearce-Hall-like mechanism (alphaPH). The influences of these two attention parameters shifted as a function of cue duration. In this model, a Mackintosh mechanism plays a more important role at the beginning of cue presentation, and a Pearce-Hall mechanism gradually takes over as cue presentation progresses. For a full description of the model please see Appendix G, which shows the number and types of trials the model was trained on, and formulae used to update learning weights, attention, and activation levels.

Figure 5.4 shows learning weights produced by the model for P and NP cues during Phase 1 of a learned predictiveness design. The top panel shows learning curves produced when cue duration was short (10 cycles), while the middle and bottom panels show data from intermediate (25 cycles) and long (40 cycles) cue durations, respectively. Note that in each condition the outcome was presented for 5 cycles after cue offset. As can be seen in the figure, the model was able to learn which cues were P and NP irrespective of cue duration. Figure 5.5 shows how activation levels of the P and NP cue changed on the last trial of Phase 1. For a given cycle, the activation level is identical across cue duration conditions. For example, over the first 10 cycles the changes in activation are the same across all durations. The only difference across conditions is that the model is given more or less time to change activity levels. For short cue durations (top of Figure 5.5), there is a clear advantage in activity for the P cue compared to the NP cue at cue offset. However, as cue duration increases the discrepancy in activity reduces (middle) and eventually reverses by the end of the longest cue duration (bottom). Therefore, the model was able to simulate attention shifting from the P to the NP cue.

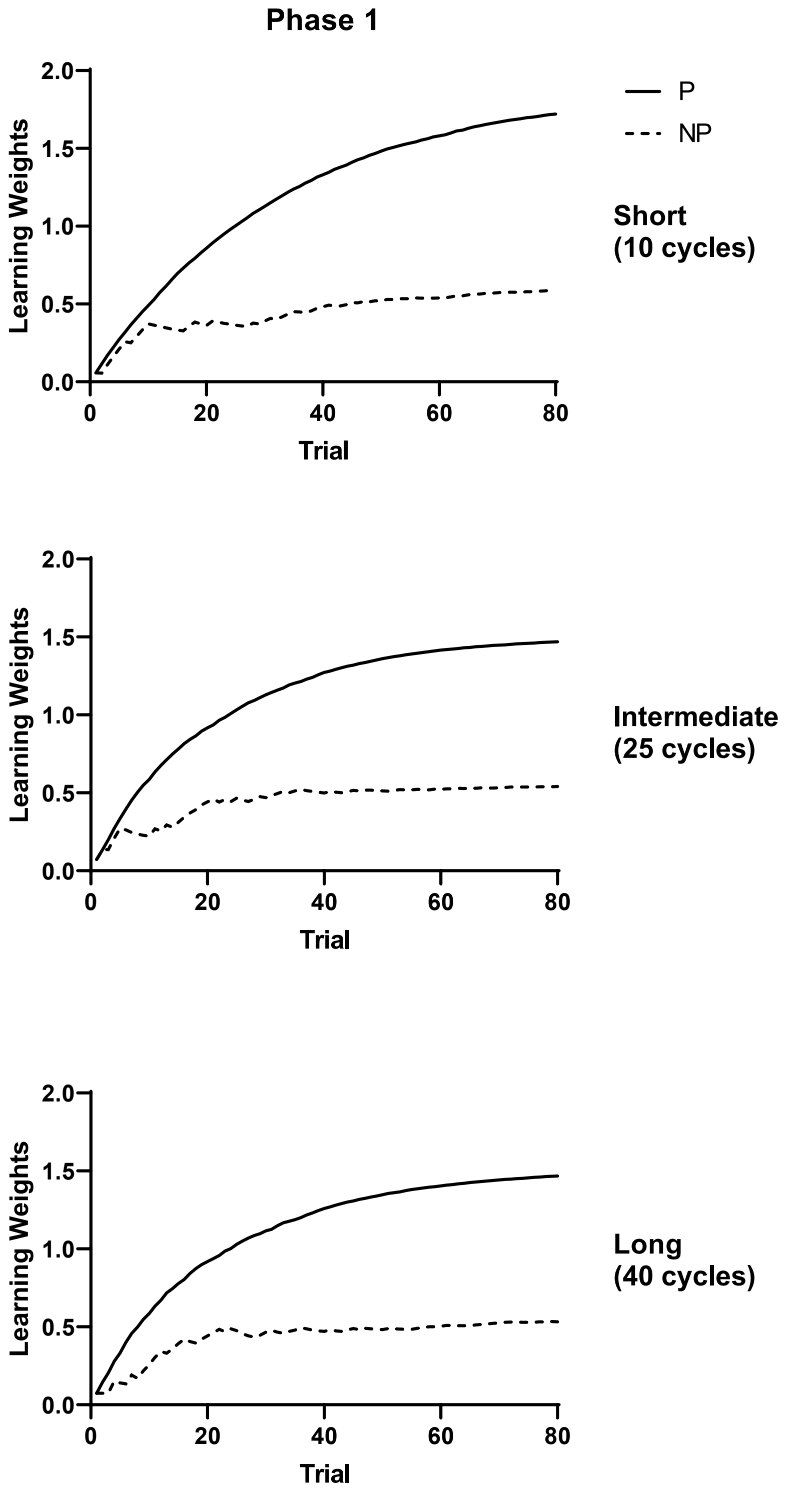

Figure 5.4: Phase 1 learning curves for P and NP cues in the short (top), intermediate (middle) and long (bottom) trial durations.

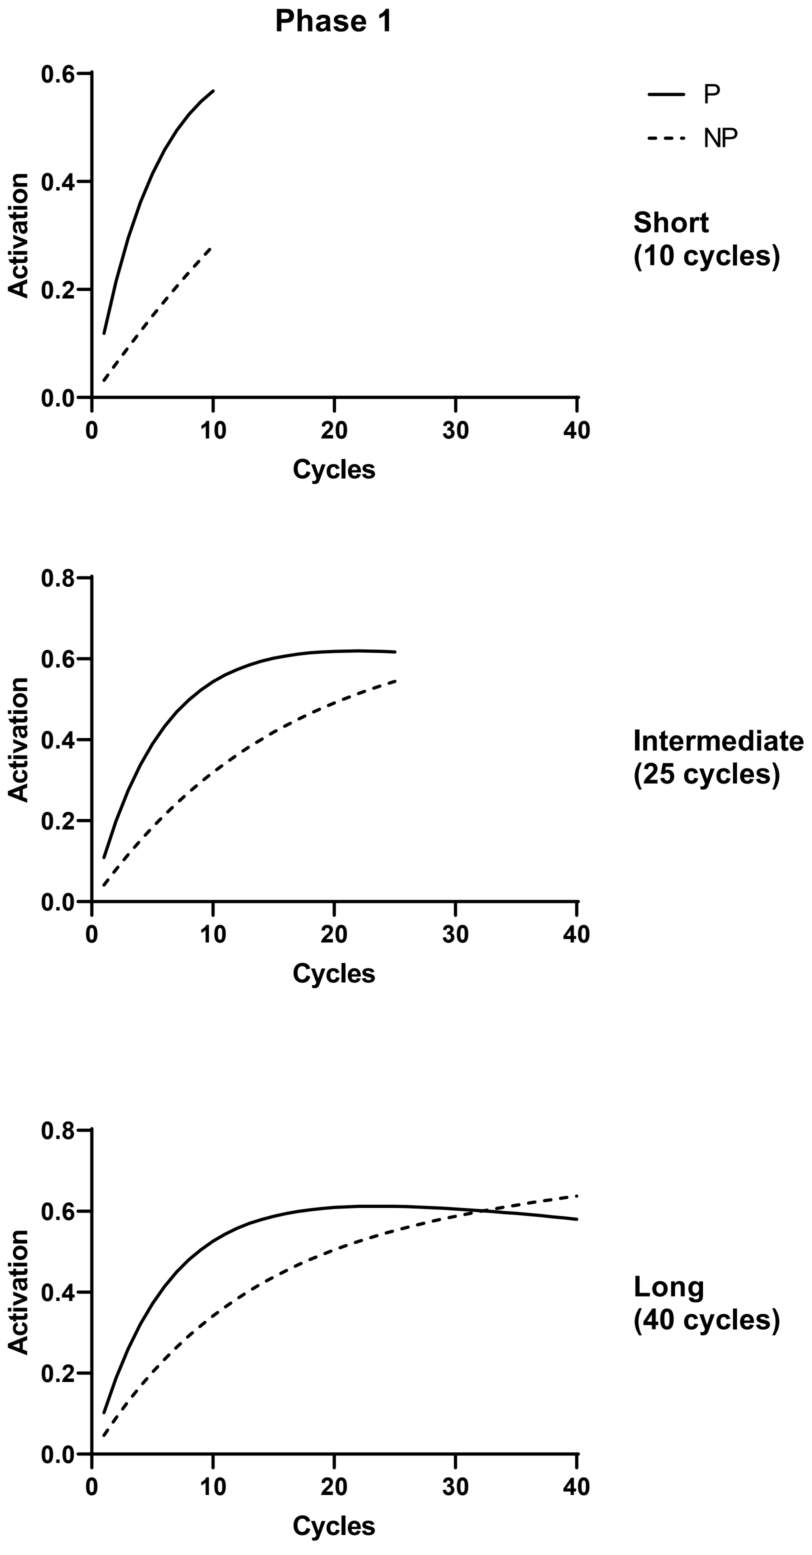

Figure 5.5. Activation of P and NP cues over time on the last trial of Phase 1 for short (top), intermediate (middle) and long (bottom) trial durations. Note that these patterns of activation are similar to the first trial of Phase 2 because the attention parameters associated with each cue transfers from Phase 1 to Phase 2.

Figure 5.6 shows learning weights produced by the model for P and NP cues during Phase 2 of the learned predictiveness design. In this phase both the P and NP cues perfectly predict new outcomes. As can be seen in the top panel, when cue duration was short (10 cycles) there is a clear learned predictiveness effect. That is, learning the P cue-outcome relationship occurs more rapidly compared to NP cue-outcome learning. By the end of Phase 2 the model had learnt to associate the previously P cues more strongly with Phase 2 outcomes compared to the previously NP cues. However, the model produced different Phase 2 learning curves for the intermediate and long cue durations. In these conditions (the bottom two panels of Figure 5.6) the P and NP cues are associated with Phase 2 outcomes more equally. In other words, when cue duration is increased the learned predictiveness effect is abolished.

The pattern of results described above can be explained as follows. At the start of Phase 2, attention shifts between the previously P and NP cues in the same way as shown in Figure 5.5 (i.e., the end of Phase 1). This is because the attention parameters associated with the cues at the end of Phase 1 are inherited at the start of Phase 2. Therefore, at the end of an early Phase 2 trial, P cues are more active than NP cues when the outcome occurs in the short duration condition (Figure 5.5, top). However, when durations are longer, the P and NP cues have similar activation levels by the end of the trial (Figure 5.5, middle and bottom). These differences in activation levels have flow-on effects when Phase 2 learning weights are calculated. The model updates learning weights as a function of prediction error (i.e., external activation of the outcome – internal activation of the outcome) and activation of the cue. Given prediction error is the same for all cues at the start of Phase 2, the activation of the cues will determine how much associative strength is attributed to each cue. When prediction error is high (e.g., at the start of Phase 2) and activation is high (e.g., a cue is present and attended) then the corresponding change in the cue-outcome learning weight will be large.

However, if prediction error is high but activation is low (e.g., because the cue is unattended) the corresponding change in the cue-outcome learning weight will be small. In the short duration condition, the activation of the NP cue is always lower than that of the P cue. This is because early in a trial the Mackintosh attention parameter (alphaM) influences the external input of the CSs more than the Pearce-Hall attention parameter (alphaPH). In other words, the external input of P cues is higher than that of NP cues so P cues are learned about faster. As a consequence of this faster learning is that the P cues are always seen as more predictive during Phase 2 compared to the NP cues. Therefore, at the end of Phase 2 the model continues to preferentially attend to the P cues early in the trial (see top panel of Figure 5.7). Conversely, in the longer duration conditions, the model is given time to attend to the NP cue because the influence of alphaPH on the external input of the CSs has increased, while the influence of alphaM has decreased. Therefore, the activation of the NP cues has become similar to that of the P cues by the end of the trial. Consequently, the model learns equally well about both cues during Phase 2. As both cues are learnt to be predictive in Phase 2, the amount of attention they receive gradually becomes equivalent. By the end of Phase 2 both P and NP cues receive similar amounts of attention (see middle and bottom panels of Figure 5.7). In summary, the model captures within-trial attentional shifts from P to NP cues, and makes a novel prediction regarding the influence of cue duration on the learned-predictiveness effect. That is, a learned-predictiveness effect should occur when Phase 2 cue durations are short, and it should be abolished when Phase 2 cue durations are long.

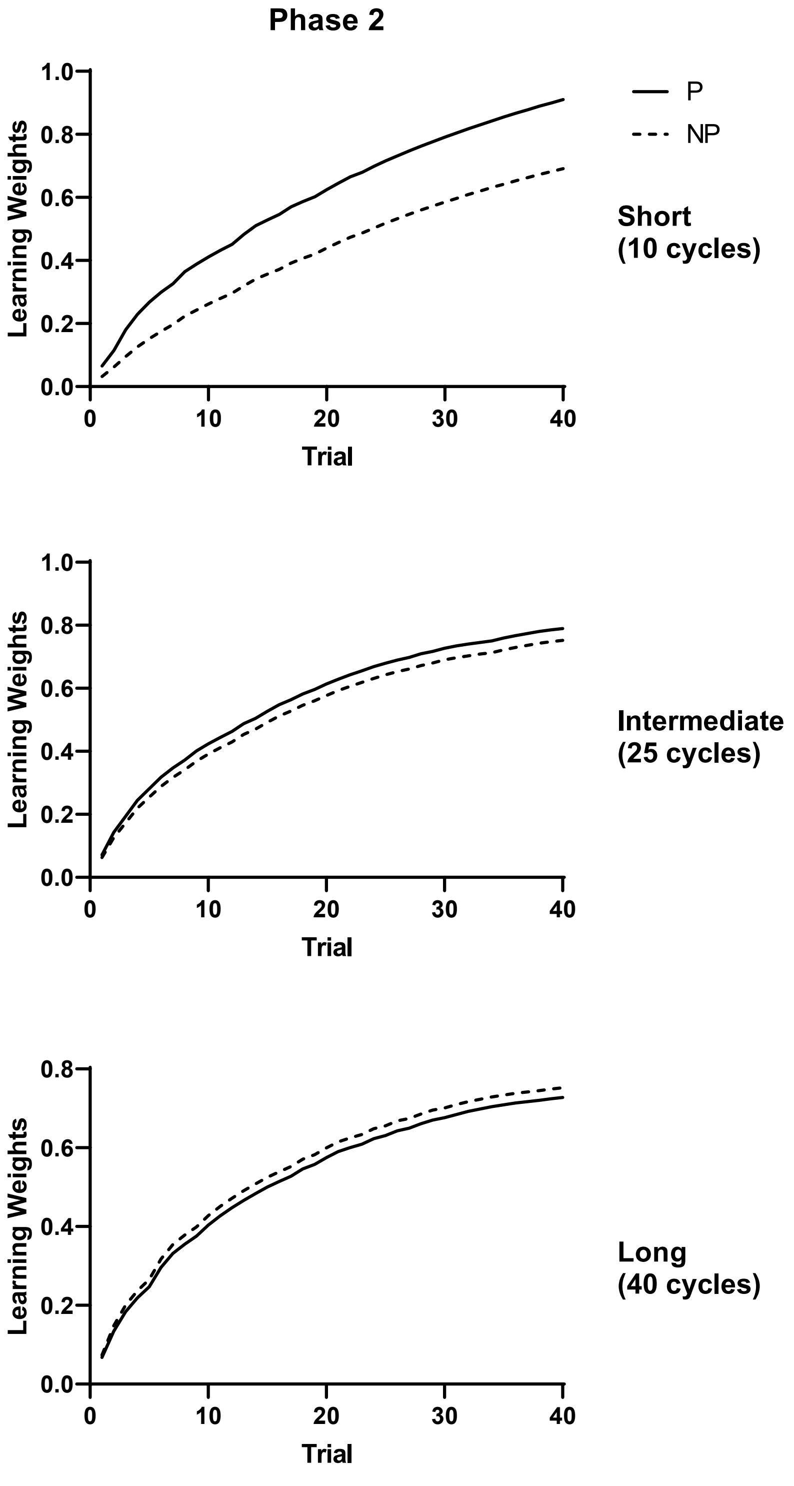

Figure 5.6: Phase 2 learning curves for P and NP cues in the short (top), intermediate (middle) and long (bottom) trial durations.

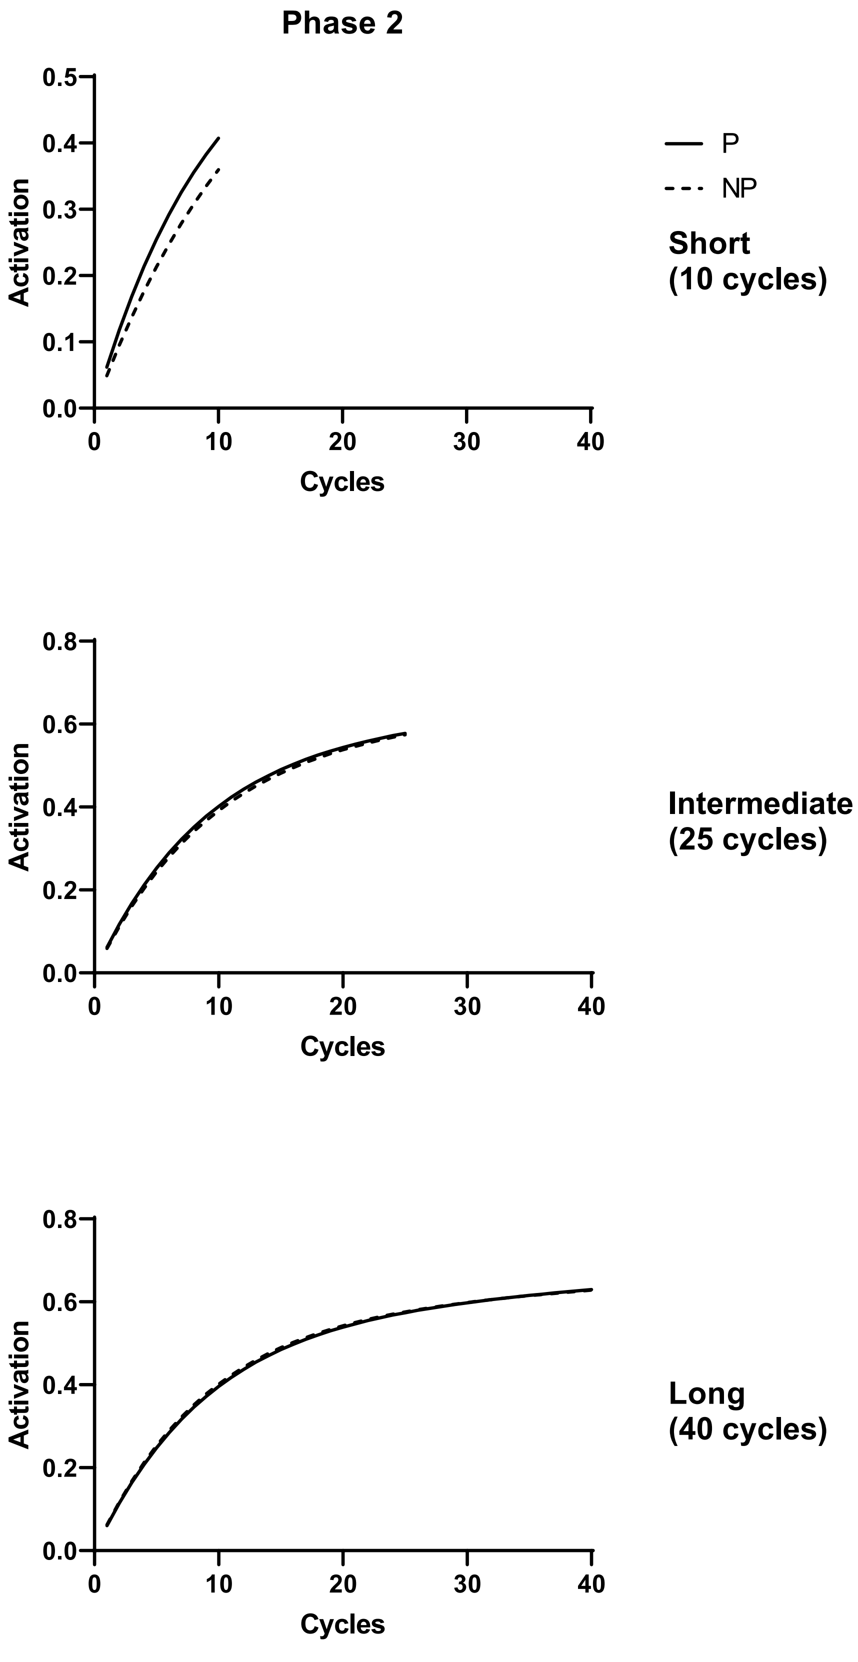

Figure 5.7: Activation of P and NP cues over time on the last trial of Phase 2 for short (top), intermediate (middle) and long (bottom) trial durations.

Future Experiments

This thesis presented one attempt to test whether the temporal dynamics of attention are important for subsequent learning. In Experiment 9, we presented two different dot probe targets (a diamond and a square) after different SOAs (early and late time windows) to see if the N2pc effects observed in the current thesis could be related to subsequent learning. That experiment did not show differences in cue-target learning as a function of SOA. However, there may be other experimental designs that can further investigate whether time is an important variable that can modulate exploitative and explorative behaviour.

For example, as suggested by the simulations presented in the previous section, cue duration could be manipulated in a learned predictiveness design. In Phase 1, participants would learn that some cues are predictive and others are non-predictive of certain outcomes. The inspection time thresholds (e.g., Vickers, Nettelbeck, & Wilson, 1972) that each participant requires to accurately complete the discrimination could also be measured. In Phase 2, participants would learn new cue-outcome relationships, with both P and NP cues now perfect predictors. Importantly, during Phase 2 half of the participants would be presented with short cue durations, while the other half would have long durations. The length of these durations could be a percentage of the inspection time thresholds calculated in Phase 1. Participants in both groups would not be shown the outcome until their cue duration had expired. If attention preferentially shifts to the P cues, then those in the short duration group should learn cue-outcome relationships better for the P cues. Therefore, these participants should show a learned predictiveness effect. Conversely, those in the long duration group would have time to correct for this bias and explore the NP cues. These participants should learn cue-outcome relationships equally well for both the P and NP cues and should not show a learned predictiveness effect. This would be evidence in favour of the idea that time is important for exploration.

All experiments presented in this thesis measured attention after learning. In fact, the majority of experiments measured attention while participants were performing a categorisation task in which only the P cue was task relevant. Therefore, there was never an opportunity for participants to learn about the NP cues. A logical next step would be to measure participants' attention while they are performing a subsequent intermixed dot-probe and categorisation task in which both previously P and NP cues are predictive of the categorisation response (i.e., analogous to Phase 2 of the learned predictiveness design). This would provide participants an opportunity to make NP cue-outcome connections. For example, in Phase 1 participants could be presented with cues on the vertical midline of a monitor and these cues would be P and NP of a left or right categorisation response. Participants would then complete a subsequent categorisation task in which the cues are presented horizontally. In this subsequent phase both the previously P and NP cues would be predictive of an up or down response. Participants would then complete an intermixed categorisation and dot probe task using the Phase 2 discrimination. If participants are continuing to use the previously P cues to solve the Phase 2 categorisation, then the same Predictiveness×SOA interaction observed on RTs in our experiments would be expected (Figure 5.8, top). If, however, the participants are using both the P and NP cues to solve this categorisation task, then there should be no difference in RTs between P and NP cues as a function of SOA (Figure 5.8, bottom). Note that a control experiment using novel cues could be conducted to determine the pattern of RT data when participants can use either one of the two cues to solve the categorisation response (i.e., Phase 2 without a P and NP bias).

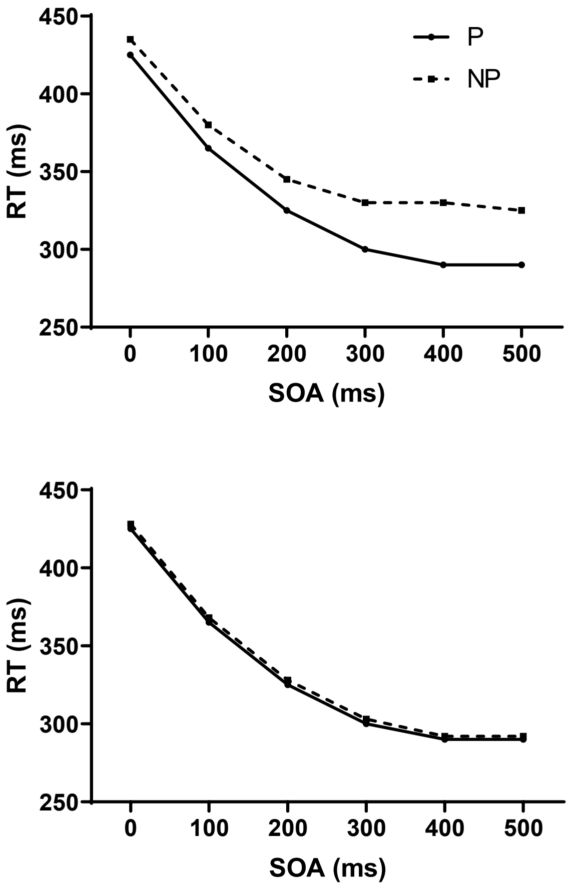

Figure 5.8: Expected patterns of data for an intermixed dot-probe and categorisation task in which both previously P and NP cues are predictive of the categorisation response: Predictiveness×SOA interaction which indicates that the participant is using the P cue to solve the categorisation task (top). No Predictiveness×SOA interaction, and no main effect of Predictiveness, which could indicate that the participant is using both the P and NP cues to solve the categorisation task (bottom).

Essentially, this design uses the presence or absence of the Predictiveness×SOA interaction (Figure 5.8,) to infer what cue(s) the participants are using (covertly) to solve Phase 2 of a learned predictiveness design. Such an experiment would be a nice addition to the literature that has shown that participants continue to look at P cues more often in Phase 2 even though both the P and NP cues are now task relevant (i.e., overt attention, for a review see Le Pelley et al, 2016). The eye gaze literature would predict that participants continue to use the P cues to solve the Phase 2 categorisation task (i.e., Figure 5.8, top). The experiment could also obtain rating data similar to that gathered in the current experiments. This could inform the debate on whether the learned predictiveness effect is driven by top down or bottom up processes. For example, participants' RT data might show that they are using both the P and NP cues to solve the Phase 2 categorisation task (Figure 5.8, bottom). Therefore, they would have learnt that the NP cues are predictive of the categorisation response in Phase 2. However, the rating data might show a learned predictiveness effect, such that the participants attribute higher ratings to the P cues compared to the NP cues. Such a pattern of results could be taken as evidence that the learned predictiveness effect is due to a rational process. That is, despite having learnt that NP cues are predictive in Phase 2, participants continue to attribute higher ratings to the P cues, perhaps because they were also predictive in Phase 1.

Interestingly, if the covert RT measure showed that participants continued to use the P cues to solve the Phase 2 categorisation (e.g., Figure 5.8, top), this would not necessarily mean that participants are engaging in this strategy through a rational process. In other words, participants could continue to use the P cues in Phase 2 automatically because that attentional behaviour was well practiced in Phase 1. Therefore, caution might be needed when using the words strategic, top-down and voluntary interchangeably (Gaspelin & Luck, 2018).

In these experiments a P cue was always presented alongside a NP cue. Therefore, the resulting RT and N2pc data might be the consequence of more attention being attracted to the P cue, or attention being repelled from the NP cue. The current experiments have no way of disentangling these two possibilities. Future studies could introduce a neutral stimulus condition to further investigate this issue. Lastly, future experiments investigating the automaticity of attentional biases could introduce tasks that ensure the attentional biases are counterproductive to the task goals. For example, Luque et al. (2020) had participants learn that certain colours were P or NP of a categorisation response. Following this, participants completed a dot probe phase in which they were instructed that the probe would appear more often over certain shapes and that colour was completely irrelevant. They found that RTs to the probe were slower, and error rates were higher, when the distractor shape in the dot probe task was rendered in a colour that was previously predictive of the categorisation response.

Conclusions

This thesis explored how associative learning affects attention and whether changes in attention over time could help reconcile exploitative and explorative attentional learning theories. The behavioural results provided clear evidence in favour of exploitation. RTs to targets were faster when they appeared over P cues compared NP cues, and participants made more errors and premature responses in the direction of the P cues. In addition, participants consistently reported that they noticed targets appear more often over P cues compared to NP cues, even though they appeared equally often over all cues. The RT results also provided evidence of strategic processing of the P cues when these cues were task relevant. However, the EEG results showed some evidence in support of the idea that after P cues are preferentially processed, they are inhibited. This inhibition could bias attention towards NP cues at late time points during cue presentation, consistent with an explorative attentional mechanism. Therefore, exploitation may be fast and occur early when cues are presented, but a slower explorative mechanism may need time to emerge. The attention modulated auto-associator presented in the final chapter demonstrates that this explanation is plausible and predicts that long cue durations should result in the elimination of the learned predictiveness effect. However, future empirical studies are required to determine the extent to which time can moderate exploitative and explorative behaviour.

References

Awh, E., Belopolsky, A.V., and 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.

Anderson, B.A. (2013). A value-driven mechanism of attention selection. *Journal of Vision, 13*, 1-16. doi: 10.1167/13.3.7.

Anderson, B.A. (2015). The attention habit: how reward learning shapes attentional selection. *Annals of the New York Academy of Sciences, 1369*, 24-39.

Anderson, B.A. (2018). Neurobiology of value-driven attention. *Current Opinion in Psychology, 29*, 27-33. doi: 10.1016/j.copsyc.2018.11.004.

Anderson, B.A., Laurent, P.A., & Yantis, S. (2011a). Value-drive attentional capture. *PNAS: Proceedings of the National Academy of Sciences of the United States of America, 108*, 10367-10371. doi: 10.1073/pnas.1104047108

Anderson, B.A., Laurent, P.A., & Yantis, S. (2011b). Learned Value Magnifies salience-based attentional capture. *PLos ONE, 6*(11), e27926. doi: 10.1371/journal.pone.0027926

Anderson, S.K., Muller, M.M., & Hillyard, S.A. (2012). Tracking the Allocation of Attention in Visual Scenes with Steady State Evoked Potentials. In M.I. Posner (Ed.), *Cognitive Neuroscience of Attention* (pp. 197-216). New York, NY: The Guilford Press.

Bacigalupo, F., & Luck, S.J. (2019). Lateralised Suppression of Alpha-Band EEG Activity As a Mechanism of Target Processing. *The Journal of Cognitive Neuroscience, 39*(5), 900-917. doi: 10.1523/JNEUROSCI.0183-18.2018

Baetu, I. & Baker, A. (2016). Human Learning About Causation. In R.A. Murphy and R.C. Honey (Eds.), *The Wiley Handbook on the Cognitive Neuroscience of Learning* (pp. 380-407). Chichester, West Sussex, UK: John Wiley & Sons Inc.

Balsam, P.D., Drew, M.R., and Gallistel, C.R. (2010). Time and Associative Learning. *Comparative Cognition & Behavior Reviews, 5*, 1-22. doi: 10.3819/ccbr.2010.50001

Baxter, M.G. & Gaffan, D. (2007). Asymmetry of attentional set in rhesus monkeys learning colour and shape discriminations. *The Quarterly Journal of Experimental Psychology, 60*(1), 1-8. doi: 10.1080/17470210600971485

Beesley, T., Nguyen, K.P., Pearson, D., & Le Pelley, M.E. (2015). Uncertainty and predictiveness determine attention to cues during human associative learning. *The Quarterly Journal of Experimental Psychology, 68*, 2175-2199. doi: 10.1080/17470218.2015.1009919

Bernstein, A.S. (1979). The Orienting Response as Novelty and Significance Detector: Reply to O'Gorman. *Psychophysiology, 16*(3), 263-273. Doi:10.1111/j.1469-8986.1979.tb02989.x

Boakes, R.A., & Costa, D.S.J. (2014). Temporal Contiguity in Associative Learning: Interference and Decay From an Historical Perspective. *Journal of Experimental Psycholgoy: Animal Learning and Cognition, 40*(4), 381-400. doi: 10.1037/xan0000040

Bornstein, M., Hedges, L.V., Higgins, J.P.T., & Rothstein, H.R. (2009). *Effect sizes based on means. Introduction to meta-analysis.* Chichester, UK: Wiley.

Boudewn, M.A., Luck, S.J., Farrens, J.L., & Kappenman, E.S. (2018). How many trials does it take to get a significant ERP effect? It depends. *Psychophysiology, 55*(6):e13049. doi: 10.1111/psyp.13049.

Brainard, D.H. (1997). The Psychophsyics Toolbox. *Spatial Vision, 10*, 433-436. doi: 10.1163/156856897X00357.

Brisson, B., Robitaille, N., & Jolicoeur, P. (2007). Stimulus intensity affects the latency but not the amplitude of the N2pc. *Cognitive Neuroscience and Neuropsychology, 18*(15), 1627-1630. doi: 10.1097/WNR.0b013e3282f0b559.

Brisson, B., & Jolicoeur, P. (2007). The N2pc component and stimulus duration. *Cognitive Neuroscience and Neuropsychology, 18*(11), 1163-1166. doi: 10.1097/WNR.0b013e3281e72d1b.

Carlson, T.A., Hogendoorn, H., Verstraten, F.A. (2006). The speed of visual attention: what time is it? *Journal of Vision, 6*(12), 1406-1411. doi: 10.11676/6.12.6

Castek, A.D., Chasteen, A.L., Scialfa, C.T., & Pratt, J. (2003). Adult age difference in the time course of inhibition of return. *The Journals of Gerontology Series B Psychological Sciences and Social Sciences, 58*(5), 256-259. doi: 10.1093/geronb/58.5

Caio, R., Wu, L., Wang, S. (2017). The different inhibition of return (IOR) effects of emergency managerial experts and novices: An event-related potentials study. *Frontiers in Behavioral Neuroscience, 11*:90. doi: 10.3389/fnbeh.2017.00090.

Carrasco, M. (2011). Visual attention: The past 25 years. *Vision Research, 51*, 1484-1525.

Cherlazzi, L., Perlato, A., Santandrea, E., & Della Libera, C. (2013). Rewards teach visual selective attention. *Vision Research, 85*, 58-72. doi: 10.1016/j.visres.2012.12.005.

Cobos, P.L., Vadillo, M.A., Luque, D., & Le Pelley, M.E. (2018). Learned predictiveness acquired through experience prevails over the influence of conflicting verbal instructions in rapid selective attention. *Plos One, 13*(9): e0200051. doi: 10.1371/journal.pone.0200051

Davis, R. (1959). The role of "attention" in the psychological refractory period. *Quarterly Journal of Experimental Psychology, 11*, 211-220.

Delorme, A., & Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods, 134*, 9-21.

Desimone, R. & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience, 18*, 193-222.

Drazin, D.H. (1961). Effects of foreperiod, foreperiod variability, and probability of stimulus occurrence on simple reaction time. *Journal of Experimental Psychology, 62,* 43-50.

Failing, M, & Theeuwes, J. (2018). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review, 25*(2), 514-538. doi: 10.3758/s13423-017-1380-y

Eimas, P.D. (1966). Effects of overtraining and age on intradimensional and extradimensional shifts in children. *Journal of Experimental Child Psychology, 3*, 348-355

Easdale, L. C., Le Pelley, M.E., & Beesley, T. (2019). The onset of uncertainty facilitates the learning of new associations by increasing attention to cues. The Quarterly Journal of Experimental Psychology, 72(2), 193-208. doi: 10.1080/17470218.

Feldmann-Wusterfeld, T., Uengoer, M., & Schubo, A. (2015). You see what you have learned. Evidence for an interrelation of associative learning and visual selective attention. *Psychophysiology, 52*(11), 1483-1497. doi: 10.1111/psyp.12514.

Gabay, S., & Henik, A. (2010). Temporal expectancy modulates inhibition of return in a discrimination task. *Psychonomic Bulletin & Review, 17*, 47-51.

Gallistel, C.R., and Gibbon, J. (2000). Time, Rate, and Conditioning. (2000). *Psychological Review, 107*(2), 289-344. doi: 10.1037//0033-295X.107.2.289

Gaspelin, N., Leonard, C. J., & Luck, S. J. (2017). Suppression of overt attentional capture by salient-but-irrelevant color singletons. *Attention, Perception, & Psychophysics*, 79(1), 45–62. DOI: https://doi.org/10.3758/s13414-016-1209-1

Gaspelin, N., and Luck, S.J. (2018). “Top-down” Does Not Mean “Voluntary”. Journal of Cognition, *1*(1)*,* 25. doi : http://doi.org/10.5334/joc.28

Gottlieb, J. (2012). Attention, Learning, and the Value of Information. *Neuron, 76*(2), 281-295. doi: 10.1016/j.neuron.2012.09.034.

Griffiths, O. & Mitchell, C.J. (2008). Selective attention in human associative learning and recognition memory. *Journal of Experimental Psychology: General, 137*(4), 626-648.

Griffiths, O., Johnson, A.M., & Mitchell, C.J. (2011). Negative Transfer in Human Associative Learning. *Psychological Science, 22*(9), 1198-1204. doi: 10.1177/0956797611419305

Gwinn, O.S., Russo, S., Baetu, I, Nicholls, M.E.R., & Griffiths, O. (in preparation). Associability shapes primary visual cortical responses: An SSVEP study.

Hall, G., & Pearce, J.M. (1982). Restoring the associability of a pre-exposed CS by a surprising event. *The Quarterly Journal of Psychology, 34B*, 127-140.

Harris, J.A. (2006). Elemental representations of stimuli in associative learning. *Psychological Review, 113*(3), 584–605. doi: 10.1037/0033-295X.113.3.584

Harris, J.A., & Livesy, E.J. (2010). An attention-modulated associative network. *Learning & Behavior, 38*(1), 1-26.

Hedges, L. (1981). Distribution theory for Glass's estimator of effect size and related estimators. *Journal of Educational Statistics, 6*, 107-128. doi: 10.2307/1164588.

Hogarth, L., Dickinson, A., Austin, A., Brown, C., & Duka, T. (2008). Attention and expectation in human predictive learning: The role of uncertainty. *The Quarterly Journal of Experimental Psychology, 61*(11), 1658-1668. doi: 10.1080/17470210701643439

Hollingworth, A., Matsukura, M., & Luck, S. J. (2013). Visual Working Memory Modulates Rapid Eye Movements to Simple Onset Targets. *Psychological Science*, 24(5), 790–796. DOI: https://doi.org/10.1177/0956797612459767

James, W. (1950). Attention. In, *The Princinciples of Psychology* (pp. 402-458). New York, NY: Dover Publications. (Original work published in 1890).

Kaye, H., & Pearce, J.M. (1984). The strength of the orienting response during Pavlovian conditioning. *Journal of Experimental Psychology: Animal Behavior Processes, 10*, 90-109.

Klein, R. (2000). Inhibition of return. *Trends in Cognitive Sciences, 4*(4), 138-146. doi: 10.1016/S1364-6613(00)01452-2

Klein, R., & Lawrence, M.A. (2012). On the Modes and Domains of Attention. In M. I. Posner (Ed.) *Cognitive Neuroscience of Attention* (2nd ed., pp. 11-28). New York, NY: The Guilford Press.

Klein, R., and MacInnes, W.J. (1999). Inhibition of return is a foraging facilitator in visual search. *Psychological Science, 10*, 346-352.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussar, C. (2007). What's new in Psychtoolbox-3. *Perception, 36*(14), 1-16. doi: 10.1068/v070821.

Kokkola, N.H., Mondragon, E., & Alonso, E. (2019). A Double Error Dynamic Asymptote Model of Associative Learning. *Psychological Review, 126*(4), 506-549.

Lang, P.J. & Bradley, M.M. (2013) Appetitive and Defensive Motivation: Goal-Directed or Goal-Determined? *Emotion Review, 5*(3), 230-234.

Lennie, P. (2003). The Cost of Cortical Computation. *Current Biology, 13*(6), 493-497.

Le Pelley, M.E. (2004). The role of associative history in models of associative learning: A selective review and a hybrid model. *The Quarterly Journal of Experimental Psychology Section B, 57*(3), 193-243. doi: 10.1080/02724990344000141

Le Pelley, M.E. (2010). Attention and human associative learning. In C.J. Mitchell & M.E. Le Pelley (Eds.) *Attention and associative learning: From brain to behaviour* (pp. 187-215). Oxford, UK: Oxford University Press.

Le Pelley, M.E., & McLaren, I.P.L. (2003). Learned Associability and Associative Change in Human Causal Learning. *The Quarterly Journal of Experimental Psychology, 56*(1), 68-79.

Le Pelley, M.E., Beesley, T. Griffiths, O. (2011). Overt attention and predictiveness in human contingency learning. *Journal of Experimental Psychology: Animal Behaviour Processes, 37*(2), 220-229. doi: 10.1037/a0021384.

Le Pelley, M.E., Beesley, T., & Griffiths, O. (2016). Associative Learning and Derived Attention in Humans. In R.A. Murphy & R.C. Honey (Eds.), *The Wiley Handbook on the Cognitive Neuroscience of Learning* (pp. 114-135). Chichester, UK: Wiley.

Le Pelley, M.E., Haselgrove, M., & Esber, G.R. (2012). Modelling attention in associative learning. Two processes or one? *Learning & Behavior, 40*, 292-304. doi: 10.3758/s1342-012-0084-4

Le Pelley, M.E., Mitchell, C.J., Beesley, T., George, D.N., & Wills, A.J. (2016). Attention and Associative Learning in Humans: An Integrative Review. *Psychological Bulletin, 42*(10), 1111-1140. doi: 10.1037/bul0000064

Le Pelley, M.E., Pearson, D., Griffiths, O., & Beesley, T. (2015). When Goals Conflict With Values: Counterproductive Attentional and Oculomotor Capture by Reward-Related Stimuli. *Journal of Experimental Psychology: General, 144*, 158-171. doi: 10.1037/xge0000037

Le Pelley, M.E., Reimers, S.J., Calvini, G., Spears, R., Beesley, T., & Murphy, R.A. (2010). Stereotype Formation Biased By Association. *Journal of Experimental Psychology: General, 139*(1), 138-161.

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.

Lopez-Calderon, J., & Luck, S.J. (2014). ERPLAB: an open-source toolbox for the analysis of event-related potentials. *Frontiers in Human Neuroscience, 8*, 213.

Luck, S.J. (2014). *An introduction to the event-related potential technique.* Cambridge, Massachusetts: The MIT Press.

Luck, S.J., Fuller, R.L., Braun, E.L., Robinson, B., Summerfelt, A., & Gold, J.M. (2006). The speed of visual attention in schizophrenia: Electrophysiological and behavioral evidence. *Schizophrenia Research, 85*, 174-195.

Luck, S.J., & Gaspelin, N. (2017). How to Get Statistically Significant Effects in Any ERP Experiment (and Why You Shouldn't). *Psychophysiology, 54*, 146-157. doi:10.1111/psyp.12639.

Luck, S.J., Girelli, M., McDermott, M.T., & Forde, M.A. (1997). Bridging the gap between monkey neurophysiology and human perception: An ambiguity resolution theory of visual selective attention. *Cognitive Psychology, 33*, 64-87. doi: 10.1006/cogp.1997.0660.

Lupianez, J., Milan, E.G., Tornay, F.J., Madrid, E., & Tudela, P. (1997). Does IOR occur in discrimination tasks? Yes, it does, but later. *Perception & Psychophysics, 59*(8), 1241-1254.

Luque, D., Vadillo, M.A., Gutierrez-Cobo, M.J., & Le Pelley, M.E. (2016). The blocking effect in associative learning involves learned biases in rapid attentional capture. *The Quarterly Journal of Experimental Psychology, 71*(2), 522-544. doi: 10.1080/17470218.2016.1262435

Luque, D., Vadillo, M., Le Pelley, M. E., Beesley, T. (2017). Prediction and uncertainty in associative learning: examining controlled and automatic components of learned attentional biases. *The Quarterly Journal of Experimental Psychology, 70*(8), 1485-1503. doi: 10.1080/17470218.2016.1188407

Luque, D., Beesley, T., Morris, R.W., Jack, B.N., Griffiths, O., Whitford, T.J., & Le Pelley M.E. (2017). Goal-Directed and Habit-Like Modulations of Stimulus Processing during Reinforcement Learning. *The Journal of Neuroscience, 37*(11), 3009-3017. doi: 10.1523/JNEUROSCI.3205-16.2017.

Martin-Arevalo, E., Chica, A.B., & Lupianez, J. (2016). No single electrophysiological marker for facilitation and inhibition of return: A review. *Behavioural Brain Research, 300*, 1-10. doi: 10.1016/j.bbr.2015.11.030.

Martin-Pichora, A.l., Mankovsky-Arnold, T.D., & Katz, J. (2011). Implicit versus explicit associative learning and experimentally induced placebo hypoalgesia. *Journal of Pain Research, 15*(4), 67-77. doi: 10.2147/JPR.S15966.

Mackintosh, N.J. (1975). A theory of attention: Variations in the associability of stimuli with reinforcement. *Psychological Review, 82*, 276-298. doi: 10.1037/h0076778.

Mackintosh, N.J. & Little, L. (1969). Intradimensional and extradimensional shift learning by pigeons. *Psychometric Science, 14*, 5-6. doi: 10.3758/BF03336395.

Macaluso, E., & Doricchi, F. (2013). Attention and predictions: control of spatial attention beyond the endogenous-exogenous dichotomy. *Frontiers in Human Neuroscience, 7*, 685. doi: 10.3389/fnhum.2013.00685

Matthews, T.J., and Lerer, B.E. (1987). Behavior patterns in pigeons during autoshaping with an incremental conditioned stimulus. *Animal Learning & Behavior, 15*, 69-75.

Mazza, V., Pagano, S., & Caramazza, A. (2013). Multiple object individuation and exact numeration. *Journal of Cognitive Neuroscience, 25*(5), 697-705. doi: 10.1162/jocn_a_00349.

McDonald, J.J., Hickey, C., Green, J.J., & Whitman, J.C. (2009). Inhibition of return in the covert deployment of attention: Evidence from human electrophysiology. *Journal of Cognitive Neuroscience, 21*(4), 725-733. doi: 10.1162/jocn.2009.21042.

McLaren, I.P.L, & Mackintosh, N.J. (2000). Associative learning and elemental representations. I: A theory and its application to latent inhibition and perceptual learning. *Animal Learning and Behavior, 26*, 211-246.

McClelland, J.L., & Rumelhart, D.E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs, and exercises.* Cambridge, MA: MIT Press.

Meigen, T., Lagreze, E., & Bach, M. (1994). Asymmetries in preattentive line detection. *Vision Research, 34*(23), 3103-3109. doi: 10.1016/0042-6986(94)90076-0.

Moore, T., & Zirnsak, M. (2017). Neural Mechanisms of Selective Visual Attention. *Annual Review of Psychology, 68*, 47-72.

Mitchell, C.J., Griffiths, O., Seetoo, J. & Lovibond, P.F. (2012). Attentional mechanisms in learned predictivness. *Journal of Experimental Psychology: Animal Behavor Processes, 38*(2), 191-202. doi: 10.1037/a0027385.

Muller, M.M., Picton, T.W., Valdes-Sosa, P., Riera, J., Teder-Salejarvi, W.A., & Hillyard, S.A. (1998). Effects of spatial selective attention on the steady-state visual evoked potential in the 20-28 Hz range. *Cognitive Brain Research, 6*(4), 249-261.

Nasser, H.M. & Delamater, A.R. (2016). The Determining Conditions for Pavlovian Learning. Psychological and Neurobiological Considerations. In R. A. Murphy & R.C. Honey (Eds.) *The wiley handbook on the cognitive neuroscience of learning* (pp. 7-46). Chichester, UK: Wiley.

Nasser, H.M. & McNally, G.P. (2013). Neural correlates of appetitive-aversive interactions in Pavlovian fear conditioning. *Learning and Memory, 20*, 220-228. doi: 10.1101/lm.029744.112.

Niv, Y., Daniel, R., Greana, A., Gershman, S.J., Leong, Y.C., Radulescu, A., & Wilson, R.C. (2015). Reinforcement Learning in Multidimensional Environments Relies on Attention Mechanisms. *The Journal of Neuroscience, 35*(21), 8145-8157.

Norcia, A.M., Appelbaum, L.G., Ales, J.M., Cottereau, B.R., & Rossion, B. (2015). The steady-state visual evoked potential in vision research: A review. *Journal of Vision, 15*(6), 1-46. doi: 10.1167/15.6.4.

Owen, A.M., Roberts, A.C., Polkey, C.E., Sahakian, B.J., & Robbins, T.W. (1991). Extra-dimensional versus intra-dimensional set shifting performance following frontal lobe excisions, temporal lobe excisions or amygdalo-hippocampectomy in man. *Neuropsychologia, 29*(10), 993-1006.

Pearce, J.M., & Bouton, M.E. (2001). Theories of associative learning in animals. *Annual Review of Psychology, 52*, 111-139. doi: 10.1146/annurev.psych.52.1.111.

Pearce, J.M., & Hall, G. (1980). A model of Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review, 87*, 532-552. doi: 10.1037/0033-295X.87.6.532

Pearce, J.M., & Mackintosh, N.J. (2010). Two theories of attention: a review and a possible integration. In C.J. Mitchell & M.E. Le Pelley (Eds.) *Attention and associative learning: From brain to behaviour* (pp. 11-40). Oxford, UK: Oxford University Press.

Pelli, D.G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision, 10*(4), 437-442. doi: 10.1163/15685689&X00366

Pierce, A., Crouse, M.D., & Green, J.J. (2017). Evidence for an attentional component of inhibition of return in visual search. *Psychophysiology, 54*(8), 1676-1685.

Posner, M.I., Rafal, R.D., Choate, L.S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology, 2*(3), 211-228. doi: 10.1080/02643298508252866.

Reynolds, J.H., Chelazzi, L., & Desimone, R. (1999). Competitive mechanisms subserve attention in macaque areas V2 and V4. *Journal of Neuroscience, 19*, 1736-1753. doi: 10.1523/JNEUROSCI.19-05-01736.1999.

Robitaille, N., & Jolicoeur, P. (2006). Effect of cue-target interval on the N2pc. *Cognitive Neuroscience and Neuropsychology, 17*(15), 1655-1658. doi: 10.1097/01.wnr.0000236859.16457.34.

Samuel, A.G., & Kat, D. (2003). Inhibition of return: A graphical meta-analysis of its time course and an empirical test of its temporal and spatial properties. *Psychonomic Bulletin & Review, 10*(4), 897-906.

Satel, J., Wang, Z., Hilchey, M.D., & Klein, R.M. (2012). Examining the dissociation of retinotopic and spatiotopic inhibition of return with event-related potentials. *Neuroscience Letters, 524*(1), 40. doi: 10.1016/j.neulet.2012.07.003

Schmajuk, N. A., Lam, Y.W., & Gray, J.A. (1996). Latent inhibition: A neural network approach. *Journal of Experimental Psychology: Animal Behavior Processes, 22*, 321-349.

Schneiderman, N., and Gormezano, I. (1964). Conditioning of the Nictitating Membrane of the Rabbit as a Function of the CS-US Interval. *Journal of Comparative and Physiological Psychology, 57*(2), 188-195.

Shepp, B.E. & Eimas, P.D. (1964). Intradimensional and Extradimensional Shifts in the Rat. *Journal of Comparative and Physiological Psychology, 57*(3), 357-361.

Silva, K.M., and Timberlake, W. (1997). A Behavior Systems View of Conditioned States during Long and Short CS-US Intervals. *Learning and Motivation, 28*, 465-490.

Silva, K.M., and Timberlake, W. (1998). The organization and temporal properties of appetitive behaviour in rats. *Animal Learning & Behavior, 26*(2), 182-195.

Silva, K.M., and Timberlake, W. (2005). A behaviour systems view of the organisation of multiple responses during a partially or continuously reinforced interfood clock. *Learning & Behaviour, 33*(1), 99-110.

Smulders F.T.Y., Miller J.O. (2012). The lateralized readiness potential. In: Luck SJ, Kappenman E, editors. *Oxford handbook of event-related potential components*. New York: Oxford University Press.

Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica, 135*, 77-99. doi: 10.1016/j.actpsy.2010.02.006.

Theeuwes, J., Olivers, C.N.L., & Belopolsky, A. (2010). Stimulus-driven capture and contingent capture. *WIREs Cognitive Science, 1*(6), 872-881. doi: 10.1016/j.visres.2009.11.013.

Theeuwes, J. & Belopolsky, A.V. (2012). Reward grabs the eye: Occulomotor capture by rewarding stimuli. *Vision Research, 74*, 80-85. doi: 10.1016/j.visres.2012.07.024

Theeuwes, J. (2018). Visual Selection: Usually Fast and Automatic; Seldom Slow and Volitional. *Journal of Cognition, 1*(1), 1-29. doi: 10.5334/joc.13.

Theeuwes, J. (2019). Goal-driven, stimulus-driven, and history-driven selection. *Current Opinion in Psychology, 29*, 97-101.

Thorndike, E.L. (1911). *Animal Intelligence*. New York, NY: Macmillian.

Thomas, D.R. (1970). Stimulus selection, attention and related matters. In J.H. Reynierse (Ed.), *Current issues in animal learning*. Lincoln, NJ: University of Nebraska Press.

Timberlake, W., and Lucas, G.A. (1989). Behavior systems and learning: from misbehaviour to general principles. In S.B. Klein and R.R. Mowrer (Eds.), *Contemporary Learning Theories: Instrumental Conditioning Theory and the Impact of Biological Constraints on Learning* (pp. 237-275). Hillsdale, NJ: Erlbaum.

Timberlake, W. (2001). Motivational Modes in Behavior Systems. In R.R. Mowrer and S.B. Klein (Eds.), *Handbook of Contemporary Learning Theories* (pp. 155-209), New York, NY: Routledge

Trobalon, J.B., Miguelez, D., McLaren, I.P.L, & Mackintosh, N.J. (2003). Intradimensional and extradimensional shifts in spatial learning. *Journal of Experimental Psychology: Animal Behavior Processes, 29*, 143-152. doi: 10.1037/0097-7403.29.2.143

Uengoer, M., Pearce, J.M., Lachnit, H., & Koenig, S. (2018). Context modulation of learned attention deployment. *Learning & Behavior, 46*, 23-37. doi: 10.3758/s13420-017-0277-y.

Vickers, D., Nettelbeck, T., & Wilsom, R.J. (1972). Perceptual indices of performance: The measurement of 'inspection time' and 'noise' in the visual system. *Perception, 1*, 263-295.

Wang, Z., & Klein, R.M. (2010). Searching for inhibition of return in visual search. A review. *Vision Research, 50*, 220-228. doi: 10.1016/j.visres.2009.11.013

Wang, L., Yu, H., & Zhou, X. (2013). Interaction between value and perceptual salience in value-driven attentional capture. *Journal of Vision, 13*(3):5, 1-13. doi: 10.1167/13.3.5.

Weiser, M.J., Miskovic, V., & Keil, A. (2016). Steady-state visual evoked potentials as a research tool in social affective neuroscience. Psychophysiology, 53(12), 1763-1775.

Woodman, G.F. (2010). A Brief Introduction to the Use of Event-Related Potentials (ERPs) in Studies of Perception and Attention. *Attention, Perception & Psychophysics, 72*(8), 2031-2046. doi: 10.3758/APP.72.8.2031

Yeshurun, Y., & Carrasco, M. (1998). Attention improves or impairs visual performance by enhancing spatial resolution. *Nature, 396* (6706), 72–75.

Appendix A: Latencies and probe locked ERPs at each SOA for Experiments 4 and 5

Table A1

CS onset N2pc median 50% Fractional Area Latencies (ms) in Experiments 4 and 5

SOA	Experiment 4	Experiment 5
PO7/PO8	326	328
O1/O2	376	378

Table A2

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 4

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms
PO7/PO8	223	203	193	186	184	182
O1/O2	223	203	192	185	182	183

Table B3

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 5

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms
PO7/PO8	223	203	193	186	184	182
O1/O2	223	203	192	185	182	183

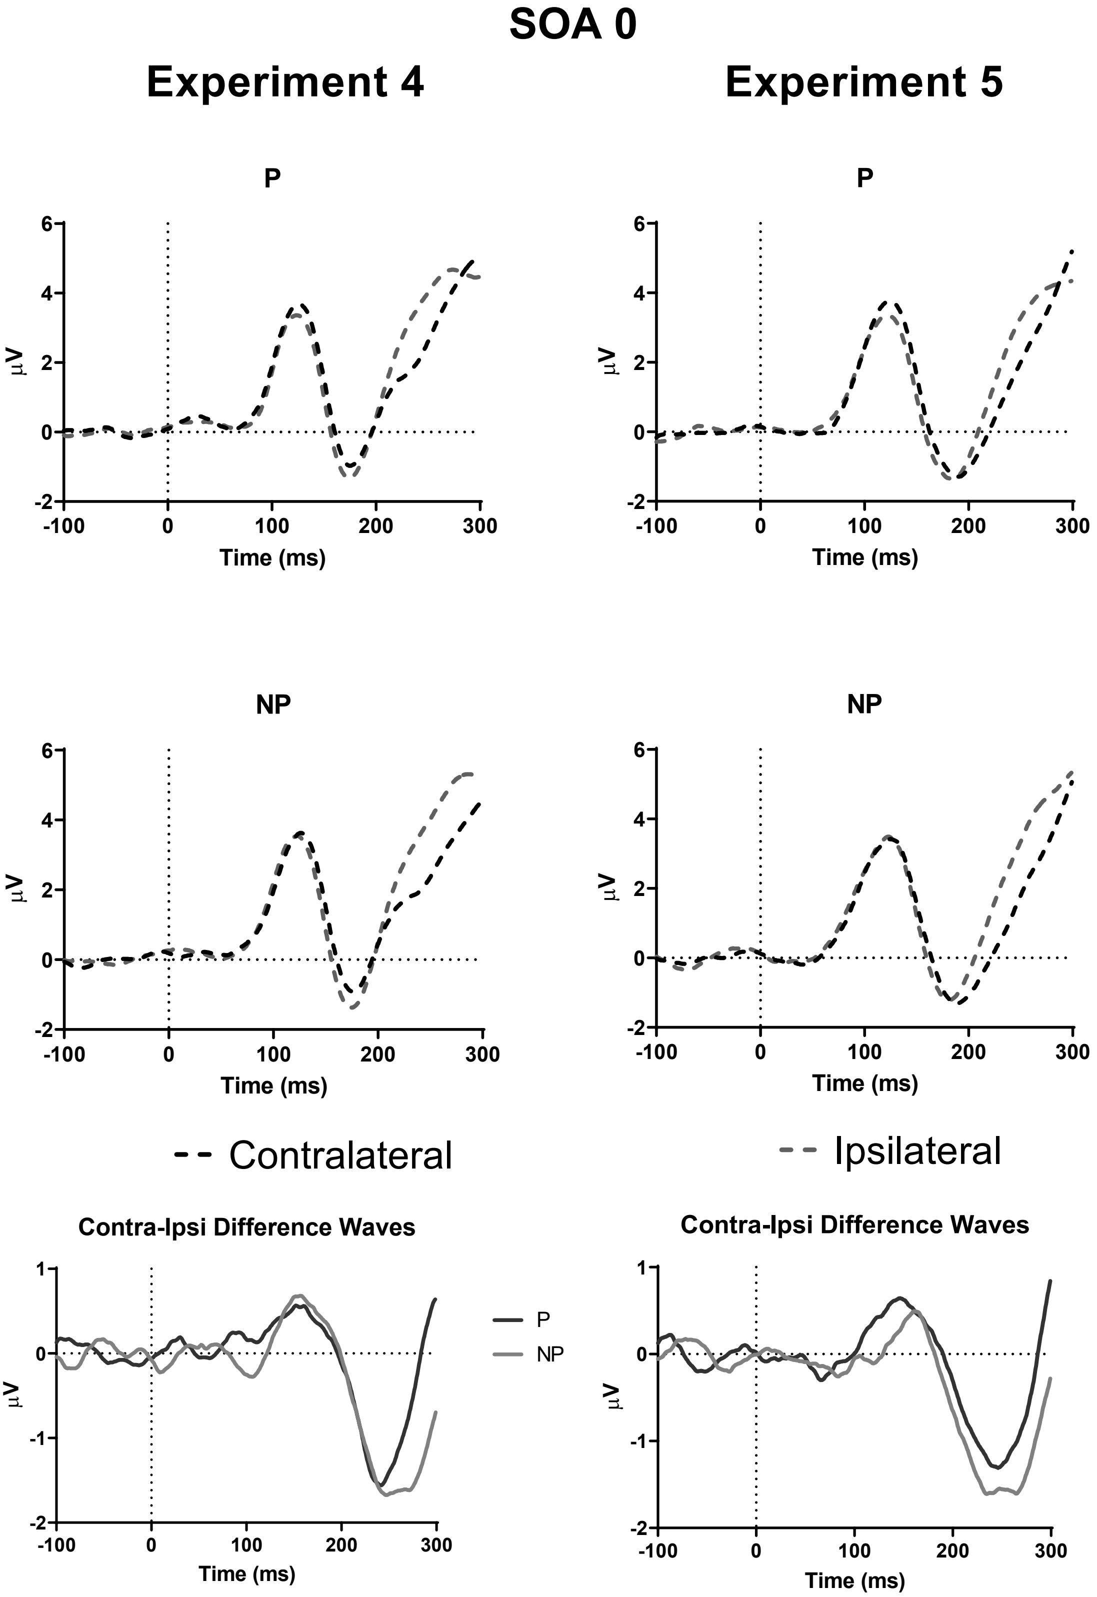

Figure A1: Grand average ERP waveforms for the 0 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

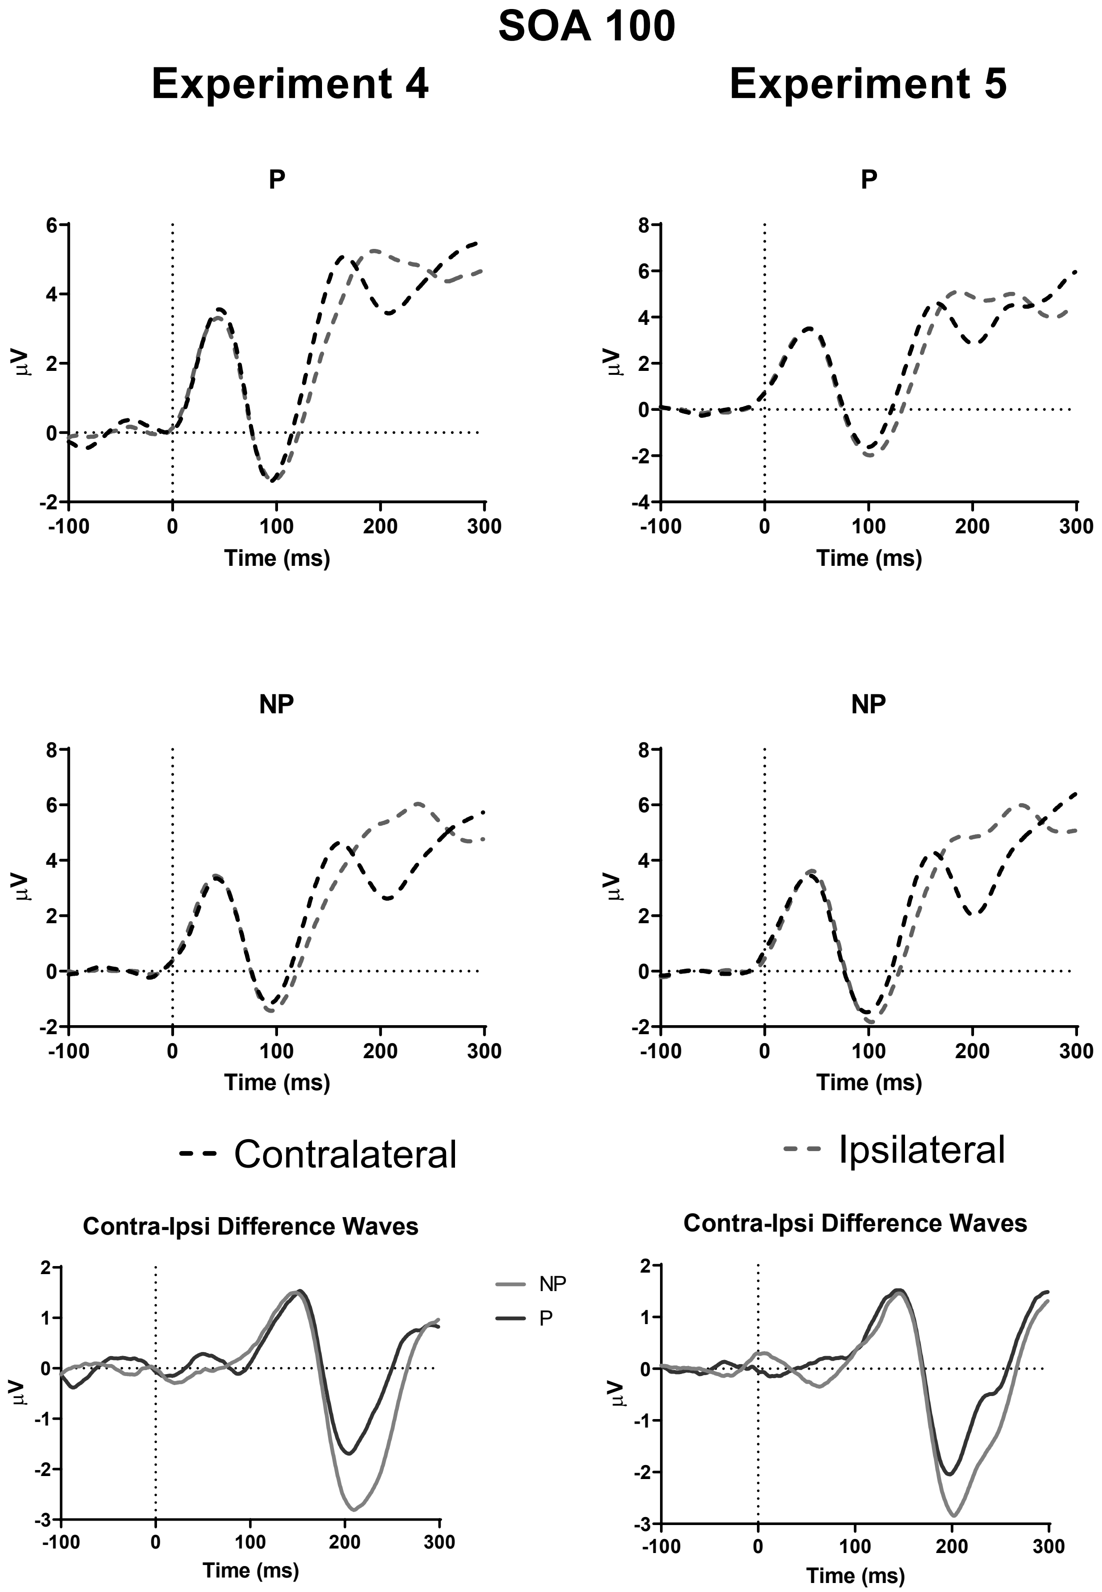

Figure A2: Grand average ERP waveforms for the 100 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

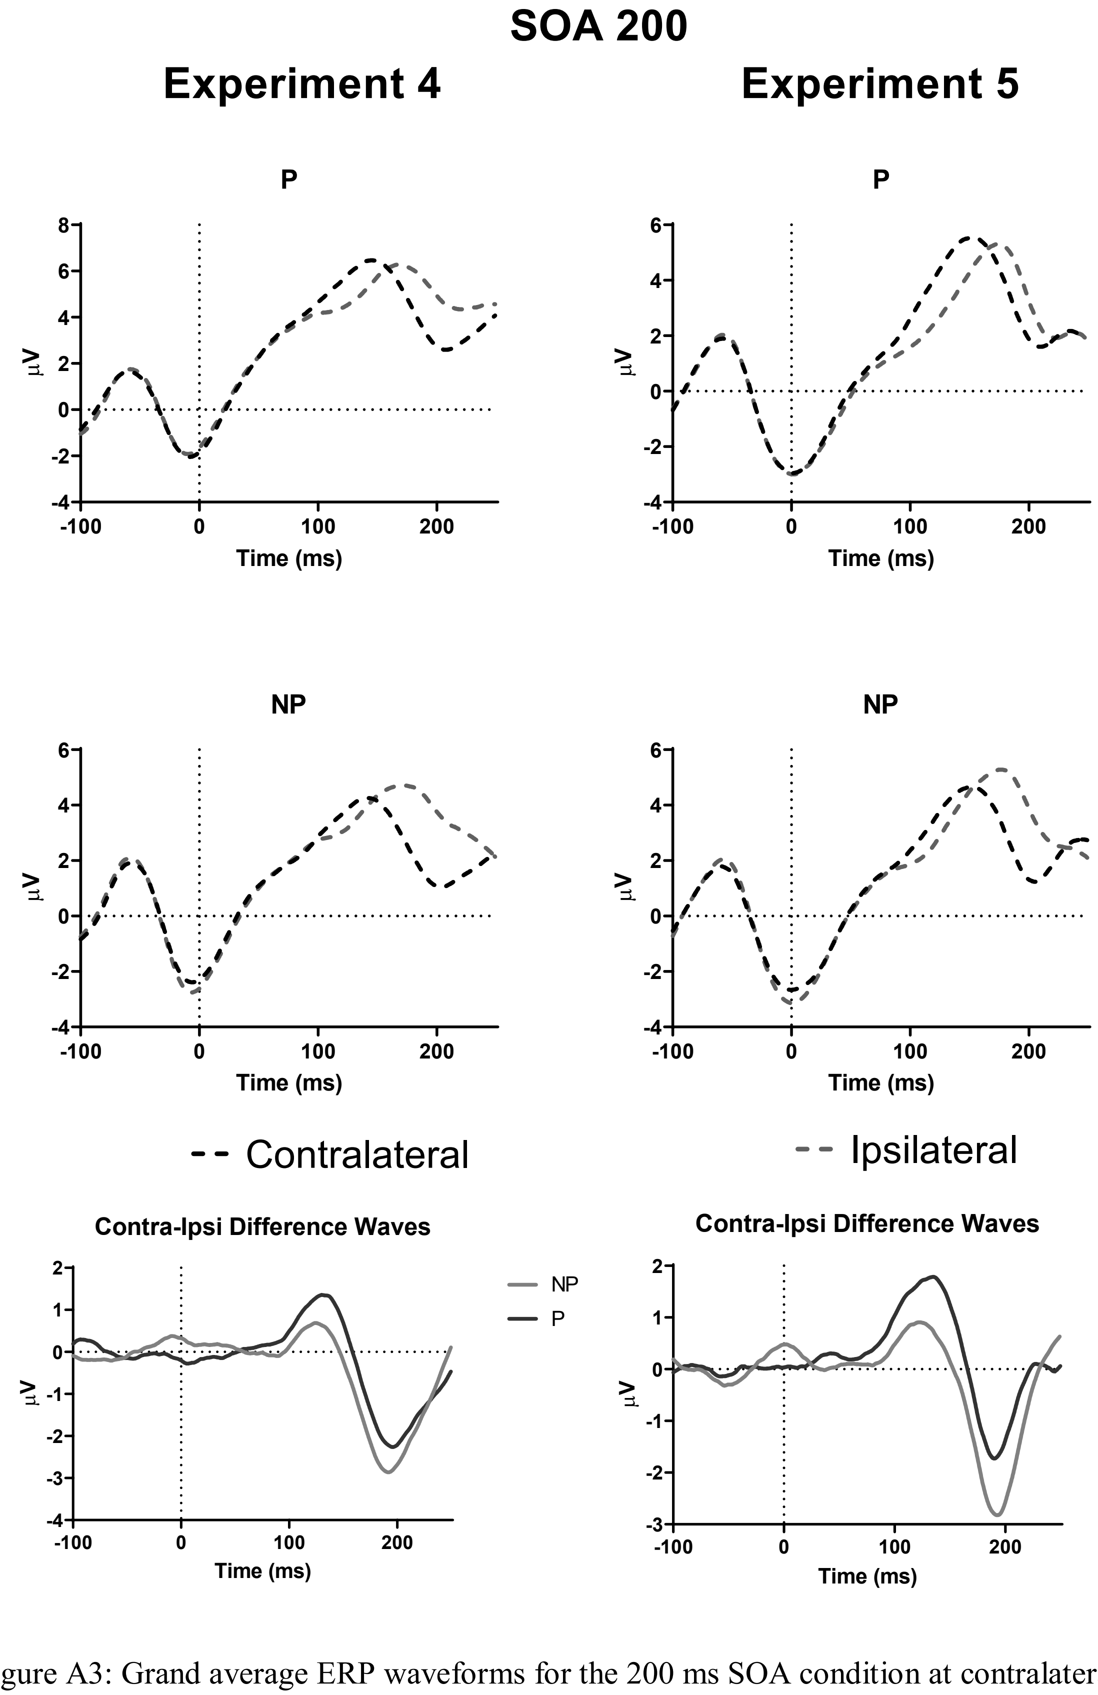

Figure A3: Grand average ERP waveforms for the 200 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

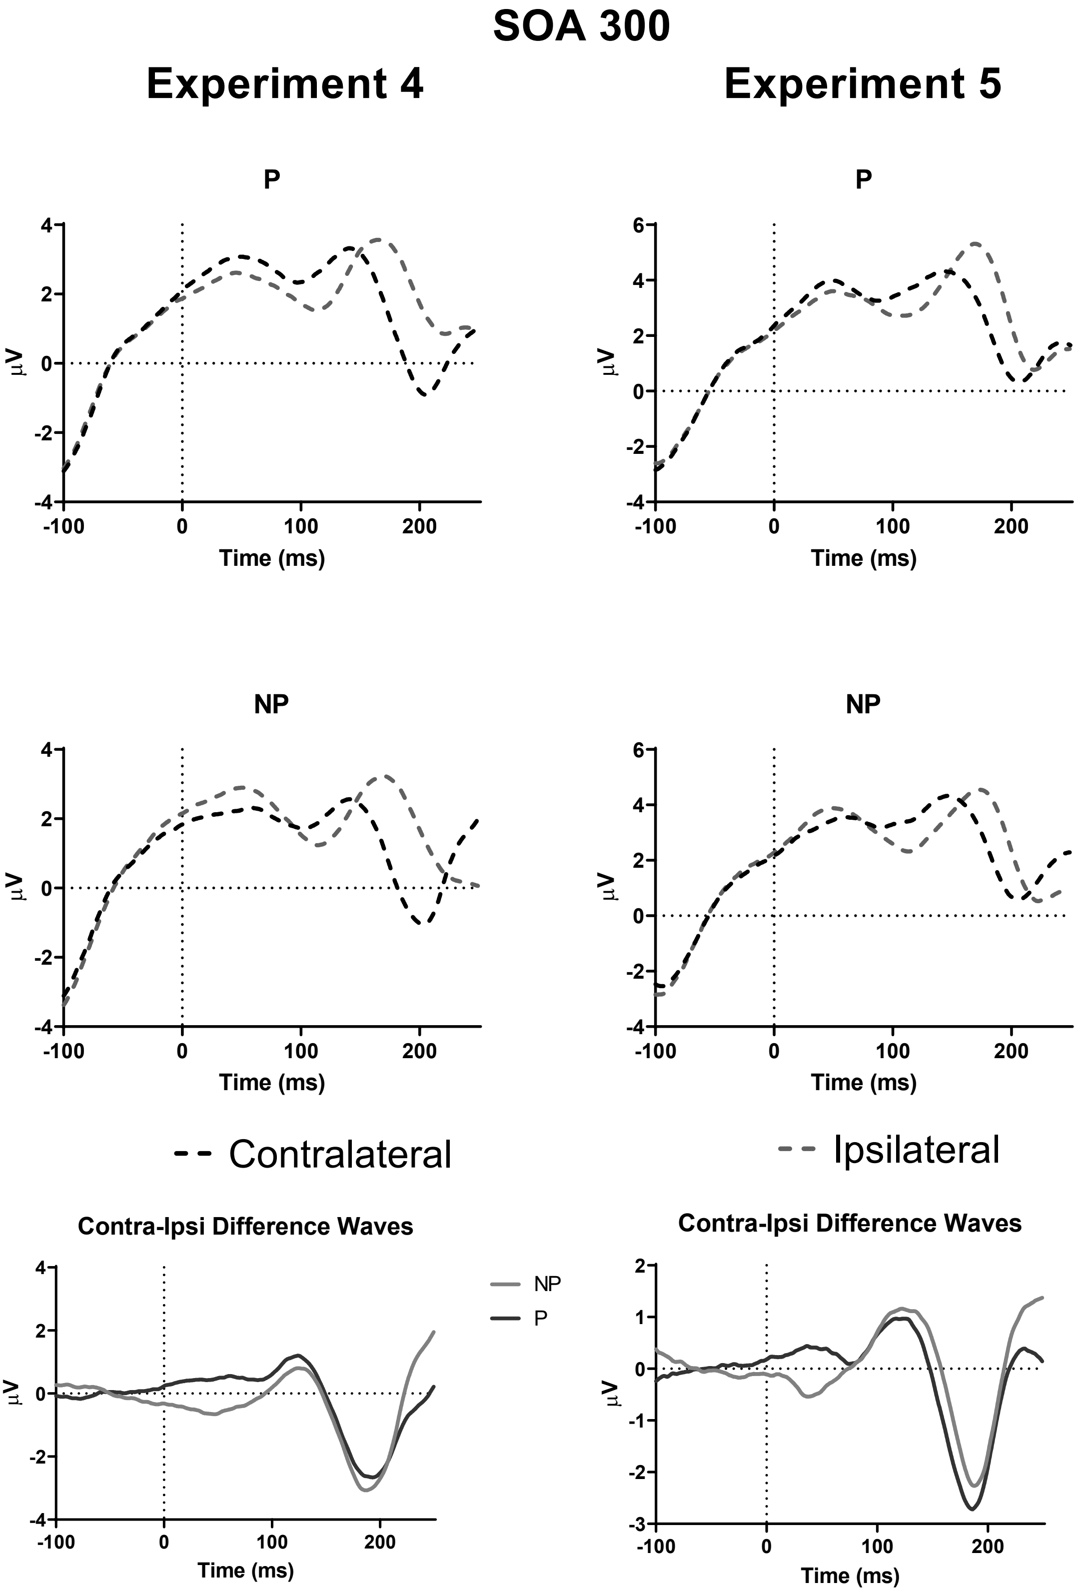

Figure A4: Grand average ERP waveforms for the 300 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

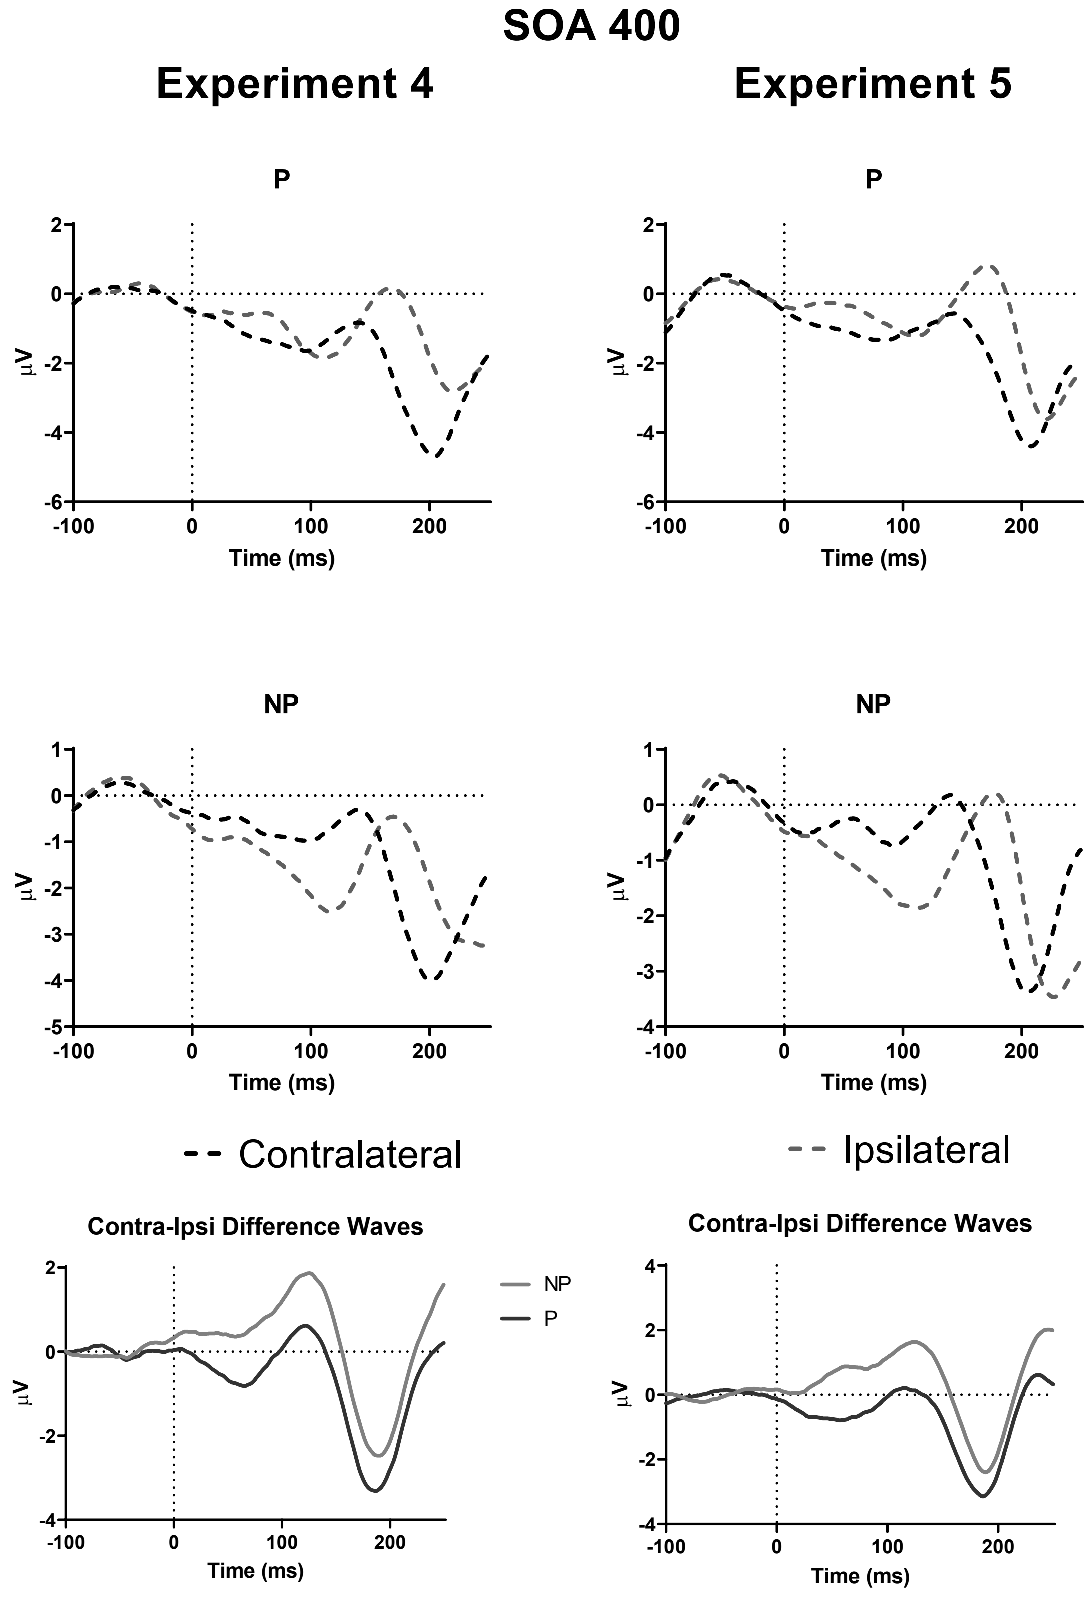

Figure A5: Grand average ERP waveforms for the 400 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

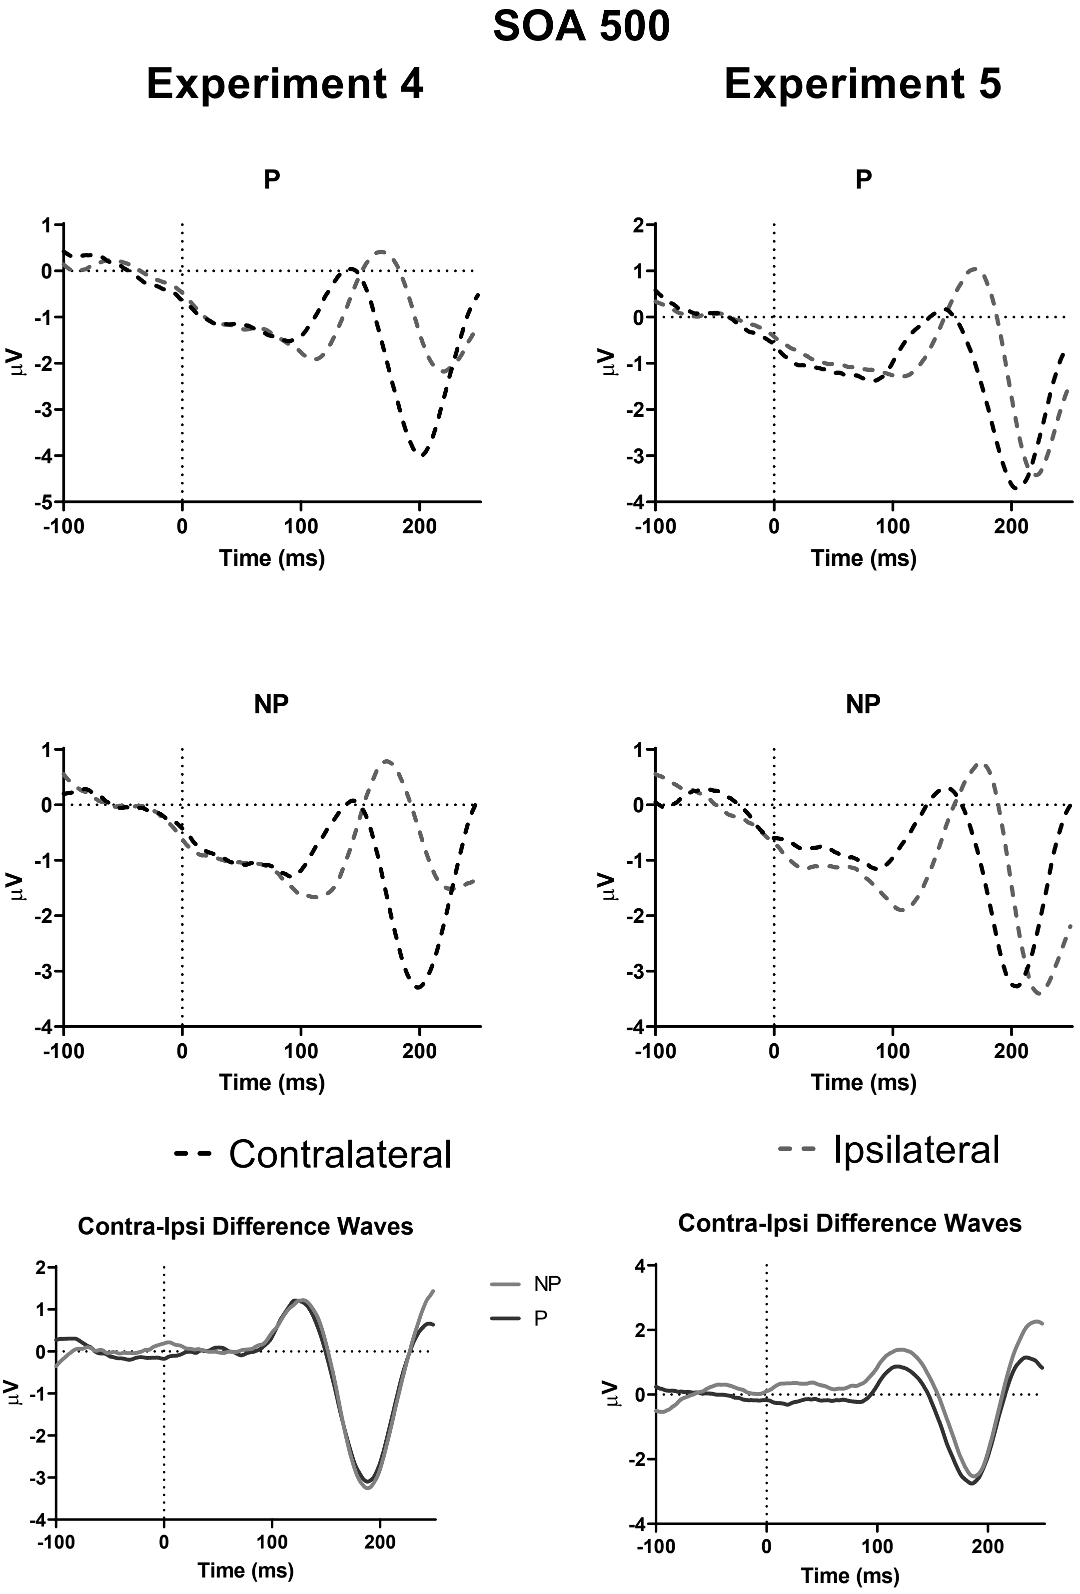

Figure A6: Grand average ERP waveforms for the 500 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

Appendix B: Latencies and probe locked ERPs at each SOA for Experiments 6

Table B1

CS onset N2pc median 50% Fractional Area Latencies (ms) in Experiment 6

SOA	Experiment 6
PO7/PO8	450
P7/P8	436
O1/O2	446

Table B2

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 6

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms	600 ms
PO7/PO8	274	244	230	231	213	215	207
P7/P8	280	247	236	230	212	219	207
O1/O2	265	244	229	227	214	215	207

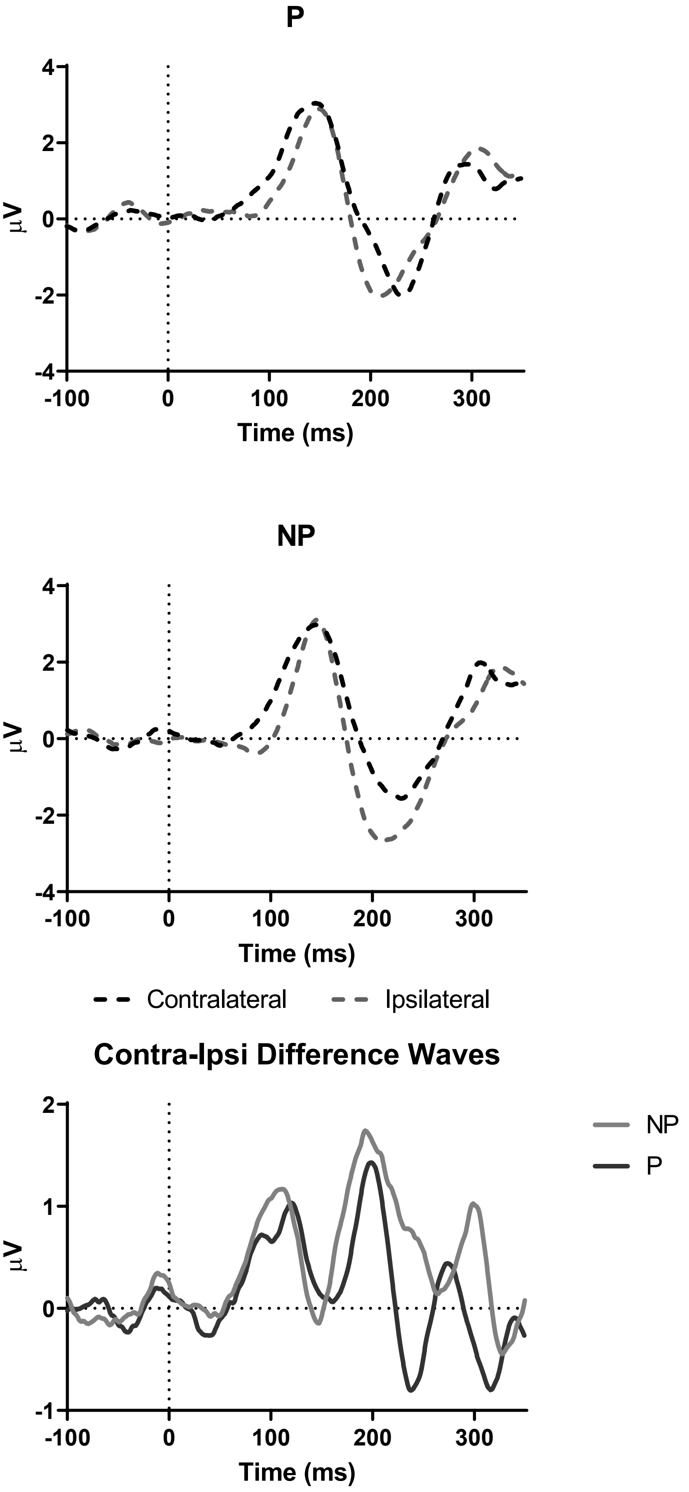
P

NP

Contra-Ipsi Difference Waves

Figure B1: Grand average ERP waveforms for the 0 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

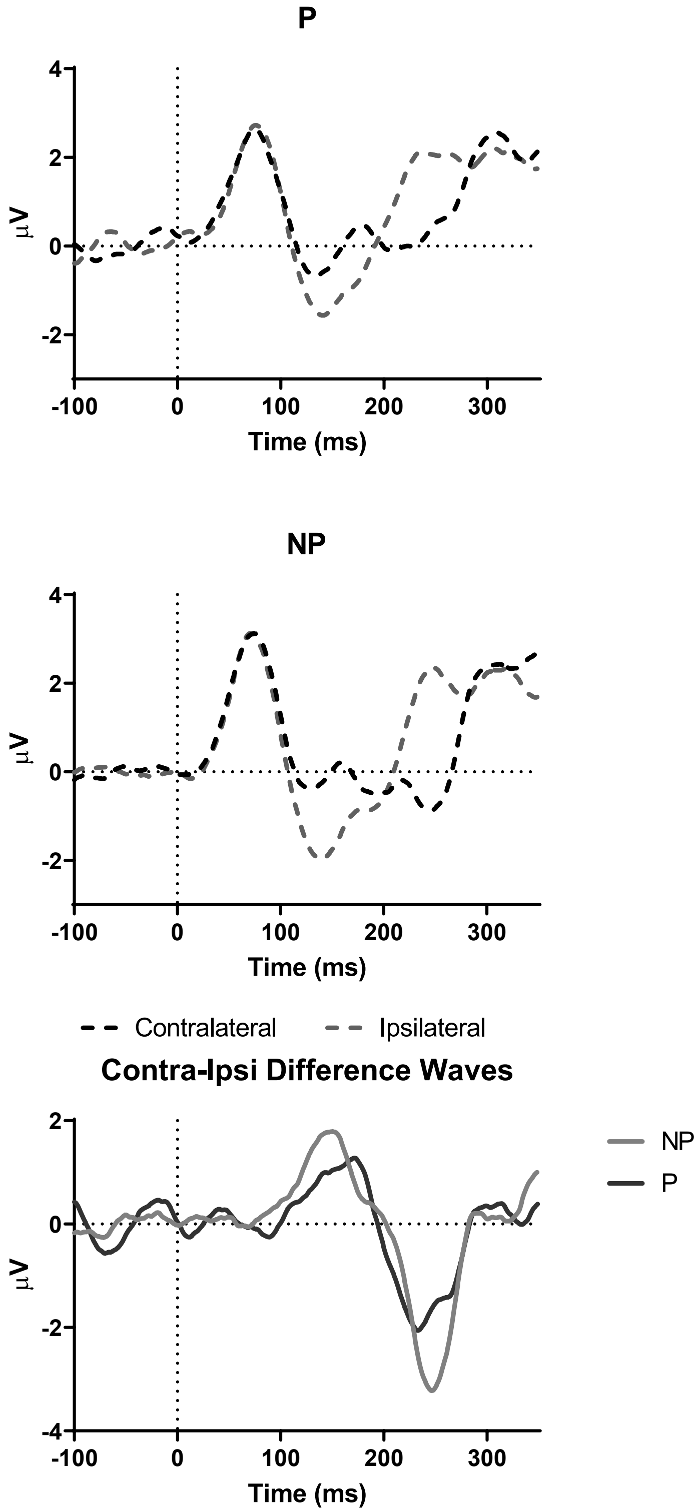

Figure B2: Grand average ERP waveforms for the 100 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

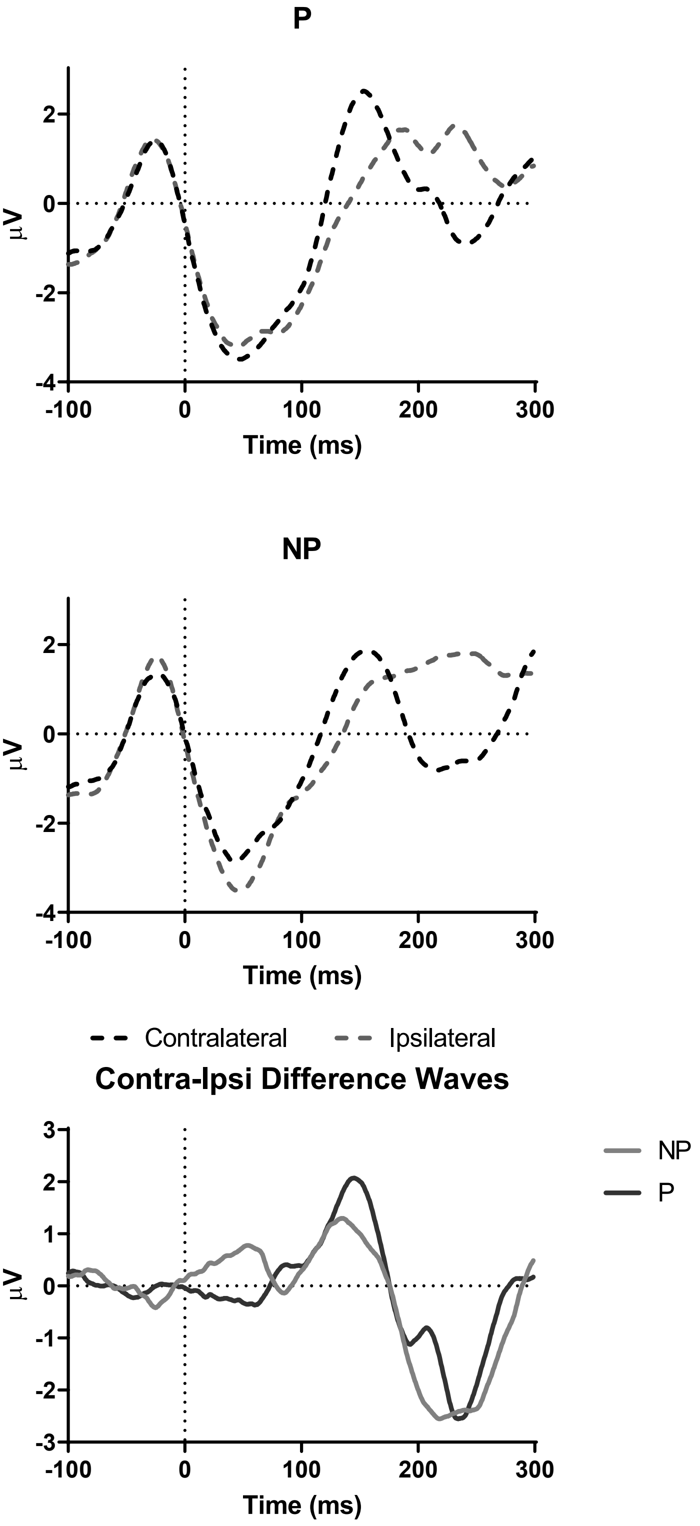

Figure B3: Grand average ERP waveforms for the 200 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

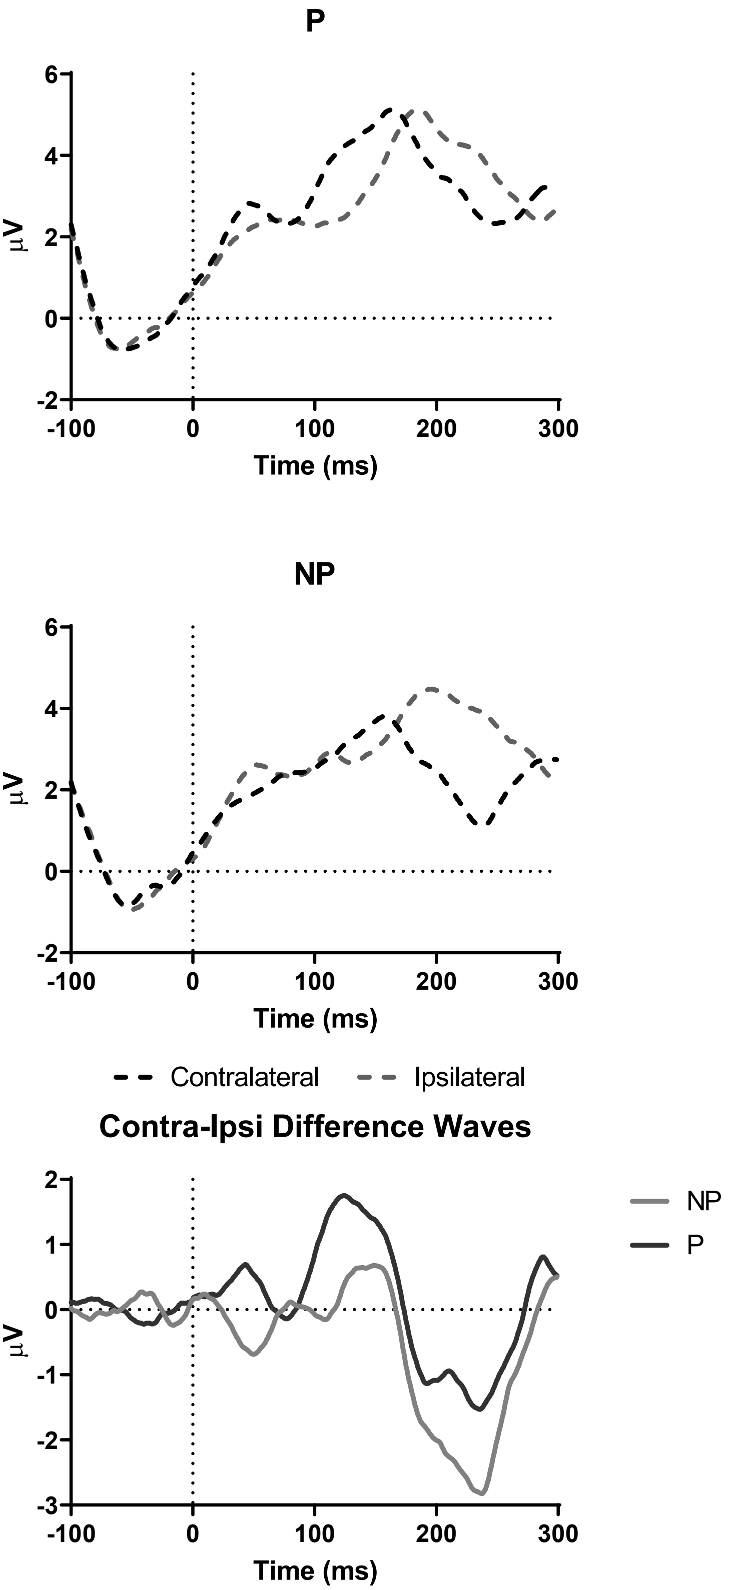

Figure B4: Grand average ERP waveforms for the 300 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

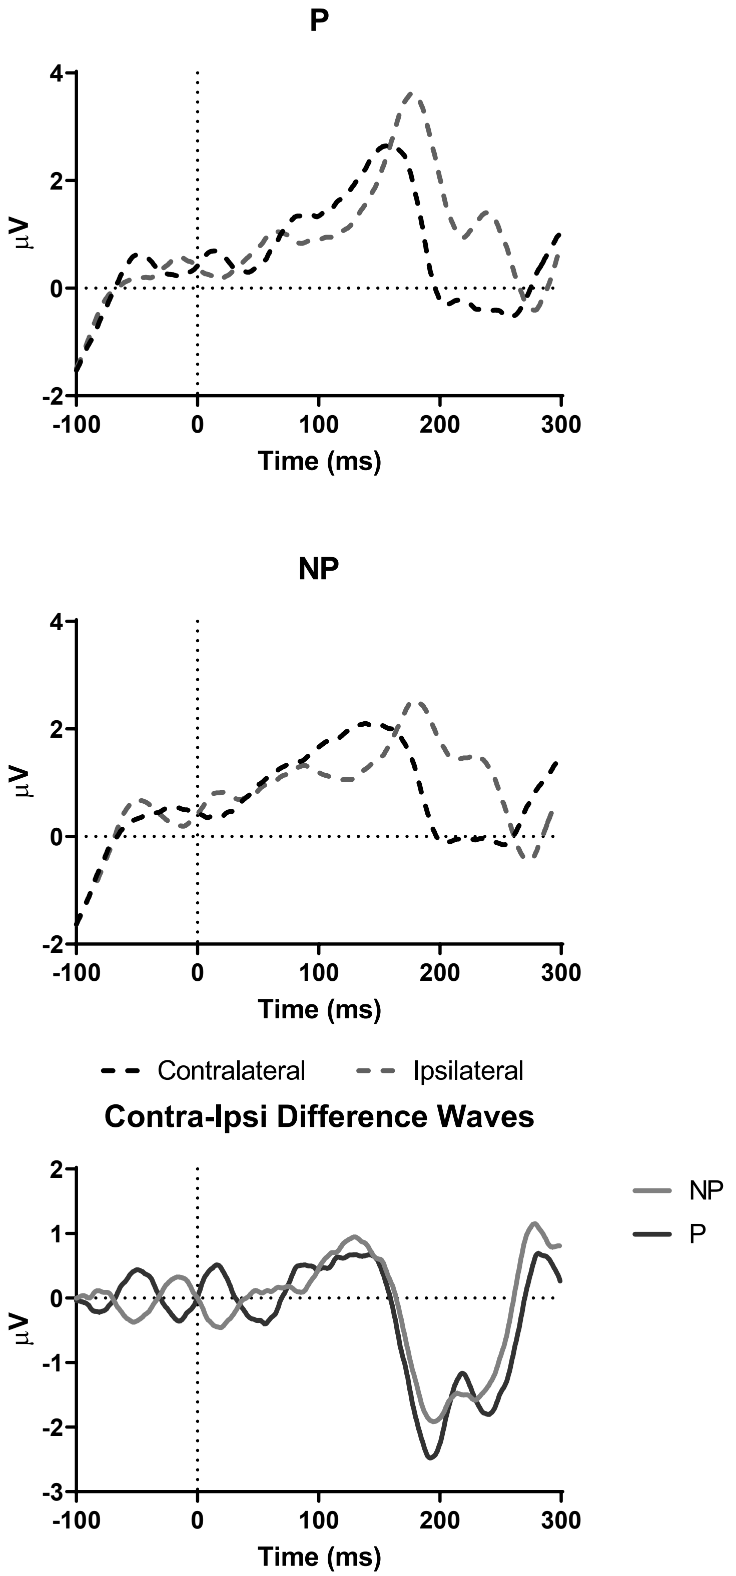

Figure B5: Grand average ERP waveforms for the 400 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

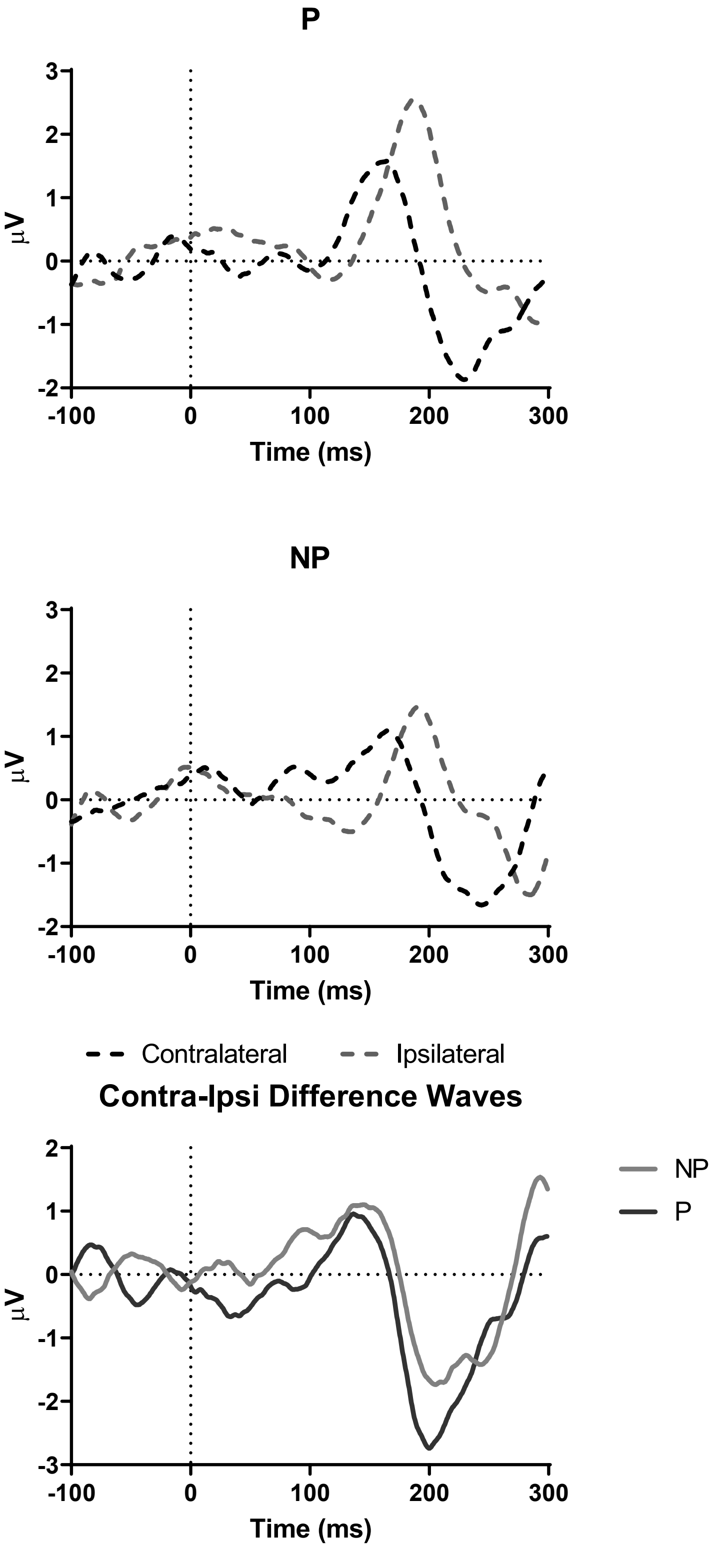

Figure B6: Grand average ERP waveforms for the 500 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

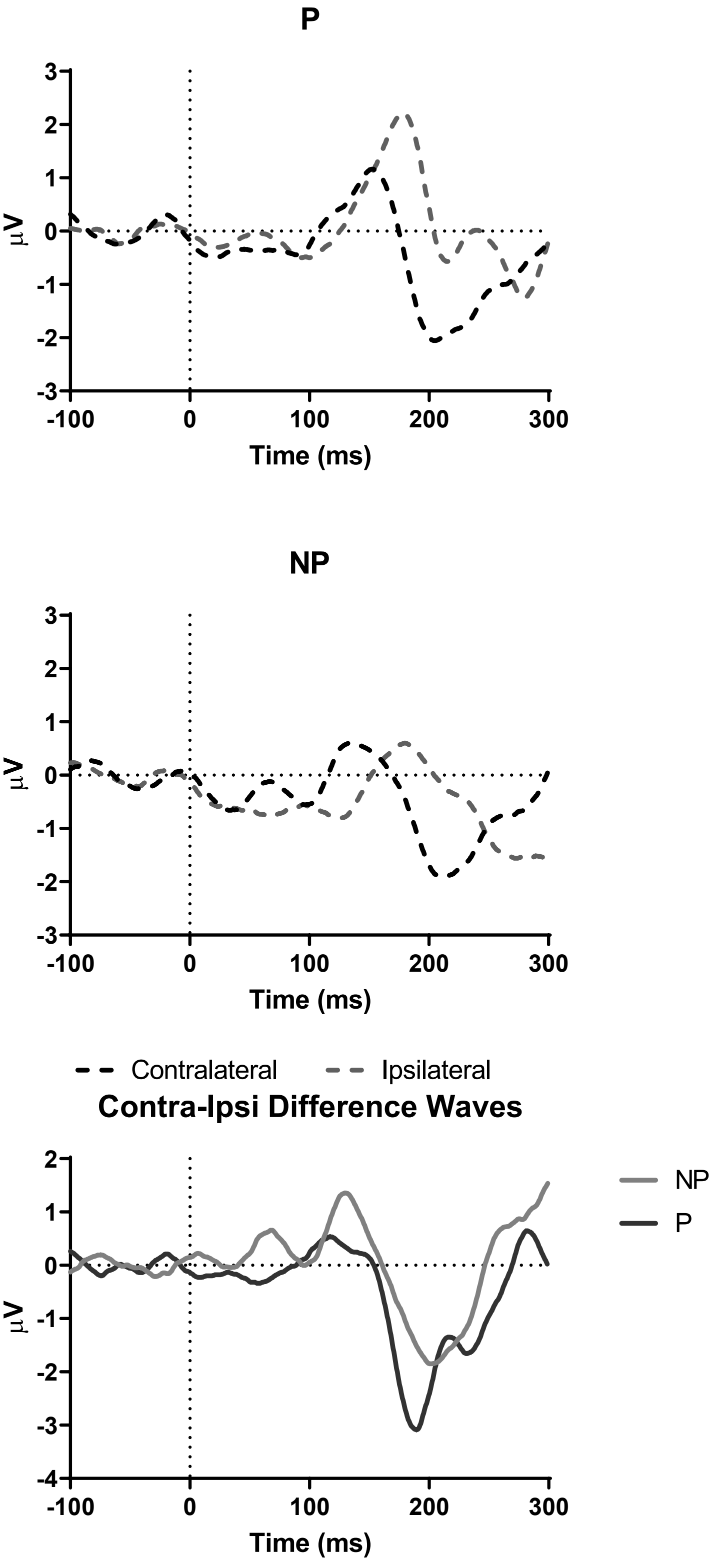

Figure B7: Grand average ERP waveforms for the 600 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

Appendix C: Discussion of cue-locked N2pc

In the published manuscript (Russo et al., 2019) that included Experiments 4 and 5 the cue locked data were presented in a supplementary materials, and it was reported that no lateralised differences were found (see Figure C1). However, the cue locked data reported in Experiment 6 did show lateralised differences in the direction expected if participants were attending towards the P cue (Figure 3.6).

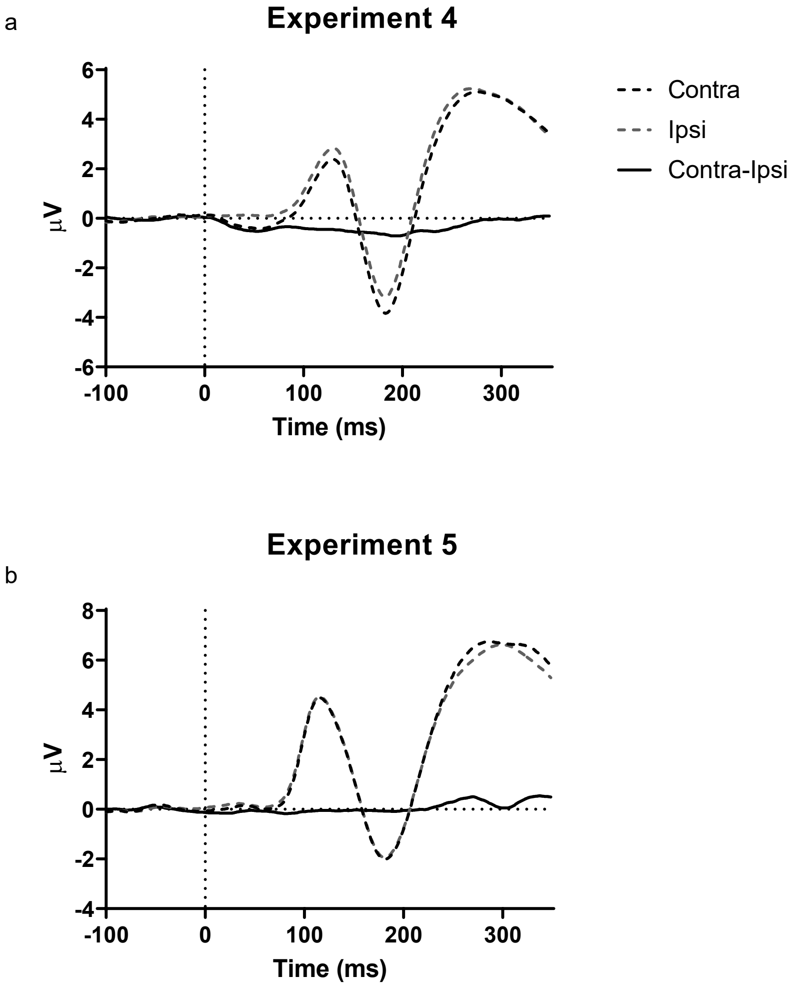

Figure C1: Grand average ERP waveforms at contralateral versus ipsilateral electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the cues during the categorisation task in (a) Experiment 4 and (b) Experiment 5. The contralateral-ipsilateral waveform is also shown.

Comparing the cue-locked data in Experiments 4 and 5 with those obtained in Experiment 6, one noticeable difference is that they have different epochs (i.e., 350 ms vs. 500 ms). The cue locked data from Experiments 4 and 5 were originally analysed using a 350 ms epoch because the probe locked data had shown lateralised differences by this time (for example, see Figure 2.11). However, it seemed appropriate to reanalyse the cue locked data from Experiments 4 and 5 with the extended epoch of 500 ms. This subsequent analysis did find voltage differences at contralateral and ipsilateral electrode sites (see Figure C2), however the direction of the effect was opposite to that found in Experiment 6. That is, the cue locked data from the categorisation task in Experiments 4 and 5 seems to show that participants were attending to the NP cue. This reanalysis was conducted after data collection was complete for all experiments contained in this thesis. Therefore, the cue locked data from the categorisation task in Experiments 7, 8 and 9 of the next chapter (Chapter 4) were also analysed with an epoch of 500 ms. Note that these subsequent experiments used the same feature-based discrimination as Experiments 4 and 5.

The cue locked waveforms can be seen in Appendices D (Experiment 7), E (Experiment 8) and F (Experiment 9). However, a meta-analysis of the N2pc amplitudes from all the cue locked data can be seen below in Figure C3. Note that Experiment 6 is shown at the top of Figure C3 but was not included in the meta-analysis because it used a conjunction discrimination compared to the feature discrimination that was used in the other experiments. From Figure C3 it is clear that the feature-based discrimination showed cue-locked N2pc amplitudes in the direction opposite to what was expected (i.e., greater negative voltages at electrodes contralateral to the NP cue). However, it is also clear from Figure C3 that the conjunction discrimination used in Experiment 6 resulted in a cue locked N2pc that was in the direction one would expect if participants were paying attention to the P cue.

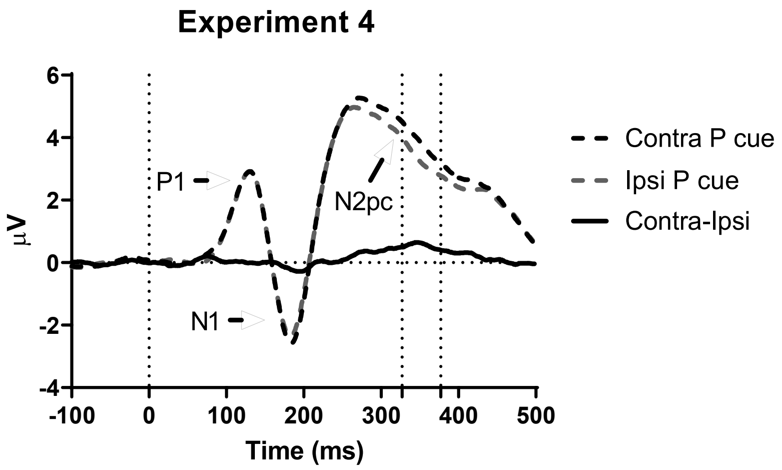

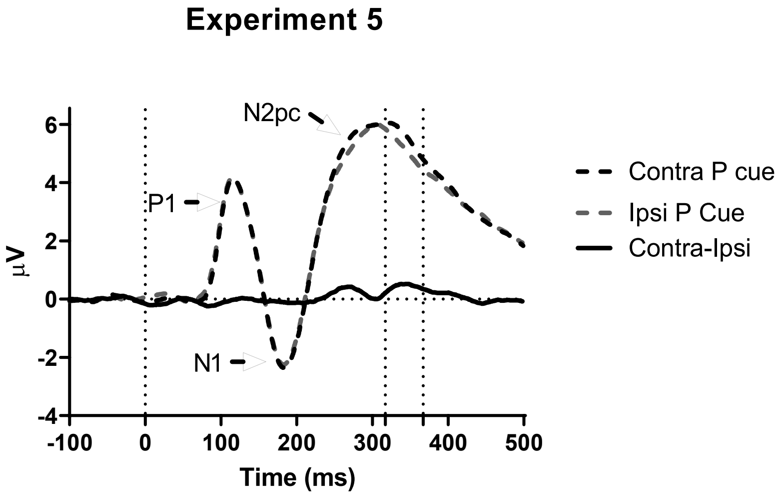

Figure C2: Grand average ERP waveforms at contralateral and ipsilateral (to the P cue) electrode sites (averaged over PO7/PO8 and O1/O2) time-locked to the onset of the cues during the categorisation task in Experiment 4 and Experiment 5 using an extended epoch duration. The vertical dotted lines represent the 50 ms time window used to measure the N2pc amplitude. The contralateral-ipsilateral waveform is also shown.

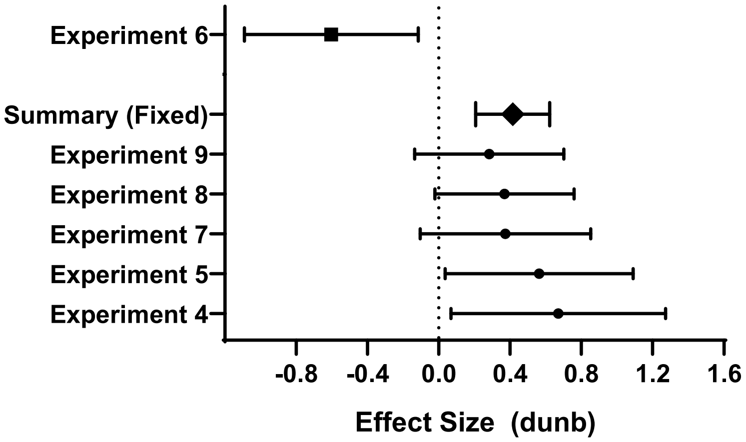

Figure C3: Meta-analysis of N2pc amplitudes from the cue locked data. Negative effect size indicates that attention was directed to the P cue and positive effect size indicates that attention was directed towards the NP cue. Experiment 6 used a conjunction discrimination, while all other experiments used a feature-based discrimination.

It is not entirely clear why the feature-based discrimination would result in a reversed cue locked N2pc compared to the conjunction discrimination. During the categorisation task of all experiments participants had to attend to the P cue in order to make a correct categorisation response. Only participants who could complete the task successfully were included in the EEG analyses. Therefore, it is unlikely that participants were paying attention to the NP cues because these cues could not inform the categorisation response. In addition, the direction of the probe locked N2pc has remained consistent throughout all experiments. That is, the early N2pc time window shows probes eliciting larger amplitudes over NP cues, and the late time window shows the reverse effect. One way to explain the N2pc amplitudes from the cue locked data comes from experiments conducted by Luck et al. (1997). In Experiment 1, Luck et al. had participants complete a conjunction task in which they were presented a green and a red T. Participants had to respond to the orientation (upright or inverted) of the T rendered in a predefined target colour. The same participants also completed Experiment 2, which was a feature discrimination task that involved responding to the presence or absence of a coloured square, while ignoring the square of a different colour. In both experiments the target and distractor were presented for 750 ms on opposite sides of the screen in either the upper or lower visual fields. In addition, on some trials the target and non-target item were surrounded by an additional grey distractor. The N2pc results showed three main effects. Firstly, the N2pc component was larger when additional distractors were present. Secondly, the N2pc component elicited by the conjunction targets was larger and longer lasting compared to the feature targets. Lastly, the N2pc component was larger for arrays presented in the lower visual field. In addition, arrays presented in the upper visual field during the feature detection task elicited N2pc amplitudes that showed a reversed polarity (i.e., electrodes contralateral to the non-target item were more negative than those contralateral to the target).

This last result may be particularly relevant to the cue locked data obtained in our feature discrimination task (for a similar result see Bacigalupo & Luck, 2019). Although our cues were presented on the horizontal midline, it is possible that feature-based targets need to be presented in the lower visual field to elicit N2pc amplitudes in the expected direction. Indeed, in Experiment 3, Luck et al. (1997) conducted another feature detection task and presented all stimuli in the lower visual field to increase the signal to noise ratio of the N2pc component. All components in that experiment were in the expected direction. In addition, the feature-based discrimination conducted by McDonald et al. (2009) in their EEG study on IOR also presented two coloured stimuli in the lower visual field. Therefore, the cue locked N2pc observed in our feature-based discrimination may be reversed because the cues were not presented in the lower visual field. An easy way to test this would be to present our cues in the lower visual field to see if the cue locked N2pc reverses back to the expected direction.

Another important point is related to how our cue-locked analyses calculated the mean amplitude of the N2pc component. In each experiment, the cue-locked analysis calculated the mean amplitude of the N2pc using different 50 ms time windows. These 50 ms time windows were centred on the median of the 50% fractional area latency calculated from each participant. This method is flawed because the mean amplitude calculated in this time window will be biased towards significant results. In other words, if an unbiased time period was chosen for our cue-locked analysis (e.g., an a-priori window of 150-300 ms based on the consensus of the N2pc literature) the N2pc mean amplitudes would not have been statistically different from zero. Therefore, the cue-locked N2pc analyses must be interpreted cautiously (i.e., directionally at best). Note that the probe locked analyses used the same method to calculate the time periods in which the N2pc amplitude was scored. However, our probe locked analyses were not biased because each of our comparisons of interest (e.g., P0 vs. NP0, P100 vs. NP100 etc.) were scored using the same time windows. This method is similar to the "collapsed localizer" approach, which is often the best method when measurement parameters cannot be set based on prior research (Luck and Gaspelin, 2017).

In summary, the reversed cue locked N2pc found using our feature discrimination tasks can be explained by cue location. In addition, biased time-windows were used when calculating the cue-locked N2pc amplitudes Therefore, the author still finds it reasonable to interpret the early probe locked N2pc as attention towards the P cue. The behavioural data (RTs, dot probe errors, premature responses, ratings) and SSVEP amplitudes from the conjunction discrimination all indicate that attention was directed towards the P cue. In addition, the early N2pc time window was extended using a conjunction task, a result that would only be expected if participants were paying attention to the P cue for longer. Therefore, it seems appropriate to attribute the earlier and larger probe locked N2pc effect observed in the current experiments as reflecting attention towards the P cues. However, future research may be required to settle the issue.

Appendix D: Latencies, CS onset and probe locked ERPs at each SOA for Experiments 7

Table D1

CS onset N2pc median 50% Fractional Area Latencies (ms) in Experiment 8

SOA	Experiment 8
PO7/PO8	392
P7/P8	390
O1/O2	384

Table D2

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 8

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms	1000 ms
PO7/PO8	247	214	193	189	187	185	182
P7/P8	252	221	196	188	187	187	183
O1/O2	244	209	193	188	186	186	182

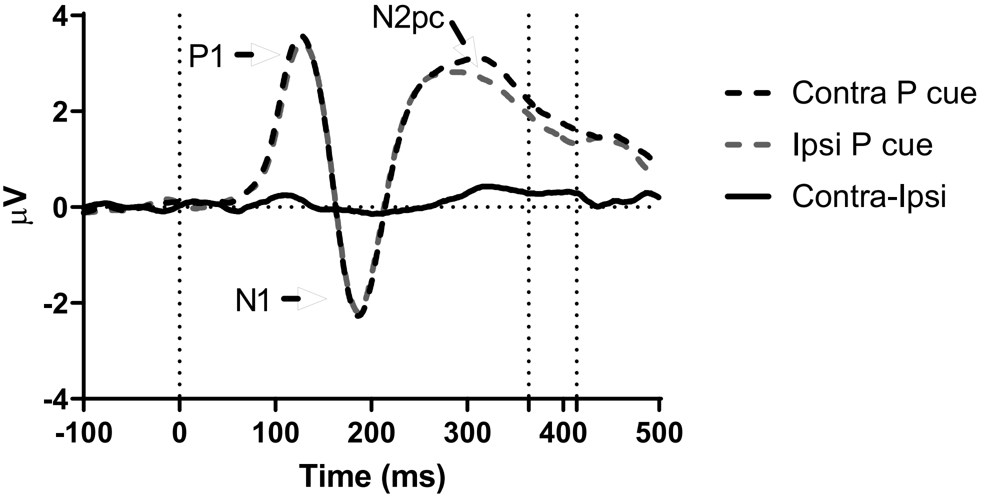

Figure D1: Grand average ERP waveforms at contralateral and ipsilateral (to the P cue) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the cues during the categorisation task in Experiment 7. The contralateral-ipsilateral waveform is also shown.

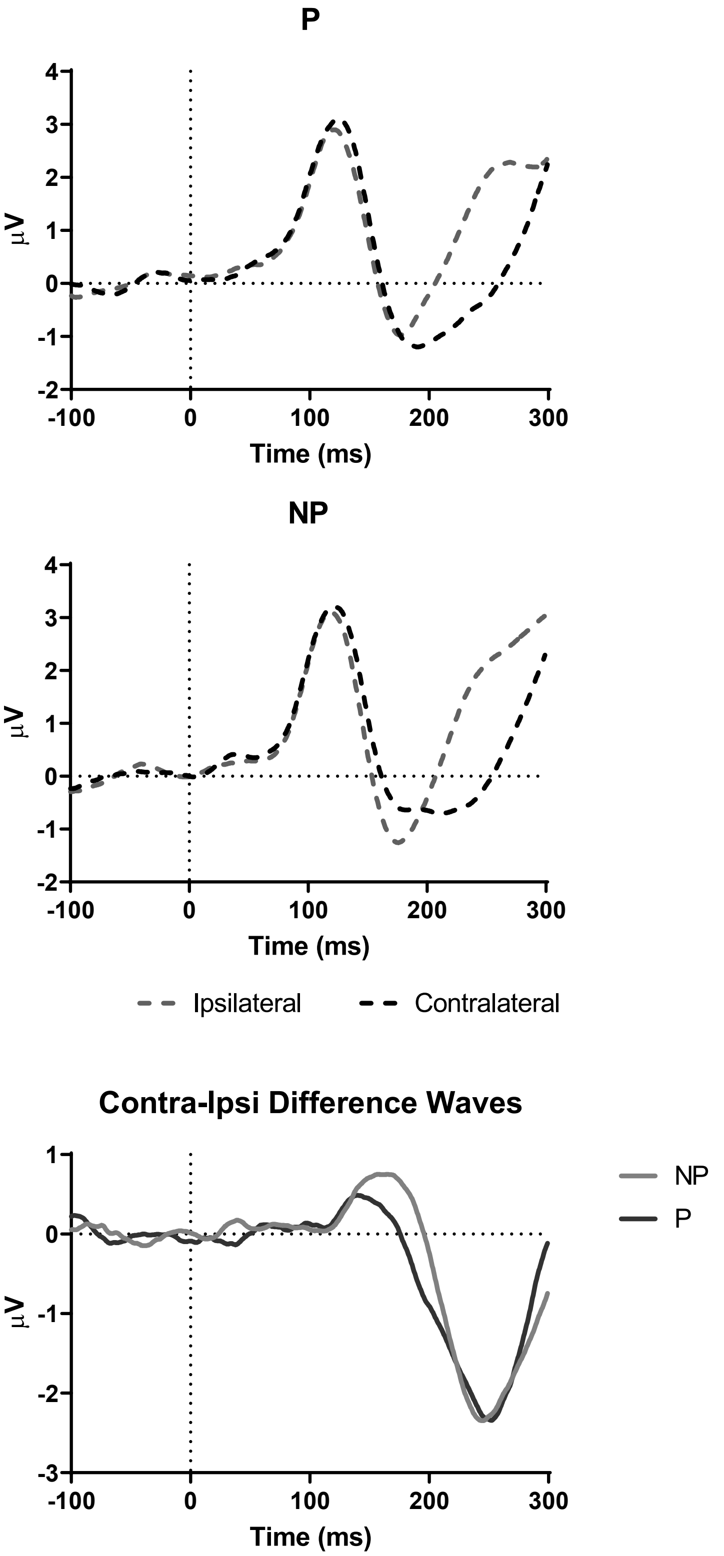

Figure D2: Grand average ERP waveforms for the 0 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

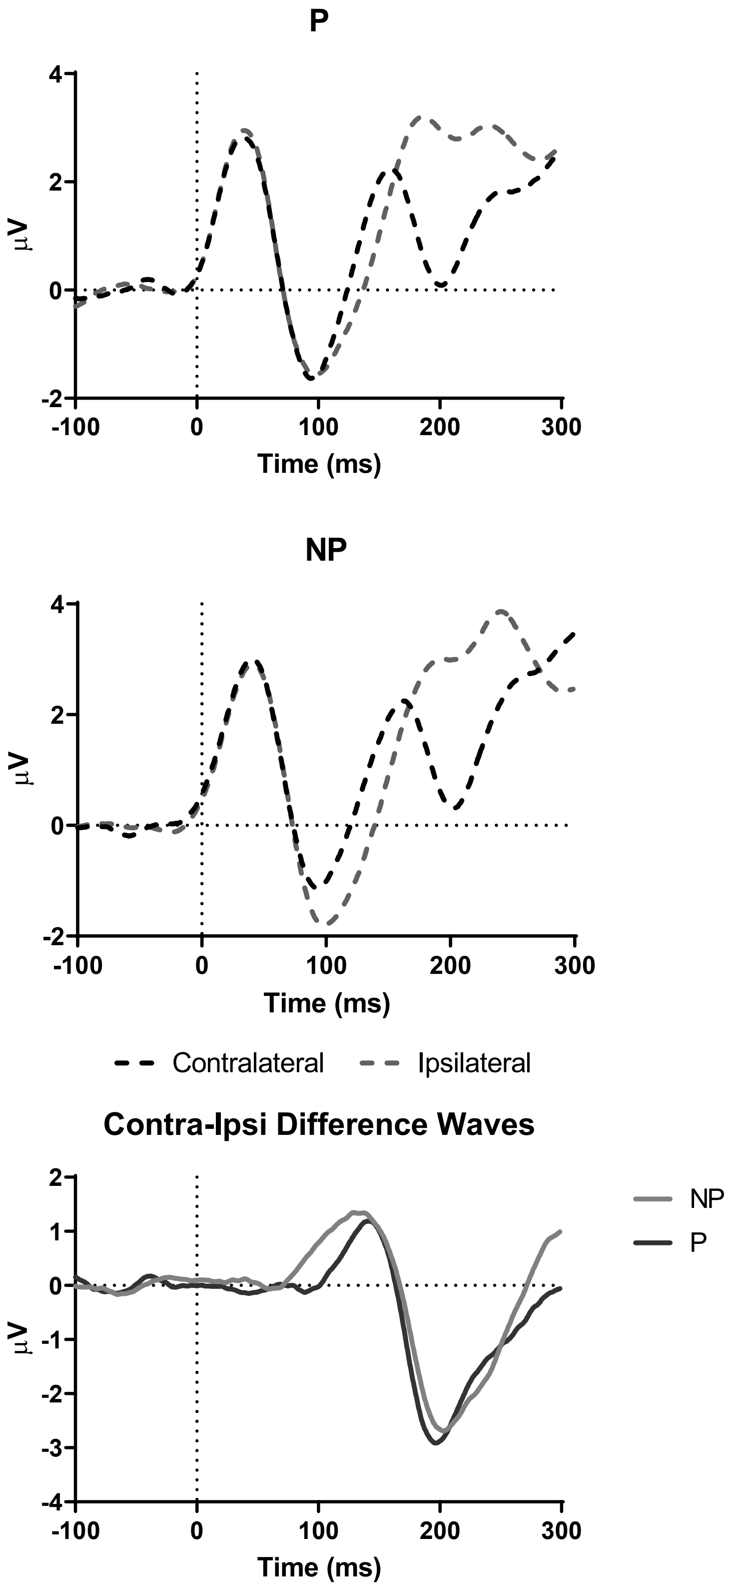

Figure D3: Grand average ERP waveforms for the 100 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

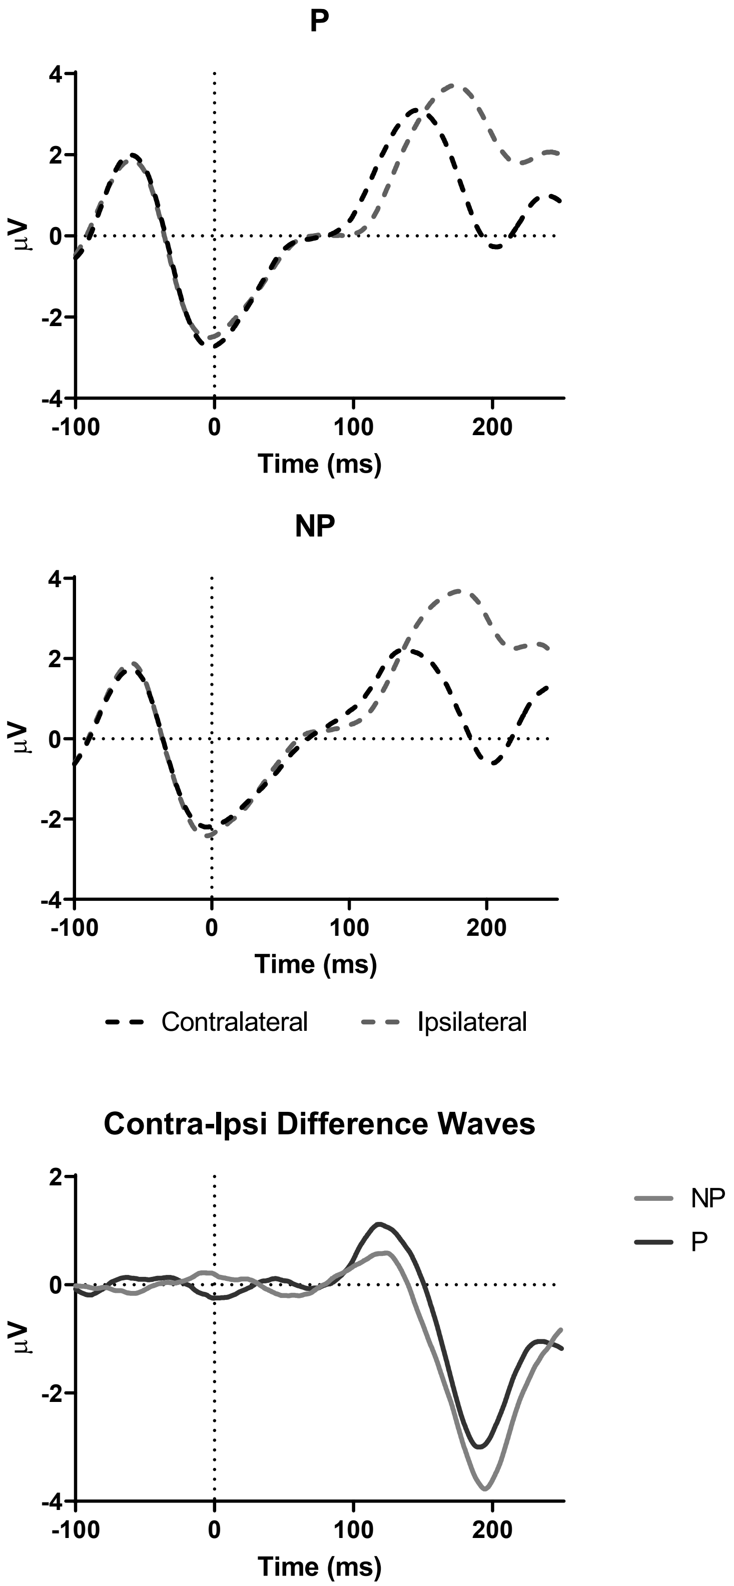

Figure D4: Grand average ERP waveforms for the 200 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

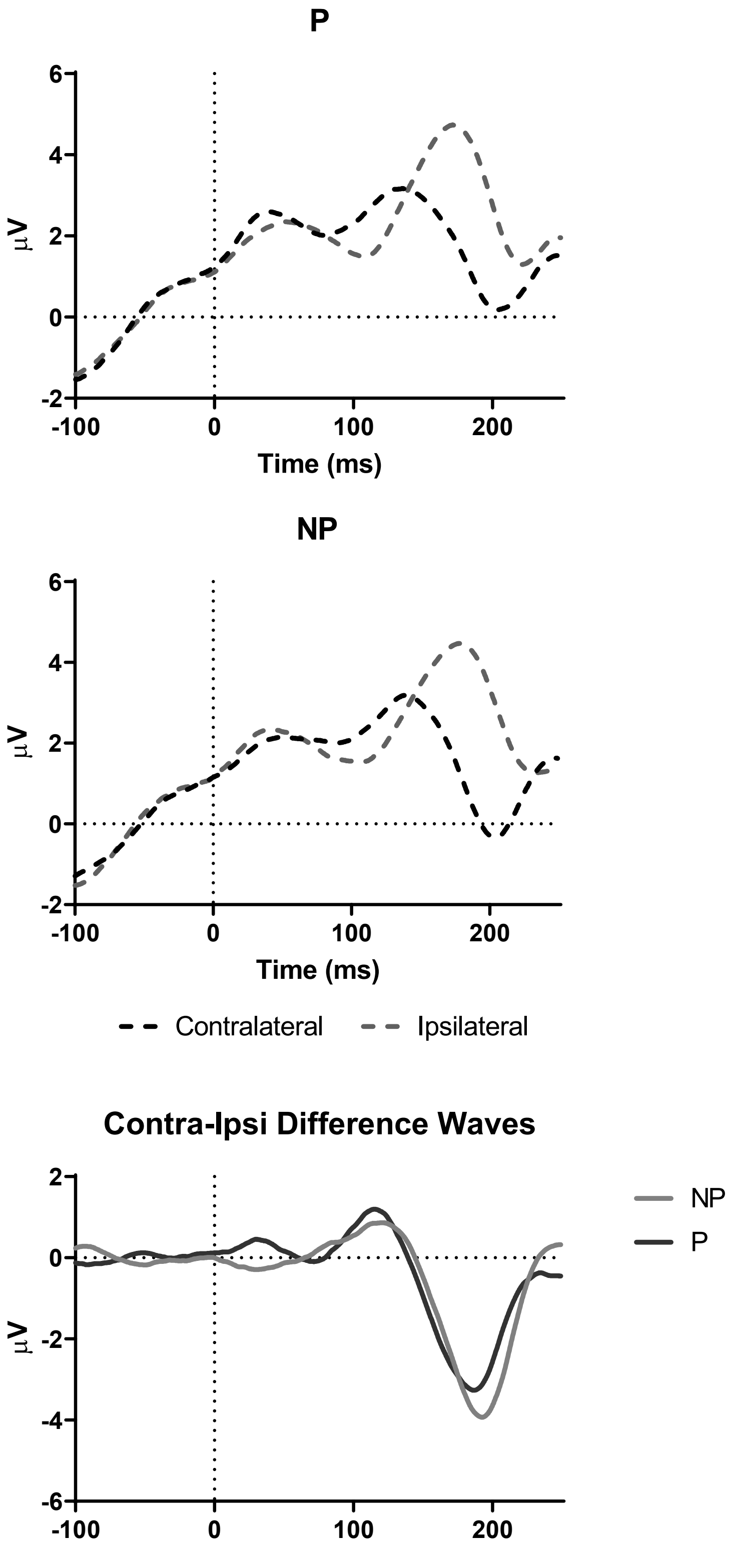

Figure D5: Grand average ERP waveforms for the 300 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

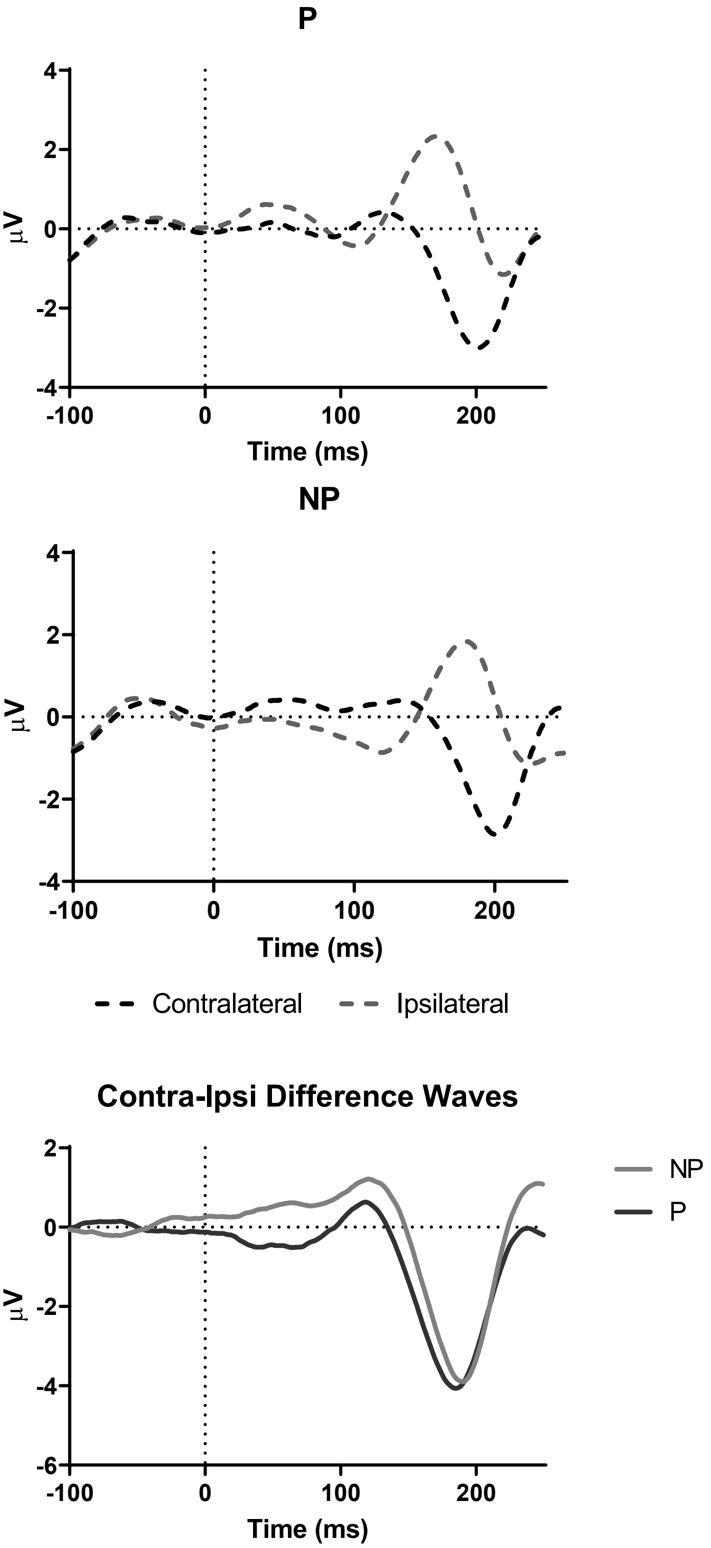

P

NP

Contra-Ipsi Difference Waves

Figure D6: Grand average ERP waveforms for the 400 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

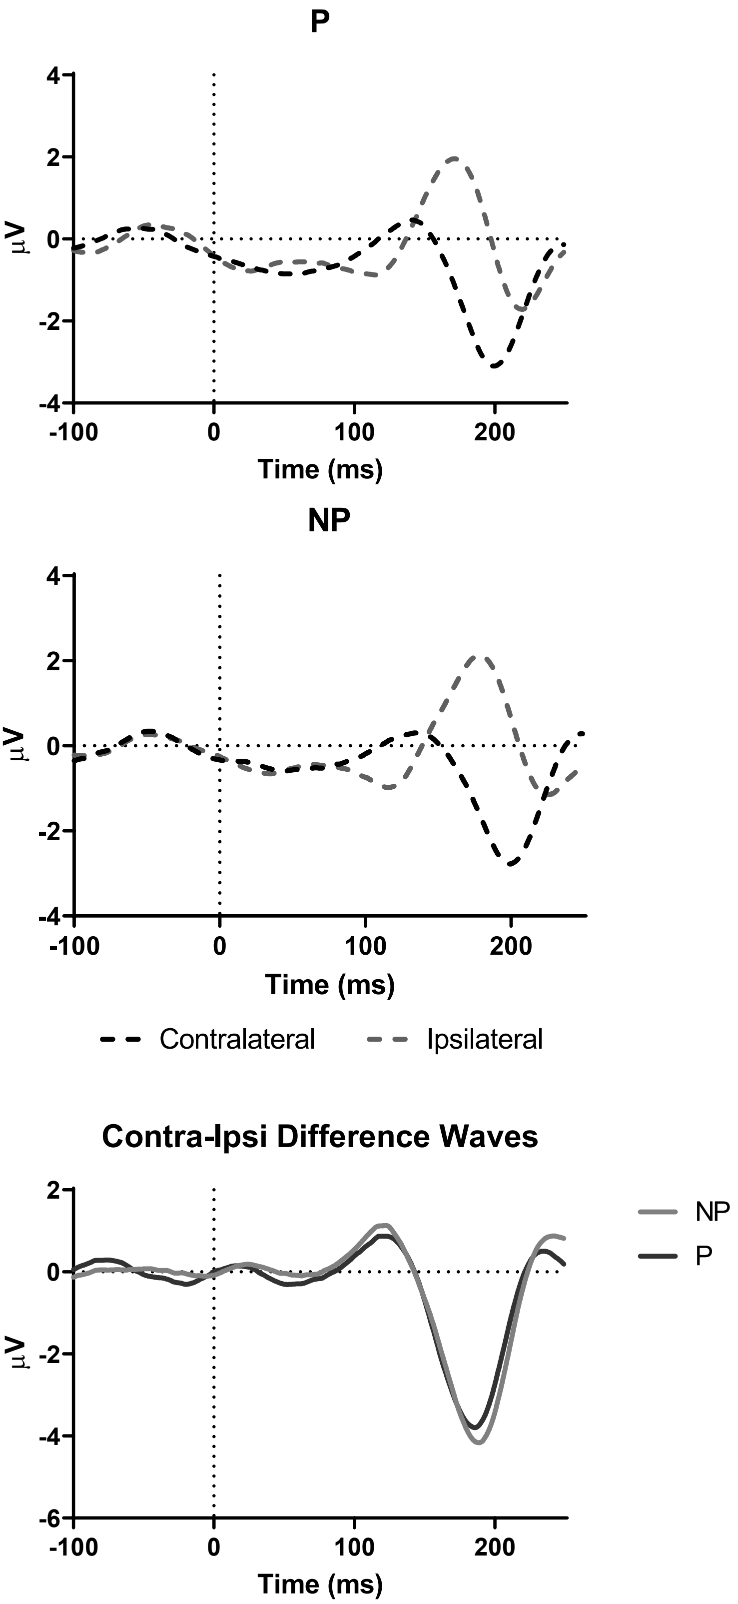

Figure D7: Grand average ERP waveforms for the 500 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

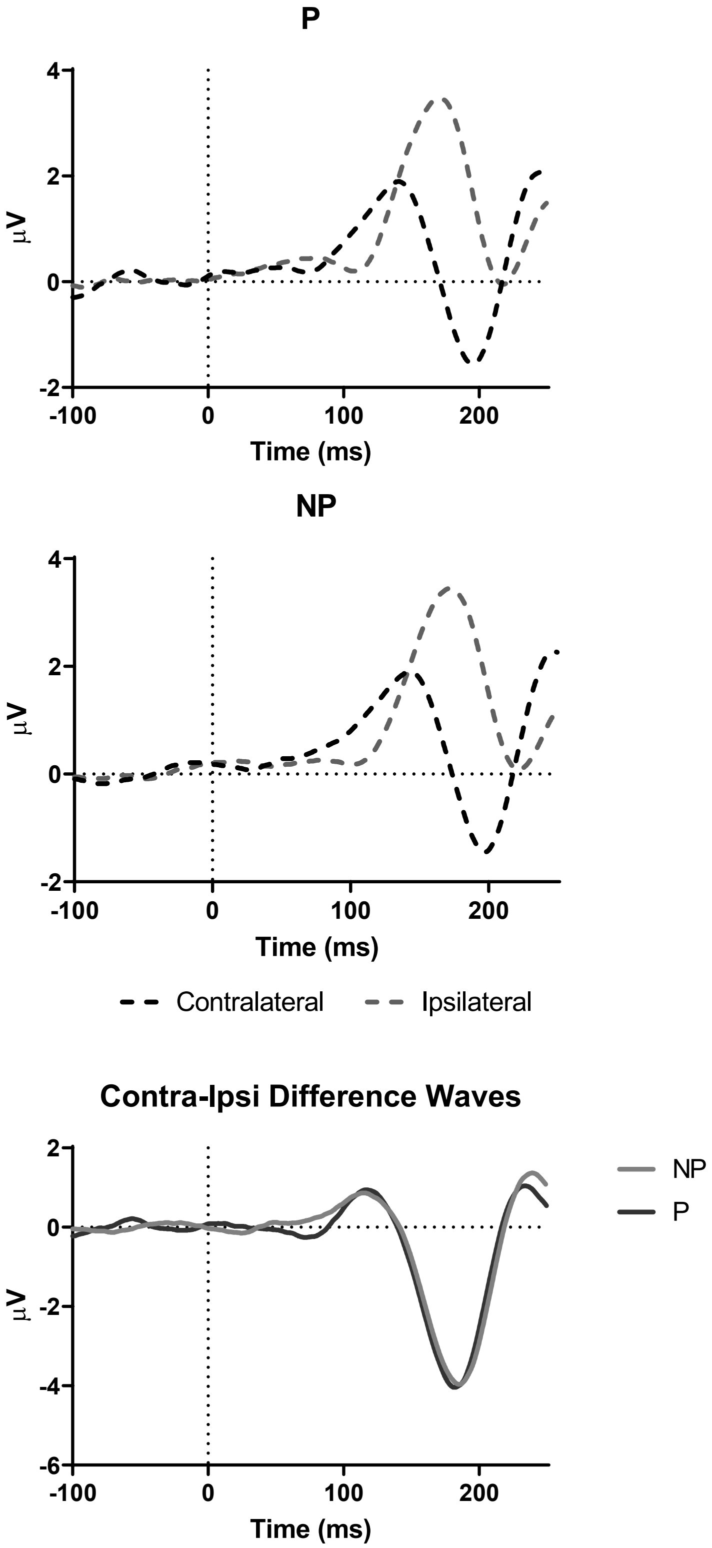

Figure D8: Grand average ERP waveforms for the 1000 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

Appendix E: Latencies, CS onset and probe locked ERPs at each SOA for Experiment 8

Table E1

CS onset N2pc median 50% Fractional Area Latencies (ms) in Experiment 7

SOA	Experiment 7
PO7/PO8	314
P7/P8	318
O1/O2	341

Table E2

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 7

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms
PO7/PO8	239	213	197	191	190	192
P7/P8	246	224	198	189	190	190
O1/O2	235	209	196	190	191	191

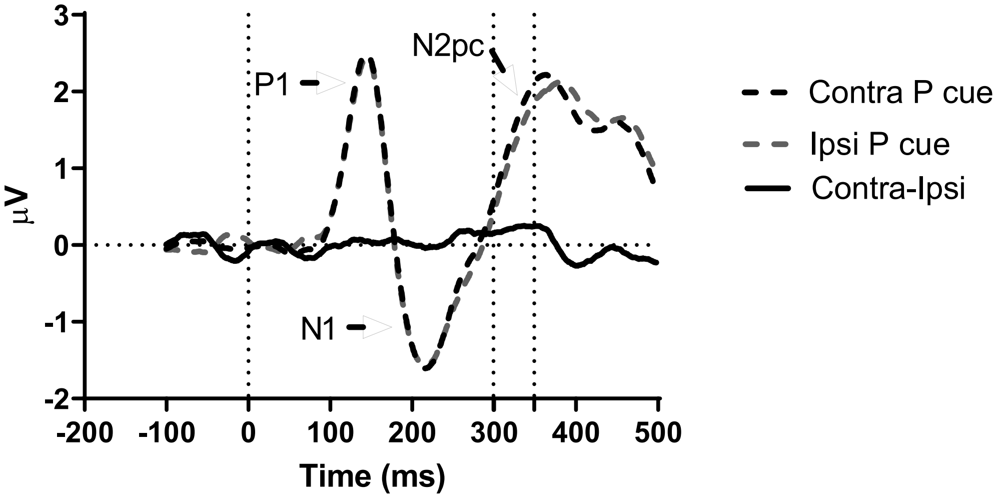

Figure E1: Grand average ERP waveforms at contralateral and ipsilateral (to the P cue) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the cues during the categorisation task in Experiment 8. The contralateral-ipsilateral waveform is also shown.

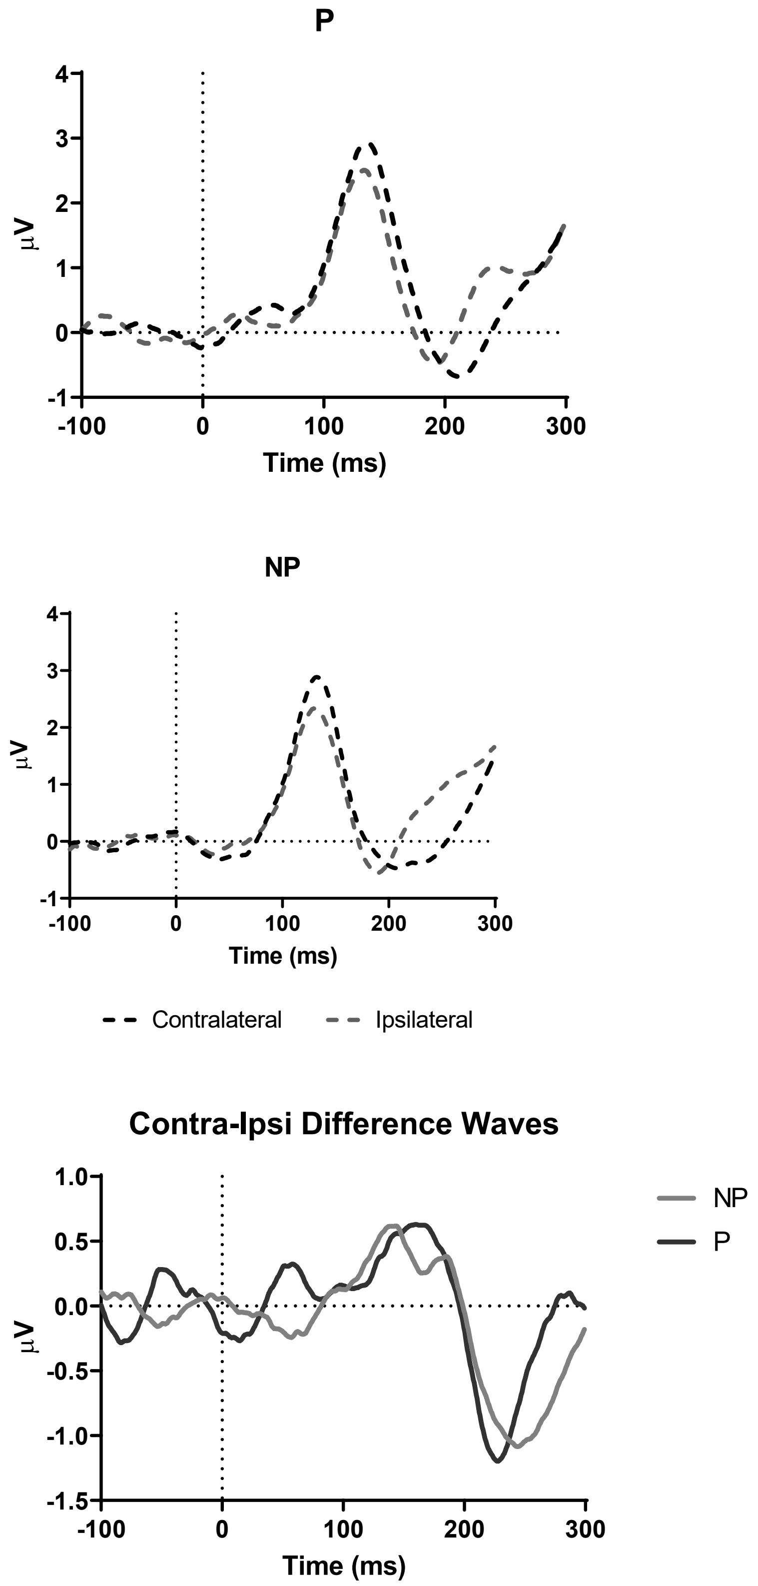

Figure E2: Grand average ERP waveforms for the 0 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

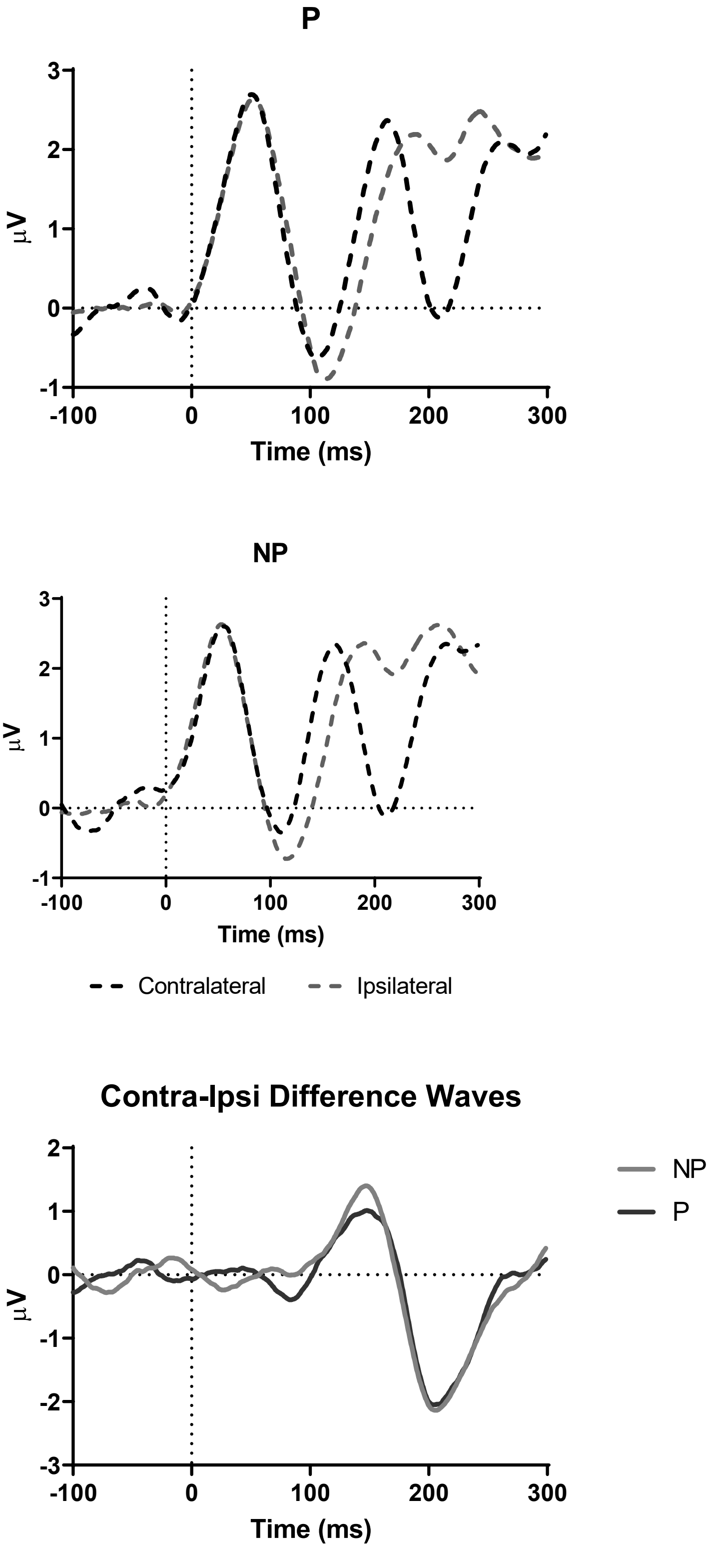

Figure E3: Grand average ERP waveforms for the 100 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

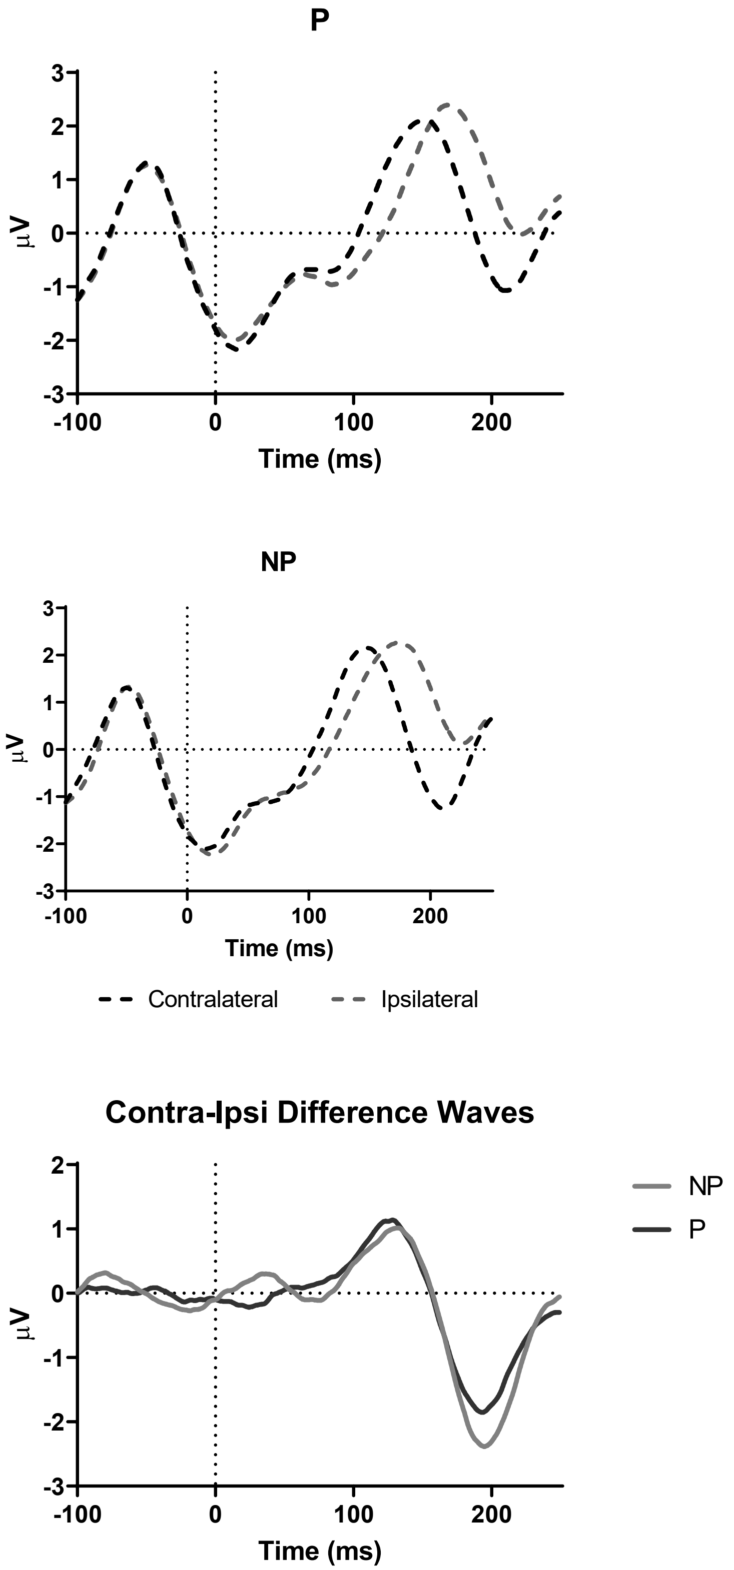

Figure E4: Grand average ERP waveforms for the 200 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

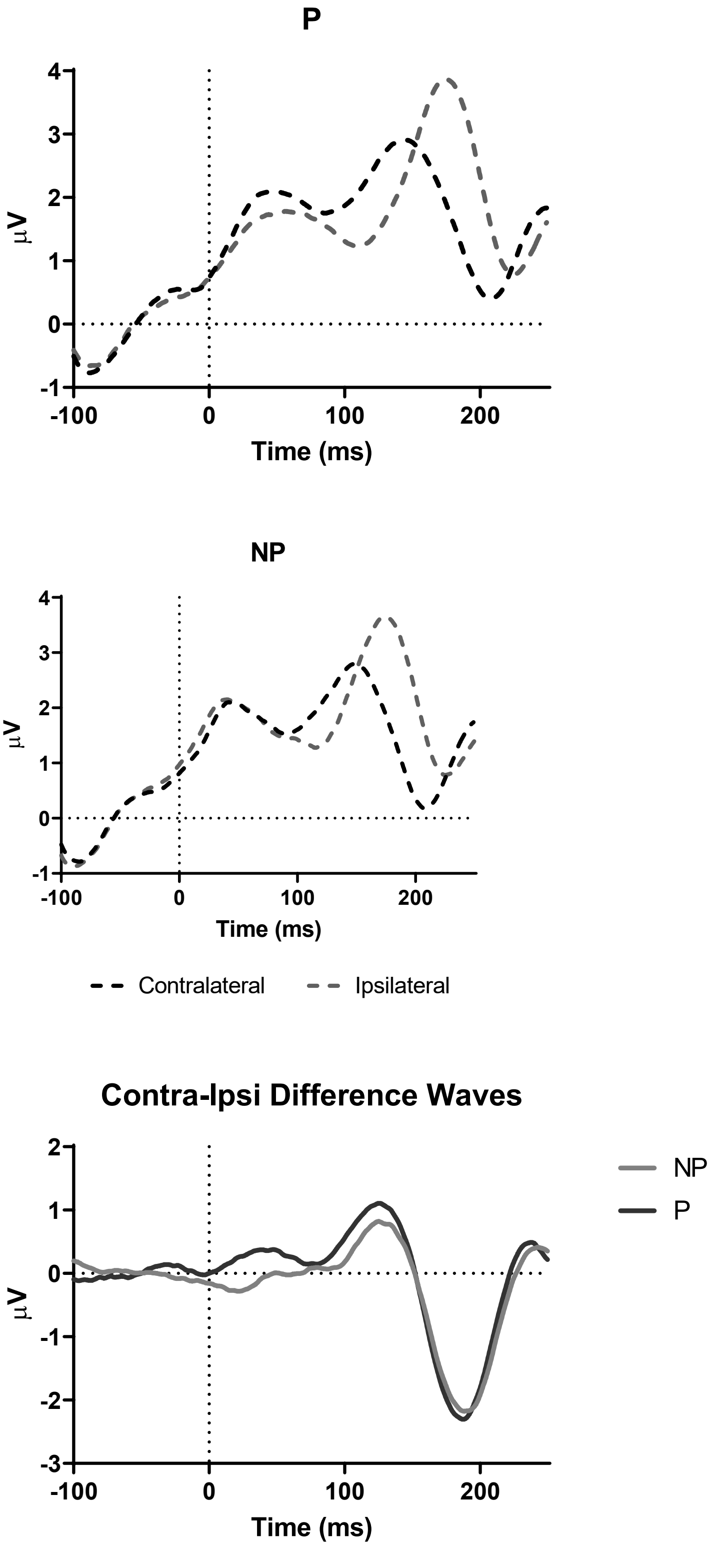

Figure E5: Grand average ERP waveforms for the 300 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

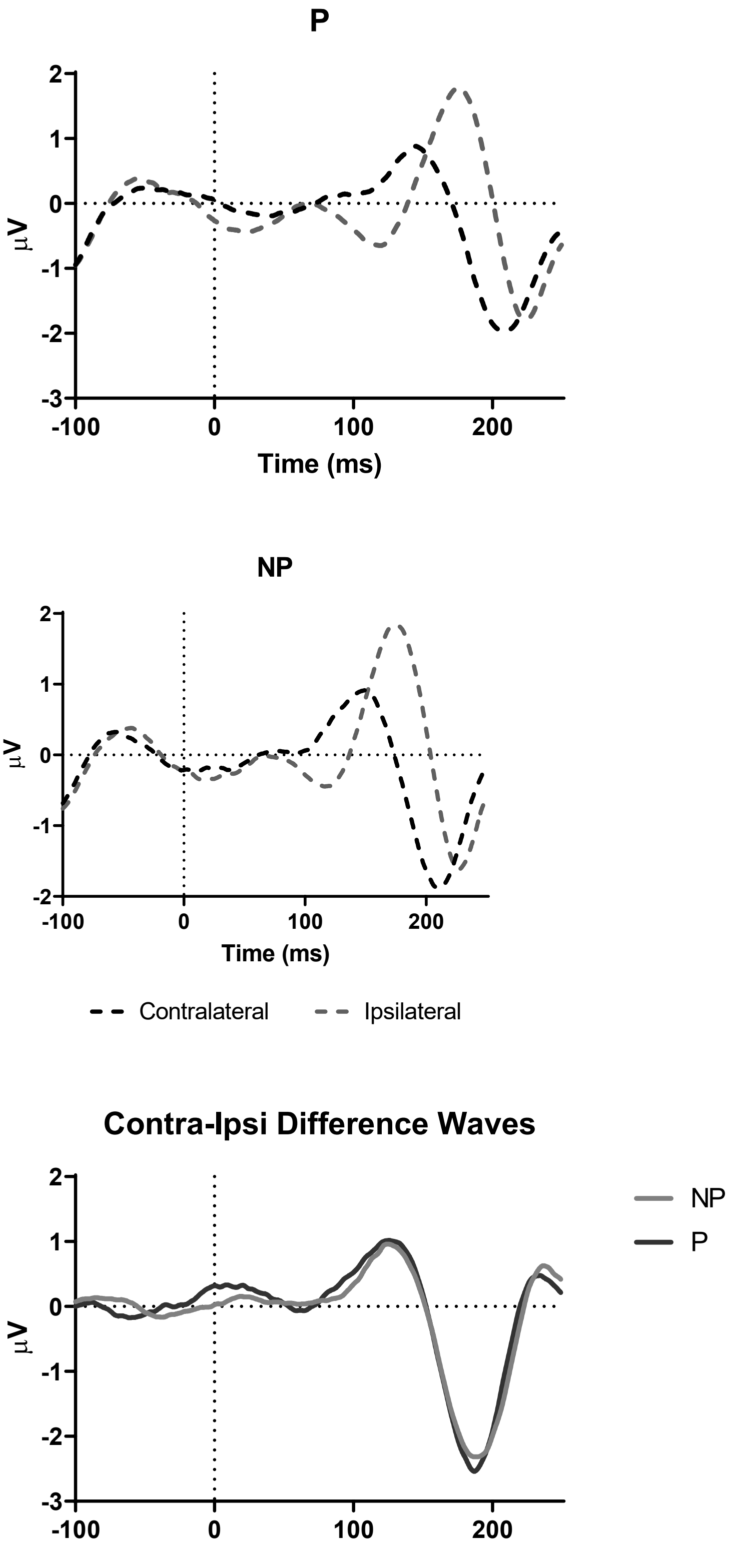

Figure E6: Grand average ERP waveforms for the 400 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

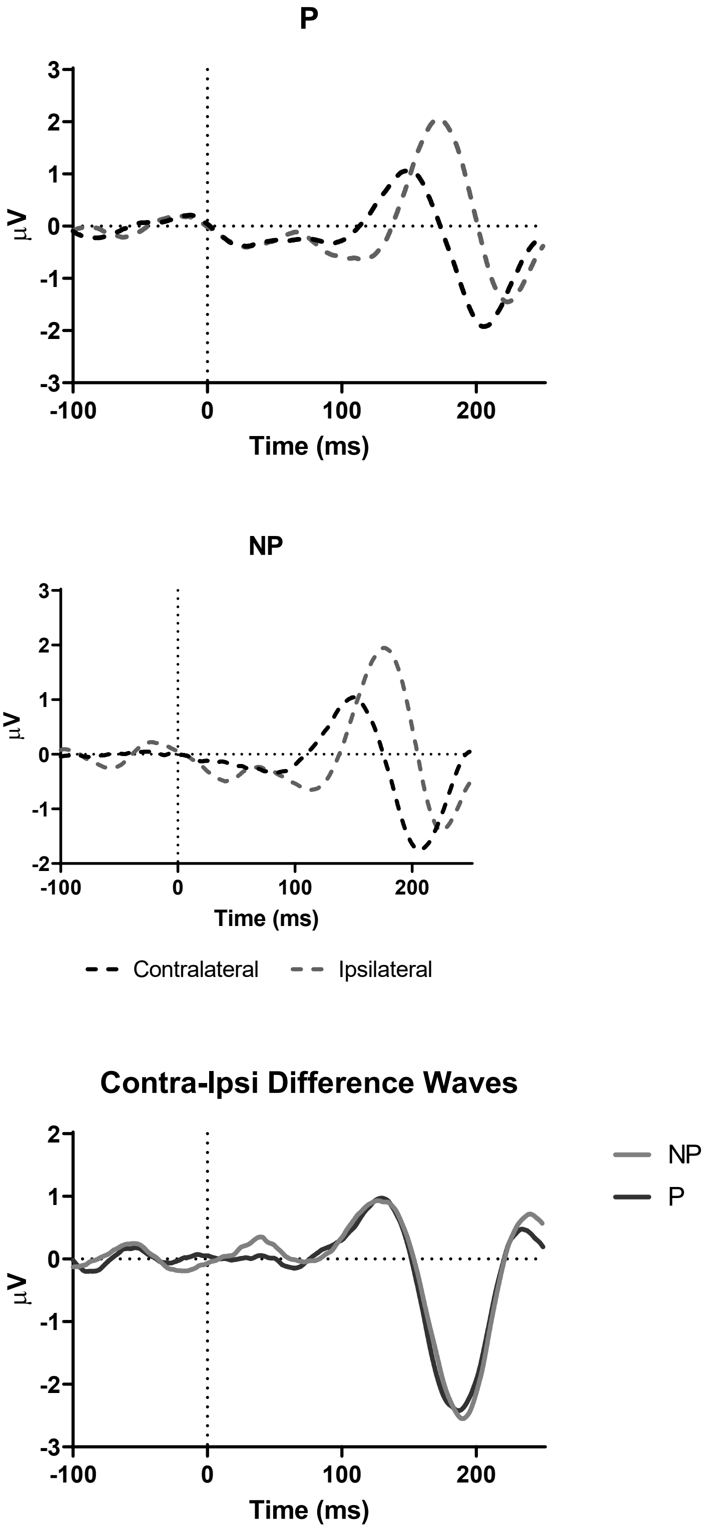

Figure E7: Grand average ERP waveforms for the 500 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

Appendix F: Latencies, CS onset and probe locked ERPs at each SOA for Experiment 9

Table F1

CS onset N2pc median 50% Fractional Area Latencies (ms) in Experiment 9

SOA	Experiment 9
PO7/PO8	349
P7/P8	365
O1/O2	353

Table F2

Dot Probe N2pc median 50% Fractional Area Latencies (ms) in Experiment 9

SOA	0 ms	100 ms	200 ms	300 ms	400 ms	500 ms
PO7/PO8	242	218	200	187	184	185
P7/P8	248	224	201	188	185	186
O1/O2	240	212	198	188	185	187

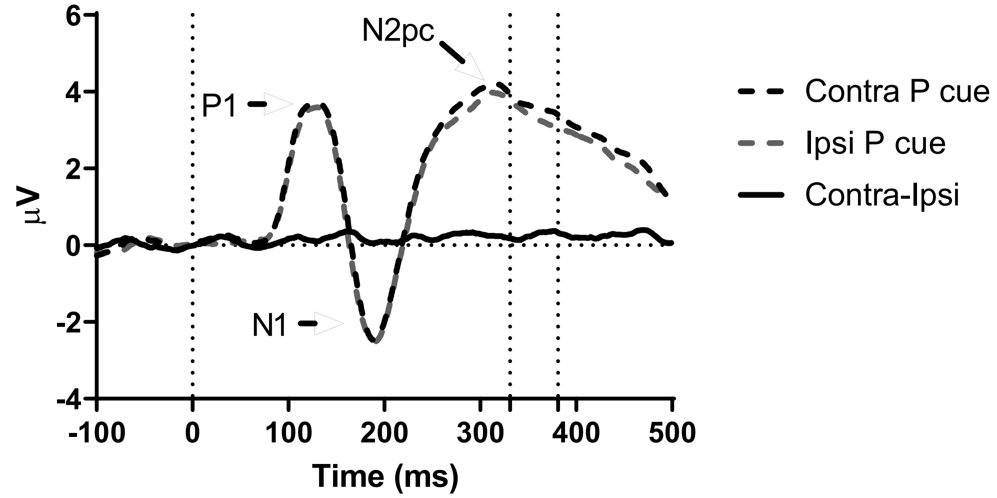

Figure F1: Grand average ERP waveforms at contralateral and ipsilateral (to the P cue) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the cues during the categorisation task in Experiment 9. The contralateral-ipsilateral waveform is also shown.

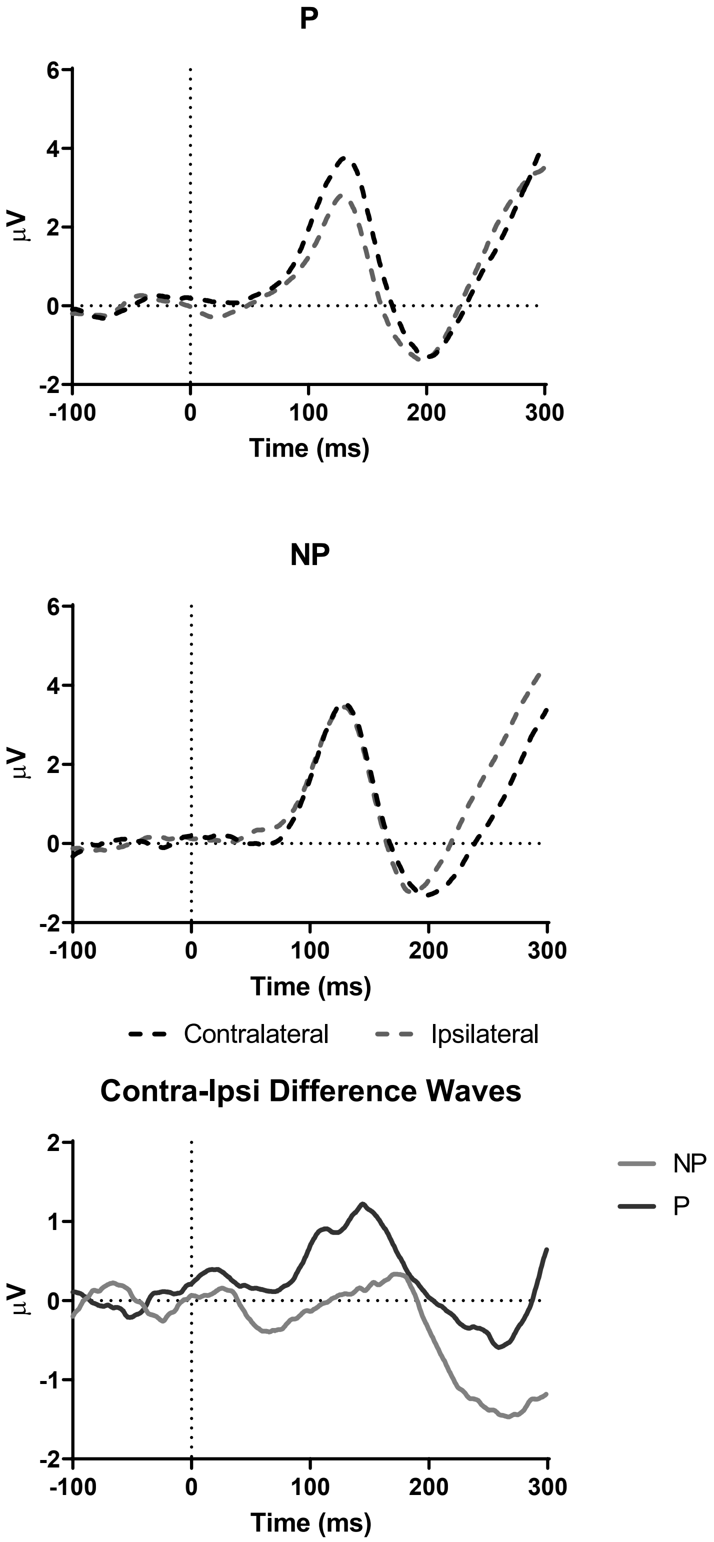

Figure F2: Grand average ERP waveforms for the 0 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

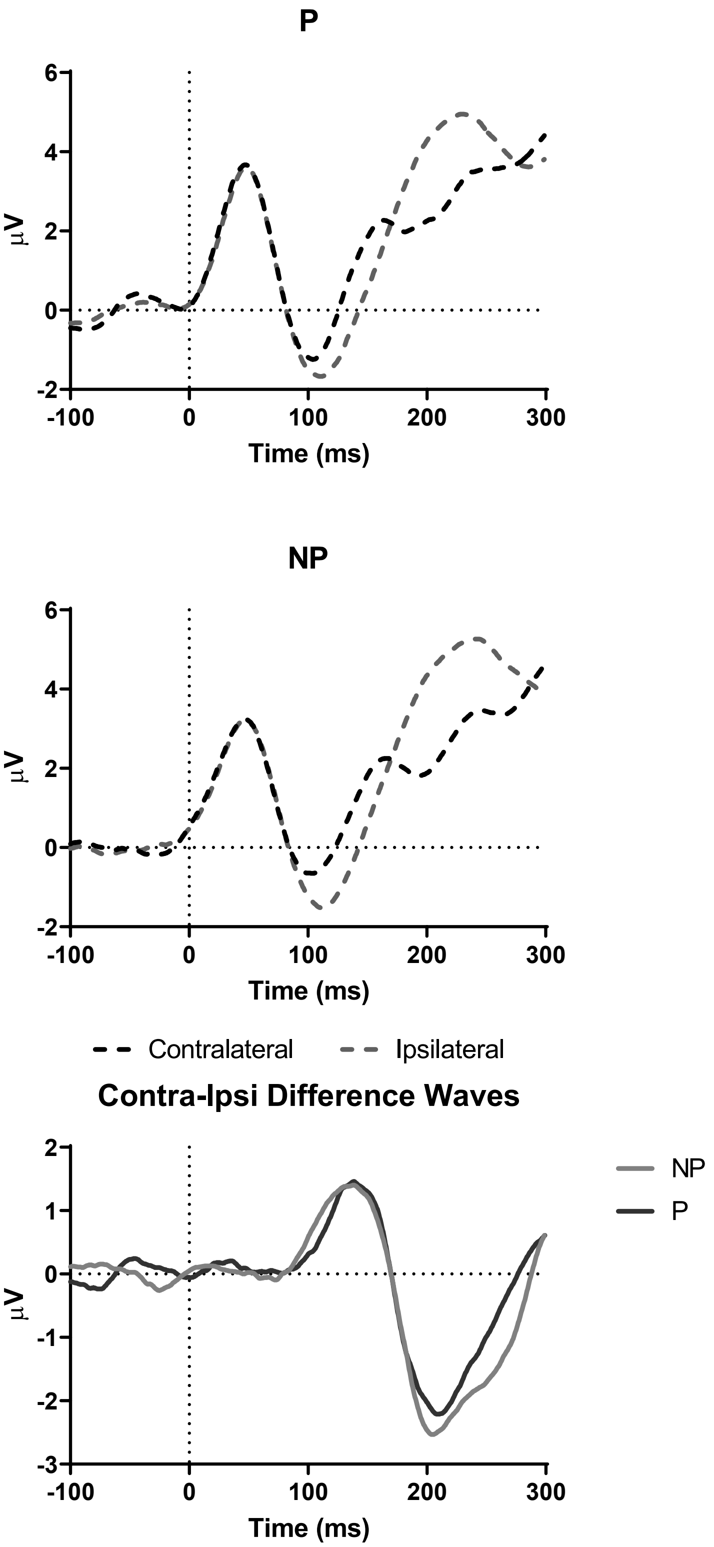

Figure F3: Grand average ERP waveforms for the 100 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

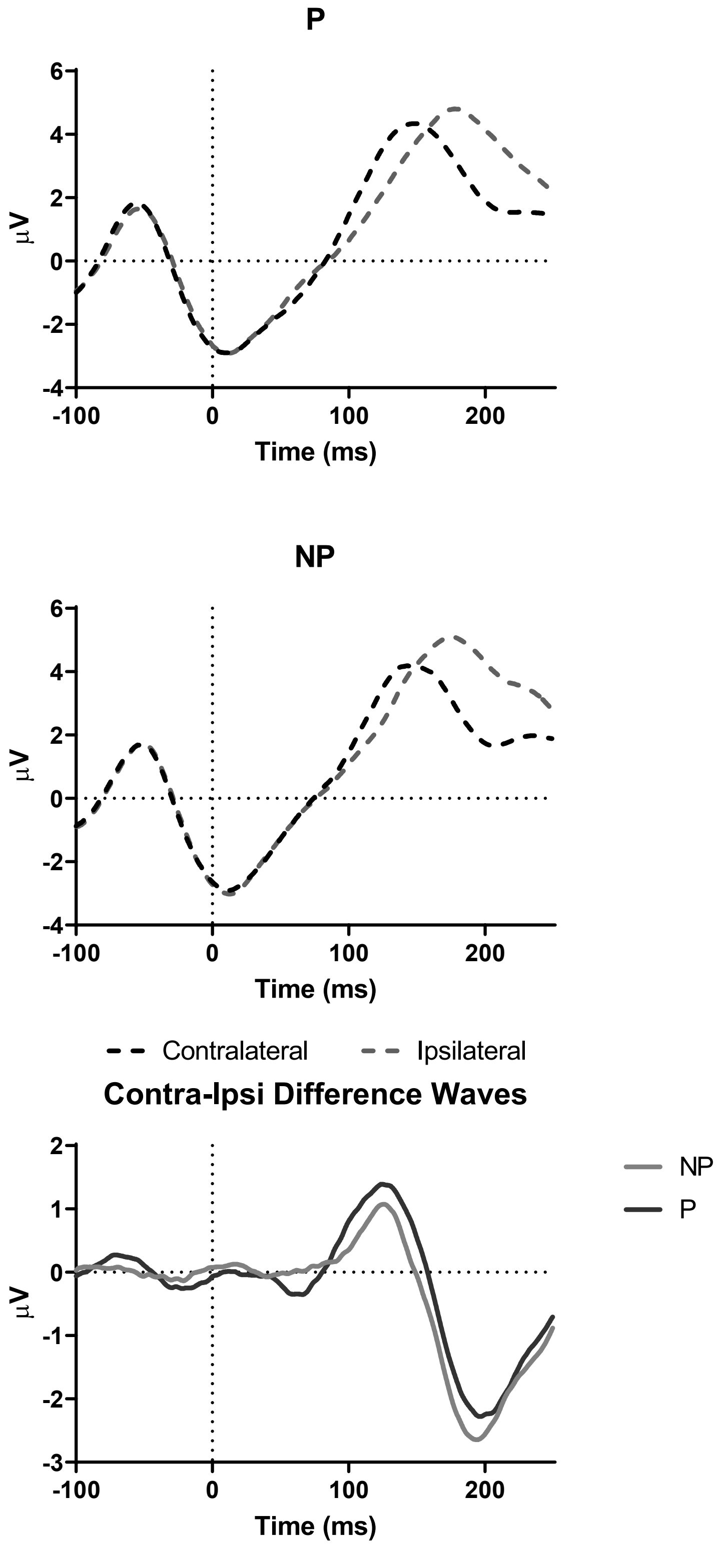

Figure F4: Grand average ERP waveforms for the 200 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

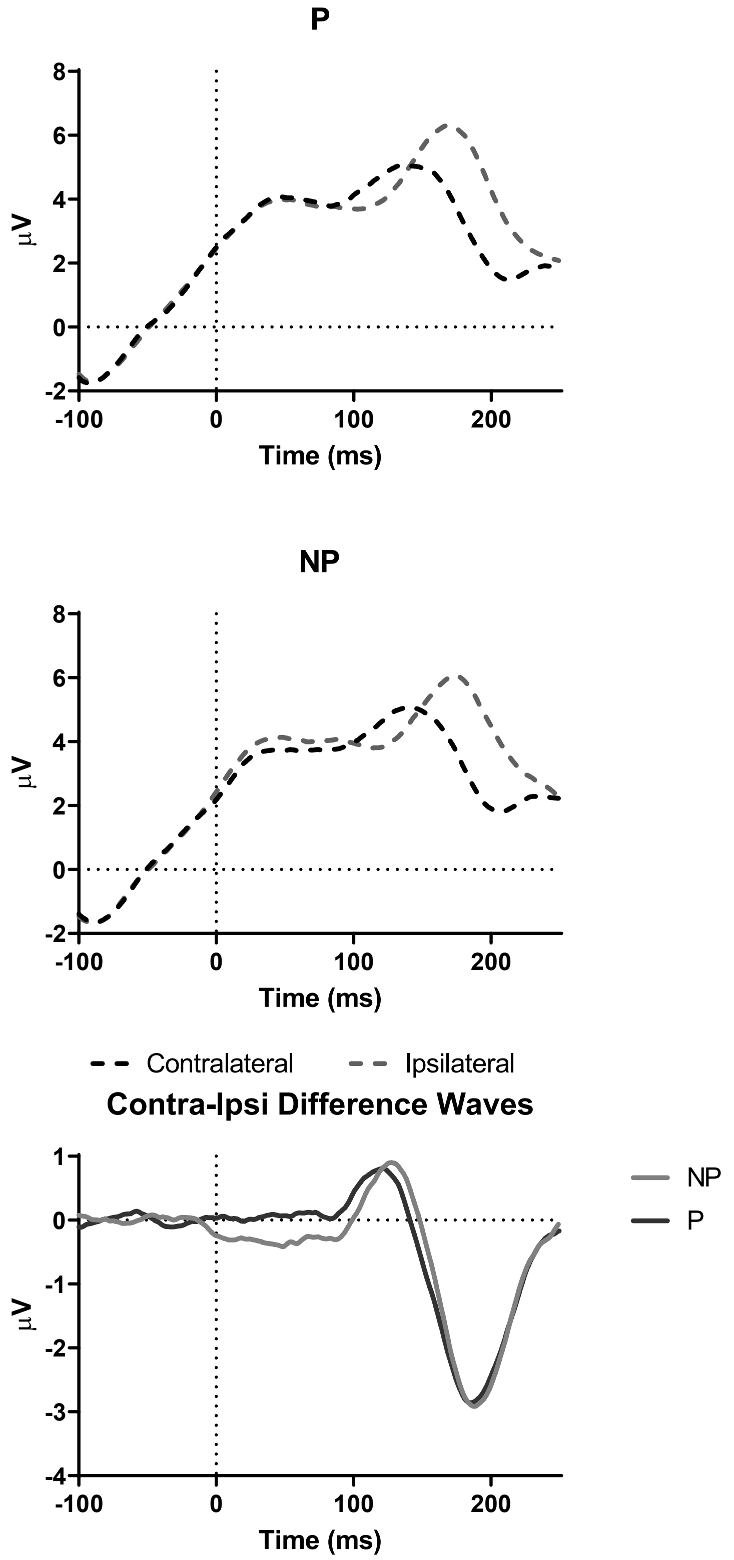

Figure F5: Grand average ERP waveforms for the 300 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

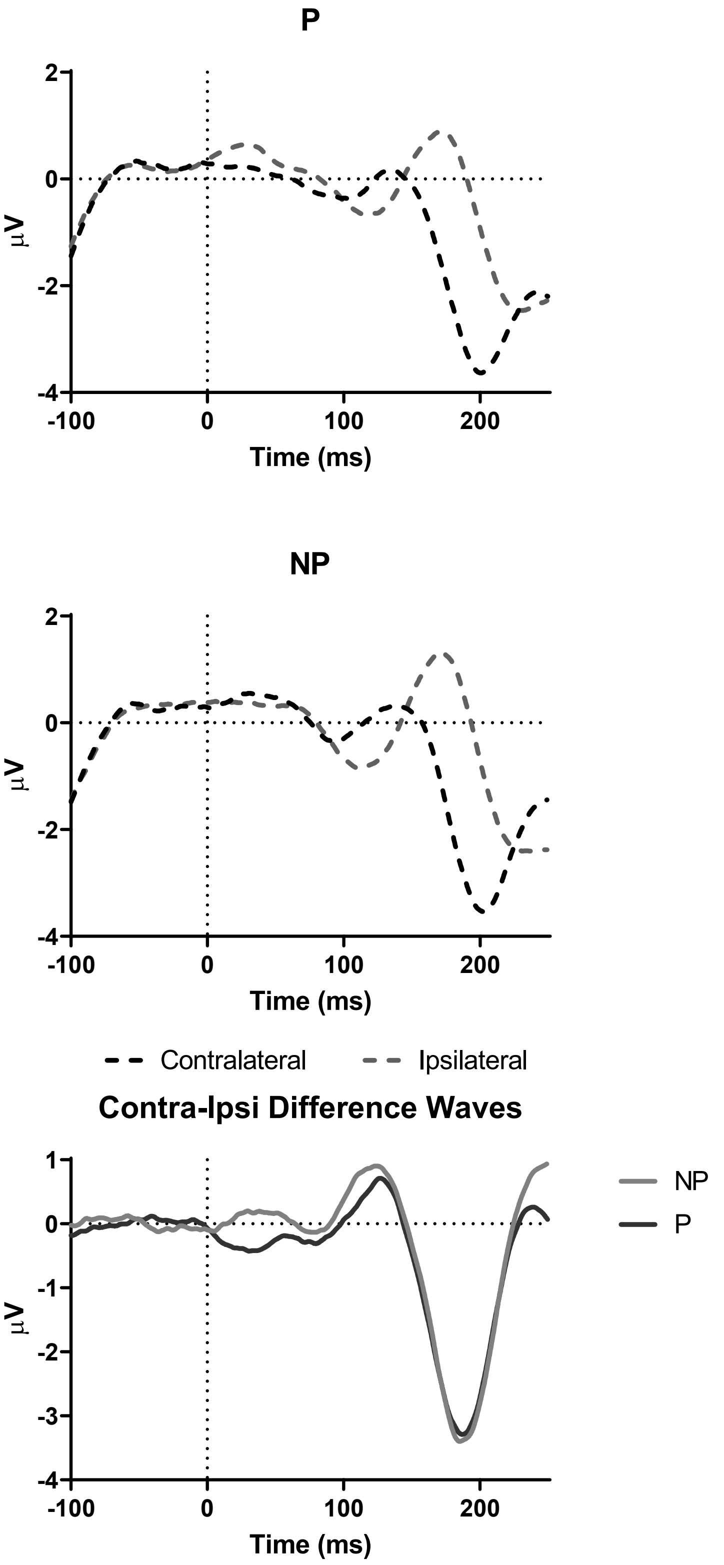

Figure F6: Grand average ERP waveforms for the 400 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

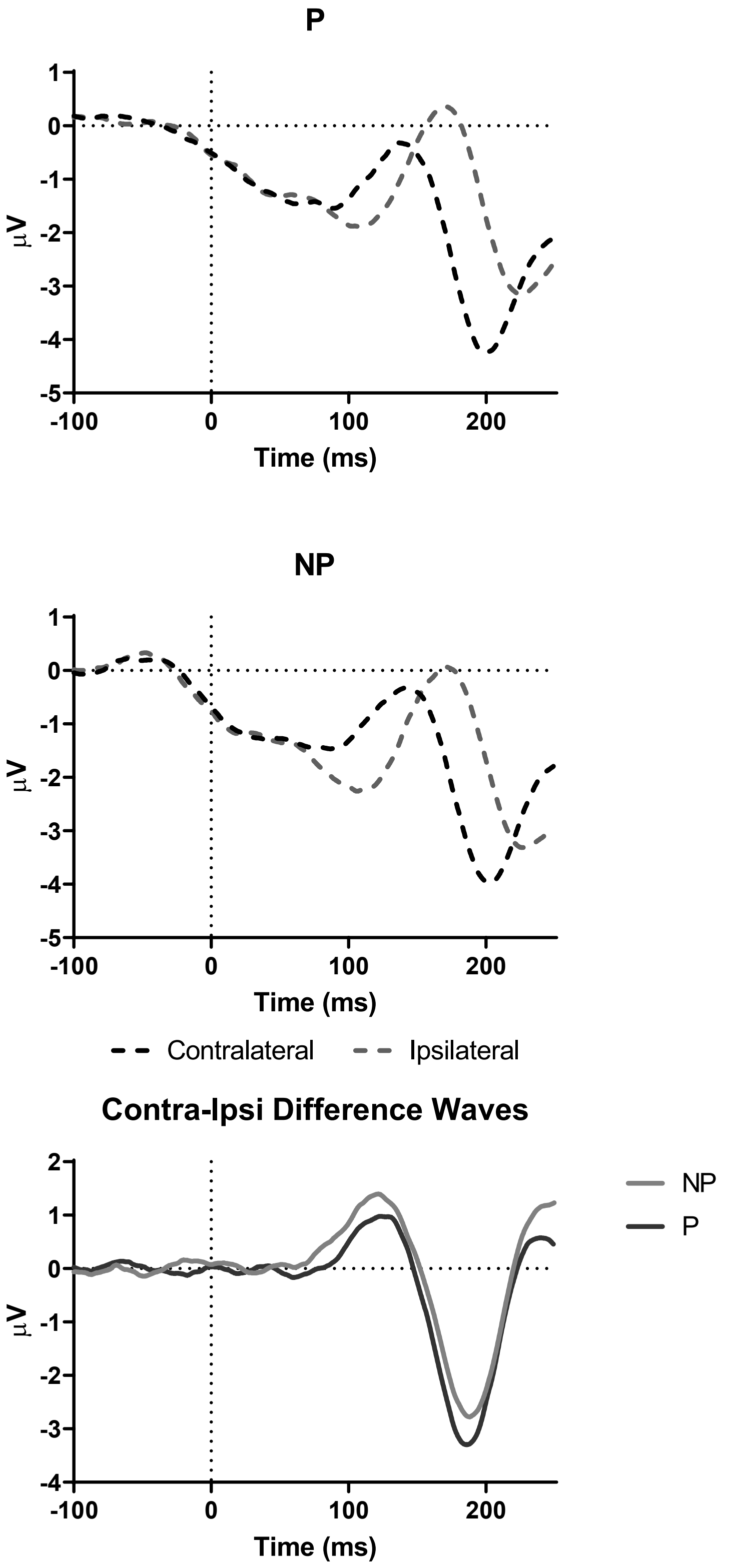

Figure F7: Grand average ERP waveforms for the 500 ms SOA condition at contralateral and ipsilateral (to the dot probe) electrode sites (averaged over PO7/PO8, P7/P8 and O1/O2) time-locked to the onset of the white square target over the P (top) and NP (middle) cue. Contralateral minus ipsilateral difference waves are also shown (bottom).

Appendix G: Attention Modulated Auto-Associator

For simplicity, we refer to the US units as *i* and the cue units as *j*. But note that the model updates the activations of all units and their associations in the same way, regardless of whether they are US or cue units.

The internal input to a US unit *i* ($intinput_i$) is the net influence that all active cue units (*j*) have on the US unit. Every active cue unit *j* might influence the input activation to US unit *i* as a function of the connection strength it has to the US unit (w_{ij}) and its own activation level (a_j).

$$intinput_i = \sum w_{ij} a_j$$

The net input to the US unit *i* ($netinput_i$) is the sum of the external input the unit receives ($extinput_i$) and the internal input from other cue units ($intinput_i$). The external input is 1 if the US is physically present and 0 if the US is absent. The external and internal inputs are sized by two parameters (*estr* and *istr*) that range from 0 to 1.

$$netinput_i = (estr)extinput_i + (istr)intinput_i$$

If $netinput_i$ is positive, then the change in activation of unit *i* (Δa_i) on that cycle is

$$\Delta a_i = netinput_i(\max - a_i) - (decay)a_{i,}$$

and if $netinput_i$ is negative, then the change in activation of unit *i* is

$$\Delta a_i = netinput_i(a_i - \min) - (decay)a_i,$$

where *min* and *max* determine the minimum and maximum activation values, which are set to -1 and 1, respectively, and *decay* is the decay rate of the activation (i.e., how fast the activation levels returns to zero).

On the last cycle of US activation, the connections between the active CS units and the US unit are updated. The weight of the connection from CS *j* to US *i* (w_{ij}) changes according to the following rule

$$\Delta w_{ij} = lrate(extinput_i - intinput_i)a_j$$

where *lrate* is a learning rate parameter. Expanding on the original model, we included attention parameters that modulate the external input of CS units depending on the current cycle (i.e., time step) of the cue duration.

$$extinput_j = alphaM(\frac{TP-cycles}{TP}) + alphaPH(\frac{cycles}{TP})$$

In the above equation, *alphaM* and *alphaPH* parameters are Mackintosh and Pearce-Hall attention parameters, respectively. *TP* refers to a constant timing parameter that controls how fast attention switches between Mackintosh and Pearce-Hall mechanisms. On the first cycle of the cue duration, the external input of the cues is weighted mostly according to their *alphaM*. However, as cue duration progresses, the external input is gradually weighted more by *alphaPH* (and less so by *alphaM*). This mechanism results in attention initially shifting towards the most predictive cues (i.e., higher *alphaM*) and subsequently shifting towards non-predictive cues (i.e., higher *alphaPH*) if the cue duration is long enough.

Following the updating of the learning weights, *alphaM* and *alphaPH* are also updated. The *alphaM* parameter is updated according to the following formula

$$alphaM = thetaM(\frac{PEk-PEj+1}{2}) + (1-thetaM)(alphaM_{n-1})$$

where PE_j is the prediction error currently assigned to cue j

$$PE_j = |1 - (w_{ij})(a_j)|,$$

PE_k is the prediction error assigned to all other active cue units, k (i.e., not unit j)

$$PE_k = |1 - \sum(w_{kj})(\ a_k)|,$$

and *thetaM* is a parameter (which can vary between 0 and 1) that determines the extent to which *alphaM* changes with respect to its previous value ($alphaM_{n-1}$). If *thetaM* is 1 then the immediately preceding trial entirely determines the value of *alphaM*.

The *alphaPH* parameter is updated according to the following formula

$$alphaPH = gammaPH(PE_j) + (1\text{-}gammaPH)(alphaPH_{n-1})$$

where PE_j is the prediction error assigned to cue j (as above) and *gammaPH* is a parameter (which can vary between 0 and 1) that determines the extent to which *aplhaPH* changes with respect to its previous value ($alphaPH_{n-1}$).

In our simulations, the parameters *estr* and *istr* were set to 0.15. The *decay* parameter was set to 0.05 and the *lrate* was set 0.3. The *thetaM* and *gammaPH* parameters were set to 0.3, and both *alphaM* and *alphaPH* started at 0.6. When cues were turned on in the short duration condition, their external input was set to 1 for the first 10 cycles. Cues in the intermediate and long duration conditions had their external input set to 1 for 25 and 40 cycles, respectively. These external inputs were modulated by the *alphaM* and *alphaPH* parameters. In all conditions, the cues were turned off by having their external input set to 0 for five cycles. During these final five cycles, the corresponding US unit had its external activation set to 1. The timing parameter (*TP*) was set to 50 for all conditions. This means that attention would switch completely from a Mackintosh mechanism to a Pearce-Hall mechanism after 50 cycles.

Table F1 shows the trials used to train the model. The model first ran twenty trials of each Phase 1 trial type (these 80 trials were randomly intermixed). After running these trials, the model ran twenty trials of each Phase 2 trial type (these 40 trials were randomly intermixed).

Table G1

Learned predictiveness trials used to train the attention modulated auto-associator

Phase I	Phase II
P1NP1 – O1	P1NP1 – O3
P1NP2 – O1	P2NP2 – O4
P2NP1 – O2	
P2NP2 – O2	