

Bangor University

DOCTOR OF PHILOSOPHY

Motivational and emotional salience effects on working memory

Thomas, Paul

Award date:
2014

Awarding institution:
Bangor University

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Motivational and Emotional Salience effects on Working Memory

Paul M. J. Thomas, BSc, MSc

This thesis is submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy, completed in the School of Psychology, Bangor University.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support, patience, ideas and encouragement of my mentors, colleagues, family and friends. I would like to thank each of the following for their contribution in helping me through the exciting process of the PhD program. In particular: My supervisor, Jane for teaching me how to write and think scientifically and encouraging me never to be complacent with a line of research. My second supervisors, David Linden and Kim Shapiro for their constructive feedback and advice, Miles Cox for chairing my PhD committee, all those who have come and gone from the lab: Madge, Yoshi, Julia, Anthony, Rashmi, Shanti and Claire, for making Bangor a fun place to work, and for lively debates on working memory. My family, for all their encouragement and support. Most especially I must thank my partner Jayne, whose patience and understanding has seen me through the rigours and challenges of the PhD process.

I also thank Bangor University for supporting me with a '125' scholarship. Experiments 1 through 5 were supported by a grant from the Biotechnology and Biological Sciences Research Council (BBSRC): BB/J020206/1.

'I would rather discover a single causal relationship than be king of Persia.' (Democritus, B.C.460-B.C.370)

TABLE OF CONTENTS

Summary	xi
PART 1. INTRODUCTION	1
Chapter 1. An introduction to Visual Working Memory	2
Models of WM	5
The slots vs. resources debate	13
Visual Binding	17
Change detection and AFC tasks	23
The role of expertise	26
Neuropsychology of WM	32
Alternative 'expert' models of WM	37
Early WM processes	42
<i>Chapter References</i>	47
Chapter 2. Emotional Faces	53
Bruce & Young model of face perception	56
The N170 ERP component	58
The cross-race effect	60
Threat and angry faces	61
Faces and WM	63
Neuropsychology of emotional faces in WM	66
<i>Chapter References</i>	72
Chapter 3. Motivational Salience	78
Visual Salience	79
Motivational Salience	82

Motivational salient items and visual WM	88
WM and time estimation	91
<i>Chapter References</i>	94
Chapter 4. General Methods	98
Change detection	100
‘K’ measures of WM capacity	101
Signal detection and the d' statistic	106
Calculation of d'	109
Measure of response bias	114
Assumptions of d'	116
ROC curves and non-parametric SDT measures	117
Why choose d' ?	122
D' for AFC tasks.....	122
<i>Chapter References</i>	127
PART 2. VISUAL WORKING MEMORY FOR EMOTIONAL FACES	130
Chapter 5. Singleton faces and the angry benefit	131
Experiment 1:	
Methods	138
Results and Discussion	143
Experiment 2:	
Methods	148
Results and Discussion	149
Chapter Discussion	158
<i>Chapter References</i>	166
Chapter 6. Early Visual Memory and Consolidation Masking	170

Experiment 3:	
	Methods173
	Results177
	Discussion179
Experiment 4:	
	Methods182
	Results183
	Discussion184
Experiment 5:	
	Methods185
	Results186
	Discussion187
Chapter Discussion188
<i>Chapter References</i>192
PART 3. VALUE LADEN FACES IN WORKING MEMORY194
Chapter 7. Motivational Salience and Value Learning195
Experiment 6:	
	Methods199
	Results212
	Discussion216
Experiment 7:	
	Methods220
	Results222
	Discussion225
Experiment 8:	
	Methods228
	Results235
	Discussion244

Chapter Discussion	248
<i>Chapter References</i>	253
PART 4. MOTIVATION AND WORKING MEMORY	255
Chapter 8. Motivational State and control of WM resources	255
Experiment 9:	
Methods	259
Results	265
Discussion and Chapter Discussion.....	271
<i>Chapter References</i>	276
PART 5. GENERAL DISCUSSION	278
Chapter 9. General Discussion	279
Synopsis of results	280
What types of stimuli cause WM to be boosted?	282
Can WM boosts be explained by attentional mechanisms?	284
What neural processes might underlie WM boosting?	290
Implications for cognitive models	295
How does the brain know when to boost a representation?	300
Limitations and future directions	309
<i>Chapter References</i>	314
Appendices	319
Appendix A - Pre-analysis treatment of eye movement data	320
Appendix B - RGB coordinates for coloured squares task stimuli	328
Appendix C - Face stimuli used in experiments	329
Appendix D - Representing error variance graphically	333

Appendix E - Comparing single probe vs. full array designs	338
Appendix F - Tables of D' , Hit rates, False Alarm rates and Bias 'c'	342

Summary

“An object of primary memory is not thus brought back; it never was lost; its date was never cut off in consciousness from that of the immediately present moment. In fact, it comes to us as belonging to the rearward portion of the present space of time, and not to the genuine past.” (James, 1890, p. 609).

Since William James’s seminal work on ‘primary’ memory, the notion that a short term memory store, now termed short term memory (STM) or working memory (WM) is a key part of human consciousness has become commonplace. This memory store is known to be subject to limitations in the amount and quality of information it can represent at any time, but is a crucial part of the cognitive toolkit humans use to navigate the physical and social world. Given these capacity limitations, there is a clear evolutionary imperative for efficient selection into and retention within WM of items of high utility over those of less importance to survival.

Here, in a series of 9 behavioural and eye tracking experiments I investigate the cognitive mechanisms that serve to prioritise items of importance in WM, and the constraints and concomitant trade-offs that bound such processes. In particular, I focus on cognitive mechanisms by which ‘boosting’ of memory traces of motivationally important visual stimuli may be achieved, what factors cause stimuli to be treated as ‘motivationally salient’ in WM, and whether when one item gets ‘special treatment’ in WM this has impact on representation of other co-present stimuli. Experiments 1 through 5 investigate these phenomena with regard to faces bearing emotional expressions. Experiments 6 through 8 investigate whether learnt reward and punishment associations of faces affect

their treatment in WM. Experiment 9 investigates whether motivational state *per se* can facilitate visual WM processes.

While motivationally salient stimuli have long been known to produce shifts in selective attention (and indeed attention has been termed the 'gatekeeper' of WM), herein I provide evidence that selective attention is but one of the mechanisms by which stimuli of importance are accorded especially robust WM representations.

PART 1 – INTRODUCTION

CHAPTER 1.

An introduction to visual working memory (WM)

The role of working memory

The term working memory (WM) has various connotations, and is often used interchangeably with the term short-term memory (STM) (e.g. Hyun, Woodman, Vogel, Hollingworth & Luck, 2009). Recently, there has been an attempt to define the difference between these two terms. For example, a distinction may be drawn between STM as a capacity limited storage buffer whose capacity may be temporarily improved via strategies such as chunking and rehearsal, and WM as a more complex construct that involves both a storage element and an attentional element, which is required for the maintenance of memory representations and manipulation of information. (Conway, Cowan, Bunting, Theriault & Minkoff, 2002). Despite this attempt at clarification, many research articles make little distinction between the two terms.

Measures of WM have been found to correlate with measures of IQ (Kyllonen & Christal, 1990). Running memory span task (see below) measures of WM have been shown to correlate substantially better with intelligence than the digit span task commonly used in intelligence tests (Cowan et al., 2005), and working memory capacity has also been shown to be a better predictor of general fluid intelligence than either short-term memory capacity or processing speed (Conway et. al, 2002).

Visual WM

Visual working memory is generally believed to operate over a relatively short time period (measured in seconds rather than minutes) and involve limited capacity store(s) capable of holding approximately 3-4 stimuli,

(Cowan, 2001). Attention mechanisms are thought to facilitate the encoding of stimuli into working memory and attention is sometimes therefore termed the gatekeeper of WM (e.g. Bays & Husain, 2008), although it has also been shown that maintaining a representation in WM leads to the selection of matching visual inputs even when there is no strategic reason to do so, presumably through biasing the deployment of attention to such matching stimuli (Downing, 2000). Thus, the relationship between attention and visual WM seems to be a two-way one with the contents of WM biasing deployment of attention and attention gating information entry into WM. Memory representations in WM are prone to decay and/or interference over time, although the extent to which decay and interference separately contribute to the loss of memory information is a topic of some dispute with some researchers (e.g. Lewandowsky & Oberauer, 2009) arguing that interference alone accounts for information loss and no temporal decay occurs, while others (e.g. Portrat, Barrouillet & Camos, 2008) suggest that WM representations decay over time even in the absence of interference from other stimuli.

Visual working memory may be thought of as a temporally intermediate store, distinct from and existing between iconic memory and long term memory (LTM). Briefly, the differences between these three memory stores may be summarised as follows: Iconic memory is a sensory based, unlimited capacity store, with a retention time span of less than a second (Coltheart, 1980). Visual WM is a non-sensory based, limited capacity store, involving an attentional element that facilitates manipulation of information, and typically lasts for up to 10 seconds or so, with retention facilitated by rehearsal and/or

chunking. LTM is an unlimited capacity store within which stimuli do not require rehearsal to keep their memory traces intact.

Visual WM is a key component in change detection processes, which occur all the time as we build visual representations of the environment around us. If we view one apple in a bowl, and our eyes saccade to another apple close by, the storage of an immediately accessible and comparable representation of the first apple in visual WM both serves to allow us to identify the two stimuli as objects that are discrete from one another, and aids us should we need to make a value judgement between the two (Hyun et al., 2009). More generally, there is evidence to suggest that visual WM is involved in the continual correction of the thousands of errant saccades made each day by the human visual system by facilitating object correspondence across saccades and allowing for an automatic gaze correction process based on this correspondence (Hollingworth, Richard & Luck, 2008).

So far I have discussed only visual WM. However, most models of WM posit a separate WM store for auditory information (with some also positing separate stores for information from the other sensory modalities). Competing and historical models of WM are discussed below.

Models of WM

While WM provides a useful and intuitive explanation for a variety of observed behavioural phenomena, the exact nature of WM is disputed; with different models each seeming to be supported by their own sets of empirical evidence.

Atkinson and Shiffrin's model

One of the earliest models of short term memory was proposed by Atkinson and Shiffrin (1968; see Figure 1.1). In this model a short term store with a rehearsal buffer used to prevent memory traces from decaying serves as a separable intermediate store between an a very early auditory store somewhat akin to iconic memory (and termed the 'sensory register' by Atkinson and Shiffrin) and LTM. However, Atkinson and Shiffrin's (1968) model only specifies an auditory-linguistic form of STM and does not speak to the notion of visual STM.¹

¹ Atkinson and Shiffrin (1968) cite lack of existing data as the reason for not proposing a visual form of STM in their model, and acknowledge that such a store could potentially exist.




Figure 1.1. Atkinson and Shiffrin's (1968) model of short term memory.

Reproduced from Atkinson and Shiffrin (1968).

Baddeley's tri-partite and multi-component models:

Perhaps the most well-known model of WM is Baddeley and Hitch's tri-partite division model (Baddeley & Hitch, 1974; Baddeley, 1996). This model posits three main dissociable but interlinked components: the 'Central Executive', the 'Phonological Loop', and the 'Visuo-Spatial Sketch Pad'. The phonological loop is concerned primarily with the short-term storage and

manipulation of auditory information. It has a limited capacity, set in part by the need for stored acoustic/speech items to be rehearsed so that they do not 'spontaneously fade away within 2 or 3 seconds' (Baddeley, 1996). By contrast, the visuo-spatial sketch pad's function is to hold and manipulate (at least) two types of visual information; visual information relating to the form and identity of objects, and spatial information concerning where an object is in relation to other objects. Thus, the visuo-spatial sketch pad serves multiple functions and is sometimes treated as two separate though interconnected memory stores. The central executive in Baddeley's model was initially a less well defined concept; essentially a construct that reflects the multitude of processes/memory stores not covered by the phonological loop and visuo-spatial sketch pad. A key part of Baddeley's model is that the central executive serves to direct attention, and can facilitate processing/encoding in either the phonological loop, or the visuo-spatial sketch pad, essentially boosting their capacity as and when required.

More recently, Baddeley (2000, 2001) has updated the model in various ways to account for the wealth of empirical information that has become available since the model's inception. The phonological loop is now better understood, and thought to contain both a verbal/acoustic temporary store and an articulatory rehearsal system which are involved in, and have evolved to support, the process of language acquisition. These two elements have been seen to be neurologically distinguishable from one another. In a similar way, the visuospatial sketchpad is now assumed to be involved in supporting visual semantics. By far the greatest change to the original tripartite model, however, is the addition of a new element – the episodic buffer.

This component is a limited capacity store, and is 'episodic' inasmuch as it "...holds episodes whereby information is integrated across space and potentially extended across time" (Baddeley, 2000 p421). It plays a role in feeding information into and out of episodic long term memory, and as such stores information in the form of multi-dimensional codes (rather than the auditory or visual codes assumed to constitute memory entries in the phonological loop and visuospatial sketchpad respectively; Baddeley, 2000). Thus, as well as episodic information, this construct may possibly serve functions such as the short-term storage of emotional/social information, and possibly information from other (tactile, gustatory, olfactory, proprioceptive, vestibular etc.) senses. It is important to note that the maintenance of information in the episodic buffer is thought to depend on the limited capacity attentional system, the central executive (Repovš and Baddeley, 2006; see Figure 1.2). Thus, the episodic buffer is not a passive store, but depends on attention to maintain codes. With the addition of the episodic buffer to the model, the role of the Central Executive has also changed – it is no longer assumed to have its own memory storage facility, but retains its role in the directing of attention (Figure 1.2).




Figure 1.2. Repovš and Baddeley's (2006) multi-component model of WM.

The model incorporates proposed links from the various parts of the WM system (top box) to aspects of LTM (bottom box). Reproduced from Repovš and Baddeley (2006).

Cowan's model of Working Memory and the 'focus of attention':

Other models of working memory (e.g. Cowan, 2001) also posit a central, amodal storage mechanism capable of supplementing the auditory and visual limited capacity stores according to task demands. It has been noted by Morey and Cowan (2005), that while it is possible for the central storage mechanism (termed the 'focus of attention' in their model) to support the modality-specific stores, it is not always necessary for this to happen in order to perform a task, and this may account for the mixed results observed by different researchers when looking for evidence of conflict between performance on two simultaneous tasks each recruiting separately visual and auditory working memory. Crucial to Cowan's model is the dual role played by

the focus of attention, with processing and storage functions of this module existing in a trade-off scenario (Bunting & Cowan, 2005). This contrasts with Baddeley's (2001) multi-component model, in which the central executive has no storage capabilities, but a storage and manipulation space is provided by the episodic buffer. In contrasting Cowan's (2005) model with Repovš and Baddeley's (2006) multi-component model (the most recent iteration of the model at time of writing), we can view Cowan's 'focus of attention' as being a construct that incorporates both the central executive and the episodic buffer (i.e. similar to the original version of the central executive in Baddeley and Hitch's (1974) version of the tri-partite model).

Perhaps a more pertinent difference between Cowan's (2005) model and Repovš and Baddeley's (2006) multi-component model is the purported nature of the representations that are stored in the limited capacity passive auditory and visual stores. Whereas Cowan's model defines these as being a combination of sensory memory lasting several seconds (e.g. Cowan, 1988) and temporarily activated LTM representations (Cowan, 2001; see Figure 1.3 below), Repovš and Baddeley's (2006) model calls for a separate storage facility that is discrete from (though linked to) both LTM and sensory memory.

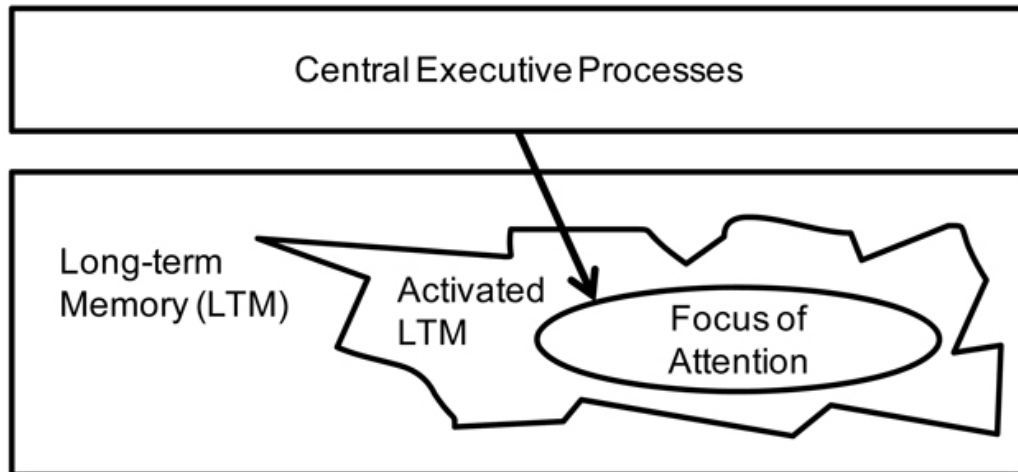


Figure 1.3. Cowan's (1988, 2001, 2005) model of WM wherein the focus of attention temporarily activates codes from LTM. Reproduced from Cowan (2010).

The purpose of this thesis is to investigate the role of visual WM in representing items of value or motivational salience. While it is generally accepted that visual working memory is a separate 'module' to auditory working memory, there are several different accounts of how the cognitive processes associated with visual WM work. I examine these theories below.

Visual WM (VWM) capacity:

The capacity of visual working memory is generally thought to be around three to four items, perhaps because, in Cartesian space, we can determine an object's location using either its projections on the three orthogonal axes or its positioning relative to a small number of other objects each resting on the same surface (most usually the ground plane; Xu & Chun, 2009). For a review of studies supporting the 'magical' number four in visual WM (VWM) capacity see Cowan (2001), but see Olsson and Poom (2005) for

an alternative interpretation. For a review of the usual methods of calculating VWM capacity see below.

When measured by the number of discrete objects/stimuli that can be retained, VWM capacity has been seen to vary as a function of the complexity of (or information load carried by) stimuli, with a VWM capacity of four to five objects only being achieved when stimuli carry the lowest level of information load possible (one piece of discrete information as in the case of simple colours; Alvarez & Cavanagh, 2004). This finding poses the interesting question as to whether the capacity limit of four items applies to objects, individual object features, or some mixture of the two.

Discrete ('slot-based') vs. continuous resource allocation models of visual working memory capacity:

While it is generally accepted that visual working memory is characterised by a storage space that is limited in its capacity, there is some debate as to how this limited capacity is allocated to representations in WM. There are two main conflicting theories that purport to account for how storage space is divided up. Flexible resource models posit that working memory can be allocated in a continuous fashion such that, in order to fulfil task demands, visual stimuli that have greater complexity are allocated a greater portion of the pool of working memory resources. Support for the flexible resource theory comes from Bays and Husain (2008), who found that the precision with which an object was remembered, measured by the distribution of errors across various changes in location or orientation of black symbols, did not decline

suddenly at a 'magic number' as set size increased (as would be predicted by a slot based theory), but declined according to a power law.

By contrast, slot-based models posit a discrete number of memory 'slots' that comprise the overall resource quotient for an individual's working memory capacity. According to discrete models, each stimulus presented to an individual will be accorded a single slot, regardless of the stimulus's complexity. Discrete allocation or slot-based models gain support from research by Barton et al. (Barton, Ester & Awh, 2009), who found that varying the complexity of certain stimuli in a 2x2 array had no effect on participants' ability to accurately detect changes between presentation and probe in complexity-constant stimuli from the same array. A continuous allocation or flexible resource model would have suggested a decrease in WM resources available to complexity-constant stimuli in conditions where other stimuli in the array had high complexity, with accompanying deterioration in change detection performance. Barton et al. (2009) interpreted this finding as showing that the complexity-constant stimuli each received a discrete portion of WM resource (or a slot).

It should be noted that the unit posited to be subject to slot allocation is most usually an object, rather than an object feature. This position is derived in part from experiments by Luck and Vogel (1997), who found that while WM capacity for their object stimuli was limited to about four objects, increasing the number of features within each object (e.g. having each square in the presentation being bi-coloured rather than mono-coloured) did not result in any significant decline in feature change detection performance, even when feature load was increased up to four features per object. As well as providing an indication that a slot-based account of VWM may be plausible, this finding

brings to focus a key issue that must be accounted for by any model of VWM; the facilitation of WM provided by the binding of features into an object. This issue has a direct parallel in the study of auditory WM, where chunking of individual items can facilitate their tenure in auditory WM (e.g. Chen & Cowan 2009). Interestingly, however, it might be possible for features that are bound into an item to exist as 'chunked information' without recourse to LTM involvement, which is often assumed to be key to the creation of chunks (e.g. Cowan, 2001). Thus there may be two possibilities in terms of grouping data for its optimal storage in WM; a semantic form of chunking reliant on LMT involvement, and a low level non-semantic binding of features together reliant on sensory processes and without LTM involvement. I discuss the issue of visual binding further below.

A compromise between the flexible resource and slot-based theories has been proposed by Zhang and Luck (2008). Their model posits a slot-averaging mechanism, where, provided the number of stimuli present for encoding is less than the number of slots in an individual's WM capacity, multiple slots may be assigned to certain stimuli. Crucially, this model specifies that the slot is the only unit of resource assignable in WM, and that no flexible pool of resources exists in addition to the slots. This bold statement is based on Zhang and Luck's findings that their measure of precision or accuracy of the representation held in WM (measured by the accuracy of the output given on a probe colour wheel when evaluating previously seen sets of colour squares, and termed 's.d. '), varied when presentation set size varied between one and three, but did not vary when presentation set size increased above three items. This result can be interpreted as showing that once slots have been used up in

memorising individual objects (i.e. once the value of K [see Chapter 4] for a participant had been exceeded), precision of the representation of individual items cannot be further improved, but when set size is sufficiently low such that there are slots 'left over' these slots can be assigned to objects already accorded a slot, so that multiple representations of the same object are stored. An output system of averaging can then be used when recall is required, resulting in an output that is an average of two or more imperfect representations, which is more accurate than an output based on a single imperfect representation stored by a single slot.

An additional experiment by Zhang and Luck (2008) set a task in which participants' optimal strategy would have been to devote the vast majority of their working memory resources to a single cued stimulus, while leaving 'only a few drops' of resource available for other stimuli present. Their results highlighted the inability of participants to devote their resources in this manner, but instead showed a pattern consistent only with a model in which the minimum resource allocatable to any given stimulus is a single slot, such that the precision or accuracy of the memory trace for a given stimulus never falls below the value predicted by one slot of memory resource (except in instances where no slots are assigned at all, such as when insufficient time is allowed for encoding).

An alternative possibility is that different stages of object representation are characterised by different capacity limits, with capacity limits at the stage of object individuation (selection of an object from the visual scene for encoding into WM) being slot-based and subserved by the inferior intra-parietal sulcus (IPS), and capacity limits at the stage of object identification (the building up of

a detailed representation of an object) being based on a continuous pool of resources, and being subserved by the superior IPS (Xu, 2008). Where such a process of individuation followed by identification takes place, object features and binding information could be represented as part of the identification process. I discuss this possibility in greater detail below.

The visual binding problem:

As discussed above, individual features may be bound into objects such that visual WM capacity is limited in terms of the number of objects, rather than the total number of features present in a display (e.g. Luck & Vogel, 1997). However, Wheeler and Treisman (2002) have suggested that a binding account of the processing of stimuli is not necessary to account for Luck and Vogel's (1997) data, and that their pattern of results could instead be explained by a parallel stores account of visual WM processes, such that different types of features (shape, colour, location etc.) get stored in different parallel (though possibly overlapping) limited capacity visual working memory stores. Luck and Vogel (1997) had found that when two squares of different colours are presented for memorisation, with one square fully surrounding the other, they can be bound together as an object, and encoded in visual WM such that they are treated as a single entity (using up only one slot of capacity, rather than the two slots that would be predicted if each feature – or colour part – were stored separately). Wheeler and Treisman (2002) were unable to replicate this finding with these colour-colour conjunction stimuli (for discussion of why this may be see below), and concluded that features in the same dimension (e.g. colour with another colour feature) are not necessarily bound together as a single

object for encoding in visual WM. In further experiments using a memorisation array comprising three coloured squares, Wheeler and Treisman (2002) manipulated the probe display such that four conditions were presented: Colour alone could vary in the probe, location alone could vary, both location and colour could vary, or the binding information could vary (i.e. two of the coloured squares could swap places such that the location information and colour information presented in the memorisation array and probe array were identical but the way these dimensions were bound together in memorisation array and probe array varied). In a replication of Luck and Vogel's (1997) finding, there was no significant difference between performance on *colour* trials (where colour alone could change at probe) and *colour-and-location* trials (trials where either colour or location could potentially change at probe) when the dimension that varied between probe and display was colour. Similarly, there was no significant difference between performance on *location* trials (where location alone could change at probe) and *colour-and-location* trials (trials where either colour or location could potentially change at probe) when the dimension varying between probe and display was location. Since the *colour-and-location* trials required the encoding of twice as much information as the shape only and location only trials, this finding suggests either that features are bound and encoded as a single object, as proposed by Luck and Vogel (1997), or that each type of dimension (in this case location and colour) has its own separate limited capacity visual WM store. If features are bound into objects and stored as such in VWM, then the final condition in which only the binding information was changed would have been expected to engender similar performances levels to the other conditions, but this condition was seen

to have significantly lower performance levels than the other conditions, indicating that the binding information – i.e. the relationship between the two features of the squares, colour and location – was not available for recall at probe. The results from this experiment therefore favour a model in which features are stored in parallel dimension-specific limited capacity VWM stores.

Wheeler and Treisman (2002) also replicated this effect with *colour-and-shape* arrays (i.e. trials on which the binding information between shape and colour could vary between memorisation array and probe). Additionally, they repeated the experiments, for both *colour-and-location* and *colour-and-shape* arrays, but substituted the three object 'same/different' probes with single object 'present/absent' probes. When a single object was used at probe, the disadvantage seen in the binding information manipulation condition disappeared, suggesting that binding information is available to participants under such conditions. It therefore seems that binding information may be recorded along with features, but that this representation is fragile, and prone to interference at time of retrieval. The mechanisms underlying this interference are not yet understood, although Wheeler and Treisman (2002) ruled out increased decision-making load as the explanation.

An interesting parallel to the type of retroactive/output interference suggested to occur when a multiple component probe array is used (Wheeler et al., 2002) is found in the Running Memory Span task. In this task, initially devised as a way of measuring how the accuracy of monitoring a digit stream may change under varying conditions (Pollack, Johnson & Knaff, 1959), participants are presented with a stream of digits, and required to memorise both the digits and their order in the stream. Recent work with this paradigm

(Bunting, Cowan & Saults, 2006), has found that recall accuracy at individual serial positions depends on the number of other serial positions to be reported first. However, the output interference arising from report of other serial positions disappeared when participants were permitted to determine for themselves the number of digits they reported (i.e. up to the maximum they were comfortable recalling on each trial) rather than being forced to recall a fixed window of five, six or seven digits. Thus, it seems that the output interference in the running memory span task is primarily a result of having to guess at earlier items in the digit stream. A similar mechanism might also explain the output interference observed by Wheeler and Treisman (2002): If we consider a trial on which the participant had successfully encoded two out of the three stimuli presented, but failed to encode the third, presentation of a single item at probe (a present/absent probe trial) results in a one in three chance that the probe item presented is the non-encoded item. If, however, on this trial the full display probe (change/no change probe) had been used, and assuming for this example a no change trial, two thirds of the information would be known to match to the initial presentation, and the final third (the non-encoded item) would be unknown. Under such conditions, we might initially assume that since the preponderance of evidence available to participants points towards a match with the initial presentation array, this would be the most likely response (see also Appendix E for a discussion of how summary statistics reflecting WM capacity measures should be applied to such data). However, in this condition, there is a necessity to guess/imagine whether the non-encoded stimulus matches to the initial display, and the cognitive mechanism underlying this guess may contribute to the output interference

observed (i.e. the loss of binding information). Such a guessing mechanism can be conceptualised as a 'call on the imagination'; given that the correct answer is not known, one must be invented to fulfil task demands, and this invention process may place cognitive demands that are responsible for the output interference observed in the above tasks.

It should be noted that the auditory suppression task used by Wheeler and Treisman (repeating 'Coca-Cola' out loud) may have engaged significantly less auditory memory resources than the one used by Luck and Vogel (remembering by auditory sub-vocal processes two numbers for a later test) which would have forced an active auditory rehearsal process. This could explain the difference between the two sets of results if it resulted in a difference in the encoding strategy used by participants in one of the conditions common to both experiments. In both Luck and Vogel's (1997) original experiment and in Wheeler and Treisman's (2002) attempt at replication, one condition used memorisation arrays comprising squares that had multiple featural (coloured) elements, while another condition used memorisation arrays in which each square was of a different (single) colour. If we allow that auditory encoding and/or rehearsal may be facilitated by subvocalising the to-be-remembered stimulus name or label (a supposition common to models of auditory WM) then it is apparent that the recruitment of auditory working memory might only be used as a strategy when there is sufficient time over the presentation period and inter-stimulus interval (ISI) to subvocally articulate the names/labels at least once. Given that the period for the combined memorisation array display (147ms) and ISI (906ms) in Wheeler and Treisman's experiment lasted just over a second, it could be that participants

were able to recruit auditory working memory resources to aid their encoding of the simple and quickly articulated *all-squares-of-different(single)-colours* (set size = 3) condition (condition 'nine' in their Experiment 1), but were not able to recruit auditory memory resources in any of the conditions where displays were more complex (including a condition where all squares were of different single colours but set size was six). By contrast, the more demanding auditory suppression task in the original Luck and Vogel (1997) experiment would have prevented any recruitment of auditory WM even in the simplest conditions, thus accounting for the absence of any advantage for the equivalent condition in their experiment. Thus, Luck and Vogel's original interpretation of their (1997) results might be supported: Binding information may serve to 'chunk' features together such that they can be encoded more efficiently as a single object. There is, therefore, still some uncertainty as to how binding mechanisms affect the way objects are encoded in VWM. If Luck and Vogel's (1997) model holds true, then this lends support for a general slot-based account of the limited capacity VWM store. If Wheeler and Treisman's (2002) model is correct, then different parallel stores exist, and the slot-based vs. continuous question must be examined in turn for each of them.

The comparison of Luck and Vogel's (1997) experiment with Wheeler and Treisman's (2002) experiment and the apparent discrepancy in the results therein is illustrative of one of the challenges in investigating visual WM; two seemingly similar experiments designed to measure the same process may in fact measure subtly different aspects of the WM system. Below I review how some of the paradigms designed to investigate WM attempt to tackle this challenge.

Investigating visual WM capacity

Change detection and alternative forced choice (AFC) tasks

One of the most frequently used paradigms for investigating visual WM is the change detection task. In this paradigm participants are shown an array of to-be-remembered stimuli (the 'memory array') for a very brief time, typically less than 2 seconds, providing sufficient time to encode the stimuli into WM but insufficient time to encode them into LTM. Next follows a retention interval, a brief period of time typically accompanied by a blank screen, in which participants must retain the information in WM. Next follows the test of the contents of WM; a new array of stimuli (the 'test array') is presented which is identical to the original memory array on some trials but differs from the memory array on other trials. Participants indicate (typically by key press) whether the test array is the same as or different from the memory array by comparing the contents of their WM to the items in the test array. Where participants have been able to encode and retain a large proportion of the information contained in the memory array, they will have better information on which to base the same/different judgement. Thus, task performance should be indicative of the ability to encode, retain and retrieve information into and from WM under whatever the specific task conditions are. This type of task is termed *change detection* since participants are in effect answering the question 'has anything changed between the memory array you saw first and the test array'.

As noted in the discussion of the differences between Luck and Vogel's (1997) and Wheeler and Treisman's (2002) experiments above, a frequently used variant of this task uses only a single item at probe, and the judgement to be made is then whether this single item was present or absent from the memorisation array. While in principle the two variants of the task are both designed to measure WM capacity, there are subtle differences not only in the cognitive processes they may engage (see discussion above), but also in the statistical summary measures that are appropriate for description of data from each variant (see Chapter 4 and Appendix E for discussion of this point).

While change detection tasks are frequently used to investigate WM, another widely used method is the alternative forced choice (AFC) task. In this task a memory array is presented and then after a retention interval a set of response options is presented. One of these response options comprises the target stimulus (the stimulus having been present in the memory array) and the other(s) are lures (not having been present in the memory array). Participants are required to pick the stimulus response option they believe to be the target stimulus. While the most frequently used AFC task is the 2-AFC task in which a single lure is presented alongside the target at test, in principle any number of lures may be presented. Thus, in any N-AFC task the number of lures is N-1. Using this notation, yes/no tasks are sometimes called 1-AFC tasks (since although two response options are possible no lures are presented at test). In principle, it is possible to use a variant of the AFC task, in which a series of response options comprising one 'target' (item that may have changed) and N-1 lures must each be compared to the entirety of the memorisation array (see Figure 1.4, panel 3). However, this design falls foul of certain practical

limitations, notably that the number of comparisons of individual entities that must be made at test grows very large as N increases beyond 1, resulting in a difficult judgement process that occurs over an extended time such that response options inspected later in the comparison-to-memory process may be compared to a more degraded version of the memory array (due to decay of the memory array representation over time). Thus, this alternative AFC design is rarely used. Differences in the way in which data from different types of AFC task should be analysed are described in Chapter 4.

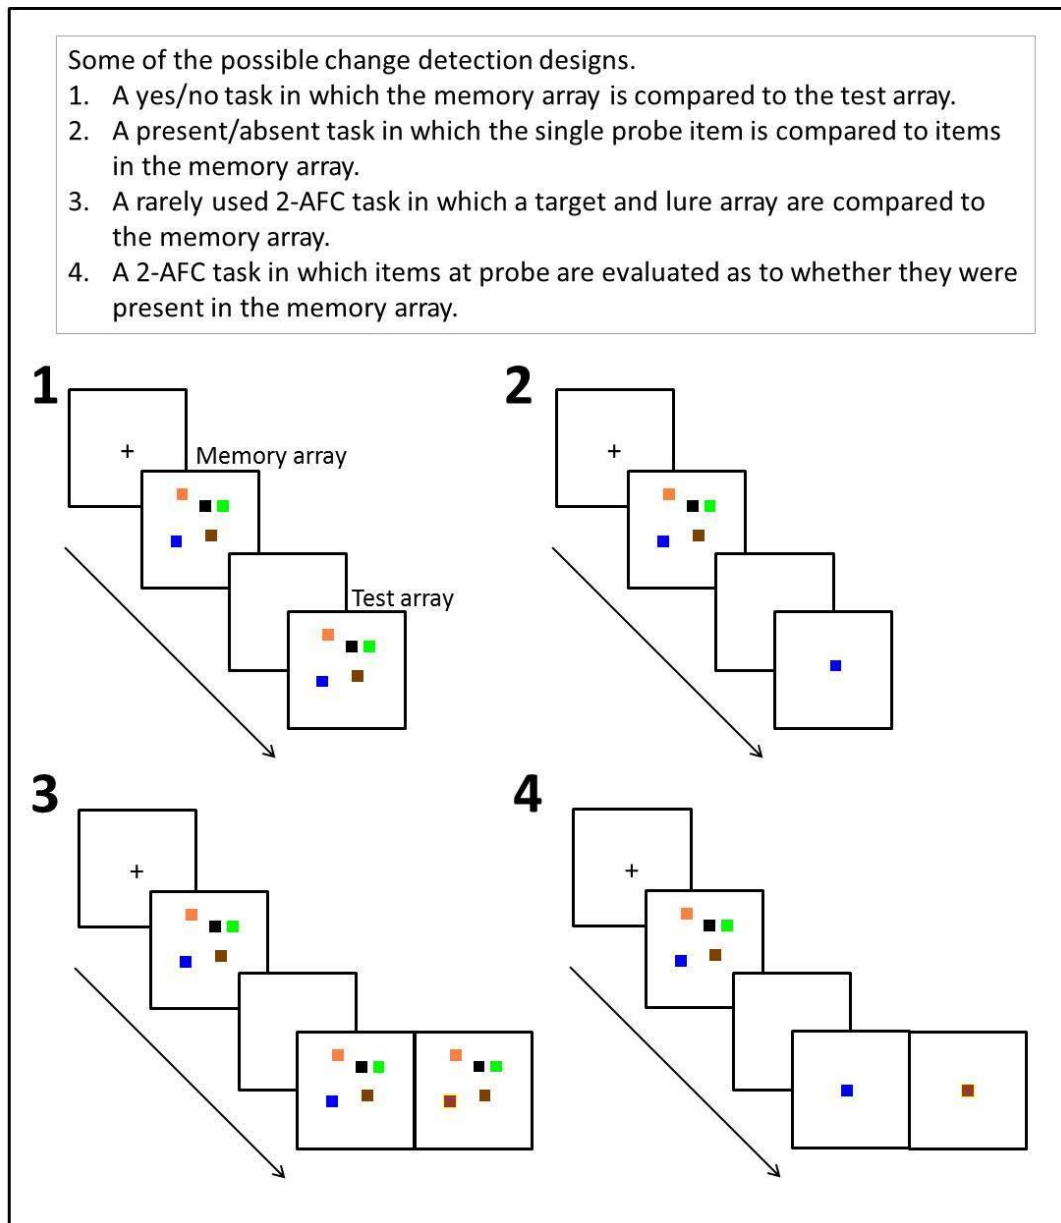


Figure 1.4. Alternative yes/no and AFC change detection designs using coloured square arrays. Other designs are also used, for example in some paradigms single probe items occupy the same spatial location as they did in the test array allowing comparison of only this probed item to its representation (when such a representation it is held in WM on a particular trial).

LTM contributions and the role of expertise

In paradigms designed to measure the capacity of WM, it is often important to block possible contributions to task performance from LTM. LTM has been implicated in the ability to group individual stimuli into chunks (Cowan, 2001), with this mechanism believed to account for some, though not all, of the WM advantage shown by experts when asked to memorise objects of expertise (Curby, Glazek & Gauthier, 2009). The game of chess provides an excellent resource for investigating the phenomenon of expertise since a quantitative rating system exists that ranks all competitive chess players in terms of their ability (the ELO system; Elo, 1978), allowing a quantitative definition of an expert. It also offers strong external validity (Gobet, 1998), as well as a mathematically formulisable environment in which task difficulty can be systematically manipulated, and a rich database of games played both by human players and by artificial programs which can be used to statistically assess elements of the game (for example, commonality of certain patterns of pieces within the 'position space' of the game).

Studies of expert chess players have shown that they are better able to remember briefly presented chess positions than non-experts, but this advantage is lost when pieces are positioned at random locations on the board rather than in positions that more commonly occur in real games (de Groot, 1965; de Groot & Gobet, 1996). For example, the positioning of chess pieces shown in Figure 1.5 below can be coded as a single entity by a player of reasonable experience but the square upon which each piece stands and the positioning of pieces relative to one another may make a call upon far greater spatial WM resources when shown to a beginner. Thus, it can sometimes be difficult to predict whether it is possible to code particular types of visual

information simply in terms of objects and object features or whether spatial WM may be called upon in any particular instance. WM coding of the same information may call upon spatial WM in one individual or situation but non-spatial (i.e. object-based) visual WM in another depending on whether the individual is able to group or 'chunk' elements together into complex single entities. It is interesting to note that the coding of such positions by experts seems not to rely on an invocation of auditory coding: While the position of the white pieces in Figure 1.5 could be verbally coded quite efficiently by the auditory code 'double fianchetto, king-side castling and central knight development', suppression of the articulatory loop has been seen to have little effect on coding of such positions, even among experienced players (Robbins, Anderson, Barker, Bradley, Henson, Hudson, & Baddeley, 1996²). Rather, this ability seems to be due to automatic chunking of a large amount of visual information into a single visual code or compound object (see also de Groot & Gobet, 1996 for evidence from eye movements supporting the efficient chunking of positional information among chess experts).

² It is also worth mentioning that this study provides an elegant illustration of the necessity of visual information in planning behaviour: In a second experiment Robbins et al. required participants to plan chess moves from a position they were shown while visual WM was selectively loaded (without loading auditory WM or the central executive). This impacted ability to plan moves as did loading of the central executive, while auditory suppression did not.




Figure 1.5. The position of the pieces may be coded almost instantaneously by a chess expert or experienced player. This particular positioning of pieces occurs relatively frequently in competitive games, allowing it to be entered into WM as a single entity rather than a selection of spatially related individual pieces.

One option to prohibit LTM contributions is to use a paradigm that requires participants to recall a serial order of stimuli, and draw from a limited set such that stimuli are repeated across many trials. Viewing of the same stimulus many times in different trials makes it more difficult for participants to rely on LTM associations for any given stimulus and so lessens the possibility of chunks being formed (Cowan, 2001). Additionally, harder-to-chunk stimulus arrays may be used (e.g., unusual chess positions rather than common positions in the case of chess experts). More generally, all humans can be considered experts with regard to the patterns of how certain stimuli are likely to be positioned relative to one another, and as such may be able to call upon LTM to facilitate coding of stimulus arrays into WM if the visual layout of such

stimuli follows a pattern that is held within LTM. For example, consider the stimulus sets shown in Figure 1.6 below. The set of clothes on the left is arranged in a manner that accords with how they would be worn by a person, whereas the same items are positioned in pseudo-random locations in the array to the right. It is easier to encode the stimulus array on the left, in part because a LTM template of a human body may be (semantically or visually) overlaid on the array during encoding, allowing certain encoding errors to be eliminated as being semantically incongruent with such a template (memorisation mistakes such as 'there were two hats present in the display' can be avoided in the left array since wearing two hats at once is incongruent with the template of a human).




Figure 1.6. Memorisation arrays of items of clothing. The array to the left presents items in a semantically ‘sensible’ pattern in accordance with existing LTM templates. The array to the right is pseudo-randomised. In both arrays items retain their upright orientations (that is, they appear in the most commonly encountered real world orientation) but in the right side array binding of the items onto a LTM template is no longer possible. The role of templates is discussed in more detail in the section on ‘alternative models of WM’ below.

Blocking auditory WM contributions

To fully utilise overall WM capacity, and fulfil task demands in a visual WM task, participants may use a strategy that assigns verbal labels to visual stimuli, and so recruits auditory WM processes (the phonological loop). When we are interested in the capacity of visual working memory alone, it is necessary to block such a contribution from auditory WM. The most common method used is *auditory suppression*. Auditory suppression involves loading

the auditory WM store with information to be rehearsed such that it is fully engaged, and there is no spare capacity available to contribute to visual task performance. Common auditory tasks include having participants repeat a word or phrase (usually one with no visual element associated with it) out loud or sub-vocally, or having participants remember two numbers for a later test by repeating them in either of these ways. While a verbal load of two syllables is common in auditory suppression tasks, it is important to note that it may not always be sufficient to fully prevent contamination of visual WM by recruitment of auditory processes (Curby et al., 2009; Marsh & Hicks, 1998). Additionally, it has been observed that a verbal memory load spoken aloud has a greater interference effect on concurrent visual working memory tasks than silently maintained rehearsal of the same auditory information, perhaps because speaking aloud forces stimuli to be continually brought back into the focus of attention (Morey & Cowan, 2005).

Neuropsychology of Visual Working Memory

Determining the neural structures underlying WM is an on-going process, and there is currently no single unifying account of the neurophysiology of visual WM.³ In part, this is due to the well-known inherent difficulties in localising any neurological function (the correlational nature of fMRI data, the macroscopic anatomical differences between primate and human brains, the scarcity of human lesion patients with focal lesions within only the particular region of interest), these difficulties being further

³ There is, however, a general agreement that whereas LTM is *weight-based*, reflecting differences in the synaptic connections between neurons, WM is *activation-based*, reflecting patterns of firing within populations of neurons.

exacerbated by the existence of multiple definitions of WM, so that different neurological studies may address different cognitive functions (e.g., short term visual storage vs. short term visual manipulation of information) each attempting to localise visual WM but each actually localising a different cognitive function. A further consideration is that visual WM may not be subserved by a single neural region but may reflect a complex circuit of regions, perhaps with local specialisation in terms of the parts of the WM process each is involved in (encoding, maintenance, manipulation, selection and retrieval). Thus, rather than attempting to localise WM *per se*, it may be more appropriate to attempt to localise each of its sub-processes. For example, a recent study of patients with dorsolateral prefrontal cortex (dlPFC) lesions (Barbey, Koenigs & Grafman, 2012) reported that while unilateral dlPFC was observed not to be necessary for WM maintenance, left dlPFC was seen to be necessary for manipulating verbal and spatial knowledge, while right dlPFC was important for verbal and spatial reasoning. Thus, it would seem that an important sub-component of WM is served by these frontal regions.

A further subtle difficulty in neurological investigation of WM processes is that the regions typically implicated in WM show a large overlap with the circuit of regions thought to be involved in attention (Fougnie, 2008; Pessoa & Ungerleider, 2004). This makes distinguishing between these two processes difficult, especially since it is hard to design tasks that selectively load WM in the absence of requiring a similar loading of at least some attentional processes.

Despite these difficulties, a wealth of neurological research exists that attempts to delineate the neurology of WM processes. While this thesis is not

concerned primarily with the neurological basis of WM, I describe below some interesting recent findings to which I make reference later in discussion of my own experiments.

Attempts have been made to functionally localise the various components of the tri-partite model. For example, Brodman areas 40 (a left hemisphere parieto-temporal area) and 44 (Broca's area) have been implicated as serving the phonological loop (Paulesu, Frith & Frackowiak, 1993), while the visuo-spatial sketchpad seems to be neurally subdivisible into occipital and inferior frontal regions associated with spatial WM, and inferotemporal areas supporting object-based WM, with both object and spatial WM areas quantitatively lateralised towards the right hemisphere (Smith & Jonides, 1997). Central Executive processes are associated with frontal regions, although precise localisation of specific executive processes is difficult (not least because executive processes are also used to apply task rules and guide attention). While there has not yet been any success in neurally localising the episodic buffer, it is quite possible that rather than being functionally localised this component depends on a distributed process of synchronous firing (e.g. see Baddeley, 2000).

An alternative, although less frequently encountered view is that the same neural circuitry that controls LTM can fully account for STM processes (note that here I am referring specifically to STM rather than WM, that is to say short term storage and retrieval rather than manipulation of short term information). This possibility was acknowledged by Atkinson and Shiffrin (1971), who noted that 'Our account of short-term and long-term storage does not require that the two stores necessarily be in different parts of the brain or

involve different physiological structures. One might consider the short-term store simply as being a temporary activation of some portion of the long-term store.' While there is generally a consensus that neural mechanisms underlying STM occupy different regions from those that support LTM (though possibly with some overlap), the view that STM and LTM may not have neurally distinct memory stores has not yet been entirely ruled out (e.g. see Ranganath & Blumenfeld, 2005, who argue for the shared circuitry account). To some extent, this view of non-separable underlying neural resources parallels Cowan's (2001) model of WM, as does the long term WM (LT-WM) theory (discussed below), where temporarily activated LTM codes serve as the contents of WM.

Recent neuropsychological evidence has raised an interesting question concerning the role of the posterior parietal cortex (PPC) in visual WM processes. Some models treat the PPC as a limited-capacity store, capable of encoding the traditional four or so items (e.g. Todd & Marois, 2004; Xu & Chun, 2006), while more recently it has been suggested that the role of the PPC is in rehearsal processes, with the memorised stimuli being 'stored' (in their traditional 'retrievable-to-consciousness-without-aid-of-cuing' form) elsewhere in the brain (Magen, Emmanouil, McMains, Kastner & Treisman, 2009). Magen et al.'s evidence for this later view is based on observations that PPC activity continues to increase as a function of set size during maintenance time, at a non-monotonical rate relative to behavioural measures (K). This result is not consistent with the traditional view of visual STM as a relatively robust but capacity-limited store, but could possibly be interpreted in relation to the model proposed by Sligte et al. (Sligte, Scholte & Lamme, 2008) of an additional

'fragile' store with a larger capacity (see below), with the PPC being involved in fragile visual STM processes.

It is also worth noting that different representations of the same visual stimulus might be stored with differing levels of detail in different brain areas, with each area constituting a different limited capacity store. Recent research (Xu et al., 2009) into the role of the intra-parietal sulcus (IPS) has implicated the inferior IPS in object individuation (the selection by spatial location of up to four objects in visual space, with object representations at this stage being coarse in nature and containing only minimal feature information such that an initial 'object file' can be set up), and the superior IPS in object identification (the addition of further information, including feature binding information to an object file such that elaboration of the object takes place). Interestingly, the object identification mechanisms subserved by the superior IPS seem to have a separate capacity limit, determined in part by object complexity (Alvarez et al., 2004), with superior IPS activity increasing in presence of greater levels of feature information (Xu, 2007).

If WM depends on an interconnected circuit of cortical regions then our understanding of this circuitry must depend on an understanding of the contribution made by each component as well as the role served by white matter fibers connecting the various cortical regions. It has recently been suggested that a parieto-frontal network may underlie WM function as well as general intelligence (Barbey, Colom, Solomon, Krueger, Forbes & Grafman, 2012; Jung & Haier, 2007). Indeed, Barbey et al. (Barbey, Koenigs & Grafman, 2013) suggest that such a network is sufficient to account for WM. However, neural regions not included in this network have also been implicated in WM

processes. For example, there is evidence that the habenula, while not part of the key circuit described above, may play a modulatory role through its involvement in the control of the release of neuromodulators such as dopamine, serotonin, noradrenaline and acetylcholine in other cortical regions (Hikosaka, 2010), while synchronisation of theta oscillations between area V4 and the PFC has been observed to predict performance in monkeys performing visual WM tasks, both reinforcing the concept that WM depends on an integrated network of cortical regions and implicating area V4 in visual WM (Liebe, Hoerzer, Logothetis & Rainer, 2012). Establishing the neural basis of the different constituent components of WM thus remains an ongoing work, though progress is being achieved at a significant pace.

Alternative models of WM

Cowan's (2001) and Baddeley's (e.g. Repovš & Baddeley, 2006) models are probably the most well-known current models of WM. However, an alternative class of models has also emerged from the field of expertise and skilled learning. This field is concerned with understanding how experts achieve enhanced performance on tasks related to their areas of expertise. Much of the research in this field has been undertaken with chess experts for the reasons discussed above, and one fundamental question that it seeks to address is how chess experts can achieve such high WM task performance when asked to memorise chess positions. I discussed above how this enhanced ability is thought to be a function of very efficient visual chunking mechanisms, rather than recruitment of auditory WM resources. However,

three alternatives to this chunking theory have been advanced in this field, each with implications for the nature of WM. These are summarised below.

SEEK (SEarch, Evaluation, Knowledge) theory (Holding, 1985, 1992)

SEEK theory allows that experts are able to convert visual information (e.g. a position on a chess board) into a system of interlocking semantic themes. Thus, where a beginner may code a group of pieces in the centre of the board as 'my knight on square E5, my pawn on square D4 and my opponents' pawns on squares D5 and F5', an expert may code this same configuration as 'an octopus⁴ not dislodgeable without exchange'. This semantic information can then be used as an aid in reconstructing the visual information originally present. SEEK theory thus specifically suggests that experts do not use visual chunking to facilitate their performance. The suggestion that experts are able to code elements within their field of expertise in semantic terms is interesting, but SEEK theory has not yet been evaluated other than in the narrow field of chess expertise.

Template theory (Gobet, 1998; Gobet and Simon, 1996).

Template theory also derives from research on chess experts. In this theory it is assumed that experts have access to some 10,000 to 100,000 chess positions in LTM. However, especially well studied positions may exist not merely as codes for a position itself but as 'templates' with a number of fillable slots corresponding to potential developments of the position (i.e. the

⁴ The importance of positioning a knight on one of the centre squares of the chess board and maintaining it there so that it threatens eight other squares is so important in chess games that such a knight has a special name (an octopus) in chess literature.

next moves). More generally, template theory allows for special templates of commonly encountered stimuli to exist in LTM which may be particularly suited to manipulation by WM resources. One can consider such a template to be constructed of a series of linked specifications, some of which are fixed in nature, and others of which are malleable. For example, a template of a face may specify the general position of the eyes, ears, mouth and nose in relation to one another as well as the general elliptical/ovoid outline of a stereotypical face while allowing for fillable slots to remain open so that the shape, texture and colour of each component part of the overall stimulus can be filled in when a particular face stimulus is presented for memorisation.

Long term working memory (LT-WM; Ericsson and Kintsch, 1995)

Rather than creating visual chunks (chunking theory), semantic chunks (SEEK theory) or accessing malleable stored templates (template theory), LT-WM theory proposes that experts are able to encode with great rapidity and precision a large amount of information concerning objects of expertise directly into LTM, as well as creating a hierarchical retrieval structure (essentially nested pointers to each part of the information in LTM) in WM with which to access this information. However, a criticism of this theory is that the ability to create and use such retrieval structures may only be possible in situations where there has been a conscious intent to improve one's memory to the level of 'expert' through learning of strategies (Gobet, 2000).

For the sake of clarity we should note again that none of the above alternative models of WM purport to offer an alternative explanation of general WM, rather they seek to account only for the unusually enhanced WM

performance of experts, and each is in agreement that experts learn through constant practice to develop an enhanced WM specific to their area of expertise. Thus, these models do not propose that general WM capacity is better in experts, and indeed experts who are able to achieve enhanced WM performance in their own field of expertise have not been observed to have any general WM enhancement; an expert chess player has enhanced WM for chess positions but normal WM on other WM tasks. Indeed the current prevailing view is that an individual's general WM capacity is fixed (though for an alternative view, that training can improve general WM through neural plasticity see Klingberg, 2000).

I will not seek to advance any argument here as to the veridicality, strengths, and weaknesses of each model above (for a review broadly in support of template theory see Gobet, 1998). However, in each experiment reported herein I compare participants' ability to perform WM tasks under varying conditions and in each instance ask the question 'does the manipulation in question allow participants to perform better than they normally would on a particular WM task?' It is therefore instructive to bear in mind these proposed models that purport to account for some of the mechanisms by which WM performance may be enhanced over 'typical' levels.

Early memory processes

So far we have more or less treated WM as a store existing between iconic memory and LTM. However, the actual processes at play may be less easily parsed into neat components. Two converging strands of research in particular suggest that the processes occurring between iconic memory and

WM are of vital importance in understanding what information gets selected into WM. These are the idea of 'fragile' WM (Sligte, Scholte & Lamme, 2008, 2009; Sligte, Vandenbroucke, Scholte & Lamme, 2010), and the proto-object/interactive approach to consolidation (Gao, Gao, Li, Sun & Shen, 2011; Xu & Chun, 2007).

Fragile working memory:

Recent research (Sligte et al., 2008, 2009, 2010, Vandenbroucke, Sligte & Lamme, 2011), has posited an additional 'fragile VSTM' store, that has the following characteristics: 1) persists for at least four seconds after stimulus offset; 2) is approximately double the capacity of traditional or 'robust' WM; 3) is 'boostable' by direction of attention to the internally held location of a stimulus encoded in this fragile store; 4) is not iconic in nature. Fragile VSTM has been disambiguated from iconic memory by studies showing that after presenting a stimulus in a memorisation array, cuing of the location of the stimulus after the decay interval of iconic memory (about 500 ms; Sperling, 1960) has passed facilitates recall of the stimulus previously presented at this location. Up to sixteen or so simple line-based stimuli can be shown on average to be memorised into fragile VSTM in such a paradigm (well in excess of the typical four or so available to WM). By contrast, when cues are not provided as to which stimulus is to be tested in this paradigm the capacity measure at test leads to a more traditional K value in the region of four items, reflecting traditional visual WM (Sligte et al., 2008). This finding, along with evidence that constraining attentional resources in similar paradigms impacts visual WM to a far greater extent than it does fragile VSTM (Vandenbroucke, et al., 2011)

dissociates fragile VSTM from visual WM. Thus, fragile WM should not be considered an alternative model of WM, rather it is an extension of existing models and is agnostic in its choice of 'main WM' model, that is to say it may be theoretically integrated with any of the existing models described above, representing as it does a stage between iconic and working memory. Recently, the existence of a separate fragile memory store has been challenged (Matsukura & Hollingworth, 2011).

Early WM, proto-objects, and consolidation

Consolidation

In a series of experiments, Potter (1976) identified that formation of durable short term memory representations of visual scenes occurred at a slower rate than simple perceptual analysis. Using the rapid serial visual presentation (RSVP) paradigm, sixteen photographs of visual scenes were presented sequentially on each trial, with a presentation rate of one, three, four, six, or eight stimuli per second. Participants were then given a single photograph and asked if it had appeared in the sequence. Recognition performance was high at slower speeds but fell at higher speeds, indicating a failure to form durable memory representations at higher speeds. However, when participants were given a picture to identify *prior to* the trial (a preview condition), identification performance rates were high even at the fastest presentation rates. This rules out the possibility that participants' performance in the recognition task was impacted by an inability to perceptually process the photographs at faster times, and indicates instead that participants were unable

to use this perceptual representation to form a durable memory trace. This is not to say that participants were unable to form *any* memory trace; merely that any such representation in STM was of insufficient robustness to survive until report. In a third condition, participants were given a word prior to the trial (a word-preview condition), and had to identify the picture that this word described. Thus, presented with (say) the word 'house', participants were able, at the fastest presentation rates, to identify the photograph showing a house. This ruled out the possibility that participants were using low level features in the preview photograph to do the task, and shows that participants were able to extract conceptual information from the photographs sufficient to make a match with the preview word. Potter's (1976) experiments show that a certain amount of time is required to form a durable representation in STM. Since this seminal paper, there have been attempts to quantify the time course of this process (e.g. Vogel, Woodman & Luck, 2006), which is termed *consolidation*.

Investigating the time course of consolidation processes

The representations of pictures formed at fast stimulus presentation rates in Potter's (1976) experiment were not durable enough to survive to later report. By durable, I mean that it is able to survive interference from competing stimuli. In Potter's (1976) experiment, subsequent picture stimuli served to mask preceding stimuli. It has been shown (e.g. Gegenfurtner & Sperling, 1993) that masking interrupts the consolidation process. Masking is therefore a useful tool to investigate the process of consolidation.⁵ By applying masks at

⁵ It should be noted that performance on the picture preview condition in Potter's experiment indicates that perceptual processing of the stimuli was not interfered with by the mask. In a similar way to more recent studies (e.g. Gegenfurtner & Sperling, 1993; Vogel, Woodman & Luck, 2006) the mask interfered with memory consolidation while perceptual processing was largely uninterrupted.

various times after stimulus onset, stimulus consolidation can be discriminately interrupted (e.g. Gegenfurtner et al., 1993). I utilise such a methodology to investigate the time course of consolidation of emotional face stimuli in Chapter 6, wherein I describe it in more detail.

Early memory and proto-objects

In discussing models of WM above, I mentioned briefly an idea put forward by Xu and Chun (2006, 2009; Xu, 2007) whereby object representations are formed in WM according to a two stage process, individuation and identification. The initially formed coarse representations have been referred to variously as preattentive objects (Wolfe & Bennett, 1997), object files (Kahneman, Treisman & Gibbs, 1992) or proto-objects (Rensink, 2000; Gao, 2011). Whereas 'object file' tends to be used to refer to both early and later representations, 'preattentive-' and 'proto-' object are terms used to refer exclusively to very early representations. Henceforth in discussing early memory I refer to such representations as proto-objects.

In a series of experiments Gao et al. (2011) revealed many of the attributes of early proto-objects and the later fully 'fleshed-out' WM representations that they can eventually become. In particular, Gao et al. suggest that: 1) Early proto-object representations are initially encoded via a process of parallel processing which occurs automatically (i.e. without the need to deploy attention); 2) Basic features such as colour and line orientation are rapidly consolidated in parallel when proto-objects are set up; 3) While these proto-object representations remain relatively stable, detailed information (such as orientation of a gap in a circle stimulus) is not consolidated in parallel but

must be consolidated later via a slower, non-automatic, attention-based process; 4) The basic features initially encoded into a proto-object representation are maintained with better veridicality than detailed information later added into the representation; 5) Once basic features and detailed information for a representation have been consolidated into visual WM, they are maintained there by distinctive processes (see also Gao, Shen, Shui, & Gao 2007; Xu, 2007; Xu & Chun, 2006, 2009).

Both Gao et al. (2011) and Xu and Chun (2009) posit an early form of memory based on setting up initial coarse representations that are then fleshed out. In contrasting this to the notion of fragile VSTM we see that they are not merely two names for the same early process: Fragile memory posits that the initial early information encoded is especially unstable and will decay unless a signal brings attentional resources to bear on a particular part of the fragile information, whereas early proto-object theory suggests that the items initially encoded (the proto-objects) are relatively stable while the detailed information that is later added into them is less stable. The proto-object theory and the fragile memory theory may represent two complementary early processes that are necessary for object based memory.

We will return to early memory processes again, specifically in the context of how they might be affected by emotional salience in Chapter 6.

A final note on terminology

Above I have noted that the STM and WM are terms that are often used interchangeably. However, some attempt has been made to functionally distinguish between these two concepts. In particular, Conway et al. (Conway,

Cowan, Bunting, Therriault & Minkoff, 2002) and Engle et al. (Engle, Tuholski, Laughlin & Conway, 1999) have used structural equation modelling⁶ to distinguish between a process that statistically loads onto a fluid intelligence variable (WM) and a nested process that operates merely for storage/retrieval (STM) which loads only onto simple memory task performance. In this hierarchical view, STM and WM are not two separate stores; rather STM is a specific subcomponent of WM. The experiments reported herein primarily focus on retrieval rather than computation/evaluation. Nonetheless, I shall henceforth exclusively use the term WM herein since in the context of the above distinction it incorporates the nested function of STM and so is the more 'general' term inasmuch as it conveys no indication of absence of evaluation mechanisms. Moreover, the majority of the stimuli I employ herein are posited to carry certain 'value associations' and are thus likely to attract the evaluation mechanisms inherent to WM rather than merely being stored for retrieval.

⁶ Structural equation modelling (SEM) combines multiple regression analysis and factor analysis to allow both the establishing of latent variables, and the testing of how well they load onto (predict) changes in dependent variables. Crucially, SEM also allows the relationship between latent variables to be estimated. In this way nesting of variables may be determined. Variables can be defined based on theoretical notions (in the case reported here STM and WM) allowing SEM to be used in a hypothesis-testing context (Bentler, 1986).

References – Chapter 1

- Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, *15*, 106-111.
- Atkinson, R. C. & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In Spence, K.W. & Spence, J.T. (Eds.) *The psychology of learning and motivation: Advances in research and theory*. Vol. 2. (89-195). New York: Academic Press.
- Atkinson, R. C. & Shiffrin, R. M. (1971). The control of short-term memory. *Scientific American*, *225*, 82–90.
- Baddeley, A. D. & Hitch, G. J. (1974). Working Memory. In G. A. Bower (Ed.), *The psychology of learning and motivation: advances in research and theory* (pp. 47-89). New York: Academic Press.
- Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, *4*, 417-423.
- Baddeley, A. D. (1996). The fractionation of working memory, *Proc. Natl. Acad. Sci. U. S. A.* *93*(24), 13468-13472.
- Baddeley, A. D. (2001). The magic number and the episodic buffer. *Behavioral and Brain Sciences*, *24*, 117-118.
- Barbey, A. K., Colom, R., Solomon, J., Krueger, F., Forbes, C., & Grafman J. (2012). An integrative architecture for general intelligence and executive function revealed by lesion mapping. *Brain*, *135*, 1154-1164.
- Barbey, A. K., Koenigs, M. & Grafman, J. (2013). Dorsolateral prefrontal contributions to human working memory, *Cortex*, *49*(5), 1195-1205.
- Barton, B., Ester, E. F., & Awh, E. (2009). Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *35*(5), 1359-1367.
- Bays, P. M. & Husain, M. (2008). Dynamic Shifts in Limited Working Memory Resources in Human Vision. *Science*, *(321)*, 851-854.
- Bentler, P. (1986). Structural modeling and Psychometrika: An historical perspective on growth and achievements. *Psychometrika*, *51*, 35-51.
- Bunting, M. F., Cowan, N., & Sualts, J. S. (2006). How does running memory span work? *The Quarterly Journal of Experimental Psychology*, *59*(10), 1691-1700.

- Bunting, M. F. & Cowan, N. (2005) Working memory and flexibility in awareness and attention. *Psychological research*, 69, 412-419.
- Chen, Z. J., & Cowan, N. (2009). Core verbal working-memory capacity: The limit in words retained without covert articulation. *Quarterly Journal of Experimental Psychology Vol 62(7)*, 1420-1429.
- Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & Psychophysics*, 27(3), 183-228.
- Conway, A. R. A., Cowan, N., Bunting, M. F., Theriault, D. J. & Minkoff, R. B. (2002). A latent variable analysis of working memory capacity, short-term memory capacity, processing speed, and general fluid intelligence. *Intelligence* 30 163-183.
- Cowan, N., Elliott, E. M., Saults, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., et al. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, 51, 42-100.
- Cowan, N. (1988). Evolving conceptions of memory storage, selective attention, and their mutual constraints within the human information processing system. *Psychological Bulletin*, 104, 163-191.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences* 24, 87-185.
- Cowan, N. (2010). Multiple Concurrent Thoughts: The Meaning and Developmental Neuropsychology of Working Memory. *Developmental Neuropsychology*, 35(5), 447–474.
doi:10.1080/87565641.2010.494985.
- Curby, K. M., Glazek, K., & Gauthier, I. (2009). A Visual Short-Term Memory Advantage for Objects of Expertise. *Journal of Experimental Psychology: Human Perception and Performance*, 35(1), 94-107.
- de Groot, A. D. (1965). *Thought and choice in chess*. The Hague: Mouton.
- de Groot, A. D. & Gobet, F. (1996). *Perception and memory in chess. Studies in the heuristics of the professional eye*. Assen: Van Gorcum.
- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, 11(6) 467-473. doi: 10.1111/1467-9280.00290
- Elo, A. (1978). *The rating of chess players, past and present*. New York: Arco.
- Engle, R. W., Tuholski, S. W., Laughlin, J. E., Conway, A. R. A. (1999). Working memory, short-term memory, and general fluid intelligence: A

latent-variable approach. *Journal of Experimental Psychology: General*, 128, 309-331.

- Ericsson, K. A. & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102, 211-245.
- Fougnie, D. (2008). The Relationship between Attention and Working Memory. In Noah B. Johansen (Ed.), *New Research on Short-Term Memory* (1-45). Nova Science Publishers, Inc. ISBN 978-1-60456-548-5
- Gao, T., Gao, Z., Li, J., Sun, Z., & Shen, M. (2011). *The perceptual root of object-based storage. Journal of Experimental Psychology: Human Perception and Performance*, 37, 1803-1823.
- Gao, T., Shen, M., Shui, R., & Gao, Z. (2007). The reconfiguration of task set has no effect on the efficiency of feature search. *Perception & Psychophysics*, 69, 345–352. doi:10.3758/BF03193755
- Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 845-866.
- Gobet, F. (1998). Expert memory: A comparison of four theories. *Cognition*, 66, 115-152.
- Gobet, F. & Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, 31, 1-40.
- Gobet, F. (2000). Some shortcomings of long-term working memory. *British Journal of Psychology*, 91, 551-570.
- Hikosaka, O. (2010). The habenula: from stress evasion to value-based decision-making. *Nature Reviews Neuroscience* 11, 503-513.
- Holding, D. H. (1985). *The psychology of chess skill*. Hillsdale, NJ: Erlbaum.
- Holding, D. H. (1992). Theories of chess skill. *Psychological Research*, 54, 10-16.
- Hollingworth, A., Richard, A. M., & Luck, S. J. (2008). Understanding the function of visual short-term memory: Transsaccadic memory, object correspondence, and gaze correction. *Journal of Experimental Psychology: General*, 137, 163-181. doi: 10.1037/0096-3445.137.1.163
- Hyun, J., Woodman, G. F., Vogel, E. K., Hollingworth, A., & Luck, S. J. (2009). The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35, 1140-1160. doi: 10.1037/a0015019

- Jung, R. E. & Haier, R. J. (2007). The parieto-frontal integration theory (PFIT) of intelligence: Converging neuroimaging evidence. *Behavioral and Brain Sciences*, 30(2): 135-187.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of files: Object-specific integration of information. *Cognitive Psychology*, 24, 175–219. doi:10.1016/0010-0285(92)90007-O
- Klingberg, T. (2010). Training and plasticity of working memory. *Trends in Cognitive Sciences* 7, 312–324.
- Kyllonen P. C & Christal R. E. (1990). Reasoning ability is (little more than) working-memory capacity?! *Intelligence* 14:389–433.
- Lewandowsky, S., & Oberauer, K. (2009). No evidence for temporal decay in working memory. *Journal of experimental psychology: Learning, memory, and cognition*, 35(6), 1545-1551.
- Liebe, S., Hoerzer, G. M. Logothetis, N. K. & Rainer, G. (2012). Theta coupling between V4 and prefrontal cortex predicts visual short-term memory performance. *Nature Neuroscience* 15, 456–462.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281
- Magen, H., Emmanouil, T-A., McMains, S. A., Kastner, S., Treisman, A. (2009). Attentional demands predict short-term memory load response in posterior parietal cortex. *Neuropsychologia* 47: 1790-1798.
- Makovski, T., & Jiang, Y. V. (2008). Proactive interference from items previously stored in visual working memory. *Memory & Cognition*, 36, 43-52
- Marsh, R. L., & Hicks, J. L. (1998). Event-based prospective memory and executive control of working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24, 346-339.
- Matsukura, M. & Hollingworth, A. (2011). Does visual short-term memory have a high-capacity stage? *Psychonomic Bulletin & Review*, 18, 1098-1104. DOI 10.3758/s13423-011-0153-2
- Morey, C. C. & Cowan, N. (2005). When Do Visual and Verbal Memories Conflict? The Importance of Working-Memory Load and Retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 31(4), 703-713.
- Olsson, H., & Poom, L. (2005). Visual memory needs categories. *Proceedings of the national academy of sciences, USA*, 102(24), 8776-8780.

- Paulesu, E., Frith, C. D. & Frackowiak, R. S. J. (1993). The neural correlates of the verbal component of working memory. *Nature*, 362, 342-345.
- Pessoa, L. & Ungerleider, L. G. (2004). Top-down mechanisms for working memory and attentional processes. In M. S. Gazzaniga (Ed.), *The New Cognitive Neurosciences, 3rd Edition*, pp. 919-930, MIT Press.
- Pollack, I., Johnson, L. B., & Knaff, P. R. (1959). Running Memory Span. *Journal of Experimental Psychology*, 57(3), 137-146.
- Portrat, S., Barrouillet, P., & Camos, V. (2008). Time-related decay or interference-based forgetting in working memory. *Journal of experimental psychology: learning, memory, cognition*, 34(6), 1561-1564.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of Experimental Psychology: Human Learning and Memory*, 2, 509-522.
- Ranganath, C. & Blumenfeld, R. S. (2005). Doubts about double dissociations between short- and long-term memory. *Trends in Cognitive Sciences*, 9(8), 374-380.
- Rensink, R. A. (2000). The dynamic representation of scenes. *Visual Cognition*, 7, 17–42. doi:10.1080/135062800394667
- Repovš, G. & Baddeley, A. (2006). The multi-component model of working memory: explorations in experimental cognitive psychology, *Neuroscience* 139 (2006), 5–21.
- Robbins, T., Anderson, E., Barker, D., Bradley, A., Fearneyhough, C., Henson, R., Hudson, S. & Baddeley, A. D. (1996). Working memory in chess. *Memory and Cognition*, 24, 83–93.
- Sligte I. G., Scholte H. S., Lamme V. A. F. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLoS ONE* 3(2): e1699. doi:10.1371/journal.pone.0001699
- Sligte, I. G., Scholte, H. S., and Lamme, V. A. (2009). V4 activity predicts the strength of visual short-term memory representations. *Journal of Neuroscience*, 29, 7432–7438.
- Sligte, I. G., Vandenbroucke, A. R. E., Scholte, H. S. and Lamme, V. A. F. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in Psychology* 1, 1-10. doi: 10.3389/fpsyg.2010.00175
- Smith, E. E. & Jonides, J. (1997). Working Memory: A View from Neuroimaging. *Cognitive Psychology*, 33, 5-42.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74, 1–29.

- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, *428*, 751–754.
- Vandenbroucke, A. R. E., Sligte, I. G., Lamme, V. A. F. (2011). Manipulations of attention dissociate fragile Visual Short-Term Memory from Visual Working Memory. *Neuropsychologia* *49*, 1559-1568.
- Vogel, E. K., Woodman, G. F. & Luck, S. J. (2006). The Time Course of Consolidation in Visual Working Memory. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 6, 1436-1451.
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, *131*, 48-64.
- Wolfe, J. M., & Bennett, S. C. (1997). Preattentive object files: Shapeless bundles of basic features. *Vision Research*, *37*, 25–44. doi:10.1016/S0042-6989(96)00111-3
- Xu, Y. (2008). Representing connected and disconnected shapes in human inferior intra-parietal sulcus. *Neuroimage* *40*, 1849-1856.
- Xu, Y. (2007). The role of the superior intra-parietal sulcus in supporting visual short-term memory for multi-feature objects. *Journal of Neuroscience*, *27*, 11676-11686.
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short- term memory for objects. *Nature*, *440*, 91–95
- Xu, Y., & Chun, M. M. (2007). Visual grouping in human parietal cortex. *Proceedings of the National Academy of Sciences, USA*, *104*, 18766-18771. doi:10.1073/pnas.0705618104
- Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in cognitive sciences*, *13*(4), 167-174.
- Zhang, W., & Luck, S. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233-235.

CHAPTER 2.

Emotional Faces

What do we perceive when we look at someone's face? Faces provide us with a wealth of social and non-social information, allowing us to more efficiently interact with others. First, provided a face image is not impoverished (blurred, occluded etc.) to too great an extent we can perceive that the image is that of a face. Indeed, basic facial expressions can be detected using only the low frequency spatial information that is available in the periphery of the visual field (e.g. Vuilleumier & Pourtois, 2007). Second, we can tell whether we are looking at an image of a 'real' or stylised face. This seemingly simple classification is a non-trivial operation in its own right. Cartoon faces share many of the features and much of the feature organisation of real faces yet humans can reliably make such distinctions with minimal effort while machine-based vision systems must invoke complex algorithms to solve such a question (Figure 2.1)⁷. Third, generic identity information is present, allowing us to identify the race, age, gender and even general health of the individual. Fourth, if we have encountered the face before we may have a stored representation of it in LTM allowing us to identify the specific individual. Fifth, information as to an individual's emotional state may be displayed in their facial expression. Such information is of particular importance as it can be used to predict an individual's behaviour. Sixth, information as to an individual's attentional state may be read from their eyes, with gaze providing an index of overt attention. Seventh, information as to an individual's alertness may also be inferred from facial features (e.g. half-closed eyes indicating sleepiness). Eighth, speech patterns may be inferred from facial movements. Relying on a blend of some of the above, various compound characteristics may be determined

⁷ Cartoon and real faces seem to share a common coding mechanism so that both are recorded within similar face space (see below) as evidenced by the fact that repeated exposure to cartoon faces may act through a process of adaptation to shift the prototypical face against which new (real) faces are compared and thus alter perceptions of attractiveness for real faces such that more 'cartoonish' faces are preferred (Chen, Russell, Nakayama, & Livingstone, 2010).

(e.g. attractiveness, Jung, Ruthruff, Tybur, Gaspelin & Miller, 2012; cuteness, Alley, 1981; Lorenz, 1971). A high level compound characteristic may also be specified; face value. Here, different faces may have different potential utility to an individual depending on current and long term goals. For example, to fulfil a goal of seeking safety or protection, adult faces have a greater potential to be able to fulfil this goal than child faces since they signify individuals able to afford such protection. Likewise, when the current goal is to mate, attractive faces are up-weighted as goal fulfilling relative to unattractive faces. Thus, the compound characteristic of value is modulated by needs and goals. I discuss this idea in more detail in the chapter on value learning.




Figure 2.1. Distinguishing between a real and cartoon face is an effortless process for humans, but is a non-trivial problem for computer-based vision systems. Left (cartoon) face reproduced from <http://www.wikihow.com/Draw-a-Face>. Right (real) face of the author.

Recognising Faces

Perhaps the most well-known model of face recognition is Bruce and Young's (1986) functional model (Figure 2.2). In this framework, the retrieval of personal identity of a familiar face is a sequential, hierarchical process which proceeds through a perceptual (structural encoding) stage, followed by semantic evaluation of the face within face recognition units, each of which holds stored structural codes associated with a particular (viewpoint and expression specific) person's face, the weighted outputs from which can then activate 'person identity nodes', which are semantic representations of known individuals. One person identity node exists in LTM for each known individual. Finally, names of individuals may be retrieved from stored LTM codes linked to each person identity node.

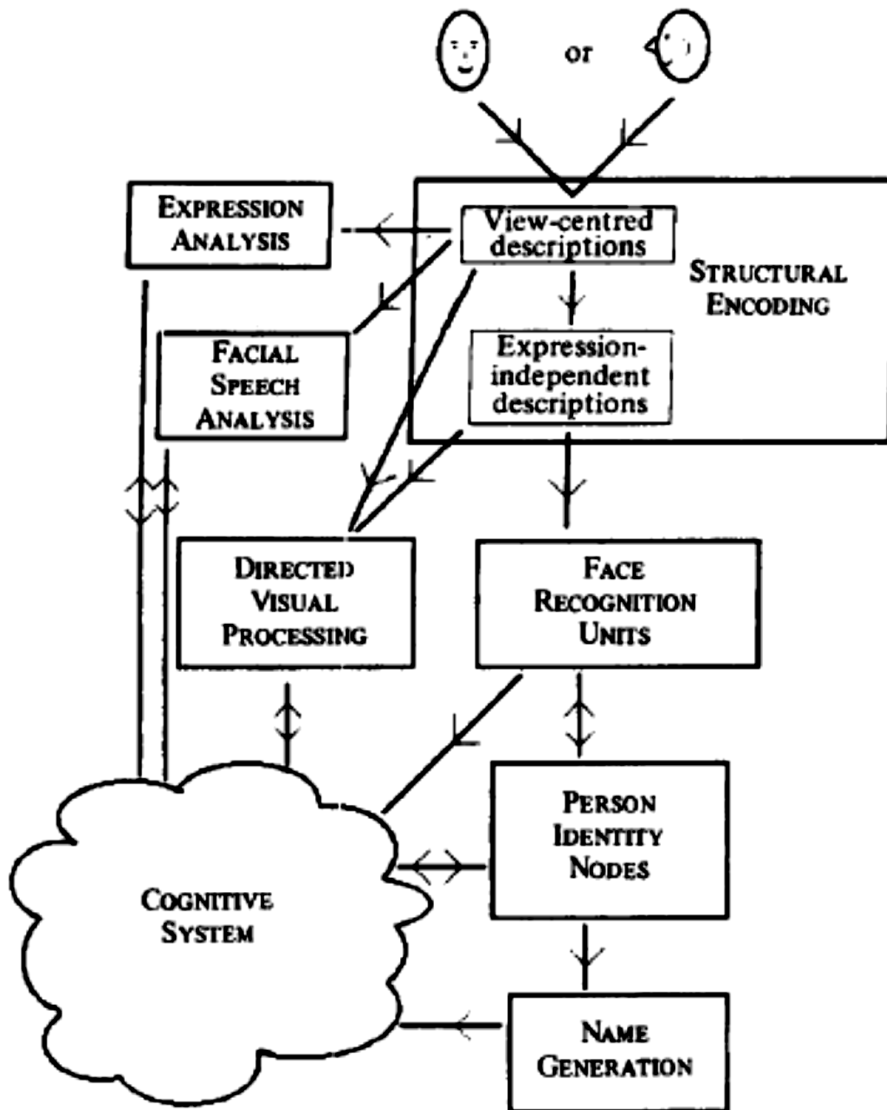


Figure 2.2. The Bruce and Young functional model of face recognition.

Reproduced from Bruce & Young (1986).

A key implication of this model is that the face identification sequence proceeds independently of other processes concerned with evaluation of the information relating to face expression (including facial speech). This independence has been observed in studies of prosopagnosics, where double dissociations have been observed between ability to identify faces and ability to

evaluate emotional expressions (e.g. Young, Newcombe, de Haan, Small, & Hay, 1993), as well as in PET and MRI studies (e.g. Sergent, Ohta, MacDonald, and Zuck, 1994). However, not all studies agree that the two pathways are entirely independent (e.g. Jackson, Wu, Linden & Raymond, 2009; Schweinberger, Burton, and Kelly, 1999; Schweinberger & Soukup, 1998) and there may be some interaction between them.

Face space

Face space is a construct wherein any face can be realised as an entity that may be represented within a multidimensional space whose dimensions reflect deviations away from a prototypical face that represents the mean of a person's experience such that faces may be judged by proximity in face space to the prototype (Valentine 1991). While face space was first proposed with dimensions that reflect the physiognomic characteristics of faces (Valentine, 1991), it is also possible to determine dimensions that represent complex social and non-social face characteristics and determine where a given face lies in face space relative to such dimensions (e.g. attractiveness; Chen et al., 2010).

The N170 ERP component – a face specific indicator?

The N170 component of an event related potential (ERP) has historically been thought to reflect neural processing of faces specifically. First described by Bentin et al. (Bentin, McCarthy, Perez, Puce & Allison, 1996), this component is right hemisphere lateralised (Rossion & Jacques, 2008), may originate from posterior superior temporal sulcus (Itier & Taylor, 2004), and seems to be specific to the structural encoding of faces rather than their identification (Bentin & Deouell, 2000). However, recent research has also suggested that this

component can be elicited by similarity between stimuli (Thierry, Martin, Downing & Pegna, 2007, though see Rossion et al., 2008; Eimer, 2011), or by cropped photographic stimuli of non-face items (Dering, Martin, Moro, Pegna, & Thierry, 2011), or by stimuli within the observer's field of expertise (Busey and Vanderkolk, 2005), including greater N170 amplitudes for more familiar faces (Caharel et al., 2002). The logic that stimuli of expertise should elicit this component, perhaps by recruiting the fusiform face area (FFA; the N170 having been posited to reflect a degree of FFA activation, e.g. Halgren, Raij, Marinkovic, Jousmaki & Hari, 2000, though see Itier et al., 2004), derives from the *subordinate level expertise model* (Tarr and Gauthier, 2000), which argues that specific brain areas such as the FFA are selectively activated by faces because faces are usually classified or identified at the subordinate (individual) level based on analysis of configural information rather than at the basic level (i.e. mere classification as being a face) as with many other stimuli. Taking the logic of this model further, it is apparent that if an observer has sufficient perceptual expertise with a particular non-face object category (i.e. is an expert in such a category) then configural information might also be used to discriminate objects at a subordinate level within the category of expertise (e.g. car experts rapidly identifying the make and model of cars presented to them as stimuli), and crucially, processing of these objects should recruit the same brain areas that are activated during face processing. Thus, in this model areas such as the FFA which are commonly observed as selectively responding to faces are not in fact dedicated purely to faces but instead to "a specific type of visual processing (recognition of individual object tokens based on configural information)" (Eimer, 2011, p338). From this point of view, what makes faces a special class of stimuli is that fact that we (humans) are experts in their

perception and are therefore able to recruit the so-called ‘face-specific’ brain areas to process them. To some extent then, this controversy in what the N170 reflects parallels the debate as to the function of the FFA⁸. As it is not the primary purpose of this thesis to speculate on the underlying neural substrates of face processing I set these questions aside here.

Ethnicity and the cross-race effect

When an observer of one particular race views faces of a race with which he or she is not familiar, such faces are hard to discriminate from one another and ‘look all the same’ (Byatt & Rhodes, 2004). Additionally, it is harder to judge the emotion portrayed by such faces, (Elfenbein & Ambady, 2002) and to distinguish between races other than one’s own⁹. This phenomenon is known as the cross-race effect or other-race effect and its occurrence seems to effect a shift in the heuristics used by observers to make memory judgements, in particular occasioning over-use of *resemblance heuristics* (evaluation of stimuli in relation to internal stereotypes of templates, a problem common to the use of eye-witness testimonies; Kleider & Goldinger, 2006). It is important to note that it is the level of

⁸ There is some debate as to the exact role of the FFA. Whereas the prevailing view is that the FFA is selectively activated in the presence of faces and therefore plays an important role in their processing (e.g. Kanwisher, McDermott & Chun, 1997) it has more recently been suggested that the FFA might also be involved in discriminating objects of expertise (Gauthier, Skudlarski, Gore & Anderson, 2000). Even if the FFA is not entirely domain specific, the evolution of an area capable of functioning *as if* it were specific to the domain of faces, as seen in many experiments, illustrates the importance of being able to rapidly process and respond to faces compared to other generic (non-face) stimuli.

⁹ Examples of this phenomenon in the social world include the issue of a politically incorrect guide ‘How to Spot a Jap’ (Caniff, 1942) in the US army guidebook issued to soldiers sent to China during the Second World War. This cartoon guide was intended as an aid in distinguishing between the facial characteristics Japanese (enemy) and Chinese (allied) people but played on racial stereotypes of the time. Additionally, recognition of the cross-race effect has long informed the way in which identity parades are carried out; for example, in many countries it is a legal requirement that participants in an identity parade have similar facial features to the suspect (e.g. the Commonwealth Crimes Act of 1914 requires that participants 1) “resemble the suspect in age, height and general appearance; and 2) “not have features that will be visible during the parade that are markedly different from those of the suspect as described by the witness before viewing the parade”.

exposure or expertise with faces of different races that seems to determine the level of the cross-race effect, and exposure to faces of another race ameliorates the difficulty in distinguishing between faces and processing expressions (Ekman & Friesen, 1976; Elfenbein & Ambady, 2002). Accordingly, all participants who took part in experiments in which faces were used as stimuli herein were asked how long they had lived in a country in which the main race was Caucasian (all face stimuli used in experiments herein were of Caucasian individuals; see Appendix C). Data for those who had resided in such a country for less than three months were excluded from all analysis.

Threat and angry faces

Anger is a facial expression assumed to have its evolutionary roots in a more complete behavioural attack response incorporating a baring of the teeth and furrowing of the brow with non-facial bodily signals such as expansion of the chest and tightening of muscles in the arms (Darwin, 1972/1998). Expression of extreme anger can therefore be predictive of direct physical threat. Moreover, processing such expressions is particularly important since they may be the only or first indication of such a threat. Indeed, in humans, faces represent a special class of stimuli in that their properties may rapidly change from denoting non-threatening to threatening information (whereas other threat items such as spiders and snakes must always be treated as threatening).

While an expression of mild anger might not usually predict physical threat from the expresser, mild anger (or annoyance) is often indicative of disapproval, which is inherently punishing since it in turn signals the possibility of social punishment, such as loss of reputation or withdrawal of social privilege.

The threat of social punishment must be taken very seriously since such punishment can be as detrimental to an individual's chance of survival as physical punishment. In particular, an individual's reputation modulates the amount of altruism or cooperation that individual receives from other non-genetically related individuals (Fehr, 2004, Panchanathan & Boyd, 2003). A mechanism supporting efficient detection of even mild expressions of anger is therefore of benefit to an individual's ability to function in the social world.¹⁰ Such a mechanism does indeed exist, and is supported by a ventral neural pathway and the amygdala (see below).¹¹

Attention and angry faces

Angry faces hold attention, with disengagement of gaze from angry faces reported to be slower than that from either happy or neutral faces (Belopolsky, Devue, & Theeuwes, 2011). However, in infants as young as 4 months, a pattern of gaze avoidance of angry and fearful faces has been reported, although the classic pattern of specific avoidance of eye regions (cf. van Honk & Schutter, 2007) in threatening faces appears to be absent at this stage of development

¹⁰ Game theory can be used to model the level to which it is provident to socially engage with others given the estimation of likely rewards and risks in doing so. In particular, this discipline illustrates one of the key functions of expressing anger. In the public goods game (e.g. Isaac, Walker & Williams, 1994) signalling of intention to cooperate can be modelled as a factor in determining cooperation of neighbours, while signalling of intention to punish non-cooperators can be modelled as a less costly alternative to actually enacting punishment. Thus, the signalling of anger can be used to correct behaviour by signalling intent to punish, and this is more economical than punishment itself. Even in the most simple of games, such as the well-known prisoners' dilemma game (Poundstone, 1992), signalling of intent to cooperate and punish defection leads to an optimal strategy given participants able to read such signals who behave according to self-interest.

¹¹ Interestingly, it is possible to bias the way in which faces are processed by this mechanism by administering the chemical oxytocin (OT; Guastella, Mitchell & Dadds, 2008). Administration of OT has been associated with lowered levels of amygdala activity in the presence of threatening (angry and also fearful) faces (Kirsch, Esslinger, Chen, Mier, Lis, Siddhanti, Gruppe, Mattay, Gallhofer & Meyer-Lindberg, 2005), and increased gaze time to the regions of faces associated with social cognition, in particular the eyes (Guastella et al., 2008). It is possible that the lowered amygdala activity results in decreased perception of social threat and thus mediates the increased gaze time to eyes and the associated enhanced ability to read emotion in faces (Domes, Heinrichs, Michel, Berger & Herpertz, 2007).

(Hunnius, de Wit, Vriens & von Hofsten, 2011). This may reflect an increase in vigilance to the environment (a broadening of attentional deployment) or a social or emotional avoidance strategy for each of these stimulus classes.

Threatening faces may sometimes be avoided attentionally, especially among older adults with higher functioning executive control systems (Isaacowitz, Toner & Neupert, 2009) who may use such avoidance as a mood regulation strategy. This negativity avoidance bias seems to be attenuated among adults whose culture reflects a self-schema of group membership rather than independence (Fung, Isaacowitz, Lu & Li, 2010), perhaps because the utility of angry faces to those who tend to operate by achieving group consensus is particularly high (group conflict needing to be resolved before group-based actions can be taken).

When investigating WM for faces, it is therefore possible that attention plays a role in determining the entry of faces of different emotions into WM in a way that may favour the encoding of some faces over others. In Experiments 2 and 7 I investigate this by monitoring eye movements (an index of overt attention) when participants view arrays of faces.

Visual WM for faces

Memory for faces involves subtle and complicated processes, and has been shown to be influenced by factors including congruency of gender between the observer and the face stimulus, gaze direction of the face, and the location of the face in the visual field, with interactions between these three variables (Vuilleumier, George, Lister, Armony & Driver, 2005). Herein, however, I restrict my investigations specifically to short term memory for faces and in particular to

faces that have motivational value either by virtue of their emotional expression or due to imbued value association occasioned by prior learning experiences.

Emotional faces have an intrinsic emotional value to neurotypical observers. Although there are many objects in daily life that may elicit emotional responses of one type or another (e.g. spiders), the emotional response experienced, and the intensity of this response will vary markedly between different people. While an arachnophobe may experience an intense negative response to a spider, an entomologist may experience an emotionally positive response. Fortunately, with emotional faces there is some degree of uniformity to the emotional response elicited so that the valence of the response, if not necessarily its intensity may usually be assumed to hold across neurologically normal participants. Responses to emotional faces are likely to be 'hardwired' for evolutionary reasons – we must be able to correctly identify anger expressed on a person's face in order to respond to this, and our own emotional response (accompanied by increased arousal) may motivate an appropriate response. There are several other advantages to using emotional faces as a stimulus class when investigating response to threat: They are subserved by specific neural substrates (e.g. the FFA), allowing for certain advantages in imaging research relative to other classes of stimuli without a unique cortical area. By varying the intensity of emotion expressed in a face stimulus (e.g. by morphic manipulation), one can vary the level of positive or negative emotional content in the stimulus. Additionally, faces are relatively complex stimuli, with many features comprising the overall face, but may be thought of as being objects of universal expertise, such that features are typically processed in parallel as a single face stimulus (Curby, Glazek & Gauthier, 2009), although this parallel, holistic processing effect

is lessened when faces encountered are of an unfamiliar race (Tanaka, Kiefer, & Bukach, 2004), or inverted (Taubert, Aporp, Aagten-Murphy, & Alais, 2011; Yin, 1969).

Although it might be thought that identifying the emotional expression of a face would rely mainly on visual WM, there is recent evidence suggesting that auditory working memory can also play a crucial role in this process. Phillips et al. (Phillips, Channon, Tunstall, Hendenstrom & Lyons, 2008) used a 2-back paradigm to vary the level of auditory working memory resources available between conditions, and found that while *discriminating* same/different emotion between two faces did not require auditory working memory, *assigning emotional labels presented on screen* to faces did require auditory working memory, perhaps because these labels consisted of words, and so required verbal working memory to be evaluated themselves before pairing with a face. Although this finding does not suggest that verbal WM is required for the detection of, or emotional response to, an emotional face, it serves as a useful cautionary note regarding the different demands on WM resources that different emotional evaluation tasks may make, particularly as it is common to load auditory WM with an auditory suppression task when investigating visual WM.

Working memory load for faces seems to be constrained to about two faces (Jackson & Raymond, 2008). Slot-based models of visual WM account for this by suggesting that each face attracts two slots to facilitate processing (Jiang, Shim & Makovski, 2008). This may be due in part to the relative complexity of this stimulus class, although Jiang et al. (2008) found that under certain conditions (when high resolution representations had to be stored for comparison to very

similar probes) simple objects (oriented lines) could exhaust working memory to a greater extent than the more complex faces.

A recent experiment by Jackson and Raymond (2008) using a conventional change-detection WM task (discussed in Chapter 5) has found a WM advantage for familiar faces over non-familiar faces in working memory. This effect was robust across a manipulation of concurrent verbal memory load (thus ruling out a strategic contribution from verbal working memory as an explanation for the familiarity effect), but was abolished by face inversion (ruling out an explanation based on low-level feature processing). The WM benefit experienced by familiar faces over unfamiliar faces therefore seems to be served by holistic representations stored in LTM. Indeed, it has been noted that visual LTM is likely to contribute to a varying degree to all short-term visual memory tasks (Jiang et al., 2008; Makovski & Jiang, 2008), and this finding would lend support to models of WM that posit temporary activations of LTM codes (e.g. Cowan, 2001).

Supporting neuropsychological evidence for this familiarity benefit being based on a holistic rather than feature based representation comes from ERP studies that have observed no difference in N170 (reflecting structural encoding) between familiar and non-familiar faces, but a difference in N400 (Eimer, 2000; Bentin & Deouell, 2000).

Neuropsychology of emotional faces in WM

Using a functional magnetic resonance imaging (fMRI) methodology and a delayed match-to-sample task for emotion and identity of faces (to facilitate active maintenance of both these types of information simultaneously in WM), LoPresti et al. (2008) found that despite evidence for face emotion and identity being initially

processed in different brain regions (e.g. Hasselmo, Rolls & Baylis, 1989), the maintenance of both these pieces of important social information is related to sustained activity in the orbitofrontal cortex (OFC), the amygdala, and the hippocampus; three highly connected neural structures. This contrasts to previous research showing delay related activity for faces in the temporal occipital cortex TOC (e.g. Druzgal and D'Esposito, 2003; Postle, Druzgal & D'Esposito, 2003; Xu and Chun, 2006) and LoPresti et al. have suggested that these two mechanisms may come into play in situations where, when multiple information strands (in this case both identity and emotion) need to be retained, the TOC processes are insufficient, and a circuit including the amygdala, hippocampus and OFC becomes active (LoPresti et al., 2008).

Emotional faces have recently been shown to have an advantage in WM tasks, at least when memory load is high (Langeslag, Morgan, Jackson, Linden & Van Strien, 2009), and a particular benefit has been observed for angry over both neutral and happy faces (Jackson et al., 2009). Two recent experiments have begun to delineate the neural substrates underlying these advantages. Jackson et al. (Jackson, Wolf, Johnston, Raymond & Linden, 2008) used fMRI to show that the benefit for angry faces observed in a previous study occurred with increased activation in the right superior temporal sulcus (STC), ventrolateral prefrontal cortex (vIPFC), and basal ganglia. Further research (Langeslag et al., 2009) into the time course of neural events underlying this advantage revealed both an increased N250r during face retrieval for emotional (both happy and angry) faces relative to neutral faces, indicating enhanced repetition (rehearsal) for the emotional faces, and a decreased P3b during retrieval – indicating reduced deployment of resources to probe items, and by extrapolation reflecting a larger

deployment of resources to the maintenance of codes in working memory. Thus, the advantage observed for emotional faces seems to be due at least in part to an increased allocation of resources to these codes such that they are better maintained (and therefore better retrieved) in working memory. There remains, however, the possibility that a maintenance advantage is only part of the story; it could be that emotional faces, or perhaps certain types of emotional face, also benefit from an encoding advantage, perhaps due to a preferential attraction of attention to themselves. In Chapter 5 I investigate whether attention can account for the boost in WM performance observed for angry faces. In Chapter 6 I investigate whether such faces benefit from a boost in the way their initial WM traces are consolidated into more durable representations.

Feldmann-Wüstefeld et al. (Feldmann-Wüstefeld, Schmidt-Daffy & Schubö, 2010) have used the N2pc component (Vogel & Machizawa, 2004) of ERPs to measure attentional shifts to angry and happy faces in a face-in-the-crowd task. The N2pc is a sustained negative voltage that arises at posterior regions on the hemisphere contralateral to the attended stimulus, about 180 to 300 ms after stimulus onset and can be taken as indicative of a shift in attention.¹² Among a high socially anxious group, they found larger and earlier onset N2pc for angry as opposed to happy faces, indicating spontaneous allocation of attention to angry faces in a face-in-the-crowd task and quicker behavioural detection of angry as opposed to happy faces. Feldmann-Wüstefeld et al. argue that this finding supports the theory that behavioural benefits in angry face processing are facilitated by such spontaneous attention shifts, at least among certain (high

¹² Additionally, the N2pc potential increases as the number of items to be remembered in WM increases and asymptotes when the number of items reaches capacity for any given person. It can therefore be used as an indicator of the number of items held in working memory at a particular time. (Vogel & Machizawa, 2004; Vogel, McCollough & Machizawa, 2005).

anxiety) participant groups. However, they also found an early posterior negativity (EPN) for angry faces that occurred earlier (160 ms after stimulus onset) than the N2pc (230 ms after stimulus onset), indicating that identification of the emotion of a face is performed at least to some level preattentively. This EPN may depend on prior tagging in the priority subcortical threat pathway (Schupp, Öhman, Junghöfer, Weike, Stockburger & Hamm, 2004), and would be a candidate mechanism for guiding/initiating the rapid attentional shifts and N2pc toward angry faces seen in the face-in-the-crowd paradigm. It is also possible that the general angry face benefit (that is, the processing benefit that angry faces exhibit across a wide variety of behavioural tasks) is at least partly facilitated by a process that occurs preattentively as indicated by the EPN (Feldmann-Wüstefeld et al, 2010).

Further evidence that automaticity of face emotion processing is greater in high trait anxious groups comes from an ERP study by Holmes et al. (Holmes, Nielsen, Tipper & Green, 2009), who found that while frontal (late positive potential; LPP) augmentations by angry faces were eliminated among a low trait anxious group when task difficulty in an *n*-back task was increased, such attenuation was not induced among high trait anxious individuals. Also of note in this study was the finding that early emotion related electrocortical effects (P1 and EPN) were unaffected by increasing task difficulty, supporting the theory that the mechanisms reflected by such effects are largely automatic, at least in a situation where stimuli fall within the focus of spatial attention and attentional shifts are not therefore required.

A special pathway for emotion processing?

The preferential access angry faces have to WM resources might be supported by the specific neuronal pathways by which threat information and more generally fear related information can be relayed to appropriate processing areas in the brain. In particular, it has long been thought that there may exist a separate, ventral, rapid, automatically activated and largely unconscious subcortical pathway by which threat information can be relayed directly to the amygdala without having to take the slower dorsal (cortical) path common to most other forms of visual information. Such a ventral stream is believed to include the superior colliculus, the pulvinar and the amygdala (Morris, Ohman, & Dolan, 1999) and be phylogenetically ancient, representing as it does a crucial part in the rapid ability to respond to threat (LeDoux, 1996). Evidence for the unconscious ventral threat pathway comes from patients with cortical blindness who are able to perceive facial expressions in the blind portion of their visual fields (de Gelder, Vroomen, Pourtois & Weiskrantz, 1999), from the finding that blocking of conscious perception of emotional faces by backward masking does not block non-conscious perception (as indexed by skin conductance changes) of their emotional expressions (Esteves, Dimberg, & Ohman, 1994), and from the observation that masked fearful facial expressions lead to an increase in amygdala activity (Whalen et al., 1998). While the routing of threat information down this pathway is often thought to be automatic in nature (i.e. it can proceed in the absence of attention, or 'pre-attentively' and be relatively impervious to conscious control; Öhman & Mineka, 2001), this view is not universally accepted (e.g. see Bishop, Jenkins & Lawrence, 2007). While information in the ventral threat stream is not thought to be directly available to consciousness, this is not to say that it does not influence the contents of consciousness. For example, by

interacting with cortical processes information in the ventral stream may influence experience of the conscious world (Morris et al., 1998, Tamietto & De Gelder, 2010). Interestingly, it seems that conscious perception of a stimulus may block (Jolij & Lamme, 2005), or at least delay (Tamietto et al., 2009) the subcortical processing that would occur in absence of conscious perception.

Although the experiments reported herein do not specifically address whether the ventral threat/negativity pathway directly facilitates encoding of threat-based information (e.g. angry faces) into WM, it is of note that the ventral pathway, which has been posited not only as being involved in the routing of negative emotional expressions but also to a lesser degree positive facial emotions (e.g. Tamietto et al., 2009), may play a role in facilitating rapid categorisation of visual stimuli into those that must be prioritised (have motivational salience – see Chapter 3) and those that are of less importance. This rapid categorisation, however it is achieved, is necessary if resources are to be rapidly differentially deployed among competing visual stimuli. I discuss this issue further in Chapter 9. Chapters 5 and 6 are devoted to investigation of causes of the boost in WM performance observed for emotional and in particular angry faces. Chapter 7 also looks at WM for face stimuli but examines whether previous pairing with reward or loss can imbue a value into such stimuli and whether retention of such a value association affects retention in WM in a similar way that expression does.

References – Chapter 2

- Alley, T. (1981). Head shape and the perception of cuteness. *Developmental Psychology*, 17(5), 650-654.
- Belopolsky, A. V., Devue, C., & Theeuwes, J. (2011). Angry faces hold the eyes. *Visual Cognition*, 19(1), 27-36. doi:10.1080/13506285.2010.536186
- Bentin, S. & Deouell, L. Y. (2000). Structural Encoding and Identification in Face Processing: ERP Evidence for Separate Mechanisms. *Cognitive Neuropsychology*, 17 (1/2/3), 23-54.
- Bentin, S., McCarthy, G., Perez, E., Puce, A., Allison, T. (1996). Electrophysiological studies of face perception in humans. *Journal of Cognitive Neuroscience*, 8, 551-565.
- Bishop, S. J., Jenkins, R., & Lawrence, A. D. (2007). Neural processing of fearful faces: effects of anxiety are gated by perceptual capacity limitations. *Cerebral cortex*, 17(7), 1595-603. doi:10.1093/cercor/bhl070
- Bruce, V., & Young, A. (1986). Understanding Face Recognition. *British Journal of Psychology*, 77(3), 305-327.
- Busey, T. A. & Vanderkolk, J. R. (2005). Behavioral and electrophysiological evidence for configural processing in fingerprint experts. *Vision Research*, 45, 431-448.
- Byatt, G., & Rhodes, G. (2004). Identification of own-race and other-race faces: Implications for the representation of race in face space. *Psychonomic Bulletin & Review*, 11, 735-741.
- Caharel, S., Poiroux, S., Bernard, C., Thibaut, F., Lalonde, R. & Rebai M. (2002). ERPs associated with familiarity and degree of familiarity during face recognition. *International Journal of Neuroscience*, 112, 1531-1544. DOI: 10.1080/00207450290026
- Caniff, M. (1943). How to Spot a Jap. In: United States. Army Service Forces. Special Service Division (Ed.). A Pocket Guide To China. Washington; War and Navy Departments.
- Chen, H., Russell, R., Nakayama, K., & Livingstone, M. (2010). Crossing the 'uncanny valley': Adaptation to cartoon faces can influence perception of human faces. *Perception*, 39(3), 378-386. <http://dx.doi.org/10.1068/p6492>
- Commonwealth Crimes Act, section 3ZM (1914). Retrieved from : http://www.austlii.edu.au/au/legis/cth/consol_act/ca191482/s3zm.html
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences* 24, 87-185.

- Curby, K. M., Glazek, K., & Gauthier, I. (2009). A Visual Short-Term Memory Advantage for Objects of Expertise. *Journal of Experimental Psychology: Human Perception and Performance*, 35(1), 94-107.
- Darwin, C. (1872/1998). The expression of emotion in man and animals. New York: Oxford University Press.
- de Gelder, B., Vroomen, J., Pourtois, G. & Weiskrantz, L. (1999). Non-conscious recognition of affect in the absence of striate cortex. *Neuroreport* 10, 3759-3763.
- Dering, B., Martin, C. D., Moro, S., Pegna, A. J. & Thierry G. (2011). Face-sensitive processes one hundred milliseconds after picture onset. *Frontiers in Human Neuroscience* 5(93). doi: 10.3389/fnhum.2011.00093.
- Domes, G., Heinrichs, M., Michel, A., Berger, C. & Herpertz, S. (2007). Oxytocin improves "mind-reading" in humans. *Biological Psychiatry* 61, 731-733.
- Druzgal, T. J., D'Esposito, M. (2003). Dissecting contributions of prefrontal cortex and fusiform face area to face working memory. *Journal of Cognitive Neuroscience*, 15, 771-784.
- Eimer, M. (2000). Event related potentials distinguish processing stages involved in face perception and recognition. *Clinical Neuropsychology*, 111, 694-705.
- Eimer, M. (2011). The Face-Sensitive N170 Component of the Event-Related Brain Potential. In: Calder, A. J., Rhodes, G., Johnson, M. H. & Haxby, J. V. (Eds.), *The Oxford Handbook of Face Perception* (pp. 329-344). Oxford University Press.
- Ekman, P., & Friesen, W. (1976). Pictures of facial affect. Palo Alto, CA: Consulting Psychological Press.
- Elfenbein, H. A. & Ambady, N. (2002). On the universality and cultural specificity of emotion recognition: A meta-analysis. *Psychological Bulletin*, 128(2), 203-235.
- Esteves, F., Dimberg, U. & Ohman, A. (1994). Automatically elicited fear: conditioned skin conductance responses to masked facial expressions. *Cognition & Emotion* 8, 99-108.
- Fehr, E. (2004). Human behaviour: Don't lose your reputation. *Nature*, 432, 449-450. doi:10.1038/432449a
- Feldmann-Wüstefeld, T., Schmidt-Daffy, M. & Schubö, A. (2011). Neural Evidence for the Threat Detection Advantage: Differential Attention Allocation to Angry and Happy Faces. *Psychophysiology*, 48, 697-707.
- Fung, H. H., Isaacowitz, D. M., Lu, A. Y. & Li, T. (2010). Interdependent Self-Construal Moderates the Age-Related Negativity Reduction Effect in Memory and Visual Attention. *Psychology and Aging*, 25(2), 321-329.

- Gauthier, I., Skudlarski, P., Gore, J. C. & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature Neuroscience*, 3(2), 191-197.
- Guastella, A. J., Mitchell, P. B. & Dadds, M. R. (2008). Oxytocin Increases Gaze to the Eye Region of Human Faces. *Biological Psychiatry*, 63, 3-5.
- Halgren, E., Raij, T., Marinkovic, K., Jousmaki, V., Hari, R. (2000). Cognitive response profile of the human fusiform face area as determined by MEG. *Cerebral Cortex*, 10, 69-81.
- Hasselmo, M. E., Rolls, E. T. & Baylis, G. C. (1989). The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behav Brain Res*, 32, 203-218.
- Holmes, A. Nielsen, M. K., Tipper, S. & Green, S. (2009). An electrophysiological investigation into the automaticity of emotional face processing in high versus low trait anxious individuals. *Cognitive, Affective, & Behavioural Neuroscience*, 9(3), 323-334.
- Hunnius, S., de Wit, T. C. J., Vrins, S., & von Hofsten, C. (2011). Facing threat: Infants' and adults' visual scanning of faces with neutral, happy, sad, angry, and fearful emotional expressions. *Cognition and Emotion*, 25(2), 193-205.
- Itier, R. J., Taylor, M. J. (2004). Source analysis of the N170 to faces and objects. (2004). *NeuroReport*, 15(8), 1261-1265.
- Isaac, R. M., Walker, J. M. & Williams, A. W. (1994). Group Size and the Voluntary Provision of Public Goods: Experimental Evidence Utilizing Large Groups. *Journal of Public Economics* 54(1), 1-36.
- Isaacowitz, D. M., Toner, K. & Neupert, S. D. (2009). Use of Gaze for Real-Time Mood Regulation: Effects of Age and Attentional Functioning. *Psychology and Aging*, 24(4), 989-994.
- Jackson, M. C., & Raymond, J. E. (2008). Visual working memory for faces. *Journal of Experimental Psychology: Human Perception and Performance*, 34(3), 556 -568.
- Jackson, M. C., Wolf, C., Johnston, S. J., Raymond, J. E., & Linden, D. E. J. (2008). Neural correlates of enhanced visual short-term memory for angry faces: An fMRI study. *PLoS ONE*, 3, e3536.
- Jackson, M. C., Wu, C., Linden, D. E. J., & Raymond, J. E. (2009). Enhanced Visual Short-Term Memory for Angry Faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(2), 363-374.
- Jiang, Y. V., Shim, W. M. & Makovski, T. (2008). Visual working memory for line orientations and face identities. *Perception and Psychophysics*, 70(8), 1581-1591.

- Jolij, J. & Lamme, V. A. F. (2005). Repression of unconscious information by conscious processing: Evidence from affective blindsight induced by transcranial magnetic stimulation. *PNAS*, *102*(30), 10747-10751.
- Jung, K., Ruthruff, E., Tybur, J., Gaspelin, N., & Miller, G. (2012). Perception of facial attractiveness requires some attentional resources: Implications for the “automaticity” of psychological adaptations. *Evolution and Human Behavior*, *33*, 241-250.
- Kanwisher, N., McDermott, J. & Chun, M. M. (1997). The fusiform face area: A module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience*. *17*(11), 4302-4311.
- Kirsch, P., Esslinger, C., Chen, Q., Mier, D., Lis, S., Siddhanti, S., Gruppe, H., Mattay, V. S., Gallhofer, B. & Meyer-Lindberg, A. (2005). Oxytocin modulates neural circuitry for social cognition and fear in humans. *Journal of Neuroscience* *25*, 11489 –11493.
- Kleider, H. M. & Gondinger, S.D. (2006). The Generation and Resemblance Heuristics in Face Recognition: Cooperation and Competition. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *32*(2), 259–276.
- Langeslag, S. J. E., Morgan, H. M., Jackson, M. C., Linden, D. E. J & Van Strien, J. W. (2009). Electrophysiological correlates of improved short-term memory for emotional faces. *Neuropsychologica*, *47*, 887-896.
- LeDoux, J. E. (1996). *The Emotional Brain*. New York: Simon & Shuster.
- Lorenz, K. (1971). *Studies in Animal and Human Behavior*. Cambridge, MA: Harvard University Press.
- LoPresti, M. L., Schon, K., Tricarico, M. D., Swisher, J. D., Celone, K. A. & Stern, C. E. (2008). Working memory for social cues recruits orbitofrontal cortex and amygdala: a functional magnetic resonance imaging study of delayed matching to sample for emotional expressions. *Journal of Neuroscience* *28*, 3718–3728.
- Morris, J. S., Friston, K. J., Buchel, C., Frith, C. D., Young, A. W., Calder, A. J. & Dolan, R. J. (1998). A neuromodulatory role for the human amygdala in processing emotional facial expressions. *Brain*, *121*, 47-57.
- Morris, J. S., Ohman, A. & Dolan, R. J. (1999). A subcortical pathway to the right amygdala mediating ‘unseen’ fear. *Proceedings of the National Academy of Sciences USA* *96*, 1680-1685.
- Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review*, *108*(3), 483–522. doi:10.1037//0033-295X.108.3.483

- Panchanathan, K. & Boyd, R. (2003). A tale of two defectors: the importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology*, 224(1), 115-126.
- Pessiglione, M., Seymour, B., Flandin, G., Dolan, R., & Frith, C. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, 442, 1042–1045.
- Phillips, L. H., Channon, S., Tunstall, M., Hedenstrom, A. & Lyons, K. (2008). The Role of Working Memory in Decoding Emotions. *Emotion*, 8(2), 184-191.
- Postle B. R., Druzgal T. J., D'Esposito M. (2003). Seeking the neural substrates of visual working memory storage. *Cortex* 39:927–946.
- Poundstone, W. (1992). *Prisoner's Dilemma*, Doubleday, NY NY.
- Rossion, B., & Jacques, C. (2008). Does physical interstimulus variance account for early electrophysiological face sensitive responses in the human brain? Ten lessons on the N170. *NeuroImage*, 39, 1959-1979.
- Schupp, H. T., Öhman, A., Junghöfer, M., Weike, A. I., Stockburger, J. & Hamm, A. O. (2004). The facilitated processing of threatening faces: An ERP analysis. *Emotion*, 4, 189–200.
- Schweinberger, S. R., Burton, A. M., & Kelly, S. W. (1999). Asymmetric dependencies in perceiving identity and emotion: Experiments with morphed faces. *Perception and Psychophysics*, 61, 1102-1115.
- Schweinberger, S. R., & Soukup, G. R. (1998). Asymmetric relationships among perceptions of facial identity, emotion, and facial speech. *Journal of Experimental Psychology: Human Perception and Performance*, 24, 1748-1765.
- Sergent, J., Ohta, S., Macdonald, B., & Zuck, E. (1994). Segregated processing of facial identity and emotion in the human brain: A PET study. *Visual Cognition*, 1, 349-370.
- Tamietto, M. et al. (2009). Unseen facial and bodily expressions trigger fast emotional reactions. *Proceedings of the National Academy of Sciences USA* 106, 17661-17666.
- Tamietto, M., & de Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature reviews. Neuroscience*, 11(10), 697–709. doi:10.1038/nrn2889
- Taubert, J., Apthorp, D., Aagten-Murphy, D. & Alais, D. (2011). The role of holistic processing in face perception: Evidence from the face inversion effect. *Vision Research*, 51, 1273–1278.

- Tanaka, J. W., Kiefer, M., & Bukach, C. (2004). A holistic account of the own-race effect in face recognition: Evidence from a cross-cultural study. *Cognition*, 93, B1-B9.
- Tarr, M. J. & Gauthier, I. (2000). FFA: A flexible fusiform area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, 3, 764-769.
- Thierry, G., Martin, C. D., Downing, P., Pegna, A. J. (2007). Controlling for interstimulus perceptual variance abolishes N170 face selectivity. *Nature Neuroscience*, 10, 505-511.
- Valentine, T. (1991). A unified account of the effects of distinctiveness, inversion, and race in face recognition. *Quarterly Journal of Experimental Psychology A*, 43(2) 161-204.
- Van Honk, E. J. & Schutter, D. J. L. G. (2007). Vigilant and Avoidant Responses to Angry Facial Expressions: Dominance and Submission Motives. In E. Harmon-Jones & P. Winkielman (Eds.), *Social Neuroscience*. New York: Guilford Press.
- Vogel, E. K. & Machizawa, M. G. (2004). Neural activity predicts individual differences in visual working memory capacity. *Nature*, 428, 748-751.
- Vogel, E. K., McCollough, A. W. & Machizawa, M. G. (2005). Neural measures reveal individual differences in controlling access to working memory. *Nature* 438, 500-503.
- Vuilleumier, P., George, N., Lister, V., Armony, J., & Driver, J. (2005). Effects of perceived mutual gaze and gender on face processing and recognition memory. *Visual Cognition*, 12, 85-101.
- Vuilleumier, P., & Pourtois, G. (2007). Distributed and interactive brain mechanisms during emotion face perception: Evidence from functional neuroimaging. *Neuropsychologia*, 45, 174–194.
- Whalen, P. J., Rauch, S. L., Etcoff, N. L., McInerney, S. C., Lee, M. B. & Jenike, M. A. (1998). Masked presentations of emotional facial expressions modulate amygdala activity without explicit knowledge. *Journal of Neuroscience* 18, 411-418.
- Xu, Y., & Chun, M. (2006). Dissociable neural mechanisms supporting visual short- term memory for objects. *Nature*, 440, 91–95
- Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*. 81(1), 141-145.
- Young, A. W., Newcombe, F., de Haan, E. H. F., Small, M., & Hay, D. C. (1993). Face perception after brain injury: Selective impairments affecting identity and expression. *Brain*, 116, 941-959.

CHAPTER 3.

Motivational Salience

Visual salience

Entities are said to be visually salient when they stand out from surrounding stimuli. Evolutionarily speaking, the detection of threat and reward (predators, food, etc.) may be served by these entities being visually salient. Indeed, many plants propagate via seed dissemination mechanisms that rely on their berries or fruit being eaten and as such have evolved brightly coloured fruit or berries to signal their presence to potential feeders. Conversely, certain species such as the poison dart frog (Summers & Clough, 2000) and cinnabar moth (Zoelen & Meijden, 1991) have evolved brightly coloured markings to signal to predators that they are poisonous, a phenomenon known as aposematism.

These efficient signalling mechanisms rely on the intended observer having a corresponding mechanism that allows bright colours to stand out in the visual field (i.e. be salient). Salience, therefore, is not strictly speaking a property of a physical stimulus, but is a phenomenon that arises when an observer and stimulus come together in certain specific situations (see below).¹³

Visual salience may be driven by both *bottom up* and *top down* factors. Bottom up factors relate to the physical properties of the stimulus. A classic example of this is the feature search task. In this task an observer is required to find among distractors a stimulus that has a particular feature or combination of features. In many instances of the visual search task the target stimulus will ‘pop

¹³ Technically, it is more accurate to say that the interaction of the stimulus with the perception systems of the observer is what results in salience; individual stimuli are in fact only salient with regards to any given observer and what may be salient to one observer may not be salient to another. As an example, a green stimulus among red stimuli may be visually salient when observed by an observer with colour vision that allows discrimination between these two colours but may not be visually salient to a colour-blind observer. Nonetheless, referring to stimuli as being salient or not is a useful shorthand and will be used herein under the assumption that it is generally understood that this in fact incorporates the notion of stimulus/observer interaction.

out' and be easily and quickly identified. However, manipulation of the type of distractor stimuli can impact the visual salience of the search target (see Figure 3.1).




Figure 3.1. Examples of visual search task. In instances where distractors share only a single feature with the target they can be processed in parallel and search is efficient. When multiple distractor types are present, each sharing different features with the target (conjunction search; Treisman & Gelade, 1980), search is inefficient and proceeds in a serial fashion. Graphics reproduced from http://www.scholarpedia.org/article/Visual_salience.

Saliency mechanisms are important to facilitate day-to-day commerce with the world; but how is visual saliency realised by the human mind? One answer is through saliency maps. Itti et al. (Itti, Koch & Neiber, 1998) provide an elegant description of this mechanism. Assume a stimulus as in Figure 3.2 below. Each rectangle is equally (or very nearly equally) intense in terms of its difference in luminance from the background. The high peaks in the intensity map shown in Figure 3.2 correspond to the large degree to which each bar stands out from the

background (note that the intensity map coordinates correspond to the two dimensional x and y spatial coordinates of the stimulus). However, because all the peaks are high, the intensity map is relatively uninformative when trying to distinguish differences between the rectangles. Therefore, a normalisation factor $\mathcal{N}(\cdot)$ is applied to the intensity map, down-weighting the peaks. A similar process is applied with regards to orientation, except that here the map *is* informative since one of the bars has different orientation to the others, so the contribution of the orientation dimension to saliency is upweighted. In this way, an overall saliency map corresponding to spatial locations is generated (essentially this is a compound of the normalised intensity and orientation maps shown in Figure 3.2).




Figure 3.2. Bottom up saliency maps for a simple display of white rectangle stimuli. Reproduced from Itti et al. (1998).

Visual salience is also driven by top down factors. The most obvious of these is conscious search; if one is searching for a particular entity (say, one's own car in a car park) then this entity may pop out from the crowd. Additionally, it is possible for less complex features to take on salience when they are situationally relevant. If I own a red car and need to find it in a car park, then the colour red may serve as an initial guiding factor to allow me to limit processing only to a consideration set of red cars. The way in which top-down factors guide visual search in this manner has been extensively studied and termed 'guided search' (Cave & Wolfe, 1990; Wolfe, 1994; Wolfe, 2007; Wolfe, Cave, & Franzel, 1989; Wolfe & Gancarz, 1996).

Each of these top down factors can be described as being motivational in nature inasmuch as each relates to a particular goal an individual has at a particular time. However, not all motivational states need be consciously held, and factors such as hunger, thirst, tiredness etc. may each play a role in determining which stimuli or stimulus features may be salient at any given time.

Motivational salience

The visual search literature is replete with examples of the effect of both bottom up and top down factors on salience and the subsequent deployment of attention.¹⁴ Indeed, the deployment of attention to a stimulus (e.g. as indexed by saccades; Pomplun, 2006) is often used to determine salience in visual search tasks. Thus, salience often denotes that a stimulus is worthy of selective attention.

¹⁴ There is an ongoing debate as to the extent to which bottom-up and top-down factors separately influence the immediate saliency of stimuli and the associated initial feed-forward sweep of visual processing and deployment of visual attention. The two poles of this argument are represented by a review from Theeuwes (2010) in which he concludes that 'top-down knowledge regarding non-spatial features of the objects cannot alter the initial selection priority', and the contingent-capture hypothesis (Folk, Remington, & Johnston, 1992), which argues that bottom-up signals must match to a top-down goal state if they are to capture visual attention.

Selective attention, however, is not a processing resource in its own right; rather it is a guiding mechanism that serves to allocate neural processing resources towards particular entities. More generally then, we can say that saliency may denote a stimulus as worthy of more processing. This is an important distinction since it also allows for the level of processing at post-selective/encoding levels (i.e. once an entity has been encoded in memory) to possibly be biased by saliency.

In the visual search example depicted in Figure 3.1 above, the colour red is salient due to its uniqueness in the display. In the natural world, bright red fruit and berries are salient both because they tend to stand out from distractors (e.g. green leaves with which they share few features), and because they have *intrinsic* value (they are food). It is also possible that features with no intrinsic value may carry learned value codes and so also result in saliency. For example, it is possible to imbue a particular colour with a positive or negative value by repeatedly pairing it with either reward or punishment (see below). Such a colour then carries associated motivational value, which may bias saliency even in situations where the colour is no longer task relevant. In Chapter 7 I imbue value into face stimuli and test whether the resulting motivational saliency impacts WM for such faces.

Value is thus a higher order dimension that may bias saliency through a variety of top down processes. Once stimulus processing has occurred, and the fact that a particular stimulus has value has been established, then value may serve to alter the saliency map such that a greater saliency is realised at the location of the value laden stimulus. This could be achieved through the construction of a value map similar to those of orientation and intensity which also

feeds in to the overall saliency map. Alternatively, value may serve to bias the weights assigned to particular features. For example, a red stimulus and blue stimulus presented together with a crowd of green stimuli will both stand out from the crowd, but if the red colour has previously been paired with (say) monetary reward then the reward association history may serve to decrease the threshold at which neurons representing the colour red (in early visual areas) fire. It is not the purpose of this thesis to distinguish between these two mechanisms, and indeed both may play a role.

In the same way that visual salience is affected by both bottom-up and top-down factors, motivational salience is also influenced by top-down goals (see below) and by bottom-up factors such as the automatic motivational saliency accorded to items of evolutionary threat (snakes, spiders; Öhman, Flykt, & Esteves, 2001; Öhman & Mineka, 2001) and items of modern threat (knives, syringes; Brown, El-Dereby & Blanchette, 2010). Stimuli such as snakes, spiders, guns and syringes can be thought of as carrying threat information within the context of an omnipresent goal; survival. As such, they may be expected to always carry such importance. However, the overall goal to survive is subserved by a number of sub-goals; sustenance, shelter, reproduction, rest and exploration¹⁵ (see Table 3.1).

¹⁵ Surprisingly, exploration is rarely thought of as being a low-level goal, more traditionally being discussed as a higher level goal (see for example Maslow's hierarchy of needs; Maslow, 1943). This is probably because strictly speaking it does not directly and immediately subserve survival. If an organism lives within a group in which food, shelter and reproductive partners are readily available then there may be no need for this individual to explore beyond the boundaries of the group in order to successfully pass on genetic material. However, if we consider low level 'second order' goals as being driven by an evolutionary pressure that works over generations rather than merely during the lifetime of a single individual, then since the long term survival of a group depends upon its eventual expansion to acquire new resources then the absence of exploration as a hard-wired goal would have negative consequences for survival. Note that the above scenario of territorial expansion is only one extreme form of exploration. More generally we can

Survival/reproduction related sub-goal	Motivating factors	Satiating stimuli	Demotivating factors
Sustenance	Hunger Thirst	Food Water	Satiation
Shelter	Cold Heat Presence of predators	Shelter	Comfort
Reproduction	Potential sexual partner	Sexual partner	Presence of sexual competitor
Rest	Tiredness	Sleep	Absence of safe resting place
Exploration	Absence of reward in current situation	'Discovery'	Reward already obtainable in current situation Perceived risk / cost of exploration

Table 3.1. Low-level goals, motivating factors, satiating stimuli, and demotivating factors.

There is, of course, an opportunity cost of pursuing any particular goal since only a limited number of goals may be pursued at one time (we cannot seek out food while asleep). Opportunity costs act in a similar way to directly demotivating factors in controlling an organism's propensity to pursue any given goal. It is evident from Table 3.1 that low-level goals may be broadly divided into two categories; threat negating and reward seeking (with shelter perhaps being a compound goal comprising elements of both). This distinction is important since different neural substrates underlie the acquisition of positive and negative value codes for motivational stimuli (see below).

define exploration as expenditure of energy to create or move into a new situation in which potential rewards may be obtainable (see Table 3.1).

So far I have discussed low-level goals common to higher organisms. Humans, however, are complex social animals and in addition to homeostatic requirements are driven to pursue a number of 'higher' goals to fulfil complex social motivations (e.g. Maslow, 1943). More specifically, it is of note that each 'need' can be satisfied by the setting and fulfilment of certain goals, as illustrated in Table 3.1 above. In this way it is possible to conceptualise almost all human behaviour as resulting from a process of goal setting and fulfilment. This is not to say that goals are necessarily usually *consciously* set. Indeed, conscious development of a goal-plan need only occur when a series of ordered steps are required: If I am hungry and I have on my person a bar of chocolate then I can consume it without deliberation, but if I do not possess food then I must formulate multi-step plan to acquire some (e.g. find money, find a shop, select food, purchase).

It is apparent from the discussion of setting goals to fulfil low level needs above that motivational salience is also affected by top-down factors. Thus, motivational salience may be regarded as a property of a stimulus (or more formally a stimulus and observer interaction) that derives from the value that that stimulus has to the observer at any particular time. Value is a complex construct, representing as it does a compound of various stimulus attributes such as valence, predictiveness of learned outcome related to stimulus, delay to outcome predicted, goal congruence, intensity of outcome predicted, action trade-off status (energy that must be expended to utilise stimulus), and rarity of stimulus (probability of encountering later if not approached/utilised on this occasion), the net result of which is modulated by current goals to derive final stimulus value or utility to the observer.

Given that WM is a crucial part of the response planning apparatus, it would be beneficial if motivationally salient items had preferential access to, were efficiently manipulated in, and were easily recovered from WM. In Chapter 2 I noted that emotional and in particular angry faces have been observed to experience a boost in WM relative to non-emotional (neutral) faces. It is possible that such a boost derives from their status as motivationally salient items rather than their threat status *per se*. To investigate the possibility that motivational salience might effect a boost in WM processing, in Chapter 7 I imbue value into expressively neutral face stimuli and test whether this affects WM processing for them. Imbuing of stimuli with value can be achieved using a value-learning paradigm. In such a procedure, different stimuli are consistently paired with either a negative, neutral, or positive outcome, usually in the form of a game in which points or small amounts of money are won and lost. In-procedure response bias analysis can be used to confirm that learning of stimulus values and reward/punishment prediction has taken place, and post-procedure stimulus evaluation tests can be used to ascertain whether this learning translates to an emotional bias across stimuli (whether, for example, stimuli paired with highly positive outcomes are then rated as more trustworthy than those paired with negative outcomes). The mechanisms underlying the acquisition of positive and negative value codes seem to be based on different neurological structures, with the amygdala being associated with acquisition of negative value codes (Kahn, Yeshurun, Rotshtein, Fried, Ben-Bashat & Hendler, 2002), and the ventral striatum and right OFC more with positive value codes (Yacubian, Gläscher, Schroeder, Sommer, Braus, & Büchel, 2006), although it is possible that these systems may in fact share some neural resources, with the amygdala playing a role in

assignment of both positive and negatively valenced codes to objects (Paton, Belova, Morrison & Salzman, 2006).

In Experiment 8 I use an operant conditioning paradigm that has been reported before (O'Brien & Raymond, 2009) to imbue value associations directly to specific face stimuli. In Experiments 6 and 7 I use a novel value learning paradigm to imbue value associations into particular classes of stimuli as defined by different colours. The notion of using colour to define value association has been successfully utilised elsewhere.¹⁶ For example, Anderson et al. (Anderson, Laurent & Yantis 2011a; 2011b) successfully imbued high and low positive value into colour categories (using an associative learning task in which colour was orthogonal to the task)¹⁷ and then, in a separate test phase, used stimuli of different value-associated colours as either distractors (2011a) or salient distractors (2011b), finding attentional capture by value associated items in both instances.

Motivationally salient stimuli in visual WM

There is an inherent advantage in having lower level visual detection and representation mechanisms that are able to quickly prioritise certain types of information for efficient evaluation and appropriate response generation. In daily life, we encounter numerous stimuli that require near instantaneous evaluation if we are to generate the optimal response quickly enough to take advantage of, or avoid a situation. One mechanism that facilitates such differential stimulus processing is attention. Attention mechanisms are differentially deployed under

¹⁶ Indeed, the ability to imbue learnt status into colour is prevalent across species. For example, pigeons can be trained to associate pecking a coloured key with reward (e.g. Lea, 1979).

¹⁷ See Chapter 7 for discussion of their method. See also Hickey et al. (Hickey, Chelazzi & Theeuwes, 2010) for an alternative colour-based value learning paradigm.

conditions of high and low arousal (e.g. Shapiro & Lim, 1989), such that in conditions of high arousal attention deployment is biased toward auditory attention over visual attention and also biased within visual attention to the periphery of the visual field. Such differential deployment of attention may serve an evolutionary purpose, increasing awareness of surroundings in situations of high arousal (danger). The degree to which particular stimuli generate positive or negative affect depends in part upon the way in which they fit current and future goals. When hungry, food may generate positive affect, but when hunger is satisfied food may become an emotionally neutral stimulus (Brendl, Marknan, & Messner, 2003). Additionally, stimuli perceived as neither instrumental nor disinstrumental to achieving a current focal goal may be subject to an unconscious 'devaluation' effect such that they are evaluated more negatively, with this devaluation thought to be due to a mechanism whereby evaluative responses are blocked rather than inhibition of attention to stimuli (Brendl et al, 2003).

Additional to the role that mood may play in directing attention, emotional states associated with specific stimuli may also serve to direct attention. In particular, negative stimuli better capture attention than positive stimuli (Pratto & John, 1991), make disengaging attention more difficult (Fox, Russo, & Dutton, 2002), and generate greater interference on detection tasks (Pessoa, Mckenna, Gutierrez, & Ungerleider, 2002).

The relationship between attention and motivational salience is bi-directional, such that attentional state can directly influence emotional evaluation of a stimulus. In particular, under conditions where a stimulus must be ignored, and an inhibitory 'tag' attached to the stimulus to facilitate task demands (non-response to the stimulus), the inhibited stimulus may be subject to a devaluation

effect so that it is rated less emotionally favourably (Raymond, Fenske, and Tavassoli, 2003). Such 'inhibitory tags' are recoverable at a later date (one possible mechanism of reinstating do-not-respond locations in interrupted visual search; Tipper, Grison, and Kessler, 2003). Devaluation effects can also be reinstated in this way and the original formation of such devaluation associated tags requires working memory (Goolsby, Shapiro & Raymond, 2009).

Recent research (Raymond & O'Brien, 2009) using a value-learning task followed by attentional blink task has demonstrated survival of memory traces for positively valenced but not negatively valenced stimuli under conditions of low attention availability at recall. This finding, along with the fact that emotional devaluation states associated with objects require working memory for their formation (Goolsby et al., 2009), suggests that a relationship may exist between emotional coding of stimuli and the way they are processed in WM. Indeed, preferential treatment in WM space of emotionally laden stimuli might have evolutionary advantages, allowing quicker or more accurate responses to such stimuli.

Both positive and negative value codes are therefore important in controlling responses to motivationally salient stimuli, but the appropriate responses to positive and negative stimuli are very different. Exploiting the presence of positive value stimuli most usually involves some form of 'approach' or 'acquire' response whereas negative stimuli may elicit either an 'approach' response (to deal with or nullify negativity) or an 'avoid' response (evade potential negative outcome), or an 'evaluation' response, to determine whether or not it is best to try to deal with or evade the negative stimulus. There is currently some debate as to whether negative stimuli elicit approach or avoid responses and how

automatic these responses are. Interestingly, knowledge of the potential presence of threatening stimuli is in itself enough to alter deployment of attention to monitor visual space for occurrence of such stimuli, even when they are task irrelevant (Devue, Belopolsky & Theeuwes, 2011).

WM and the estimation of time

Converging evidence for the importance of salience in determining deployment of WM resources comes from research into temporal perception. In rats and pigeons there is evidence that WM is a crucial component in evaluating the passage of time. A similar view governs the perception of time in humans (Wearden & McShane, 1998, Wearden, 2003). Furthermore, there is evidence in rats (Buhusi, 2012) and in pigeons (Buhusi, Sasaki & Meck, 2002) that the salience (illumination level in pigeons, sound level in rats) of distractors presented during a peak-interval (PI) procedure (Catania, 1970) proportionally affects the ability to judge the passage of time. It appears that the level of salience of a distractor affects the amount of WM resource it co-opts away from WM, resulting in a greater passive decay of WM information due to an associated lack of WM resource left over to maintain the contents of WM (Buhusi & Meck, 2006) when more salient distractors are presented. Interestingly, this effect has also been observed for affectively salient distractors (white noise conditioned to be affectively aversive via pairing with a foot shock; Matthews, He, Buhusi & Buhusi, 2012); when affective value of the distractor was manipulated by administration of nomifensine (a norepinephrine and dopamine reuptake inhibitor with anti-depressant properties which has been shown to facilitate down-weighting of fear responses) affective distractors no longer co-opted WM resource to the same

degree. It would thus appear that, at least in rats, the level of affective salience of a distractor determines its ability to co-opt WM resources to itself.

Similar research exists concerning the effect of emotion on the perception of time in humans. In humans arousing stimuli, including both happy and angry faces, have been shown to lead to an increase in subjective time estimation when viewed during the period to be estimated, presumably by causing arousal and thus increasing the firing rate of an internal pacemaker. Conversely, viewing a face showing an expression of shame has been shown to lead to an underestimation of the passage of time, presumably due to its causing of introspection, an emotion that engages WM resources (Gil & Droit-Volet, 2011). Thus, it appears that temporal estimation is influenced by two processes, the firing rate of an internal clock or pacemaker which increases in situations of arousal and the level of WM resources available to store and maintain the count of the 'clicks' of this internal clock, a factor influenced by concurrent demands on WM (see Figure 3.3).

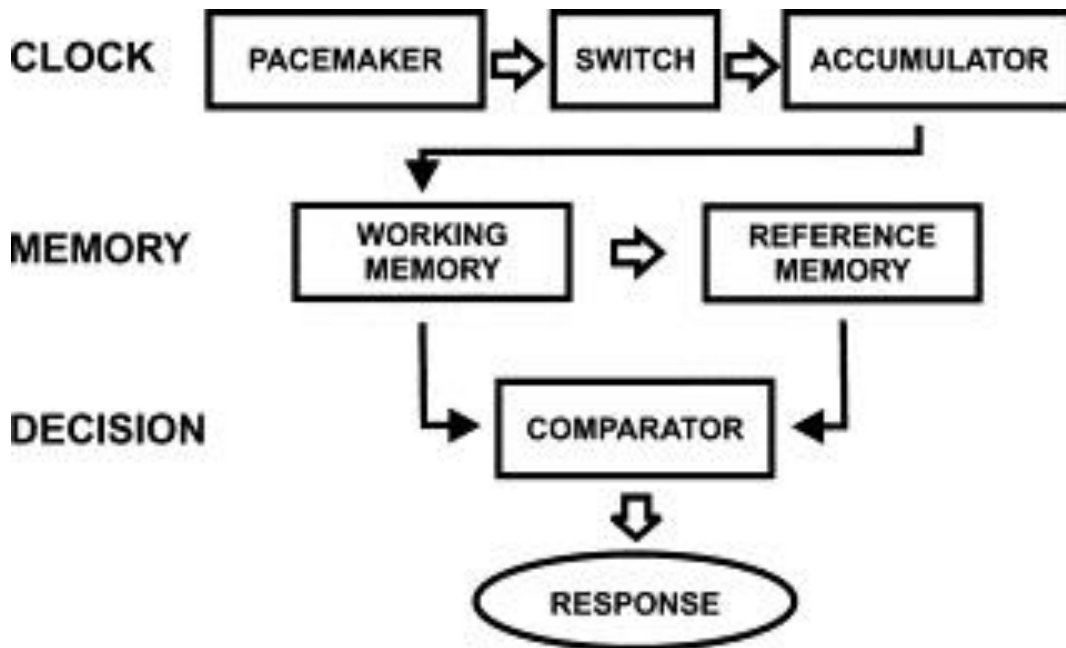


Figure 3.3. The temporal information-processing model. Upon start of a timing interval, the switch closes, allowing pulses (ticks of the pacemaker) to accumulate. These pulses are stored in WM, and compared with a value retrieved from reference memory (LTM) appropriate to the pre-specified target interval being estimated (reference memory representations of interval length having been built up via prior functioning of the clock during a training phase in experimental settings or via experience *in vivo*). Reproduced from Yarrow, Haggard & Rothwell (2004).

This research on time estimation provides converging evidence that the salience, including emotional learnt salience, of stimuli may play a crucial role in determining the level of WM resources that they attract. Additionally, the notion that WM resource is required to monitor time has bearing on the results observed for several experiments herein (see Chapters 6 and 9).

References – Chapter 3

- Anderson, B.A., Laurent, P.A., & Yantis, S. (2011). Learned value magnifies salience-based attentional capture. *PLoS ONE* 6(11): e27926. doi:10.1371/journal.pone.0027926
- Anderson, B.A., Laurent, P.A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences USA*, 108, 10367-10371.
- Brendl, C.M., Markman, A.B., & Messner, C. (2003). The devaluation effect: Activating a need devalues unrelated objects. *Journal of Consumer Research*, 29, 463-473.
- Brown, C., El-Deredy, W., & Blanchette, I. (2010). Attentional modulation of visual-evoked potentials by threat: investigating the effect of evolutionary relevance. *Brain and cognition*, 74(3), 281-287. doi:10.1016/j.bandc.2010.08.008
- Buhusi, C. V. (2012). Time-sharing in rats: effect of distracter intensity and discriminability. *Journal of Experimental Psychology: Animal Behaviour Processes*, 38, 30-39.
- Buhusi, C.V. & Meck, W.H. (2006). Interval Timing With Gaps and Distracters: Evaluation of the Ambiguity, Switch, and Time-Sharing Hypotheses. *Journal of Experimental Psychology: Animal Behavior Processes*, 32(3), 329-338.
- Buhusi, C.V., Sasaki, A. & Meck, W.H. (2002). Temporal integration as a function of signal and gap intensity in rats (*Rattus norvegicus*) and pigeons (*Columba livia*). *Journal of Comparative Psychology*, 116(4), 381-390.
- Catania, A.C. (1970). Reinforcement schedules and psychophysical judgements: A study of some temporal properties of behavior. In Schoenfeld, W.N. (Ed.), *The theory of reinforcement schedules* (1-42). New York: Appleton-Century-Crofts.
- Cave, K. R., & Wolfe, J. M. (1990). Modeling the role of parallel processing in visual search. *Cognitive Psychology*, 22, 225-271.
- Devue, C., Belopolsky, A. V., & Theeuwes, J. (2011). The role of fear and expectancies in capture of covert attention by spiders. *Emotion*, 11(4), 768-775. doi:10.1037/a0023418

- Folk, C. L., Remington, R. W., & Johnston, J. C. (1992). Involuntary covert orienting is contingent on attentional control settings. *Journal of Experimental Psychology: Human Perception and Performance*, 18, 1030-1044.
- Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition and Emotion*, 16(3), 355-379.
- Gil, S., & Droit-Volet, S. (2011). How do emotional facial expressions influence our perception of time? In S. Masmoudi, D. Yan Dai & A. Naceur (Eds). *Attention, Representation, and Human Performance: Integration of Cognition, Emotion and Motivation* (1-12) London: Psychology Press, Taylor & Francis.
- Goolsby, B.A., Shapiro, K.L., & Raymond, J.E. (2009). Distractor devaluation requires visual working memory. *Psychonomic Bulletin & Review*, 16(1), 133-138.
- Goolsby, B.A., Shapiro, K.L., Silvert, L., Kiss, M., Fragopanagos, N., Taylor, J.G., Eimer, M., Nobre, A.C., & Raymond, J.E. (2009). Feature-based inhibition underlies the affective consequences of attention. *Visual Cognition*, 17, 500-530. doi:10.1080/13506280801904095
- Hickey, C., Chelazzi, L. & Theeuwes, J. (2010). Reward Changes Salience in Human Vision via the Anterior Cingulate. *The Journal of Neuroscience*, 30(33) 11096-11103.
- Itti, L., Koch, C. & Niebur, E. (1998). A Model of Saliency-Based Visual Attention for Rapid Scene Analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1254-1259.
- Kahn, I., Yeshurun, Y., Rotshtein, P., Fried, I., Ben-Bashat, D., & Hendler, T. (2002). The role of the amygdala in signalling prospective outcome of choice. *Neuron*, 33, 983–994.
- Laurent Itti (2007). Visual salience. Scholarpedia, 2(9):3327., revision #72776
- Lea, S. E. G. (1979). Foraging and reinforcement schedules in the pigeon: Optimal and non-optimal aspects of choice. *Animal Behaviour*, 27(3), 875-886.
- Maslow, A. H. (1943). A Theory of Human Motivation. *Psychological Review* 50(4), 370-396.
- Matthews, A.R., He, O.H., Buhusi, M. & Buhusi, C.V. (2012). Dissociation of the role of the prelimbic cortex in interval timing and resource allocation: beneficial effect of norepinephrine and dopamine reuptake inhibitor

- nomifensine on anxiety-inducing distraction. *Frontiers in Integrative Neuroscience* 6, 1-12.
- Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion drives attention: Detecting the snake in the grass. *Journal of Experimental Psychology: General*, 130, 466-478.
- Öhman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: Toward an evolved module of fear and fear learning. *Psychological Review*, 108, 483-522.
- Paton, J.J., Belova, M.A., Morrison, S.E., & Salzman, C.D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature*, 439, 865–870.
- Pessoa, L., Mckenna, M., Gutierrez, E. & Ungerleider, L.G. (2002). Neural processing of emotional faces requires attention. *PNAS*, 99, 11458-11463.
- Pomplun, M. (2006). Saccadic selectivity in complex visual search displays. *Vision Research*, 46, 1886-1900.
- Pratto, F. & John, O.P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61(3), 380-391.
- Raymond, J. E., Fenske, M. J., & Tavassoli, N.T. (2003). Selective attention determines emotional responses to novel visual stimuli. *Psychological Science*, 14, 537-542.
- Raymond, J. E. & O'Brien, J. L. (2009). Selective Visual Attention and Motivation: The Consequences of Value Learning in an Attentional Blink Task. *Psychological Science*, 20(8), 981-988.
- Shapiro, K. L. & Lim, A. (1989). The impact of anxiety on visual attention to central and peripheral events. *Behaviour Research and Therapy*, 27, 345-351.
- Summers, K. & Clough, M. (2000). The evolution of coloration and toxicity in the poison frog family (Dendrobatidae). *Proceedings of the National Academy of Sciences of the United States of America* 98(11), 6227–6232. doi:10.1073/pnas.101134898
- Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica*, 135, 77–99.
- Tipper, S.P., Grison, S, & Kessler, K. (2003). Long-term inhibition of return of attention. *Psychological Science*, 14, 19-25.

- Treisman, A. & Gelade, G. (1980). A feature integration theory of attention. *Cognitive Psychology* 12, 97-136.
- van Zoelen, A. M. & van der Meijden, E. (1991). Alkaloid concentration of different developmental stages of the cinnabar moth (*Tyria jacobaeae*). *Entomologia Experimentalis et Applicata*, 61(3), 291–294.
- Wearden, J. H., & McShane, B. (1988). Interval production as an analogue of the peak procedure: Evidence for similarity of human and animal timing processes. *Quarterly Journal of Experimental Psychology*, 40B, 363-375.
- Wearden, J. H. (2003). Applying the scalar timing model to human time psychology: Progress and challenges. In H. Helfrich (Ed.), *Time and mind II* (21-39). Göttingen: Hogrefe & Huber.
- Wolfe, J. M. (1994) Guided Search 2.0: A Revised Model of Visual Search. *Psychonomic Bulletin & Review*, 1(2), 202-238.
- Wolfe, J. M. (2007). Guided Search 4.0: Current Progress with a model of visual search. In W. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 99-119). New York: Oxford.
- Wolfe, J. M., Cave, K. R., & Franzel, S. L. (1989). Guided Search: An alternative to the Feature Integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 419-433.
- Wolfe, J. M. & Gancarz, G. (1996). Guided Search 3.0: A model of visual search catches up with Jay Enoch 40 years later. In V. Lakshminarayanan (Ed.), *Basic and clinical applications of vision science* (pp. 189-192). Dordrecht, Netherlands: Kluwer Academic.
- Yacubian, J., Gläscher, J., Schroeder, K., Sommer, T., Braus, D.F., & Büchel, C. (2006). Dissociable systems for gain- and loss-related value predictions and errors of prediction in the human brain. *Journal of Neuroscience*, 26, 9530-9537.
- Yarrow, K., Haggard, P. & Rothwell, J.C. (2004). Action, arousal, and subjective time. *Consciousness and Cognition*, 13(2), 373-390.

CHAPTER 4.**General Methods**

In this chapter I discuss the statistical techniques and measurements used throughout this thesis, including their strengths and weaknesses. Recent reviews of published psychology papers illustrate a surprising incidence of misuse of statistics, including failure to report effect sizes (Schatz, Jay, McComb & McLaughlin, 2005; Scheff, 2011); increased chance of type 1 error, inappropriate use of p values, and inappropriate use of null hypothesis testing (Schatz et al., 2005); non-independent correlation between brain (grouped fMRI voxel) activation and personality measures (Vul, Harris, Winkielman & Pashler, 2009); and inclusion of inappropriate control variables in multiple regression (Spector & Brannick, 2011) to name but a few. While it is not the primary purpose of this thesis to address purely mathematical issues, the surprising incidence of such errors provides a cautionary lesson for researchers in all areas of psychology, and recommends a 'less haste, more speed' approach.

My statistical analyses herein typically proceed along 'classical' inferential lines (ANOVAs, t -tests, simple linear regression and correlation between groups / conditions of interest etc.). These are well known techniques and as such do not need to be elaborated further here. However, one issue does deserve further elaboration at this stage; use of summary statistics in the representation of WM performance. While it may seem unusual to devote a chapter to what are, after all, relatively widely used statistics in this field, I believe it is of the utmost importance to demonstrate the strengths and limitations of these summary measures before using them to present evidence of psychological processes. Many of the limitations of these measurements are rarely acknowledged, yet alone discussed in the psychological reports that employ them.

Change detection

One of the most common ways of investigating short term memory is the change detection task. In this paradigm an array of to-be-remembered items is presented briefly followed by a short interval and then a probe comprising either a single test item (to which the participant responds 'present' or 'absent' depending on whether it had been present in the memory array) or a test array (to which the participant responds 'change' or 'no-change' depending on whether the array is the same or different to the memory array).

There are various ways in which performance on change detection memory tasks can be portrayed. The most simple is to report the proportion of correct decisions made for each condition in the task. While this is an intuitive measure and is easy to calculate, it is hampered by the fact that it does not reflect sensitivity to change alone but also incorporates participants' response biases. Suppose a participant in a change detection task chooses to answer "change" to every item. This strategy results in 100% correct across change trials. However, it also results in 0% correct on the no-change trials. This is problematic for the researcher because if the no-change condition is not of interest (that is, if the researcher is primarily interested in how accurately the participant is able to detect change when it occurs) and trials for the no-change condition are therefore not analysed, then an apparent 100% correct change detection score is arrived at for the participant. Does this really mean that that the participant was able to discriminate change with 100% accuracy on these trials? Clearly it does not; it is a function of the participant's response bias (tendency to respond "change" vs. "no-change" when unsure) rather than ability to detect change that results in this

score. Indeed, the participant need not have attended to the task at all in order to achieve such a performance level.

The ability to accurately detect change cannot therefore be derived from performance on change trials alone. Performance must instead be interpreted with respect to the participant's response bias. This response bias can be derived from responses to no-change trials. Likewise, the reverse is true; performance on no-change trials can only be interpreted in the context of response bias as measured by performance on change trials. Thus, any summary statistic used as a descriptor of performance in change (or signal) detection tasks must account for response biases by incorporating performance across both change and no-change trials. The various capacity and signal detection measures described below each do this in slightly different ways.

K measures (Pashler's K, Cowan's K, K_{\max} , K_{average} , K_{it} and Bangor K)

An alternative to percent correct is to try to establish *how many item representations* are held in WM. The construct K (sometimes reported as lower case 'k') was first developed by Pashler (1988) as an estimate of the number of items held in WM in a change detection task, and took the form:

$$K = \frac{N(H-g)}{g-1} \quad (4.1)$$

Where N is the number of items in the presentation array (termed the set size), H is the hit rate, and g is the ‘guessing rate’ or the respondent’s propensity to answer ‘change’ on trials in which he or she did not successfully memorise the memory array item probed (i.e. in absence of information at test). However, Pashler’s (1988) formula does not adequately account for guessing in single probe WM designs (see Appendix E for detailed discussion). In particular, it makes the assumption that $g = FA$ (the false alarm rate) and so fails to factor in that performance on no-change trials is aided by memorisation of items. In his seminal review of visual WM capacity, Cowan (2001) reformulated the equation to provide an estimate of K that is now frequently used to determine the number of discrete items that can be recorded into visual WM simultaneously when single probes are presented at test. Rendered in its most simple form, Cowan’s formula reads:

$$K = I * [H - FA]^{18} \quad (4.2)$$

Where I is the number of items presented in the memory array (the ‘load’ or ‘set size’), H is the proportion of absent trials in which the stimulus presented at probe is correctly identified as not having been present in the memorisation array, and FA is the proportion of present trials in which the participant erroneously indicates the stimulus as having been absent from the original array.¹⁹

¹⁸ Note that set size I is sometimes denoted ‘ N ’ or ‘ S ’ and false alarm rates are sometimes denoted ‘ F ’ in this formula.

¹⁹ Sometimes the Hit and FA rates are defined in the reverse; Hits being the correct identification that the item presented at test had been present in the memorisation array on ‘present’ trials and FAs being failures to report change on ‘absent’ trials. This is, perhaps, a more ‘traditional’ definition, coming from the

When a participant takes part in an experiment in which there are several different levels of load (e.g. two objects to be remembered in half of trials and four objects to be remembered in the other half), it is common to extrapolate from the values of K obtained on each load a single K value for each participant as a summary of their WM capacity. There are several ways in which this can be done: K_{\max} is the K value obtained for the load that has the highest K value. As such, is not compromised by levels of (low) load at which participants' performance is 'at ceiling'. However, as K_{\max} is based on a single load and thus only a subset of the total data, it may overestimate the level of capacity (since the more loads that are run, the greater the statistical probability of obtaining by chance on one or more load levels a value that exceeds the true level of capacity for any given participant). K_{average} is the mean value of K across all loads. However, as a load of one allows a maximum K value of 1, and a load of 2 allows a maximum K value of 2 etc., K_{average} may result in an underestimation of capacity. A solution to the above may be found in the measure $K_{\text{iterative}}$ (K_{it} ; Jackson, Wu, Linden & Raymond, 2009), which is calculated as follows: Cowan's K is computed for a participant at each load condition and then for each participant the mean K value across all levels is calculated. Next, if this mean value is lower than or equal to the value of K at the lowest load, then the mean is taken to be the value (for this participant) of K_{it} . If the mean is higher than the K value at the lowest load, then the value of K at the lowest load is discarded, and a new average calculated based only on the remaining loads. This new average is then compared to the K value of the lowest load used in its calculation, and the process repeated iteratively until either the

background Hit rates have in signal detection theory (see below). However, as memory based change detection tasks typically ask participants to look for and report *changes* in signal rather than *presence of* signal it seems more correct to define Hit and FA rates in the way I have done here.

calculated average (mean) is lower than K for the lowest load used in the calculation, or all loads are exceeded (in which case the value of K_{it} is taken to be the value of K at the highest load). This gives a K_{it} value for each individual participant, and these values can then be averaged across experimental groups or conditions.

In certain circumstances it can be useful to establish the load K at which accuracy falls below a certain level. An example of this approach, termed *Bangor K*, in which K values are interpolated from least squares lines fit to performance data at each load has been reported by Jackson and Raymond (2008). However, while this approach utilises data points from multiple loads in its calculation and so maximises data use, it also has certain constraints, requiring multiple set sizes with asymptotic performance at higher loads.

While measures of K can give an estimation of the number of item representations held in working memory, it is also possible to conceptualise differences in WM performance in terms of the way in which the representations held in WM may be utilised. An argument for this approach is that measures of K assume that WM representations can be perfectly compared to test items, that is, if a test stimulus is presented and this matches to a representation held in WM, it is assumed that the participant will then report that the test item matches the held representation. There are several reasons why this may not always be so. First, matching of the perceptual representation of the test item to the WM representation is unlikely to be an error free process and the amount of error involved may vary across situations. Second, the WM representation may not be stored with sufficient resolution to effect a match to a particular stimulus. For example, suppose a participant holds in WM a representation of a face but this

representation is held with imperfect precision. When presented with a test item of another face the participant may be unable to report with confidence that the test item does not match the memory item. However, when presented with an unambiguously distinct test item (say, a car rather than a face), such a 'non-match' discrimination might be made with confidence. In defining WM capacity in terms of the number of items, measures of K are therefore suggestive of a purely slots-based model of WM, a contention that is neither universally accepted in the literature (see Chapter 1) nor supported by the empirical findings of this thesis.

A further theoretical difficulty with conceptualising WM change detection task data in terms of capacity is that it is possible, under certain circumstances, for a participant to achieve a negative K score. From a capacity point of view this makes no sense; it is not possible to hold a negative number of items in WM. While such an occurrence would normally indicate something untoward, such as a participant misinterpreting task instructions, it is also possible that in a condition of extreme task difficulty (where performance is likely to be close to chance) natural chance variability in scores across participants would result in some participants showing negative K scores.

These theoretical and practical issues with the various K capacity measures lead us to require an alternative measure that reflects performance on a WM change detection task. An appropriate measure of performance will be one that reflects sensitivity to (ability to detect) change without being confounded by response biases. Conceptually, such a measure reflects sensitivity to change rather than the number of items stored and so circumvents the issues with K discussed above since it assumes neither error free matching, nor slot-based storage, and has interpretable meaning as a negative number. One such

measure that is commonly reported in the literature and is used throughout this thesis is *d prime* (d') which derives from signal detection theory (SDT).

Signal Detection Theory and the d' statistic

In 1932 Bartlett published his classic work on the nature of memory. A crucial part of Bartlett's theory was to suggest that memory processes are inherently reconstructive rather than reproductive, and that therefore remembering was an inherently imprecise process:

“Remembering is not the re-excitation of innumerable fixed, lifeless and fragmentary traces. It is an imaginative reconstruction or construction, built out of the relation of our attitude towards a whole active mass of organised past reactions or experience, and to a little outstanding detail which commonly appears in image or in language form. It is thus hardly ever really exact...”

In recognising the limitation of accurate recall Bartlett pre-empted later models of both LTM and WM that acknowledge that accuracy of recall (or recognition) is constrained by a certain amount of *noise* in memory representations. That memory representations (based as they are on neural activity) are inherently noisy underlies the application of SDT to studies of human memory.

Signal Detection Theory

SDT was originally developed as a military application for the detection of radar signals (Marcum, 1947), and was developed and popularised in psychology by Green & Swets (1966) as a method of evaluating ability to detect sensory stimuli. Crucial to SDT is the notion that all sensory signals occur in the presence of noise. Such noise has a number of internal and external sources. The classic example of SDT in practice is the detection of a flash of light in an otherwise dark environment. In this example, the light is a signal stimulus, and the darkness absence of signal. However, even in absence of signal a certain level of neural activity of equivalent form to that which a signal would elicit will be present (e.g. from chance spontaneous increases in the graded potentials of rod cells in the retina). Thus, the absence of signal condition may be referred to as the 'noise distribution' and the presence of the signal as the 'signal distribution' or the 'signal plus noise distribution' (since uncertainty as to presence or absence of signal exists in this case in much the same way as it does when the signal is absent). Figure 4.1 below shows these two distributions. In this classic example, external noise results from factors including the statistical distribution of photons that actually reach the retina and result in photopigment isomerisations when viewing a stimulus. Internal noise reflects the inherently noisy neuronal functions that underlie memory consolidation, maintenance and retrieval processes. When there is a very large difference in signal between the presence and absence of a stimulus (i.e. when what is perceived in presence of the stimulus is very different to the 'noise only' or 'background' sensory experience) then presence of the stimulus can be accurately ascertained. Likewise, when the level of 'background sensory noise' is very small in comparison to the change in sensory experience elicited by the stimulus this also makes the presence/absence judgement less

prone to error. Thus, the presence of noise in both external and internal signals results in a certain level of *uncertainty* experienced in perception of a stimulus. In Figure 4.1 this uncertainty is reflected by the overlap between the two distributions; when the level of sensory experience (abscissa value) is in this region the subject's experience of the stimulus is uncertain and can be interpreted as either absent or present depending on the exact level of sensory experience and the response bias (see below).




Figure 4.1. Determination of sensitivity (d'), bias (c), and likelihood ratio (β) from noise and signal plus noise distributions. Reproduced from Stanislaw & Todorov (1999).

It is readily apparent from Figure 4.1 that both the shape and separation of the signal and signal plus noise distributions determine how much overlap exists between them. Narrower distributions and/or distributions that are further apart from one another will have a lesser degree of overlap with one another and are therefore easier to discriminate since there is a smaller chance that the experience of the stimulus will come from the overlap area of uncertainty. One application of SDT is to quantify this level of uncertainty and so establish an index of sensitivity to change irrespective of response bias. SDT achieves this by partitioning the overall pattern of response to reflect an element of sensitivity and an element of response bias separately. Herein I use SDT to determine an index of sensitivity (d') and report bias in appendix F (see below).

In the example above I have discussed the use of STD to discriminate between signal-present and signal-absent distributions. However, it should be readily apparent that the same principles can be applied when the two distributions reflect two different signal distributions that must be discriminated from one another. SDT can therefore be applied to detection of *change in signal* as well as detection of signal. It is to this use that I put SDT herein in the calculation of an index of sensitivity to change in signal.

Calculation of d'

The meaning of the d' statistic is best understood with reference to Figure 4.1. We have discussed that the separation between the two distributions and the width of the distributions determine the level of uncertainty experienced.

Sensitivity is, in essence, the inverse of this uncertainty; the more sensitive the

participant the less uncertainty experienced. Thus, we can derive an index of sensitivity from the same two pieces of information. In practice, this is done by first z-transforming each of the target-present and target-absent distributions (such that each has a standard deviation of 1) and then measuring the distance between their means. This measure of separation in terms of standard deviation is d' . More formally, d' is calculated according to Equation 4.3 below.

$$d' = z(H) - z(FA) \quad (4.3)$$

Where $z(H)$ is the z-transform of the Hit rate and $z(FA)$ is the z-transform of the FA rate (see section on K statistics above for our definition of what constitutes a Hit and what a false alarm). Note that z is denoted Φ^{-1} in some texts; this is a functionally equivalent mathematical transformation (technically, inverse phi is the transformation that converts a probability into a z score, but it is clear enough for our purposes to use 'z' to denote both the transformation and the resulting score).

Practically, calculation of d' can be done in several ways. There are software programmes that will calculate d' automatically, or alternatively it can be done manually using tables of z-scores such as can be found in most statistics text books, or 'semi-manually' using the NORMSINV function in Microsoft Excel. As a cautionary note, and to illuminate a particular limitation of the d' statistic, I next describe a nicety of the manual and semi-manual calculation procedures.

I have discussed above that d' is calculated by subtracting the mean of the z-transform of the FA rate from the mean of the z-transform of the Hit rate. A problem occurs, however, when either or both the Hit or FA rate is 0.00 or 1.00 for a participant. These maximum and minimum values are sometimes empirically

found for individuals, particularly within conditions that have relatively few trials or, in the case of a Hit rate of 1.00, within a task condition that an individual finds particularly easy. However, such values cannot be subjected to z-transformation (they would equate to z values of infinite magnitude). Therefore, a correction to the Hit or FA proportion must be made prior to calculating the z-transform. The size of such a correction is based on two factors; the number of change and no-change trials undertaken by the participant, and common sense. To illustrate, I discuss below the treatment of Hit rates of 1.00. However, similar logic is also applied to Hit rates of 0.00 and FA rates of 1.00 and 0.00.

Trial numbers must be taken into consideration since a Hit rate of 1.00 calculated from ten trials is supported by less evidence of 'actually being very close to 100%' than a Hit rate of 1.00 calculated from one hundred trials. Essentially, we ask the question 'what would happen if we were to ask the participant to continue to do this task...? Would the participant continue to achieve 100% Hits? We assume that eventually the participant's performance would drop from this perfect level. How soon it would likely be before the participant's performance dropped is the criterion used to make our pre-z-transform correction. If we have used only 10 trials to calculate the Hit rate then the possibility exists than on the eleventh trial that participant would make a mistake. Therefore, we can adjust the Hit rate from $10/10 = 1.00$ to $10/11 \approx 0.91$. However, if we used 100 trials to calculate to Hit rate then our adjusted Hit rate would be $100/101 \approx .99$.

More formally, for maximum (1.00) values:

$$p = N/(N+1) \tag{4.4}$$

And for minimum (0.00) values:

$$p = 1/(N+1) \quad (4.5)$$

Where p is the adjusted proportion of Hits or FAs and N is the number of trials from which the original unadjusted proportion (the 1.00 or 0.00) was calculated.

The formulae for the maximum and minimum corrections are sometimes given as $p = (N-1)/N$ and $p = 1/N$ respectively (Healy, Berger, Aberson, Saw, & Romero, 2003). However, an examination of the rationale above shows that the corrections in formulas 4.4 and 4.5 above are the more accurate estimations. Alternative correction formulae have also been proposed; an alternative to assuming a change in response on the $N+1$ th trial is to assume a change on the $(N+2)$ th trial (Wixted & Lee, n.d.). This is a less conservative estimate of performance and results in correction formulae of $p = 1 - (1/2N)$ for the maximum and $p = 1/2N$ for the minimum. For a review of other correction techniques see Miller (1996) and Macmillan, Rotello & Miller (2004). While there is some debate as to the best correction to use, throughout this thesis unless otherwise stated I use formulas 4.4 and 4.5 above in such corrections. Common sense also comes into play here. Suppose a participant has indicated that a particular condition was so easy as to be trivially challenging and has said that he or she would never make a mistake on this condition no matter how many trials we ask him or her to do. Suppose further that converging evidence supports this argument (e.g. all participants find the particular condition easy and achieve 100% correct). It may now be more parsimonious to assume that performance is indeed near to 100% correct even if this cannot be concluded from the number of trials alone. For example, if every single participant gets a Hit rate of 1.00 across twenty trials, and all indicate that the task is not found to be challenging it may be more appropriate

(that is, give a better estimate of reality) to reduce their Hit rates prior to z-transform to .99 rather than to .95, as might otherwise be suggested by the purely mathematical 'trial number' corrections above.

While the theoretical maximum and minimum values of d' are $+\infty$ and $-\infty$ respectively, with a score of zero indicating complete inability to discriminate, use of the .01/.99 correction described above results in a practical maximum value of just over 4.65 and a practical minimum of just under -4.65. (Macmillan & Creelman, 2005).

Treatment of negative d' scores

Another decision that must sometimes be made when using the d' statistic is how one treats negative scores. A negative d' score indicates worse than chance performance and can be indicative of a participant misunderstanding task instructions. Alternatively, if a task (or condition within a task) is so difficult that performance is essentially at chance, then natural within-condition experimental variance may cause some participants to have negative scores while others have positive scores for a particular condition. In some instances it may be appropriate to exclude a participant from a study altogether if, for example, performance is so low across conditions as to show complete lack of engagement with the task. However, the question remains how to treat negative d' scores obtained in certain conditions by participants who have achieved higher scores in other conditions. Such participants may be clearly engaged with a particular task, but find one particular condition difficult. There are two schools of thought here: One can take the approach that a negative d' score is essentially meaningless; in WM tasks it equates to a negative Cowan's/Pashler's K score and so represents a negative

number of items in WM. If this approach is taken, then such scores may be corrected to the lowest meaningful score (zero). Alternatively, one can refrain from such correction and allow the negative score to 'drag down' the average score across participants for the condition in question. I have adopted the latter approach throughout as I believe it is better justified: 'Correcting' such scores to 0 has an unjustifiable effect on data averaged across participants. As discussed above, such a score may not reflect failure to understand task instructions but instead reflect variance around a chance or near chance level of performance resulting from task difficulty within an experimental condition. Thus, changing the score to 0 before averaging scores across participants results in an overestimate of the 'true' (i.e. population) score for the condition in question. Similarly, it results in an underestimation of sample variance in the condition in question.

Additionally, there are several situations in which negative d' scores are meaningful: It may be theoretically sound to conceptualise of a negative number of items stored in WM if we allow that a wrong item representation, held with sufficient strength, should be considered a negative version of a correct item (see Meese, Baker & Summers, 2012 for a paradigm in which this occurs). For example, the employment of imperfect neural mechanisms to the determination of target/pedestal separation has been shown to lead to a meaningful negative d' scores under specific conditions in which a pedestal (mask) and target become inappropriately combined after an initial process in which the target exerts a suppressive effect on the pedestal (so-called dilution masking; Baker, Meese & Georgeson, 2013).

Measure of bias

As well as determining sensitivity, SDT also can be used to determine a measure of bias. Bias is a participant's propensity to answer either 'present' or 'absent' when uncertain. For example, one participant may decide to give a random response when unsure while another may choose to always respond 'absent' in such situations. When performing a change (or signal) detection task, a participant will at any given time have set a criterion level. Such a criterion reflects the level of sensory experience, or in graphical terms the location on the decision variable abscissa (see Figure 4.1) below which it shall be reported that the stimulus came from the noise distribution and above which it shall be reported that the stimulus came from the signal plus noise distribution. The purpose of SDT as used in this thesis is to derive an index of sensitivity that is independent of bias. As such, any reported measure of bias should be orthogonal to that of sensitivity (e.g. d'). There are two commonly reported ways of measuring bias: ' c ' and β . The c measure reflects the distance between the criterion position on the decision variable axis and the 'neutral point' where the distributions meet. The β measure reflects a likelihood ratio between the signal and noise distribution values at the level of criterion. Essentially, use of β assumes that responses are made based on whether the likelihood ratio (height of the signal distribution divided by height of the noise distribution) on any given trial exceeds the criterion likelihood ratio value denoted by β ; see Figure 4.1). More formally, c can be calculated as:

$$c = - \frac{\Phi^{-1}(H) + \Phi^{-1}(F)}{2} \quad (4.6)$$

Where H is the Hit rate and F the False Alarm rate, and β is calculated as:

$$\beta = e^{\left\{ \frac{[\Phi^{-1}(F)]^2 - [\Phi^{-1}(H)]^2}{2} \right\}} \quad (4.7)$$

Although calculation of bias is computationally trivial, I have chosen not to report measures of bias in the main text of the experiments reported herein, but instead list bias (c) along with Hit and False Alarm rates separately in appendix F. There are several reasons for this. First, bias is not commonly reported in change detection tasks in this field, so the benefit of comparing bias to studies in the literature is lost. Second, my hypotheses do not predict changes in bias but instead predict changes in sensitivity. Third, no reliable systematic changes in bias were found across conditions in any study reported herein.

Assumptions of d'

D' is the most widely reported signal detection measure in memory research, yet it is a parametric statistic based on two assumptions; normality and homogeneity of variances between the two underlying signal and signal plus noise distributions.²⁰ This is problematic when we consider what these distributions represent. I have noted above that the distributions represent a 'decision variable' that is affected by both external and internal factors and reflects what the participants actually *experience*. As such these distributions are inherently empirically not measurable with the result that we cannot check that the

²⁰ When the variances of the two distributions are known but are unequal then an alternative statistic d_a can be employed. D' can thus be thought of as a special case of d_a in which variances are known to be (or believed to be) equal.

parametric assumptions hold true for any given case. This is important as violations can lead to systematic inaccuracies in the statistic. For example, if the homogeneity of variances assumption is violated then a change in the criterion can lead to a change in the d' index even when sensitivity itself has not changed (Verde, Macmillan & Rotello, 2006).

There are three ways of dealing with this issue. The first involves taking multiple observations at different levels of criterion, plotting a receiver operating characteristic (ROC) curve (see below), and taking the area under this curve as a non-parametric index of sensitivity. While this is theoretically sound practice, it is not always practical as it requires an increase in the number of observations collected. An alternative solution is to attempt to derive an ROC curve from a single point and again take the area under this estimated curve. This has the advantage of requiring fewer observations but has theoretical drawbacks (see below). The third alternative is to use d' as an index but remain aware of its limitations. In Chapter 8 I report an experiment designed to create an ROC curve using multiple criteria. Throughout the rest of this thesis I utilise d' as my primary measure of sensitivity to change.

ROC curves and non-parametric SDT measures

For a given level of sensitivity (that is, for any given participant under certain conditions), a participant is free to adjust their criterion (c or β) from trial to trial. It is possible to design experiments that encourage participants to adopt one criterion on some trials and another on others (for example by differentially incentivising the relative rewards / punishment associated with Hits and FAs).

As the placement of the criterion is moved, the hit and false alarm rates increase or decrease in tandem. When the criterion is set particularly low both the Hit and FA rates increase. The relationship between the Hit and FA rates as the criterion moves can be plotted for each possible level of sensitivity. This is known as an ROC curve since its shape is dependent on the individual 'receiver' of the signal (the participant). Examples of ROC curves that would be associated with levels of sensitivity corresponding to d' values between zero and four are shown in Figure 4.2.




Figure 4.2. ROC curves corresponding to different d' values.

ROC curves are generated by plotting a data point representing the level of Hits and FAs for each level of criterion (criterion varies from low to high as one moves from left to right across any of the lines representing different levels of sensitivity). These points can then be joined and the rest of the curve completed via extrapolation (either through linear extrapolation though this tends to underestimate sensitivity, or by assuming normality in the underlying decision

variable and using this assumption to fit a curvilinear function; Stanislaw et al., 1999). The more 'bowed' an ROC curve (the more it tends towards the top left corner of probability space – see Figure 4.2), the greater the area under it, and thus the greater the separation between the signal and noise probability distributions and so the greater the receiver's sensitivity.

An alternative option to manipulating different placements in criterion with reward/punishment incentives is to ask participants to rate on an ordinal scale how confident they are that they have made the correct response decision after each trial. I utilise this methodology in Experiment 9. Ratings of confidence can be used to infer criterion changes in the following way: Responses that have been made at the high end of the confidence scale reflect either an indication of belief in definite presence of change (when response is [change; very sure]) or else of definite absence of change (when response is [no-change; very sure]). Thus, confidence ratings taken alongside present/absent decisions can be transformed into a scale where each end represents an extreme of confidence (i.e. definitely present to definitely absent). An alternative to this method is to provide such a Likert-like scale directly to participants as the instrument on which they respond. This has the advantages of not requiring the simple data transformation described above but in some instances may be a cognitively more complex task for participants than the simple 'confidence in response' approach. Next, cumulative frequency proportions for Hits and FAs can be computed for each of the confidence ratings. These Hit and FA proportion pairs can then be plotted as the points on an ROC curve taking the proportion associated with 'definitely change' as the first point, the cumulative fraction associated with the next most sure-of-change level as the next point and so on. Note that since the proportions are

cumulative they are necessarily constrained to increase or remain constant as the observer's 'confidence that it was a change trial' criterion is relaxed. Figure 4.3 below illustrates this process with example data for a five point Likert-like scale tumour detection experiment.

Table 1.6. Observer response data from an experiment with a test set of 40 diseased and 60 disease-free cases

	Rating of the confidence that the case is diseased					Total cases
	5 (definitely present)	4 (probably present)	3 (possibly present)	2 (probably absent)	1 (definitely absent)	
Diseased cases	20	8	6	4	2	40
Disease-free cases	3	6	6	9	36	60

Table 1.7. The cumulative frequencies of the observer response data

	Cumulative frequencies				
	5	5 + 4	5 + 4 + 3	5 + 4 + 3 + 2	5 + 4 + 3 + 2 + 1
Diseased cases	20	28	34	38	40
Disease-free cases	3	9	15	24	60

Table 1.8. The cumulative fractions of the observer response data

	Cumulative fractions				
	5	5 + 4	5 + 4 + 3	5 + 4 + 3 + 2	5 + 4 + 3 + 2 + 1
Diseased cases	0.50	0.70	0.85	0.95	1.00
Disease-free cases	0.05	0.15	0.25	0.40	1.00

Figure 4.3. Example data using confidence ratings to plot an ROC curve. Reproduced from Journal of the International Commission on Radiation Units and Measurements (ICRU).

The proportion of area that is under the ROC curve provides an accurate, non-parametric estimate of sensitivity. This measure can take different names, depending on the method used to generate the ROC curve. When only a single point on the graph is known (i.e. when criterion has not been systematically manipulated as described above), it is still possible to plot an ROC

curve since any data point on the curve has implications for the shape of the curve. However, when only a single data point is available then there are a number of possible 'proper' ROC curves (i.e. curves consistent with the use of an optimal decision strategy based on evaluation of likelihood ratios; Egan, 1975) that could pass through this single point and the 'non-parametric' summary statistic (e.g. A' , A'' , A_z or A) must then be calculated from an *estimated* ROC curve extrapolated from this single data point. The area under such a curve does not constitute a non-parametric index since estimation of the curve necessarily makes assumptions as to the underlying decision criteria distributions (Zhang et al., 2005). For example, the A_z^{21} statistic assumes Gaussian distributions and the A' statistic assumes Gaussian distributions with equal variances). This is of particular importance since historically A' has been widely used as a non-parametric substitute for d' (e.g. Grier, 1971; Pollack & Norman, 1964), but this claim has been proven to be untrue (Macmillan & Creelman, 1996; Pastore, Crawley, Berens & Skelly, 2003).

I should acknowledge here that the various measures of area under a single-point extrapolated ROC curve can still be of use in certain situations, indeed A' and the other estimated area measures (e.g. A'' : Smith, 1995; A : Zhang & Mueller, 2005; A_z : Swets, 1986) are 'pleasant' statistics inasmuch as their ranges correspond to intuitive representations of performance, with 0.0 representing perfectly poor performance, 1.0 corresponding to perfect performance and 0.5 indicating chance performance. However, such measures have been widely

²¹ The 'z' in the A_z index implies an assumption of normality (Gaussian noise) for the underlying distributions. A_z can be used both to estimate the area under an ROC curve derived from multiple points (Stanislaw et al., 1999) or single points (Verde et al., 2006).

misused (Macmillan et al., 1996; Pastore et al., 2003; Smith, 1995), and so care must be taken in their interpretation and use.

Why choose d' ?

Given the constraints on the use of d' , and the inherent difficulty in evaluating the assumptions underlying it, why have I chosen to use d' my primary index of sensitivity? We have noted that non-parametric SDT measures rely on calculation of the area under an ROC curve. When only a single data point in ROC space is available for the calculation of the ROC curve then the estimate of the area under it cannot be described as being truly non-parametric. While the area under an ROC curve derived from multiple data points may have advantages as an index of sensitivity it is not always practical to calculate. Therefore, I use d' as a measure of performance in various WM change detection tasks. Recognising the limitations of the d' statistic as described above, I note here that throughout, unless otherwise stated, analysis of percent correct and A' as alternative measures would not have resulted in rejection of null hypotheses that would have otherwise been rejected using d' .

D' for AFC tasks

I described the difference between yes/no and AFC tasks in Chapter 1. Just as it is possible to compute an index of sensitivity in yes/no tasks (as

described above), it is likewise possible to derive a similar measure when the task is AFC²². For a 2-AFC task, the formula for d' is

$$d' = 1/\sqrt{2} \cdot [z(H)-z(FA)] \quad (4.8)$$

Comparison of Equations 4.3 and 4.8 renders it immediately apparent that, given similar Hit and FA rates d' in a 2-AFC task will be calculated as less than that in a present/absent task. This down-weighting of 2-AFC task scores occurs to correct for such tasks being 'easier' than a similar present/absent task (Macmillan & Creelman, 2005, p.168). At first glance this seems an odd assertion. After all, an experimenter can manipulate task difficulty in any number of ways and different participants may find different tasks more or less difficult if, for example, they have already been able to practice one particular type of task. Indeed, while 2-AFC LTM tasks are widely believed to be easier than present/absent LTM tasks (e.g. see Jang, Wixted & Huber, 2009) this may well not be so when the task specifically probes WM (Makovski, Watson, Koutstaal & Jiang, 2010). How is it, then that we can say that one type of task is more difficult than another? More to the point, how is it possible to compute a relative difficulty for one task 'as if it had been another' by the simple expedient of multiplying sensitivity scores by a constant ($1/\sqrt{2}$). The answer comes from understanding just what Macmillan and Creelman (2005) mean by 'difficult'. Difficult in this sense refers to the information theoretic difficulty level of the task rather than a practical difficulty level: Consider a visual WM task in which a stimulus (A) is presented, followed by a brief retention interval, followed by a response array comprising two stimuli, the 'old' stimulus just

²² While d' is most usually described as 'sensitivity to change' on present/absent tasks, for AFC tasks this is not such an appropriate label since the cognitive process tested by such tasks is the ability to identify targets from lures rather than make a judgement as to the presence/absence of a target. A better description of d' in AFC tasks might be 'discrimination ability' (that is, the ability to discriminate between lures and the representation in memory of the target).

previously presented for memorisation (A_1), and a 'new' stimulus (B) as a lure (see Figure 4.4). This is one of the simplest versions of a visual WM 2-AFC task, and varies from a traditional equivalent yes/no task only in the presentation of an additional response option.




Figure 4.4. A simple present/absent WM task and an 'equivalent' 2-AFC WM task. From an information theoretic viewpoint the 2-AFC task is easier, while from an experiential viewpoint it is not possible to determine which task is easier based only on the task structure.

How does the presence of this lure response option make the task 'theoretically easier'? The answer comes from the stimulus-to-memory comparison processes the participant is assumed to carry out in such tasks. In the case of the present/absent (1-AFC) task the participant must compare the internal representation of stimulus A with the test item and make a same vs. different judgement. However, in the case of the 2-AFC task the participant can still perform this comparison process for the target stimulus but may also compare

stimulus B (the lure) with the internal representation of the memory item A. The comparison of the internally held representation with the lure provides additional information on which to make the final dichotomous judgement decision, thus the 2-AFC task is theoretically easier (see Macmillan & Creelman, 2005, p. 168-170 for a graphical proof that the scaling constant is $1/\sqrt{2}$; see Stone, 2010 for a step by step mathematical proof of this relationship)²³.

More generally, it is possible to compute a d' score for any M-AFC task according to the relationship

$$P_c = \int_{-\infty}^{\infty} \phi(t - d') \cdot \Phi(t)^{M-1} dt \quad (4.9)$$

Where P_c is the proportion correct, M is the number of response options in the AFC task, t is the sensory magnitude of the target stimulus representation on any given trial, $\Phi(t)$ is the probability that t will be greater than the sensory magnitude of a single equivalent noise representation (or lure representation depending on task) on the trial in question, $\Phi(t)^{M-1}$ is the probability that t will be greater than the sensory magnitudes of *each of* the noise or lure representations from $M-1$ positions (i.e. the probability that using the decision rule 'choose the representation with the highest sensory magnitude' [an assumption of this model] the signal or target representation will be chosen over *any of* the lures), $\phi(t-d')$ is the relative likelihood of t , and dt signifies integration over all possible values of t (Kingdom & Prins, 2009). In practice, since Equation 4.9 is not easily invertible

²³ This relationship is not always properly acknowledged or understood. A quick online search will reveal a number of tutorials in use of SDT that mistake the formula for calculation of d' in a present/absent task as being the appropriate calculation to apply to determine d' for a 2-AFC task (e.g. see <http://openwetware.org/wiki/Beauchamp:dprime>)

this is most usually done either with an iterative search algorithm which tests values of d' until one is found that results in a P_c score suitably close to the input P_c value (Kingdom & Prins 2009) or from M-AFC tables (Hacker and Ratcliff, 1979).

In Chapter 7 I describe an experiment which combines a 4-AFC WM task with a change/no-change WM task. Although proportion correct would be a reasonable measure of sensitivity in the 4-AFC task reported, it would not be appropriate for the change/no-change task data (since it fails to account for bias). As it makes better intuitive sense to compare like with like, the 4-AFC proportion correct data therein are transformed into d' scores using Hacker and Ratcliff's (1979) M-AFC conversion tables so that they may be readily compared with d' scores from the change/no-change task (see Chapter 7).

The possibility of calculating d' for both yes/no and the various types of AFC task mean that in theory it should be possible to compare sensitivity levels across different tasks. However, it is important to treat such comparisons with great caution; even though one can correct for the *information theoretic difficulty* of AFC vs. present/absent (change/no-change), using for example the $\sqrt{2}$ rule described above for 2-AFC, the fact that yes/no and M-AFC WM tasks place different cognitive demands on participants even when such tasks are designed to measure the same thing (Makovski et al., 2010) means that one can rarely correct for the *actual or experiential task difficulty* unless one also has empirical data that quantitatively define this relationship for any two particular tasks.

Chapter references – Chapter 4

- “Report 79: Receiver Operating Characteristic Analysis in Medical Imaging.”
Journal of the ICRU (2008) 8 (1), 57-62. doi: 10.1093/jicru/ndn017
- Baker, D. H., Meese, T. S. & Georgeson, M. A. (2013). Paradoxical psychometric functions ("swan functions") are explained by dilution masking in four stimulus dimensions. *i-Perception (4) 1.* 17-35.
- Bartlett, F. C. (1932). *Remembering: A Study in Experimental and Social Psychology.* Cambridge University Press.
- Beauchamp Lab Notebook.* (2010).
<http://openwetware.org/wiki/Beauchamp:dprime> (accessed April 22, 2013).
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences 24,* 87-185.
- Egan, J. P. (1975). *Signal Detection Theory and ROC Analysis* (Academic Press, New York)
- Grier, J. B. (1971). Nonparametric indexes for sensitivity and bias: Computing formulas. *Psychological Bulletin, 75(6),* 424-429.
- Hacker, M. J. & Ratcliff, R. (1979). A revised table of d' for M-alternative forced choice. *Perception & Psychophysics, 26(2),* 168-170.
- Healy, M. R., Berger, D. E., Aberson, C. L., Saw, A., & Romero, V. L. (2003). *Signal Detection theory Tutorial.* Claremont Graduate University, Web Interface for Statistics Education (WISE), from <http://wise.cgu.edu/util/> (accessed September 01, 2012).
- Jackson, M. C., & Raymond, J. E. (2008). Visual working memory for faces. *Journal of Experimental Psychology: Human Perception and Performance, 34(3),* 556 –568.
- Jackson, M.C., Wu, C., Linden, D.E.J., & Raymond, J.E. (2009). Enhanced Visual Short-Term Memory for Angry Faces. *Journal of Experimental Psychology: Human Perception and Performance, 35(2),* 363-374.
- Jang, Y., Wixted, J. T. & Huber, D. E. (2009). Testing Signal-Detection Models of Yes/No and Two-Alternative Forced-Choice Recognition Memory. *Journal of Experimental Psychology: General, 138(2),* 291-306.
- Kingdom, A. A. & Prins, N (2009). *Psychophysics: A Practical Introduction.* Academic Press. ISBN: 9780123736567
- Macmillan, N. A., & Creelman, C. D. (1996). Triangles in ROC space: History and theory of “nonparametric” measures of sensitivity and response bias. *Psychonomic Bulletin & Review, 3,* 164-170.

- Macmillan, N. A. & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, NJ: Erlbaum.
- Macmillan, N. A., Rotello, C. M., & Miller, J. O. (2004). The sampling distributions of Gaussian ROC statistics. *Perception & Psychophysics*, *66*, 406-421.
- Makovski, T., Watson, L. M., Koutstaal, W. & Jiang, Y. V. (2010). Method matters: Systematic effects of testing procedure on visual working memory sensitivity. *Journal of Experimental Psychology: Learning Memory and Cognition*, *36*(6), 1466–1479.
- Marcum, J. I. (1947). *RM-754, A Statistical Theory of Target Detection by Pulsed Radar*. RAND Corporation research memorandum report.
- Meese, T. S., Baker, D. H. & Summers, R. J. (2012). *Fake Tilt Shift Miniaturization Causes Negative D-Prime for Detecting Reality*. In: Predicting Perceptions: Proceedings of the 3rd International Conference on Appearance. Lulu Press, Edinburgh UK, p. 61. ISBN 978-1-4716-6869-2
- Miller, J. (1996). The sampling distribution of d' . *Perception & Psychophysics*, *58*, 65-72.
- Pashler, H. (1998). Familiarity and visual change detection. *Perception & Psychophysics*, *44*, 369-378.
- Pastore, R. E., Crawley, E. J., Berens, M. S., & Skelly, M. (2003). "Nonparametric" A' and other modern misconceptions about signal detection theory. *Psychonomic Bulletin and Review*, *10*, 556–569.
- Pollack, I. & Norman, D. A. (1964). Non-parametric analysis of recognition experiments. *Psychonomic Science* *1*, 125-126.
- Schatz, P., Jay, P. A., McComb, J. & McLaughlin, J. R. (2005). Misuse of statistical tests in *Archives of Clinical Neuropsychology* publications. *Archives of Clinical Neuropsychology*, *20*, 1053–1059
- Scheff, T. (2011). The Catastrophe of Scientism in Social/Behavioral Science. *Contemporary Sociology: A Journal of Reviews* *40*(3), 264-268.
- Smith, W. D. (1995). Clarification of sensitivity measure A' . *Journal of Mathematical Psychology* *39*, 82-89.
- Spector, P. E., & Brannick, M. T. (2011). Methodological urban legends: The misuse of statistical control variables. *Organizational Research Methods*, *14*, 287-305.
- Stanislaw, H. & Todorov, N. (1999). Calculation of signal detection theory measures. *Behaviour Research Methods, Instruments, & Computers*, *31*(1), 137-149.
- Swets, J. (1986). Form of empirical ROCs in discrimination and diagnostic tasks: Implications for theory and measurement of performance. *Psychological Bulletin* *99*, 181-198.

- Verde, M. F., Macmillan, N. A., & Rotello, C. M. (2006). Measures of sensitivity based on a single hit rate and false alarm rate: The accuracy, precision, and robustness of d' , A_z , and A' . *Perception & Psychophysics*, 68, 643-654.
- Vul, E., Harris, C., Winkeilman, P. & Pashler, H. (2009). Puzzlingly High Correlations in fMRI Studies of Emotion, Personality, and Social Cognition. *Perspectives on Psychological Science*, 4(3), 274-290.
- Wixted, J & Lee, K. (n.d.). *Signal Detection Theory*.
<http://psy2.ucsd.edu/~kang/sdt/sdt.htm> (accessed September 01, 2012).
- Zhang, J. & Mueller, S. T. (2005). A note on ROC analysis and non-parametric estimate of sensitivity. *Psychometrika*, 70(1), 203-212.

**PART 2 – VISUAL WORKING MEMORY FOR EMOTIONAL
FACES**

CHAPTER 5.**Singleton Emotional Faces and the Angry Benefit²⁴**

²⁴ The two experiments reported in this chapter are in press in *Journal of Experimental Psychology: Human Perception and Performance*. They are presented here with expanded discussion.

Visual cues that signal social threat are often fleeting, such as a brief frown or a transient glare. Nevertheless, encoding and remembering their presence is critical for planning appropriate social behaviour within a current social episode (e.g., a conversation). To do so, one must first deploy visual attention to selectively encode relevant social information, and then engage visual working memory (WM) to retain the information after it is no longer present in the sensory array. Whereas visual WM is a limited capacity, short-term dynamic memory system that retains information for brief periods (Cowan, 2001; Luck & Vogel, 1997), selective visual attention, also thought to be a limited resource, is a set of neural mechanisms that bias the competition among sensory representations to gain access to high-level cognitive processes such as WM (Bundeson, 1990; Desimone & Duncan, 1995; Kastner & Ungeleider, 2000). Attentional biases can originate from 'top down' mechanisms that code current behavioural goals and make stimulus predictions, or from task-independent mechanisms sensitive to salience in the visual array that can arise from sensory contrasts, or from emotional (Vuilleumier, 2005) or motivational content (Anderson, Laurent & Yantis, 2011; Raymond & O'Brien, 2009). Not only is selective attention generally viewed as the 'gatekeeper' of WM (e.g., Bays & Husain, 2008), numerous studies suggest that it plays a critical role in selecting information from WM (e.g., Griffin & Nobre, 2003), and in prioritising maintenance (Makovski, Sussman & Jiang, 2008).

These putative links between WM and selective attention predict that stimuli affording strong task-independent attentional biases should be especially successful as potential WM memoranda, even when top-down biases such as task relevance do not especially favour them. Here, I investigate this possibility using a conventional visual WM change detection task. Task-independent attentional

biases were manipulated by using social stimuli (faces) with and without emotional expressions. Specifically, I measured WM for arrays of four faces, where one expressed happiness or anger and the other three were affectively neutral. Memory for each face was equally likely to be tested, making each face equally task-relevant, regardless of its expression. Using this experimental set-up I could (1) examine the effects of irrelevant emotional salience on WM by comparing performance for emotional versus neutral faces seen at the same time; (2) investigate the effect of emotional valence on WM by comparing memory for positive versus negative emotional singleton faces (seen in different trials); and (3) assess whether WM for non-emotional faces in a scene benefits or suffers when a positive versus negative face is concurrently present by comparing WM performance for neutral faces when the accompanying emotional singleton was positive versus negative.

There is widespread empirical evidence that facial expressions displaying threat are especially likely to bias selective attention (Eastwood, Smilek, & Merikle, 2003; Feldmann-Wüstefeld, Schmidt-Daffy & Schubö, 2011; Fox & Damjanovic, 2006; Hahn, Carlson, Singer & Gronlund, 2006; Huang, Chang & Chen, 2011; Öhman, Flykt, & Esteves, 2001; Pratto & John, 1991), make disengaging attention more difficult (Fox, Russo, & Dutton, 2002), and generate greater interference on simple detection tasks (Pessoa, Mckenna, Gutierrez, & Ungerleider, 2002). If attention plays a role in facilitating visual WM, then these findings predict that a single threatening face appearing in a crowd of other non-threatening faces should be remembered better than other neutral faces seen with it, and better than a singleton face displaying a positive expression. Moreover, considering the theoretical notions of biased competition and WM's severe capacity limitation, WM

for neutral objects should be weakened as the strength of the attentional bias toward an emotional object in a to-be-remembered array goes up. Thus if an angry singleton face could bias attention more than a happy one, then memory for a neutral face should be worse when accompanied by an angry face than by a happy face.

Other related theories make similar predictions. One such notion is that multiple memory items compete for resources, causing mutual suppression of each other's representation (Bahcall & Kowler, 1999). In this view, if one item were more successful at acquiring WM resources, then it should also be more powerful at suppressing other representations. Applied here it predicts poorer memory for neutral objects when other objects in a scene are emotional. Another view making a similar prediction is that arousal acts to sharpen attentional competition during encoding, and that this, then, enhances memory for attended stimuli and weakens memory for ignored stimuli (Mather & Sutherland, 2011). If threatening faces are more arousing than happy faces, and are better at biasing attention, then this view predicts better WM for angry versus happy singletons and also greater costs to neutral faces seen with angry versus happy singletons.

However, other considerations make the opposite prediction for the WM fate of neutral stimuli encoded at the same time as emotional stimuli. First a common sense approach suggests that a threat processing strategy that suppresses concurrent monitoring of non-threat items seems non-optimal because other scene information could reasonably be expected to inform the behavioural and social planning needed to deal with threat. This suggests that although threat-related stimuli might capture attention initially, such attentional prioritization should eventually be suppressed so that high-level processing of other relevant

information in the scene could be enabled (Becker, 2009). Moreover, attentional selection and subsequent processing should be biased by motivational factors (Anderson, Laurent & Yantis, 2011; Raymond & O'Brien, 2009) so that information needed to meet current goals is processed as a priority, overriding momentary attention capture by task-independent features in stimuli. Supporting this possibility is a study by Becker and Detweiler-Bedell (2009) who showed that when passively viewing an array of faces comprised of three neutral and one angry or fearful face, gaze is directed away from the expressive negative face. This finding predicts that WM for neutral faces might be unaffected or indeed might benefit from the presence of a threatening face.

A similar prediction is suggested by several studies that report an attention-initiated boost for visual scene memory (Lin, Pype, Murray, & Boynton, 2010; Makovski, Swallow, & Jiang, 2011; Swallow & Jiang, 2010; 2011). In these studies, participants were given the task of remembering a series of briefly presented visual scenes whilst at the same time monitoring for the appearance of a simple target. The general finding is that scene memory, including short-term memory (Makovski et al., 2011), is better on trials in which a target versus a distractor object was concurrently presented. The explanation of this so-called attentional boost effect is that the presentation of a goal relevant, or motivationally salient, item produces a brief generalised enhancement of processing of all concurrent stimuli. This effect suggests that the presentation of an emotional face might have a similarly positive scene-wide boost effect. Angry faces are often considered to be motivationally salient, activating approach responses (Carver & Harmon-Jones, 2009; Wilkowski & Meier, 2010), whereas happy expressions, on the other hand, may be viewed as less motivating because they communicate

approval and do not imply that a change in behaviour is necessary. If the presentation of a motivationally salient object could induce an attention boost for WM, then WM for neutral faces presented concurrently with an angry face might be better than that for neutral faces presented concurrently with a happy face.

Several previous studies have shown that negative emotional expression facilitates visual WM for face identity (Jackson, Wolf, Johnston, Raymond, & Linden, 2008; Jackson, Wu, Linden, & Raymond, 2009; Sessa, Luria, Gotler, Jolicoeur, & Dell'Acqua, 2011). These studies measured WM performance using a simple change detection task in which a study array comprising a small number of different people's faces each bearing the same emotional expression was presented for several seconds for encoding. After a one or two second retention interval, a single test face was presented and participants judged if it was from the study array or not. WM for angry (Jackson et al., 2008; 2009) or fearful (Sessa et al., 2011) faces was found to be significantly better than that for neutral faces. Jackson et al. (2008; 2009) also showed that performance was better for angry versus happy faces and that there was no difference in WM for happy versus neutral faces. One interpretation of these findings is that WM capacity is enhanced by negative emotional information (Jackson et al., 2009). Another is that negative emotion on a face encourages a combination of greater precision during encoding and better maintenance during the delay period (Sessa et al., 2011), processes that could potentially be independent of selective attention and operate in a non-competitive way. These contrasting viewpoints make different predictions about the fate of affectively neutral information present in a scene with negative emotional content. If anger were to enhance WM capacity, *per se*, then one might expect a WM boost for other concurrently viewed stimuli, i.e., a scene-

wide benefit. The enhanced precision view predicts WM enhancement for the negative object only (i.e., an object-based benefit), leaving WM for any other affectively neutral or positive objects in the scene unaffected.

To address these issues, I conducted two studies, each using a modification of Jackson et al.'s (2009) face WM task that involved presenting a single emotional face (angry or happy) and three neutral faces as the study array. I always presented four faces so that WM capacity (about 2.5 faces on average, Curby & Gauthier, 2007; Jackson & Raymond, 2008) would be exceeded, preventing performance from reaching ceiling. In the first experiment, I also presented uniform trials in which all faces in the study array had the same expression (all angry, all happy or all neutral), replicating the conditions of Jackson et al. (2009). To ascertain whether an angry benefit to WM could be found with a single face, I compared WM for angry versus happy singletons, so as to control for any purely singleton effects, i.e., a benefit accrued from being an atypical item in the array. I also compared singleton WM performance with that in the corresponding uniform condition (all angry or all happy). If selective attention could boost WM, then performance might be better for singletons in the singleton condition than for any one face in the uniform condition, and this benefit should be bigger for angry than for happy faces. To address the question of the fate of neutral faces in singleton arrays, I compared WM performance for neutral faces seen in each emotional singleton condition and in the neutral uniform condition. To anticipate, I found, as expected, that WM for angry faces was better than that for happy faces in both the singleton and uniform conditions. Surprisingly, I found no effect of an emotional singleton, angry or happy, on WM for concurrently viewed neutral faces compared to performance for the uniform neutral condition.

The second experiment used only singleton study array conditions and monitored gaze position during each study array interval to assess the possibility of differential encoding. Not only did I fully replicate the relevant results from Experiment 1, I found that fixation patterns were not different when the emotional singleton was angry versus happy, suggesting that gaze behaviour cannot account for the observed negative face benefit to WM.

EXPERIMENT 1

The primary goal of this experiment was to allow a comparison between uniform study array conditions and singleton study array conditions. In the former, study arrays comprised four faces, each bearing the same expression. Test faces always had the same expression as faces seen in the preceding study array. In the latter, singleton condition, one face was expressive whilst the other three were neutral. In both conditions, each face in the study array was equally likely to be presented at test.

Method

Participants.

Participants for both experiments were recruited through Bangor University and received course credit or money in exchange for participation. All reported normal or corrected to normal visual acuity. Twenty-five adults (11 females, mean age = 22 years s.d. = 4 years) participated in Experiment 1.

Apparatus.

In this and the second experiment, stimuli were presented on a Mitsubishi Diamond Plus 20 inch colour monitor (resolution = 1280 x 1024 pixels, refresh rate

= 100 Hz), using a Dell Optiplex GX400 computer running Microsoft Windows XP. Viewing distance was 60 cm. E-Prime software (Version 1.0; Schneider, Eshman, & Zuccolotto, 2002) generated stimuli and recorded responses acquired via a keyboard.

Stimuli.

Face stimuli used in both experiments were grayscale bitmap images (depth 24; sized to 73 x 84 pixels; subtending approximately 2.2 x 2.4 degree of visual angle) of six male adults bearing angry, happy or neutral expressions (Ekman & Friesen, 1976). Hair but not neck was visible. Face images were the same as those used in Jackson et al. (2009) wherein it was reported that these faces are equally discriminable when angry, happy or neutral, and are shown in Appendix C. Study arrays comprised four faces arranged in a grid centred on the screen (as shown in Figure 5.1). Horizontal and vertical spaces separating the faces were approximately 0.2 degrees of visual angle. The test face was presented centrally.

Procedure and Design.

Each trial began with a 3000 ms central fixation interval. Midway through, the fixation cross was briefly enlarged (for 1000 ms) to encourage fixation. The fixation display was followed immediately by a 2000 ms study array of four different faces (drawn randomly without replacement from a set of six face identities), a blank screen for 1000 ms (retention interval), and then the presentation of a single test face that remained visible until response (see Figure 5.1). Response times (RTs) were recorded. Participants indicated by pressing the “i” or “e” key if the test face had been present (no-change) or not (change) in

the study array, respectively. For each condition of the study array, half the trials were change trials and half were no-change trials.

The experimental session comprised a total of 352 trials of which 96 had uniform study arrays (all faces had the same expression) and the remaining 256 trials had study arrays with one expressive face and three neutral faces. The facial expression in the uniform condition was equally likely to be angry, happy or neutral; the expression of the singleton face on singleton trials was equally likely to be angry or happy. These study array conditions were fully crossed with the test conditions of change/no-change and were presented in a random order over eight blocks of 44 trials each. For the uniform conditions, each combination of expression (angry, happy, neutral) and test type (change, no-change) was presented 16 times. For the singleton condition, each combination of singleton expression and test type was presented 64 times. The location of the emotional singleton was fully counterbalanced across these trials. On no-change trials, the test face was equally likely to be any of the faces from the study array. This meant that for singleton trials the test face matched the singleton in identity and expression on 25% of no-change trials, and matched one of the neutral faces in identity and expression on 75% of trials. On change trials, the test face did not match any of the study faces in identity but had the same expression as the singleton on 25% of trials and was neutral on 75% of trials. In total, participants were exposed to 320 angry faces, 320 happy faces and 640 neutral faces. Each face identity was seen approximately the same number of times.

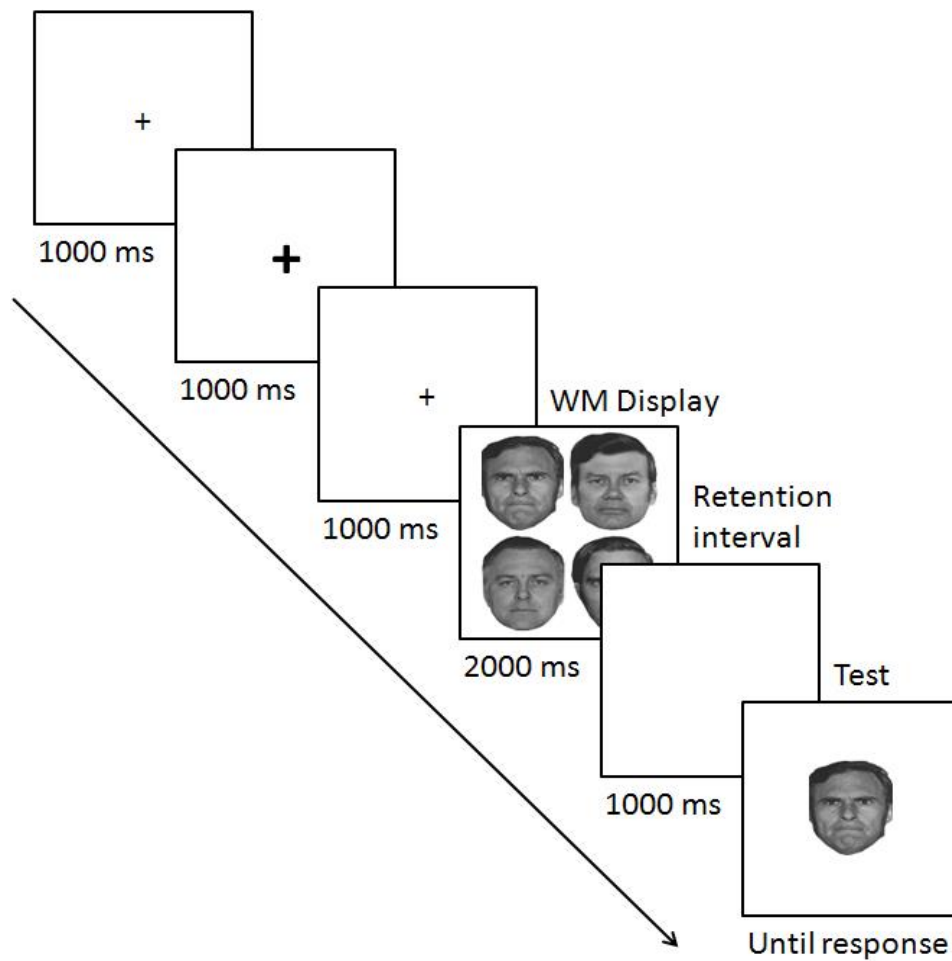


Figure 5.1. An example trial sequence. After a 3000 ms interval, an array of four faces was presented. A single test face appeared after a 1000 ms blank interval. The task was to report whether the test face was one of the faces seen in the preceding study array or not. Images are not drawn to scale.

To suppress verbal encoding and verbal working memory, two letters were presented at the beginning of each block and the participant repeated them aloud (monitored by the experimenter) throughout the block. At the end of the block, two letters were presented and the participant indicated whether they were the same as or different from those they had been repeating. Performance on this task was

perfect barring one error from one participant. A practice block of 12 trials preceded the experimental session.

Data analysis.

Change detection performance was quantified using d' . This was computed for each participant for each condition by subtracting the Z-transform of the proportion of False Alarms (FA; no-change trials on which participants incorrectly responded 'change') from the Z-transform of the proportion of Hits (change trials for which participants correctly responded 'change'). To assess memory for singleton faces in the singleton conditions, d' was calculated using the FA and Hit rates obtained on trials when an emotional test face was presented, using no-change and change trials, respectively. To assess memory for neutral faces in the singleton conditions, d' calculations used the FA and Hit rates obtained on trials when a neutral test face was presented, using no-change and change trials, respectively. An ANOVA using study array condition (happy, angry, neutral) as a within-subject factor was used to analyse the d' scores from the uniform trial conditions. To analyse the d' scores from the singleton trial conditions I conducted an ANOVA using singleton study array expression (happy, angry) and test expression (emotional, neutral) as within-subject factors. Planned comparisons used within-subject t-tests. A similar ANOVA was conducted on RTs. For this analysis I excluded trials with RTs longer than 3000 ms or less than 200 ms, leading to the exclusion of all the data from one participant who had too few trials in some conditions and excluding 5.2% of the data for remaining participants. Alpha levels were set at .05.

Results and Discussion

WM Performance: Uniform conditions.

As can be seen in the leftmost bars of Figure 5.2, when all faces in the study array had the same expression, group average d' scores were higher in the angry condition ($M = 1.65$, $SD = .87$) than in the happy ($M = 1.14$, $SD = .74$, $t(24) = 2.80$, $p = .010$) or neutral ($M = 1.14$, $SD = .63$; $t(24) = 2.77$, $p = .011$) conditions, $F(2,48) = 5.38$, $p = .008$, $\eta_p^2 = .183$. The difference in performance for the happy versus neutral condition was non-significant [$t(24) = .033$, $p = .974$]. Comparable performance for happy and neutral faces argues against the notion that proactive interference may play a role in this task. Proactive interference occurs when current stimulus processing is impaired by prior processing of the same stimuli and is more likely to happen when the test item matches an item from a memory set on a previous trial (see Hartshorne, 2008). In the task reported here, stimuli from a small set of faces (18 in total: 6 identities x 3 expressions) were repeated across many trials, and it is therefore possible that greater repetition of neutral than emotional faces from trial to trial led to increased proactive interference. However, even though happy faces were seen less often than neutral faces throughout the entire experiment and would therefore have incurred less proactive interference, WM performance for happy versus neutral faces in the uniform conditions did not differ. Happy and angry faces were seen equally often and should have produced the same amount of proactive interference, yet I observed a significant difference in performance between these two conditions. These results replicate Jackson et al. (2009) in all respects, and do not support a proactive interference account of the negative face benefit to WM.

WM Performance: Singleton conditions.

When the WM study array comprised only one emotional face and three neutral faces, a singleton benefit (i.e., d' for singleton faces minus d' for neutral faces) was clearly observed. Group mean d' was significantly better for the emotional singleton than for any one of the neutral faces [$F(1,24) = 4.90$, $p = .037$, $\eta_p^2 = .169$] (see Figure 5.2). However, this effect was largely driven by results from the angry singleton condition: I observed a significant interaction of singleton expression in the study array (happy, angry) and expression in the test face; $F(1,24) = 4.65$, $p = .041$, $\eta_p^2 = .162$. When the angry singleton was tested, mean d' ($M = 1.63$, $SD = .56$) was 0.41 units higher [$t(24) = 3.10$, $p = .005$] than when a neutral face from this study array condition was tested ($M = 1.22$, $SD = .54$). In contrast, when the happy singleton was tested, mean d' ($M = 1.31$, $SD = .66$) was only 0.04 units higher (a non-significant difference, $p > .5$) than when a neutral face ($M = 1.27$, $SD = .55$) from the same study array condition was tested. The difference in the magnitude of angry and happy singleton effects yields an index of the negative face benefit that is distinct from singleton effects per se (e.g., Theeuwes, 1993; Yantis, 2000) and reflects an emotion-specific process. The negative face benefit observed here (mean = 0.37 d' units) is further supported by the observation that WM performance for the angry singleton was 0.32 d' units higher than for that for the happy singleton [$t(24) = 2.76$, $p = .011$]. These effects cannot be explained by differences in proactive interference for angry versus happy stimuli because all emotional stimuli were viewed an equal number of times.

WM for neutral faces was unaffected by emotional expression of the accompanying singleton ($p > .5$). WM performance for the neutral faces in each of

the two singleton conditions was not different from that in the neutral uniform condition (both cases, $p > .5$). The boost in performance found for angry versus happy singletons does not appear to have had any consequences (facilitatory or inhibitory) for concurrently viewed neutral faces, indicating that this effect is object-specific, not scene-wide.

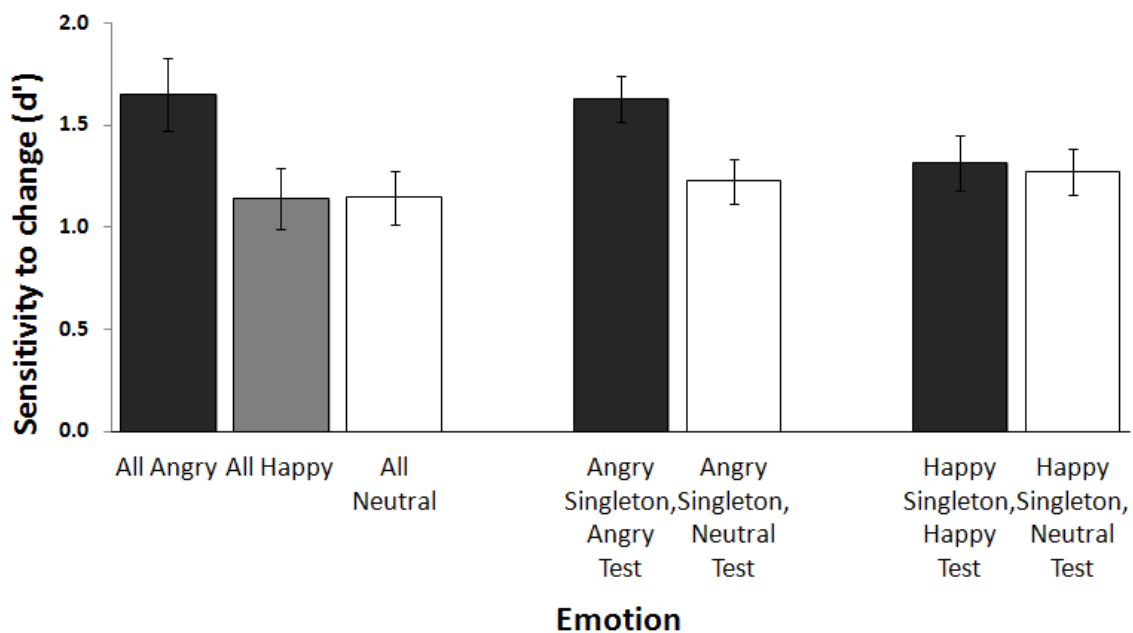


Figure 5.2. Mean d' scores for the experimental conditions of Experiment 1. The three leftmost bars represent average WM performance when all faces in the study array had the same expression (angry, happy, or neutral). The four rightmost bars represent average WM performance when one face in the study array was expressive (angry or happy) and the remaining faces were neutral. Error bars represent ± 1 normalised standard error of the mean (Cousineau, 2005; see Appendix D).

I then asked whether the magnitude of the boost from a single angry face in a study array was different from that obtained when all faces in the study array had angry expressions. Although the latter condition provided a larger overall signal of threat by presenting four angry faces in the study array, WM for angry singletons was not significantly different from WM for faces tested in the angry uniform condition ($p > .5$). Thus it appears that a singleton angry face accrues the same magnitude of benefit to its representation in WM as an angry face seen within an array of other angry faces. A similar comparison for the happy conditions also showed a non-significant difference ($p > .25$).

Response Time.

If selective attention plays a role in WM retrieval (Griffin & Nobre, 2003), then the singleton face in WM should have biased attention during retrieval, making RTs in the WM task faster when these versus neutral faces were probed. Moreover, this effect should be greater with angry versus happy faces, if indeed, angry faces attract attention more than happy faces. I observed that RTs were 45 ms faster when the test face had an emotional versus neutral expression [$F(1, 23) = 17.227, p < .001, \eta_p^2 = .428$], indicating a singleton effect that could have arisen from selective attention to the singleton item stored in WM. The main effect of singleton condition (angry versus happy in WM array) was non-significant ($F < 1$) and this factor did not interact significantly with expression in the test face ($F < 1$), lending no support for the notion that a specifically threat-related attentional bias operating at retrieval could account for the angry benefit. (See Table 5.1 for RT means.)

		Exp 1	Exp 2
Angry Singleton	Angry Test	1127 (49)	1064 (33)
	Neutral Test	1175 (63)	1151(37)
Happy Singleton	Happy Test	1123 (59)	1079 (39)
	Neutral Test	1165 (54)	1147 (38)

Table 5.1. Mean RTs in the WM task for Experiments 1 and 2. SE in parentheses.

EXPERIMENT 2

The primary findings of Experiment 1 do not readily support the notion that the angry singleton drew resources away from concurrently presented neutral stimuli in this task. However, it remains a possibility that an attention-based reduction in processing of neutral faces in angry singleton-containing arrays occurred in tandem with a scene-wide boost in processing resulting from the presence of the angry singleton. Acting in opposite directions, these processes could have cancelled each other out, resulting in our observed null effect on WM for neutral faces. To investigate, I re-tested the singleton conditions of Experiment 1, (using the same stimuli) but this time monitoring eye movements during the study array presentations. Considering gaze to be a reasonable index of attention in this situation, I predicted that if the angry singleton was capturing attention to a greater extent than the other faces in its array, then this should be evident as longer or more fixations to that face than to the concurrently viewed neutral faces. Importantly, if greater selective attention to angry than happy faces during encoding was the basis for the negative face benefit in WM performance, this bias

of gaze toward the angry singleton should be greater than any corresponding bias toward the happy singleton.

Method

Participants.

Thirty-two adults (21 females, mean age = 22 years, s.d. = 4.1 years) participated in Experiment 2. None had participated in Experiment 1.

Apparatus.

This was that same as for Experiment 1 except that an EyeLink® 1000 desktop mounted eye-tracker (SR Research Ltd., 2002; sampling frequency of 500 Hz) was used to record eye movements monocularly from the left eye. A chinrest was used to maintain head position.

Procedure and Design.

The procedure was that same as that described for the singleton trials of Experiment 1, except that the verbal suppression task was omitted. Four blocks of 64 trials each were presented. The singleton was equally likely to be angry or happy, to occur in any of the four face locations in the study array, and to be followed by a change or no-change test condition. The sequence of trials presented within each block was individually randomized. A practice block of 16 trials preceded the experimental session. The eye-tracker was calibrated (using standard nine point calibration/validation) before each block of trials.

Data Analysis.

Data from trials with response times (RT) of less than 200 ms and longer than 3000 ms were excluded from the analysis accounting for 3.24% of the data. Eye position was extracted using the EyeLink® Data Viewer (SR Research Ltd., 2002) software package. Fixations were defined as periods when the pupil could be detected and no saccade was underway. Saccades were defined as periods when two criteria were met: gaze position had both changed by more than 0.1 deg and was either accelerating by at least $8000^{\circ}/\text{sec}^2$ or exceeded a velocity of $30^{\circ}/\text{sec}$. All the data from two participants were excluded due to difficulty monitoring their gaze position due to thick spectacles. For remaining participants, WM performance data were excluded on a trial by trial basis if less than two fixations were successfully recorded during presentation of the study array, a process that eliminated a further 0.95% of data. Four vertical rectangular (79 X 90 pixels) regions of interest (ROI) were defined; their centers were coincident with the center of each face in the study array. The mean number of fixations, mean dwell time (i.e., sum of durations of all fixations), and proportion of first fixations landing within each ROI were determined for each participant for each study array condition. ANOVAs were conducted on each eye movement measure using study array condition (angry, happy) and test expression (expressive, neutral) as within-subjects factors. ANOVAs on RT and d' were calculated as for singleton conditions in Experiment 1.

Results & Discussion

WM Performance.

As can be seen in Figure 5.3, the pattern of results replicated that obtained in the singleton conditions of Experiment 1²⁵. As before, there was a significant main effect of test expression (emotional, neutral) [$F(1, 31) = 64.964$, $p < .001$, $\eta_p^2 = .677$] indicating a singleton benefit. The main effect of singleton expression (angry, happy) [$F(1, 31) = 5.191$, $p = .030$, $\eta_p^2 = .143$] was also significant. Importantly, as evidence of a negative face benefit, the interaction of test expression and singleton expression was also significant [$F(1, 31) = 10.821$, $p = .003$, $\eta_p^2 = .259$]. Angry singletons ($M = 2.09$, $SD = 0.71$) were remembered markedly better (by 1.05 d' units) than neutral faces in the same study array condition ($M = 1.04$, $SD = 0.51$; $t(31) = 9.616$, $p < .001$). WM performance for happy singletons ($M = 1.64$, $SD = 0.81$) was also better (by 0.51 d' units) than that for neutral faces presented in the same study array ($M = 1.13$, $SD = 0.52$; $t(31) = 3.585$, $p < .001$). Subtracting this happy singleton effect from the angry singleton effect yields a negative face benefit of 0.44 d' units, a value comparable to the 0.37 d' units measured in Experiment 1. As before, performance was significantly better for angry versus happy singletons [$t(31) = 3.082$, $p = .004$], and differences in performance for neutral faces presented in the angry versus happy study array conditions were non-significant [$t(31) = 1.311$, $p = .200$].

²⁵ I conducted an ANOVA on the WM d' scores for the singleton conditions using experiment as a between group factor and study array expression and test expression as within-subjects factors. The main effect of experiment, the interaction of experiment and study array expression, and the triple interaction (experiment X study array expression X test expression) were all non-significant (all F 's < 1). However, the interaction of experiment and test expression was significant [$F(1, 55) = 15.552$, $p < .001$, $\eta_p^2 = .220$], indicating that the singleton effect was greater in the second experiment, an effect largely due to the angry singleton condition [$t(55) = 2.727$, $p = .01$]. The corresponding difference for happy singletons was marginally significant [$t(55) = 1.668$, $p = .10$]. Between-experiment differences in WM for the neutral face conditions were non-significant ($p > .30$).

Response Time.

As for Experiment 1, I analysed the RTs for each condition and found a similar pattern of results (see Table 5.1). RTs were 77 ms faster when the test face had an emotional versus neutral expression [$F(1, 31) = 23.883, p < .001, \eta_p^2 = .435$]. As before, this effect did not interact significantly with singleton condition (angry, happy; $F < 1$), lending no support for the notion that a specifically threat-related attentional bias could account for the angry benefit. The interaction term was also non-significant ($F < 1$).

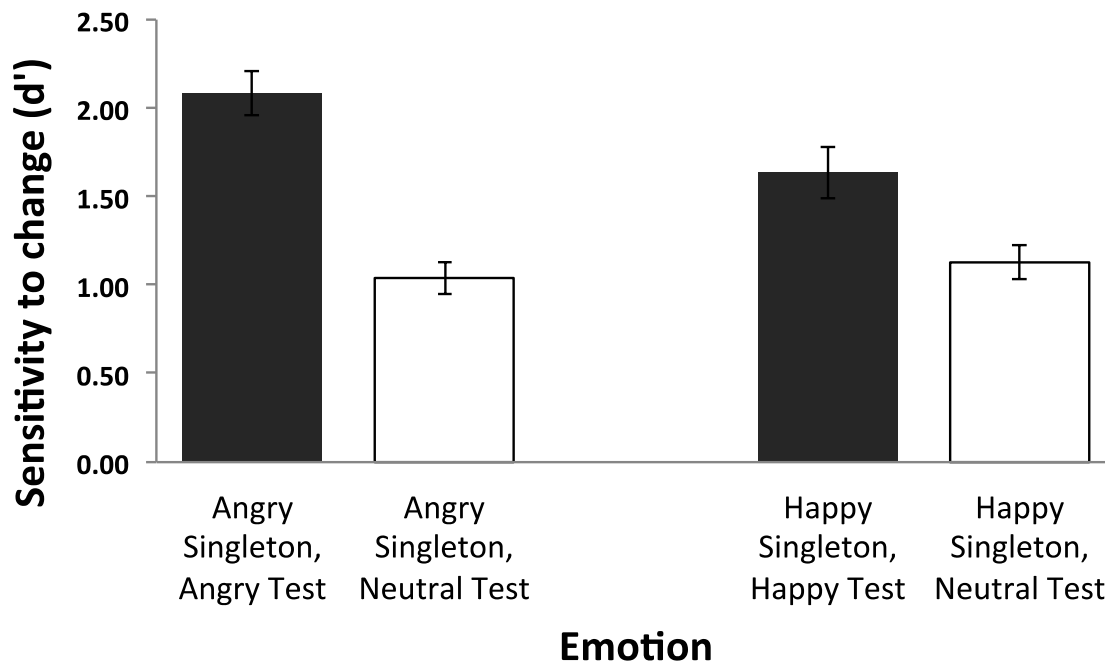


Figure 5.3. Mean d' scores for the experimental conditions of Experiment 2. Black bars represent performance when memory for the emotional singleton was tested and white bars represent performance when memory for a neutral face was tested. Error bars represent ± 1 normalised standard error of the mean (Cousineau, 2005; see Appendix D)

Eye-movements.

The singleton face was fixated more frequently than any other single neutral face [$F(1, 31) = 12.296, p < .001, \eta_p^2 = .284$] and was also looked at longer [$F(1, 31) = 14.393, p = .001, \eta_p^2 = .317$]. See Table 5.2 for group mean eye movement measures for each condition. However, this effect did not interact with the emotion of the singleton in the display for either measure [fixations: $F(1, 31) = 1.256, p = .271, \eta_p^2 = .039$, dwell time: $F < 1$], and the main effect of singleton condition (angry, happy) was non-significant for each measure [both F 's < 1].

Condition	Number of fixations (SD)	Dwell time, ms (SD)
Angry Singleton	1.60 (0.21)	477 (66)
Happy Singleton	1.61 (0.26)	483 (78)
Neutral (with Angry)	1.52 (0.16)	436 (49)
Neutral (with Happy)	1.50 (0.16)	436 (52)

Table 5.2. Group mean eye movement data measured during the study interval of Experiment 2. Values shown for neutral faces are averaged across all the three neutral faces seen in each study array.

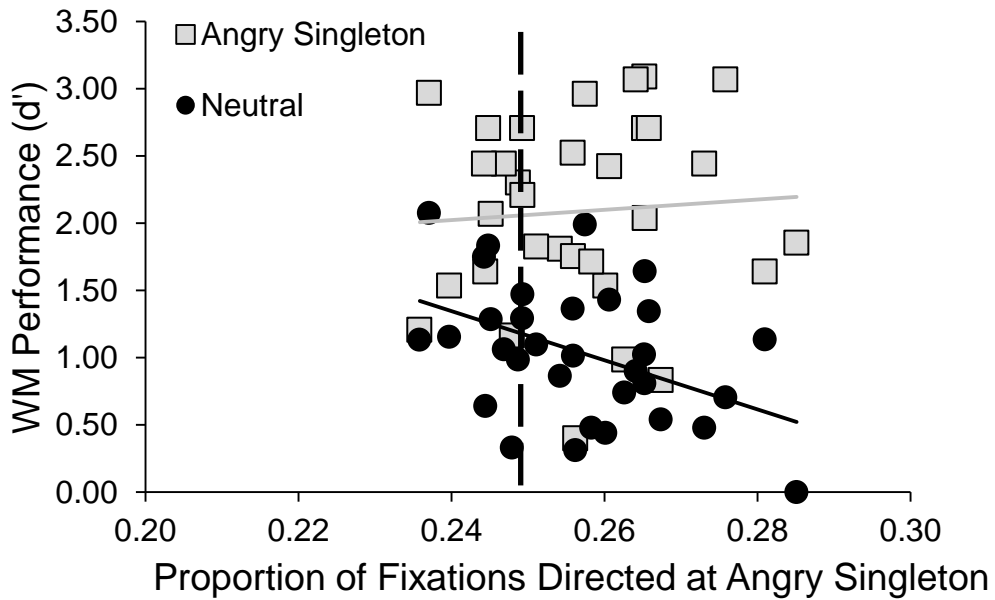
Thus, while singletons attracted more looking, their emotional valence did not influence this behaviour. Pertinent here is that the correlation between each individual's difference in the number of fixations for angry versus happy singletons and their difference in WM d' for angry versus happy singletons was 0.17, suggesting that differences in fixation behaviour had little impact on differences in WM performance resulting from the expression of the singleton. A similar correlation calculated using dwell time was 0.09.

I next examined correlations between eye movements and WM performance using the number of fixations directed at the singleton face (happy or angry) expressed as a proportion of the total number of fixations directed at faces during the study interval. I used this eye movement measure²⁶ because it can be viewed as an index of an inappropriate orienting bias. Considering that all faces were equally likely to be tested, orienting to the singleton face more than any other face would be non-optimal, presumably resulting from poor top-down control over task-independent biases arising from the singleton face. I found that the proportion of fixations to the angry singleton was not significantly correlated with d' for the angry face ($r = .07$). The corresponding correlation for the happy singleton was also non-significant ($r = .00$). However, the proportion of fixations to the angry singleton was significantly and negatively correlated with the d' value for neutral faces seen in the angry singleton condition ($r = -.43$, $p = .01$). The corresponding correlation for the happy singleton condition was also negative but non-significant ($r = -.17$, ns.). Scatter plots illustrating these effects (and the distribution of d' scores) for individual participants are shown in Figures 5.4A and 5.4B. There it

²⁶ I redid all the analyses reported here using (1) raw number of face fixations that were directed at the angry face, (2) dwell time to the angry face, and (3) proportion of total dwell time spent looking at the angry face. In all cases, the results were as reported here using proportion of singleton fixations regardless of singleton expression.

can be seen that a handful of participants showed a propensity to disproportionately fixate the singleton, a behaviour associated with poor memory for the less fixated neutral faces.

A.



B.




Figure 5.4. Individual mean WM (d') scores obtained in Experiment 2 when the test stimulus was an emotional (gray squares) or neutral (black circles) face plotted as a function of the proportion of fixations directed toward the emotional singleton face versus other neutral faces in the four-face study array. Scatter plots for the angry face (A) and happy face (B) singleton condition are shown separately. The dashed vertical lines shows the proportion of fixations expected (.25) if gaze was used to optimize encoding an array of four faces and emotion was treated as irrelevant. Linear regression lines for each data set are shown as solid lines.

To probe whether this relationship between fixation behaviour and WM performance could account for the WM boost seen for angry relative to happy faces and singletons relative to neutrals, I split the participants into two groups (using the median) depending on the proportion of fixations made toward the singleton faces (averaging scores for angry and happy singleton conditions). I found that the group with the smaller, more optimal proportion of singleton fixations (mean = .249, where .250 is optimal for a four-face display) did overall better than the group with more singleton fixations (sub-optimal fixators; mean = .265). (Overall $d' = 1.40$ and 1.05 , for optimal and sub-optimal groups respectively; between-group t-test, $p = .02$). The correlation between the proportion of fixations directed at the singleton faces and the overall session mean d' was negative and significant ($r = -.413$, $p < .02$). It remains unclear however whether eye movements determine overall working memory performance or whether participants with poorer WM have more difficulty suppressing inappropriate orienting responses to emotional stimuli.

Of interest here is that the optimal versus sub-optimal fixators did not differ significantly on d' performance for singletons, but the optimal fixators outperformed the sub-optimal group when memory for neutral faces was tested, $p < .05$. This result shows that memory for the neutral faces depended on time spent orienting to them. For the optimal fixators, who did not fixate the singleton face disproportionately to the other neutral faces, memory performance for the angry singleton ($d' = 2.12$) was nevertheless significantly better than for the accompanying neutral faces ($d' = 1.23$, $p < .001$), showing clearly that the negative face benefit to WM did not rely on greater overt orienting to this face. The corresponding difference for the happy singleton condition was only marginally significant (happy singleton, $d' = 1.69$; neutral faces, $d' = 1.36$; $p = .09$), suggesting that benefits accrued from singleton status, per se, may depend on biased overt orienting to the singleton. Indeed the sub-optimal fixators who showed this behaviour showed large significant singleton benefits for both angry and happy singletons (versus neutrals, $p < .001$, in both cases). Their negative face benefit (group mean angry minus happy singleton d' scores = 0.46,) was similar in size to that found in the optimal fixators (mean = 0.43 d' units). To summarise, biased orienting may play a role in promoting a singleton benefit, but does not appear to contribute to the negative face benefit. Even when participants who obviously favoured orienting to singletons are removed from the analysis, the negative face benefit to WM was still observed.

I next analysed the distribution of initial fixations made after the study array was presented to determine if the emotional nature of the singleton had an immediate effect on gaze. Angry singletons attracted the first fixation on 26% of angry singleton trials ($SD = 4\%$), whereas on happy singleton trials, the singletons

attracted the first fixation on 28% of trials (SD = 4%). Both values are modestly greater than chance (25%), suggesting that there may have been an attentional bias toward the emotional singletons during the early stage of encoding.

However, happy singletons were more effective at capturing the first fixation than angry singletons [$t(31) = 2.037, p = .050$], arguing against the notion that preferential orienting to the angry face singleton accounts for their advantage in WM.

Although this analysis of eye movements during the presentation of the study array indicates that the emotional singletons were modestly more effective at attracting gaze than their neutral counterparts and this effect was more pronounced in some individuals than others, neither fixation frequency nor dwell time were influenced by emotional expression. These findings lend no support for the notion that the boost in WM performance for the angry singleton found here and in Experiment 1 derives from an enhanced opportunity for encoding. Interestingly, our eye movement results do not replicate those reported by Becker and Detweiler-Bedell (2009) who used a very similar display (an emotional singleton and three neutral faces displayed in a grid, as used here) in a passive viewing task. In contrast to our finding, they showed that angry (and fearful) faces are less likely to be the target of the first fixation than any of the other neutral faces in the array. Perhaps the difference in our results is that here, each face was relevant for performance in the task whereas in their study, there was no specific reason to control visual orienting. Nevertheless, in neither data set is there any evidence of preferential overt visual orienting to angry versus happy singletons.

CHAPTER DISCUSSION

The aim of the experiments reported here was to determine how stimuli affording unusually strong task-independent attention biases affect WM for themselves and for other equally task-relevant, but less compelling objects in the same scene. Specifically, I asked whether positive versus negative emotional content in a face would affect WM for it and other affectively neutral faces. In two experiments I show that an emotional singleton face is remembered better in WM than other concurrently viewed neutral faces, and that this effect is significantly stronger when the singleton face has a negative versus positive expression. The results of Experiment 2 showed that these effects were found even when eye movements did not favour the singleton during the study period. Surprisingly, I also found that the benefit accrued to the emotional singleton appeared to have no negative consequence for WM of the neutral faces that accompanied them. Performance for a neutral face was the same regardless of whether it was accompanied by a happy singleton, an angry singleton, or no singleton (i.e., all faces in the study array were neutral). In summary, there are three effects from these studies that require explanation: A singleton benefit that is independent of singleton emotion; a negative face benefit (defined here as the difference in the size of the singleton benefit measured with angry versus happy singletons); and the independence of neutral face WM from the attention or emotional salience of concurrently viewed faces.

Finding a singleton benefit to WM is most easily explained as a straightforward consequence of selective attention (Theeuwes, 1993; Yantis, 2000). The singleton face, a singleton by virtue of its expression only (all faces were male but all faces had unique identities), appears to have afforded sufficient

sensory and/or emotional salience to bias attentional competition at some stage in the processes that determine WM performance. This effect is probably indexed best in the current studies as the difference between singleton face d' and neutral face d' in the happy singleton condition only because here effects are not augmented by the negative face benefit. The eye movement data of Experiment 2 suggests that singletons attracted modestly more overt orienting than non-singletons, raising the possibility that these eye movement biases can account for the singleton effect. However, participants who did not show an orienting bias to the singletons still produced 0.33 d' units of singleton benefit in the happy singleton condition, suggesting that biased orienting leading to better encoding cannot fully account for this effect. Another, not mutually exclusive account, is that selective attention facilitated retrieval processes during the WM task. An emotional face at test would serve to bias attention to the only emotional face in WM, whereas a neutral face at test would bias attention less specifically, requiring potentially more than one representation to be examined. This notion is supported by the finding that RTs on the change detection task were faster when singleton versus non-singleton items were tested. This effect was found for both angry and happy singleton conditions and was present for optimal and non-optimal fixators under both conditions (all p 's < .05). This RT advantage indicates that the expression information that made the singleton unique was available to WM and could be used to aid search through WM items. In summary, these data provide good evidence for the contention that selective attention is the mechanism mediating the singleton benefit to WM.

The second effect demanding explanation is the negative face benefit. As with the singleton effect, it is tempting to assume that this occurs in the singleton

conditions of Experiment 1 and in Experiment 2 as a result of threat biases acting on selective attention (e.g. Eastwood et al., 2003). However, problematic for this explanation is our finding that the singleton's emotional expression had no obvious impact on gaze fixations patterns during the study interval, nor appeared to have any influence on RT in the WM task. If angry faces had been more attentionally compelling than happy faces, then I would have expected to see more fixations toward angry versus happy singletons and faster RTs when angry versus happy faces were presented at test, but I did not. Moreover, in Experiment 1 I found that the negative face benefit was robust even when, at study, all faces were negative versus all happy in expression (see also Jackson et al., 2009, Sessa et al., 2011). This uniform array condition eliminates any potential for selective bias among memory items. Furthermore, finding a negative face benefit with uniform arrays excludes explanations based on item distinctiveness or semantic relatedness, factors thought to provide a cognitive basis for emotional enhancement of memory (EEM) effects in the short term (Talmi & McGarry, 2012). Another finding in the data that is problematic for a selective attention account of the negative face benefit found in the singleton condition is that WM performance for neutral faces was not affected by the singleton's emotional expression. This is the third interesting effect from our experiments that needs explaining. Accounting for it and developing an explanation for the negative face benefit are intertwined so will now be discussed together.

Different theories addressing the impact of emotional, especially threatening stimuli on perception and memory make different predictions for the fate of an accompanying, task-relevant but affectively neutral object in WM. Arguing for a benefit are notions of generalised arousal and recent findings that

scene-wide cognitive processing can be transiently boosted by the presentation of motivationally relevant (target) stimulus (e.g., Swallow & Jiang, 2010; Makovski et al., 2011). Predicting a detrimental effect are theories proposing a role for biased competition in memory (e.g., Mather & Sutherland, 2011) and a range of findings related to EEM effects showing that emotional stimuli promote a narrowing of processing, leading to poorer processing of stimuli peripheral to the emotional epicentre of a scene (Easterbrook, 1959). However, I found no support for either prediction; instead, WM for neutral faces was the same whether a single angry or happy face appeared in the scene or all faces were neutral. Although interpretation of null effects may be difficult when the method of measurement has questionable sensitivity, observing replicable expression-dependent effects on WM for singletons in the same experiments on the same participants in which these null effects were found suggests that the change detection task (as used here) is, in fact, sufficiently sensitive to permit cautious interpretation.

An explanation for why angry faces enhance memory with no obvious cost to concurrently viewed neutral stimuli must be rooted in object-based processes, i.e., processing that operates on stimulus information after it has been used to formulate representations of distinct objects. It cannot be explained by scene-wide processes, such as generalised arousal induced in response to emotional stimuli, because this type of process should have modulated WM for all stimuli in the same way. Also arguing against a generalised arousal account of the negative face benefit is that the magnitude of this effect obtained in the singleton condition of Experiment 1 was not different from that observed in the all-angry uniform condition of the same experiment. Arousal would be expected to be greater with four angry faces than with just one; if arousal accounted for the

negative face benefit, then the magnitude of the WM boost would be expected to reflect this. My rejection of a generalized arousal account is consistent with Jackson et al. (2009) who showed that music induced arousal (as measured using the Self-Assessment Manikin, Bradley & Lang, 1994) did not improve WM for non-angry faces, that calming music did not abolish the angry face benefit, and that explicit ratings of arousal for the angry faces used in their (and the current) study were not different from those for happy faces. It is possible that arousal reached some maximum level with exposure to a single face, precluding the observation of greater effects with additional threatening faces. However, this seems unlikely given the relatively low arousal nature of face photographs viewed on a computer screen, as used here.

However, another possibility is that arousal was elevated at retrieval. In the current experiments, test faces always matched the study array face in expressions so that the angry face singleton was tested with an angry face and neutral faces were tested with neutral faces. If arousal was elevated at retrieval by the former but not by the latter, and more so by angry than happy faces, this could account for what appears to be an object-based effect on WM performance. However, this is unlikely because Jackson, Linden & Raymond (in press), using a similar WM single-probe change detection procedure, report that when study array faces are emotional but the test face is neutral, WM for face identity is still better when the faces seen at study are angry versus happy. They also showed that WM for neutral study faces is not better when an angry versus happy expression is seen in the test face.

Nevertheless, arousal has long been thought to play a role in producing the well-studied EEM effect on LTM, and this rich literature may provide a clue to

explain the object-based negative face advantage effects we describe here. EEM effects are widely thought to be mediated by activity in the basolateral amygdala (BLA, see McGaugh, 2004 for a review), a structure that is able to modulate slow acting memory consolidation processes in multiple brain locations. Although these same processes are unlikely to mediate short-term memory effects, the heightened amygdala responses they depend on may also have immediate effects on perceptual processes that could mediate the emotion-specific WM effects reported here. The amygdala produces rapid responses (within a half second) to emotional (Brieter et al., 1996) and motivationally-salient (Paton, Belova, Morrison & Salzman, 2006) stimuli. The dense reciprocal interactions between visual cortex and amygdala (Amaral & Price, 1984; Stefanacci & Price, 2002) allow the possibility that amygdala activation could boost visual encoding of the angry singleton face during the lengthy (2 s) study intervals used here. This would lead to richer, more precise sensory encoding of angry versus happy or neutral faces and thus yield direct benefits for WM in our change detection task (Sessa et al., 2011). Our data show that this object-specific boost did not come at a cost for other items in the display. Although object-based models of selective attention might predict that such costs would be evident, the eye movement analysis of Experiment 2 suggests that during the long study interval, selective attention was moved serially to each face in the study array, allowing any costs for encoding that may have arisen from attentional competition to be overcome. Thus, each face would have successively become the focus of attention, but only the angry face would have attracted high resolution encoding due to its ability to heighten amygdala activation.

A question that arises is why this effect is anger-specific, especially when other studies have shown that happy faces are also able to heighten amygdala activation (Brieter et al., 1996). One possibility is that the amygdala does not specifically code emotional content, but rather codes motivational salience (Holland & Gallagher, 2004; Paton et al, 2006), a factor that was confounded with face expression under the passive viewing conditions used by Brieter et al. (1996). Perhaps angry faces had greater motivational salience than happy faces in our study because they signal disapproval and carry an implicit instruction to alter behaviour or “do better” to avoid unpleasant consequences, as opposed to the approval or signal to maintain current behavioural strategies implied by a smiling face. This motivational interpretation of expression information is consistent with the view that anger in another’s face or gesture activates approach responses (Carver & Harmon-Jones, 2009; Wilkowski & Meier, 2010) and also with a recent finding that an angry but not a happy or fearful expression can overcome attention-induced inhibition that dulls sensitivity to subtle facial expression (Gomez-Cuerva & Raymond, 2011). Additional support for the idea that motivation may play a role in WM was reported by Raymond & Thomas (2012; see Chapter 7) who showed that stimuli previously associated with high reward value were better remembered in a WM task than stimuli previously associated with losses or no outcome.

The finding that WM can be boosted for one item in a scene without incurring costs or benefits to other concurrently viewed items also has implications for current theories concerning capacity limitations of visual WM. WM capacity for faces is about 2.5 on average (Jackson & Raymond, 2008), so even though my experiments only used four faces in each study array, WM capacity was probably

exceeded for all individuals, allowing me to make inferences about capacity limitations from the data. Some argue that WM capacity is limited by a fixed number of 'slots' and that only a single object representation along with all its attributes can be encoded in each slot (Luck et al., 1997; Zhang & Luck, 2008). Others posit that capacity is constrained by the size of a pool of available resources that can be flexibly allocated to a small or large number of items by trading off representational precision (Bays et al., 2008; Frick, 1988; Wilken & Ma, 2004). Finding that WM for neutral faces was unaffected by the presence of an emotional singleton is inconsistent with a purely resource-based model. It is perhaps better accommodated by recent two-factor views that posit an initial slot-based process accompanied by a post-selection, flexible adjustment of representational precision within each slot (Awh, Barton & Vogel, 2007; Xu & Chun, 2009; Barton, Ester, & Awh, 2009; Fukuda, Awh, & Vogel, 2010).

In summary, I show that negative emotional content in face images enhances visual WM for those stimuli, without causing either costs or benefits to other neutral, concurrently viewed stimuli. These findings, especially when considered in light of the eye movement patterns observed during the study period and the RTs in the change detection task, indicate that selective attention with its putative mechanism of competition among stimulus representations cannot account for the negative face advantage found here and in previous studies. Instead, it appears that negative emotional information may serve to heighten encoding precision for an object's representation thereby enhancing WM performance for it but not other stimuli.

References – Chapter 5

- Amaral, D. G. & Price, J. L. (1984). Amygdalo-cortical projections in the monkey (*Macaca fascicularis*). *Journal of Comparative Neurology*, 230, 465-469.
- Anderson, B. A., Laurent, P. A., & Yantis, S. (2011). Value-driven attentional capture. *Proc. Nat. Acad. Sci. USA*, 108, 10367-10371.
- Awh, E., Barton, B., & Vogel, E. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, 18(7), 622-628.
- Barton, B., Ester E. F., and Awh, E. (2009). Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 35(5), 1359-1367.
- Bays, P. M., & Husain, M. (2008). Dynamic Shifts in Limited Working Memory Resources in Human *Journal of Vision Science*, 321, 851-854.
- Becker, M. W. (2009). Panic search: fear produces efficient visual search for non-threatening objects. *Psychological Science*, 20(4), 435-437.
- Becker, M. W., & Detweiler-Bedell, B. (2009). Early detection and avoidance of threatening faces during passive viewing. *The Quarterly Journal of Experimental Psychology*, 62(7), 1257-1264.
- Bradley, M. M. & Lang, P. J. (1994). Measuring emotion: The self-assessment manikin and the semantic differential. *Journal of Behavioral Therapy and Experimental Psychiatry*, 25(1), 49-59.
- Breiter, H. C., Etcoff, N. L., Whalen, P. J., Kennedy, W. A., Rauch, S. L., Buckner, R. L., Strauss, M. M., Hyman, S. E., & Rosen, B. R. (1996). Response and habituation of the human amygdala during visual processing of facial expression. *Neuron*, 17, 875-887.
- Bundesen, C. (1990). A theory of visual-attention. *Psychological Review*, 97, 523-547.
- Carver, C. S., & Harmon-Jones, E. (2009). Anger is an approach-related affect: Evidence and implications. *Psychological Bulletin*, 135, 183-204.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences*, 24, 87-185.
- Curby, K.M. & Gauthier, I. (2007). A visual short-term memory advantage for faces. *Psychonomic Bulletin & Review*, 14(4), 620-628
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18, 193-222.

- Duncan, J. (1984). Selective Attention and the Organization of Visual Information. *Journal of Experimental Psychology: General*, 113(4), 501-517.
- Easterbrook, J. A. (1959). The effect of emotion on cue utilization and the organization of behavior. *Psychological Review*, 66, 183-201.
- Eastwood, J. D., Smilek, D., & Merikle, P. M. (2003). Negative facial expression captures attention and disrupts performance. *Perception & Psychophysics*, 65(3), 352–358.
- Ekman, P., & Friesen, W. V. (1976). *Pictures of facial affect*. Palo Alto, CA: Consulting Psychologists Press.
- Feldmann-Wüstefeld, T., Schmidt-Daffy, M., & Schubö, A. (2011). Neural evidence for the threat detection advantage: Differential attention allocation to angry and happy faces. *Psychophysiology*, 48(5), 697-707
doi: 10.1111/j.1469-8986.2010.01130.x
- Fox, E., & Damjanovic, L. (2006). The eyes are sufficient to produce a threat superiority effect. *Emotion*, 6(3), 534–539.
- Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition and Emotion*, 16(3), 355-379.
- Frick, R. W. (1988). Issues of representation and limited capacity in the visuospatial sketchpad. *British Journal of Psychology*, 79, 289–308.
- Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current Opinion in Neurobiology*, 20, 177-182.
- Gomez-Cuerva, J., & Raymond, J. E. (2011). Perception of facial expression depends on prior attention. *Psychonomic Bulletin & Review*, 18 (6), 1057-1063. DOI 10.3758/s13423-011-0170-1.
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience*, 15(8), 1176 –1194.
- Hahn, S., Carlson, C., Singer, S., & Gronlund, S. D. (2006). Ageing and visual search: Automatic and controlled attentional bias to threat faces. *Acta Psychologica*, 123, 312-336.
- Hartshorne, J. K. (2008). Visual working memory capacity and proactive interference. *PLoS ONE*, 3(7), e2716.
- Huang, S-L., Chang, Y-C., & Chen, Y-J. (2011). Task-irrelevant angry faces capture attention in visual search while modulated by resources. *Emotion*, 11(3), 544-552.
- Hyun, J., Woodman, G. F., Vogel, E. K. Hollingworth, A., Luck, S. J. (2009). The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35 (4), 1140-1160.

- Jackson, M. C., Linden, D. E. J., & Raymond, J. E. (in press). Angry expressions strengthen the encoding and maintenance of face identity representations in visual working memory, *Cognition & Emotion*.
- Jackson, M. C., & Raymond, J. E. (2008). Familiarity enhances visual working memory for faces. *Journal of Experimental Psychology: Human Perception and Performance*, 34 (3), 556-568.
- Jackson, M. C., Wu, C., Linden, D. E. J., & Raymond, J. E. (2009). Enhanced Visual Short-Term Memory for Angry Faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(2), 363-374.
- Kastner, S. & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Review of Neuroscience*, 23, 315–341.
- Kristjánsson, A., Óladóttir, B., & Most, S. B. (2012). “Hot” facilitation of “cool” processing: Emotional distraction can enhance priming of visual search. *Journal of Experimental Psychology: Human Perception & Performance* [EPub ahead of print; PMID: 22642218].
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature* 390, 279–281.
- Mather, M. & Sutherland, M. R. (2011). Arousal-biased competition in perception and memory. *Perspectives in Psychological Science*, 6(2), 114-133.
- Makovski, T., Sussman, R., & Jiang, Y. V. (2008). Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 34(2), 369 – 380.
- Makovski, T., Swallow, K. M., & Jiang, Y. V. (2011). Attending to unrelated targets boosts short-term memory for color arrays. *Neuropsychologia*, 49(6), 1498-1505. doi:16/j.neuropsychologia.2010.11.029
- Öhman, A., Flykt, A., & Esteves, F. (2001). Emotion Drives Attention: Detecting the Snake in the Grass. *Journal of Experimental Psychology: General*, 130(3), 466-478.
- Paton, J.J. Belova, M. A. Morrison, S. E. & Salzman, C. D. (2006). The primate amygdala represents the positive and negative value of visual stimuli during learning. *Nature*, 439, 865-870.
- Pessoa, L., McKenna, M., Gutierrez, E., & Ungerleider, L. G. (2002). Neural processing of emotional faces requires attention. *Proceedings of the National Academy of Sciences*, 99, 11458-11463.
- Phelps, E. A., Ling, S., & Carasco, M. (2006). Emotion facilitates perception and potentiates the perceptual benefits of attention. *Psychological Science*, 17(4), 292-299.
- Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology*, 61(3), 380-391.

- Raymond, J. E., & O'Brien, J. L. (2009). Selective visual attention and motivation: the consequences of value learning in an attentional blink task. *Psychological Science, 20*(8), 981-988.
- Raymond, J.E. & Thomas, P. M. J. (2012). Reward associations facilitate visual working memory. *Abstracts of the Psychonomics Society Meeting, 17*, 38.
- Schneider, W. X., Eshman, A., & Zuccolotto, A. (2002). *Eprime User's Guide*. Pittsburgh, PA: Psychology Software Tools, Inc.
- Sessa, P., Luria, R., Gotler, A., Jolicoeur, P., & Dell'Acqua, R. (2011). Interhemispheric ERP asymmetries over inferior parietal cortex reveal differential visual working memory maintenance for fearful versus neutral facial identities. *Psychophysiology, 48*, 187-197. DOI: 10.1111/j.1469-8986.2010.01046.x
- Stefanacci, L. & Amaral, D. G. (2002). Some observations on cortical inputs to the macaque monkey amygdala: An anterograde tracing study. *Journal of Comparative Neurology, 451*(4), 301-323.
- Swallow, K. M., & Jiang, Y. V. (2010). The attentional boost effect: Transient increases in attention to one task enhance performance in a second task. *Cognition, 115*, 118-132.
- Talmi, D. & McGarry, L. M. (2012). Accounting for immediate emotional memory enhancement. *Journal of Memory and Language, 66*(1), 93-108.
- Theeuwes, J. (1993). Visual selective attention: A theoretical analysis. *Acta Psychologica, 83*, 93-154.
- Vuilleumier, P. (2005). How brains beware: neural mechanisms of emotional attention. *Trends in Cognitive Sciences, 9*(12), 585-594.
- Wilkowski, B. M. & Meier, B. P. (2010). Bring it on: Angry facial expressions potentiate approach-motivated motor behaviour. *Journal of Personality and Social Psychology, 98*(2), 201-210.
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision, 4*, 1120-1135.
- Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in Cognitive Sciences, 13*(4), 167-174.
- Yantis, S. (2000). Goal directed and stimulus driven determinants of attentional control. In S. Monsell & J. Driver (Eds.), *Attention and performance XVIII: Control of cognitive processes* (pp. 71-104). Cambridge, MA: MIT Press.
- Zhang, W., & Luck, S. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453*, 233-235.

CHAPTER 6.

Early Visual Memory and Consolidation Masking

In Chapter 1, we visited the notions of consolidation, fragile VSTM, and ‘individuated’ proto-objects as important aspects of the early processes by which stimuli are memorised. In this chapter I ask whether and how emotional content of faces may have an influence on how they are encoded at this early stage. Perhaps emotional content can influence early processes and so facilitate the encoding of emotional faces into visual WM. There are good reasons for supposing this might be the mechanism by which the angry boost effect occurs: First, as I discussed in Chapter 2, angry faces and to a lesser extent happy faces are believed to be able to activate pre-attentive processing via activation of subcortical networks, especially when attentional resources are directed elsewhere. Second, in order to be treated differently in WM from non-emotional faces, emotional faces must be encoded differently. Thus, we would expect some difference in early WM processes to at least ‘tag’ emotional faces for preferential treatment by later WM processes even if the early consolidation process is not directly responsible for the WM boost observed in the previous chapter.

The time-course of consolidation

The process of consolidation has been shown to be disrupted by masking, and thus interrupting the process by presenting masks at different times can be used to differentially disrupt the process and establish an index of the time required to consolidate a particular representation into a durable WM trace (Gegenfurtner and Sperling, 1993). Following a similar masking procedure, Vogel et al. (Vogel, Woodman & Luck, 2006) suggest that the time required to consolidate simple stimuli (coloured squares) is a mere 50 ms per item. However, consolidation time for different stimulus features (colour, orientation, shape, etc.) may not be constant (Woodman & Vogel, 2008). Additionally, the perceived

complexity, that is, the physical complexity of a stimulus moderated by the degree of expertise of the perceiver (Jackson & Raymond, 2005; Zimmer & Fu, 2008) has been shown to modify the amount of time required to consolidate, with stimuli of higher perceived complexity requiring longer consolidation times (Sun, Zimmer & Fu, 2011).

If there is something 'special' about the consolidation process for angry faces then two possibilities could account for the angry WM boost: If angry faces are encoded with more detail than happy faces then it is possible that the consolidation process might take longer for angry faces as the object file for such faces may stay open for longer in order that additional detail be consolidated into the angry face representation. Alternatively, if angry faces are encoded with greater efficiency than happy faces then the consolidation process for angry faces might proceed more quickly such that 'closing' of the object file occurs at an earlier time.

To investigate whether the time-course of consolidation is different for angry and happy faces, I used a paradigm that built on the WM task reported in Chapter 5. I asked participants to memorise an array of faces and report the presence/absence of a single face presented at probe. However, during the retention interval and at a time that varied between blocks of trials I presented masks (scrambled faces) at the location of the to-be-remembered stimuli in order to disrupt the consolidation process²⁷. If angry faces are encoded with greater resolution, accompanied by a longer period of consolidation, then we would expect to find a particular mask onset presentation time at which consolidation of

²⁷ It is reported in Sun et al. (2011) that masks need to be presented at the location of the to-be-remembered stimuli rather than just besides these locations in order to interfere with consolidation, presumably because consolidation operates on a retinotopically organised perceptual trace.

an angry face is disrupted but consolidation of a happy face is not (as the mask would be presented during consolidation for the angry face but post consolidation when the stimulus to be memorised is a happy face). Conversely, if the angry WM boost is occasioned by more efficient processing of angry faces and a shorter consolidation time then we would expect to find a time of mask presentation at which encoding of happy faces is selectively disrupted.

EXPERIMENT 3

Method

Participants

40 students (26 female, mean age = 23 years, s.d. = 5.9 years) reporting normal or corrected-to-normal vision completed the experiment in exchange for course credit. Informed consent was obtained prior to participation.

Apparatus

Stimuli were presented on a Mitsubishi Diamond Plus 20 inch colour monitor (resolution = 1280 x 1024 pixels, refresh rate = 100Hz), using a Dell Optiplex GX400 computer running Microsoft Windows XP. Viewing distance was 60 cm. E-Prime software (Version 1.0; Schneider, Eshman, & Zuccolotto, 2002) generated stimuli and recorded responses acquired via a keyboard.

Stimuli

Face stimuli used in each experiment reported in this chapter were grayscale bitmap images (depth 24; sized to 73 x 84 pixels; subtending

approximately 2.2 x 2.4 deg. of visual angle) of six male adults bearing angry, happy, or neutral expressions. Neither hair nor neck was visible.²⁸ Face images were of the same individuals as those reported in Chapter 5 and are shown in Appendix C. Masks were scrambled face images constructed via a segmenting of face images not used here but drawn from the same stimulus set (Ekman & Friesen, 1976) into twenty one square portions, pseudo-random allocation of each portion to a position in the new mask stimulus and cropping of the resulting stimulus such that it subtended approximately the same visual angle as the face stimuli. The masks used are shown in Appendix C. Study and mask arrays comprised four locations arranged in a grid centred on the screen (as shown in Figure 6.1). Horizontal and vertical spaces separating the locations were 0.2 deg. of visual angle. For study arrays, at each location either a face or a scrambled face (one of the mask stimuli) appeared (see procedure). For mask arrays, four mask stimuli were present, one at each of the location in the grid. Probe faces were presented centrally.

Procedure

The experiment comprised five blocks, each comprising 48 trials (with an additional 12 practice trials to start each block) and each using different timing intervals (see below). Block order was counterbalanced between participants according to a balanced Latin Square procedure. Prior to each block, participants were required to study for a minimum of two minutes a diagram showing what

²⁸ Not having hair visible in the face stimuli makes the task harder than it would otherwise have been since the stimuli have less information on which to be distinguished from one another. However, this change in stimuli from those used in Chapter 5 was necessary since I needed the mask stimuli to exactly overlay the spatial locations of the stimuli presented in the study array and thus needed to crop each face stimulus such that each subtended the same degree of visual angle rather than continuing to use the more 'natural' stimuli as used in Experiments 1 and 2.

each of the six individuals in our stimulus set looked like when bearing a happy, angry, or neutral expression (see Appendix C). Participants were briefed on the task and told that being able to determine the identity of an individual irrespective of their emotional expression would help them complete the task. This process was done because a pilot study had revealed that the task was extremely difficult without its inclusion.

On each trial, participants performed a WM change detection task (Figure 6.1). A central fixation cross (3000 ms) briefly enlarged to indicate trial onset. Next, a study array comprising two faces and two scrambled faces was presented²⁹. Two face stimuli selected at random (without replacement) from the set of six individuals occupied two grid locations selected at random. On 50% of trials both faces had happy expressions, on the remainder angry. Two scrambled faces selected at random (without replacement) from the set of six scrambled faces occupied other locations. Participants were instructed to ignore scrambled faces and only attempt to memorise the face stimuli. After 2000 ms the retention interval began. Retention intervals lasted 1000 ms. At a certain point during the retention interval, varying as a function of block, a mask array lasting 200 ms was presented. Timings from memory array offset until mask onset were: 17 ms, 117 ms, 317 ms, 500 ms, and 700 ms. At times during the retention interval other than during the presentation of the mask the screen was blank. Finally, a probe face

²⁹ I used study arrays with two rather than four faces here as: (1) I expected the consolidation masking and changes to face stimuli from Experiments 1 and 2 as noted above to result in a harder task than the task reported in Chapter 5 and so wanted to compensate for this by reducing the set size. (2) The angry WM benefit has been more strongly observed in arrays comprising two rather than four faces (Jackson, Wu, Linden, & Raymond, 2009) and I wanted as strong as possible an angry benefit in order to better show the effect of interrupted consolidation on it. Scrambled faces were included in the study array so that the task might better be later transported into an ERP setting.

was presented until response. Probe faces bore neutral expressions.³⁰ On 50% of trials probe faces matched in identity to a face presented in the memorisation array, on remaining trials they did not. Participants indicated by key press whether the individual presented at probe had been present ('e' key) or absent ('i' key) in the study array. Within each block, probe presence/absence and study array expression were fully crossed. Accuracy and response time (RT) were recorded. Each trial started automatically on completion of the previous trial.

A verbal suppression task, identical to that reported for Experiment 1, ran throughout.

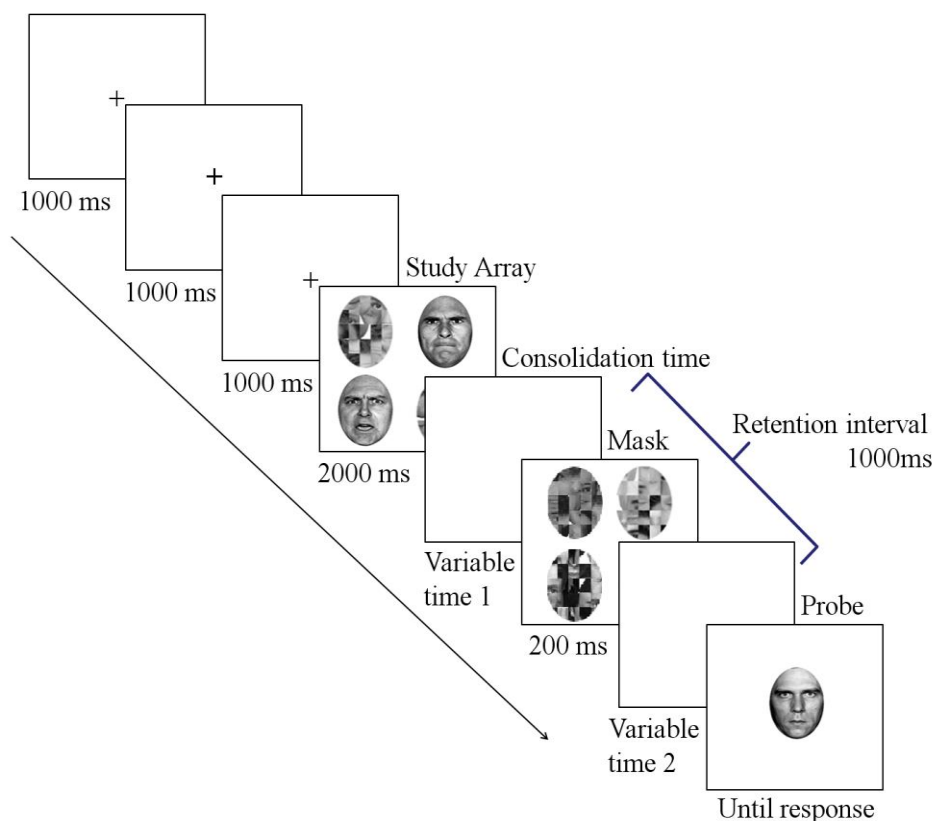


Figure 6.1. Trial procedure for Experiment 3.

³⁰ Use of probe faces bearing neutral expressions with a match to identity rather than probe faces bearing emotional expressions with a direct match to memorisation array allows us to rule out the possibility that arousal at retrieval underpins the angry benefit. This issue has been reported in greater detail elsewhere (Jackson, Linden & Raymond, in press).

Data Analysis

Four participants showed ceiling performance across the task and so were excluded. Another six participants found the task too difficult, either achieving near chance performance throughout or exhibiting a sudden drop in performance level indicative of giving up half way through. Accordingly, these participants were also excluded from further analysis. For remaining data, trials on which anticipation errors (RT less than 200 ms) or task disengagement (RT greater than 10,000 ms) had occurred were excluded. This accounted for 0.2% of data. D' scores were calculated as an index of sensitivity to change according to the procedure outlined in Chapter 4. A within-subjects analysis of variance (ANOVA) was conducted with expression at encoding (angry; happy) and mask onset delay (time between memory array offset and mask onset: 17 ms; 117 ms; 300 ms; 500 ms; 700 ms) as factors. Trend analysis (within subject contrasts) investigated the shape of the progression of performance across mask onset delays for angry and happy faces separately. Alpha levels were set at .05.

Results

Overall performance was better when faces were angry, replicating the angry face WM benefit observed in Chapter 5. However, in the 17 ms conditions only, performance was better for happy ($M = 1.70$, $SD = .64$) than angry ($M = 1.55$, $SD = .67$) faces while for the 117 ms (happy $M = 1.68$, $SD = .99$; angry $M = 1.91$, $SD = .95$), 317 ms (happy $M = 1.59$, $SD = .81$; angry $M = 1.82$, $SD = .90$), 500 ms (happy $M = 1.67$, $SD = .90$; angry $M = 1.80$, $SD = .77$), and 700 ms (happy $M =$

1.51, SD = .84; angry M = 1.76, SD = .93) conditions performance was better for angry faces (see Figure 6.2).




Figure 6.2. Performance across conditions in Experiment 3. Error bars represent within subjects normalised standard error (Cousineau, 2005).

Tests of within subject contrasts (trend analysis) for angry and happy faces data separately revealed that a quadratic function approached significance ($p = .062$) in fitting the angry face data, while no function-fit approached significance for the happy data. Accordingly, I ran a within subjects ANOVA with the mask onset timing factor defined by two levels, early (17 ms) and late (data for 117 ms to 700 ms combined) and the emotion factor also comprising two levels

(happy; angry).³¹ This revealed a non-significant effect of mask onset time [$F(1,29) = 1.261, p = .271, \eta_p^2 = .042$], a non-significant effect of emotion [$F(1,29) = .125, p = .727, \eta_p^2 = .004$], but a significant interaction between these factors [$F(1,29) = 5.930, p = .021, \eta_p^2 = .170$]. Paired samples t-tests between emotions at the early and late onset times as defined above revealed a non-significant difference in the early mask onset time interval [$t(29) = -1.118, p = .273$] but a significant difference for the late onset time intervals [late happy $M = 1.61, SD = .69$; late angry $M = 1.82, SD = .66$; $t(29) = 2.774, p = .010$].

The ANOVA and t-tests reported above are useful in terms of exploring the data, but the ‘rolling up’ of data points into new ‘early’ and ‘late’ data points for this analysis is justified only inasmuch as the trend analysis on the data series for angry faces suggests that there is a change in performance between the 17 and 117 ms mask onset delay timings. A stronger justification for comparing the very early (17 ms) timing to later timings would be provided by showing that there is a change in the ‘angry face benefit’ (the relatively better performance for angry vs. happy faces) across timing intervals. To investigate this, I plotted the relative angry benefit (the difference between memorisation performance for angry and happy faces) as a function of timing interval (see Figure 6.3).

³¹ A similar ANOVA using just the 17 ms and 117 ms conditions rather than a rolled up late condition also revealed a significant interaction between emotion and onset time [$F(1,29) = 5.384, p = .028, \eta_p^2 = .157$].

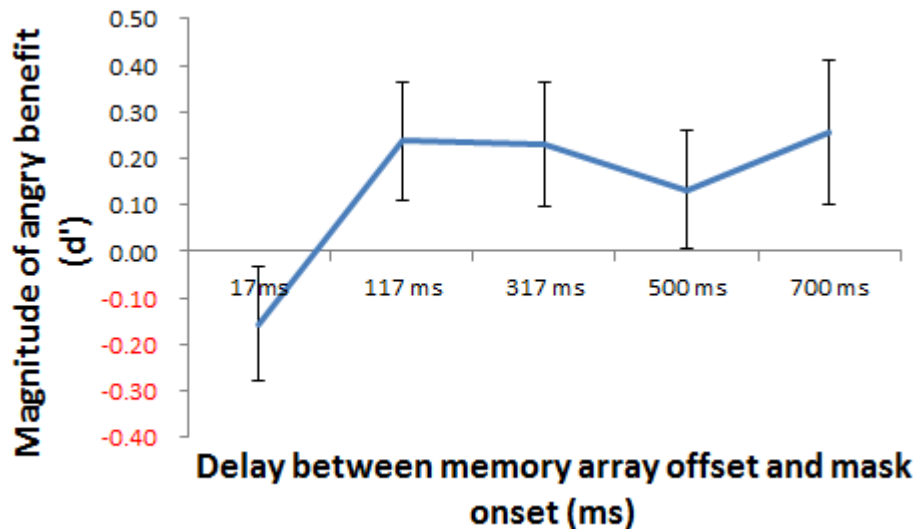


Figure 6.3. Magnitude of angry face benefit at different mask onset intervals. Error bars represent normalised standard error of the mean (Cousineau, 2005).

If the angry benefit is different at different timing intervals then we would expect to see a trend in the relative benefit data. However, trend analysis revealed that neither linear nor quadratic nor cubic nor 'order 4' trends fit this data significantly ($p > .15$ in each case). Thus, while our pairwise comparisons provide some indication that there might be a difference in the way angry and happy faces are treated at different masking intervals and direct us to investigate this further with regard to masking intervals between 17 ms and 117 ms (see Experiment 4 below), we must be cautious in the conclusions we draw from these comparisons given that no significant trend could be fit to the angry benefit (angry minus happy face performance) across timing intervals (Figure 6.3).

Discussion

The angry face WM boost observed in Chapter 5 was also observed here, but when masks were presented sufficiently early in the retention interval the

angry benefit appears to be abolished. This supports my hypothesis that elaborated consolidation processes underpin the angry face WM boost and that such processes are disrupted by early masking but finish relatively early after stimulus offset such that later masks do not disrupt them. That only angry face retention was disrupted when early masks were presented suggests that the process disrupted is one of consolidation rather than perception (otherwise we would also have observed an impact of early making on happy face performance).

This result is also interesting since it suggests that certain consolidation processes may be initiated by stimulus offset. While I had presented study arrays for 2000 ms in order to allow for complex face stimuli to be fully perceived, it was possible that this extended viewing time might also have been enough time for consolidation processes to be completed. This appears not to be the case. Perhaps this is due to a process of efficient encoding that only initiates consolidation processes when there is reason to do so: For as long as stimuli remain in the visual scene, the effortful process of creating robust, durable WM representations for them may be suspended in order to allow WM to be utilised for other functions (thoughts, planning, etc.). This is not to suggest that no WM representation for entities present in the visual scene is formed until offset, rather that such representations remain fluid, and are perhaps best conceptualised as object files that have yet to be closed (Kahnemann, Treisman, & Gibbs, 1992; Sun, Zimmer & Fu, 2010) or as proto-objects (Gao, Gao, Li, Sun & Shen, 2011).

The cut-off point at which the angry boost effect seems to be disrupted appears be somewhere between 17 ms and 117 ms. In order to better define where this point might be I ran Experiment 4, in which I used more finely grained early mask onset timings. Experiment 4 also included a condition in which no

mask was presented to act as a baseline condition. Additionally, in Experiment 4 I chose to randomise rather than block the presentation of masking intervals in order to eliminate the possibility that a strategic element contributed to the findings of Experiment 3.

EXPERIMENT 4

Method

Except as below, Experiment 4 replicated Experiment 3 in all regards.

Participants

20 students (12 female, mean age = 20 years, s.d. = 2.1 years) participated.

Procedure

Prior to the experiment participants completed a block of 14 practice trials. The main experiment comprised a single block of 384 trials with enforced rest breaks every 48 trials. Memory array offset to mask onset intervals were 17 ms, 33 ms, 67 ms, 117 ms, and 317 ms. Additionally, a final baseline condition had no mask present during the retention interval. Trial types were intermixed (randomised) throughout, subject to the constraint that an equal number of each trial type occurred within each bin of 48 trials.

Data analysis

No participants made any anticipation errors or showed signs of disengagement with the task. Accordingly no data were excluded from analysis.

Results

In the 17 ms (happy $M = 1.31$, $SD = .71$; angry $M = 1.29$, $SD = .69$), 33 ms (happy $M = 1.42$, $SD = .68$; angry $M = 1.40$, $SD = .75$), and 67 ms (happy $M = 1.39$, $SD = .69$; angry $M = 1.43$, $SD = .74$) conditions performance was very similar between angry and happy face arrays, while in the 117 ms (happy $M = 1.37$, $SD = .62$; angry $M = 1.50$, $SD = .66$), 317 ms (happy $M = 1.40$, $SD = .97$; angry $M = 1.47$, $SD = .76$), and no mask (happy $M = 1.42$, $SD = .62$; angry $M = 1.54$, $SD = .88$), conditions performance was slightly better for angry faces (see Figure 6.4).

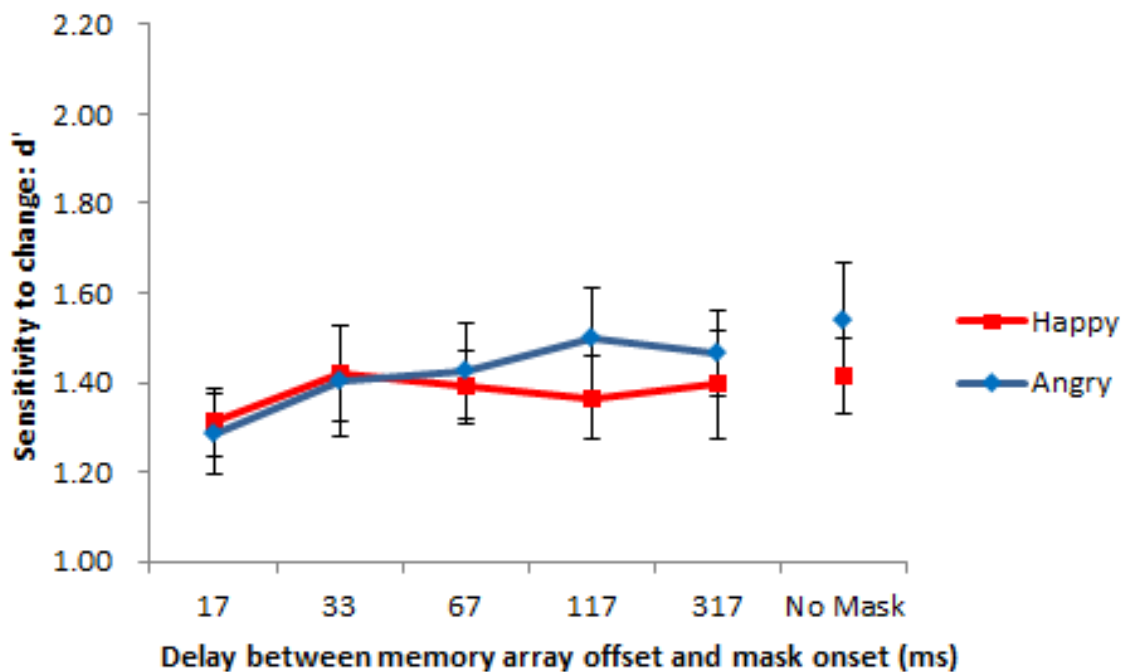


Figure 6.4. Performance across conditions in Experiment 4. Error bars represent within subjects normalised standard error (Cousineau, 2005).

Tests of within subject contrasts (trend analysis) for angry and happy faces data separately revealed no functions that fit either emotional data set well. A repeated measures ANOVA on all data save the no-mask condition revealed that neither the main effect of emotion ($p > .5$) nor mask onset delay ($p > .7$) nor the interaction ($p > .9$) approached significance. A paired samples t-test on the no-mask baseline condition revealed that the performance difference by emotion here did not approach significance ($p > .5$). Thus, it seems that the angry face WM benefit is abolished throughout under these conditions.

Discussion

Surprisingly, I did not observe the angry boost effect at any mask onset timing condition or even when no mask was presented. One possibility for this result is that the shift from an experiment in which conditions were blocked to one in which they were randomised on a trial by trial basis might have introduced an element of uncertainty as to the time available to consolidate after memorisation array offset. Such uncertainty might have caused a 'rush to consolidate' so that participants failed to activate the elaborated consolidation mechanism posited as an explanation for the result observed in Experiment 3. Alternatively, perhaps uncertainty as to consolidation time available results in maintenance of a more complex set of task rules in WM and this recruitment of WM resources leaves less available to bring to bear on memorisation. In this impoverished state insufficient resource may be left for recruitment to facilitate the angry face boost. I discuss the issue of sharing WM resources between representations and task rules in more detail in Chapter 9 (General Discussion).

In order to determine whether the change from blocked to randomised presentation of mask onset delay intervals could have abolished the angry face WM boost, I conducted a third experiment in which in some blocks mask onset delay intervals were held constant and in others were varied on a trial by trial basis. If a rush to consolidate is caused by ambiguity in time available to consolidate then we would expect to see abolition of the angry face WM benefit in trials with long mask onset delays when such trials occurred as part of a mixed timing block, but no abolition of the angry face WM benefit when trials with long mask onset delays occurred all together in a single block.

EXPERIMENT 5

Method

Except as below, Experiment 5 replicated Experiment 3 in all regards.

Participants

36 students (29 female, mean age = 20 years, s.d. = 4.1 years) participated.

Procedure

Prior to the experiment participants completed a practice block of 32 trials comprising trials of each trial type. Prior to each block of trials in the main experiment participants completed 14 practice trials of the timing type appropriate to the block. There were 5 blocks of trials, the order of which varied according to a balanced Latin Square across participants. One block used a mask onset timing

of 17 ms, another used 117 ms, and another had no masks. Each of these blocks comprised 64 trials. The remaining two blocks comprised 96 trials each, divided equally and randomised in presentation between the two mask onset delays and the no mask condition. Thus, in total each experimental condition (see Figure 6.5) had 32 trials

Data analysis

Eight participants were observed to have very low d' scores either overall or in the latter parts of the experiment indicating that they did not remain motivated in later blocks. Data from these participants was therefore excluded from further analysis. Anticipation errors accounting for .06% of remaining data were excluded from the remaining data.

Results

When the mask onset delay was 17 ms, performance was slightly better in the blocked happy condition ($M = 2.22$, $SD = .89$) than in the blocked angry condition ($M = 2.06$, $SD = .84$), while performance in the randomised happy ($M = 2.11$, $SD = .98$) and angry ($M = 2.10$, $SD = .88$) conditions for this timing was similar. When mask onset was 117 ms, performance was marginally poorer in the blocked happy ($M = 2.02$, $SD = .96$) vs. blocked angry ($M = 2.09$, $SD = .87$) condition while performances were similar between the randomised happy ($M = 2.22$, $SD = .83$) and angry conditions for this timing. When no masks were present, performance was markedly poorer when the trials were blocked (happy $M = 1.89$, $SD = 1.06$; angry $M = 2.02$, $SD = .80$) than when they were randomised (happy $M = 2.19$, $SD = .80$; angry $M = 2.29$, $SD = .78$), see Figure 6.5.

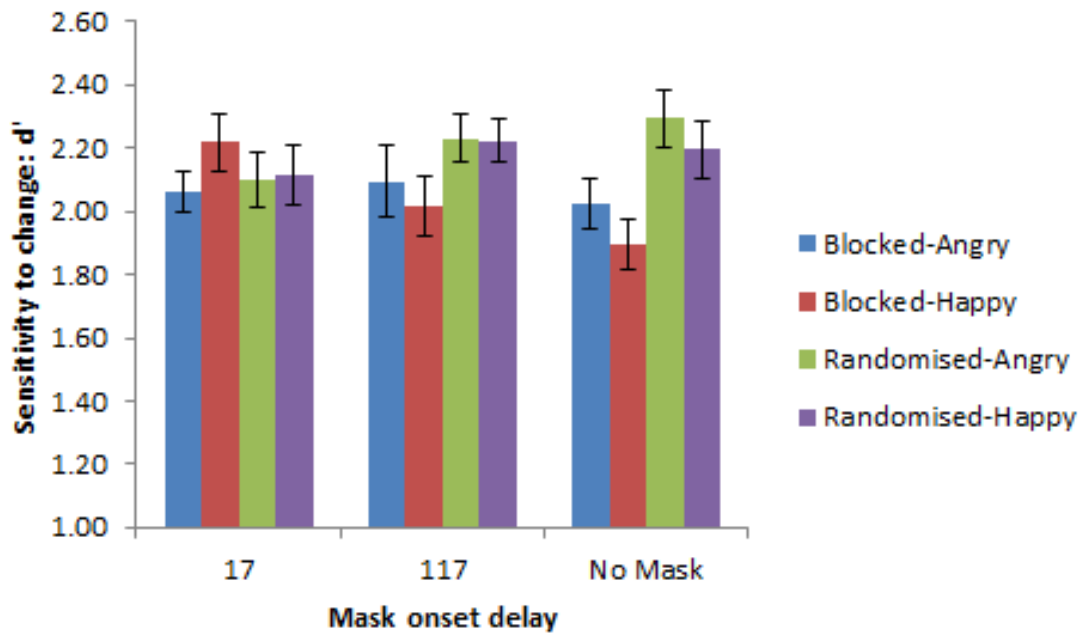


Figure 6.5. Performance by condition in Experiment 5. Error bars represent within subjects normalised standard error (Cousineau, 2005; see Appendix D).

A repeated measures ANOVA with onset delay (17 ms; 117 ms), emotion (happy; angry), and block type (blocked; randomised) as factors revealed no main effects or interactions that approached significance. A repeated measures ANOVA on the data for the baseline (no-mask) condition with trial type and emotion as factors revealed a non-significant effect of emotion [$F(1,27) = 1.541$, $p = .225$, $\eta_p^2 = .054$], a significant effect of trial type [$F(1,27) = 8.388$, $p = .007$, $\eta_p^2 = .237$], and a non-significant interaction between these two factors [$F(1,27) = .030$, $p = .864$, $\eta_p^2 = .001$].

Discussion

This experiment did not work as well as hoped. In interpreting the performance for the 28 participants whose data were included in analysis, we must bear in mind that the pattern of performance among eight other participants revealing disengagement from that task, resulting in their exclusion from analysis. This is suggestive of a fatiguing task in which difficulty was high and must be borne in mind in terms of the credibility one can place on the patterns of performance observed in the remaining data (since a strategy of disengagement over a few trials then reengagement with the task cannot be ruled out, is hard to detect, and is one possible strategy that participants might take when a task is particularly tiring). In particular, the unexpected finding that performance was poorer across both happy and angry no-mask trials when such trials were blocked vs. randomised might reflect a withdrawal of effort in this condition due to its relative ease compared to each other condition. This provides a cautionary lesson in identifying and controlling for participant engagement in tasks involving multiple blocks.

CHAPTER DISCUSSION

In the experiments reported in this chapter I set out to determine whether an early consolidation process might underpin the angry face WM benefit observed in Chapter 5. There is some indication that this might be the case: The abolition of the angry face benefit occasioned by very short mask onset intervals designed to constrain consolidation time (Experiment 3) along with absence of the effect in Experiment 4 where ambiguity of consolidation time available might have had a role in disrupting elaborated processing suggests that the angry face WM boost is vulnerable to masking during the retention interval perhaps due to an interruption of consolidation processes.

In each of the experiments reported in this chapter I gave ample time (2000 ms) for full perception of the face stimuli to occur. Given this long encoding interval, if consolidation occurs in an automatic fashion during stimulus presence then we might expect consolidation processes to also be fully completed during the study interval. It is interesting to note that the process that is disrupted in the case of angry faces at short mask onset delays must be something that is still running after stimulus offset. There are two possible accounts of what this process might be. First, it is possible that the offset of the stimulus triggers the initiation of a process of closing of an object file, i.e. that consolidation of angry faces is extended until stimulus offset. Second, it is possible that a process of elaboration/consolidation for angry faces continues shortly after stimulus offset and it is this process rather than the process of closing the object file that is disrupted. That object files for angry faces seem not to be closed until shortly after stimulus offset might reflect a process that allows effective monitoring of such faces: If keeping an object file open while a stimulus is present allows for rapid updating of such a stimulus then this might well be a beneficial response to angry faces given their dynamic nature and the importance of monitoring changes in facial expression outside the experimental setting.

Consolidation has generally been framed in terms of a process whereby perceptual features are integrated into or made more durable within an existing WM representation which becomes higher resolution or more robust through this process. However, it is worth noting that angry faces have semantic connotations as well as perceptual features and such information might also be subject to a 'firming up' process after stimulus offset. This notion fits within the framework of template memory (see Chapter 1) and conceptual STM (CSTM; see General

Discussion, Chapter 9), both of which suggest that semantic aspects of a stimulus can quickly determine the way it is treated by WM processes. It is also possible that stimulus offset may play a special role as an initiation point for such semantic elaboration in the same way that it seems to do for perceptual consolidation. This idea gains support from a recent study by Swallow et al. (Swallow, Zacks & Abrams, 2009) who showed that event boundaries (time points where one observed goal-directed visual action was judged to have ended or another begun in a naturalistic setting) seem to have a special role in determining when the contents of WM are updated. In particular, Zacks et al., suggest an event segmentation theory (EST; Zacks, Speer, Swallow, Braver, & Reynolds, 2007; Zacks & Swallow, 2007), a key aspect of which is the predictive way in which the contents of WM determine when semantic event boundaries occur. For example, suppose an observer sees an actor boil a kettle, open a cupboard, take out a mug, place a tea bag into the cup and then pick up the kettle. The next logical step in the sequence described is for the actor to pour hot water into the mug in order to make a cup of tea. If this next action does occur then a semantic event boundary will occur after the action of pouring the water, upon offset of the tea-making goal and initiation of the next goal. However, if rather than pouring the water into the cup the actor instead pours it into the sink then the event boundary occurs before this action, since the tea-making goal has been interrupted and a new goal initiated. In a similar way, the observation that uncertainty as to time available to consolidate impacts the angry face boost (Experiment 4; where mask onset delays were randomised) may be interpreted either as a rush to consolidate, or in terms of event segmentation theory: Perhaps ambiguity as to the onset of an event occasions a breakdown in, or inability to form, the event model. From this point of

view, the boost usually occasioned to angry faces must be at least partially semantic in nature (since event segmentation theory describes event boundaries as being defined by semantic interpretation of the visual world). That consolidation might include a semantic aspect, rather than being a purely visual process is also supported by a finding from Jolicoeur and Dell'Acqua (1998), who showed that consolidation is an amodal process, relying on a central resource, as illustrated by the observation that consolidation of visual stimuli interfered with later response selection/initiation in an auditory task. Additionally, the notion that semantic information can be derived very quickly (within 100 ms of stimulus onset) has been suggested as a key part of the theory of CSTM (Potter, 1999; see Chapter 9).

In conclusion, it appears that angry face stimuli may occasion a process of elaboration that results in improved representational precision, and thereby produces better WM performance. In this view, threat in faces may activate emotional processes that facilitate elaboration. My data suggest that such processes occur relatively early after stimulus offset, are negated by backward masking when this occurs early enough after stimulus offset, and are vulnerable to ambiguity in the time available to consolidate memory representations, which may cause a rush to consolidate or lead to difficulty in establishing an event boundary, either of which might result in difficulty in elaborating on the contents of an object file.

References – Chapter 6

- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology*, 1(1), 42-45.
- Ekman, P., & Friesen, W. (1976). Pictures of facial affect. Palo Alto, CA: Consulting Psychological Press.
- Gao, T., Gao, Z., Li, J., Sun, Z., & Shen, M. (2011). *The perceptual root of object-based storage. Journal of Experimental Psychology: Human Perception and Performance*, 37, 1803-1823.
- Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance*, 19, 845–866.
- Jackson, M.C.J., Linden, D.E.J., & Raymond, J.E. (in press). Angry expressions strengthen the encoding and maintenance of face identity representations in visual working memory. *Cognition & Emotion*.
- Jackson, M. C., & Raymond, J. E. (2005). Visual working memory: Capacity is dependent on perceived, not physical, stimulus complexity. *Journal of Vision*, 5, 621.
- Jackson, M. C., Wu, C., Linden, D. E. J., & Raymond, J. E. (2009). Enhanced Visual Short-Term Memory for Angry Faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(2), 363-374.
- Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology*, 36, 138–202.
- Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of files: Object-specific integration of information. *Cognitive Psychology*, 24, 175–219. doi:10.1016/0010-0285(92)90007-O
- Potter, M. C. (1999). Understanding sentences and scenes: The role of Conceptual Short Term Memory. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp.13-46). Cambridge, MA: MIT Press.
- Schneider, W. X., Eshman, A., & Zuccolotto, A. (2002). *Eprime user's guide*. Pittsburgh, PA: Psychology Software Tools, Inc.
- Sun, H., Zimmer, H. D. & Fu, X. (2010). The influence of expertise and of physical complexity on visual short-term memory consolidation. *The Quarterly Journal of Experimental Psychology*, 64(4), 707-729.
- Swallow, K. M., Zacks, J. M., & Abrams, R. A. (2009). Event boundaries in perception affect memory encoding and updating. *Journal of Experimental Psychology: General*, 138(2), 236-257.

- Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, *32*, 1436-1451.
- Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin & Review*, *15*(1), 223-229.
- Zacks, J. M., Speer, N. K., Swallow, K. M., Braver, T. S., & Reynolds, J. R. (2007). Event perception: A mind/brain perspective. *Psychological Bulletin*, *133*, 273-293.
- Zacks, J. M., & Swallow, K. M. (2007). Event segmentation. *Current Directions in Psychological Science*, *16*, 80-84.
- Zimmer, H. D., & Fu, X. (2008, July). *Working memory capacity and culture-based expertise*. Paper presented at the XXIV International Congress of Psychology, Berlin, Germany.

PART 3 – VALUE LADEN FACES IN WORKING MEMORY

CHAPTER 7.

Motivational Salience and Value Learning

So far, we have seen that the emotional expression of a face may affect how it is treated by the WM system. However, the motivational salience of a face stimulus can depend not only on its expression but on prior learning experiences when the face has been previously encountered. To investigate whether the value association of a face (whether it is likely to be predictive of reward or punishment) might also affect how it is treated in WM, I conducted the three experiments in this chapter. In Experiments 6 and 7, I imbued value into colour-defined categories of faces. In Experiment 8 I imbued value into individual face stimuli.

Rationale

Only recently has the notion of reward association begun to be integrated into theories of bias in early visual processing, salience, and attentional deployment. For example, neither the Feature Integration Theory (Treisman & Gelade, 1980) nor the most current theory of guided search (Wolfe, 2007) have a place for motivational association of items, except inasmuch as they suggest that attention should be guided towards task relevant items. Thus, while such theories may account for the strategic effects of reward and punishment associations, they do not account for automatic effects of reward or punishment association of a stimulus (while reward may make an item task-relevant, the reward or punishment association of an item may in some situations be orthogonal to a task).

Despite the absence of reward (and punishment) from current models of salience, there are several reasons to suppose that reward/punishment association might in fact bias motivational salience, perhaps by directly biasing visual salience. First, Raymond and O'Brien (2009) have shown that stimuli with a history of reward pairing are able to escape the attentional blink, even when

reward association is orthogonal to the attentional blink task. Second, two studies by Hickey et al. (Hickey, Chelazzi & Theeuwes, 2010a; 2010b) using the attentional singleton paradigm (Theeuwes, 1991) have shown that prior reward association of a stimulus both facilitates response towards it when it is a target and makes it harder to ignore when a distractor (Hickey, Chelazzi & Theeuwes, 2010a). Such effects are more pronounced among individuals with reward seeking personalities (Hickey, Chelazzi & Theeuwes, 2010b). Of particular interest in these studies is the finding that the size of reward-based priming effects was correlated with activity in the anterior cingulate cortex which is believed to index reward related activity, as indexed by a medial frontal negativity (MFN; Gehring & Willoughby, 2002) ERP component. Thus it appears that the level to which a reward stimulus is experienced as rewarding, and the level to which it retains its reward association when encountered in a task where such a association is unrelated to task demands, may determine its ability to determine stimulus salience and thus bias attention capture.

A similar effect has also recently been demonstrated by Anderson et al. (Anderson, Laurent & Yantis 2011a; 2011b), who imbued high and low positive value into colour categories and showed subsequent attentional capture by high reward associated items relative to low reward associated items. This attentional capture occurred both when high reward associated items were presented as 'inconspicuous' distractors (a circle of a colour previously paired with reward, presented among other distractor coloured circles in a search paradigm where the item to be reported was a line orientation bounded by a coloured square target stimulus; 2011a) and when they were presented as 'salient' distractors (a single coloured circle previously paired with reward among black distractor circles where

the item to be reported was a line orientation bounded by a black square target stimulus; 2011b).

The above studies suggest that the reward association of a stimulus may affect the way in which it is perceived, attended to and perhaps the way in which it is memorised. However, it is also possible that punishments (negative events reliably paired with presentation of a stimulus during learning instances) might similarly affect these processes. We observed in Chapter 5 that angry faces are better memorised than both happy and neutral faces, and that this WM boost seems to be item specific, with little impact on concurrently presented neutral faces. One possibility is that rather than anger *per se* underlying this advantage it may be part of a more general negativity bias. In a variety of experimental situations it has been demonstrated that humans respond to equivalent negative and positive events (e.g., wins and losses) in such a way as to exaggerate the relative weighting of the loss over the gain event. This is known as loss aversion (Kahneman & Tversky, 1979).

If angry faces (and fearful faces; Sessa, Luria, Gotler, Jolicoeur, & Dell'Acqua, 2011) have greater motivational salience than happy faces due to a general negativity bias, perhaps loss associated faces might similarly have greater motivational salience than the gain associated faces due to a general loss aversion bias (that is, they may be internalised as more 'important' or 'relevant' entities than positively associated faces and so become salient). If this is so then we would expect loss associated faces to receive a greater WM boost than gain associated faces. However, since both loss and gain associated faces have motivational salience, we would expect both to be better memorised than no-outcome associated faces. We should note, however, that loss aversion has not

been universally found in previous studies (e.g., Rozin and Royzman, 2001; Yechiam & Hochman, 2013).

While the above studies suggest that the salience of items might be affected by their value association, they do not directly address the issue of how such a moderation of salience might affect how such items are treated in WM. To investigate this I conducted the three experiments reported below.

EXPERIMENT 6

Experiment 6 was designed to test the hypothesis that face stimuli into which a learned value association had been imbued might be treated differently in WM from faces in which no learned value had been imbued. Additionally, I was interested in determining whether any such differences might parallel the singleton effects observed for angry, and to a lesser extent happy faces in Chapter 5 and whether any such effects would be dissociable from deployment of overt attention. Accordingly, the WM task in Experiment 6 paralleled that of Experiment 2 (see below).

Method

Participants

Adult participants were recruited through Bangor University and received course credit or money in exchange for participation. All reported normal or corrected to normal colour vision and visual acuity. 31 participants (15 females, mean age = 19 years s.d. = 1.26 years) participated. The experiment was terminated part way through for three further participants, one due to excessive

tiredness and two due to an inability of the eye tracking apparatus to accurately record their eye movements (both of these participants wore thick glasses which, combined with certain shapes of eye and/or small frequent head movements can result in difficulty recording accurate fixation data, especially when fine differentiation of gaze position is required, i.e. when stimuli are close together on screen as in our setup here). Data for these participants were discarded without inspection.

Apparatus

Stimuli were presented on a Mitsubishi Diamond Plus 20 inch colour monitor (resolution = 1024 x 768 pixels, screen refresh rate = 100Hz), using a Dell Optiplex GX400 computer running Microsoft Windows XP. Viewing distance was 60 cm. E-Prime software (Version 1.0; Schneider, Eshman, & Zuccolotto, 2002) generated stimuli and recorded responses acquired via a keyboard. In Experiment 6, an EyeLink® 1000 desktop mounted eye-tracker (SR Research Ltd., 2002) recorded eye movements monocularly from the left eye with a sampling frequency of 500 Hz and a chinrest was used to maintain head position. Standard nine point calibration/recalibration was done prior to each phase and at intervals of every forty trials in phase 3.

Stimuli

'Tile' stimuli were squares (side length = 90 pixels) subtending approximately 2.7 degrees of visual angle, with a mottled light and dark gray interior and a darker gray surround giving them a three dimensional appearance (see Figure 7.1). Tile arrays comprised four tiles, separated from one another by 10 pixels (approximately 0.3 degrees visual angle) in each of the vertical and

horizontal planes. Face stimuli were bitmap images (depth 24; sized to 73 x 84 pixels; subtending approximately 2.2 x 2.4 deg of visual angle) of six male adults bearing neutral expressions (Ekman & Friesen, 1976). These were first rendered in grayscale and then overlaid with either a yellow, red, green or blue tint (see Appendix B for RGB coordinates) giving a total of twenty four face stimuli (6 of each colour). Face arrays comprised four faces, separated horizontally and vertically from one another by approximately 0.5 degrees of visual angle in phase 1 and 0.2 degrees visual angle in phase 3. Arrays and singly presented faces were all presented centrally (see Figure 7.3).

Procedure and Design

The experiment comprised four phases. In phase 1, participants learnt to associate positive, negative or neutral values with particular colours (tints applied to faces) in a novel selection/revelation task. In phase 2, learning of these values was tested. In phase 3 faces tinted with the value associated colours were presented as memoranda in a WM task. Phase 4 was a replication of phase 2 and served both to test whether values remained learnt throughout phase 3 and to provide an incentive for participants to retain such values. The procedure for each phase is described below.

Phase 1 ('Learning' phase)

Participants were told they had the chance to win a small amount of money in a game. They were instructed that they would start the game with 100 points, that each point was worth one penny, that there was the opportunity to win or lose points on each trial, and that 'win', 'loss' and 'nothing' outcomes were determined by colour tints of faces.

After a 1000 ms fixation cross, a display of four tiles was presented with a running points total beneath and the text 'Choose a tile...' above. Participants selected one of the tiles using the '4', '5', '1' and '2' keys on the number-pad of a standard qwerty keyboard. Key positions corresponded naturally to tile positions (i.e. the '4' key represented the top left tile and the '2' key the bottom right tile etc.). Upon response the tile chosen was 'revealed' (replaced with a face stimulus of a particular colour positioned coincident with the centre of the removed tile) and the running points total was replaced with a question mark. Next, the text 'GAIN' or 'LOSS' or 'NOTHING' was presented on screen rendered in bold, 18 point 'Courier New' text, and coloured green when 'GAIN', red when 'LOSS' or black when 'NOTHING'. The running points total was adjusted accordingly (+10 points for 'GAIN'; -10 points for 'LOSS'; no points change for 'NOTHING'). The text presented and points change was determined by the tint of the face present. One tint corresponded to 'GAIN', another to 'LOSS' and the remaining two tints to 'NOTHING'. Tint assignment to outcome was rotated among participants. Finally, the three tiles not selected were also 'removed' to reveal faces of the remaining three tints. Identities of these faces were randomly selected with the constraint that no identity be repeated in the array. There were 48 trials. From the participants' point of view, the task appeared to be based on luck (choosing the tile with the win face under it and avoiding choosing the tile covering the loss face). However, while trials were presented in random order they comprised a set such that each face stimulus would be 'chosen' twice. Thus, each participant was exposed to the same number of wins (12), as losses and to a total of 24 'nothing' outcomes (12 for each of the two 'nothing' outcome tints), and ended up with total winnings of 100 points. To add to the verisimilitude (that the task involved

chance), the experimenter recorded how much the participant had 'won' at the end of phase 1 on a piece of paper. Upon debrief after the entire experiment participants generally expressed surprise that the task was not in fact random.

Phase 1 is shown in Figure 7.1 below.

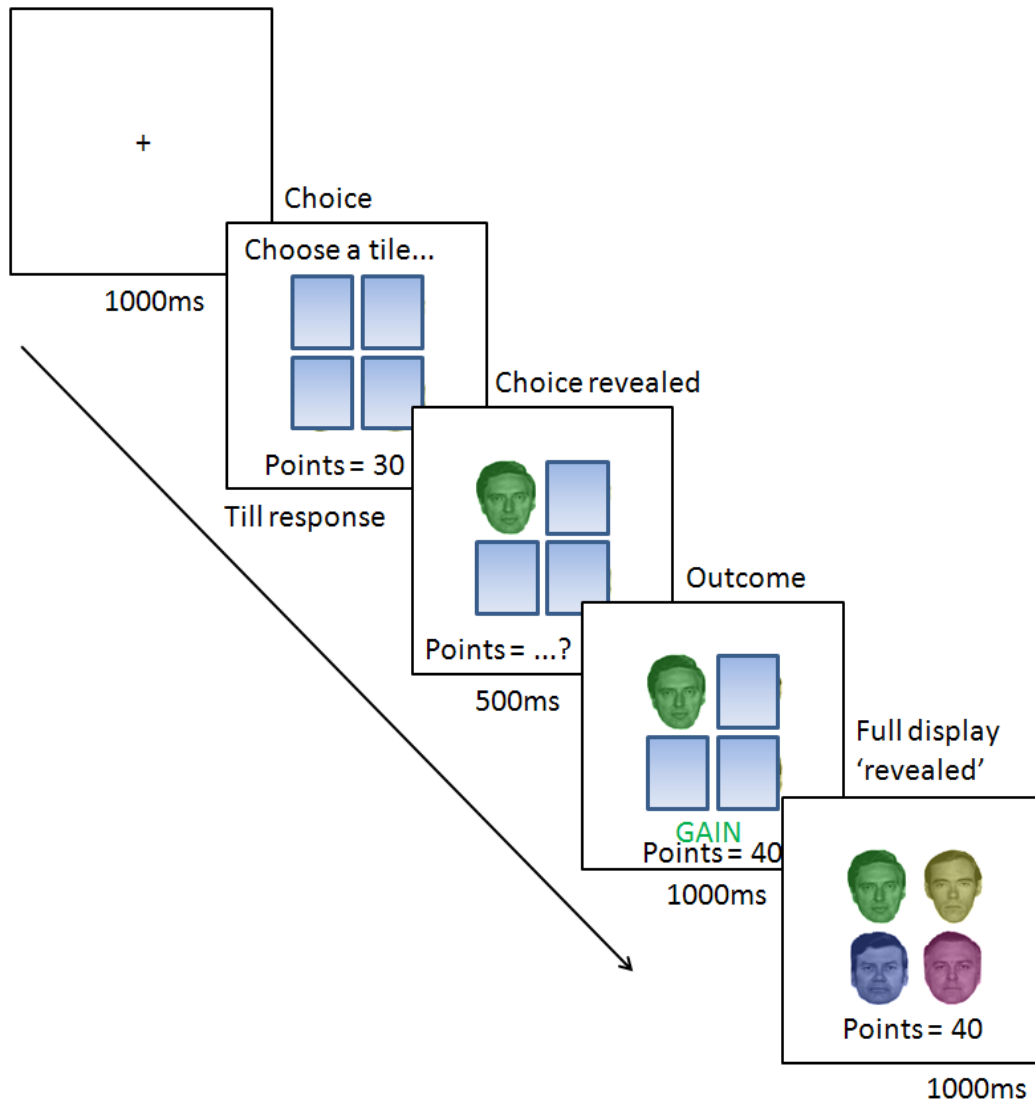


Figure 7.1. Trial sequence for phase 1. In this example the green face has been revealed and for this participant results in a win.

Phase 2 ('Value test' phase)

Phase 2 was described to participants as a 'prediction task'. The trial sequence was explained to them and they were told that the faces retained the same win/loss/nothing tint-based contingencies as in phase 1. Phase 2 comprised 24 trials each having the following sequence (also see Figure 7.2): A fixation cross briefly changed size, indicating trial onset. Next, a single tinted face was presented centrally for 2500 ms (across the 24 trials each face identity with colour tint combination was presented once). Next, with the face remaining on screen, the text 'Do you think you will... Win, Lose or Nothing?' was presented at the top of the screen in bold 20 pt 'Courier New' font. Participants indicated by key press ('z' key = win; 'x' key = loss; 'c' key = nothing) whether they thought the face present would result in a win, a loss, or no outcome. Upon response, an outcome event occurred consistent with the tint of the face presented. Outcomes took the same form as in phase 1 (appropriately coloured text showing 'GAIN', 'LOSS' or 'NOTHING' and a 10 point adjustment to the points total if appropriate) and were presented for 1000 ms.

Upon completion of this phase, participants were instructed that at the end of the experiment they would have to repeat this phase. This gave participants an incentive to remember the value contingencies whilst performing phase 3.

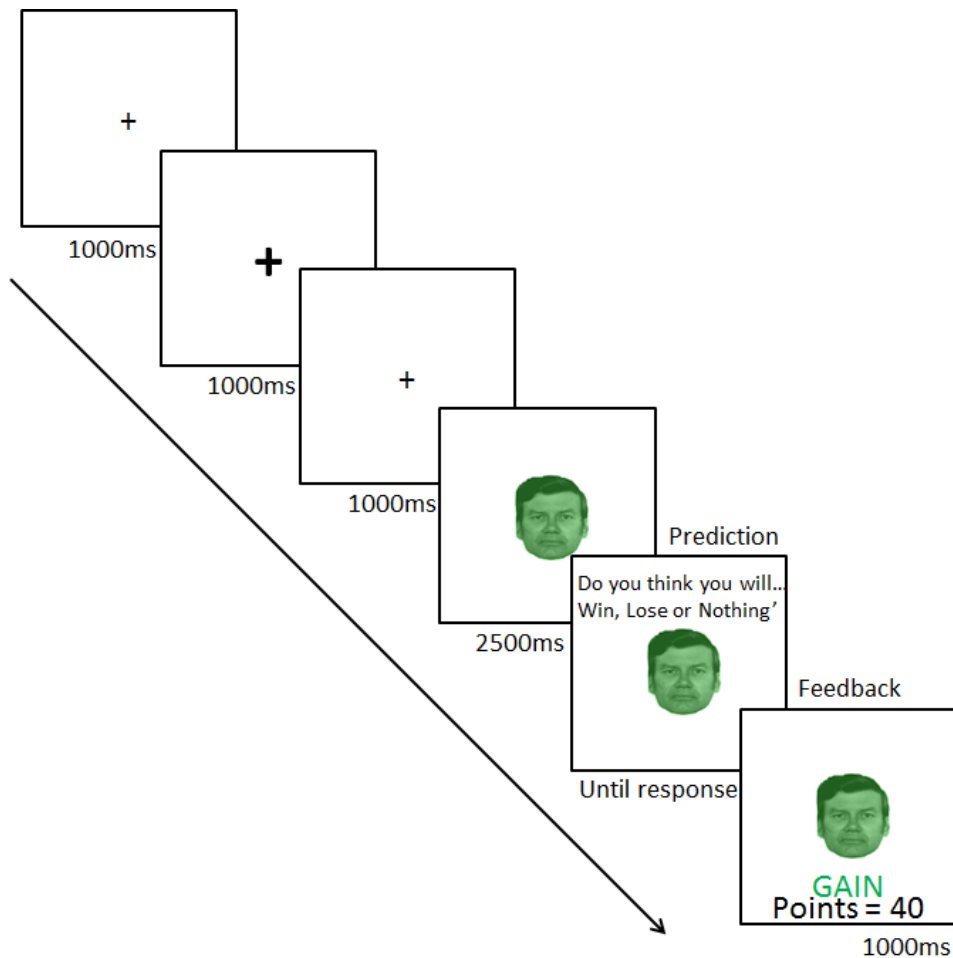


Figure 7.2. Trial sequence for phases 2 and 4. In this example, a green face is shown, the participant predicts whether this will result in a win, a loss or no outcome. For this participant green faces resulted in wins in phase 1, and this is reinforced here by confirmation of the outcome to face assignment.

Phase 3 (WM phase)

Each trial comprised the following: A fixation cross briefly changed size, indicating trial onset. Next, four faces were presented in a two by two array (see Figure 7.3). One of these faces was of a colour tint that had denoted either win, loss or neutral (nothing) contingency in phases 1 and 2. Hereafter I refer to these

value-learned items as the 'win face', 'loss face' and 'neutral faces' respectively. When the singleton face was of win or loss value the three neutral faces present in the array took with equal frequency either of the two possible neutral colour tints. When the singleton face was neutral the non-singleton faces took the alternative neutral-associated tint. Thus, there were three types of memory array: 'win', comprising a win face and three neutrals; loss, comprising a loss face and three neutrals; neutral, comprising a neutral face and three neutral faces of a different tint. The memory array was presented for 2000 ms. This was followed by a retention interval (blank screen) of 1000 ms. Finally, a single probe face was presented centrally. On 50% of trials the probe face had been present in the study array. On 50% of trials the probe face shared a colour with the singleton face.³² Participants responded by key press ('i' = absent, 'o' = present) to indicate presence or absence of the probe from the study array based on identity of the face (i.e. irrespective of colour).

³² It is more usual in studies of this type to allow each item in the study array to be probed with equal likelihood. However, my planned comparisons of interest were between singleton faces of different expected value and between neutral faces co-presented alongside different value faces. I was not particularly interested in examining the effects of being a singleton *per se* (i.e. making comparisons between memory performance for value faces and their respective co-presented neutral faces) since visual colour-based pop out effects could contribute to any such effects. By allowing frequency of probes to differ between singleton and non-singleton items in this way, while I may have introduced a strategic bias in favour of processing singleton stimuli (over and above that already expected to be driven by visual pop out), I was able to acquire twice as many trials in singleton probe conditions as would otherwise have been possible.

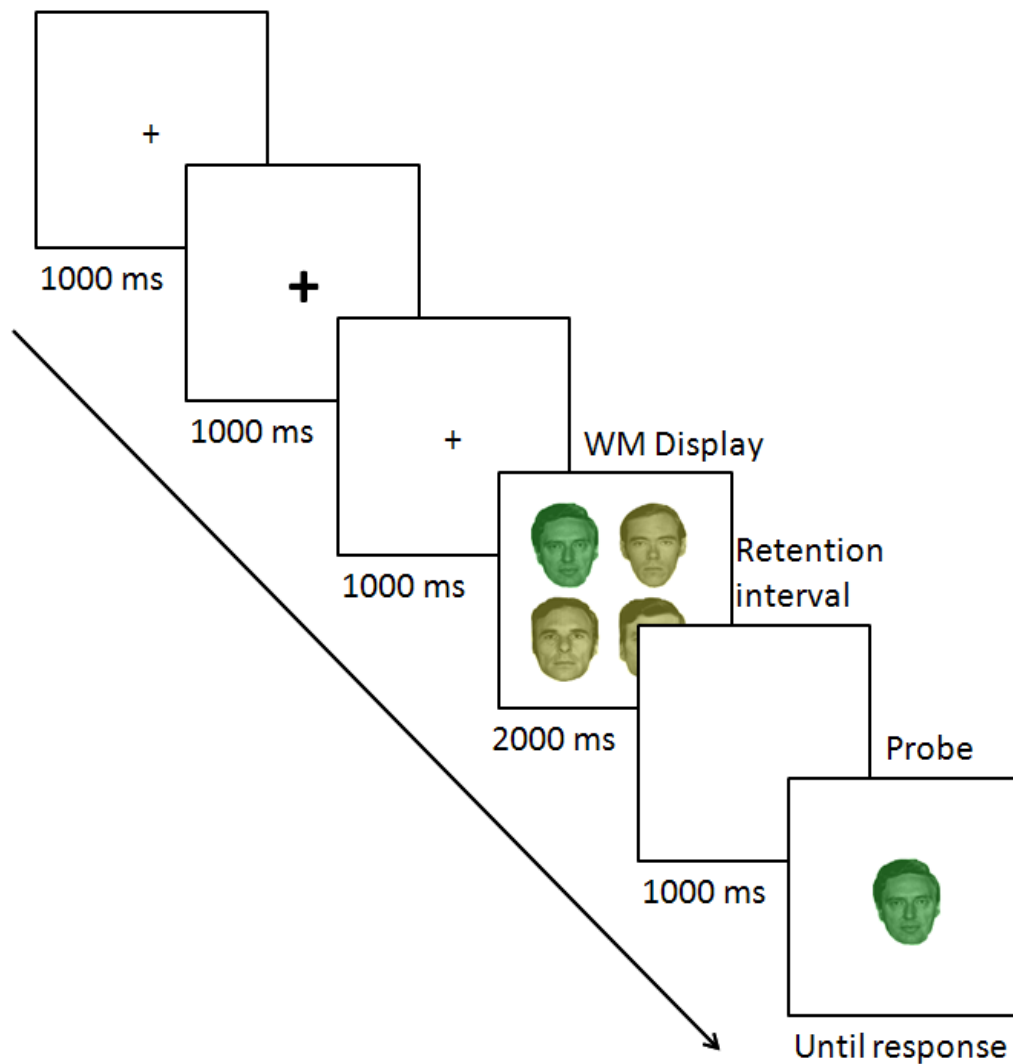


Figure 7.3. Trial sequence for phase 3. In this example memory for the singleton green face is tested with a green probe. The trial shown is a no-change trial as the probe face was present in the memory array.

Phase 3 was preceded by sixteen practice trials and comprised two blocks, each of 120 trials. Within each block there were three broad types of study array; positive singleton, negative singleton, and neutral singleton. There were two possible iterations for each of these types of study array, reflecting the fact that non-singleton faces could take one of the two possible neutral-associated

tints (see Figure 7.4). Each of the four possible colours of singleton appeared equally often. This meant that there were twice as many trials for neutral singleton arrays as for either loss singleton or gain singleton arrays, but this was necessary to allow singleton colours to occur with equal frequency. Array types were presented in random order. The position of the singleton faces within array types was fully counterbalanced within each block. Enforced rest breaks of at least 30 seconds preceded eye tracker recalibration and occurred after every forty trials.

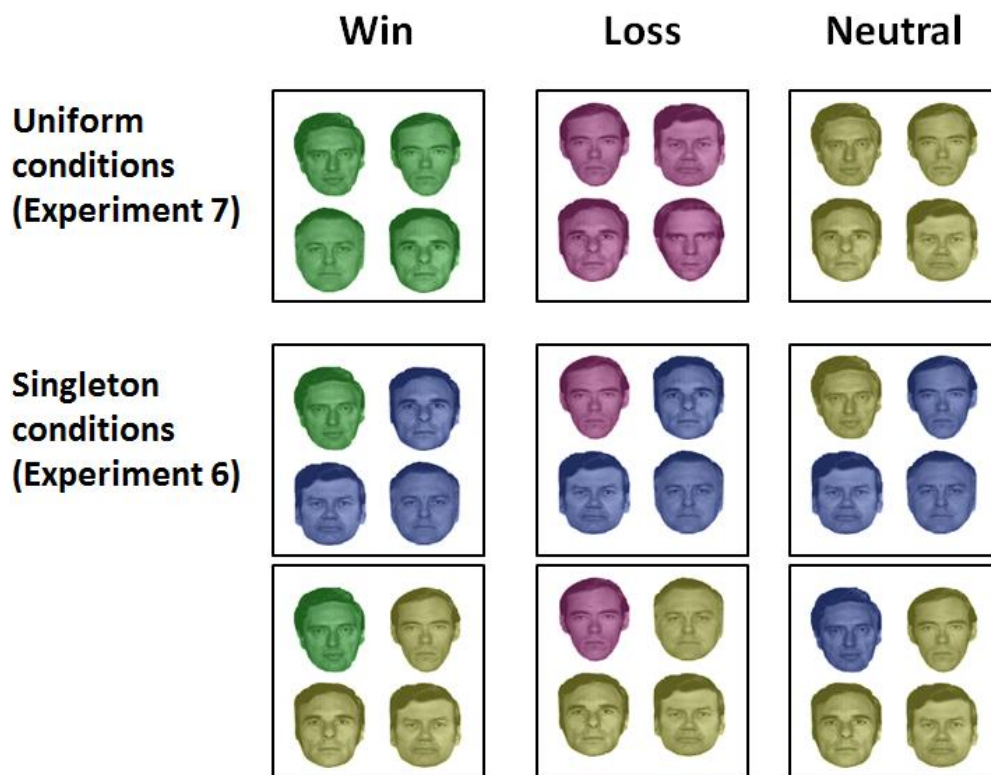


Figure 7.4. Array types used in phase 3 (Experiment 6) and phases 3a and 3b (Experiment 7). In this example the green face is associated with gain, the red with loss and the yellow and blue with no outcome. Each of the neutral singleton conditions was shown with twice the frequency as the gain and loss singleton conditions in Experiment 6 in order that each colour of singleton occur with equal frequency (see text). Only one of the two value-learnt neutral colours

(in this example yellow) was used generate neutral face arrays used in Experiment 7.

Formally, phase 3 was a 3 (singleton type; win, loss, neutral) x 2 (non-singleton colour; colour 1, colour 2) x 2 (probe colour; match to singleton, match to non-singleton) x 2 (probe type; present, absent) design. However, collapsing across non-singleton colours and across probe types allowed analysis to follow a more intuitive 3 (singleton type; win, loss, neutral) x 2 (probe; singleton tested, non-singleton tested) design.

Phase 4 ('value retention' phase)

The procedure for this phase was identical to that of phase 2 above.

Data analysis

Each participant's percentage of correct responses to each value type (colour) in phases 2 and 4 was inspected. Surprisingly, only 12 participants in phase 2 and 20 participants in phase 4 demonstrated perfect learning for each of the colour to reward association contingencies. Loss associations (phase 2 proportion correct: $M = .87$, $SD = .21$; phase 4: $M = .88$, $SD = .23$) appeared to be learnt slightly less well than either neutral (phase 2: $M = .88$, $SD = .21$; phase 4: $M = .94$, $SD = .15$) or gain (phase 2: $M = .94$, $SD = .20$; phase 4: $M = .98$, $SD = .09$) associations. An interesting possible explanation for this apparently weaker learning of loss association was revealed upon interviewing the participants after the experiment. A number of participants indicated that they felt an aversion to pressing the key denoting loss association. Comments such as "I knew it was the

loss colour, but I pressed one of the other keys to see if I could cheat the task” were given by approximately three quarters of the participants who failed to achieve 100% accuracy on this task, despite the fact that no monetary outcomes were associated with performance in these phases. Similar comments were made by participants taking part in Experiment 7. Bearing in mind this interesting dissociation between participants’ ability to learn the colour to reward/loss contingencies and their ability to report it, I also examined an additional measure of learning; how the preceding trial affected response times on trials in phase 1 (see below) in order to have a second measure to confirm learning of contingencies was occurring.

Working memory data

Trials in which reaction times (RTs) were lower than 200 ms (indicating anticipation errors) were excluded from the analysis, accounting for 0.4% of data points. A further 1.2% of trials were excluded from analysis after inspecting their associated eye-movement data (see below). Change detection performance was quantified by d' . This was computed for each participant for each condition by subtracting the Z-transform of the proportion of False Alarms (FA; no-change trials on which participants incorrectly responded ‘yes’) from the Z-transform of the proportion of Hits (change trials for which participants correctly responded ‘yes’). For Experiment 6, d' for singleton faces was calculated using the FA and Hit rates obtained on trials when a test face matching the colour of the singleton was presented. For neutral valenced faces presented alongside the singletons, d' calculations used the FA and Hit rates obtained on trials when a test face matching the colour of the co-presented neutrally valenced faces was presented.

Eye movement data

Fixations were defined as periods when the pupil could be detected and no saccade was underway. Saccades were defined as periods when two criteria were met: gaze position had both changed by more than 0.1 deg and was either accelerating by at least $8000^{\circ}/\text{sec}^2$ or exceeded a velocity of $30^{\circ}/\text{sec}$. Next, drift correction and fixation cleaning processes were applied to fixations on a trial by trial basis according to the procedure outlined in Appendix A. This corrected for small head movements that may have been made within each block of trials while accounting for changes in observation strategy (where to look first and what order to view stimuli in) used by participants. Trials were then excluded from analysis on a trial by trial basis if less than two fixations were successfully recorded during presentation of the study array or if total detected dwell time (sum time of all fixations) over the study array was less than 200 ms (indicating excessive blinks, gaze not on areas of interest at all or other temporary loss of detection of the eye across the majority of the trial). This resulted in removal of 1.2% of trials (as above). Four vertical rectangular (79 X 90 pixels) regions of interest (ROI) were defined; the center of each ROI was coincident with the center of each face in the study array. The mean number of fixations landing within each ROI, the mean dwell time (i.e., sum of durations of all fixations landing within each ROI), and the proportion of first fixations landing within each ROI were determined for each participant for each study array condition.

Repeated measures ANOVAs were conducted on d' data and each eye movement measure using study array condition (positive, negative, neutral) and singleton status (singleton, non-singleton) as within-subject factors. Planned comparisons used within-subject t-tests. Alpha levels were set at .05.

Results

Learning phase

To see whether the value codes were experienced differently, I examined how the preceding trial on phase 1 data affected response times (RT) on each immediately subsequent trial (see Figure 7.5). RTs were fastest after loss trials ($M = 830$, $SD = 366$), slowest after win trials ($M = 992$, $SD = 500$), and in between after neutral trials ($M = 901$, $SD = 406$). An ANOVA using preceding trial type as a within participant factor indicated a significant difference [$F(1.53, 45.74) = 7.79$, $p = .003$, $\eta_p^2 = .206$, Greenhouse-Geisser corrected], and Bonferroni corrected post hoc t-tests revealed that the difference between neutral- and loss-preceded trials was marginally significant ($p = .110$), as was the difference between neutral- and win-preceded trials ($p = .065$), while the difference between win- and loss-preceded trials was significant ($p = .011$).

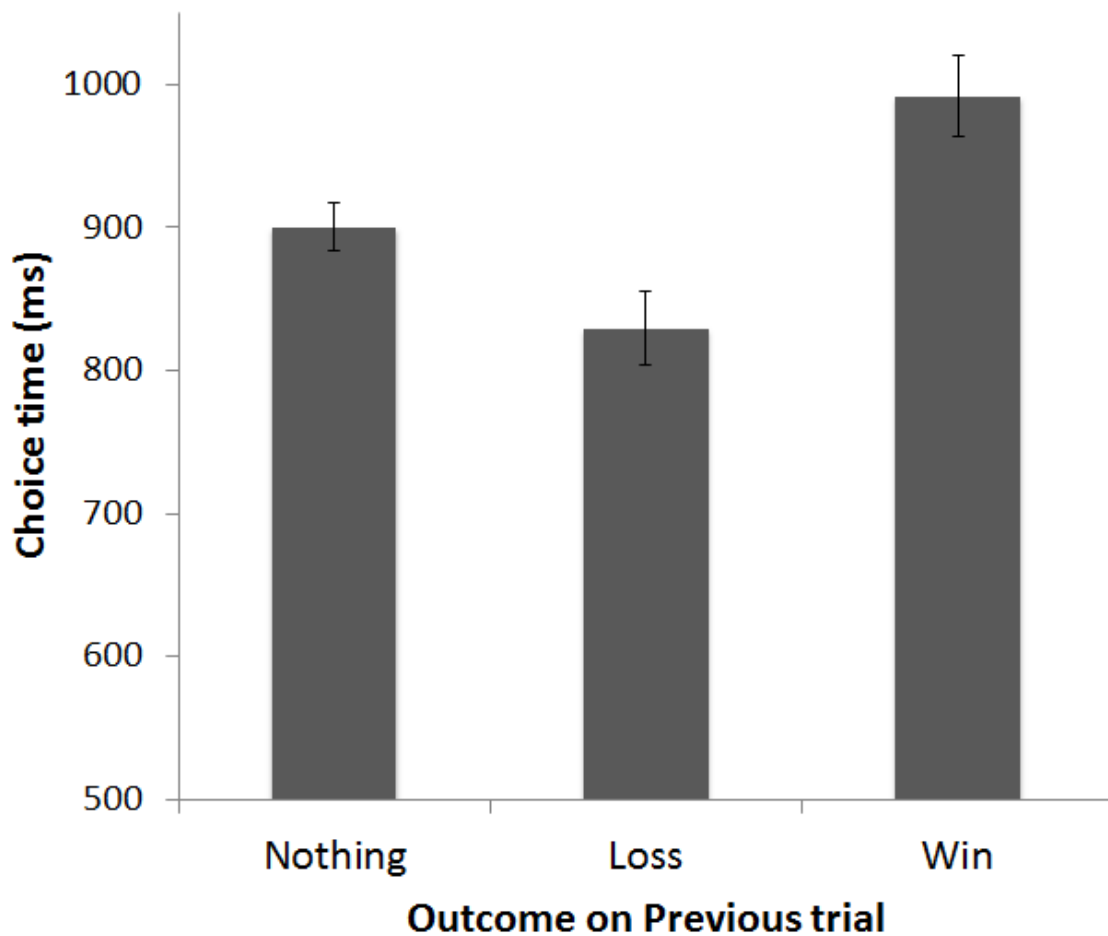


Figure 7.5. Effect of previous trial in learning phase 1 on immediately subsequent trial. Error bars represent within subject standard error of the mean (Cousineau, 2005; see Appendix D). Loss and gain trials affect RT on the next trial in opposite ways.

Effect of valence on WM for singletons

WM performance (d') for singleton faces tinted with a win associated colour ($M = 2.04$, $SD = .81$), was better than that for either loss ($M = 1.74$, $SD = .83$), or neutral associated ($M = 1.70$, $SD = .72$) colours (see Figure 7.6). A within

participants ANOVA indicated a significant effect [$F(2, 60) = 4.14, p = .021, \eta_p^2 = .121$]. Planned comparisons showed that win stimuli were significantly better remembered than either neutral [$t(30) = 2.682, p = .012$] or loss [$t(30) = 2.069, p = .047$] stimuli, while the difference between loss and neutral stimuli was non-significant [$t(30) = .278, p = .783$].




Figure 7.6. WM performance (d') for singleton faces. Error bars represent within subject standard error of the mean. Better performance was observed for faces tinted in colours previously associated with reward.

Effect of singleton valence on WM for co-presented (neutral) non-singletons

Average WM performance (d') for neutral outcome associated tinted faces presented alongside neutral singletons ($M = 1.36, SD = .59$), loss singletons ($M =$

1.24, SD = .80) and win singletons (M = 1.47, SD = .92) did not differ significantly [$F(1.665, 49.945) = 1.454, p = .242, \eta_p^2 = .046$], suggesting that the win singleton WM boost did not spill over to co-presented neutrally valenced faces, and was not therefore due to arousal (see Figure 7.7).




Figure 7.7. Effect of singleton status on co-present non-singletons. Error bars represent within subject standard error of the mean.

Effect of singleton on gaze³³

³³ Pupil diameter has been associated with arousal, with losses resulting in larger pupil dilations than equivalent gains; an expression of the general negativity bias or loss aversion (Granholt & Steinhauer, 2004; also see Hochman and Yechiam, 2011). Therefore, I also measured pupil diameter during phases 2 and 4 (when only a single stimulus was observed at any one time) to see if win, loss and no-change

The number of fixations made during the encoding (memory array presentation) interval to neutral ($M = 1.54$, $SD = .25$), loss ($M = 1.54$, $SD = .23$), and win ($M = 1.53$, $SD = .26$) associated singleton faces did not vary significantly [$F(2, 60) = .187$, $p = .830$, $\eta_p^2 = .006$], see Figure 7.8. Measures of dwell time (sum of length of all fixations) followed a similar non-significant pattern $F(2, 60) = .079$, $p = .924$, $\eta_p^2 = .003$. This suggests that the WM boost for the win associated tinted faces was not driven by selective attention capture at encoding.

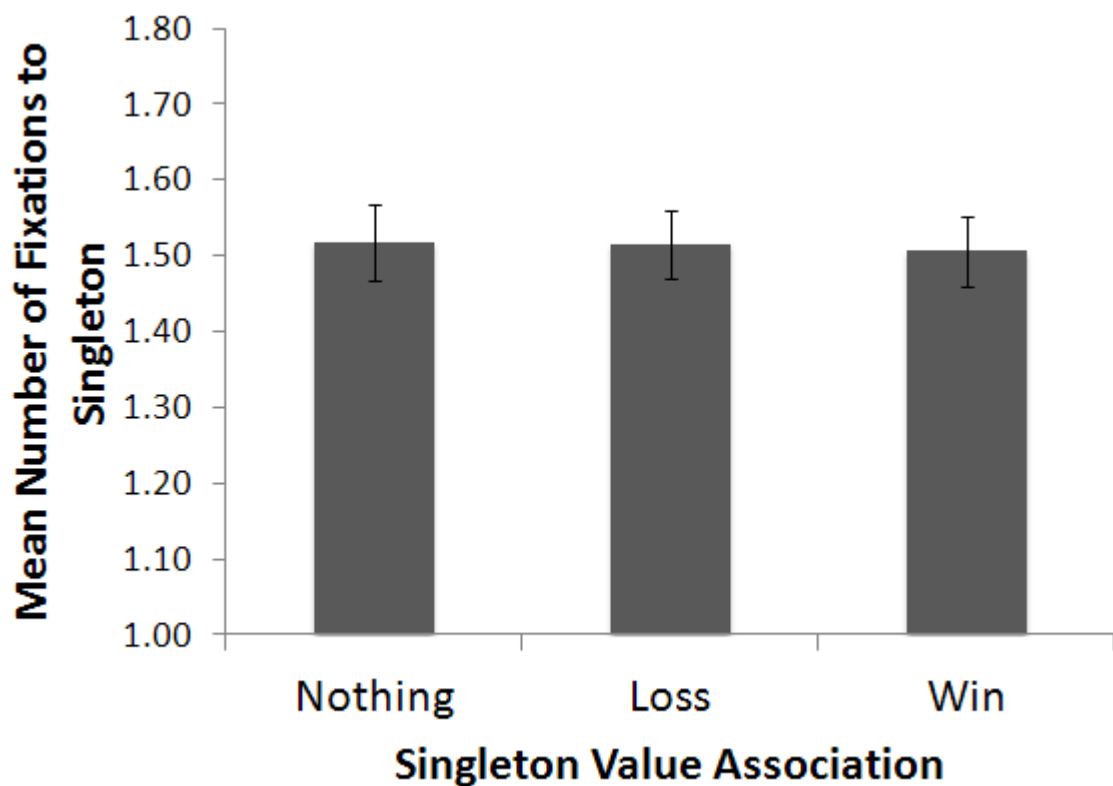


Figure 7.8. Fixations made to singletons. Error bars represent within subject standard error of the mean.

Discussion

associated faces would have any effect on pupil diameter. However, I did not observe any significant differences in pupil size between the three valence types.

Experiment 6 was conducted to ascertain whether the motivational salience of faces in a monetary win/loss scenario would have lasting effects on the way in which they were later encoded in a subsequent WM task in which previous motivational status was orthogonal to the task. The pattern of results observed suggests that singleton faces presented in colours previously paired with 'win' outcomes are better encoded in WM than those presented in colours previously paired with 'loss' outcomes or 'no change' outcomes. The boost in WM performance for win associated faces did not occur at the expense of other, neutral, co-presented non-singleton faces (as might be expected if it was driven by redeployment of selective attention to such singletons), nor did this performance boost result in facilitation of such faces (as might be expected if it resulted from generalised arousal). Further evidence that the boost to WM performance for singleton faces presented in positive association colour was not caused by a redeployment of overt attention comes from the observation that the prior value associations of singleton faces did not determine the extent to which they attracted and held gaze.

Note that the value associations of stimuli were orthogonal to the WM task. In other words, the difference in how stimuli of different value association were processed was driven by a stimulus dimension (prior pairing) that was task-irrelevant. This is indicative of an inability to fully discount the effects of value association on performance on a task in which it is no longer relevant and suggests that the WM boost observed for faces presented in a previously positively associated colour seems to represent an effect of 'retained' value

association³⁴ rather than an effect of goal-based motivational salience. This effect of value association on WM parallels the effect of value association on attention capture (Anderson et al., 2011a, 2011b) discussed above.

Also speaking against the notion of a purely goal-based motivational salience effect is the observation that negative value codes did not lead to a facilitation of WM for their associated face stimuli. If a purely motivational salience account could explain our results then we would expect both positive and negative value codes to result in facilitation of WM since both valences of value code signify that a change in state of the environment is about to occur and both therefore have motivational importance in terms of how the observer should best interact with the world around them.

However, it is not so easy to fully discount motivational salience as the underlying cause of the positive colour associated face WM boost. Although value association (as signified by colour) was orthogonal to the WM task, this is not to say that participants were able to treat it as such. Indeed, the finding that participants are unable to exclude from automatic attention capture high value laden distractors even when such distractors are inconspicuous (Anderson et al., 2011a) can be explained in one of two ways: Either value association leads directly to an automatic capture of attention (the option offered by Anderson et al., 2011a), or perhaps value association leads to a reinstatement of previously held goal states.³⁵ I discuss this possibility further in the chapter discussion below.

³⁴ Value association being 'retained' in the sense that it is still present after it ceases to be of immediate utility (i.e. is no longer task relevant). By contrast, goal-based motivational salience refers to situations in which the value association of the stimulus in question must be processed to perform on the task (e.g. when the task is to choose the item of highest reward association or lowest punishment association).

³⁵ While it was not necessary to explicitly form a goal during the learning task (since the learning paradigm was a classical rather than operant learning task) it is possible that goal like states are an efficient way in

A possible confounding variable in Experiment 6 is that faces with neutral value association were presented more frequently in the WM task (i.e. presented as non-singletons as well as singletons). If familiarity caused lack of engagement with stimuli then this could account for the lower WM performance associated with memory for neutral singleton faces relative to gain associated singletons, although this would not account for the better performance of gain over loss associated singleton faces. Additionally, it is possible that the WM boost associated with gain associated faces was driven by a generalised arousal boost at retrieval (i.e. in the presence of the gain associated probe stimulus) rather than at encoding or during maintenance. To rule out these possibilities, I conducted an experiment that used uniformly valenced memorisation arrays (all faces of the same value associated colour on each trial). On half of the trials in this experiment the test (probe) face was of the colour of the faces in the memory array, on the other half it was rendered in grayscale. In this way, I was able to investigate whether an arousal boost at retrieval could fully account for the increased WM performance associated with gain related faces, and whether arrays of all gain faces were better remembered than arrays of neutral or loss associated faces. Note that the replacement of coloured test faces with grayscale faces had to be accompanied by a switch to uniform arrays since singleton status in the memory arrays in Experiment 6 was defined by colour and the ability to discriminate between singleton statuses (whether a singleton or not) would otherwise be lost with the introduction of a greyscale probe.

which to realise value codes in the brain since the upweighting accorded to a positive value stimulus in terms of its ability to capture attention (Anderson et al., 2011a, 2011b) and, here, to facilitate WM is likely to be of use to the observer in a later situation in which the stimulus then forms part of an 'approach' goal.

To anticipate, surprisingly for uniform arrays the gain associated colour boost was not observed.

EXPERIMENT 7

Method

The method for Experiment 7 was identical to that of Experiment 6 except as follows.

Participants

Forty-eight³⁶ participants (36 females, mean age = 20 years, s.d. = 3.5 years) participated in Experiment 7.

Procedure

Eye movements were not recorded and no chinrest was used.³⁷ In phase 3, memory arrays comprised four faces each of the same colour (see Figure 7.4). For each participant one loss associated, one gain associated and one neutral (no outcome) associated colour were used to form these uniform arrays (the additional neutral associated colour learnt in phase 1 was not used in this phase). Phase 3 was split into two blocks, '3a' and '3b', the order of which was fully crossed with colour to value pairings and counterbalanced between participants. In phase 3a test faces were of the same colour as faces in the memory array for their

³⁶ Initially 24 participants participated. When I failed to observe the anticipated WM boost to positively value associated coloured faces I repeated the experiment with another 24 participants to confirm that no effect was observable. Results for the first and second batches of participants were similar with no significant differences in variables of interest observed in either batch. Therefore, data for the full 48 participants are reported here as a single data set.

³⁷ This allowed participants to speak aloud, permitting a concurrent verbal suppression task to be run.

respective trial. In phase 3b test faces were rendered in grayscale. Each of phases 3a and 3b comprised ninety-six trials and were preceded by twelve practice trials. To suppress verbal encoding and working memory, two letters were presented at the beginning of each block and the participant repeated them aloud (monitored by the experimenter) throughout the block. At the end of the block, two letters were presented and the participant indicated whether they were the same as or different from those they had been repeating.

Data Analysis

One participant's data indicated clear disengagement from the task (chance performances throughout) and these were removed from further analysis.

Learning phases

In interpreting learning performance we must bear in mind the difficulty participants had in completing self-report measures of learning in Experiment 6, and the similar difficulty observed in Experiment 7. This was indexed both by participants self-report of the way they chose to respond to the task in phases 2 and 4 (an unwillingness to acknowledge in responses the negative association of the negatively associated colour, see Experiment 6 above for examples) and by the pattern of performance on the phase 2 and 4 prediction tasks mirroring these statements (phase 2 prediction proportions correct: Neutral $M = .87$, $SD = .24$; Loss $M = .81$, $SD = .27$; Gain $M = .93$, $SD = .18$; phase 4 proportions correct: Neutral $M = .92$, $SD = .18$; Loss $M = .87$, $SD = .24$; Gain $M = .91$, $SD = .24$). I therefore took as my main measure of learning in Experiment 7 the effect of preceding trial outcome on choice time in phase 1 trials in the same way as is reported above for Experiment 6.

WM phases

Less than 0.01% of data points in phase 3a (colour probes) and less than 0.01% in phase 3b (grayscale probes) were excluded from analysis due to RTs less than 200 ms. A repeated measures ANOVA was conducted on d' scores using array condition (positive, negative, neutral) and test face type (colour, grayscale) as within subjects factors.

Results

Learning phase (phase 1)

To see whether the value codes were experienced differently, I again examined how the preceding trial affected response times (RT) on each trial (see Figure 7.9). RTs were fastest after neutral trials ($M = 751$, $SD = 299$), markedly slowest after win trials ($M = 959$, $SD = 422$), and in between after loss trials ($M = 791$, $SD = 325$). An ANOVA using preceding trial type as a within participant factor indicated a significant difference [$F(1.63, 74.88) = 14.93$, $p < .001$, $\eta_p^2 = .245$, Greenhouse-Geisser corrected], and Bonferroni corrected post hoc t-tests revealed that the difference between neutral- and loss-preceded trials was not significant ($p = .515$), while the differences between neutral- and win-preceded trials ($p < .001$), and between loss- and win-preceded trials were significant ($p = .002$). This is indicative of good learning of the win contingencies.




Figure 7.9. Effect of previous trial in learning phase on immediately subsequent trial. Error bars represent within subject standard error (Cousineau, 2005; see Appendix D). Win trials lead to a delayed response on the subsequent trial.

WM phases (phases 3a and 3b)

Three participants made a single mistake each on the verbal suppression task (across both phases 3a and 3b). All others were at 100% performance on this task.

When probes had the same colour as the preceding memorisation array (phase 3a) performance was slightly poorer on win associated colour trials ($M = 1.03$, $SD = .59$) than on neutral ($M = 1.11$, $SD = .66$) or loss associated trials ($M =$

1.11, $SD = .64$). When probes were rendered in grayscale (phase 3b), performance was slightly better on trials where the study array had been rendered in a neutral associated colour ($M = 1.17$, $SD = .64$) rather than a loss ($M = 1.04$, $SD = .65$) or win ($M = 1.06$, $SD = .69$) associated colour (see Figure 7.10).




Figure 7.10. Change detection performance across the six experimental conditions. Error bars represent within subject standard error of the mean. d' scores are similar across experimental conditions.

A within subjects ANOVA on d' scores with array condition (positive, negative, neutral) and test face type (colour, grayscale) as within subjects factors and block order (colour probe block first, grayscale probe block first) as a between subjects factor revealed a non-significant effect of array condition [$F(2, 90) = .925$, $p = .400$, $\eta_p^2 = .020$], a non-significant effect of test face type [$F(1, 45) = .005$, $p =$

.942, $\eta_p^2 < .001$], and a non-significant interaction between these factors [$F(2, 90) = .355, p = .702, \eta_p^2 = .008$]. The effect of block order was non-significant [$F(1, 45) = .106, p = .746, \eta_p^2 = .002$] as were interactions of block order with array condition [$F(2, 90) = .170, p = .844, \eta_p^2 = .004$], and test face type [$F(1, 45) = .035, p = .853, \eta_p^2 = .001$]. The triple interaction was also non-significant [$F(2, 90) = .084, p = .919, \eta_p^2 = .002$].³⁸

To investigate the possibility that the absence of a significant effect of value associated colour on WM performance might simply be the result of some participants not properly attempting the task, I split the participants into higher and lower performing participants using a median split on overall (combined probe types) d' . Even among the higher performance group there was no effect of value association [d' scores by study array type: Neutral $M = 1.43, SD = 0.47$; Loss $M = 1.38, SD = .49$; Gain $M = 1.35, SD = .59$; repeated measures ANOVA on array type: $F(2, 46) = .180, p = .836, \eta_p^2 = .008$]. Thus, it appears that the memory boost occasioned by positive colour associated singletons to themselves (Experiment 6) does not necessarily occur when such positive associated stimuli are not presented as singletons, perhaps because in this context colour is an irrelevant feature that carries no useful information (see chapter discussion).

Discussion

While positive value associated stimuli seem to receive an 'item specific' WM boost when they are singletons (Experiment 6), this boost is absent when they are to be memorised as part of a uniform display in which all items share the

³⁸ A similar pattern of null effects was observed for RTs.

positive value associated tint. This is a surprising finding and marks a departure from the pattern of results observed for angry faces (Chapter 5), where within uniform displays of all angry faces each face received a WM boost, and suggests that perhaps the underlying mechanism for the boost to angry singleton faces and positive value-laden singleton faces is different. Whereas I have so far suggested an account whereby the emotion of value association of face stimuli may determine their motivational salience and this salience effect affect the way they are treated by WM, perhaps the qualitatively different nature of signals of anger and 'positivity' result in different types of WM boost. I discuss this in more detail in the chapter discussion below. To more fully understand the effect of value association on the way a stimulus is encoded into WM I conducted Experiment 8.

EXPERIMENT 8

Thus far, I have investigated the effects that memorisation and maintenance of representations of motivationally salient stimuli have on representation of other concurrently presented stimuli. However, in order to disentangle the effects of biased competition at encoding (the competing for attentional and perhaps consolidatory resources among concurrently presented stimuli) from the effects of maintaining a representation in WM on other to-be-remembered stimuli, it is necessary to allow one WM representation to be fully formed before a requirement to add an item to WM is initiated. Here I ask what effect a value laden visual WM representation may have on the ability to perform a secondary WM task. One possibility is that value laden representations take up more 'memory space' than non-value laden representations and thus have a

negative impact on the ability to perform a secondary WM task.³⁹ Alternatively, maintenance of a value laden representation may cause an associated maintenance of a heightened state of short term arousal and vigilance (Yeichiam & Hochman, 2013) which may facilitate WM performance on a secondary task. Alternatively it may have no effect.

Another interesting question is whether a value laden WM representation is more 'robust' (better at surviving interference) than a non-value laden one. One way of investigating this is to put variable amounts of 'pressure' on the maintenance of WM representations. By varying the difficulty (set size) of a WM task that must be performed while maintaining a value laden visual WM representation it is possible to differentially recruit WM resources for the secondary task and so differentially constrain the WM resources available to maintain the value laden representation. Experiment 8 was designed to answer these questions by first imbuing value associations into a selection of face stimuli and then presenting these stimuli as memoranda in a WM sandwich task where face stimuli were to be held in memory while a second concurrent visual WM task of varying difficulty (set sizes) was performed. In this way, the effect of holding face stimuli of different value associations in WM on a secondary task could be assessed, as could the effect of loading WM resources on maintaining representations of face stimuli of different value association.

³⁹ Since the boost to processing resources in Experiments 1, 2 and 6 seems to occur in an item specific fashion we would not expect this to be the case. However, in Experiments 1, 2 and 6 the allocation of additional resource was determined in a situation where there was competition for processing resource between the various stimuli in the memorisation array. In a situation where no such competition occurs during encoding (i.e. where only a single value laden item is encoded) it is possible that rather than bringing online additional resource from an otherwise 'idle' pool, the value of a stimulus might instead determine how much 'natural resource' (that is, resource from an 'unboosted' pool; see Chapter 8) it attracts.

Method

Participants

Participants were randomly assigned to the Gain or Loss group. 33 and 38 participants were assigned to the gain and loss groups, respectively (Gain group: 27 female; mean age = 20 years, s.d. = 1.8 years; Loss group: 24 female; mean age = 21 years, s.d. = 4.4 years). Participants reported normal neurological histories and normal or corrected to normal vision. Informed consent was obtained from all participants.

Apparatus

The same apparatus as in Experiment 7 were used. Viewing distance was approximately 60 cm.

Stimuli

For the learning and face memory tasks, four male faces with neutral expressions and hair and neck showing were selected from the Karolinska Directed Emotional Faces image bank (KDEF, Lundqvist, Flykt, & Öhman, 1998). These were rendered in greyscale as bitmaps with a bit depth of 8 and re-sized to 105 (width) x 150 (height) pixels, so that they subtended approximately 3.2 by 4.1 degrees of visual angle. During the learning task, two faces were presented on a white field, arranged above and below a central fixation point along the vertical meridian, separated by 2.2 degrees.

For the squares memory task, filled colour squares were presented on a uniform grey background. Each square was rendered as a bitmap with a bit depth of 24 and side length of 29 pixels (approximately .84 degrees). Square colours

were: Black, blue, brown, green, orange, pink, purple, red, turquoise, and white. RGB coordinates for the squares and background are given in Appendix B. On each trial, different coloured squares were presented in random locations within a 350 by 350 pixel (10.5 horizontally by 9.6 degrees vertically) area in the centre of the screen. Minimum inter-square separation was 130 pixels (approximately .84 degrees).

Design

The experiment comprised two phases, a learning phase and a WM test phase. Half the participants earned points⁴⁰ during the learning phase (Gain Group) and remaining participants lost points from an initially awarded starting points total of 1000 points during the learning phase (Loss group). In all other ways, both groups were treated similarly. Stimuli in the learning phase were assigned one of four Outcome Probabilities (1.00, .70, .30 and .00). The measure of interest in the learning phases was probability of correct choice, i.e., choosing the item that either maximized the chance of a win or minimized the chance of a loss. Each then served as the to-be-remembered face in the memory test phases, making Outcome Probability a within-subject variable.

The WM sandwich task used a 2 (Group; Gain, Loss) x 4 (outcome probability associated with face presented; 1.00, .70, .30 and .00) x 3 (set size; two, four, six) x 2 (array probe; same, different) x 4 (face probe location; left, centre left, centre right, right) design.

Procedure

⁴⁰ Points were redeemable after the experiment for money at a value of one point = one penny. This was done to better motivate participants and maximise the value associated with winning and losing points.

Learning phase

For the learning phase, I followed a procedure similar to that of Pessiglione et al. (Pessiglione, Seymour, Flandin, Dolan, & Frith, 2006). Each trial began with a 1000 ms fixation cross, followed by a pair of vertically arranged faces that remained present until response (see Figure 7.11). Participants chose the upper or lower faces using the 't' or 'b' key, respectively, having been instructed to maximize winnings or minimize losses (for Gain and Loss Groups, respectively) by choosing one of the two faces on each trial. Probability of an outcome (win, loss) depended on the face chosen and on the outcome probability assigned to each face at the beginning of the experiment. Outcome probabilities were 100% (100-Win, 100-Loss); 70% (70-Win, 70-Loss); 30% (30-Win, 30-Loss) or 0% (0-Win, 0-Loss). The default was no outcome in all cases.⁴¹ Face-to-outcome probability assignment was counterbalanced across participants to avoid stimulus effects and remained unchanged throughout the experiment. After choice, the face pair remained visible and the word 'GAIN', 'LOSS', or 'NOTHING' appeared at fixation for 1000 ms accompanied by a 'chimes', 'buzz' or no sound respectively, presented through headphones. A running points total was updated (+10 points for a win, -10 points for a loss) and shown after each choice was made. Participants in the Gain group started with zero points and those in the Loss group with 1000 points. The next trial started automatically.

On successive trials, each face appeared equally often in each location and was paired equally often with each other face. Trial order was randomised for

⁴¹ The rationale for using such differentially predictive stimuli is that the associability of a stimulus is determined by how accurately it predicts reinforcement, in this case positive or negative change in points total (Mackintosh, 1975).

each participant. The learning phase comprised 144 trials, such that each face was presented 72 times.

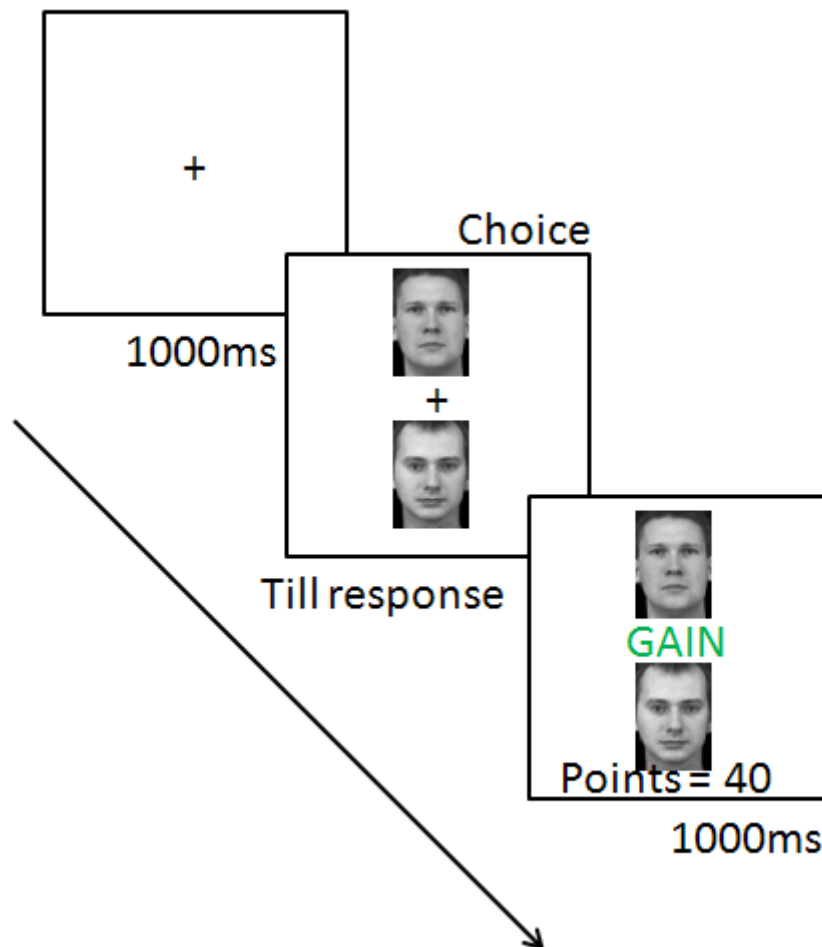


Figure 7.11. Procedure for the learning phase in the gain group. On the trial shown the participant selects a face and receives a gain (+10 points) outcome.

After all choice trials were complete, participants were presented with each face one more time and asked to rate 'how lucky' the face was on a scale of zero to nine using number keys on the keyboard. The order of presentation of

faces by Outcome Probability in this question phase was counterbalanced across participants within each group.

WM sandwich task

Each trial followed the following sequence: A 1000 ms fixation cross preceded the presentation of a single face stimulus presented centrally for 1000 ms. Next, another 1000 ms fixation cross preceded the initial coloured squares study array, which was present for 100 ms and comprised two, four or six different squares. Next, a 900 ms blank screen preceded the second (probe) coloured squares array, presented for 2000 ms. Probe arrays were identical to initial arrays except that on 50% of trials one of the squares changed colour between the two arrays ('different' trials). A prompt screen displaying the words 'SAME OR DIFFERENT' followed, and participant responded here using the 's' key if they thought no change had occurred between the initial and probe arrays, or the 'k' key if they detected a difference. The prompt screen was presented until response. Finally, there followed a test screen for memory of the face. This contained the text 'Which face have you just seen?' and after 1000 ms the four possible faces were presented in a random order such that they occupied four locations along a horizontal grid below this text. Each location was labelled from 1 to 4 sequentially from left to right. Participants responded using the number keys '1' to '4' at the top of their keyboard, indicating which face was the one they had seen at the start of the trial. Pressing the space bar then initiated the next trial. The trial sequence along with the verbal suppression task described below is shown in Figure 7.12.

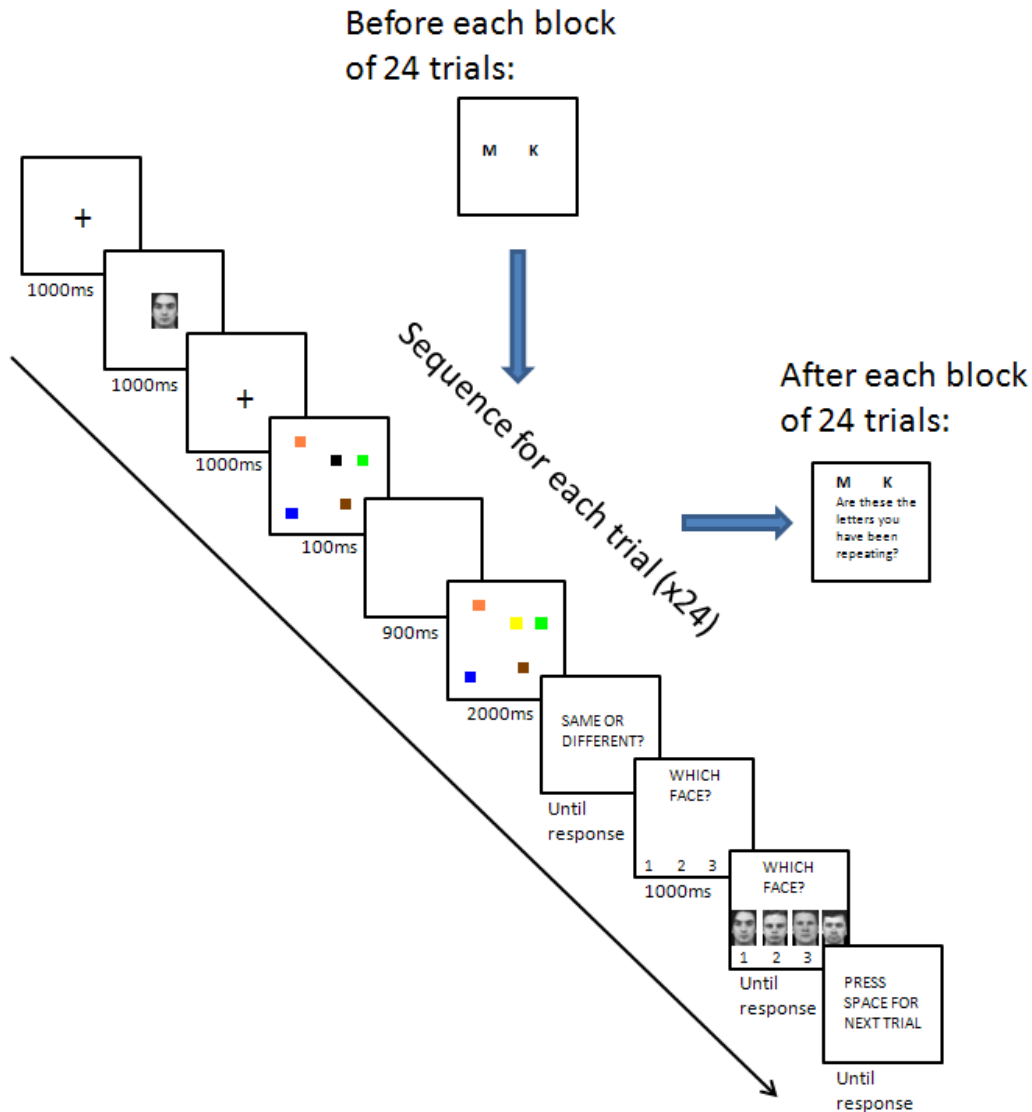


Figure 7.12. Procedure for the sandwich and verbal suppression tasks.

On the trial shown the correct response is to select the far left face (press '1') on the face WM task and select 'different' on the coloured squares WM task

A practice phase of ten trials preceded the main experiment. Practice trial types were randomly selected from the twenty four possible trial types. The main experiment had five blocks of trials, each comprising twenty four trials (one instance of each trial type per block). An enforced rest of 30 seconds occurred

between blocks. A verbal suppression task (repeating two letters given at the start of each block aloud and reporting whether a second pair of letters given at the end of the block matched these) ran throughout, monitored by the experimenter.

Data Analysis

Trials where reaction times on either the coloured squares task or the face WM task were lower than 200 ms were rejected as anticipation errors. Trials where reaction times to the coloured squares task were in excess of 3000 ms were also rejected, partly as long reaction times might indicate distraction on the task, but also because an extended reaction time here resulted in a longer retention interval for the face to be remembered. In total this resulted in the removal of 6% of data points from the Gain group data and 7% of data points from the Loss group data.

Examination of the remaining WM data revealed one participant in the Loss group with too few trials remaining to accurately analyse data. Examination of the learning data revealed two participants in the Gain group and one in the Loss group who had failed to engage with the task (having approximately chance performance). These participants were removed from further analysis.

To establish that learning of the face reward/loss outcome contingencies had taken place, tests of orthogonal within subjects contrasts were performed on face ratings in the learning data sets for the Loss and Gain groups separately.

D' scores for the face memory task were calculated for each condition according to the procedure outlined in Chapter 4 for M-AFC tasks using conversion tables (Hacker & Ratcliff, 1979). D' scores for the coloured squares

task were calculated using the standard formula for present/absent tasks [$z(H)$ - $z(FA)$]; see Chapter 4].

To examine the effect of amount of information entered into WM (indexed by coloured square task set size) on retention in WM of expected value (EV) faces, a repeated measures ANOVA with Outcome Probability (1.00, .70, .30 and .00) and coloured squares task set size (2, 4, 6) as within subjects factors and valence (Gain group; Loss group) as a between subjects factor was conducted on d' scores for WM for EV faces, using only data from trials where response on the coloured squares WM task was correct (i.e. those trials where evidence existed the coloured squares array had been memorised).

To examine the effect of retaining in WM faces of different EV on ability to do the coloured squares task, a repeated measures ANOVA using the same factors was conducted on d' scores for WM on the coloured squares task using only data from trials where response on the faces WM task was correct (i.e. those trials where evidence existed that the EV face was present in WM). Planned comparisons used within subject t tests. Alpha levels were set at .05.

Results

Learning tasks

In the Gain group, faces that were 100% predictive of winning were rated luckiest ($M = 8.55$, $SD = 1.06$) followed by those with 70% ($M = 6.81$, $SD = 1.19$), 30% ($M = 3.87$, $SD = 1.16$) and 0% ($M = 2.06$, $SD = .93$) outcome predictiveness respectively. In the Loss group, a trend in the opposite direction was observed;

faces that were 0% predictive of losing were rated luckiest ($M = 8.33$, $SD = .29$) followed by those with 30% ($M = 5.56$, $SD = 1.70$), 70% ($M = 3.86$, $SD = 1.68$) and 100% ($M = 3.08$, $SD = 2.08$) outcome predictiveness respectively (see Figure 7.13).




Figure 7.13. Luckiness ratings for the various outcome predictive faces in the Loss and Gain groups. Loss predictive faces are rated with increasing luckiness as their predictiveness of a loss event decreases. Gain predictive faces are rated with increasing luckiness as their predictiveness of a gain event increases. Error bars represent standard error of the mean.

Tests of orthogonal within subjects contrasts (trend analysis) revealed a linear [$F(1, 30) = 680.1$, $p < .001$, $\eta_p^2 = .958$] and cubic [$F(1, 30) = 4.36$, $p = .045$, $\eta_p^2 = .127$] trend in the Gain group data and a linear [$F(1, 35) = 85.2$, $p < .001$, $\eta_p^2 = .709$] and quadratic [$F(1, 35) = 15.8$, $p < .001$, $\eta_p^2 = .310$] trend in the Loss group data. Of particular relevance here are the two linear trends observed, indicative

as they are of ability to correctly identify in general which faces are most associated with reward or which are least associated with punishment.

To provide converging evidence that face contingencies had been correctly learnt, I examined the selection choices made during the learning experiment. If faces were learnt correctly then we would expect to see a greater than chance selection of the optimal choice face when individual pairs of faces were presented. Figure 7.14 shows the proportion of choices that were optimal when each face pair was presented. Within the Gain group, the 100-win face was more frequently chosen than the 70-win face when paired with it (proportion optimal choices made: $M = .84$, $SD = .25$) and was also optimally chosen when paired with the 30-win ($M = .94$, $SD = .13$) and 0-win ($M = .94$, $SD = .09$) faces. The 70-win face was optimally chosen vs. both the 30-win ($M = .81$, $SD = .20$) and 0-win ($M = .81$, $SD = .17$) faces. The 30-win face was optimally chosen vs. the 0-win face ($M = .60$, $SD = .17$). Within the Loss group, each pairing of faces likewise resulted on average in the optimal choice being made (proportion optimal choices made: 100 vs. 70, $M = .60$, $SD = .16$; 100 vs. 30, $M = .82$, $SD = .17$; 100 vs. 0, $M = .96$, $SD = .11$; 70 vs. 30, $M = .77$, $SD = .22$; 70 vs. 0, $M = .96$, $SD = .09$; 30 vs. 0, $M = .92$, $SD = .20$).




Figure 7.14. Mean proportion of optimal choices made for each face pair.

Error bars represent standard error of the mean.

It is apparent from Figure 7.14 that the majority of face contingencies were learnt reasonably well. However, participants in the Gain group had trouble distinguishing between the 30-win and 0-win faces while participants in the Loss group had trouble distinguishing between the 100-loss and 70-loss faces. This reflects a general weakness in learning the contingency pairing of the second least optimal face in each group, as revealed by both the relative weakness of performance in choosing such faces when paired with the third least optimal face (4th red and 4th blue bars from left in Figure 4.14)⁴², and by the quadratic function

⁴² Note that the discrimination of the second least optimal faces from the most optimal faces (100-win vs. 30-win in the gain group, 70-loss vs. 0-loss in the loss group) was strong in both groups. However, this would seem to reflect a strong learning of the most optimal face contingencies and a reliance on this knowledge alone to make selection when such pairings were presented (i.e. when the most optimal face – the 100-win or the 0-loss – is presented it can be safely selected without need to attempt to identify the outcome probability of the face it is paired with).

found in rating data for the Loss group.⁴³ Post experiment informal interviews with participants revealed that they generally found these second least optimal faces the hardest to learn. This is likely due to the relatively few times such faces were chosen in the learning task (being relatively unrewarding, they were generally avoided as a choice except when paired with the least positive outcome face, resulting in fewer opportunities to learn that these faces were not 100% predictive of negative outcome).

WM sandwich task

No participants made any mistakes on the verbal suppression task in either group.

Coloured squares maintenance task

Means and standard deviations for performance (d') on the coloured squares task by condition are given in Table 7.1 below. A repeated measures ANOVA with Set Size and Outcome Probability as within subject factors and Group as a between subjects factor revealed that the overall effect of Set Size was significant [$F(1.79, 116.2) = 160.9, p < .001, \eta_p^2 = .712$, Greenhouse-Geisser corrected], reflecting poorer performance as set size increased, while effects of Group and Outcome Probability and all interactions were non-significant (all p values $> .1$). That the effect of set size was significant is hardly surprising given the widely reported increasing difficulty of this task when performed at higher set sizes (e.g. Luck & Vogel, 1997), but is important here since it indicates that this task successfully loaded WM to a greater extent at higher set sizes.

⁴³ This is indicative of a 'sharpening' of discrimination between faces as the linear trend progresses (see Figure 7.13).

Outcome Probability		100			70			30			0		
Set Size		2	4	6	2	4	6	2	4	6	2	4	6
Gain Group	Mean	2.39	2.01	1.54	2.50	1.89	1.25	2.22	2.07	1.39	2.28	1.80	1.52
	StdDev	0.32	0.62	0.74	0.21	0.66	0.86	0.45	0.48	0.81	0.36	0.58	0.88
Loss group	Mean	2.40	1.89	1.47	2.25	1.95	1.28	2.36	2.00	1.43	2.30	2.03	1.21
	StdDev	0.34	0.60	0.76	0.56	0.67	1.02	0.39	0.52	0.75	0.45	0.63	0.61

Table 7.1. Mean and standard deviation d' scores for performance on the coloured squares WM task.

EV face maintenance task

Means and standard deviations for performance (d') on the Faces WM task by condition are given in Table 7.2 below. Given that the manipulation of set size differentially loaded WM (see above), we would have expected to see an effect of set size on ability to retain EV faces. Surprisingly, a repeated measures ANOVA with Set Size and Outcome Probability as within subject factors and Group as a between subjects factor revealed that the overall effect of Set Size on EV face WM non-significant ($p > .8$), while the effect of Outcome Probability approached significance [$F(3, 195) = 2.23, p = .086, \eta_p^2 = .033$], the interaction of Outcome Probability with Group was significant [$F(3, 195) = 3.51, p = .016, \eta_p^2 = .051$], and the three way interaction of Outcome Probability with Set Size and Group was also significant [$F(6, 390) = 2.66, p = .015, \eta_p^2 = .039$].⁴⁴ The effect of Group and double interactions of Set Size with Group and Set Size with Outcome Probability were not significant (all p values $> .35$).

⁴⁴ This triple interaction was problematic to investigate further since conditionallising the face WM task performance data on correct performance on the coloured squares task resulted in too few trials within each sub-condition (EV type by set size within group) to allow robust analysis at this most granular level. Therefore, I did not pursue the triple interaction further.

Outcome Probability		100			70			30			0		
Set Size		2	4	6	2	4	6	2	4	6	2	4	6
Gain Group	Mean	3.05	3.47	3.10	3.08	2.91	2.84	2.86	2.91	3.08	2.91	2.68	2.77
	StdDev	0.81	0.69	0.94	0.94	1.05	1.05	1.02	0.84	0.94	0.91	0.89	1.00
Loss group	Mean	2.89	2.50	3.00	2.77	2.91	2.70	3.14	2.97	2.90	2.80	2.92	2.81
	StdDev	0.99	0.85	0.94	0.98	1.02	0.96	0.93	0.88	0.96	0.99	0.92	0.93

Table 7.2. Mean and standard deviation d' scores for performance on the Faces WM task.⁴⁵

The significant effect of Outcome Probability by group indicates that it is appropriate to treat the two groups separately and look at the pattern of data within each.⁴⁶ Within the gain group, the 100-win face ($M = 3.00$, $SD = .61$) was associated with the highest level of performance, followed by the 70-win face ($M = 2.77$, $SD = .82$), then the 30-win face ($M = 2.68$, $SD = .57$), and then the 0-win face ($M = 2.6$, $SD = .67$; see Figure 7.15).

⁴⁵ Note the apparent discrepancy between the estimation of sensitivity (d') obtained by combining together the performance values of each set size within each expected value (EV) condition vs. the mean combined values reported in the text and in Figure 7.15. This illustrates a constraint of the d' statistic; ceiling effects within sub-conditions may give rise to an overestimation of d' within such sub-conditions especially when trial numbers are relatively low as they are at the level of granularity represented by the cells in Table 7.2. The degree of such overestimation depends on the correction term applied to ceiling and floor values (see Chapter 4 for discussion).

⁴⁶ An alternative approach would be to 'flip' the data for one of the groups so that Outcome Probability (the likelihood of a change in the status of the world – i.e. a change in points total on the learning phase) is replaced by valence. In this way the performance data for the 100-win and 0-loss faces would have the highest valence rating, the 70-win and 30-loss the next highest and so on. While this approach would lend statistical power to the investigation of a valence effect across all data points, it is not theoretically justified if we consider reward and punishment contingencies to reflect different cognitive processes.

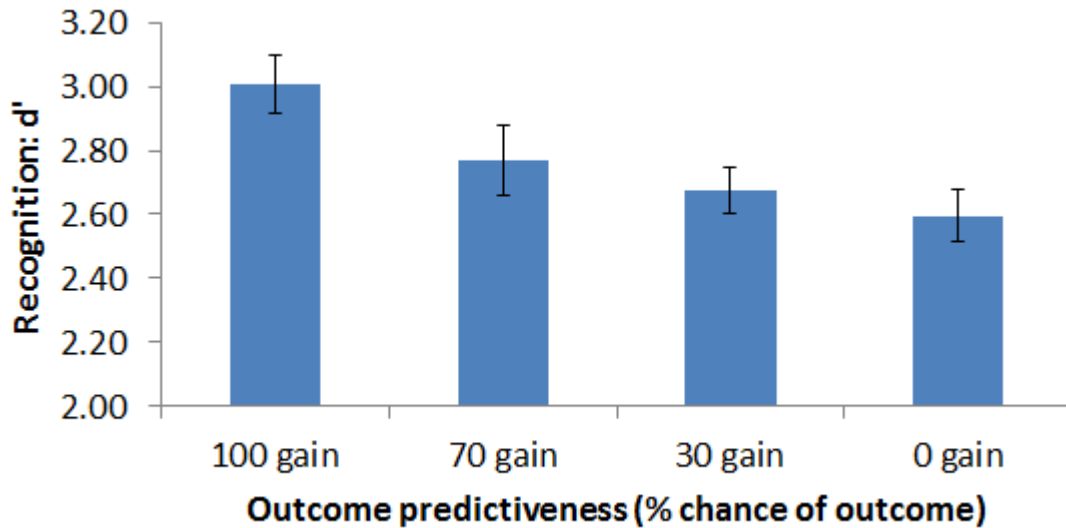


Figure 7.15. Performance on the EV Faces task within the gain group.

Error bars represent normalised standard error of the mean.

A repeated measures ANOVA on Outcome Probability revealed a significant effect [$F(3, 90) = 3.02, p = .034, \eta_p^2 = .091$]. Within subjects contrasts indicated that a linear relationship best described the pattern of data [$F(1, 30) = 9.65, p = .004, \eta_p^2 = .243$]. Bonferroni corrected t-tests revealed that the difference between the 100-win and 0-win faces was significant ($p = .018$) while no other pairwise comparison were significant (all p values $> .125$). The linear trend observed in this data could not be accounted for by a speed/accuracy trade-off: Responses to the 100-win face were fastest ($M = 1767, SD = 352$), followed by responses to the 70-win face ($M = 1871, SD = 319$), then the 0-win face ($M = 1922, SD = 353$) and finally the 30-win face ($M = 1963, SD = 379$; see Figure 7.16).⁴⁷

⁴⁷ A repeated measures ANOVA on these response times revealed a significant effect of outcome probability [$F(3, 90) = 4.01, p = .009, \eta_p^2 = .120$]. Within subjects contrasts revealed a linear relationship in




Figure 7.16. Time taken in ms to make the present/absent face judgement (gain group only). Error bars represent normalised standard error of the mean.

Within the loss group, the 30-loss face was associated with the highest level of performance ($M = 2.79$, $SD = .66$), followed by the 0-loss face ($M = 2.65$, $SD = .70$) and then the 100-loss ($M = 2.59$, $SD = .64$) and 70-loss ($M = 2.58$, $SD = .67$) faces (see Figure 7.17).

the opposite direction to that observed for accuracy (d') on this task [$F(1, 30) = 14.28$, $p = .001$, $\eta^2 = .322$]. Bonferroni corrected pairwise comparisons revealed that the difference between the 100-win and 30-win faces was significant ($p = .002$), as was the difference between the 100-win and 0-win faces ($p = .040$).




Figure 7.17. Performance on the EV Faces task within the loss group.

Error bars represent normalised standard error of the mean.

A repeated measures ANOVA on Outcome Probability revealed no significant effects [$F(2.28, 79.7) = 1.70, p = .186, \eta_p^2 = .46$, Greenhouse-Geisser corrected], nor did linear, quadratic or cubic relationships describe the data well (all within subjects contrasts p values $> .121$).

Discussion

The purpose of this experiment was to investigate three questions; whether the outcome probability or expected value of a face representation would impact the degree of accuracy with which it could be remembered in WM, whether this outcome probability of a face encoded in WM would impact the ability to do a concurrent visual WM task, and whether short term memory representations of

faces with different outcome probabilities would be differentially impacted by loading WM with a concurrent visual WM task.

With regard to the first of these questions, whether prior outcome probability influences the way stimuli are encoded into WM, qualitatively different patterns of behaviour were observed when face stimuli had been predictive of gain or loss. When face stimuli had been differentially predictive of gain events, the strength of such predictiveness was observed to modulate WM performance in a linear fashion such that stimuli that had been more predictive of gain (i.e. higher valence stimuli) were better remembered than those less strongly predictive of gain events. No such pattern of performance was observed for stimuli that had been differentially predictive of loss events. This provides converging evidence (along with that reported earlier in this chapter) that positive but not negative learned stimulus associations bolster WM performance for such stimuli and indicates that such modulations of WM performance are graded according to the level of association as indexed by outcome probability.

Surprisingly, the coloured squares and face WM tasks did not seem to interact with one another. Despite differentially loading WM at different set sizes (as evidenced by poorer performance at higher set sizes), the coloured squares task was observed to have little impact on memory for face stimuli already encoded into WM. In a similar fashion, the outcome contingency associated with faces had no overall effect on WM for the coloured squares.⁴⁸ There are several possible explanations for the lack interaction between the two WM tasks. First, the two WM tasks used stimuli from two very different visual domains. While coloured square arrays comprised visually simple entities, face stimuli by contrast

⁴⁸ Though see discussion on triple interaction above.

were visually complex. One possibility is that representations generated when these two very different stimulus classes are encoded are stored in qualitatively different visual WM stores. In particular, the process by which representations of the two different types of stimuli are brought into being may be very different: Face stimuli had received prior exposures both in the learning task and in the preceding trials on the WM task, and therefore it is likely that LTM representations for them would have existed. This could facilitate a WM process whereby a temporarily activated LTM representation of a face presented on a trial could facilitate (or comprise) its representation in WM. This notion of temporarily activated LTM codes acting as part of the WM representation process is a crucial part of Cowan's model of WM (Cowan, 2001; also see Chapter 1). By contrast, it is far less likely that coloured square arrays can call upon such LTM codes; not only do they comprise simple components with little 'real-world' relevance outside the experimental setting (a status that would have discouraged formation of LTM codes for them), but squares were arranged in pseudo-random locations on each trial with the result that from trial to trial the spatial composition of the array was unique and unlikely to match to any LTM template. Thus, it is likely that the process of encoding a coloured squares array was constructive rather than 'retrievalistic' and so may have resulted in a qualitatively different type of memory trace than that for face stimuli. Additionally, the temporal nature of the task combined with the fact that the memoranda came from two discrete stimulus classes might have contributed to the lack of task interaction observed. In typical temporal WM tasks (e.g. N-back tasks, see Chapter 1) stimulus representations are presented one by one to the receiver (participant) and added one by one to a stimulus set for later memory test. However, such tasks typically use stimuli from

the same or similar domains (e.g. numbers, words, simple objects) whereas here I have used stimuli from two very different domains. It is possible that in my task participants fully consolidated the memory representation of the face stimulus before performing the coloured squares WM task, while in other temporal WM tasks a 'supra-object file'⁴⁹ for the encoded stimulus set might remain open while the task is being performed in order that additional stimuli can be added to it (allowing such an supra-object file to remain open during encoding rather than starting a new file for each new stimulus would allow chunking to take place). In this view, in most temporal WM tasks stimuli compete for space within a limited supra-object file defined resource pool, whereas in my task no such interference would take place as the face stimuli were encoded first on each trial and then a separate object file opened to receive input from the coloured squares task.

That the outcome probability of faces held in WM does not impact on WM performance on the coloured squares task (despite an effect of outcome probability on WM for faces in the gain group) can be accounted for both by the fact that the two tasks may have recruited different WM processes as discussed above and by the consideration that boosts to value faces (experiment 6) and emotional faces (Experiments 1 and 2) seem to occur in an item-specific fashion, with little effect on concurrently presented memoranda. Additionally, the lack of task interaction suggests that the high gain face WM advantage is not due to a generalised and sustained arousal boost.

One question that is not answered by this experiment is whether it is the positive valence of high gain EV faces that makes them better memoranda or

⁴⁹ An object file that receives as its input representations of multiple objects and holds them as a chunked whole.

whether it is the predictiveness of positive change in a world-state that is the dimension that determines memory performance for stimuli that predict positive events. Valence and predictiveness of outcome cannot be disentangled within the current data since rewards on win trials all had the same value, making predictiveness equivalent to valence.

CHAPTER DISCUSSION

In Experiment 6 we observed an item-specific WM performance boost to face singletons tinted in a colour previously associated with monetary gain events. In Experiment 7, in a WM task where study arrays of faces were always uniform with regard to the colour they were tinted with, no such boost was observed for positive colour tinted faces, regardless of whether probe faces were tinted in the colour of the study array or presented in grayscale. Together, these findings suggest that rather than value association having a direct effect on memorability, there is something special about being a singleton face of a positively associated colour that results in better short term memory for such items. Interestingly, the WM boost accorded to positive value singletons appears not to be driven by a preferential allocation of overt attention, as evidenced both by the fact that such singletons do not preferentially capture gaze relative to other, non-positively associated colour singletons and by the lack of a cost to WM performance for concurrently presented non-singleton faces relative to non-singleton faces presented alongside negatively associated or no-change associated colour faces. Therefore, the lack of a WM boost when all faces in the array are tinted in the positively associated colour (uniform displays; Experiment 7) is not likely to be due to a sharing of attention among the faces in the study array. Instead, we must

consider why the combination of both being a singleton and being associated with a gain event is necessary to achieve a boost to WM.

Why no effect of positive value association in uniform conditions (Expt 7)?

When study arrays comprised all positive tinted faces, all face stimuli present had equal 'potential utility' to the participant. By potential utility I mean that while outcome contingencies were orthogonal to the WM task demands, if outcome contingencies were to be reinstated then the utility of each face would be equal. It is this potential for value codes to be useful at a later time that allows the different value codes to be retained even when they are no longer being paired with outcome events (a common observation in classical and instrumental learning) and evidenced here by retention of value association in phase 4 for Experiments 6 and 7. It is worth considering how in day to day life such outcome values relate to acquisition of a reward. In a great many instances when a reward-predictive stimulus occurs an action must be taken to actively realise the reward (e.g. I hear the music of the ice cream van, a stimulus predictive of potential reward, but must take a series of actions [approach, purchase] to actually then acquire an ice cream). It is possible that when study arrays had a singleton positive face this initiated a goal-like 'approach the singleton' state in participants since such a state is the mechanism by which rewards would be realised if associated value codes were to become active. By contrast, when study arrays comprised all positive value associated faces the implementation of such a goal-like state would be unnecessary. It is possible that this may account for why no WM boost was observed for positive value faces in uniform arrays. This is in contrast to the effect of face emotion on WM (Experiment 1; see also Jackson, Wu, Linden & Raymond, 2009), where array wide but item specific boosts were

seen for angry faces. However, angry faces within a uniform array might induce a 'retreat' goal state (or alternatively an 'approach' goal state depending on the individual, e.g. see Fung, Isaacowitz, Lu & Li, 2010) since an action must be taken or at least be prepared in order to deal with angry individuals, which would account for the WM boost effect observed for uniform displays of angry faces.

An alternative interpretation of these results is that the colours present in singleton-containing study arrays act as biasing agents for a spatial saliency map such that positive value associated colours bias spatial allocation of resources to the retinotopically mapped locations they occupy in the memory representation. Although the lack of a significant effect of gaze direction during study array in Experiment 6 speaks against a differential deployment of overt attention, it is possible that the saliency of locations instead serves to indicate that items at these locations should be the subject of a boost in processing at a later stage in the memory process, perhaps during consolidation. In this view, the WM boost effect is signalled by a spatial cue exogenous to the stimulus representation to be boosted.

A third alternative is that a combination of task utility and outcome association are required for stimuli to be boosted. In this view it is not merely the singleton status of a value associated face that renders it eligible for a WM boost; rather, it is the utility that singleton status has in terms of task demands that interacts with its value association and so produces a boost. Recall that in Experiment 6 singleton faces were more likely to be probed than any given non-singleton face. Because of this, face colour carried information that could aid task performance (by facilitating devotion of more processing resource to singletons). By contrast, in Experiment 7 the colour of faces was entirely orthogonal to task

requirements (i.e., it had no utility in terms of task demands) and so seems to be fully discounted despite any association with prior reward or loss.

Effect of increasing positive valence

In Experiment 8, we observed that the propensity of a face to predict a positive (win) event in a previous task impacted the way in which such a face would be memorised in a WM task, with faces that were more predictive of gain being better memorised, while the propensity of faces to predict a negative (loss) event had no such effect. Interestingly, we also observed that a secondary concurrent WM task (change detection between two arrays of coloured squares) seemed not to interact with ability to do the primary face WM task, suggesting that qualitatively different types of WM might be utilised in the two tasks.

Conclusions

The outcome predictiveness of face stimuli was observed to have qualitatively different effects on how they were treated in WM depending on whether the outcome valence predicted was positive or negative. While positive association benefits memorisation under certain circumstances (Experiment 6) and has increasing benefits to memorisation for stronger positively predictive associations (Experiment 8), negative associations do not give rise to such effects on WM performance. I discuss the implications of this modulation of WM by positive value stimuli in relation to the existing literature in Chapter 9 (general discussion).

Whereas in the first two experiments reported in this chapter value was associated with face stimuli in a WM task by tinting them in different colours that

were themselves associated with positive or negative outcomes, in Experiment 8 values were directly associated with individual face stimuli. In both instances, stimulus values were established in learning phases prior to the memory test and were orthogonal to the memory task demands. In both instances we have evidence for the effect of positive value association on WM representation when value associations are encoded in LTM. An alternative possibility is to signal the value association of stimuli immediately before they are presented in a WM trial such that their value associations must be retained in WM rather than in LTM. In the next chapter I report an experiment which took this approach.

References – Chapter 7

- Anderson, B.A., Laurent, P.A., & Yantis, S. (2011). Learned value magnifies salience-based attentional capture. *PLoS ONE* 6(11): e27926. doi:10.1371/journal.pone.0027926
- Anderson, B.A., Laurent, P.A., & Yantis, S. (2011). Value-driven attentional capture. *Proceedings of the National Academy of Sciences USA*, 108, 10367-10371.
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology*, 1(1), 42-45.
- Ekman, P., & Friesen, W. (1976). Pictures of facial affect. Palo Alto, CA: Consulting Psychological Press.
- Fung, H. H., Isaacowitz, D. M., Lu, A. Y. & Li, T. (2010). Interdependent Self-Constraint Moderates the Age-Related Negativity Reduction Effect in Memory and Visual Attention. *Psychology and Aging*, 25(2), 321-329.
- Gehring, W. J. & Willoughby, A. R. (2002). The medial frontal cortex and the rapid processing of monetary gains and losses. *Science* 295, 2279-2282.
- Granholm, E., & Steinhauer, S.R. (2004). Pupillometric measures of cognitive and emotional processes. *International Journal of Psychophysiology*, 52, 1-6.
- Hacker, M. J. & Ratcliff, R. (1979). A revised table of d' for M-alternative forced choice. *Perception & Psychophysics*, 26(2), 168-170.
- Mackintosh, N. J. (1975). A theory of attention: variations in the associability of stimuli with reinforcement. *Psychological Review*, 82, 276-298.
- Hickey, C. Chelazzi, L. & Theeuwes, J. (2010a). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*, 30, 11096-11103.
- Hickey, C. Chelazzi, L. & Theeuwes, J. (2010b). Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLoS One*, 5(11), 1-5.
- Hochman, G., & Yechiam, E. (2011). Loss aversion in the eye and in the heart: The Autonomic Nervous System's responses to losses. *Journal of Behavioral Decision Making*, 24, 140-156.
- Jackson, M. C., Wu, C., Linden, D. E. J., & Raymond, J. E. (2009). Enhanced Visual Short-Term Memory for Angry Faces. *Journal of Experimental Psychology: Human Perception and Performance*, 35(2), 363-374.

- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263-291.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, *390*(6657), 279-81. doi:10.1038/36846
- Lundqvist, D., Flykt, A. & Öhman, A. (1998). *The Karolinska Directed Emotional Faces*. Psychology section, Department of Clinical Neuroscience, Karolinska Institute, Stockholm, Sweden.
- Pessiglione, M., Seymour B., Flandin G., Dolan R. & Frith C. (2006). Dopamine-dependent prediction errors underpin reward-seeking behaviour in humans. *Nature*, *442*, 1042-1045.
- Raymond, J. E. & O'Brien, J. L. (2009). Selective visual attention and motivation: the consequences of value learning in an attentional blink task. *Psychological Science*, *20*, 981-988.
- Rozin, P., & Royzman, E.B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, *5*, 269-320.
- Sessa, P., Luria, R., Gotler, A., Jolicoeur, P., & Dell'Acqua, R. (2011). Interhemispheric ERP asymmetries over inferior parietal cortex reveal differential visual working memory maintenance for fearful versus neutral facial identities. *Psychophysiology*, *48*, 187-197. DOI: 10.1111/j.1469-8986.2010.01046.x
- Theeuwes, J. (1991). Cross-dimensional perceptual selectivity. *Perception & Psychophysics* *50*, 184-93.
- Treisman, A. & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, *12*(1), 97-136.
- Wolfe, J. M. (2007). Guided Search 4.0: Current Progress with a model of visual search. In W. Gray (Ed.), *Integrated Models of Cognitive Systems* (pp. 99-119). New York: Oxford.
- Yechiam, E., & Hochman, G. (2013). Losses as modulators of attention: review and analysis of the unique effects of losses over gains. *Psychological bulletin*, *139*(2), 497–518. doi:10.1037/a0029383

PART 4 – MOTIVATION AND WORKING MEMORY

CHAPTER 8.**Motivation and Working Memory**

Can incentives improve visual WM?

In the preceding chapters we have observed that under certain conditions visual WM capacity, or the resolution with which item representations are stored, can be boosted, and that such boosts are flexible enough to occur in an item specific fashion and appear to derive in part from modulation of consolidation processes occurring early after stimulus offset. Given that it is possible for WM boosts to occur driven by bottom up motivational salience (angry faces), and ‘imbued motivational salience’ (positive EV faces), perhaps it is also possible to exert conscious top-down control over whatever mechanism serves to boost WM performance. In situations of high motivation, perhaps people can decide to boost WM in order to better perform a task. Such a process has a clear evolutionary advantage; we have discussed the possibility that the item specific nature of ‘singleton boosting’ serves to make encoding optimally neurally efficient, in a similar way it seems reasonable that the activation of a WM boosting mechanism might only occur when there is a reward to be gained or a punishment to be avoided. Such situations need not always require the presence of intrinsically motivationally salient items. For example, if I pick up a magazine that contains an ‘intelligence test’ with a WM component and in a moment of leisure complete such a test, my motivation to perform well may be less than if I am given the same test as an evaluation of ability during a job interview. In the latter case, a specific incentive exists to boost my performance, over and above the more ‘natural’ level of the first example. However, there is as yet no evidence that deliberate boosts in visual WM can be effected at will.⁵⁰

Monetary rewards used as incentives in an attention deployment task have been shown to enhance perceptual sensitivity (d') by facilitating both

⁵⁰ Although see Matsukura & Hollingworth (2011) for evidence that extensive training on a WM task can facilitate better performance.

orienting and reorienting of exogenous visual attention (Engleman and Pessoa, 2007), a process described by Engleman et al. as a 'sharpening' of exogenous attention. Herein we investigate whether similar 'sharpening' of visual WM might occur under conditions of motivation. There are various reasons for supposing this might be so. First, the interlinked nature of, and shared neural circuitry between the visual attention and visual WM systems (Awh & Jonides, 2001; Pessoa & Ungerleider, 2004, though see Fougne, 2008) suggests that they may be affected in similar ways by certain manipulations. Second, clear advantages exist to the boosting of WM under conditions of motivation since doing so allows the rapid processing of items of utility and so grants a competitive advantage in obtaining such items over others who process them less quickly. Third, if visual WM rehearsal and/or retrieval processes are facilitated by an internal attention-like system (Awh et al., 2001) (that is, if the direction of attention internally to the representation of a visual stimulus can aid in its retention in WM) then a motivationally induced sharpening of this internally directed attention may lead to better retention of visual information in WM. This notion is supported by studies that have found similar ERP wave components (N2pc potentials with similar time courses and scalp distributions) for searching perceptual and VSTM space, suggesting that the neural mechanisms for these two processes overlap. (Nobre, Griffin & Rao, 2008; Kuo, Rao, Lespien & Nobre, 2009).

Here I ask whether knowledge of potential rewards or the short term motivational state this knowledge may engender can improve performance on a simple WM task or whether mechanisms mediating WM performance are constrained by "hard wired" capacities. I gave participants a simple WM test under two conditions; high reward possibility and low reward possibility. The WM

task used was a coloured squares change detection task that has been widely reported elsewhere (e.g. Luck & Vogel, 1997). The basic change detection task was modified by adding a screen signalling to participants prior to each trial whether a high or low reward (an adjustment to a running total of monetary winnings) was obtainable. If the possibility of obtaining (high) reward facilitates a processing boost then better WM performance should be seen on those trials where a high reward is obtainable. Alternatively, if *experiencing* a reward results in such a short term boost (if, for example short term arousal can facilitate memorisation in a similar way as has been suggested in the domain of LTM for 'flashbulb' memories; e.g. Brown & Kulik, 1977) then we would expect to see better performance on trials immediately following a high win. Alternatively, if motivation to obtain a reward does not alter WM ability (if WM capacity is hard-wired) then no performance change would be observed between such conditions.

Method

Participants

Participants were recruited through Bangor University and received course credit in exchange for participation. All reported normal or corrected to normal colour vision. 26 adults (16 females, mean age = 24.4 years s.d. = 3.2 years) participated.

Apparatus

Stimuli were presented on a Mitsubishi Diamond Plus 20 inch colour monitor (resolution = 1280 x 1024 pixels, refresh rate = 100Hz), using a Dell

Optiplex GX400 computer running Microsoft Windows XP. Viewing distance was 60 cm. E-Prime software (Version 1.0; Schneider, Eshman, & Zuccolotto, 2002) generated stimuli and recorded responses acquired via a keyboard.

Stimuli

Coloured square stimuli, rendered as a bitmaps with a bit depth of 24 and side length of 29 pixels, were presented in arrays on a grey background. Within each array squares were presented at random locations within a 350 by 350 pixel area in the centre of the screen with no square displayed less than 130 pixels from any other square. Square colours were: Black, blue, brown, green, orange, pink, purple, red, turquoise, and white. RGB coordinates for the squares and background are given in Appendix B.

Procedure and Design

Each trial began with a 1000 ms central fixation cross followed by a reward cue (the text 'HI' or 'LO') signalling whether a large reward (ten pence) or a small reward (one penny) could be obtained with a correct change detection response. Another 1000 ms fixation cross was followed by a 100 ms presentation array of six coloured squares. After a blank screen retention interval lasting 900 ms, a 1000 ms test array was shown. On half of trials ('change trials') one of the squares changed colour between the memory and test arrays. On remaining trials ('no-change trials') the memory and test arrays were identical. Next, a screen with the text 'Same or Different' was presented until response. Participants responded using the 's' key to denote same and the 'k' key to denote different. After response, a screen querying surety of the judgement made appeared. This comprised the text 'How sure are you of your answer?' presented centrally at the

top of screen and the response options 'Not very sure', 'Somewhat sure', and 'Very sure' presented below from left to right, each with their associated keyboard response key (1, 2, and 3 respectively) under the text. This screen was present until response. Finally, a screen appeared giving feedback as to whether the change detection response made had been correct or not, the amount of money won or lost on the trial, and the running total of winnings (see Figure 8.1 below). Correct responses resulted in wins determined by the trial type (HI or LO as above). Incorrect responses were punished by loss of one penny (see Table 8.1 below).

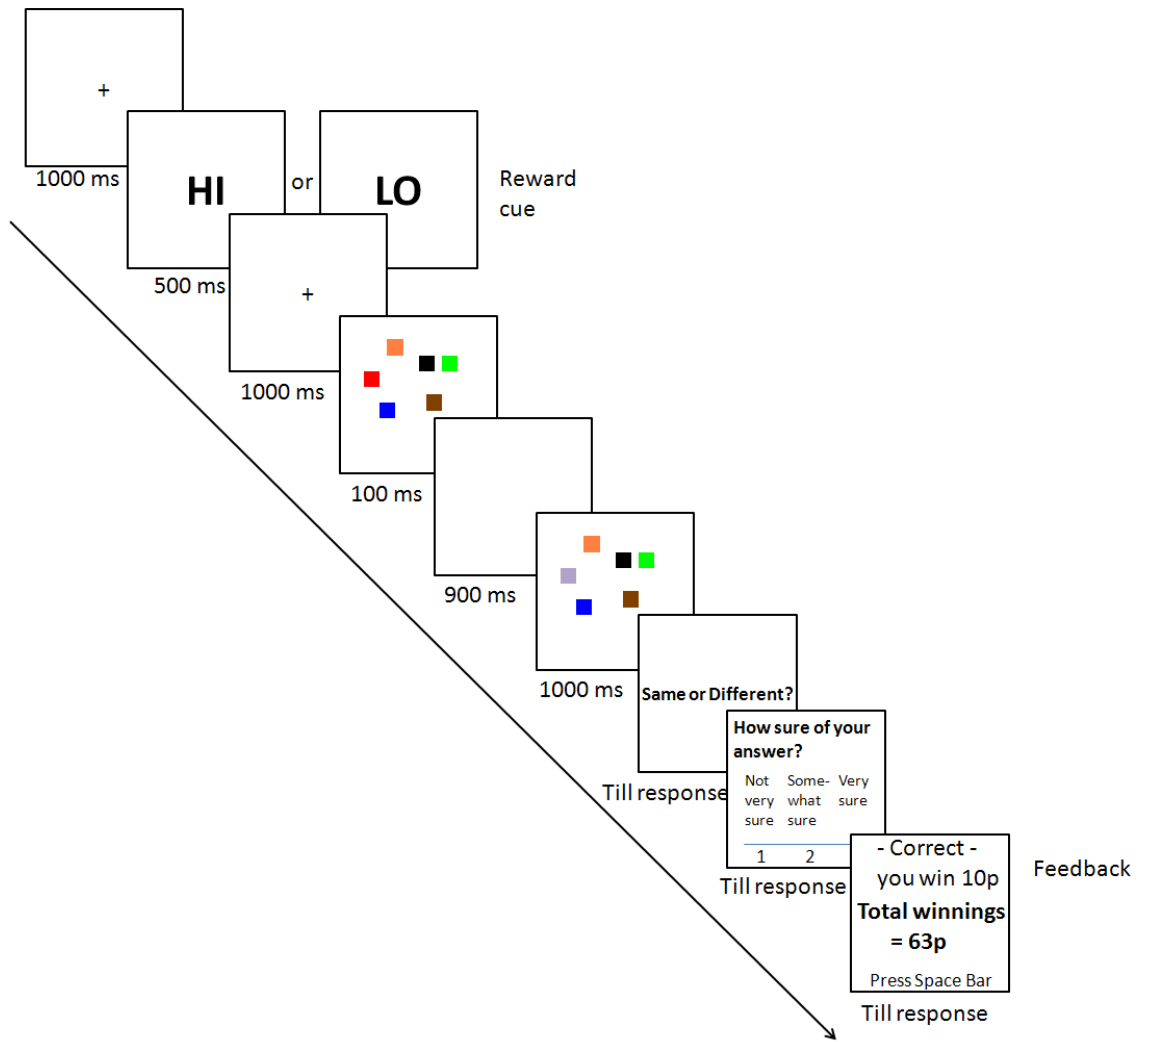


Figure 8.1. Trial sequence. In this example a high win trial is shown, and the participant has made a correct response (reporting 'change') and has therefore won 10p.

Reward Cue	Response	Feedback
High	Correct	Correct – you win 10p
High	Incorrect	Wrong – you lose 1p
Low	Correct	Correct – you win 1p
Low	Incorrect	Wrong – you lose 1p

Table 8.1. Reward and punishment contingencies.

There were 120 trial in total (60 'HI' cued; 60 'LO' cued) presented in a single block. A practice block of eight trials preceded the experimental block after which initial winnings were reset to zero.

Data analysis

Data from one participant indicated non-engagement with the task (performance at chance across all conditions), so these data were excluded from further analysis. Data from trials with response times of less than 200 ms (anticipation errors) or greater than 6000 ms were excluded from the analysis, accounting for 3.1% of remaining data.

My initial plan was to use participants' ratings of their confidence of response to plot a receiver operating characteristics (ROC) curve and use the area under this curve as an index of sensitivity to change (see Chapter 4). However, participants rarely (on average 12% of the time) chose to use the 'Not very sure' rating, resulting in insufficient data to plot the corresponding data point

for many participants.⁵¹ Rather than excluding this 12% of data (and basing my SDT measure on an area under a graph plotted from data points for the remaining surety ratings), I instead chose to collapse all confidence ratings together and quantify change detection performance by d' (see Chapter 4).⁵²

My two theoretical factors of interest were Reward Availability Signal (HI, LO), and Previous Trial Outcome (a concatenation of potential reward available on the previous trial and the actual outcome for it; HI-win, HI-loss, LO-win, LO-loss). However, because participants achieved on average 81% correct across all trial types there were insufficient trials for robust analysis of the HI-loss and LO-loss levels of the Previous Trial Outcome factor. Accordingly, I first ran a within-

⁵¹ It might seem tempting to try to 'force' participants to use the full scale of confidence responses available (e.g. by asking them to try to use each surety response approximately an equal number of times and showing them a running total of the number of times they have made each surety response). However, this is not to be advised since the purpose of the scale is to try to accurately reflect the confidence with which the judgement is being made. However, a longer scale might allow for better discrimination between points by participants. It is possible that the end points of the scale I used are interpreted as representing extremes of confidence and so are not used whereas a scale with, say five options might allow the end values only to be interpreted as extreme and all three central values to be used.

⁵² This illustrates a practical issue in the choice of SDT measure deployed in the analysis of such data. While measures derived from areas under a graph may be desirable since they make fewer assumptions about the underlying distribution there is always a chance, especially when trial numbers are low, that either through response bias (unwillingness to use a part of a scale provided as seems to be the case in my experiment here), or through random chance, that one of the desired data points on the graph will not receive enough contributory scores to be plotted with robustness. In such cases, three ways of handling the data may be attempted: (1) The problem may be ignored. This has the advantage of inclusion of all data points but the disadvantage that the point plotted on the SDT graph may be inaccurate, leading to a less accurate measure of the area under the ROC graph. (2) The particular data point in question may be excluded and the area under the graph calculated from remaining points. When a very low percentage of data contribute to the calculation of a point on the graph and there are sufficient remaining data points then this may be the optimal solution since very little of the overall data is lost. (3) Data may be collapsed and a measure that makes assumptions about the underlying distribution (A' , d' , K or even percent correct) used. The advantage of this approach is that it does not discard any data. An alternative approach would be to assign data from the 'underweight' point to its nearest neighbours, e.g. assign data contributing to the 'Not very sure' point to the 'Somewhat sure' point. While this approach makes use of the full data set and stays on course for a graphical signal detection theory (SDT) measure it makes the assumption that participants treat points on the surety response scale as being equidistant from one another – i.e. that they treat the response scale as being truly scalar. This assumption is unwarranted and cognitive 'distance' between points is likely to vary across participants – a particular problem if the underweight data point we want to remove by data reassignment constitutes mainly data points from only a few participants.

subjects t-test on Reward Availability Signal (to test the prediction that Reward Availability Signal might influence WM performance using the maximum amount of data available to us for this analysis), and then ran a repeated measures ANOVA on d' scores using current trial Reward Availability Signal (HI, LO) and Prior Win Experienced (HI-win, LO-win) as within-subject factors, i.e. using only those data where the previous trial outcome was a win. Planned comparisons used within-subject t -tests. Alpha levels were set at .05.

Results

To test my first hypothesis, that the possibility of obtaining a high reward facilitates performance, I looked at how performance on trials on which a high reward was obtainable differed from performance on trials on which only a low reward was available. WM performance on trials where Reward Availability Signal was 'HI' ($M = 2.04$, $SD = .53$) was slightly better than when Reward Availability Signal was 'LO' ($M = 1.84$, $SD = .60$; see Figure 8.2). However, a within subjects t -test revealed that this difference was not significant [$t(24) = 1.513$, $p = .143$]. It seems, therefore, that magnitude of potential reward available may not alone be enough to bolster performance significantly.

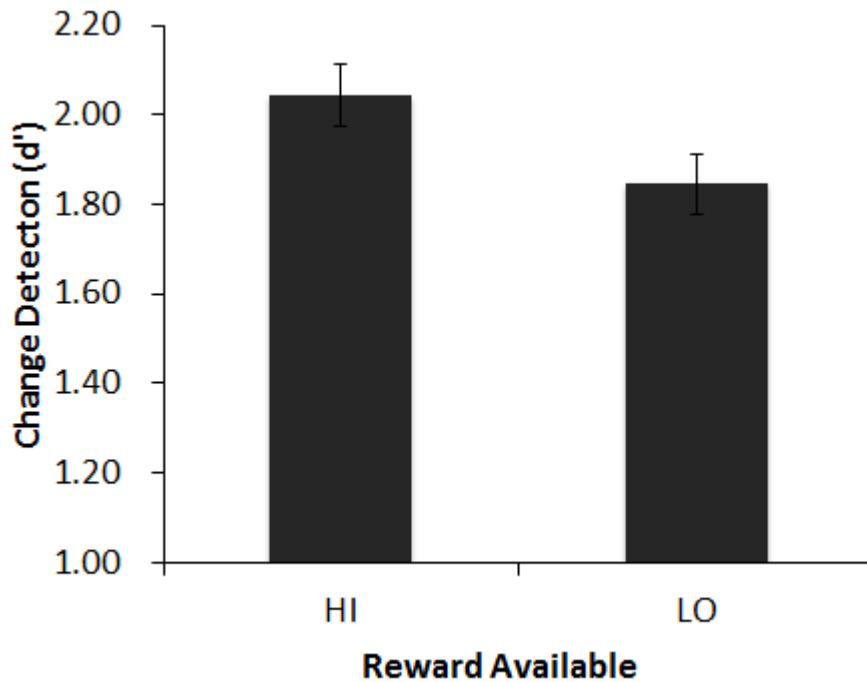


Figure 8.2. WM performance split by reward available on each trial. Error bars represent within subject normalised standard error of the mean (Cousineau, 2005; see Appendix D).

To test my second hypothesis, that experiencing a reward on a previous trial might facilitate performance, I split performance data into trials preceded by big wins, small wins, high reward available losses, and low reward available losses.⁵³ However, due to overall performance being better than I had anticipated, only 19% of trials fell into either of the preceded-by-loss categories. With too few trials available for robust analysis of these loss trial categories, I confined my analysis to trials preceded by big wins vs. those preceded by small wins.

⁵³ The first trial of the experiment for each participant was excluded from this analysis since it was not preceded by any trial event.

A different pattern of results was observed on trials preceded by big wins vs. those preceded by small wins (see Figure 8.3). When a big win preceded a trial, if that trial was of the 'HI' reward obtainable type, performance was higher ($M = 1.97$, $SD = .54$) than if it was of the 'LO' reward obtainable type ($M = 1.63$, $SD = .74$). However, if a small win had preceded a trial then performance on 'HI' trials ($M = 1.75$, $SD = .68$) was weaker than on 'LO' trials ($M = 1.88$, $SD = .63$; see Figure 8.3). A within-subjects ANOVA with Reward Availability Signal (HI, LO) and Prior Win Experienced (HI-win, LO-win) as within-subject factors revealed a non-significant effect of Prior Win Experienced [$F(1, 24) = .015$, $p = .904$, $\eta_p^2 = .001$], a non-significant effect of Reward Available [$F(1, 24) = .671$, $p = .421$, $\eta_p^2 = .027$], but a significant interaction between these factors [$F(1, 24) = 7.094$, $p = .014$, $\eta_p^2 = .228$]. Planned comparisons revealed that the interaction was of the following form: When a trial was preceded by a big win, performance was significantly better if the reward available cue on the trial was 'HI' vs. 'LO' [$t(24) = 2.243$, $p = .034$], but for trials preceded by small wins the reward availability cue had a non-significant effect [$t(24) = -.816$, $p = .423$]. Thus, it seems that experiencing a big win empowers the following reward availability cue with the properties necessary for it to exert an influence on performance.




Figure 8.3. Effect of prior win received (big win vs. small win) on performance for 'HI' and 'LO' potential reward trials. Error bars represent normalised within subject standard error (Cousineau, 2005).

Within any WM task, there are always going to be participants who are performing at their very best, being fully engaged with the task, motivated and alert, but also other participants who for one reason or another are performing sub-par relative to what they are truly capable of. It is reasonable to assume that this variance across participants might parallel to some extent a similar pattern of performance on tasks outside the experimental setting (i.e. in participants' daily lives). Thus, inclusion of participants performing at a suboptimal level along with those performing optimally may result in better ecological validity of experimental findings. However, in the context of a task designed to boost performance levels this distinction becomes more important since a boost in performance under

incentivised conditions may be the result either of facilitation of some process across all participants (e.g. bringing online otherwise idle cognitive resources in a similar way to that described in Chapter 5) or may be the result of generating temporary task re-engagement among those participants who otherwise perform at a sub-par level relative to their true capability. The difference between these two explanations of the boost we have observed above is of crucial theoretical importance: One account suggests that WM capacity is not a fixed resource but can be boosted under the right conditions, the other suggests that incentivisation should be able to facilitate the performance only of those participants performing at a level lower than their 'true' (maximal) capability. To investigate this, I performed a median split on overall WM performance (performance across all conditions) to divide participants into a low WM capacity group and a high WM capacity group.⁵⁴ If the WM 'boost' we observed is due to temporary re-engagement with the task by lower performing participants then we would expect to see that the effect of the boost holds for these participants only, whereas if it reflects a boost of WM to some 'special', normally unobtainable level then we would expect to see boosted performance (on trials where a big win precedes a HI reward obtainable signal vs. trials where only a small reward precedes a HI reward obtainable signal) across both the high and low capacity groups.⁵⁵

⁵⁴ It is worth noting that this nomenclature is found throughout the literature, but may be suspect for the reasons discussed in the text: While the division of participants into high and low WM 'capacity' groups is commonly achieved by measuring performance on simple WM tasks (e.g. Mall & Morey, 2013) such measurements are affected by factors such as task engagement, concentration and so on, which vary across participants. Thus, it may be more appropriate when making such a split of participants to make the distinction between high and low performing participants rather than high and low capacity participants.

⁵⁵ Although scores will necessarily be lower in the low WM group by virtue of its definition, we are interested in seeing if there is any difference in the pattern of scores within the low and high WM groups. While an alternative way of defining high and low WM capacity groups would be to conduct a test of capacity in a separate task, this alternative has the drawback that it cannot be performed 'online', i.e. at the exact time of the task of interest. This is particularly important here as the state of 'being in an

Figure 8.4 shows performance on trials preceded by high wins for the high and low WM groups. Participants in the low WM group achieved better performance on HI cued trials ($M = 1.79$, $SD = .60$) than LO cued trials ($M = 1.19$, $SD = .70$), whereas participants in the high WM group achieved similar performance on HI cued ($M = 2.09$, $SD = .41$) and LO cued ($M = 2.06$, $SD = .55$) trials.




Figure 8.4. Effect of prior reward available cue (HI vs. LO) on performance. A median split divides participants into those achieving high overall performance and those with low overall performance. Error bars represent standard error of the mean.⁵⁶

experiment in which rewards are obtainable' may in and of itself affect general motivational state, rendering measures of WM capacity performed under different conditions less valid as the basis upon which to perform a split into high and low performing groups.

⁵⁶ As we wish to make comparisons both within groups of participants (i.e. across conditions in the low WM group and across conditions in the high WM group) as well as between the high and low WM groups, use of normalised standard error bars as in the preceding two figures is not appropriate here since they would hide the between groups component of the overall variance (see appendix D).

From the analysis above we know that that the main effect of Reward Available for trials preceded by big wins is significant, and (by definition) I expected the high and low WM groups to show a significant difference in performance, but of interest is whether there is an interaction between these factors. An ANOVA with Reward Available as a within subjects factor and WM Group as a between subjects factor revealed that the effect of group was indeed significant [$F(1, 22) = 10.035, p = .004, \eta_p^2 = .313$] and that the interaction between WM group and Reward Available was marginally significant [$F(1, 22) = 3.726, p = .067, \eta_p^2 = .145$]. This is suggestive, though not conclusive evidence that high and low performing participants are affected by incentivisation after experiencing a big win in different ways, perhaps due to a refocusing of resources on the task in the case of lower performing participants. As I was specifically interested in this point I followed up the marginal interaction effect with paired sample t-tests within the high and low WM groups separately. This revealed a significant performance difference between HI and LO incentive trials following big wins in the low WM group [$t(11) = 2.397, p = .035$] but a non-significant difference in the high WM group [$t(11) = .181, p = .859$].

Discussion

I set out to investigate whether incentivisation could boost performance on a simple WM task. I found that under specific conditions, namely when a large win on one trial was followed by a cue signalling availability of another potential large win, task performance was facilitated. This effect was driven by participants whose overall task performance was weaker, while participants whose

performance was better across that task tended not to show this pattern of results. One explanation for this result can immediately be ruled out: None of our participants performed at ceiling on any condition. Therefore a ceiling effect cannot account for the similar performances across conditions among the participants in the high performance group. Rather, this pattern of results suggests that this group of participants had reached their 'natural' level of performance (see below) under these conditions.

It is possible that the low capacity group may have 'boosted' their WM in the presence of potential high reward after a big win, perhaps by drawing upon otherwise undeployed resources in a similar way to that described in Chapter 5. Alternatively, they may have 'withdrawn' resources or 'partially switched off' in the conditions where performance was weaker. However, rather than treating these two explanations as dichotomous, it is perhaps more appropriate to draw a distinction between the 'natural WM capacity' (the capacity which a person generally calls upon in day to day tasks) and 'true WM capacity' (the maximum level of capacity which a person can achieve under fully optimised circumstances, i.e. when boosted by provision of incentivising stimuli, and in the absence of distraction⁵⁷) that each participant may have. If we conceptualise WM performance in this way, then the interaction of reward experienced and reward available for lower performing participants serves to shift their performance towards their 'true' WM capacity (or create conditions under which natural capacity is nearer to true capacity). It is interesting to note that levels of performance when a big win was followed by a high reward available signal were broadly similar between the low and high WM capacity groups. This suggests that the low

⁵⁷ Distraction in this sense may include both presence of formally presented distractor stimuli but also any other distraction from the task, e.g. thoughts extraneous to the task.

performance group was capable of performing to this level but needed the very specific conditions of a big win followed by the potential of another big win to reach this level of performance.

In Chapters 5 and 6 I have argued for the existence of a pool of resources, perhaps equivalent to space in Baddeley's (Baddeley, 2000, 2001; Repovš and Baddeley, 2006) episodic buffer (see Chapter 9, General Discussion) that are only brought online in certain circumstances such as when attracted by a motivationally salient item. Here, I suggest a refinement to this model, specifically that individual differences may play a role in determining what type of stimulus or motivational state is required to bring online these reserves of resources.

How might the WM 'boost' be effected?

It has been shown that individuals are capable of selectively attending to one part of a WM representation when asked to do so and that if the stimulus feature internally attended to has emotional valence (e.g., the teeth rather than the eye in a picture of a wolf) then this will induce an emotional response as indexed by both self-report and the late positive potential (LPP; Schupp et al., 2000) EEG component (Thiruchselvam, Hajcak & Gross, 2012). This shows that the contents of WM are subject to top-down control, and suggests that it is possible to actively allocate resource to salient parts of an internal representation. This comes as little surprise given the effects we observed in Chapters 5 and 6, where salient stimuli appeared to receive a boost in WM. Here, however, this finding offers a possible explanation for the array-wide boost observed for participants in the low WM group. Specifically, if it is possible to deploy an internal attention-like resource (see below) to a specific part of a visual WM representation at the expense of

other parts of the representation then it may also be possible to deploy such a resource away from extraneous non-task relevant WM representations or other cognitive processes that make use of the WM system⁵⁸ in order to boost the representation of the whole array on trials where motivation is high.

An alternative interpretation is that what is enhanced in our experiment is the process of comparing the internal representation of the memory array with the array presented at test. At first glance this seems an unlikely explanation since the process of comparing the two arrays, which differ only in the colour of one element, seems trivial if one imagines the two arrays overlaid one on top of the other. However, if during the comparison process it is necessary to switch attention between the internal representation and the external stimulus (test array) then this switching of attention might be subject to similar limitations as have been observed in the change blindness (Rensink, O'Regan & Clark, 1997) and attentional blink (AB; Broadbent & Broadbent, 1987; Raymond, Shapiro & Arnell, 1992) paradigms. In particular, the failure to spot changes in the change blindness paradigm might be paralleled by a similar inability to spot changes between the internal array representation and the external test array. If this is a limiting factor in the performance on this standard WM task then it is possible that the attentional switching costs during the comparison process are what are ameliorated by motivational incentives. This notion is supported by the observations that attentional orienting and reorienting are facilitated under conditions of high motivation (Engleman et al., 2007) and that motivationally salient items are able to escape the AB (Raymond & O'Brien, 2009). Additionally,

⁵⁸ For example, it has been observed that participants with low WM have more difficulty than high WM participants in maintaining focus on a cognitive task as they become distracted by secondary issues such as monitoring the progress of time (Woehrle & Magliano, 2012)

visual attention appears to be a crucial component in the array comparison process as after the initial change is detected, visual attention is then directed to the area of the change in order to confirm the change/no-change decision with further processing of the stimulus (Pessoa & Ungerleider, 2004a).

Conclusion

The purpose of this experiment was to investigate whether performance on a simple WM task could be facilitated either by the experience of reward, or the presence of a signal indicative of high reward availability. While neither reward magnitude experienced nor magnitude of reward available was found to directly affect performance, the interaction of these factors did so for a group of lower performing participants, such that experiencing a high reward outcome facilitated performance only when another high potential reward was immediately available. This facilitation of performance may be driven by a temporary automatic re-engagement with the memory task among those participants whose natural level of engagement is otherwise lower than optimal.

References – Chapter 8

- Awh, E., & Jonides, J. (2001). Overlapping mechanisms of attention and spatial working memory. *Trends in Cognitive Science*, 5, 119–126.
- Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4, 417–423.
- Baddeley, A. D. (2001). The magic number and the episodic buffer. *Behavioral and Brain Sciences*, 24, 117–118.
- Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception and Psychophysics*, 42, 105-113.
- Brown, R. & Kulik, J. (1977). Flashbulb Memories. *Cognition*, 5(1), 73-99.
- Engleman, J. B. & Pessoa, L. (2007). Motivation Sharpens Exogenous Spatial Attention. *Emotion*, 7(3). 668-674.
- Fougnie, D. (2008). The Relationship between Attention and Working Memory. In N. B. Johansen (Ed.), *New Research on Short-Term Memory*, 1-45, Nova Science Publishers, Inc. ISBN: 978-1-60456-548-5
- Kuo, B-C., Rao, A., Lepsien, J. & Nobre, A. C. (2009). Searching for Targets within the Spatial Layout of Visual Short-Term Memory. *The Journal of Neuroscience* 29(25). 8032-8038.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-281.
- Mall, J.T. & Morey, C. C. (2013). High Working Memory Capacity Predicts Less Retrieval Induced Forgetting. *PLoS ONE* 8(1): e52806. doi:10.1371/journal.pone.0052806
- Matsukura, M. & Hollingworth, A. (2011). Does visual short-term memory have a high-capacity stage? *Psychonomic Bulletin & Review*, 18, 1098-1104. DOI 10.3758/s13423-011-0153-2
- Nobre, A. C., Griffin, I. C. & Rao, A. (2008). Spatial attention can bias search in visual short-term memory. *Frontiers in Human Neuroscience*, 1, 1-9.
- Pessoa, L. & Ungerleider, L.G. (2004a). Neural Correlates of Change Detection and Change Blindness in a Working Memory Task. *Cerebral Cortex*, 14, 511-520.

- Pessoa, L. & Ungerleider, L.G. (2004b). Top-down mechanisms for working memory and attentional processes. In M. S. Gazzaniga (Ed.), *The New Cognitive Neurosciences, 3rd Edition*, 919-930, MIT Press.
- Raymond J. E. & O'Brien, J. L. (2009). Selective visual attention and motivation: The consequences of value learning in an attentional blink task. *Psychological Science, 20*(8), 981-988.
- Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink?. *Journal of Experimental Psychology: Human Perception & Performance, 18*(3), 849-860.
- Rensink, R. A., O'Regan, J. K. & Clark, J. J. (1997). To see or not to see: the need for attention to perceive changes in scenes. *Psychological Science, 8*, 368–373.
- Repovš, G. & Baddeley, A. (2006). The multi-component model of working memory: explorations in experimental cognitive psychology, *Neuroscience 139* (2006), 5–21.
- Schneider, W. X., Eshman, A., & Zuccolotto, A. (2002). *Eprime user's guide*. Pittsburgh, PA: Psychology Software Tools, Inc.
- Schupp, H. T., Cuthbert, B. N., Bradley, M. M., Cacioppo, J. T., Ito, T., & Lang, P. J. (2000). Affective picture processing: The late positive potential is modulated by motivational relevance. *Psychophysiology, 37*, 257-261.
- Thiruchselvam, R., Hajcak, G. & Gross, J. J. (2012). Looking Inward: Shifting Attention Within Working Memory Representations Alters Emotional Responses. *Psychological Science, 23*(12), 1461-1466. DOI: 10.1177/0956797612449838
- Woehrle, J. L. & Magliano, J. P. (2012). Time flies faster if a person has a high working-memory capacity. *Acta Psychologica, 139*(2), 314-319

PART 5 – GENERAL DISCUSSION

CHAPTER 9.
General Discussion

"Selection is the very keel on which our mental ship is built. And in the case of memory its utility is obvious. If we remembered everything, we should on most occasions be as ill off as if we remembered nothing."

- William James, Principles of Psychology, 1890

The human short term visual memory system represents the world subject to capacity and resolution constraints and as such represents only a fraction of available information at any time. Here, in a series of nine experiments I have addressed the way in which two general classes of stimuli, emotional faces and value laden faces, may bias the deployment of the limited capacity resources that underlie short term memory representation of visual stimuli.

Experiments 1 through 5 addressed resource deployment with regard to emotional faces while Experiments 6 through 8 addressed this issue for value laden (positive and negative event predictive) faces. Experiment 9 addressed a more general question of whether motivational state may facilitate resource deployment. A common theme to all experiments reported herein is the investigation of whether it is possible to 'boost' WM performance or whether WM capacity (as indexed by performance) is fixed by hardwired constraints.

In Experiment 1, I examined the effect of emotionality of a face on WM for that face when it appeared as a singleton in an array of 4 faces. Not only were singleton faces bearing angry expressions remembered better than those bearing happy expressions, but this WM boost was not accompanied by bettering or worsening of performance for accompanying neutrally expressive faces. This suggests a mechanism whereby the presence of an angry face brings online otherwise idle resources to bolster WM performance in a way that is 'item specific',

i.e. the additional resources brought online serve to boost the WM representation of the stimulus that occasioned their deployment but do not contribute to the representation of other co-present stimuli.

In Experiment 2, I replicated the effect observed in Experiment 1, and also ruled out shifts in overt attention as a causal factor for the angry face WM boost by taking gaze direction during study array presentation as an index of overt attention. Additionally, I observed a weaker boost to happy faces (over and above neutral co-present faces).

In Experiments 3, 4 and 5, using a backward masking paradigm I observed the angry boost effect to be driven by a process that occurred very early after stimulus offset, most probably during the consolidation period. This process was disrupted not only by presentation of a mask during it, but also by the mere expectation of such a mask. This suggests that elaborated consolidation of angry faces requires a short period of time unconstrained by expectation of another stimulus.

In Experiment 6, I demonstrated an item-specific boost to singleton face stimuli that had previously been paired with positive outcomes, but no such singleton boost to faces previously paired with negative outcomes. Gaze tracking was used to rule out shifts in overt attention as a causal factor. Experiment 7 revealed that the singleton positive face WM boost effect observed in Experiment 6 does not generalise to situations where all faces in an array hold the same positive outcome predictive status, perhaps due to failure to activate an 'approach' mechanism by arrays of all reward predictive faces.

In Experiment 8, I reported further evidence for the effect of positive learned predictiveness on WM memorability. Predictiveness of occurrence of a positive (win) event was linearly related to WM performance with higher predictiveness of positive event resulting in better WM performance. No such relationship was observed for loss predictive face stimuli. Additionally, the lack of any interaction between a primary face WM task and a secondary coloured squares WM task suggests that different types of stimulus might make use of different types of WM store.

In Experiment 9, I examined the effect of motivational state on memory for arrays of coloured squares. Provision of a high incentive to perform well on a trial was observed to boost WM performance, but only for participants whose overall performance level was generally lower and only when preceded by a large win event on the preceding trial.

In this chapter, I reflect upon how each of these findings tie into an overall theme of boosting of WM performance as well as how each relates to the current literature and what modification they may suggest for existing cognitive models of WM.

What types of stimuli cause WM performance to be boosted?

Herein we have observed that angry faces reliably receive boosts in WM performance, in a way that is item-specific (i.e. not coupled with a spilling over of performance boost to co-presented items nor accompanied by a deterioration in performance for such co-presented items; Experiment 1), is not explained by

attention capture (Experiment 2) and depends on a brief period of consolidation time after stimulus offset unconstrained by the expectation of an immediate mask onset (Experiments 3 and 4).

Additionally, it appears that positively- but not negatively predictive face stimuli receive a WM boost, but only under certain conditions. Specifically, it appears that it is necessary for a positively value laden face stimulus to appear either as a singleton among other non-positively associated faces (Experiments 6 and 7) or as a singly presented item for memorisation (Experiment 8) where the level of the WM boost experienced depends on the level to which such a stimulus is predictive of reward.

Are boosts universally experienced?

That only our participants in the low WM performance group Experiment 9 showed boosted performance in the presence of high reward cues suggests that it is not always possible to delve into an 'idle resource pool', rather it is only in situations where a certain amount of 'memorisation capacity' remains untapped that the presence of potential reward results in bringing online additional resources. More generally, it is possible that boosts in performance are not universally experienced. Given the wide ranging evidence for differential attentional responses to threat-based items among different groups of participants (older vs. younger adults: Isaacowitz, Toner & Neupert, 2009; participants from different cultural backgrounds: Fung, Isaacowitz, Lu & Li, 2010; high vs. low anxiety groups: Fox, Russo, & Dutton, 2002) as well as evidence of differential deployment of visual resource to reward predictive items by different personality types (Hickey, Chelazzi & Theeuwes, 2010), and the susceptibility of participants

with low WM capacities to encoding distractor items (Vogel, McCollough & Machizawa, 2005), it is possible that boosts to WM performance occur for certain types of people and/or only under very specific conditions. Further research must be launched to evaluate this issue.

The finding that simple coloured square stimuli can be the subject of a WM 'boost', albeit under very specific conditions, is important as it illustrates that there need not be something 'special' about the stimulus to be boosted, rather a general boosting mechanism might be activated by each of the various conditions under which we have observed performance boosts herein. This construct is better termed a boosting rather than prioritisation mechanism since prioritisation implies movement of processing resource away from one area to another, a finding unsupported by the 'item-specific' boost we have observed here. Furthermore, 'boosting' implies that the item is at least partially encoded prior to the activation of the mechanism, a contention supported by our observation that the boosting effect seems to operate orthogonally to selective attention.

Can boosts to WM be explained by enhanced attentional processes?

Given the strong links between WM and attention, and the role of selective attention in highlighting visual items for elaborated processing, could it be that attentional processes are responsible for WM boosts? This is a particularly tricky question, since its answer depends on suitably defining both attention and WM. In particular, differing views exist as to the ways in which attention interacts with visual WM with some (e.g. Bays & Husain, 2008) considering attention to be a gatekeeper, responsible for selecting items for entry into WM and others (e.g.

Courtney, 2004) regarding attention as an emergent property of WM, arising directly from the processing of information in WM. These two views should not be regarded as mutually exclusive; rather, they may reflect different aspects of attentional processes. Thus, it is perfectly reasonable that the contents of WM serve to guide the processes that select further information (whether from the visual scene or from LTM) for coding into WM (e.g. Downing, 2000).

WM and covert attention:

It has been found that items that are the focus of covert attention (for example, those about to be fixated) develop a stronger working memory trace than other items in a display (Bays et al., 2008). More generally, shifts in covert attention reliably precede near immediate shifts in gaze to attended locations or objects (Findlay & Gilchrist, 2003; Hoffman, 1998; Reichle, Pollatsek, Fisher, & Rayner, 1998), a phenomenon that makes it possible to take shifts in gaze as a reasonable index of shifts in attention, as I have done in analysis of Experiments 2 and 6. Bays and Husain (2008) suggest that this difference in memory trace strength for to-be-fixated items (measured by the precision with which a change in location or orientation of a stimulus can be detected in their study) is accounted for by redeployment or memory resources away from the items that are no longer the focus of covert attention. In other words, attention brings limited encoding resources to bear on a particular item even before that item has been foveated and in this way acts as the gatekeeper of WM.

Recall from our discussion in Chapters 1 and 6 the notion that WM encoding proceeds according to a two stage process with proto-object representations set up first in a parallel encoding fashion and then detail for items

filled in when they become the subject of attention. It is possible that covert attention initiates the beginning of filling in of detailed information for representations. Another possibility is that the proto-object representation is strengthened by covert attention to the stimulus, perhaps by an increase in the proportion of neurons in the coding population that synchronise to represent the low level features of the stimulus and/or a decrease in the level of neural noise in surrounding populations. A third alternative interpretation of Bays and Husain's (2008) finding would be that covert attention plays no direct role in facilitating visual WM but instead simply serves to guide internal (post encoding) attention to the particular region of the (retinotopically mapped) internalised iconic representation that is left after stimulus offset (Coltheart, 1980). In this way the iconic (or fragile; Sligte, Vandenbroucke, Scholte & Lamme, 2010) representation of a to-be-saccaded-to stimulus might be prioritised for consolidation into visual WM, or prioritised for elaborated post-encoding processing within WM.

The data I have reported herein illustrate that, while attention may indeed act to facilitate encoding of stimuli into WM (e.g. as evidenced by overall greater attentional capture and better WM performance in memory for singleton vs. non-singleton faces regardless of valence association in Experiment 6), factors independent of attentional capture seem to contribute to boosting of motivationally salient and emotional faces: First, the WM boost to angry face singletons in Experiment 1 was observed not to be correlated with attention (gaze dwell time) to the face in question. Second, arrays of all angry faces were observed to have a boosted WM performance level relative to arrays of all happy or all neutral faces in this experiment. Third, a similar boosting of positive outcome associated singleton faces relative to negative outcome associated singleton faces and neutrally (no

outcome) associated faces was observed in Experiment 6 despite no significant difference in attention deployment (gaze dwell time) between these stimulus types. Fourth, WM performance was observed to vary as a function of predictability of positive outcome for faces that were presented singly on screen (i.e. in a presentation method where no other stimulus competed with them for attention at encoding) in Experiment 8. Fifth, among low WM performance participants, incentivisation following a win event was observed to effect a WM boost for an entire array of coloured squares (another presentation scenario where there were no other 'lower importance' stimuli present from which attention could have been withdrawn and redeployed).

It seems, therefore, that something other than redeployment of a limited attentional resource must account for the boosted performances we observed across different experiments. One possibility is that the motivational salience of a stimulus is rapidly evaluated (see below) and this 'tags' the stimulus representation as being worthy of elaborated processing. Such elaboration might be subserved by two discrete processes; a propensity to capture exogenous attention during encoding as evidenced by the widely reported attentional capture by threat related items, and an orthogonal process of post encoding elaboration that proceeds even when task demands result in no differences in deployment of attention at encoding between motivationally salient and non-motivationally salient items. The nature that such post encoding elaboration process might take is discussed further below.

Representational precision and attention

A recent study by Fougne et al. (Fougne, Suchow & Alvarez, 2012) fitted fixed-precision and variable precision models to histograms of errors made by individual participants on 'remembered item' trials (see below) on a standard WM colour wheel task. In this task, participants viewed a memorisation array comprising sets of circles of different colours. Shortly after memory array offset the location of one of these memory items was cued. Next, a continuous circular colour wheel was presented and response was made by indicating what colour the cued memory item was by selecting an appropriate corresponding location on the wheel. Two types of errors are typically made in this task; 'unremembered' and 'remembered' item errors (see also Zhang and Luck, 2008). Unremembered item errors reflect cases where the cued item is not in WM and are characterised by selection of random locations on the colour wheel. Remembered item errors reflect instances where the item is in WM but is held with imperfect precision. Such errors cluster around the correct response location on the colour wheel. Fougne et al. (2012) found that a variable precision model better fit patterns of errors on remembered trials than did a fixed precision model, suggesting that individual representations of simple stimuli (coloured circles) are encoded into visual WM with varying levels of quality, such that the variability in the precision with which a stimulus can be recalled (the ability to accurately plot its colour on the colour wheel) varies across representations of multiple stimuli encoded on a single trial in a stochastic manner. While this finding is not necessarily surprising in its own right (given that WM items are encoded by neural populations, and stochastic physiological processes are known to operate within neural populations in many instances), what is particularly interesting is Fougne et al.'s (2012) failure to find any impact of the precision with which one stimulus representation is held on other

representations encoded into visual WM on the same trial⁵⁹. This suggests that it is quite possible to increase the precision with which one stimulus is encoded without impacting the precision of others. This view has parallels with my interpretation of the likely mechanism underlying the angry (Chapters 5 and 6) and positive value (Chapter 7) face boost effects, as well the mechanism by which general incentive may serve to boost WM performance (Chapter 9): In each of these cases I have argued that neurally efficient processing of stimuli involves a trade-off between the strength of representations in WM (or even their creation in visual WM in the first place) and the 'neural energy' or level of cognitive resources that must be expended to facilitate such representation. In each case, I have further argued that the situational relevance, motivational salience, or importance of stimuli, whether modulated by face expression, prior outcome learning or incentive availability signalling can determine the extent to which a stimulus is able to recruit from some central pool of resources some of the commodity that allows it be represented with precision appropriate to its relevance to the participant.

In another parallel with the data I have reported here, Fougne et al. (2012) did not find any effect of deployment of attention among to-be-remembered stimuli on the relative precision with which they were encoded, a finding paralleling our null effects of attention on determining precision of representations of face stimuli (Chapters 5 and 7).

A third finding of Fougne et al.'s (2012) study is that the precision with which stimuli are represented in visual WM seems to be available as meta-mnemonic information to participants; i.e. participants are aware of the relative

⁵⁹ While this was a null effect finding, Monte Carlo simulations showed that if an effect of trading off precision of one representation against another had been present their paradigm would most probably have been able to detect it.

quality of each concurrently held WM representation. This was evidenced by the observation that when allowed to report the colour of the circle they remembered best, participants performed better than when the memory item to be reported was chosen at random for them. This meta-mnemonic ability might well play an important part in the determination of recruitment of internal resources to support representation; participants may be able to monitor the precision with which a representation is held, and compare this to the relative importance of the generating stimulus (perhaps utilising a mechanism involving CSTM; see below). In this way, representation strength determined during encoding could be modulated to accord with the appropriate 'importance level' of the representation to be held.

The process suggested above might very well serve to guide exogenous attention, and indeed there is ample empirical evidence that top-down goals and bottom up visual salience effects serve to direct visual attention. However, the experiments reported herein and Fougne et al.'s (2012) data illustrate that modulation of exogenous attention is not the only process by which precision of representations may be modulated. Indeed, attention need not be exogenously oriented at all in order to generate a representation in visual WM, since an instruction to imagine a particular visual entity can lead to generation of a representation of it in visual WM.

What neural processes might underlie boosting of WM?

In each instance where I have observed boosted WM performance herein, I have argued that 'resources' are brought online from a central pool. While I have

discussed above the way in which such resources may be deployed in an item-based fashion, I have not yet discussed the possible nature of such resources in neural terms. The experiments reported herein were not designed to directly address this point, but it is instructive to consider several possible neural mechanisms by which the cognitive 'boosting' processes described might be realised.

Leaving aside questions of functional localisation, disparate neural coding for slots vs. detail, visual binding and the various other unresolved debates in the neurophysiology of WM, we can in the broadest sense make a non-controversial statement; that the representation of an item in WM is coded for by a population of neurons working together. However, even here the issue of how some items might be better remembered in WM than others is open to multiple interpretations. One possibility is that larger populations of neurons are activated in instances of better remembered items. In this way a certain amount of signal redundancy might be achieved. Alternatively, it is possible that better memory is a function of reduced neural noise among a particular coding population, i.e. a lower proportion of neurons in the coding population are firing in such a way as to fail to code for the stimulus attribute it is optimal that they code for. A third possibility is that better memory is achieved by a resource acting at an early stage such that a more veridical representation is encoded in the first instance. In this case neurons in the coding population do not differ in their ability to maintain a representation accurately but the representations of better remembered items are encoded and held with a greater degree of veridicality. That there is a natural variability in the veridicality with which representations are held is supported by Fougine et al.'s (2012) study discussed above. A fourth possibility comes from the observation

that different neural structures may be recruited for WM maintenance depending on whether demands are high or low, with ventrolateral PFC (vlPFC) recruited in tasks with low WM demands, but additional dorsolateral PFC (dlPFC) required for tasks with higher demands (Rypma & D'Esposito, 2003). Here, we have the possibility that an additional brain area is recruited to subserve elaborated processing. Note that the angry face benefit, at least for uniform emotion arrays, has been observed to be accompanied by increased activity in the right superior temporal sulcus (STC), vlPFC, and basal ganglia (Jackson, Wolf, Johnston, Raymond & Linden, 2008) but not the dlPFC.

It would be nice if we could localise exactly the brain areas that might play a role in situations where visual WM is boosted above usual levels, and indeed some research has been done in this regard. In Chapter 1 I discussed the notion of fragile memory, a memory store posited to occur between iconic and short term memory. Research on fragile memory has used retro-cues to direct the attentional gatekeeper to the fragile VSTM representation to be encoded into visual WM, where its presence can then be tested for (Sligte, Scholte & Lamme, 2008, 2009; Sligte, Vandembroucke, Scholte & Lamme, 2010; Vandembroucke, Sligte & Lamme, 2011; see Chapter 1). A recent study (Kuo, Yeh, Chen & D'Esposito, 2011) using a similar methodology has illustrated that the behavioural performance boost accorded to items that receive such retro-cues may, at least in the case of cues presented during the iconic memory decay interval, be due to top-down processes that cause a temporary increase in the functional coupling of posterior visual and more anterior (parietal and frontal) cortical areas. Kuo et al. (2011) utilised a paradigm in which retro-cues signalled which item from the memory array (an array of coloured squares) would be tested at probe. When

cues were presented early in the retention interval (200 ms after stimulus offset) behavioural performance was boosted and fMRI indicated a functional coupling of middle frontal gyri and frontal eye fields with regions in left and right occipital cortex (the seed regions of the coherence analysis). This functional coupling followed a phase relationship which was stable across participants, with frontal regions showing increased activity earlier than occipital regions, and was correlated with behavioural performance. Thus, presentation of early retro-cues seems to boost frontal activity, which in turn creates a top-down modulation of activity in occipital (and some parietal) regions.

Unfortunately, the presentation timing of cues in Kuo et al.'s (2011) study does not allow us to distinguish whether they act on iconic or fragile memory representations.⁶⁰ Kuo et al. (2011) argue that the simple visual cues presented at this very early stage in the visual memory process effect a top-down modulation of WM performance (as evidenced by the involvement of frontal as well as parietal regions). If this is so, then it is suggestive of a very early acting evaluation process which identifies items of importance to fulfilling task demands (i.e. something akin to CSTM; see below). However, it is possible that the process by which a cued representation is boosted in fact occurs in a bottom-up fashion, with cues simply serving to direct internal attention to the region in the retinotopically mapped contents of iconic- or fragile memory where the representation of the stimulus in question lies. In this way cued items might be better consolidated into visual WM than uncued items by virtue of being the beneficiaries of such internal attention.

⁶⁰ Cues presented 800 ms after stimulus offset did not boost performance. However, such cues occurred only 100 ms before onset of the test array in Kuo et al.'s study, so they may not have had time to facilitate WM performance before the test array was presented, a limitation acknowledged by Kuo et al. (2011).

Additionally, bearing in mind the role of the middle frontal gyrus (a region roughly corresponding to dlPFC) in supporting WM (Andrews, Hoy, Enticott, Daskalakis & Fitzgerald, 2011; Courtney, 2004), it is possible that it is the increased activity in this brain region rather than those activated later that directly modulates WM task performance and that the activity then engendered in other brain regions, while observed reliably, might be incidental to task performance.

A further difficulty in determining the neural region that supports boosts in WM comes to light when we consider the notion that stimuli of motivational importance may be 'tagged' for elaborated processing. Such neural tags are likely to be realised in an executive function supporting brain area while not necessarily themselves forming part of the population of neurons that directly codes for the stimulus representation that they refer to. In this regard, activation of frontal areas in WM tasks is likely to represent elements of cognitive control (e.g. Chatham et al., 2011; Yeung, Nystrom, Aronson & Cohen, 2006) rather than direct maintenance structures. Thus, we must so far remain cautious in ascribing any particular area the role of the 'direct booster of visual WM'. Further research needs to be launched to establish which interpretations are most appropriate.

In speculating how boosted WM performance may be achieved neurally, we must also bear in mind the time-course of events that occur following perception of a visual stimulus. Whereas traditional models of visual processing postulate a hierarchical feedforward sweep across visual areas, more recently models of visual object recognition have begun to incorporate the notion of feedback from higher visual areas. An interesting example of this is a study by Wyatte, Curran & O'Reilly (2012) in which a model incorporating feedback from higher visual areas (inferior temporal cortex, V4 and V3) to V1 predicted object

identification performance when stimuli were visually degraded (objects that were occluded and/or presented in low contrast). This model specifically incorporates a reintegrative process so that degraded signals can be strengthened by iterative recurrent processing between visual areas. It is possible that an increase in such recurrent processing may be responsible for the boosts in WM we have observed herein if such boosts reflect increases in the resolution of stimuli. In this view, the boosts we have observed might reflect elaborated recurrent processing whereby salient representations are 'sharpened in focus' so that neural firing which does not code for the stimulus in question is suppressed. How such elaborated processing might be directed in an item specific manner may involve direction of elaborated processing via a saliency map. Interestingly it has been hypothesised that saliency maps may be realised as early as area V1 (Li, 2002).

What implications do these findings have for existing cognitive models of WM?

Item-specific memory boosts

In Experiments 1 and 2 we observed that when angry singleton faces receive a boost in WM this occurs without either positive or negative impact on memory for concurrently presented neutral faces. Likewise, in Experiment 6 we observed a similar pattern of results when singleton faces had been associated with a positive outcome. In Experiment 8 positivity of a face stimulus was observed to mediate its memorability without impact on a secondary WM task. Together, these results form the basis of my argument that additional WM resource may be brought online in an item-specific fashion (see Figure 9.1).

Support for this idea comes from recent evidence which suggests that among high socially anxious individuals presentation of angry faces leads to sustained prioritized sensory processing of the angry face relative to happy or neutral faces, perhaps facilitated by re-entrant connections from the amygdaloid complex to sensory areas including V1, but such increased sensory processing is not at the expense of other concurrently occurring faces (Wieser, McTeague & Keil, 2011) nor at the expense of attentional resources required to perform a separate discrimination task (Wieser, McTeague & Keil, 2012⁶¹), a finding directly paralleling the suggestion I make Chapter 5; that angry faces do not compete for resources at the sensory processing level but instead lead to both an increase in the overall level of resource available and an attraction of this boosted resource to themselves such that other co-present stimuli do not suffer in processing.

Do 'resources' deployed to boosted stimuli simply reflect the 'focus of attention' or episodic buffer?

Perhaps we need look no further than existing cognitive models of WM for an explanation of the mechanism by which certain WM items receive a boost in representational strength. In Cowan's model of WM the 'focus of attention' is a central resource that facilitates WM encoding and maintenance and is shared across visual and auditory WM stores. Importantly, Morey and Cowan (2005) have argued that it is possible to do certain WM tasks without drawing on this resource. This provides an explanation for the discrepancies found by different researchers when conducting dual visual-auditory WM tasks in which participants

⁶¹ Interestingly, this study found no decrease in behavioural performance for a Gabor patch temporal discrimination task when angry faces relative to happy or neutral faces were present, but did find a decrease in activation of neural areas associated with processing of the Gabor patch under such conditions, suggesting an increase in general processing efficiency on presentation of an angry face.

must maintain visual and auditory WM information concurrently. In such tasks one would expect to see a trade-off between visual and auditory capacity if the capacity for each modality store could be supported by the focus of attention separately but no trade-off if this central resource could not be switched between modality stores. Whereas some studies have found such trade-offs, others have not. If, therefore, it is possible to do certain visual WM tasks without calling overmuch on the focus of attention then this implies that the focus of attention remains idle, and is a suitable candidate mechanism by which WM representations identified as valuable might then be boosted.

A further consideration in support of this interpretation is that the focus of attention seems to trade off maintenance with manipulation of information (Bunting & Cowan, 2005). Manipulation of information at the most general level might be conceptualised as 'thoughts'. If we consider that in the majority of experimental settings it is quite possible that participants have thoughts of their own, related to subjects not relevant to the experiment then suppression of these thoughts when stimuli coded as 'important' are presented for memorisation would lead to a lower load on the focus of attention, resulting in its increased ability to aid in maintaining WM representations. To put this in simpler terms, important stimuli may cause participants to focus in the task in hand. Of course, this explanation does not offer a rationale for how object representation get boosted in an item specific way, merely that it is likely that idle WM resource of this type is commonly present to be called upon when required.

In a similar way, one can interpret the discrepancies in findings in the case of trade-off of visual and auditory information in dual tasks in the context of Repovš and Baddeley's (2006) model of WM. Here, the crucial component to

consider is the episodic buffer. While the episodic buffer is a non-modality specific store, capable of supporting information derived from both auditory and visual domains in a non-modality specific format, it is quite probable that not all WM tasks result in full use of the buffer, such that only when the buffer begins to become full will trade-offs be observed between the auditory and visual components of dual auditory-visual tasks. In this interpretation it is the recruitment of additional amounts of the episodic buffer that result in boosting of important stimulus representations, a process that seems qualitatively similar to the explanation of boosted WM performance of experts in their field of expertise offered by template theory (Gobet & Simon, 1996; see Chapter 1).

The resources-within-slots model

Although Cowan's central executive and Baddeley's episodic buffer can each account for the presence of 'idle' resource, deployable when required to boost representations, neither specifically posit a mechanism for holding resource in abeyance nor do their models contain mechanisms for resource deployment to one representation without cost to others (except in cases where memory load is low enough for a slot to be assigned to each representation). Figure 9.1 below illustrates a simple resources-within-slots model which incorporates both of these mechanisms.

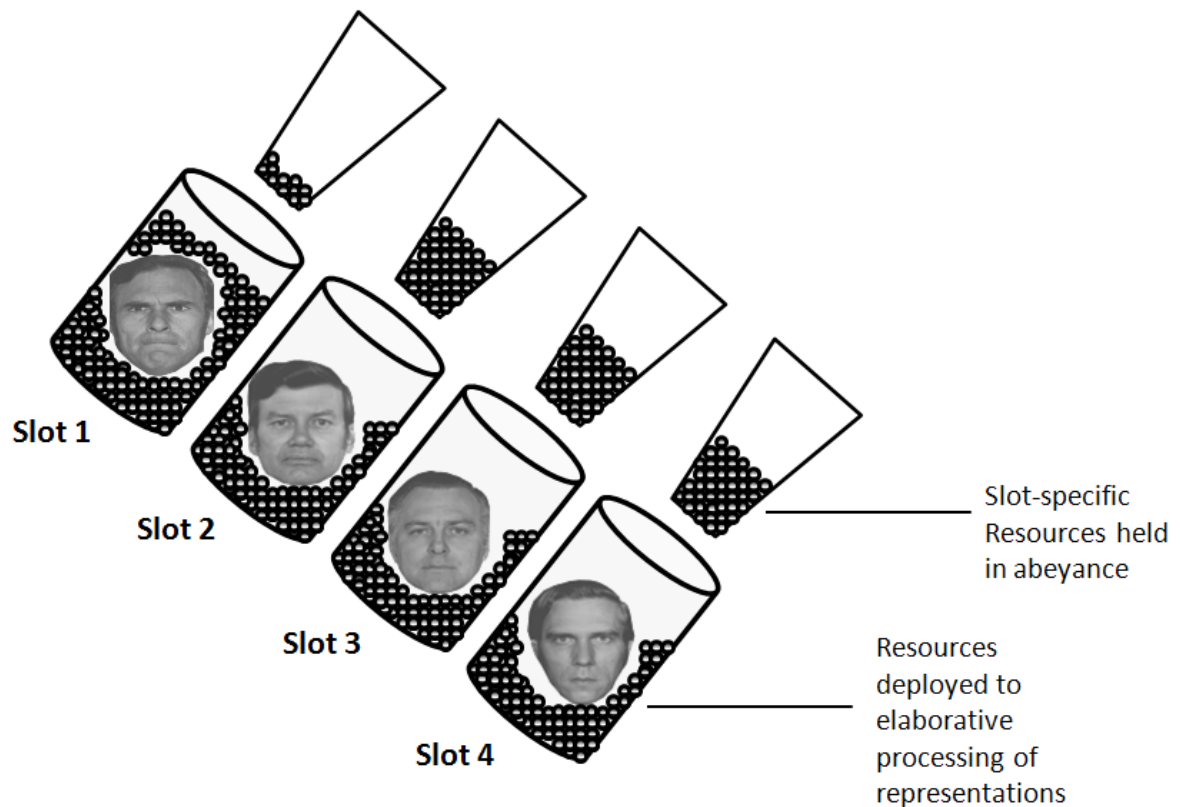


Figure 9.1. A resources within slots model of WM. Cylinders represent individual slots, the contents of which comprise individual object representations (in this example face stimuli from Experiment 1 herein). Funnels represent allocation mechanisms which permit deployment of resource to individual slots. Small circles represent memory 'resources' which may represent populations of stimulus coding neurons or higher, more executive level neural populations which may play a role in limiting neural noise in populations which code for individual objects. Here, Slot 1 is shown containing a motivationally salient (angry) face and so has attracted additional resource to facilitate elaborated processing, resulting in a more detailed WM representation. Bringing resource to bear is effortful so for each slot a certain amount of potential resource remains undeployed.

The resources held ‘in abeyance’ in Figure 9.1 are shown within four separate slot-specific ‘funnels’. However, a variant of this model which would fit equally well with my data is one wherein the slot-specific funnels are merged into a single central resource provisioner. This model is more in keeping with existing models of WM that posit a central, often modality agnostic resource boosting mechanism such as the central executive (e.g. Baddeley, 1996) or the ‘focus of attention’ (Cowan, 2001). While both the variants of the model I have described are of a resource-within-slots type, we must for now remain agnostic as to which of them (or what form of hybrid between them) might best represent the visual WM processes observed herein. The holding back of resource illustrated by the less than full deployment of resource available into slots 2, 3 and 4 in Figure 9.1 is also supported by our observation that motivation can cause deployment of additional resource among participants not otherwise performing at their optimal level (Experiment 9).

Whereas most slot-based models of WM portray slots simply as somewhere where a representation is held in memory, our observation of item specific WM boosting by angry and positively valenced faces suggests that slots may serve a second function, as delineators of representations that serve to mark out the bounds of one stimulus from another and so ensure that item-specific resource provisioning can occur.⁶²

⁶² Possibly a better description of slots in this regard is as ‘buffers’ between representations, preventing as they may do the spilling over not only of resource allocated but of ‘representation signal’ between representations (i.e. serving to prevent contamination of one representation by other concurrently memorised representations). However, buffer has a meaning both as a physical separator and as a temporary storage device (e.g. Baddeley’s episodic buffer). So, to avoid this ambiguity ‘delineator’ seems a better description here.

The key features of the model described in Figure 9.1 that distinguish it from existing models are the allocation of a flexible central resource differentially between slots and the holding back of a portion of the full amount of this resource available.⁶³ This model therefore parallels to some extent the observations made by Xu and Chun (2006, 2007, 2009; Xu, 2007, 2008) that object ‘individuation’ (setting up proto-objects) may be slot-based and subserved by the inferior IPS, and object identification may be based on a continuous pool of resources, subserved by the superior IPS. However, this model so far lacks a mechanism to determine where (to which slot(s)) resource is to be allocated. Any such mechanism must be evaluative in nature; that is, it must differentially determine resource allocation to different stimuli based on their characteristics. The type of stimuli that may be evaluated as worthy of boosting in WM are discussed above. How information may be rapidly extracted from stimuli to determine whether they are of this ‘boost-worthy’ type is discussed in the next section below.

How does the brain ‘know’ when to boost a representation?

Implications of results for early memory processes

In Chapter 6, I suggested that a rush to consolidate involving expedited closing of object files might underlie the failure to observe a significant effect of face emotion on stimulus memorability in Experiment 4 where there was ambiguity as to the amount of time available to consolidate memory representations of face stimuli. However, recall from our discussion in Chapter 3 that time estimation in

⁶³ Note that although four slots are shown in the exposition of the model in Figure 9.1, the model itself makes no specific claim as to the number of slots, nor does it suggest that a representation cannot be held in multiple slots simultaneously (the slots-averaging mechanism of Zhang and Luck, 2008; see Chapter 1).

humans requires WM resource to store the 'clicks' of an internal pacemaker or clock (Gil & Droit-Volet, 2011) and that there exists, therefore, a trade-off between accurate estimation of time and performance on non-time-based, concurrent WM tasks. It is possible that ambiguity in time available leads to an over-deployment of processing resource to the process of time estimation (i.e. an attempt on each trial to estimate the amount of time that passed between memory array offset and mask onset). This is especially likely since participants were unaware of the differences in timing between blocks of trials and universally reported not noticing any difference in timings between blocks when questioned after the task. Such a deployment of processing resource would be an optimal strategy in situations where there was no variance in the length of consolidation intervals between different trials within a block (as in Experiment 3) since it allows for the establishing of an accurate estimation in LTM of the time available for consolidation on each trial and so facilitates closing object files before mask onset. However, in situations where there is variance between consolidation intervals (as in Experiment 4), the internal LTM estimate of available time to consolidate must be constantly reset on each trial where observed time available does not match to the currently held internal estimate.⁶⁴ If the process of resetting (re-estimation and entry into LTM of the new estimate) requires resources that could otherwise be deployed to support prioritised memorisation of angry faces then this may contribute to the failure of angry faces to accrue unto themselves a boost in processing in Experiment 4.

⁶⁴ It is reasonable to assume a model in which more current trials are given greater weighting in their effect on estimation of the interval time as this allows changes to be quickly encoded. It is beyond the scope of this thesis to quantitatively address likely weighting scenarios, which may change from task to task in any case.

Whereas monitoring the passage of time and updating thoughts, task rules and goal states may involve accessing elements of LTM, maintaining stimulus representations can involve encoding of changes in the visual world (i.e. sensory information). Nonetheless, a shared single fronto-parietal network including supplementary motor area (SMA), parietal, left inferior frontal junction, and middle frontal gyrus seems to be activated when WM is updated regardless of the source (LTM vs. sensory) of the new information (Roth & Courtney, 2007) or whether the new information relates to the representations within WM or the rules governing relationships between them (Montejo & Courtney, 2008). This shared neural architecture suggests that a trade-off might have to be made between maintaining and updating the ‘thought-like’ (semantic) contents of WM, including time interval estimation, and the stimulus-type contents of WM (be they visual or articulatory traces).⁶⁵ Indeed, some models of WM posit mechanisms to shield the goal state maintained in WM from disruption by perceptual input (Botvinick, Braver, Barch, Carter & Cohen, 2001; Frank, Loughry & O’Reilly, 2001; O’Reilly, Braver & Cohen, 1999)

The ability of a stimulus to fulfil a task rule or satisfy a goal state depends on the context within which the stimulus is encountered. Food can only satisfy a hunger-based goal when it is encountered in a context where it can be consumed; a pen is only useful if I have paper to write on. Emotionality has been implicated in the facilitation of binding emotional stimuli to their context (Hadley & MacKay, 2006) in ‘immediate’ memory (LTM tested within a few minutes of memorisation) and this appears to be at the expense of concurrently presented neutral stimuli,

⁶⁵ This is not to say that there is no neural difference between updating of these different information types: For example, updating of rules has been seen to preferentially activate certain parts of the circuit such as the left inferior frontal junction (IFJ), whereas updating of numbers in WM has been seen to preferentially activate other areas such as bilateral intraparietal sulcus (IPS).

revealing that the mechanisms by which such binding occurs are both limited in nature and biased towards emotional stimuli. However, in order that a stimulus representation be bound to its context, it must first be evaluated with regard to its context. For example, efficient memorisation of the position of a chess piece in relation to other pieces on the board (binding of the stimulus to its context) will, at least for reasonably experienced players able to evaluate positions, involve an analysis of the influence that the piece in question has on the overall position (evaluation of stimulus with regard to context). A proposed mechanism by which such evaluations are made is conceptual short term memory (CSTM; Potter, 1993)

Conceptual short term memory (CSTM)

“...most cognitive processing occurs on the fly, without review of material in standard short-term memory and with little or no conscious reasoning.” (Potter, 1999, pp 13-14),

In Chapter 1, I discussed briefly several models of expert memory. While these models (the chunking model, LT-WM and template theory) each propose different mechanisms underlying the enhanced behaviour of experts, they all agree that such enhancements are (1) in some way facilitated by rapid access to LTM, (2) occur for experts only and, (3) depend on the presence in LTM of conceptual information rather than merely visual information.⁶⁶

However, another model exists that posits similar rapid LTM access to conceptual information in a way that is not specific to experts. This idea is called conceptual short term memory (CSTM; Potter, 1993, 1999, 2010; O’Connor &

⁶⁶ Even chunking theory, which suggests the representations are held in visual form suggests that semantic associations in LTM facilitate the formation of chunks.

Potter, 2002). This model seeks to account for the very rapid access to semantic information observed in rapid serial visual presentation (RSVP) paradigms such as the attentional blink (AB; Shapiro, Arnell, & Raymond, 1997) by positing a complementary mental buffer sitting alongside 'classic' WM and iconic memory, in which representations and temporarily activated associated LTM codes can be both briefly held and subjected to largely non-conscious semantic processing such that patterns may be identified, representations categorised, and goal-congruent representations highlighted. Representations in CSTM are of a conceptual, non-modality-specific nature, as evidenced by the fact that conceptually related but modality-different (pictorial vs. verbal) representations presented rapidly in sequence to a participant have the effect of a strengthened, integrated overall representation (O'Connor et al., 2002). Furthermore, this model suggests that information that fails to be integrated into some semantic concept, either because it fails to activate LTM codes or is not goal congruent will be rapidly discarded or forgotten from the store. CSTM models are agnostic as to whether such discarded information ever reaches consciousness (Potter, 1999). CSTM is assumed to operate within the first several hundred milliseconds following stimulus presentation (O'Connor et al., 2002).

Evidence for the ability to actually process and define patterns in CSTM comes mainly from RSVP studies using words in a sentence (Forster, 1970; Potter, 1993). In particular, a study by Potter, et al. (Potter, Kroll, Yachzel, Carpenter, & Sherman 1986) has shown that it is possible when given a sentence presented one word at a time in an RSVP stream to extract both syntactic and semantic plausibility from the sentence, whereas when presented with words at a similar rate but not in sentence structure only a few words can be recalled in a test

immediately post-presentation, indicating that semantic structures or patterns within information allows it to be retained as an integrated whole, while absence of such semantic 'glue' leads to decay of the individual volatile memory traces.

The ability to very rapidly derive semantic information from a stimulus, through processing in CSTM may facilitate the process of divvying up later WM processing resources between competing stimuli by quickly determining which stimuli are 'worthy' of resource. There are numerous candidate mechanisms by which such a divvying up of resources could be encoded. In Chapter 3 I assessed notions of saliency and the idea that saliency maps guide the allocation of resources. It could be, therefore, that the contents of CSTM serve to bias saliency maps, resulting in greater resource allocation to items of conceptual priority. While I assessed the notion of saliency for guiding deployment of attentional resource in the visual field, one thing we have not considered thus far is whether saliency maps may remain active at a post-perceptual stage, that is, after offset of a stimulus. Some process must determine which elements of iconic (or fragile) memory are selected for consolidation into WM, or which aspects of detailed information are selected to be added to the proto-objects set up in the process of object individuation. A saliency map might seem a reasonable mechanism for making these determinations, but the natural objection to this is that setting up a saliency map takes time and fragile early information exists only transiently. One possibility is that CSTM serves to modify a continuous (i.e. continually active) saliency map on the basis of both the contents of CSTM and the contents of WM. While models of CSTM do not currently include such a link to WM, it does not seem unreasonable to posit such a link so that CSTM may be 'aware' of the contents of WM. Indeed, one of the features of CSTM is its ability to rapidly

determine the goal-congruency of a stimulus, and if current goals are held in WM, then the addition of this type of link between CSTM and WM seems a natural addition to the model. Indeed, maintaining a conceptual representation of an entity in WM leads to an attentional blink when a pictorial representation matching the maintained representation is presented in a RSVP stream (Pashler and Shiu, 1999); again suggesting that CSTM has access to the contents of WM.

The notion of CSTM fits nicely with much of our data. It can account for the angry face benefit and the item-specific nature of this benefit revealed in Chapter 5 via a process in which rapidly activated LTM associations with anger (driven by the connections that angry face representations in CSTM have to LTM codes) flag the face for additional processing. It can account for the WM benefit for positively value-laden faces by allowing for rapid matching with a previous (though still partially activated) goal state. That the rapid evaluation of stimuli depends in part upon analysis of how they fit into a particular context offers an interesting explanation for the finding that faces with positively predictive outcomes receive a WM boost only when they are singletons: Being a singleton within an array allows comparison to other non-singleton faces and thus a richer context within which to evaluate a face stimulus.

Perhaps more pertinently, CSTM has bearing on our observation that masking can disrupt the angry benefit effect in a fashion that is orthogonal to its disruption of perceptual processing. If the process of backward masking prevents access to CSTM then this offers an explanation as to why backward masking used in Experiment 3 suppresses the angry face benefit even when such masking is not seen to interfere with perceptual consolidation of our stimuli; if it is not possible for the memory system to determine at an early stage that the face should be given a

processing boost, then the face in question will be treated just like a neutral (or happy) face would be in this regard. This is particularly interesting since it suggests that it is vital that the face is both identified as angry and then evaluated in the context of the contents of WM (in the case of the task reported in Experiment 6, other, co-present faces) *at this very early stage after stimulus offset*.⁶⁷ An angry face consolidated into WM without this additional coding in CSTM does not benefit from the usual angry boost effect. Perhaps once WM resources are already engaged in maintenance it is too late to alter the level of maintenance resources that are allocated to each representation and this can only be determined during very early memory processes. While this process would explain my results, the notion that it is access to CSTM that is blocked by backward masking in our task is an assumption that requires further research.

It is worth briefly noting the similarities in the proposed functions of CSTM and Baddeley's (2000) episodic buffer. Both are suggested to hold stimulus representations on a temporary basis and facilitate rapid evaluation of stimuli within the context of both task demands and other co-present stimuli. The main difference between these two constructs is the stage at which they are proposed to act in the WM process, in particular the very early proposed onset (<100 ms after stimulus onset) of CSTM processes.

Mathematical and neurological modelling of attentional and VSTM selection processes serving a similar function to rapid evaluation in CSTM has been undertaken by Bundesen, Habekost and Kyllingsbæk (2005, 2011). In this model (the Neural Theory of Visual Attention; NTVA) rapid evaluation and

⁶⁷ Note that ample time was provided to make the determination that the face was angry in the task reported in Experiment 6. However, it is possible that CSTM only evaluates stimuli with regard to context upon stimulus offset in order not to overtax cognitive resources (see discussion in Chapter 6).

categorisation of all stimuli in the visual scene modulates their weights in a winner-takes-all race to encode, which determines their likelihood of entry into WM. Of particular interest in this model is the idea that likelihood-of-encoding weights are derived in part from a stimulus's *pertinence*; the momentary importance that the stimulus has to the observer. This is perhaps one mechanism by which boosts in WM performance for items of motivational salience may be realised.

Limitations and directions for future research

While I have advanced the understanding of some of the cognitive mechanisms that may underlie differential performances in WM tasks, the ecological validity of these findings are constrained by certain limitations imposed by the nature of the stimuli and paradigms used.

Stimulus limitations

The stimuli I presented to participants in the tasks herein were all devoid of visual context; that is, they were presented as to-be-memorised entities but connections between them and other entities need not necessarily be formed as part of the task demands. Generally, stimuli outside the experimental setting are rarely encountered devoid of semantic context, and one aspect of motivational salience that we have not addressed here is how conceptual links between memoranda and context might be modulated by salience and emotionality.

Additionally, I have made an assumption throughout that happy expressions are generally positive in nature. Interestingly, this assumption may be invalid for certain groups of people. Schultheiss et al. (Schultheiss, Pang, Torges,

Wirth & Treynor, 2005) have recently shown that for highly 'power-motivated' (social dominance seeking) individuals faces expressing joy and anger both interfere with learning of visuomotor sequences, while for affiliation-motivated (social acceptance seeking) individuals only angry faces interfere with learning. This suggests that it may be more appropriate to think of facial expressions in terms of their motivational salience (that is, their signalling of an expresser's likelihood of facilitating or frustrating the perceiver's social goals) than their valence.

The face stimuli I used throughout were faces that 'look' directly out of the screen, towards the observer. However, gaze perception may be vital in understanding the emotional expressions of others (Baron-Cohen, & Cross, 1992; Baron-Cohen & Ring, 1994) at least in children over about 3 years of age (Doherty, 2006). In particular, the interpretation of anger differs according to gaze since eye contact may indicate anger at the observer whereas when gaze is directed in another direction this may indicate anger toward an entity in the direction in question. Indeed, processing of emotional faces when gaze is towards the observer rather than averted has been seen to preferentially activate the anterior region of the superior temporal gyrus (STG; Wicker, Perret, Baron-Cohen & Decety, 2003), perhaps due to automatic initiation of theory-of-mind computations. Gaze has also been seen to modulate the perception of attractive faces, with averted gaze associated with reduced activation of the ventral striatum, a dopaminergic region strongly linked to reward prediction (Kampe, Frith, Dolan, & Frith, 2001). Additionally, detection of direct gaze (eye contact) has been hypothesised to involve a 'fast-track' subcortical pathway allowing its rapid detection (Senju & Johnson, 2009). Further research would profit from

investigating whether the angry face WM boost effect is invariant across gaze directions or whether it might be the social implications of directed-to-the-observer anger that underlie the boost.

Paradigm limitations

One limitation of each of the experiments reported herein is the specificity with which behavioural tasks can obtain a pure measure of visual WM uncontaminated by either contributions from verbal WM or constraints imposed by limited attention. It has been suggested that performance on even the simplest of visual WM tasks, such as the coloured squares task used herein, actually reflects a mix of WM and LTM processes with LTM being a crucial part of the processes by which WM representations are rehearsed (Jeneson, Wixted, Hopkins & Squire, 2012). While it is not universally agreed that LTM underlies rehearsal, especially for items that have no prior representation in LTM (i.e. novel stimuli), this is an interesting idea that is in accordance of Cowan's (2001) model of WM (where WM representations depend on temporary activations of LTM representations).

Theoretical limitations

My data do not reveal whether value associated colour boosts WM in an object-centered way (for the face it is associated with in the WM task) or in a spatial way (by directing internal attention or 'central executive' resources to the region within a retinotopically organised WM array representation that includes the face representation in question, thus facilitating consolidation or maintenance). There are at least two possible mechanisms by which colour could 'carry' value following a learning task such as those reported in Experiments 6 and 7; either as a stimulus attribute semantically coded to value; "all red items are valuable" or as a spatially co-occurring value signal; "all items in a red region are valuable".

Additionally, while we have established that early WM mechanisms are necessary for emotional faces to receive a WM boost (Chapter 6), it remains an open question as to whether the boost itself is effected by these early mechanisms or whether they merely serve to ‘tag’ a stimulus and designate it worthy of elaborated processing at a later stage.

We must also apply a caveat to the theoretical claims made herein with regard to rewarding and punishing stimuli: While it is possible in an experimental setting to apply a learning paradigm backed by reasonable rewards and punishments, punishments and rewards beyond a certain level cannot be applied for ethical reasons. We must remain aware that particularly punishing stimuli might develop pairing with emotions such as genuine fear rather than mild displeasure and behave in a quantitatively different manner (i.e. in the case of the studies reported herein we cannot rule out a WM boost applied to stimuli with learned very high punishment contingencies).

Future directions

“You can’t play 20 questions with nature, and win.”⁶⁸ (Newell, 1973)

Understanding of the neurology of cognitive systems is increasing and future research must link cognitive, neurological and mathematical models of WM together and identify ways in which WM capacity (performance ability on WM tasks) varies across different scenarios within an integrated overarching framework that includes not only WM but also cognitive control systems, emotion systems and rapid semantic evaluation systems so that we can gradually build up

⁶⁸ Newell’s point, made in 1973 and again in his book “Unified Theories of Cognition” (1990) is that while experiments that answer binary questions are valuable, such answers obtained should always be evaluated in light of a total systems model of the mind / brain rather than being treated in isolation.

a picture of how the *total system* of the mind and brain allows us to function in day to day life.

References – Chapter 9

- Andrews, S., Hoy, K., Enticott, P., Daskalakis, Z. & Fitzgerald, P. (2011). Improving working memory: The effect of combining cognitive activity and anodal transcranial direct current stimulation to the left dorsolateral prefrontal cortex. *Brain Stimulation*, 4, 84-89.
- Baddely, A. D. (1996). The fractionation of working memory, *Proceedings of the National Academy of Sciences U.S.A.* 93(24), 13468-13472.
- Baddeley, A. D. (2000). The episodic buffer: A new component of working memory? *Trends in Cognitive Sciences*, 4, 417-423.
- Baron-Cohen, S., & Cross, P. (1992). Reading the eyes: Evidence for the role of perception in the development of a theory of mind. *Mind and Language*, 7, 182-186.
- Baron-Cohen, S., & Ring, H. (1994). A model of the mindreading system: Neuropsychological and neurobiological perspectives. In C. Lewis & P. Mitchell (Eds.), *Children's early understanding of mind: Origins and development* (pp. 183-207). Hove, UK: Lawrence Erlbaum Associates Ltd.
- Bays, P.M. & Husain, M. (2008). Dynamic Shifts in Limited Working Memory Resources in Human Vision. *Science*, (321), 851-854.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological Review*, 108, 624-652.
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2005). A neural theory of visual attention: Bridging cognition and neurophysiology. *Psychological Review*, 112, 291-328.
- Bundesen, C., Habekost, T., & Kyllingsbæk, S. (2008). A neural theory of visual attention and short-term memory (NTVA). *Neuropsychologia*, 49, 1446-1457.
- Bunting, M. F. & Cowan, N. (2005) Working memory and flexibility in awareness and attention. *Psychological research*, 69, 412-419.
- Chatham, C.H., Herd, S.A., Brant, A.M., Hazy, T.E., Miyake, A., O'Reilly, R. & Friedman, N.P. (2011). From an executive network to executive control: a computational model of the n-back task. *Journal of Cognitive Neuroscience*, 23, 3598-3619.
- Coltheart, M. (1980). Iconic memory and visible persistence. *Perception & psychophysics*, 27(3), 183-228.
- Courtney, S. M. (2004). Attention And Cognitive Control As Emergent Properties Of Information Representation In Working Memory. *Cognitive, Affective, and Behavioral Neuroscience*, 4(4), 501-516.

- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences* 24, 87-185.
- Doherty, M. J. (2006). The Development of Mentalistic Gaze Understanding. *Infant and Child Development*, 15, 179-186.
- Downing, P. E. (2000). Interactions between visual working memory and selective attention. *Psychological Science*, 11(6), 467-473.
- Findlay, J. M., & Gilchrist, I. D. (2003). *Active vision: The psychology of looking and seeing*. New York: Oxford University Press.
- Forster, K. I. (1970). Visual perception of rapidly presented word sequences of varying complexity. *Perception & Psychophysics*, 8, 215-221.
- Fougnie D., Suchow, J. W. & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, 3, pp 1229 doi: 10.1038/ncomms2237.
- Frank, M. J., Loughry, B., & O'Reilly, R. C. (2001). Interactions between frontal cortex and basal ganglia in working memory: A computational model. *Cognitive, Affective, and Behavioral Neuroscience*, 1, 137-160.
- Fung, H. H., Isaacowitz, D. M., Lu, A. Y. & Li, T. (2010). Interdependent Self-Construal Moderates the Age-Related Negativity Reduction Effect in Memory and Visual Attention. *Psychology and Aging*, 25(2), 321-329.
- Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition and Emotion*, 16(3), 355-379.
- Gil, S., & Droit-Volet, S., (2011). How do emotional facial expressions influence our perception of time? In S. Masmoudi, D. Yan Dai & A. Naceur (Eds). *Attention, Representation, and Human Performance: Integration of Cognition, Emotion and Motivation* (1-12) London: Psychology Press, Taylor & Francis.
- Gobet, F. & Simon, H. A. (1996). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, 31, 1-40.
- Hadley, C. B. & Mackay, D. G. (2006). Does Emotion Help or Hinder Immediate Memory? Arousal Versus Priority-Binding Mechanisms. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32(1), 79-88.
- Hickey, C. Chelazzi, L. & Theeuwes, J. (2010b). Reward guides vision when it's your thing: Trait reward-seeking in reward-mediated visual priming. *PLoS One*, 5(11), 1-5.
- Hoffman, J. E. (1998). Visual attention and eye movements. In H. Pashler (Ed.), *Attention* (pp. 119–153). Hove, UK: Psychology Press.

- Isaacowitz, D. M., Toner, K. & Neupert, S. D. (2009). Use of Gaze for Real-Time Mood Regulation: Effects of Age and Attentional Functioning. *Psychology and Aging, 24*(4), 989-994.
- Jackson, M. C., Wolf, C., Johnston, S. J., Raymond, J. E., & Linden, D. E. J. (2008). Neural correlates of enhanced visual short-term memory for angry faces: An fMRI study. *PLoS ONE, 3*, e3536.
- James, W. (1950). Principles of psychology (Vols. 1 & 2). New York: Dover. (Original work published 1890).
- Jeneson, A., Wixted, J. T., Hopkins, R. O. & Squire, L. R. (2012). Visual working memory capacity and the medial temporal lobe. *Journal of Neuroscience, 32*, 3584–3589.
- Kampe, K. K. W., Frith, C. D., Dolan, R. J. & Frith, U. (2001). Reward value of attractiveness and gaze - Making eye contact enhances the appeal of a pleasing face, irrespective of gender. *Nature, 413*, 589-589.
- Kuo, B-C, Yeh Y-Y, Chen AJ-W, D'Esposito, M. (2011). Functional connectivity during top-down modulation of visual short-term memory representations. *Neuropsychologia, 49*, 1589-96.
- Li, Z. (2002). A saliency map in primary visual cortex. *Trends in Cognitive Sciences, 6*(1), 9-16.
- Mary C. Potter (2010) Conceptual short term memory. *Scholarpedia, 5*(2):3334
- Montejo, C. A. & Courtney, S. M. (2008). Differential Neural Activation for Updating Rule versus Stimulus Information in Working Memory. *Neuron, 59*, 173-182.
- Morey, C. C. & Cowan, N. (2005). When Do Visual and Verbal Memories Conflict? The Importance of Working-Memory Load and Retrieval. *Journal of Experimental Psychology: Learning, Memory, and Cognition 31*(4), 703-713.
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In W. G. Chase (Ed.), *Visual information processing* (pp. 283-308). New York: Academic Press.
- Newell, A. (1990). Unified theories of cognition. Cambridge, MA: Harvard University Press.
- O'Connor, K. J., & Potter, M. C. (2002). Constrained formation of object representations. *Psychological Science, 13*, 106-111.
- O'Reilly, R. C., Braver, T. S., & Cohen, J. D. (1999). A biologically based computational model of working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 375–411). Cambridge, England: Cambridge University Press.

- Pashler, H., & Shiu, L. P. (1999). Do images involuntarily trigger search? A test of Pillsbury's hypothesis. *Psychonomic Bulletin & Review*, 6, 445–448
- Potter, M. C. (1999). Understanding sentences and scenes: The role of Conceptual Short Term Memory. In V. Coltheart (Ed.), *Fleeting memories: Cognition of brief visual stimuli* (pp.13-46). Cambridge, MA: MIT Press.
- Potter, M.C. (1993). Very short-term conceptual memory. *Memory & Cognition*, 21, 156-161.
- Potter, M. C., Kroll, J. F., Yachzel, B., Carpenter, E. & Sherman, J. (1985). Pictures in sentences: Understanding without words. *Journal of Experimental Psychology: General*, 115, 281-294.
- Reichle, E. D., Pollatsek, A., Fisher, D. L., & Rayner, K. (1998). Toward a model of eye movement control in reading. *Psychological Review*, 105, 125–157.
- Repovš, G. & Baddeley, A. (2006). The multi-component model of working memory: explorations in experimental cognitive psychology, *Neuroscience* 139 (2006), 5-21.
- Roth, J. K. & Courtney, S. M. (2007). Neural system for updating object working memory from different sources: Sensory stimuli or long-term memory. *NeuroImage*, 38, 617-630.
- Rympa, B., & D'Esposito, M. (2000). Isolating the neural mechanisms of age-related changes in human working memory. *Nature Neuroscience*, 3, 509-515.
- Schultheiss, O. C., Pang, J. S., Torges, C. M., Wirth, M. M. & Treynor, W. (2005). Perceived Facial Expressions of Emotion as Motivational Incentives: Evidence From a Differential Implicit Learning Paradigm. *Emotion*, 5(1), 41-54.
- Senju, A., Johnson, M. H. 2009. The eye contact effect: mechanisms and development. *Trends in Cognitive Sciences*, 13(3), 127-134.
- Shapiro, K.L., Arnell, K.M., & Raymond, J.E. (1997). The attentional blink. *Trends in Cognitive Sciences*, 1, 291–296.
- Sligte I. G., Scholte H. S., Lamme V. A. F. (2008). Are There Multiple Visual Short-Term Memory Stores? *PLoS ONE* 3(2): e1699. doi:10.1371/journal.pone.0001699
- Sligte, I. G., Scholte, H. S., and Lamme, V. A. (2009). V4 activity predicts the strength of visual short-term memory representations. *Journal of Neuroscience*, 29, 7432–7438.
- Sligte, I. G., Vandenbroucke, A. R. E., Scholte, H. S., & Lamme, V. A. F. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in psychology*, 1, 175. doi:10.3389/fpsyg.2010.00175

- Wieser, M. J. McTeague, L. M. & Keil, A (2011). Sustained preferential processing of social threat cues - bias without competition? *Journal of Cognitive Neuroscience*, 23, 1973-1986.
- Wieser, M. J. McTeague, L. M. & Keil, A. (2012). Competition effects of threatening faces in social anxiety. *Emotion*, 12, 1050-1060.
- Wicker, B., Peret, D. I., Baron-Cohen, S. & Decety, J. (2003). Being the target of another's emotion: a PET study. *Neuropsychologia* 41, 139-146.
- Wyatte, D., Curran, T. & O'Reilly, R. (2012). The limits of feedforward vision: Recurrent processing promotes robust object recognition when objects are degraded. *Journal of Cognitive Neuroscience*, 24, 2248-2261.
- Yeung, N., Nystrom, L.E., Aronson, J. A. & Cohen, J. D. (2006). Between-task competition and cognitive control in task switching. *Journal of Neuroscience*, 26, 1429-1438.
- Xu, Y. (2008). Representing connected and disconnected shapes in human inferior intra-parietal sulcus. *Neuroimage* 40, 1849-1856.
- Xu, Y. (2007). The role of the superior intra-parietal sulcus in supporting visual short-term memory for multi-feature objects. *Journal of Neuroscience*, 27, 11676-11686.
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short- term memory for objects. *Nature*, 440, 91-95
- Xu, Y., & Chun, M. M. (2007). Visual grouping in human parietal cortex. *Proceedings of the National Academy of Sciences, USA*, 104, 18766-18771. doi:10.1073/pnas.0705618104
- Xu, Y., & Chun, M. M. (2009). Selecting and perceiving multiple visual objects. *Trends in cognitive sciences*, 13(4), 167-174.

APPENDICIES

APPENDIX A: Pre-analysis treatment of eye movement data

Experiments 2 and 6 both used gaze tracking measures. After defining areas of interest (Aols; see individual experiments), two processes were necessary before I could analyse patterns of gaze fixations. First, it was necessary to transform the individual samples (each 'frame' recorded using our sampling rate of 500Hz) into periods of fixation and periods of non-fixation. Second, it was necessary examine each individual trial and correct for artefacts and ambiguities. I discuss these processes in detail below.

Fixation definition

Fixations were defined as any time in which the pupil was detected (i.e. participant not blinking/looking away) but a saccade was not occurring. Saccades were defined according to three threshold criteria; motion, velocity and acceleration. The motion threshold and either or both of the other thresholds had to be exceeded for a saccade to be recognised. The motion threshold defined the minimum total movement in degrees that had to be exceeded for a saccade to be recognised. Higher motion thresholds serve a function of eliminating progressively shorter saccades from data (and result in correspondingly lengthened fixation times, see Box 1 below). The acceleration threshold, measured in degrees per second squared, detected the saccade onset. The velocity threshold, measured in degrees per second, denoted the gaze-movement velocity required to accept a movement as a saccade.

For my experiments, these criteria were set as follows:

Saccade motion threshold = 0.1° (see box 1)

Saccade acceleration threshold = 8000°/sec²

Saccade velocity threshold = 30°/sec

These are the default 'cognitive' settings suggested by the Eye-Link 1000 manufacturer and are designed to be relatively conservative in saccade recognition.

Box 1

The motion threshold essentially has two functions: First, it is a way of preventing very small saccades from being recorded. Second, it can be used to maximise fixation detection. For psychophysical research this is usually set to 0 so that every small saccade is recorded and the full length of each saccade is detected. However, for cognitive research where fixation detection needs to be maximised and saccades are of less interest setting a motion threshold shortens (marginally in our case of 0.1°) saccades and maximises fixation detection - i.e. the first 0.1° movement of the saccade will be assigned to its preceding fixation rather than the saccade. By setting the motion threshold to 0.1° we thus specify that only when an eye movement of 0.1 degrees has been observed are the acceleration and velocity thresholds then allowed to be checked to see if a saccade is underway (thus eliminating exceptionally short saccades), and that the time taken for the first 0.1° movement of any saccade be assigned to the previous fixation.

Fixation data cleaning and drift correction:

Once fixations had been determined and allocated to spatial locations relative to the established areas of interest (Aols), fixation data were cleaned on a trial by trial basis by an experimenter (me) blind to the location of the singleton face according to the processes below. The first process details drift correction, the second the removal of certain artefacts from the data.

Drift correction and treatment of initial fixations:

Drift correction is the process of moving an entire array of fixations on a single trial to correct for divergence between recorded and actual gaze positions

and can be performed online or post hoc (see box 2). Such divergence can be estimated by looking at the recorded gaze position at a time when participants are asked to look at a fixation cross and adjusting the location of this fixation such that it is coincident with the actual (screen) location of the fixation cross. All other fixations in the array can then be adjusted accordingly by an equal number of degrees.

Box 2

In some experimental setups drift correction can be performed online (that is, while a participant is doing the experiment). However, such a process depends on detection of gaze during an interval in which a fixation cross is present. This can be problematic for experimental designs such as ours where experiments comprise large numbers of separate trials as on a small number of trials gaze position may be briefly lost during the fixation cross interval (due to blinking or other artefacts) resulting in a pause in the trial while the program waits for gaze to be detected in order to perform online drift correction. Such pauses would break up the 'flow' of trial presentation and introduce confounds into the experiment. In a pilot study I attempted online drift correction using the emotional face singleton paradigm described in Experiment 2. This revealed that online drift correction did indeed result in a 'disjointed' presentation of trials for certain participants. Accordingly I performed post hoc rather than online drift corrections.

For each of my experiments employing gaze tracking, visual inspection of the maps of fixations across trials revealed that participants adopted two different strategies for the tasks. On the majority of trials participants fixated on the fixation cross as instructed and then on onset of the study array moved their gaze to inspect a number of stimuli. These trials were cleaned using the process below.

The fixation immediately prior to onset of the memory array (i.e. a fixation occurring while a fixation cross was displayed) was moved to lie over the centre of the display (coincident with the fixation cross location) and the other fixations in the trial drift corrected accordingly such that their distances from the central

fixation were maintained. I observed that the fixation immediately prior to onset of the memory array usually continued into the memory array (i.e. disengagement from the fixation point was not immediate.). Accordingly, since this fixation did not represent a deliberate attempt to fixate on one of our study array stimuli I excluded it from analysis of those fixations occurring during the interval of interest (the study array presentation time). An example of this process is shown in Figure A.1 below.

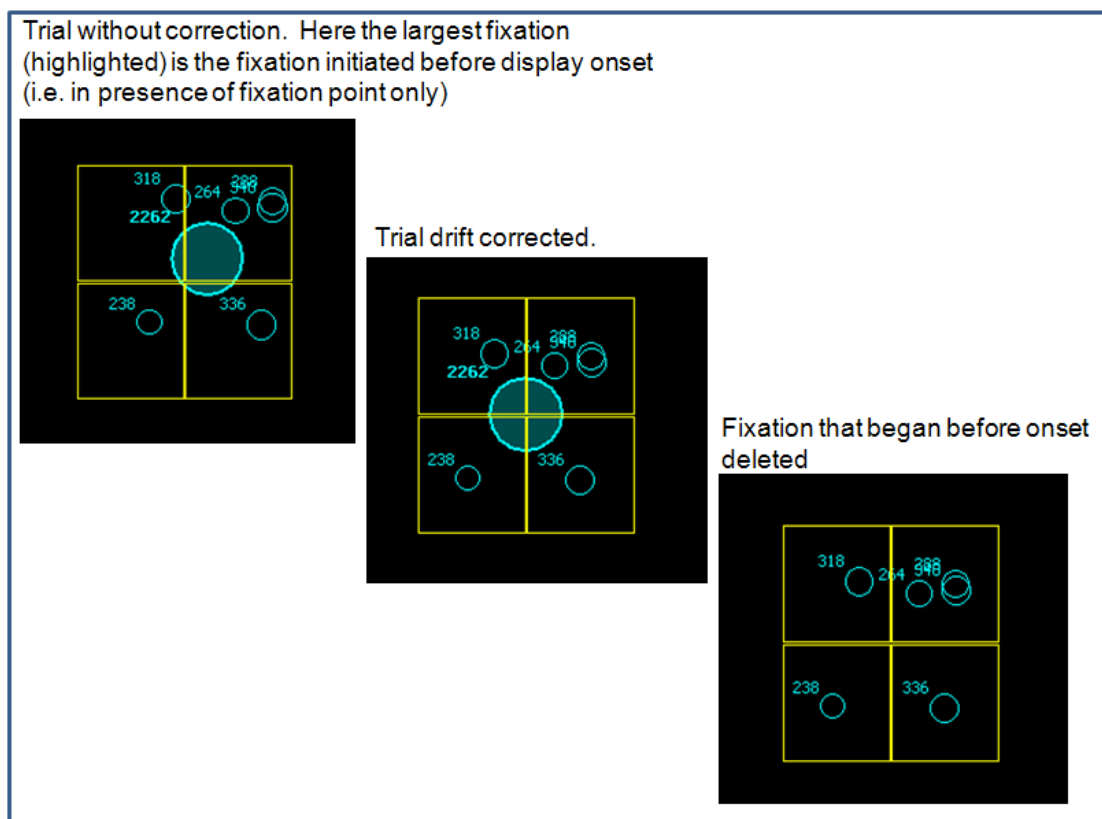


Figure A.1. Example of 'usual' drift correction process

While all participants made use of the strategy described above on the majority of their trials, some participants also switched to an alternative strategy for short periods of time before resuming the common 'fixate on the fixation cross' procedure. These participants moved their eyes to one potential stimulus location prior to memory array onset (i.e. while the fixation cross was still on screen). This

might have been a strategy to reduce total number of eye movements required or 'steal a march' on the array onset by being in the right location to immediately process one of the stimuli upon memorisation array onset rather than needing to make a saccade towards it. It is this temporary change in viewing strategy that necessitated manual inspection and cleaning of data on a trial by trial basis rather than running a cleaning algorithm. When it was obvious that a group of at least five consecutive trials showed evidence of this strategy the fixations on such trials were not drift corrected, and the initial fixation started during the fixation cross period was not deleted but was instead shortened such that its start time coincided with onset of the memory array. An example of such a trial is shown in Figure A.2 below. In Experiment 6, such trials accounted for approximately 3% of the total trials and were limited to 10 participants. In Experiment 2 such trials accounted for approximately 5% of trials and were limited to 12 participants.




Figure A.2. Example or trial from a set of 5 or more trials that did not qualify for drift correction (showed alternative eye movement strategy). Note that the eye position upon onset (shown highlighted here) is in the top right stimulus location

and the other fixations in this trial form a clear 'square pattern' with fixations to each stimulus.

Post drift correction cleaning

After drift correction (when applied), the distribution of fixations to Aols was examined on a trial by trial basis. While the majority of fixations fell in one or other of the Aols, some overlapped multiple Aols and so an element of ambiguity existed for such data points. Treatment of such fixations is described below.

Partially ambiguous fixations:

From time to time a fixation point was 'partially ambiguous' in that it was observed to be overlying two Aols equally. When this occurred and drift correction had been performed, such a data point was marginally 'nudged' back into the Aol it occupied prior to drift correction. This is a slightly more conservative process than failing to adjust such fixations back to their original Aols: It is possible that on any given trial a participant's gaze might shift slightly to one part of the fixation cross prior to onset of the study array. Such a shift should not result in a drift correction, yet is indistinguishable from genuine drift. Thus, our 'conservative nudging' strategy as outlined above should have worked against such erroneous drift corrections and resulted in more accurate data than allowing the partially ambiguous fixation to remain 'as is'. More formally, where the (drift) correction term we applied resulted in ambiguity we chose to reduce the correction term for the ambiguous point slightly.

On trials where no drift correction had been performed such ambiguous points were allowed to remain 'as is' and so be assigned randomly to one of the

two interest areas overlapped. While such data points constituted a source of 'noise' in the data, allowing them to remain was preferable to removing them from the analysis since on average 50% would be assigned to one Aol and 50% to another, but none would be assigned to either of the two Aols not overlapped by the fixation.

Fully ambiguous fixations:

Fixations that were located 'dead centre' (directly over the fixation point) were excluded from analyses as we are not able to identify which (if any) Aol it would be most appropriate to assign them to. Most fixations fulfilling this criterion were made within the first 100 ms after memory array onset (i.e. represented a lingering on the fixation point prior to investigating the stimuli). An example of such exclusions is shown in Figure A.3 below.

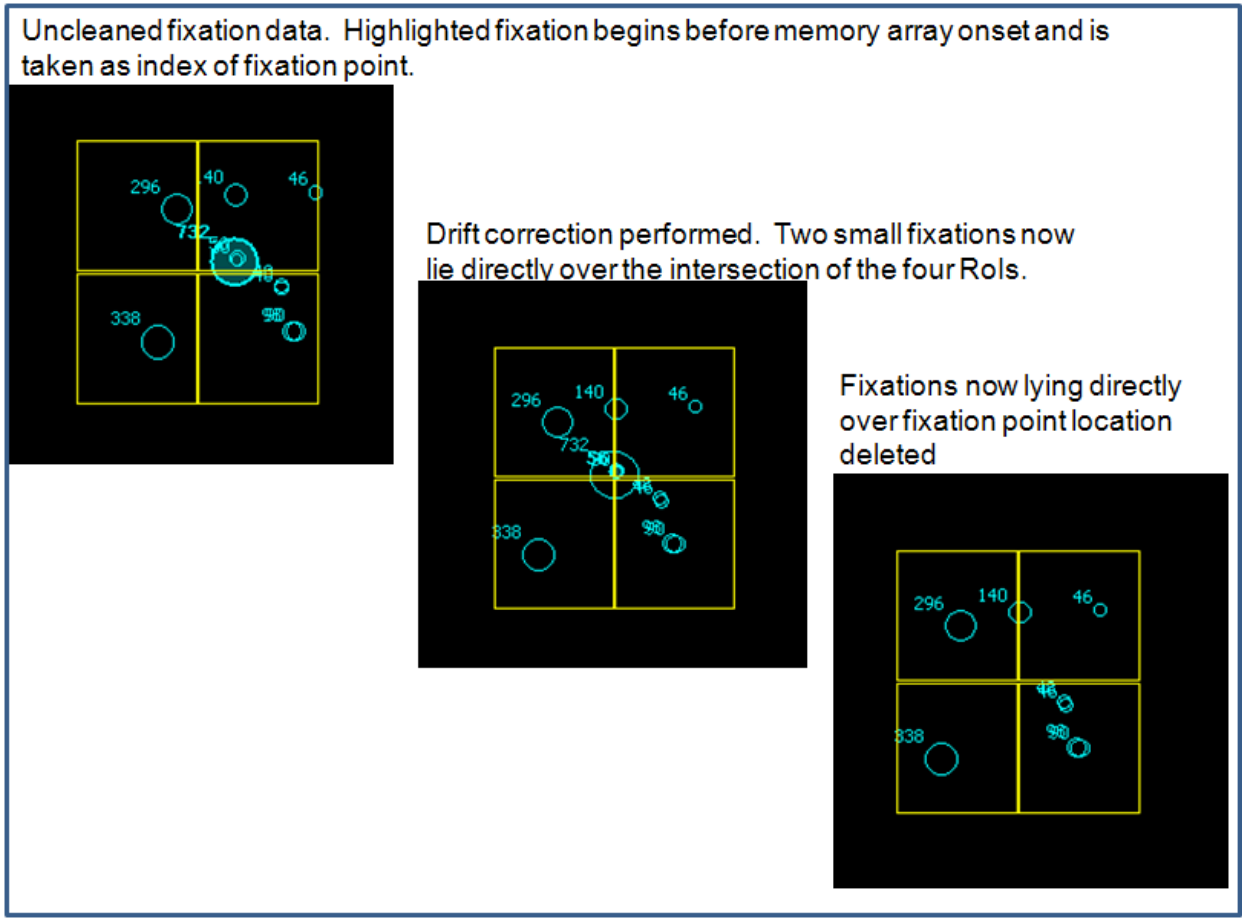


Figure A.3. Example of trial where ‘fully ambiguous’ fixations were removed from analysis.

APPENDIX B: RGB coordinates for coloured squares task stimuli.

Stimulus	Red	Green	Blue
black	0	0	0
blue,	0	0	254
brown	128	64	0
green,	0	254	0
orange,	255	128	64
pink	255	0	255
red	254	0	0
purple,	128	0	255
turquoise	0	255	255
white	254	254	254
yellow	254	254	0
grey background.	192	192	192

Table B.1. RGB coordinates for stimuli used in coloured squares arrays (Experiments 8 and 9). Stimuli were created using the RGB colour mode in Adobe Photoshop. All stimuli were rendered as bitmap images, with a bit depth of 24. Squares had side length of 29 x 29 pixels.

APPENDIX C: Face stimuli used in experiments

Figure C.1 shows the face stimuli used in Experiments 1 and 2, drawn from the Ekman and Friesen (1976) stimulus set. Figure C.2 shows the face stimuli used in Experiments 3, 4 and 5. Figure C.3 shows the scrambled face mask stimuli used in Experiments 3, 4 and 5. Figure C.4 shows the colour tinted face stimuli used in Experiments 3, 4 and 5. Figure C.5 shows the colour tinted face stimuli used in Experiments 6 and 7. Figure C.6 shows the four faces drawn from the Karolinska Face Set (KDEF, Lundqvist, Flykt, & Öhman, 1998) used as value-paired stimuli in Experiment 8. All figures not to scale.



Figure C.1. Faces used in Experiments 1 and 2. Angry, happy and neutral expressions shown.




Figure C.2. Faces used in Experiments 3, 4 and 5. Angry, happy and neutral expressions shown. Faces are arranged in the grid formation shown to participants prior to experimental trials in Experiments 3, 4 and 5.



Figure C.3. Scrambled face mask stimuli used in Experiments 3, 4 and 5.



Figure C.4. Examples of tints used to designate value faces in Experiments 6 and 7.




Figure C.5. The four male faces from the Karolinska Face Set (KDEF, Lundqvist, Flykt, & Öhman, 1998) used as value items in Experiment 8.

References – Appendix C

- Ekman, P., & Friesen, W. (1976). *Pictures of facial affect*. Palo Alto, CA: Consulting Psychological Press.
- Lundqvist, D., Flykt, A. & Öhman, A. (1998). *The Karolinska Directed Emotional Faces*. Psychology section, Department of Clinical Neuroscience, Karolinska Institute, Stockholm, Sweden.

APPENDIX D: Representing error variance graphically

In repeated measures designs there are, broadly speaking, three different functions that may be served by different types of error bars. The first is simply to describe the variance within the conditions by plotting the standard error of the mean (SEM). This allows the reader to see ‘at a glance’ the variability within each condition. There is something to be said for this method since error bars representing SEM are commonly reported in between groups designs and so, by paralleling this for repeated measures, the meaning of the error bars is intuitively understood. However, a weakness of this method is that it is possible that the relative size of the error bars may give intuitively misleading information as regards the statistical differences between conditions. This issue was noted by Loftus and Masson (1994) and relates to the degree of sphericity among conditions. Specifically, the lower the sphericity in the data, the less informative the error bars become as an indication of likelihood of differences between conditions being statistically significant. That SEM error bars (SE of the condition mean) cannot be used to infer significance in within groups (repeated measures) designs is a surprisingly poorly understood issue. Belia et al. (Belia, Fidler, Williams & Cumming, 2005) administered a practical online test designed to mimic graphical representations in published journals. Of 159 respondents comprising authors of published psychology, behavioural neuroscience, and medical journals only 11% identified that such SEM error bars could not be used to infer significance, the rest following instructions to position the mean scores on a graph “[so] that the... means are just significantly different (by conventional t-test, two-tailed, $p < .05$)...”.

The second function to which error bars are put is in paralleling null hypothesis significance testing (NHST). Despite The *Publication Manual of the American Psychological Association* (APA) recommending reporting confidence intervals since they “combine information on location and precision and can often be used to infer significance levels...” (APA, 2001 p. 22), NHST (and inferential statistics generally) remains the predominant analytical technique in reporting psychological research. The idea of showing error bars when NHST is used is that the reader’s primary interest is in whether the conditions represented are statistically different from one another (with reference to a certain pre-determined alpha level) and that error bars should be calculated in such a way so that the degree of overlap between them reflects this. Loftus and Masson (1994) suggest using confidence intervals based on the denominator of the mean square from the ANOVA analysis as error bars. This has the advantage of being unaffected by between subjects variance (since this is partitioned out in the ANOVA analysis) and so provides a measure that is representative of the likelihood of conditions being statistically significantly different from one another. A disadvantage of this method is that each error bar is of equal size, hiding any heterogeneity of variance. An alternative has been proposed by Cousineau (2005). Here, the data are first normalised by subtracting each participant’s mean performance score from their scores in each individual condition and the SEM of these normalised values then used to determine error bars. This provides a representation of variance that is independent of participant effects (between participant variance) but allows error bars to vary in size depending on within subject variance.

Formally, Cousineau’s normalisation process can be described as follows:

For each mean score in each condition for each participant (X_{ij} being the i th

condition by the j th participant), a new normalised score Y is created according to the formula

$$Y_{ij} = X_{ij} - \bar{X}_1 + \bar{X} \quad (\text{D.1})$$

Where \bar{X}_1 is the participant's mean score across conditions and \bar{X} is the group mean (see Cousineau, 2005 for details). This is easily set up with statistics software and indeed Cousineau provides SPSS code for the process.

Alternatively, the same process can be set up in a spreadsheet which has the advantage of allowing inspection of individual participants' contributions to variance in each condition. Note that if this equation is to be used merely to calculate normalised scores from which error bars are to be derived (rather than also used to generate means around which to plot these error bars) then it may be simplified by excluding the \bar{X} constant.

Using the error bars of Loftus and Masson (1994)⁶⁹ or Cousineau (2005), or Morey (2008; a refinement of Cousineau's method), the reader is able to examine the degree of overlap between error bars, and make an intuitive evaluation as to the likelihood of two conditions being significantly different to one another. Indeed, an alternative to these methods is to use inferential confidence intervals (Goldstein & Healy, 1995; Tryon, 2001), whereby whether or not the error bars overlap can be used as a direct test of statistical significance.

The third function error bars can serve is to evaluate the assumptions that underlie hypothesis testing. In between group designs, examination of SEM error bars provides a visual indication as to the violation or otherwise of the assumption of homogeneity of variances. However, neither the within-subjects confidence

⁶⁹ Note that Loftus & Masson specifically designed their confidence interval calculation method to discourage the use of hypothesis testing. Rather, they envisioned the use of descriptive CIs as an alternative way of evaluating data to 'traditional' null hypothesis significance testing.

interval error bars proposed by Loftus and Masson (1994), nor the normalised error bars of Cousineau (2005) and Morey (2008) provide any similar indication as to the violation of the assumption of sphericity for within-subjects designs. To overcome this limitation, Franz and Loftus (2012) have proposed representing pairwise differences as error bars to visually represent the assumption of sphericity (termed circularity in their paper) in within-subjects designs. While this application may be of use in guiding researchers through their own data, its application is limited to designs with relatively few conditions (since the number of pairwise comparisons increases exponentially as the number of conditions increases). Additionally, its utility to a reader is limited to situations where sphericity is both grossly violated and of interest in and of itself since for moderate violations of this assumption statistical corrections based on adjustment of the degrees of freedom can be made (Greenhouse & Geisser, 1959; Huynh & Feldt, 1976), and for major violations non-parametric tests used.

Thus, error bars can be used to describe variance within individual conditions, provide a visual representation of hypothesis testing, or to evaluate sphericity. Each of these uses is appropriate in different situations. In this thesis I have used the former two applications and evaluated sphericity with statistical tests. Specifically, in Figures 7.13, 7.14, and 8.4 where between group comparisons are made I have represented the SEM, whereas in all remaining figures in which error bars are shown I have used the 'normalisation' method proposed by Cousineau to calculate within subject error bars.

References – Appendix D

- American Psychological Association. (2001). *Publication manual of the American Psychological Association* (5th ed.). Washington, DC: Author.
- Belia, S. Fidler, F., Williams, J. & Cumming, G. (2005). Researchers Misunderstand Confidence Intervals and Standard Error Bars. *Psychological Methods*, 10(4), 389-396.
- Cousineau, D. (2005). Confidence intervals in within-subject designs: A simpler solution to Loftus and Masson's method. *Tutorial in Quantitative Methods for Psychology*, 1(1), 42-45.
- Franz, V. H. & Loftus, G. R. (2012). Standard errors and confidence intervals in within subject designs: Generalizing Loftus & Masson (1994) and avoiding biases of alternative accounts. *Psychonomic Bulletin & Review*, 19(3):395-404. doi: 10.3758/s13423-012-0230-1.
- Greenhouse, S. W. & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, 24, 95-112.
- Goldstein, H., & Healy, M. J. R. (1995). The graphical presentation of a collection of means. *Journal of the Royal Statistical Society*, 158A, Part 1, 175-177.
- Huynh, H. & Feldt, L. S. (1976). Estimation of the Box correction for degrees of freedom from sample data in randomized block and split-plot designs. *Journal of Educational Statistics*, 1, 69-82.
- Loftus, G. R. & Masson, M. E. J. (1994). Using confidence intervals in within-subjects designs. *Psychonomic Bulletin & Review*, 1, 476-490.
- Morey, R. D. (2008). Confidence Intervals from Normalized Data: A correction to Cousineau (2005).
- Tryon, W. W. (2001). Evaluating Statistical Difference, Equivalence, and Indeterminacy Using Inferential Confidence Intervals: An Integrated Alternative Method of Conducting Null Hypothesis Statistical Test. *Psychological Methods*, 6(4), 371-386.

APPENDIX E: A cautionary note on comparing performance between single probe and full-array-at-probe change detection paradigms.

In Chapter 4 I discuss some of the issues concerning the interpretation of data from AFC and yes/no change detection tasks. In this appendix I discuss a further point that has relevance to this; specifically, the distinction that must be drawn between the different types of change detection process that may arise from paradigms involving single item probes, and paradigms that utilise a full array probe.

Throughout this thesis I have used d' as my default measure of change detection performance for reasons covered in Chapter 4. However it would have been possible to report measures of WM *capacity* rather than sensitivity to change by using one of the various K statistics described in Chapter 4. Recently, a distinction has been drawn between when it is appropriate to use Pashler's K (Pashler, 1998) as a measure of capacity and when Cowan's K (Cowan, 2001) should be used. Specifically, Rouder et al. (Rouder, Morey, Morey & Cowan, 2011) note that while Cowan's K is appropriate for single probe designs, Pashler's K should be used for full array displays. Briefly, the justification for this is as follows: When a single item is presented at probe it may be compared to the single item in the memory array that occupied the same (retinotopically mapped) spatial location. If the pertinent item from the memory array has been successfully encoded into WM (and assuming a perfect comparison and report process), then a correct present/absent judgement will be made. If the pertinent item is not in memory then the participant must guess and the chance of making a correct response (assuming 50% of trials are of the 'present' type and the remainder are

'absent' type) is 50%. By contrast, when a full array of items is presented for comparison at probe and the participant is required to make a same/different judgement, if all items in the memory array have been correctly encoded then performance will be at 100%. However, if only a partial set of the items have been encoded into WM then if the participant identifies that one of the encoded items has changed between memorisation and probe performance will be 100%, but if the participant compares the items of the partial set memorised to their respective items in the probe array and fails to detect any difference between them then the participant must now guess as to whether a change has occurred. However, this guess is not made in the absence of information but is informed by the fact that comparison of each of the items from the partial set memorised to their (spatially) corresponding items in the test array has failed to identify a change. In such instances this failure to detect a change based on partial information constitutes evidence that the trial in question is likely to be a no-change trial, and this information can therefore be used to inform the guess made. Pashler's K takes account of the informed guessing process assumed to occur when a full array probe is presented (for mathematical proof see Rouder et al., 2011).⁷⁰

⁷⁰ This situation is complicated by the fact that two types of single probe yes/no design exist; probes may either be presented at the retinotopic location of the relevant memorisation array item tested (in which case Conan's K is an appropriate statistic) or may be presented centrally, encouraging a process of comparison to each of the array items successfully encoded into WM in a similar way to that engendered by full array probes. Thus, when probes are presented centrally and only a partial set of the memorisation array items are successfully encoded into WM, guesses on trials in which no change has been detected will be informed in a similar way to the informed guesses made in the case of the full array at probe paradigm. A further complication arises in the case of presentation of complex objects for memorisation where it may be possible to make an informed guess as to presence/absence of such an item in the memorisation array based on partial information if such a complex object is encoded with lower than optimal resolution (that is, when some but not all of the information present in the stimulus is encoded as part of the representation). The ability to make an informed guess in this instance would require good meta-memory for the precision with which object representations are held, and indeed it has been demonstrated that, at least under certain conditions, participants do have a very good meta-memory for this type of information (Fougnie, Suchow & Alvarez, 2012).

This finding has implications for the way in which other summary measures of WM performance should be calculated with regard to the two different change detection paradigms. Unfortunately, while suitable corrections to the d' formula exist to account for the different information present at probe in yes/no vs. AFC tasks (see Chapter 4), no such corrections have yet been established to account for the difference between informed and uninformed guessing in single vs. full array probe tasks respectively. It is beyond the scope of this thesis to determine new formulations for d' in these instances. Rather, I merely acknowledge here the difficulty that arises from this current issue, and in particular note that extreme caution must be applied when comparing change detection performance from tasks using paradigms that differ in their probe type; e.g. in the case of comparing the experiments of Luck and Vogel (1997) and Wheeler and Treiman (2002) as discussed in Chapter 1.⁷¹

⁷¹ In fact, both these studies reported percent correct as their summary measure. However, percent correct is also affected by the issue discussed here, as well as being affected by response biases.

References – Appendix E

- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioural and Brain Sciences* 24, 87-185.
- Fougnie D., Suchow, J. W. & Alvarez, G. A. (2012). Variability in the quality of visual working memory. *Nature Communications*, 3, pp 1229 doi: 10.1038/ncomms2237.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279-281
- Pashler, H. (1998). Familiarity and visual change detection. *Perception & Psychophysics*, 44, 369-378.
- Rounder, N. J., Morey, R. D., Morey, C. C. & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*, 18(2), 324–330.
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131, 48-64.

Appendix F

The following tables give hit and false alarm rates as well as decision criteria 'c' (propensity to say change when unsure; see Chapter 4) for each experimental condition reported herein.

Experiment 1

Condition	D'	Hits	FAs	c
Uniform Angry	1.65	.74	.23	-.05
Uniform Happy	1.14	.57	.19	-.38
Uniform Neutral	1.14	.63	.26	-.18
Angry Singleton, Angry Probe	1.63	.78	.27	.08
Angry Singleton, Neutral Probe	1.22	.73	.31	.07
Happy Singleton, Happy Probe	1.31	.71	.29	.01
Happy Singleton, Neutral Probe	1.27	.72	.29	.01

Experiment 2

Condition	D'	Hits	FAs	c
Angry Singleton, Angry Probe	2.09	.90	.24	.23
Angry Singleton, Neutral Probe	1.04	.69	.32	.00
Happy Singleton, Happy Probe	1.64	.79	.26	.10
Happy Singleton, Neutral Probe	1.13	.70	.31	.03

Experiment 3

Condition	D'	Hits	FAs	c
Angry face, onset delay 17 ms	1.55	.76	.24	-.01
Angry face, onset delay 117 ms	1.91	.76	.15	-.16
Angry face, onset delay 317 ms	1.82	.77	.18	-.09
Angry face, onset delay 500 ms	1.80	.80	.21	.02
Angry face, onset delay 700 ms	1.76	.75	.18	-.12
Happy face, onset delay 17 ms	1.70	.81	.24	.07
Happy face, onset delay 117 ms	1.68	.77	.23	.00
Happy face, onset delay 317 ms	1.59	.73	.21	-.09
Happy face, onset delay 500 ms	1.67	.76	.22	-.03
Happy face, onset delay 700 ms	1.51	.76	.25	.01

Experiment 4

Condition	D'	Hits	FAs	c
Angry face, onset delay 17 ms	1.29	.75	.32	0.1
Angry face, onset delay 33 ms	1.40	.75	.28	.05
Angry face, onset delay 67 ms	1.43	.75	.27	.05
Angry face, onset delay 117 ms	1.50	.78	.27	.07
Angry face, onset delay 317 ms	1.47	.75	.26	.02
Angry face, no mask	1.54	.75	.24	.00
Happy face, onset delay 17 ms	1.31	.74	.29	.05
Happy face, onset delay 33 ms	1.42	.76	.29	.10
Happy face, onset delay 67 ms	1.39	.75	.28	.05
Happy face, onset delay 117 ms	1.37	.76	.29	.07
Happy face, onset delay 317 ms	1.40	.75	.30	.09
Happy face, no mask	1.42	.75	.28	.07

Experiment 5

Condition	D'	Hits	FAs	c
Blocked block, angry face, onset delay 17 ms	2.06	.82	.16	-.04
Blocked block, happy face, onset delay 17 ms	2.22	.88	.17	.09
Randomised block, angry face, onset delay 17 ms	2.10	.85	.19	.05
Randomised block, happy face, onset delay 17 ms	2.11	.86	.19	.06
Blocked block, angry face, onset delay 117 ms	2.09	.81	.14	-.07
Blocked block, happy face, onset delay 117 ms	2.02	.83	.18	.03
Randomised block, angry face, onset delay 117 ms	2.23	.84	.15	-.01
Randomised block, happy face, onset delay 117 ms	2.22	.87	.16	.06
Blocked block, angry face, no mask	2.02	.80	.15	-.08
Blocked block, happy face, no mask	1.89	.79	.18	.06
Randomised block, angry face, no mask	2.29	.86	.14	-.01
Randomised block, happy face, no mask	2.19	.87	.17	.07

Experiment 6

Condition	D'	Hits	FAs	c
Neutral singleton, singleton probe	1.70	.82	.28	.20
Loss singleton, singleton probe	1.74	.87	.35	.46
Win singleton, singleton probe	2.04	.90	.33	.52
Neutral singleton, non-singleton probe	1.36	.79	.34	.21
Loss singleton, non-singleton probe	1.24	.80	.40	.32
Win singleton, non-singleton probe	1.43	.79	.33	.20

Experiment 7

Condition	D'	Hits	FAs	c
Neutral faces, colour probe	1.11	.59	.23	-.30
Loss faces, colour probe	1.11	.58	.21	-.34
Win faces, colour probe	1.03	.58	.23	-.30
Neutral faces, grayscale probe	1.17	.63	.23	-.22
Loss faces, grayscale probe	1.04	.62	.26	-.18
Win faces, grayscale probe	1.06	.59	.23	-.27

Experiment 8

Condition (coloured square maintenance task)	D'	Hits	FAs	c
Gain group: set size 2 overall	3.27	.94	.04	-.04
Gain group: set size 4 overall	2.45	.82	.08	-.21
Gain group: set size 6 overall	1.74	.68	.13	-.36
Loss group: set size 2 overall	3.25	.94	.04	-.04
Loss group: set size 4 overall	2.54	.85	.10	-.10
Loss group: set size 6 overall	1.63	.68	.16	-.27

Experiment 9

Condition	D'	Hits	FAs	c
High win available overall	2.04	.76	.11	-.30
Low win available overall	1.84	.74	.14	-.22
Big win then... 'HI' available reward	1.97	.77	.09	-.26
Big win then... 'LO' available reward	1.63	.73	.15	-.20
Small win then... 'HI' available reward	1.75	.76	.14	-.19
Small win then... 'LO' available reward	1.88	.77	.13	-.17
Low WM capacity group, Big win then... 'HI' available reward	1.79	.75	.14	-.19
Low WM capacity group, Big win then... 'LO' available reward	1.19	.67	.25	-.10
High WM capacity group, Big win then... 'HI' available reward	2.09	.78	.04	-.33
High WM capacity group, Big win then... 'LO' available reward	2.06	.78	.04	-.32