

University of Southampton Research Repository ePrints Soton

Copyright © and Moral Rights for this thesis are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holders.

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given e.g.

AUTHOR (year of submission) "Full thesis title", University of Southampton, name of the University School or Department, PhD Thesis, pagination

UNIVERSITY OF SOUTHAMPTON

FACULTY OF SOCIAL, HUMAN AND MATHEMATICAL SCIENCES

Psychology

Processing Affective Images in the Absence of Visual Awareness

Nicholas Andrew Alexander Hedger

Thesis for the degree of Doctor of Philosophy

July 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF SOCIAL AND HUMAN SCIENCES

Psychology

Thesis for the degree of Doctor of Philosophy

Processing Affective Images in the Absence of Visual Awareness

Nicholas Andrew Alexander Hedger

Given capacity limits, the visual system must prioritize the processing of sensory inputs that are most critical to successful interactions with the environment. Neurocognitive theories suggest that humans have evolved mechanisms that operate without awareness that selectively prioritize threatening stimuli in subsequent allocation of processing resources and access to awareness. Evidence for this ‘standard hypothesis’ comes from paradigms that dissociate visual input from awareness. This thesis combines a narrative review, a meta-analytic review and three empirical studies to examine the extent to which emotionally salient stimuli are prioritized in the absence of awareness.

A general introduction and review of the literature is provided in *Chapter 1*. The meta analysis of previous literature (*Chapter 2*) reveals that evidence for an unconscious processing bias for threat is undermined by insufficiently rigorous awareness measures and inadequate control of low-level confounds. *Chapter 3* reveals that autonomic arousal and attentional orienting to visual threats are eliminated under conditions where observers are objectively unaware of stimuli. *Chapter 4* reveals that prioritized processing of fearful faces is parsimoniously explained by effective contrast: the relationship between their Fourier spectrum and the contrast sensitivity function. Importantly, this explanation does not require or involve unconscious processing mechanisms that are sensitive to threat. *Chapter 5* reveals that prioritized processing of emotional face stimuli is restricted to conditions of awareness, and may be parsimoniously explained by simple low-level variability between emotional and neutral face stimuli.

Previous and present findings and analyses are considered together in the discussion (*Chapter 6*). It is concluded that evidence for emotion-sensitive visual processing that operates without awareness is weak and that uncritical acceptance of the standard hypothesis is premature.

Contents

ABSTRACT	i
Contents	iii
List of Figures	ix
List of Tables	xi
Declaration of Authorship	xiii
Acknowledgements	xv
Chapter 1.....	1
Literature Review: Emotion, Vision and Awareness	1
1.1. General Introduction.....	1
1.1.1. The “standard hypothesis” of emotional prioritization in human vision.	1
1.1.2. Emotion and awareness: Clarification of terms.....	5
1.1.3. Studying the relationship between emotion and awareness.....	6
1.1.4. Visual awareness: measurement and theory.....	7
1.2. Findings from Backward Masking.....	10
1.2.1. Physiological studies.....	10
1.2.2. Visual probe studies.....	13
1.2.3. Neuroimaging studies.....	17
1.3. Limitations of Backward Masking.....	19
1.3.1. Individual differences in masking susceptibility.....	19
1.3.2. Ecological validity.....	20
1.3.3. Variability of presentation methods.....	20
1.4. Findings from Binocular Rivalry.....	21
1.4.1 Classic binocular rivalry and initial percept paradigms.....	23
1.4.2. Neuroimaging findings.....	26
1.5. Limitations of Binocular Rivalry.....	26
1.5.1. Ecological validity.....	26
1.5.2. Perceptual ambiguity and partial awareness.....	27
1.5.3. Temporal uncertainty.....	28
1.6. Findings from Continuous Flash Suppression.....	28
1.6.1. The breaking continuous flash suppression (bcFS) paradigm.....	29
1.6.2. Physiological studies.....	31
1.6.3. Neuroimaging studies.....	31
1.7. Limitations of Continuous Flash Suppression.....	32
1.7.1. The influence of conscious processes are hard to control.....	32
1.7.2 Low level confounds.....	33
1.8. Other findings: Visual Crowding and Patients with Affective Blindsight.....	34

1.8.1. Visual crowding.	34
1.8.2. Patients with affective blindsight.	36
1.9. Summary and Structure of Thesis.	37
Chapter 2	41
Are Visual Threats Prioritised Without Awareness?: A Critical Review and Meta Analysis Involving 3	
Behavioural Paradigms and 2696 Observers.	41
2.a. Abstract.	41
2.b. Publication note.....	41
2.c. Contributions.	42
2.1. Background.	42
2.2. The Present Review: Justification and Scope.	43
2.3. Definitions.	45
2.3.1. What is a threatening stimulus?	45
2.3.2. What is ‘awareness’ and how is it manipulated and measured?	46
2.4. Experimental Paradigms.	47
2.4.1. The masked visual probe paradigm.....	48
2.4.2. Binocular rivalry.	50
2.4.3. Breaking continuous flash suppression (bCFS) Paradigm.	53
2.5. Meta Analyses: Inclusion and Coding Decisions.	55
2.5.1. Inclusion criteria.....	55
2.5.2. Other coding and inclusion decisions.....	55
2.5.3. General search and coding strategies.	56
2.6. Meta Analysis: Methods.	57
2.6.1. Effect size metric.....	57
2.6.2. Standardizers for d.	57
2.6.3. Model and analysis decisions.....	58
2.6.4. Handling dependency among effect sizes.	58
2.7. Meta Analysis: Results.....	59
2.7.1. The MVP paradigm	59
2.7.2. Summary of MVP findings.	70
2.7.3. Binocular Rivalry	72
2.7.4. Summary of BR findings.....	79
2.7.5. Breaking continuous flash suppression	81
2.7.6. Summary of bCFS findings.....	89
2.8. Discussion.	90
2.8.1. Summary of outcomes.....	90
2.8.2. Differences between paradigms.	90
2.8.3. Which threat stimuli receive prioritized processing?	91
2.8.4. Low level confounds.	92
2.8.5. Assessment of awareness and response criteria.	93

2.8.6. Threat-related biases and anxiety.	95
2.8.7. Future directions.	96
2.8.8. Summary and Conclusions.	101
Chapter 3.....	103
Autonomic Arousal and Attentional Orienting to Visual Threat are Predicted by Awareness.	103
3.a. Abstract.	103
3.b. Publication Note.	103
3.c. Contributions.	104
3.1. Introduction.	104
3.2. Method.	107
3.3. Results.	110
3.4. Discussion.	114
Chapter 4.....	119
Fearful Faces Have a Sensory Advantage in the Competition for Awareness	119
4.a. Abstract.	119
4.b. Publication note.	119
4.c. Contributions.	119
4.1. Introduction.	120
4.2. Image Analyses.....	122
4.2.1. Effective contrast.	122
4.2.2. Stability across viewing distances.	123
4.3. Behavioural Experiment 1: Access to Awareness from CFS.	126
4.3.1. Methods.	126
4.3.2. Results.....	128
4.4 Behavioural Experiment 2: Access to Awareness from Visual Masking.....	129
4.4.1. Methods.	130
4.4.2. Results.....	131
4.5. Local Image Analyses.	134
4.6. Discussion.....	135
Chapter 5.....	139
Facial Emotion and Implicit Attention	139
5.a. Abstract.	139
5.b. Contributions.	139
5.1. Introduction.	140
5.2. Method.....	144
5.2.1. Participants.....	144
5.2.2. Stimuli.....	144
5.2.3. Questionnaire measures.	145
5.2.4. Procedure.	145

5.2.5. Presentation conditions.	145
5.2.6. Stimulus pairing conditions.....	147
5.2.7. Noise trials.....	148
5.2.8. Summary.	149
5.3. Results.	149
5.3.1. Sample characteristics.	149
5.3.2. Assessing observers' awareness of the face stimuli.	150
5.3.4. Visual probe data.....	150
5.4. Discussion.	154
5.4.1. Attentional capture by emotionally salient stimuli.	155
5.4.2. Attentional preference for 'face like' stimuli.	155
5.4.3. Implications for paradigms used to manipulate awareness.	157
5.4.4. Implications for assessment of awareness.....	157
5.4.5. Implications for control of low level confounds.	158
5.4.5. Conclusion.....	158
Chapter 6.....	161
Discussion.....	161
6.1. Motivation for Thesis.	161
6.2. Key Findings, Implications and Future Research.	164
6.2.1. Inadequate awareness measures in previous work.	164
6.2.2. No evidence for adaptive responses to emotionally salient stimuli when stringent, objective awareness measures are employed.	165
6.2.3. The processing bias for fearful faces is best explained by low-level properties.....	167
6.2.4. Implications for paradigms used to manipulate awareness.....	168
6.2.5. Implications for anxiety and other disorders.....	169
6.2.6. Implications for neuroimaging.....	171
6.3. Concluding Remarks.....	172
Chapter 7.....	173
Appendices.....	173
Appendix A: Chapter 2. Summary of Excluded Paradigms.	173
Appendix B: Chapter 2. Search Terms and Omissions.....	174
MVP paradigm.	174
BR paradigm	175
bCFS paradigm.....	176
Appendix C: Chapter 2. Decisions Regarding Standardisers for d	177
The MVP paradigm.....	177
The BR paradigm	177
The bCFS paradigm	177
Appendix D: Chapter 2. Decisions Regarding Missing Information.....	178

Appendix E. Chapter 2. Methods for Imputing Missing Moderator Values.	178
Appendix F. Chapter 2. Explaining Sources of Heterogeneity.	179
Appendix G. Chapter 2. Summary of included effects: The MVP paradigm.	180
Demographic information.	180
Appendix H. Chapter 2. Main Effects and Interactions: The MVP paradigm.	181
Appendix I. Chapter 2. Summary of Included Effects: the BR Paradigm.	184
Demographic information.	184
Appendix J. Chapter 2. Main Effects and Interactions: The BR paradigm.	186
Appendix K. Chapter 2. Summary of Included Effects: the bCFS Paradigm.	188
Demographic information.	188
Appendix I: Main effects and Interactions: The bCFS Paradigm.	190
Chapter 8.	191
References.	191

List of Figures

Figure 1.1. Two processing routes to the human amygdala.....	3
Figure 1.2. Four concepts of visual consciousness.	9
Figure 1.3. Backward masking.....	10
Figure 1.4. The masked visual probe paradigm.	14
Figure 1.5. Binocular rivalry.	22
Figure 1.6. Piecemeal rivalry	25
Figure 1.7. Continuous flash suppression (CFS) and the breaking continuous flash suppression (bCFS) paradigm.....	29
Figure 1.8. Control experiments and stimuli for the bCFS paradigm.	34
Figure 1.9. Visual crowding.	35
Figure 2.1. Number of PubMed citations that include the terms ‘backward masking’, ‘binocular rivalry’ and ‘continuous flash suppression’ in the title and / or abstract as a function of publication date.....	43
Figure 2.2. Schematic of generic trial sequence from a masked visual probe task.....	49
Figure 2.3. Schematic of typical stimuli, percepts and response options in a binocular rivalry paradigm.....	51
Figure 2.4. Schematic of typical stimuli, percepts and response options in a breaking continuous flash suppression paradigm.	54
Figure 2.5. Outcomes of MVP analyses.....	65
Figure 2.6. Main effects from the MVP paradigm.....	68
Figure 2.7. Predicted effect sizes from the final model, plotted as a function of actual (observed) effect sizes from the MVP paradigm.....	70
Figure 2.8. Outcomes of BR analyses.	75
Figure 2.9. Main effects for the BR paradigm.	77
Figure 2.10. Predicted values from the final model, plotted as a function of actual (observed) effect sizes from the BR paradigm.....	79
Figure 2.11. Outcomes of bCFS analyses.	86
Figure 2.12. Effects from the bCFS paradigm.	88
Figure 3.1. Schematic of trial sequence from the experimental block.	109
Figure 3.2. SCR data from the CFS block, plotted as a function of trial category.	111
Figure 3.3. Data from the attentional cueing task, plotted as a function of stimulus and trial category.....	113
Figure 4.1. Effective contrast analyses and data.	124
Figure 4.2. Experiment 1 data.	129

Figure 4.3. Schematic of trial sequence in the masking task	130
Figure 4.4. Experiment 2 data.....	133
Figure 4.5. Local variations in effective contrast.	135
Figure 5.1. Face stimuli presented in the normal and control configuration.	144
Figure 5.2. Schematic of trial sequences for the three presentation conditions.	147
Figure 5.3. Schematic of the two stimulus pairing conditions.	148
Figure 5.4. Attentional biases in emotion-neutral bias trials.	151
Figure 5.5. Attentional biases in normal – control bias trials.	152
Figure 5.6. Association between overall d' and overall attentional bias score.....	154
Figure A1. PRISMA flow diagram for the MVP search.	174
Figure A2. PRISMA flow diagram for the BR search.....	175
Figure A3. PRISMA flow diagram for the bCFS search.....	176

List of Tables

Table 2.1. Coding of Individual Effects in the MVP Paradigm.....	61
Table 2.2. Summary of Effects Included in the MVP Analyses	62
Table 2.3. Coding System for Individual Effects in the BR Paradigm	73
Table 2.4. Summary of Effects Included in the BR Analyses.....	74
Table 2.5. Coding System for Individual Effects in the bCFS Paradigm	84
Table 2.6. Summary of Effects Included in the bCFS Analyses.....	85
Table 5.1. Summary of Sample Characteristics	149
Table A1. Summary of Excluded Paradigms	173
Table A2. Summary of Included Effects in the MVP Paradigm.....	180
Table A3. Moderators of Threat-related Bias in the MVP Paradigm: Main Effects.	181
Table A4. Moderators of Threat-related Bias in the MVP Paradigm: Two-way Interactions.	182
Table A5. Summary of Included Effects in the BR paradigm	185
Table A6. Moderators of Threat-related Bias in the BR Paradigm: Main Effects.....	186
Table A7. Moderators of Threat-related Bias in the BR Paradigm: Two-way Interactions.	187
Table A8. Summary of Included Effects for the bCFS paradigm.....	189
Table A9. Moderators of Threat-related Bias in the bCFS Paradigm: Main Effects.....	190
Table A10. Moderators of Threat-related Bias in the bCFS Paradigm: Two Way Interactions.....	190

Declaration of Authorship

I, *Nicholas Andrew Alexander Hedger* declare that the thesis entitled *Processing Affective Images in the Absence of Visual Awareness* and the work presented in it are my own and have been generated by me as the result of my own original research. I confirm that:

- This work was done wholly or mainly while in candidature for a research degree at this University;
- Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- Where I have consulted the published work of others, this is always clearly attributed;
- Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- Either none of this work has been published before submission, or parts of this work have been published as:

Chapter 2: Hedger, N. Gray, K.L.H., Garner, M., & Adams, W. J. (2016). Are visual threats prioritized without awareness? A critical review and meta analysis involving 3 behavioral paradigms and 2696 observers. *Psychological Bulletin*.

Chapter 3: Hedger, N., Adams, W. J., & Garner, M. (2015). Autonomic arousal and attentional orienting to visual threat are predicted by awareness. *Journal of Experimental Psychology. Human Perception and Performance*, 41(3), 798–806. <http://doi.org/10.1037/xhp0000051>

Chapter 4: Hedger, N., Adams, W. J., & Garner, M. (2015). Fearful faces have a sensory advantage in the competition for awareness. *Journal of Experimental Psychology. Human Perception and Performance*, 6, 1748-1757.

Signed:

Date:.....

Acknowledgements

A student that has just finished their thesis is indebted in many ways - not just intellectually, but emotionally (and in a few months – probably financially). It's therefore appropriate that this section does justice to more than just an intellectual debt. I would like to thank my parents, my peers and many staff members at Southampton for their support and patience. I have been fortunate enough to make many friends here (students and staff) that I could not have done without. I couldn't do any serious justice to this section without thanking Jen for having looked after me - more than I have ever looked after myself. I would like to thank Professor Wendy Adams, Dr Matthew Garner, Dr Erich Graf and Dr Katie Gray for teaching me what a good scientist should be, and for leading by example. If this work is seen to continue in the tradition of their work, I will consider it a success.

The research in this thesis was funded by a Mayflower scholarship.

Literature Review: Emotion, Vision and Awareness

1.1. General Introduction.

Critical to the success of any organism is its capacity to rapidly detect and respond to significant events in the environment. However, a major constraint on our sensory systems is their limited capacity. Visual awareness is selective - only a subset of sensory signals give rise to a conscious percept. Visual processing of natural scenes involves the ongoing selection from multiple sensory inputs, which are either omitted from, or consigned to conscious perception. Because, at any one moment, only one reduced set of stimuli can achieve this privileged status of “fame in the brain” (Dennett, 1993), sensory inputs must compete for access to our awareness. To enable adaptive functioning, the visual system must selectively prioritize the sensory inputs that are most important to successful interactions with the environment.

1.1.1. The “standard hypothesis” of emotional prioritization in human vision.

A wealth of literature suggests that human vision has evolved special capacities to prioritize the processing of emotionally significant signals (for a review, see Vuilleumier, 2005). The pervasive “standard hypothesis” (Pessoa & Adolphs, 2010) holds that because emotion-laden information has implications for survival and adaptive functioning, evolution has driven adaptations in the form of a functionally ‘specialized’ sub-cortical visual pathway to the amygdala (Ohman & Mineka, 2001). This pathway is thought to operate in a rapid, automatic fashion and evaluate emotionally salient signals without conscious awareness (Ohman, 2005). The reasoning is that since it may take hundreds of milliseconds for visual stimulation to generate a conscious percept (Koch, 2004), engaging emotion sensitive processes independently of awareness would confer rapid responses that promote survival and adaptive functioning.

The theory underpinning the standard hypothesis is hard to attribute to one single author or landmark paper. Instead, it is based on a long history of observations, which are loosely connected

by the basic principle that regions responsible for ‘emotion’ and ‘cognition’ can be separated in the brain. The first explicit articulation of the idea that humans have a rapid, independent sensory pathway for responding to emotional signals can be traced back to as early as 1885 (Lange, 1885). Danish physician Carl Lange suggested that “simple” emotional stimuli can automatically trigger responses in a subcortical “vasomotor” center via a direct pathway, whereas when a “mental process” is involved, cortical sites must be activated before reaching the vasomotor center.

emotions which are due to a simple sense impression; a loud noise, a beautiful color combination etc., the path to the vasomotor center must be quite direct... The matter becomes more complicated when those affectations are involved which are produced not by a simple impression upon some sense organ, but by some ‘mental process’ (Lange, 1885)

More famously, in 1937, James Papez proposed the similar idea that sensory input diverges at the thalamus and projects upstream or downstream into two separate processing streams for “thought” and “feeling” respectively (Papez, 1937). Later, Paul Maclean extended these ideas with the influential idea of a separable “visceral brain” or “limbic system” as the seat of primitive emotional responses such as fear. Maclean also suggested that this comprised of evolutionarily older, sub cortical structures such as the amygdala (Maclean, 1949), which operate unconsciously.

Considered in the light of Freudian psychology, the visceral brain would have many of the attributes of the unconscious id... eludes the grasp of intellect because its animalistic and primitive structure makes it impossible to communicate in verbal terms (Maclean, 1949, p.348)

The amygdala lies at the heart of current thinking about emotion processing and is part of a network of brain structures that mediate sympathetic nervous system responses (e.g. increased heart rate, blood pressure, muscle tension) that are adaptive in threatening situations (Adolphs, 2013). Additionally, the amygdala coordinates spatially directed responses, by modulating attention via efferent projections to sensory cortical areas (Amaral, Behnia, & Kelley, 2003; Padmala & Pessoa, 2008). In 1998, a landmark paper by Joseph LeDoux reviewed a series of fear conditioning studies conducted on rats. These influential observations led to the proposed existence of two independent sensory pathways to the amygdala. In fear conditioning, mundane stimuli become fear-inducing when they occur in temporal proximity with a naturally threatening event. For

instance, if a rat is exposed to an auditory tone that precedes an electric shock, after several of these pairings, it will begin to respond fearfully to the tone in isolation (e.g. by increased heart rate). Using brain lesion techniques, LeDoux, Sakaguchi and Reis (1984) observed that destruction of the rat auditory cortex did not impair autonomic and behavioral fear responses to a tone paired with an electric shock. The authors reasoned that since the rats were de-corticated, emotional responses based on simple stimuli can be mediated by a separate, subcortical pathway that bypasses sensory cortical areas. This idea was subsequently extrapolated to the visual system of humans, where it was suggested that the amygdala receives information from two visual pathways (LeDoux, 2000) (see Figure 1.1).




Figure 1.1| *Two processing routes to the human amygdala*

LeDoux (1996) suggests that there are two alternative routes for the processing of emotionally salient stimuli in human vision.

The primary visual pathway or “high road” projects from the retina via the lateral geniculate nucleus (LGN) to the primary visual cortex (V1). From V1, fibres project to the extrastriate cortex and along the dorsal or ventral pathway. This standard “object recognition” pathway has a progressive, hierarchical processing route to the amygdala (LGN- V1-V2-V4-TEO-TE-AMG).

Le Doux additionally suggests that emotional stimuli may be initially processed via a coarsely tuned subcortical pathway or “low road” that bypasses the visual cortex by projecting via the superior colliculus and pulvinar to the amygdala (SC-PUL-AMG). Unlike the “high road”, the “low road” emphasises a fast, crude “quick and dirty” processing strategy. This pathway is thought to mediate the non-conscious processing of emotionally salient stimuli. However, whether this pathway exists, or is functional in humans remains controversial (Pessoa & Adolphs, 2010).

LGN: lateral geniculate nucleus. V1: primary visual cortex. SC: superior colliculus. PUL: pulvinar.

Figure 1.1. Two processing routes to the human amygdala.

The cortical pathway, or “*high road*” projects from the retina, via the lateral geniculate nucleus of the thalamus to the primary visual cortex (V1). From V1, fibers project via the extrastriate cortex to the amygdala. Neuroimaging studies converge on the notion that this pathway is characterized by a hierarchical processing scheme (Schroeder, Mehta & Givre, 1998; Smith, Singh, Williams & Greenlee, 2001). Early regions in the hierarchy (the primary and secondary visual cortices: V1 and V2) have narrowly tuned receptive fields and code basic stimulus dimensions such as orientation, 2D position and stereoscopic disparity (Livingstone & Hubel, 1988). ‘High level’ regions in the temporal cortex (e.g. fusiform face area) have larger receptive fields and generalize over these dimensions, specializing in objects, faces and more abstract stimulus properties (Haxby et al., 2001; Kanwisher, McDermott, & Chun, 1997). The amygdala receives highly processed input at later stages of this processing hierarchy, from the temporal cortex (Amaral, Behnia, & Kelly, 2003).

In contrast, the subcortical pathway or ‘*low road*’ projects to the amygdala without cortical relay, bypassing the visual cortex and instead projecting via the superior colliculus and pulvinar (LeDoux, 2000; Morris, Ohman, & Dolan, 1999). As this pathway is assumed subcortical, the input to the amygdala would consist mostly of magnocells, which are coarsely tuned to global, low spatial frequency input (De Valois & DeValois, 1980). Magnocells operate with speed at the expense of detail, since they have fast transmission properties and large receptive fields, thereby granting this pathway ‘quick and dirty’ temporal advantages relative to the slower, cortical pathway (LeDoux, 2000; Vuilleumier, Armony, Driver, & Dolan, 2003). This low road is thus thought to mediate rapid, nonconscious processing of emotionally salient stimuli (Garrido, 2012; Tamietto & deGelder, 2010). Despite the existence of analogous visual pathways in lower mammals (Grieve, Acuna, & Cudeiro, 2000), there is no direct anatomical evidence for the existence of such a retino-collicular- pulvinar-amygdala pathway in human vision (though see Tamietto, Pullens, de Gelder, Weiskrantz, & Goebel, 2012 for an interesting diffusion tensor imaging study).

To summarize, the standard hypothesis has two essential components; the first is that the processing of emotionally salient stimuli occurs independently of awareness, the second is that this is mediated by an independent subcortical pathway to the amygdala. However, there are some notable discrepancies with the view that emotional visual stimuli are processed without awareness and the role of an independent pathway therein (Pessoa, 2005; Pessoa & Adolphs, 2010). As a result, this standard hypothesis has been the subject of much dispute.

To investigate this issue, psychologists have developed various techniques to dissociate visual input from awareness, such as backward masking and binocular rivalry (Kim & Blake, 2005). A large and rapidly growing body of evidence suggests that emotional visual stimuli suppressed from awareness by these methods can nonetheless induce changes in neural activity and behavioral performance consistent with emotional arousal (Koudier & Dehaene, 2007; Lin & He, 2009). The present review attempts to critically evaluate these findings in the context of the following empirical question: *to what extent are emotional visual stimuli processed without conscious awareness?*

1.1.2. Emotion and awareness: Clarification of terms.

Although most people have an intuitive grasp of what ‘emotion’ and ‘conscious awareness’ mean, they are concepts that are hard to define in a manner precise enough to make them amenable to scientific experimentation. Both concepts are often vaguely described and have long been a source of confusion, due to the different meanings they can have in different contexts. (Le Doux, 2013; Pessoa, 2008).

Briefly, *emotion* can be viewed as a central state of an organism that mediates between a stimulus predicting reinforcement and a set of responses (Pessoa, 2008). Emotion can be inferred from these responses (Bradley & Lang, 2000), which can be separated into i) the conscious, reportable experience of emotion, or its ‘*feeling*’ and ii) the functional behavioral and physiological concomitants of the *emotion state* (e.g. sweating, screaming, shifting attention, running away). The latter component is the focus of this review.

Conscious awareness is an equally ambiguous concept that similarly has two definitions (Dehaene & Changeux, 2011); it has an *intransitive* use, where it refers to a *general state* (e.g. ‘the coma patient had no conscious awareness of his surroundings’). In its *transitive* use, it refers to the *registration of a specific piece of information* (e.g. ‘I was unaware that a snake was in the picture you just showed me’). Again, the latter meaning is the focus of this review. Some authors further distinguish between “subliminal” “pre-conscious” and “non-conscious” to refer to specific conditions of unawareness (see Dehaene, Changeux, Naccache, Sackur & Sergent, 2006 for a full discussion). For the purposes of this article, I use the terms ‘unaware’ or ‘unconscious’ to subsume these terms. In this review I consider various techniques that allow visual stimuli to be dissociated from a conscious percept, such as backward masking, binocular rivalry and continuous flash suppression. Stimuli presented under these conditions will be referred to as ‘perceptually invisible’.

This terminology is used to indicate that stimuli are physically present, but generate no conscious percept.

1.1.3. Studying the relationship between emotion and awareness.

The dominant experimental logic in the study of emotion and awareness is well modeled by what is referred to as a “dissociation paradigm” (Wiens & Ohman, 2007). The rationale is simply as follows: if observers display no conscious perception of a stimulus and yet it elicits a response (neural, physiological, behavioral) consistent with emotional arousal- this would provide evidence that emotion processing does not require awareness - thus emotional perception and awareness are dissociated (Schmidt & Vorberg, 2006). Conclusive evidence for a dissociation of emotional perception and awareness would have far-reaching implications that extend beyond the standard hypothesis (Box 1). However, firstly, a successful dissociation paradigm requires some measurement of awareness.

Box 1| Why is the study of visual awareness important ?

Understanding the division of labour between conscious and unconscious processes has attracted considerable interdisciplinary effort (Kanai & Tsuchiya, 2012; Koch & Reid, 2012). Intuition tells us that we are not aware of all aspects of visual processing. For instance, we cannot report the “low level” operation of individual sensory receptors in the eye, but these still process sensory information. This is uncontroversial. The interest in unaware perception instead arises from claims that “the meaning” of a visual stimulus can be registered without awareness and that this can influence subsequent perceptual or cognitive operations. This matter is controversial, perhaps because intuition tells us that all important mental events are yoked to our conscious experience and cognitive control.

In practical terms, the study of visual awareness may be crucial for understanding other faculties of vision, such as attention. For instance, models of attention are rooted in the idea that there is some “filtering” of sensory input prior to conscious perception. Understanding the the level of processing that occurs prior to conscious perception is crucial to understanding the level at which information is “selected” in attention, which shapes our perception of the world (Itti & Koch, 2001).

Historically, the study of awareness has always been central to psychological theories of how we experience emotion. The James-Lange theory proposed that physiological changes precede the conscious, reportable experience of emotion, whereas the Cannon-Bard theory states the opposite order of events (Cannon, 1927). Although these theories have since been revised, it has been argued that most modern affective scientists still implicitly embrace aspects of the James-Lange theory (Dagleish, 2004).

A futher application is to clinical disorders, such as anxiety, which are thought to arise from dysfunction in early stages of information processing, via dysfunction in mechanisms that operate beneath conscious awareness and shape subsequent conscious perception (Bishop, 2007). Mapping these mechanisms and understanding how they are implemented in the brain may provide impetus for improved psychological/ pharmacological treatments of these disorders.

1.1.4. Visual awareness: measurement and theory.

The debate about the validity of different measures of awareness has an extended history (Cheesman & Merikle, 1986; Eriksen, 1960; Wiens, 2008). In the main, this centers on a tension between measures that index subjective or objective thresholds. The most lenient form of measurement is to use an observer's subjective report to index whether a stimulus is perceived (e.g. 'did you see a snake?'). Historically, this *subjective* approach is predicated on the idea that only observers themselves have access to inner states and that this is the only reliable source of information about conscious experience (James, 1890). With the development of signal detection theory (SDT: Green & Swets, 1996), came the concern that subjective measures are prone to response bias, such as reluctance to report a signal if it is weak, partial or brief. As such, according to stricter *objective criteria*, observers are only unaware when their forced choice performance in discriminating alternative stimulus states is at chance performance (Macmillan, 2005). In statistical terms 'chance performance' refers to a criterion (e.g. the binomial distribution) that defines an upper limit of performance consistent with zero discrimination ability (i.e. guessing between stimulus states).

For illustration, suppose an observer is given 100 trials where they are presented with two alternative locations (e.g. left or right of a fixation cross). In one of these locations, a stimulus is briefly presented and in the other, nothing is presented. The participant is then asked to indicate which location contains the stimulus on 100 independent trials. The reasoning is that if the observer is genuinely unaware of the stimulus, they should fail to discriminate between these stimulus states and not respond correctly on many more than 50% of the trials (50% is the performance expected by guessing). Hence, according to a binomial test, if they are only able to respond correctly on less than 60 trials, they are termed 'objectively unaware' of the stimulus. However, if they answer correctly on 60 or more trials they are termed 'objectively aware', since there is a low probability that this level of performance could have been obtained by guessing alone. A current view is that both objective and subjective measures have their own conceptual and practical limitations (Box 2) and so both should be used in combination to comprehensively characterize visual awareness (Szcepanowski & Pessoa, 2007; Wiens, 2006).

Box 2| Limitations of subjective and objective criteria for measuring awareness

Subjective Criteria

Prone to individual response biases (Pessoa, 2005).

Exhaustiveness: may not measure all aspects of awareness. Subjects may not report stimuli, but display above-chance discrimination (Szczepanowski & Pessoa, 2007).

Observers are often over-confident about the reliability of their perceptual experiences (Moore & Healy, 2007).

Objective Criteria



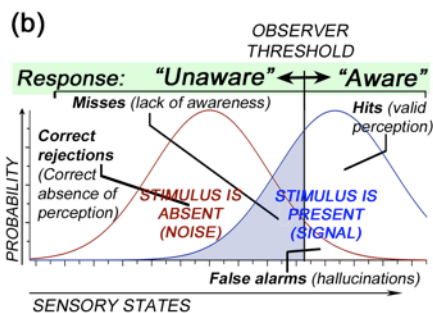

Depend on a mathematically arbitrary criterion (alpha level) for defining “significantly above chance” performance (Wiens, 2005).

Require a lot of trials/ observations to attain the adequate statistical power needed to detect small deviations from null sensitivity (Hannula, Simons, & Cohen, 2005).

May logically preclude unconscious processing, which entails a measurable response that discriminates between the presence and absence of a signal (Wiens & Ohman, 2002).

Exclusiveness: may measure aspects that are not relevant to awareness. Subjects may display above-chance discrimination, but not report stimuli (Wiens, 2005).

Despite this protracted debate concerning the ‘superior’ method for measuring awareness, the conception of awareness they are designed to measure is rarely well defined in experimental psychology. This is important, as fundamentally different models for understanding visual awareness have been proposed (see Figure 1.2), which illustrate the broader space of methodological challenges involved in establishing unconscious perception.

Figure 1.2| *Four concepts of visual consciousness.*

(a) Early, **threshold concepts** of visual consciousness described unconscious perception by analogy of a discrete sensory threshold that is internal to an observer. If a signal is sufficiently weak, then observers indicate that they are unaware of the stimulus. An experimental manipulation (x) may be used to eliminate all conscious processing of the stimulus.

(b) **Signal detection** approaches to visual awareness instead assume an internal continuum of sensory states (Macmillan & Creelman, 2005). Due to neural noise, no signal may induce a large sensory response and a strong signal may induce a small sensory response. In practical terms, reports of awareness are probabilistic statements that depend on the threshold one uses to determine sufficient evidence that a signal was present.

(c) A recent perspective states that an observer can be **phenomenally** aware of (or experience) more information than can be **accessed** (or reported) at any one time. This model derives its motivation from a study by Sperling (1960), wherein observers were presented with a brief grid of letters. Although unable to report more than 4 items, many believed that they had "seen all the letters". Thus, some information was available in phenomenal terms, but not reportable, or as described by Block (2005) "phenomenology overflows access".

(d) **The partial awareness hypothesis (Koudier, de Gardelle, Sackur, & Dupoux, 2010)** states that observers can access high (objects, forms) and low (contrast energy) levels of representation independently. Awareness corresponds to all levels of representation being accessible, whereas unawareness corresponds to no levels of representation being accessible. In the many intermediate stages, an observer is "partially aware". For instance, an observer may notice patches of contrast, or luminance changes without discriminating the meaning of the stimulus. This observer would be described as "*partially aware*".

Figure 1.2. *Four concepts of visual consciousness.*

1.2. Findings from Backward Masking.

Backward masking (Figure 1.3) is an experimental paradigm wherein a briefly presented target stimulus is quickly replaced by the presentation of a salient, co-located ‘mask’ stimulus (typically after ~30 ms). If the presentation parameters of the mask are manipulated appropriately, observers indicate being aware of the mask, but not the preceding target stimulus, thereby ‘masking’ it from conscious perception. Backward masking of visual stimuli appears to disrupt processing at the level of the primary visual cortex (Rolls, Tovee, & Panzeri, 1999). Theories suggest that the brief presentation reduces the effective contrast of the target at early stages of processing, resulting in its invisibility (Kim & Blake, 2005). Others suggest that the weak visual signal produced by masking eliminates “re-entrant feedback” from later stages of processing, which is critical for maintaining a representation in awareness (Dehaene et al., 2006). As masking is relatively simple to administer, it remains a prominent method for manipulating visual awareness.




Figure 1.3| **Backward masking.**

In a backward masking, target stimuli are presented briefly (typically ~30 ms) followed by a salient masking stimulus (Typically, the masking stimulus consists of a “scramble” of the target image, or some other meaningless pattern). If the duration of the target and stimulus onset asynchrony (SOA) between stimulus and mask are sufficiently brief, detection of the target may be difficult, or impossible.

Figure 1.3. Backward masking.

1.2.1. Physiological studies.

Early, highly cited studies employed a dissociation paradigm to measure physiological responses to stimuli rendered perceptually invisible by backward masking. Ohman and Soares (1994) used skin conductance responses (SCRs) as a physiological index of emotional arousal. Sympathetic nervous system activity is associated with elevated moisture levels on the surface of

the skin, which increases its electrical conductance. The resulting changes in physiological arousal can therefore be measured by passing a small electrical current between two electrodes next to the skin, such that when an observer becomes aroused, the skin facilitates the passing of current between the electrodes (Boucsein et al., 2012). In an influential study, Ohman and colleagues (1994) presented a sample of snake and spider phobic observers with threatening and neutral stimuli that were masked after 30 ms. Despite chance performance at determining the content of the masked images (snakes, spiders, mushrooms), observers exhibited greater skin conductance responses to masked pictures of feared animals than to pictures of non-feared or neutral stimuli. Thus, this finding is considered evidence that physiological fear responses do not require conscious recognition of the feared stimulus. Similar effects have since been reported for masked ontogenetically relevant threats, such as images of guns directed towards an observer (Flykt, Esteves, & Ohman, 2007) and less directly threatening stimuli such as words (Van den Hout, De Jong & Kindt, 2000).

In the same year, Esteves, Dimberg and Ohman (1994) used a conditioning paradigm to demonstrate that angry faces paired with an electric shock elicit elevated skin conductance responses when later presented under conditions where they were masked and rendered perceptually invisible. Critically, this finding suggests that recently learned threat associations can be registered without the involvement of conscious processes. This effect appears to be specific to stimuli that are biologically ‘prepared’ by evolution to elicit phobic responses (Seligman, 1971), as conditioned ‘fear-relevant’ stimuli (faces, snakes, spiders) induce the same effect, whereas conditioned ‘fear-irrelevant’ stimuli (flowers, mushrooms) do not (Morris et al., 1998; Ohman & Soares, 1998; Katkin, Wiens, & Ohman, 2001). More recently, a startle paradigm has been used to assess physiological responsivity to masked stimuli. In such experiments, a masked stimulus is presented, before being followed by an aversively loud noise. The general finding is that defensive physiological reactions (eye blink magnitude, skin conductance) are enhanced when the startle is preceded by a masked negative image (Reagh & Knight, 2013; Ruiz Padial & Villa, 2007).

Facial electromyography (EMG) is another measure that has been used to determine physiological responses to backward masked stimuli. When viewing facial expressions, observers spontaneously react with muscular responses that mimic the expression, which can be recorded by electrodes placed on facial muscles (Dimberg, 1982). Research employing this technique shows that backward masked facial expressions elicit spontaneous facial reactions of mimicry in observers, eliciting smiles in response to perceptually invisible happy faces and frowns in response to perceptually invisible angry faces, as determined by concurrent zygomatic and corrugator muscle activity respectively (Dimberg, Thunberg, & Elmehed (2000). Importantly, this finding

suggests that more nuanced, socially complex physiological responses, which distinguish between positive and negative stimuli can be evoked by backward masked stimuli and that such effects are not limited to a basic fear response.

Considering these observations altogether, there is considerable evidence that signals of emotion are processed (as indexed by adaptive physiological responses) independently of their conscious registration, which is therefore consistent with the standard hypothesis outlined above.

There are, however, some notable inconsistencies in this literature. For instance, Peira, Golker, Ohman, Anders and Wiens (2012) found that the heart rate of spider phobics in response to backward masked spider stimuli was strongly modulated by awareness. It was found that differential (feared v non feared) heart rate decreased with the stimulus onset asynchrony (SOA) between target stimulus and mask such that there were no differential responses at SOA's briefer than 30 ms. This is consistent with carefully controlled work demonstrating that various physiological responses (including SCR and EMG) require emotionally salient stimuli to be presented for at least 80 milliseconds to be reliably measurable across observers (a duration at which they are also reliably perceived) and that exposure durations of 25 milliseconds essentially eliminate emotional reactions in any measure (Codispoti, Mazzetti, & Bradley, 2009).


It is also the case that, in many of the aforementioned studies, the assumption that observers are genuinely unaware of masked stimuli was only indirectly supported. To take the study conducted by Ohman and Soares (1994) as an example, the authors demonstrated that a sample of participants were at chance level for detecting stimuli that were masked after 30ms, yet this data was extrapolated to a subsequent, independent sample from which physiological responses were recorded. Thus, even though they may have had different detection thresholds, the assessment of awareness was never actually applied to the observers who exhibited physiological responses to the masked stimuli. Moreover, in the subsequent identification test, despite being unable to recognize the target stimuli, observers were still able to detect "rocks" or "grass" (p.237). In signal detection terms, this indicates discrimination between the presence and absence of a signal and so may be classed as 'aware' according to stricter objective criteria. Such responses are also indicative of partial awareness- access to a low- level of representation (textures), as opposed to complete unawareness (no information is accessible). To take another example, in the post-test awareness check conducted by Ruiz-Padial and Villa (2007), observers were required to indicate which stimulus, in a grid of 12 stimuli, was presented during the masked startle trials. As Grillon and Cornwell (2007) note, this is not an appropriate awareness measure since: "It is quite possible that subjects perceived *a* flower and *a* spider, eliciting the appropriate affective state to modulate startle, even though they could not later identify *the specific* exemplars previously presented" (p.1). In

addition, 6 subjects correctly identified the correct stimulus in the awareness check, but their data was not removed from the analysis of ‘unaware’ trials. These awareness checks are thus highly prone to both overestimating the number of unaware subjects and the extent of unconscious processing.

Another issue that relates specifically to studies employing conditioning paradigms is trial order confounds. Two influential studies; Ohman & Soares (1998) and Katkin et al. (2001) had a restrictive trial order structure, such that an electric shock was absent following stimulus presentation on no more than two consecutive trials. By illustration, if an experimental block began with two trials that were not accompanied by a shock, the next trial would *always* be accompanied by a shock. Thus, a conditional SCR (larger SCR preceding a shock than preceding no shock) may have resulted through an overly restrictive and predictable trial order, rather than subjects being able to evaluate the conditioned threat relevance of the masked stimulus (Wiens, Katkin, & Ohman, 2003). These methodological caveats and aforementioned inconsistencies with earlier work preclude any firm conclusions about the extent to which physiological fear responses are elicited without awareness of the eliciting stimulus.

1.2.2. Visual probe studies.

In addition to physiological responses, emotionally salient stimuli also attract behavioral responses that are spatially directed. The visual probe paradigm provides behavioral evidence for this effect. The trial sequence begins with an emotional stimulus and a neutral stimulus either side a central fixation cross. After a short period, these are removed and a briefly presented probe appears, at the location preceded by either the left or right image. The logic is that if an observer’s detection of a probe is faster (as indexed by response time) when the probe is preceded by an emotionally salient stimulus, then spatial attention has been preferentially drawn to the location of the emotional stimulus (Chica, Martin-Arevalo, Botta, & Lupianez, 2014). The visual probe task therefore provides a measure of the extent to which emotional stimuli modulate attention. Researchers have capitalized on this paradigm by combining it with backward masking to investigate attentional biases towards stimuli presented outside of awareness (Figure 1.4).



$$\text{Emotion bias} = \text{Invalid RT} - \text{Valid RT}$$

Figure 1.4| **The masked visual probe paradigm**

The masked visual probe paradigm combines backward masking with an attentional cuing paradigm.

An emotional and neutral stimulus are presented either side of fixation for a brief duration, followed by a mask. A probe then appears at the location preceded by the emotional stimulus (valid) or neutral stimulus (invalid).

It follows that spatial attention can be assessed by comparing response latencies in valid trials to invalid trials. If response times are shorter in valid trials, this indicates that attention has been preferentially drawn to the location of the masked emotional stimulus.

Figure 1.4. The masked visual probe paradigm.

Evolutionary theories claim that the amygdala mediates orienting responses to both detailed and crude representations of threatening faces (LeDoux, 1998). This has been widely supported by studies demonstrating that angry and fearful faces rendered perceptually invisible by backward masking are prioritised in attention in visual probe tasks (Fox, 2002; Mogg & Bradley, 1999; Mogg & Bradley, 2002) and that this effect is associated with amygdala responses to these stimuli (Carlson, Reinke, & Habib, 2009). Such findings are consistent with the notion that the amygdala processes and prioritises emotional signals non-consciously to enhance their rapid detection. Similar findings have been reported for masked negative word stimuli (Mogg, Bradley, & Hallowell, 1994), International Affective Picture System (IAPS) images (Carlson, Fee, & Reinke, 2009) and stimuli paired with the occurrence of an aversively loud burst of white noise (Mogg, Beaver, & Bradley, 2005).

This masked visual probe literature can be criticised on many grounds. It is important to note that although many of these studies have included a separate block of trials to assess objective awareness, many *have not conducted any check* to verify that masking successfully eliminated awareness (Carlson & Reinke, 2008; Fox, Cahill, & Zougkou, 2010; Sutton & Altarriba, 2011; Wirth & Schultheiss, 2007) and *none* have included checks that monitor detection of the stimuli in the actual visual probe task (i.e. the trials in which behavioural responses to the probes are made). Interestingly, it is also true that only one study applied signal detection techniques to assess

awareness of stimuli (Koster, Verschuere, Burssens, Custers, & Crombez, 2007) and this found that observers exhibited above chance discrimination of stimuli masked after just 14ms.

Notwithstanding these issues with awareness checks, it is also important to understand what drives the prioritisation of threatening stimuli in such studies. Is it the emotional valence of the stimuli, or their low-level stimulus properties, such as contrast and spatial frequency? If the latter is not controlled, then such studies cannot unequivocally provide evidence for the processing of emotion, or threat outside of awareness (see Box 3). For instance, an alternative theory emerging from the behavioural literature is that the modulations of attention underlying the prioritisation of fearful faces may instead be explained by low-level stimulus properties that vary between facial expressions (e.g. luminance, contrast, spatial frequency profile) (Hedger, Adams & Garner, 2015a; Lee, Susskind & Anderson, 2013). Even when global luminance and contrast are equated across face stimuli, local variations in contrast can still influence the salience of a face. For example, the prioritisation of fearful faces in visual probe tasks could be more parsimoniously explained by the patches of high contrast around the eye region that accompany this expression, rather than the extraction of any emotional meaning.

In an attempt to control for these confounds, Carlson and Reinke (2008) used phase scrambled fearful faces as control stimuli. Phase scrambling is a technique that maintains the amplitude spectra (the amount of contrast energy at each spatial scale) of an image, but removes its recognisable structure. It was observed that there was no attentional bias for phase scrambled fear faces, thus suggesting that the “fear advantage” is not explained by unlocalised amplitude information. However, the vast majority of studies reporting this fear advantage have not included such controls, which makes it hard to rule out the contribution of low level confounds.

Box 3| What is meant by “low-level confounds?” Why do they matter?

If an unconsciously presented emotional stimulus is found to receive a processing / detection advantage over a neutral stimulus, this is often interpreted as evidence that 'high level' processing of emotion has occurred unconsciously. However, there may be low-level differences between emotional and neutral stimuli that may result in biases unrelated to the higher level processes meant to be studied. For instance, humans are not equally sensitive to contrast at all spatial frequencies, as evidenced by their contrast sensitivity function (De Valois & De Valois, 1990). Thus, a processing advantage for emotional stimuli may be predicted by simple differences in the spatial distribution of contrast (Figure 1 in this box).




Figure 1| (a) Humans require less contrast to detect gratings at low spatial frequencies. (b) depicts a typical contrast sensitivity function (blue), plotted with the amplitude spectra (contrast energy as a function of spatial frequency) of a fearful (red) and neutral face (green). The fear face has more contrast energy at spatial frequencies corresponding to the peak of our contrast sensitivity.

It could be argued that these low level properties drive a processing advantage precisely because they are associated with emotion (for instance, in fear faces, the patch of high contrast around the eye region could be associated with 'fear' and elicit an emotional response). Fortunately however, the low-level and emotional properties of stimuli are dissociable. Creating a “control” stimulus, by inverting a normal image spatially and reversing its luminance profile maintains the same mean luminance, contrast and amplitude spectra as the normal image, but severely disrupts its recognisable emotional content (Figure 2 in this box). Thus, if effects are driven by emotion sensitive processes, we would predict effects to be present for normal stimuli, but not control stimuli. However, if effects are driven by low level variability, we would expect equivalent effects with both normal and control stimuli.




Figure 2| (a) Examples of normal and control face stimuli (b) Shows the expected results if the recognisable emotional content of the face stimuli was responsible for the prioritised detection of fear (c) Shows the expected pattern if the prioritised detection was explained by low level variability between stimuli.

Why does all this matter? Firstly, if low level confounds were found to explain the prioritisation of emotionally salient images, this would cause us to re-evaluate the notion that 'high level' processing occurs unconsciously. The general shape of the contrast sensitivity function can be explained by a combination of the eyes optics, receptor spacing and the size and quantity of receptive fields in early cortical cells (De Valois & De Valois, 1990). Effects explained by variations in contrast therefore need only imply that a very primitive, early stage of image processing occurs without awareness. Secondly, some studies have used emotive face stimuli to demonstrate that 'threat detection' (as defined by enhanced detectability of threatening facial expressions) is unimpaired in elderly samples (Mather & Knight, 2006) and in samples with Aspergers syndrome (Ashwin, Wheelwright, & Baron-Cohen, 2006). By using stimuli that potentially contain low-level confounds, all these studies may actually be showing is that these patients have a normal contrast sensitivity function. Researchers who do not explicitly characterise the influence of low level confounds in such studies may thus be wasting research efforts that could instead go towards developing stimuli with legitimate diagnostic potential.

It is also possible to question the generality of many these findings; firstly, some studies have reported that an attentional bias for masked fearful faces is only observed when fearful faces are presented in the left visual field (Carlson et al., 2009), or in subjects reporting high levels of anxiety (Mogg & Bradley, 1999), or both (Fox, 2002). Moreover, the largest scale masked visual probe study to date ($N=104$: Fox et al., 2010), found no evidence for a bias towards masked IAPS images and across 3 experiments, Koster et al (2007) found no evidence that masked emotional faces attracted attention relative to neutral faces. It appears that attentional biases to masked threat stimuli may be fragile, depending fundamentally on the stimulus configurations used or sample that is tested. This is somewhat surprising, given that an effective threat detection system should be equipped to process a wide range of stimuli, in a diverse range of viewing conditions.

1.2.3. Neuroimaging studies.

Neuroimaging techniques such as functional magnetic resonance imaging (fMRI) have also provided influential evidence that suggests that emotional responses to masked stimuli transpire independently of conscious awareness. fMRI can measure changes in blood oxygen level dependent (BOLD) activity in different regions of the brain to make inferences about the neural structures involved in processing visual stimuli. This technique has been widely used to record responses from emotionally responsive brain structures when observers are presented with masked stimuli. An early, widely cited paper employed fMRI and masking to render images of happy and fearful faces perceptually invisible and recorded a stronger response to fearful faces from the amygdala (Whalen et al., 1998). Moreover, in the same year, Morris and colleagues employed a conditioning paradigm to pair angry faces with an electric shock. It was found that the recorded amygdala signal was larger to conditioned than unconditioned masked angry faces (Morris et al., 1998). Taken together with later work (Liddell et al., 2005; Whalen et al., 2004), these findings strengthen the view that the amygdala can evaluate the emotional significance of stimuli in the absence of awareness.

More recent research has challenged the validity of these findings. In a seminal study, Pessoa and colleagues used stringent signal detection criteria to assess awareness of masked fearful faces (Pessoa, Sturman, & Ungerleider, 2006). It was found, firstly, that observers varied widely in their ability to detect 33ms targets. Secondly, when those observers who could detect 33 millisecond targets above chance were analysed, the data showed a differential amygdala response

to fearful over neutral faces. Lastly and critically, for the observers who could not detect 33ms stimuli at above chance level, fear and neutral face stimuli did not elicit differential amygdala activation. Thus, emotionally sensitive amygdala responses were dependent on conscious detection of stimuli. Similarly, more recent research employing signal detection protocols by has found that the differential amygdala response to highly arousing IAPS images was eliminated under objectively unaware conditions (Hoffman et al., 2012). The discrepancy between these later and earlier findings may relate to differences in the way that awareness was characterised between studies. Given that Whalen et al (1998) and Morris et al (1998) did not assess objective thresholds of observers on an individual basis, when considered with the data of Pessoa et al., 2006, it is unclear whether the amygdala responses in these studies were genuinely obtained from conditions of unawareness. By extension, in signal detection terms, the responses may instead reflect ‘weakly conscious’ processes that were below the observers internal criterion of subjective report.

In fact, failures to establish awareness at the level of each individual observer appear to be widespread in the neuroimaging literature, with awareness check data extrapolated from previous research (Liddell et al, 2005; Morris et al., 1999), or established from a separate sample before or after the neuroimaging procedure was conducted (Nomura et al., 2004; Phillips et al., 2004). Therefore these studies cannot unequivocally support claims of unconscious perception.

Other issues relate to the choice of stimulus used to ‘mask’ conscious perception itself. Neutral faces are a stimulus that are commonly used to mask perception of fearful faces, which may prevent conscious recognition of the fearful faces, but also result in abrupt luminance changes or flicker. Given that the amygdala is sensitive to a wide range of stimulus properties, including motion (LaBar, Crupain, Voyvodic, & Macarthy, 2003), temporal unpredictability (Herry et al., 2007) and sharpness of contours (Bar & Neta, 2007) responses that are thought to be elicited by emotional valence of the fearful face may instead be elicited by the quick transition from target stimulus to mask. For instance, it has been demonstrated that masked fearful eyes induce stronger amygdala activation than masked happy eyes, but that this can be attributed to interactive effect of a fearful target and neutral mask, rather than the masked fear in itself (Straube, Dietrich, Mothes-Lasch, Mentzel, & Miltner, 2010). Thus, it is possible that the amygdala responses to masked fearful faces may reflect general amygdala sensitivity to changes in the eye region (Adams, Gordon, Baird, Ambady, & Kleck, 2003), such as changes in the size of the scleral field (Demos, Kelley, Ryan, Davis, & Whalen, 2008) rather than any threat-sensitive effect. Future work should aim to eliminate these potential confounds.

Another issue related to the interpretation of fMRI data is that single cell recordings and electrophysiological components recorded from the amygdala are generally inconsistent with rapid, unconscious processing (Box 4). Faced with these difficulties, in the absence of convergent data, it is impossible to draw firm conclusions about the extent to which BOLD responses in the amygdala represent unconscious, emotion processing driven by a rapid subcortical pathway.

Box 4| Rapid unconscious processing?

One of the core components of the standard hypothesis is that unconscious processing of emotional stimuli is *faster* than general-purpose visual processing due to the advantage of a direct subcortical route to the amygdala. The evidence for this “speed” component is very mixed. Since fMRI has poor temporal resolution, single-cell recordings provide the best test of this idea since they can estimate response latencies in the amygdala.

Intracranial studies in humans generally indicate that the earliest single-unit responses to emotional stimuli in the amygdala begin at around 200ms (Krolak-Salmon et al., 2004). Critically, these latencies are comparable to those recorded from the temporal lobe (Mormann et al., 2008). These amygdala latencies are thus consistent with the idea that they are generated by activity transmitted via a cortical pathway (i.e. the “high road”) without the need to invoke a rapid, subcortical pathway.

One interesting single unit study observed that amygdala responses that differentiate between faces and non faces occur at several hundred milliseconds from stimulus onset. Moreover, neurons were particularly selective to whole faces and were even sensitive to deletions of very small components of faces (Ruishauser et al., 2011). The latency and sensitivity of these responses is thus inconsistent with “quick and dirty” processing.

Some electrophysiological studies have indicated that responses in the amygdala can be modulated by the affective content of stimuli in as little as 30ms (Luo et al., 2010), which seems indicative of fast pathway. However, whether the activity of “deep” sub-cortical structures such as the amygdala can be reliably estimated via current electrophysiological techniques is debated (Pessoa, 2010). In addition, these latencies conflict with recordings from the thalamus, which receives *direct* retinal input (i.e. receives information *before* the amygdala), but responds on average around 33ms from stimulus onset (Lamme & Roelfsema, 2000).

1.3. Limitations of Backward Masking.

1.3.1. Individual differences in masking susceptibility.

The validity of the backward masking paradigm has been questioned due to the substantial individual differences in its ability to prevent conscious awareness. For instance, Pessoa, Japee, Sturman and Ungeleider (2005) assessed the detection of fearful faces using stringent signal detection criteria and found that around 60% of observers could reliably detect masked, 33

millisecond targets at above chance level. Similarly, Szczepanowski and Pessoa (2007) found that the majority of observers exhibited above chance performance for detecting fearful faces at durations of just 25 milliseconds. Such findings are important, given that behavioural or neural effects observed in studies employing ~ 30 millisecond masked presentations are still cited as seminal examples of “unconscious” processing (Mogg & Bradley, 1999; Morris et al., 1999; Whalen et al., 1998).

This issue is particularly salient when implicitly assuming a universal threshold for preventing awareness, by conducting an awareness check on an independent sample of participants to those in the main experiment (Carlson et al., 2009), determining awareness at the group level, or drawing on awareness check data from another study (Lidell et al., 2005). In reality, the success of masking is likely to be a function of both the stimulus onset asynchrony between target and mask and the detection threshold of the observer in question. Neglecting the latter element is likely to result in an over-estimation of unconscious processing.

1.3.2. Ecological validity.

An optimal comparison of conscious and non-conscious perception would require that stimuli are identical on every dimension other than their conscious detection. However, by definition, backward masking imposes conditions that render target stimuli spatially and temporally altered from those that are consciously viewed. Backward masked stimuli are degraded, brief signals that are typically confined to laboratory contexts. The rapid succession of perceptually different stimuli in masking may induce a similar effect to when observers make rapid eye movements from one area of a scene to another (Henderson, 2003). However, backward masking in a laboratory context typically involves no more than two stimuli in circumscribed locations, which are presented for a predictable duration. Thus, it is a legitimate concern that masking may be functionally irrelevant to understanding unconscious processing in natural viewing conditions. Moreover, masking involves creating conditions where an observers visual processing is weakened from the outset. In other words, reducing presentation time directly affects the signal to noise ratio. This may mean that unconscious processes may appear fragile, or non-existent, simply because of the weakened signal strength, rather than the genuine absence of unconscious processing.

1.3.3. Variability of presentation methods.

As successful masking procedure relies critically on a fast and reliable presentation method. Cathode ray tube (CRT) monitors operate by progressively activating areas of a phosphor layer via a CRT beamer. When all regions are activated, the beamer returns to its starting position and the process repeats (a refresh cycle). The rate of refresh cycles is termed the 'refresh rate', which typically ranges from 60hz to 160hz. However, more time may be added to a refresh cycle by a higher screen resolution, which may not support the same refresh rate, owing to the increased number of pixels that need to be activated by the CRT beamer. Equally, higher screen resolutions can result in lower image luminance, due to the lower relative time available to fully activate all phosphors. Thus, the exact strength of visual signal presented may vary markedly between studies (Hannula et al., 2005).

Another critical issue is that CRT monitors are also unsuitable for use in fMRI machines, due to interference with the magnetic scanning process. In many cases, this has prompted researchers to use liquid crystal displays (LCD) or thin film transistor (TFT) displays (e.g. Carlson et al 2009; Whalen et al, 1998), which operate via continuously filtering a constant light source. However, research using well-controlled photodiode measurements suggests that LCD or TFT monitors may be unreliable for brief picture presentations, such that images are not presented at full luminance unless they are presented for more than 50 milliseconds, or fail to produce any image at all (Wiens et al., 2004). Notably, the average discrepancy between desired and actual picture presentation for a TFT monitor was ~49 milliseconds. Thus, these monitors may be less suitable for brief presentations than is claimed by manufacturers, or assumed by the experimenters that use them.

As outlined above, there is substantial error associated with brief picture presentation. By extension, this may undermine the efficacy of backward masking as a method for manipulating awareness. Thus, in attempting to produce a reliable or replicable masking protocol, researchers should attempt to validate the presentation method used.

1.4. Findings from Binocular Rivalry.

Under normal viewing conditions, our two eyes usually receive only slightly different views of the world. In such cases, the brain is able to combine these images into a coherent percept via binocular fusion. However, when the two eyes receive entirely different images at corresponding retinal locations, binocular rivalry may occur (Figure 1.5). In such cases, the visual system cannot combine the two images into a coherent percept. Instead of perceiving both images as being

superimposed, perception alternates between the images presented to each eye (Alais & Blake, 2005). At any given time, one stimulus is dominant and visible, whilst the other is suppressed and perceptually invisible before a perceptual switch occurs, with the previously suppressed image becoming dominant and so on. Thus, although the images that are physically specified on each retina remain constant, one image is intermittently prevented from reaching awareness by interocular suppression.

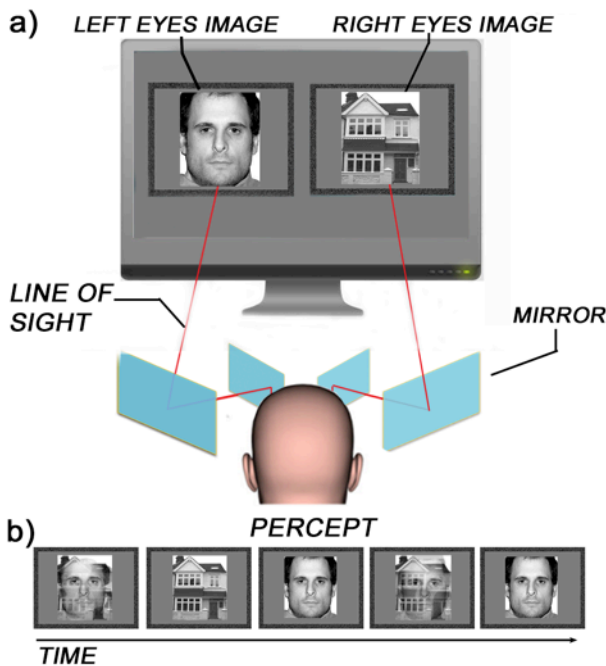


Figure 1.5| *Binocular rivalry*

In binocular rivalry (a) incompatible images are presented to each eye, via a mirror stereoscope, or coloured anaglyph. As opposed to normal binocular fusion, the resulting perceptual experience is an alternation between the images presented to each eye.

Each image undergoes brief, intermittent periods of suppression (perceptual invisibility) and dominance (visibility), despite the retinal image remaining constant (b). Thus, visual input is dissociated from awareness. The extended and invariant retinal stimulation is rather different from masking, which relies on rapidly changing retinal input.

Figure 1.5. *Binocular rivalry.*

At a neural level, this perceptual effect can be attributed to reciprocal inhibition between neural populations representing the two eyes stimuli at distributed stages of the visual processing hierarchy (Blake & Logothetis, 2002). The neural population exerting strongest inhibition achieves access to awareness, thereby suppressing the other from awareness. Neural adaptation of the dominant population then causes perceptual alternations by progressively weakening the inhibition of the suppressed stimulus, resulting in the previously suppressed stimuli reaching dominance and so on (Alais, Cass, O'Shea, & Blake, 2010). A range of stimulus properties have been found to influence the perceptual dominance of a stimulus in rivalry, including both low-level characteristics such as luminance and contrast (Alais & Blake, 2005; Baker & Graf, 2010) and more complex properties such as spatial context and surface organisation (Graf & Adams, 2008)

Collectively, the stimulus properties found to influence dominance of a stimulus in rivalry are referred to as possessing ‘stimulus strength’ (Levelt, 1965).

Although rivalry has been used as a research tool by vision scientists for more than 170 years, a landmark paper by Crick and Koch (1998) stimulated a renewed programme of research into rivalry by promulgating the notion that it offers opportunity to study the neural correlates of consciousness, owing to its capacity to dissociate visual input from awareness (Baker, 2010). Like backward masking, rivalry offers the potential to assess the extent to which unconscious (suppressed) emotional stimuli are processed by the visual system.

1.4.1 Classic binocular rivalry and initial percept paradigms

In a classic binocular rivalry paradigm, two images are presented binocularly at corresponding retinal locations for a fixed trial period. Observers are then required to indicate (via button presses) which stimulus is dominant throughout the course of the trial. It has been proposed that a longer total dominance duration of an emotional stimulus indicates has been processed during suppression and received prioritised access to subsequent awareness (Alpers & Gerdes, 2007; Bannerman et al., 2008).

The earliest study to use this paradigm in the context of emotion processing was conducted by Coren and Russel (1992) who presented observers with competing dichoptic presentations of different facial expressions. It was found that observers reported a longer total dominance duration for extreme positive (happy) or negative (angry) expressions over less valenced expressions (e.g. surprise), suggesting that more emotive stimuli were prioritised in perceptual selection. This early observation is hard to interpret however, since perception was not actually reported ‘on-line’ during the rivalry trial itself, but was recalled on a post hoc basis after the rivalrous presentation had been terminated. Thus, reports of prolonged perception of emotional stimuli may have resulted from response biases induced by memory limitations. Employing a more sensitive method, Alpers and Pauli (2006) instead concurrently measured observers perceptual reports of IAPS stimuli engaged in rivalry. It was observed that emotional images (negative or positive) were reported to be perceived for longer than neutral images. Again, due to lack of appropriate controls, it is hard to determine whether this effect was driven by emotional valence, or whether it reflects the influence of basic low-level differences between stimuli (luminance, contrast) that robustly impact on rivalry dominance.

One way to dissociate low-level and emotional properties of a stimulus is to pair perceptually similar stimuli with different valences. For example, Alpers et al (2005) used simple grating stimuli (horizontally or vertically oriented) paired with an electric shock (CS+), or no electric shock (CS-). It was found that the CS+ grating showed a trend in dominating over the CS-grating. Similarly, Anderson, Siegel, Bliss-Moreau and Barret (2011) used an affective learning paradigm to pair emotionally neutral faces with positive, negative or neutral “gossip”. It was found that faces previously paired with negative information (e.g. “he threw a chair at a classmate”) dominated more than faces associated with neutral, or positive information. Thus there is some preliminary evidence to suggest that when low-level confounds are eliminated, the affective content of stimuli can modulate access to awareness.

One important methodological concern with measuring the total dominance duration of stimuli in rivalry is that this measure is not sensitive enough to isolate the influence of unconscious processes. For instance, on the one hand, prolonged dominance of emotional stimuli could reflect unconscious processes acting on a suppressed emotional stimulus that influence its return to dominance. However, on the other hand, it may also reflect the influence of conscious processes operating on a dominant emotional stimulus that maintain its sustained awareness. One way to reduce such contamination by conscious processes is to instead record only the *initially dominant* percept as the dependent variable. With this measure, since only the first stimulus to reach dominance is reported, it is more suited to studying the initial, bottom-up, unconscious stages of perceptual competition (Carter & Cavanagh, 2007; Ooi & He, 1999).

Another issue is the possibility that response biases were not adequately controlled in many of the above studies. For instance, the large stimulus size in some studies (Alpers & Gerdes, 2007: 9.5 x 8.6 degrees of visual angle (DVA); Bannerman, Milders, De Gelder, & Sahraie, 2008: 9.8 x 13.2 DVA) increases the risk of piecemeal rivalry (periods of a ‘mixed percept’, where exclusive visibility of one stimulus or the other is difficult to report – see Figure 1.6). Piecemeal rivalry occurs because rivalry may transpire independently in separate receptive fields (Kovacs, Papathomas, Yang, & Feher, 1996) meaning that large stimuli, which stimulate more receptive fields, will result in less coherent alternations. Consequently, this may give rise to response biases, such that an emotional stimulus is reported, when in fact elements of both stimuli are visible. The likelihood of piecemeal rivalry can be reduced by using smaller stimuli or presenting stimuli in the peripheral vision, where receptive fields are larger. The prevalence of piecemeal rivalry can also be characterised to some extent by including a ‘mixed percept’ response option. However, regardless of the available response options, an observers’ individual response criteria will still determine what they deem sufficient evidence that a stimulus is ‘mixed’ or ‘dominant’.

Another, somewhat critical concern is that, in some studies, the presentation conditions may not have satisfied a critical condition for rivalry; that the two images are sufficiently different from one another. For instance, several studies have attempted to engage uniformly oriented fearful and neutral faces in rivalry (Alpers & Gerdes, 2007; Amting, Greening & Mitchell, 2010; Bannerman et al., 2008). Given the structural similarity of these images, it is possible that *fusion* may have occurred instead of rivalry. This would lead to observers perceiving a fused percept of fearful and neutral features, the composite of which would appear fearful. In fact, one study indicated that fearful and neutral faces presented in this way failed to instigate any rivalry in any observers (only fearful faces were perceived for the entire trial) and so the data were unusable (Bannerman et al., 2008: Experiment 2B).

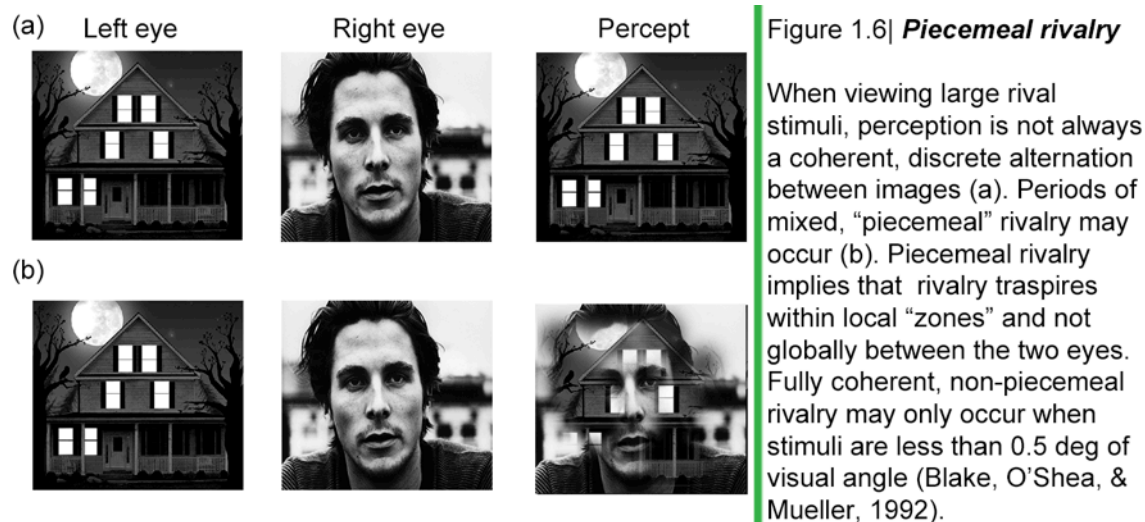


Figure 1.6. Piecemeal rivalry

Using an initial percept task, Gray, Adams and Garner (2009) built on the limitations of previous research by taking a more rigorous approach to the stimulus design. Face stimuli were presented considerably smaller (2.5 DVA) than in previous research to reduce the risk of piecemeal rivalry and were rotated 30 degrees in opposite directions to reduce featural alignment and the risk of fusion. It was found that observers reported fear and happy faces as the first percept more frequently than neutral faces. Gray and colleagues also included a spatially inverted face condition, to provide a control for low-level image properties. The logic of this manipulation is that it reduces

the recognisable emotional content of stimuli but retains low-level stimulus properties. Critically, Gray et al found more frequent perception of fear and happy expressions for both the upright and inverted stimuli, suggesting that the low-level properties of faces may drive the effect of expression.

1.4.2. Neuroimaging findings.

Two influential studies, which were published in the same year, employed rivalry with fMRI to investigate the neural signatures of unconscious face processing. Williams et al (2004) presented face and house stimuli via coloured anaglyphs to induce rivalry. The likelihood that the face stimuli were suppressed over a short presentation was increased by presenting the faces at lower contrast and ensuring the hue of the house image exploited the observers individual biases for perceiving red or green light. It was found that amygdala activity was increased on presentations of fear faces, relative to neutral faces, regardless of whether they were suppressed or dominant. In a similar study, Pasley, Mayes and Schultz (2004) found that suppressed fearful faces elicited significant amygdala activation relative to suppressed images of neutral objects. The main conclusion stemming from these studies is that amygdala responses to emotionally salient stimuli appear to occur independently of awareness. Two details of the awareness checks used in these studies are relevant in this context. Firstly, the awareness check conducted by Williams et al consisted of a “one-back” task, where observers reported any repeated presentations of stimuli (there was one such repetition per epoch). Thus awareness was not monitored on a trial-by-trial basis, meaning that it cannot be verified that faces were suppressed on every trial they were expected to be. Secondly, neither study used an objective criterion for measuring awareness. Thus, the evidence for the suppression of stimuli in these studies is not particularly robust.

1.5. Limitations of Binocular Rivalry.

1.5.1. Ecological validity.

Many papers introduce the concept of binocular rivalry with a sentence to the effect of: ‘when discrepant images are presented to each eye independently, they induce binocular rivalry’. In fact, such statements are very misleading. Such conditions *can* induce binocular rivalry, but in natural viewing conditions, this is rarely the case. Placing a hand ~5cm in front of one eye whilst reading this sentence with the other will not induce rivalry, but clearly this is a case where the two

eyes receive discrepant images. Such selective obstructions of one eye may occur often in natural scenes, but rarely induce rivalry because human eyes can only accommodate to focus an image if it is beyond 10cm away. As a result, an image that selectively obstructs one eye is almost always very blurred, meaning that it has low contrast and little energy in the mid- to high spatial frequencies. The signal with higher stimulus strength is therefore usually the un-obstructed eye, meaning that rivalry rarely occurs (Arnold, 2011). Thus although ‘binocular rivalry stimuli’ (discrepant monocular images) occur in natural environments, the characteristic perceptual ‘switches’ observed in a typical binocular rivalry experiment rarely do. Again, like masking, this raises the question as to whether the dynamics of binocular rivalry are functionally equivalent to how stimuli compete for awareness in normal circumstances.

1.5.2. Perceptual ambiguity and partial awareness.

Though binocular rivalry is seen as a powerful tool for inducing perceptual disappearance, the extent to which stimuli are unambiguously suppressed from awareness by this method still remains an empirical question. Specifically, it should be noted that for larger rival stimuli ‘mixed’ states are possible (Blake, O’shea, & Mueller, 1992), where an observer cannot unambiguously report visibility of one stimulus in particular. Under these conditions, it is difficult to measure an observer’s perception with accuracy, particularly given that under ambiguous viewing conditions, reports of a dominant image are likely to be strongly biased by an observer’s response criterion.

Participants may also retain some residual sensitivity to, or partial awareness of stimuli in the suppression phase. For instance, colours of objects suppressed under rivalry can nonetheless appear as a diffuse ‘cloud’ superimposed on the dominant image (Hong & Blake, 2009) and suppressed, drifting gratings can still give an impression of movement, when only the dominant image is visible (Zabood, Lee, & Blake, 2011). Moreover, when two flickering forms engage in rivalry, they can be temporally integrated into ‘beats’ despite observers only being consciously aware of one form (Carlson & He, 2000). These examples of ‘stimulus fractionation’ are widespread in the rivalry literature and suggest that fusion and rivalry can co-occur, such that some aspects of a stimulus may be suppressed (form) but others may be fused (colour, motion, temporal information). These phenomena are indicative of partial awareness. This, in turn, suggests discretion when equating rivalry dominance and suppression to awareness and unawareness respectively.

1.5.3. Temporal uncertainty.

In rivalry and other bistable phenomena, distributions of periods for which stimuli are dominant often resemble a gamma distribution (Levelt, 1965). This essentially means that short and long percepts are relatively infrequent and the distribution has a right skew. Although this probability distribution can be analyzed statistically, it cannot be predicted precisely, nor can individual dominance durations be predicted purely based on modeling adaptation and inhibition processes. Instead, noise appears to play a vital role in determining the length of a percept (Brascamp, van Ee, Noest, Jacobs, & van den Berg, 2006), meaning that binocular rivalry is essentially a stochastic process (Lehky, 1995). Given the unpredictable nature of perceptual switches, this makes binocular rivalry generally unsuitable for reliably presenting stimuli outside of awareness, particularly for long durations. This issue has been partially surmounted with the advent of new techniques such as flash suppression and continuous flash suppression.

1.6. Findings from Continuous Flash Suppression.

Continuous flash suppression (CFS: Tsuchiya & Koch, 2005) is a variant of binocular rivalry in which a stimulus presented to one eye is suppressed from awareness by a competing dynamic noise pattern presented to the other eye. Suppression during CFS is more potent than during traditional BR (as defined by contrast detection thresholds; Tsuchiya, Koch, Gilroy, & Blake, 2006). Temporally, the periods of suppression induced by CFS can last about 10 times longer than suppression induced by traditional BR (Tsuchiya & Koch, 2005). Another attractive property of CFS is that perceptual suppression of a target stimulus can reliably be induced from the onset of a trial. Thus, in comparison to traditional rivalry, CFS allows for more controlled, predictable and prolonged manipulations of awareness.

The relative strength of suppression induced by CFS may be due to a number of factors; the dynamic nature of the mask may reduce the neural adaptation that causes frequent perceptual switches in traditional rivalry (Shimaoka & Kaneko, 2011). Moreover, the spatiotemporal structure of the mask may exploit human sensory sensitivity; the mask can be selected to maximise human contrast and flicker sensitivity (Yang & Blake, 2012). It is currently disputed as to whether CFS constitutes a particularly robust form of binocular rivalry or whether it results from separate mechanisms (Shimaoka & Kaneko, 2011).

1.6.1. The breaking continuous flash suppression (bCFS) paradigm.

A popular application of CFS has been to use the length of the initial suppression period in CFS as a correlate of the unconscious salience of the suppressed image (the bCFS paradigm- see Figure 1.7). This is usually measured by the time it takes for an observer to report the presence or location of the initially suppressed stimulus. This logic is that if one stimulus breaks suppression more rapidly than another, then it has received prioritised unconscious processing during suppression and has received prioritised access to awareness.

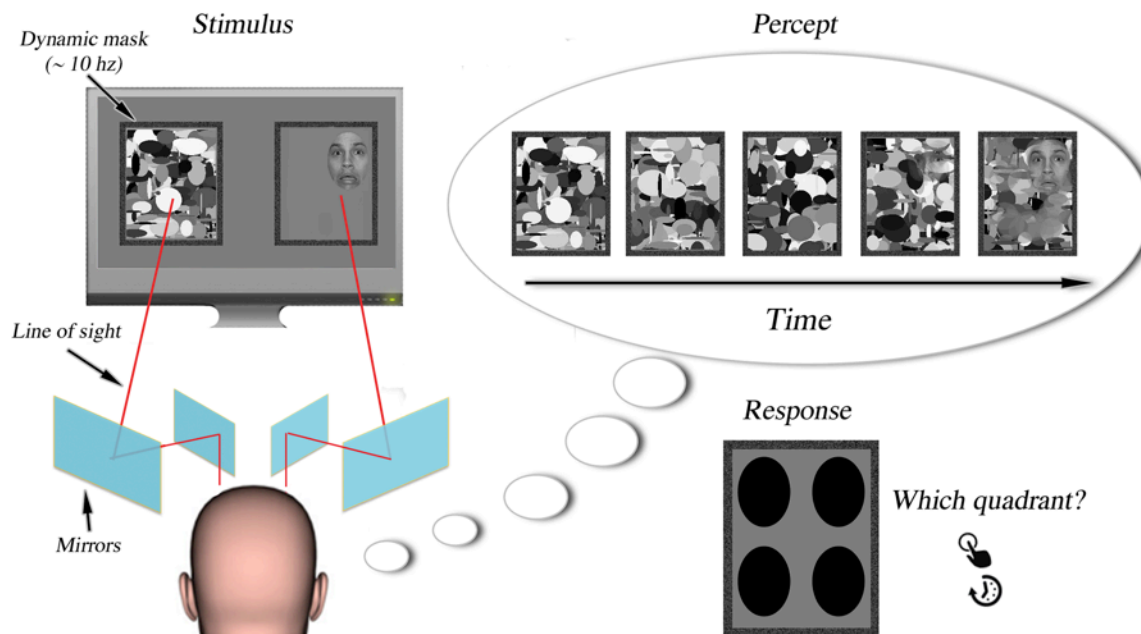


Figure 1.7| **Continuous flash suppression (CFS) and the breaking continuous flash suppression (bCFS) paradigm**

In CFS, a dynamic masking pattern is presented to one eye, which suppresses conscious perception of input presented to the other for long durations.

In the bCFS paradigm, a target image is presented to the suppressed eye and observers are asked to respond when it breaks suppression to become visible.

The logic of the bCFS task is that the time it takes for an initially invisible stimulus to reach awareness reflects its unconscious processing advantage.

Figure 1.7. Continuous flash suppression (CFS) and the breaking continuous flash suppression (bCFS) paradigm

With respect to studying unconscious perception, this paradigm has several advantages over the classic binocular rivalry paradigm. Firstly, the confounding influence of conscious evaluation

during dominance is reduced, as the trial ends as soon as the target stimulus becomes dominant. Secondly, the possibility of “mixed” states or piecemeal rivalry is reduced, as observers respond as soon as they detect any deviation from the (perceptually dissimilar) masking stimulus.

The bCFS technique was first exploited by (Jiang, Costello, & He, 2006), who found that observers reported upright faces breaking CFS faster than inverted faces. This finding suggests that the visual system has some unconscious sensitivity that prefers and promotes ‘face like’ configurations. Subsequently, Yang, Zald and Blake (2007) found that observers reliably detected fearful facial expressions from CFS faster than other expressions. The authors interpreted this ‘fear advantage’ as being consistent with the notion that threatening facial signals are processed and prioritised independently of awareness. A possibly related finding is that face stimuli with a gaze directed at an observer reach awareness faster than those with an averted gaze, thus suggesting that some sensitivity to specific social facial signals is retained outside of awareness (Stein, Senju, Peelen, & Sterzer, 2011).

More recently, the bCFS paradigm has also been exploited to study clinical differences in fearful face processing. Specifically, Sylvers, Brennan and Lillienfeld, (2011) used CFS to investigate the origin of psychopathic deficits in fear recognition ability. Using a bCFS paradigm, Sylvers and colleagues observed that higher scores on the ‘callous/unemotional’ psychopathy subscale predicted a reduced ‘fear advantage’. The authors findings support the idea facial expression recognition occurs without awareness and by extension, that fear recognition deficits in psychopathy may be mediated by unconscious mechanisms.

Recent, decisive data from a bCFS paradigm has questioned the extent to which the ‘fear advantage’ in bCFS reflects unconscious, threat sensitive processes. A patient (S.M) with bilateral lesions to the amygdala and impaired fear recognition ability displays a ‘fear advantage’ well within the normal range of control subjects (Tsuchiya, Moradi, Felsen, Yamakazi, & Adolphs, 2009). Given that S.M has severe deficits in perceiving and responding to threat in a wide range of situations (Feinstein, Adolphs, Damasio, & Tranel, 2011) this suggests that the fear advantage may not be mediated by threat specific processes and may be better explained by low-level properties of the expression. This is a possibility that is returned to later on.

Several findings in the bCFS literature appear to conflict with the notion that unconsciously presented threats are selectively prioritised. For instance, negative words, angry, dominant and untrustworthy faces are all negatively valenced, but have actually been reported to reach awareness *slower* relative to their neutral counterparts (Gray et al., 2013; Stewart et al., 2011; Yang & Yeh, 2011). These findings are hard to reconcile with the idea of an adaptive threat sensitive mechanism, since slowed perception of threatening stimuli is clearly maladaptive. Interestingly, although one

study has shown that negative words are perceived slower than neutral (Yang & Yeh, 2011), another shows that they are perceived faster than neutral (Sklar et al., 2012). Moreover, angry faces appear to be perceived slower than neutral if they are photographic stimuli (Gray et al., 2013), but not if they are in schematic form (Stein & Sterzer, 2012). The standard hypothesis can therefore only be strongly supported by a selective and uncritical review of the bCFS literature.

1.6.2. Physiological studies.

In a fear conditioning study, Raio, Carmel, Carrasco and Phelps (2012) used CFS to present two images of fearful faces outside of awareness. For one of the images (CS+), 50 % of the presentations co-terminated with an electric shock and for the other image (CS-), no shock occurred. It was found that skin conductance responses on subsequent non-reinforced trials were larger to the CS+ than to the CS-, suggesting that fear acquisition can occur in the absence of awareness. Although relatively stringent subjective and objective awareness measures were employed in this study, a puzzling aspect of the data is that the authors sub-divided the data into ‘early’ and ‘late’ trials and the differential SCR was only observed for early trials. Moreover, the differential SCR was actually substantially smaller than that obtained from a control experiment, wherein the CS+ and CS- were fully visible. Thus, an alternative interpretation of the data is that fear learning was actually substantially modulated by, rather than independent of, awareness. The following year, Lapate, Rokers, Li and Davidson (2013) observed larger skin conductance responses to fearful face stimuli than to images of flowers presented under CFS. Moreover, the magnitude of SCR predicted subjective evaluation of a subsequently viewed neutral face, such that larger SCRs led to more negative evaluations. It is notable, however, that images of spiders were also presented in this experiment and that they only yielded a differential SCR when they were consciously presented. Thus, like the results of Raio et al (2012), the are data mixed with respect to whether they support the notion that adaptive physiological responses occur without awareness of the eliciting stimulus.

1.6.3. Neuroimaging studies.

Jiang and He (2006) employed fMRI in conjunction with CFS to identify the aspects of face processing that remain effective in the absence of awareness. Remarkably, it was found that subcortical amygdala responses did not differ in magnitude between conscious and unconscious

perception of fearful faces. These findings have been widely cited as evidence that the amygdala mediates nonconscious emotional perception of faces (Axelrod, Bar, & Rees, 2015; Jiang et al., 2009, Tamietto & de Gelder, 2010, Hassin, 2013). Such findings are however incompatible with single cell recordings recorded from the amygdala under (single) flash suppression (Kreiman, Fried, & Koch, 2002), which fail to show any measurable response to emotional face stimuli. This inconsistency between single cell and fMRI responses to emotionally salient stimuli and their dependency on visual awareness is a complication that leaves this issue unresolved. A further complication is the aforementioned study of S.M by Tsuchiya et al (2009), which demonstrates that, whatever amygdala activity is elicited by fearful faces, it is probably unrelated to the fear advantage in breaking CFS. By extension, it is noteworthy that none of these neuroimaging studies have assessed concomitant changes in behaviour in response to suppressed stimuli, making it hard to infer strong links between amygdala activity and prioritised processing of threat stimuli.

1.7. Limitations of Continuous Flash Suppression.

1.7.1. The influence of conscious processes are hard to control.

bCFS studies use suppression duration (indexed by response time) as a measure of differential unconscious processing (Sylvers, et al., 2011; Yang et al, 2007;). It is important to note that this response time comprises of both an unconscious component that reflects ‘pure’ suppression duration, but is also confounded by the time it takes observers to report the stimulus once it becomes visible (Stein, Hebart, & Sterzer, 2011). It may therefore be inappropriate to conceptualise differential response times in these studies as genuine differences in unconscious processing.

In order to strengthen the conviction that detection times in bCFS tasks purely reflect unconscious specific processes, researchers have employed conscious control experiments (figure 1. 8a-8d). A typical control experiment requires observers to detect stimuli superimposed on the masking pattern whose contrast is ramped in at a constant rate (Jiang, Costello, & He, 2007; Yang, Zald, & Blake, 2007). The logic is that if differences in detection time are observed in a CFS condition, but not in this control experiment, then effects should be attributed to unconscious processes. Whilst appealing in its simplicity, this logic is unsound, because it is hard to mimic the temporal uncertainty (Brascamp, van Ee, Noest, Jacobs & van den Berg, 2006) and piecemeal visibility (Yang, Rose, & Blake, 1992) experienced in bistable perception. As a result, the

appearance of stimuli in a conscious control condition is much more easily anticipated. Thus, conscious control experiments are not perceptually comparable to the viewing conditions of CFS and may not be optimally suited for isolating unconscious processing.

1.7.2 Low level confounds

It is also important to reflect on the fact that detection of stimuli from CFS is strongly influenced by low-level stimulus properties such as luminance, contrast and spatial frequency profile. Unless these properties are equated across stimulus categories, it is impossible to claim that prioritised awareness is driven by unconscious processes sensitive to emotion. Gray, Adams, Hedger, Newton and Garner (2013) investigated this issue by creating facial expression stimuli that maintained the low-level image properties (mean luminance, contrast) of normal facial expressions, but vastly reduced their recognisable emotional content. This was achieved by reversing the luminance profile of expressions and inverting them spatially (figure 1.8e). Interestingly Gray et al replicated the same fear advantage as Yang et al. However, these apparently ‘emotional’ modulations of awareness were maintained for control stimuli, whose recognisable emotional content was vastly reduced. Thus the authors reasoned that the fear advantage is parsimoniously explained by the low-level variability between facial expressions (which is maintained in the control condition), rather than unconscious processes sensitive to emotional meaning.

Related to this issue, Yang and Blake (2012) have investigated the spatial properties of CFS by manipulating the characteristics of suppressed stimuli. One key finding was that the length of CFS suppression was strongly modulated by the spatial frequency profile of the suppressed images. On this basis, to demonstrate that emotion- specific mechanisms underlie prioritised access to awareness from CFS, the spatial profile of suppressed stimuli should be controlled, or a CFS mask developed that is not biased towards the suppressing certain spatial profiles. This is likely to be a difficult challenge for future research.

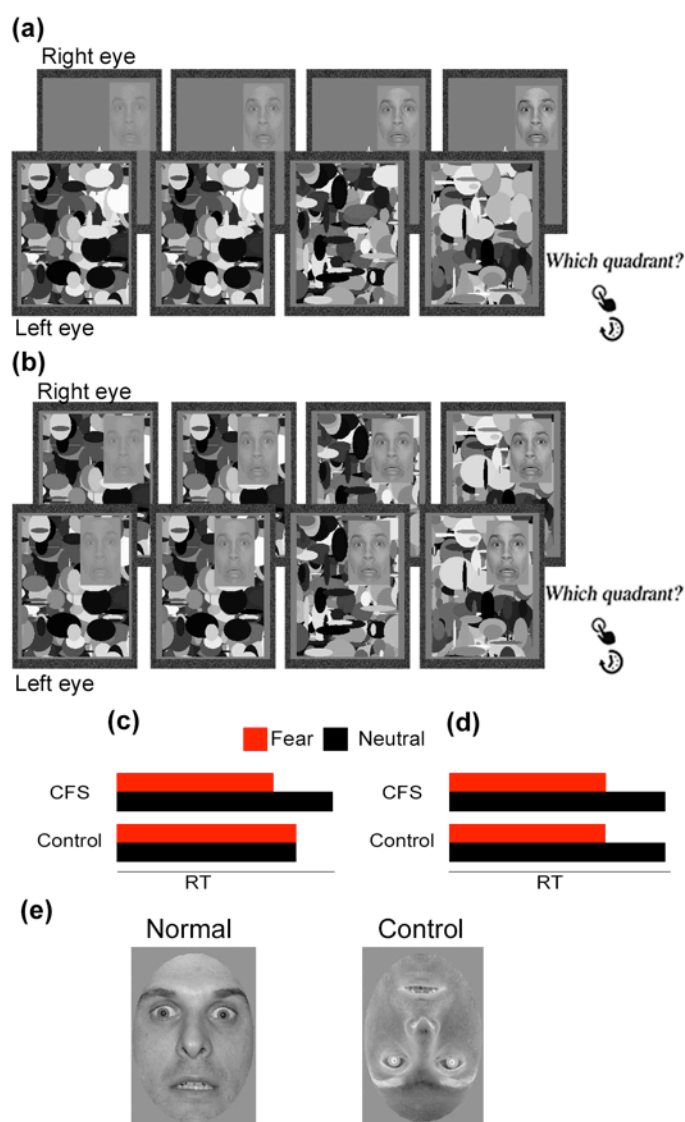


Figure 1.8| *Control experiments and stimuli for the bCFS paradigm.*

The figure in (a) shows a typical bCFS trial. Some researchers have devised control experiments to confirm that differences in detection in bCFS experiments are due to unconscious processes and not simple, inherent differences in detectability (i.e. differences when the stimuli are fully visible).

The trial sequence for a control experiment is shown in (b). Usually this consists of presenting both the masking pattern and target stimulus to both eyes. (c) shows the expected pattern of data if the fearful faces genuinely break into awareness faster than neutral faces. (d) shows the expected pattern of data if the 'fear advantage' can be attributed to differences in response times when the stimuli are fully visible.

(e) shows the normal and control stimuli employed by Gray et al., (2013). Luminance profile reversal and spatial inversion disrupt emotion recognition, but maintain mean luminance, contrast and spatial frequency profile.

Figure 1.8. Control experiments and stimuli for the bCFS paradigm.

1.8. Other findings: Visual Crowding and Patients with Affective Blindsight.

1.8.1. Visual crowding.

Visual crowding refers to the deleterious influence of nearby contours on visual discrimination (Whitney & Levi, 2011). In peripheral vision, an object that is easily recognised when isolated becomes unrecognisable and indistinct if surrounded by other objects (Figure 1.9). Critically, crowding does not impair *detection*, it impairs *recognition*. In this sense, crowding can be used to manipulate awareness of the emotional *meaning* of a stimulus, but does not cause it to disappear from awareness altogether. Crowding constitutes an appealing alternative method to masking and interocular suppression methods, since it occurs frequently in normal viewing conditions and cluttered natural scenes. Moreover, unlike CFS and masking there is absolutely no time constraint in terms of how long stimuli can be presented for.

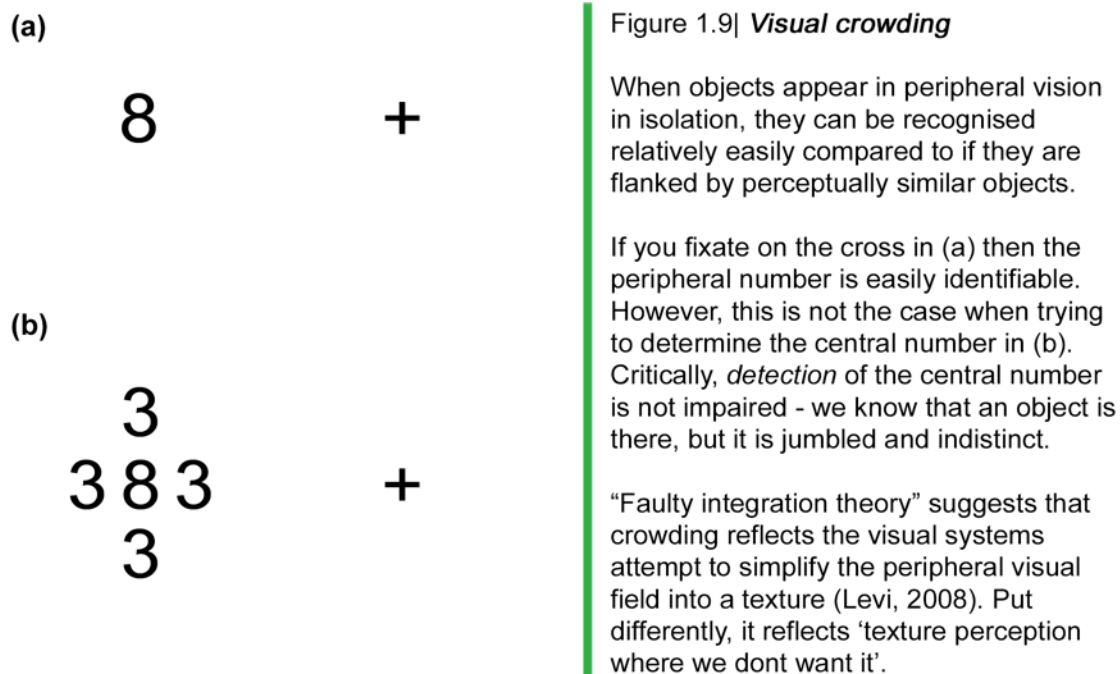


Figure 1.9. Visual crowding.

Koudier, Berthet and Faivre (2011) presented observers with happy and angry faces in peripheral vision, flanked by scrambled faces to induce crowding. After this presentation, observers were asked to make affective judgements (‘pleasant or unpleasant?’) about a subsequently presented Chinese character. It was found that, despite being at chance level for discriminating whether happy and angry faces were presented, observers reported that the character was pleasant more frequently if it was preceded by a happy face. This novel study suggests that emotional features that cannot be consciously discriminated are nonetheless able to affect subsequent decisions.

Building on these initial observations, Faivre, Berthet and Koudier (2012) conducted a series of experiments that attempted to replicate the same effect using crowding, masking and CFS paradigms to prevent awareness of the faces. Despite replicating the effect of crowded facial expressions on subsequent preferences, these effects were eliminated when the stimuli were rendered perceptually invisible by masking or CFS. Importantly, this implies that the null effects for masking and CFS may not reflect evidence for the absence of unconscious processing, but rather the methodological limitations of these methods. In this context, crowding may provide a potentially powerful, ecologically valid method for future research.

1.8.2. Patients with affective blindsight.

The human primary visual cortex (V1) is organised retinotopically, such that lesions to areas of V1 will lead to 'blindness' in a corresponding area of the visual field. Under certain conditions, emotionally salient stimuli presented to the 'blind' region of the visual field can be accurately identified, or located, despite not being consciously perceived. The first report of this kind was published by Beatrice deGelder and co-workers in 1999 (deGelder, Vroomen, Pourtois, & Weiskrantz, 1999). Experiments were conducted on one individual (patient GY) who had blindness in his right visual field following damage to the left occipital lobe. The patient was required to 'guess', under forced choice conditions, which emotional expression was presented in his blind hemifield. It was found that, under many of the conditions that were studied, GY was able to discriminate the facial expressions at significantly above chance level, despite no indication that he could 'see' them. The main conclusions stemming from this finding are conceptually similar to those emanating from previous lesion studies conducted on rats. The discrimination of emotional signals in the absence of the visual cortex suggests that alternative (possibly subcortical) processing pathways may mediate nonconscious emotional perception. Later reports indicated that affective blindsight is also associated with selective responses in the amygdala (Pegna, Khateb, Lazeyras, & Seghier, 2005), facial mimicry and pupil dilatory responses (Tamietto et al., 2009).

One recent, fascinating study employed diffusion tensor imaging (DTI) to investigate structural brain differences between an individual with affective blindsight and age-matched controls. Briefly, DTI allows the non-invasive, in-vivo mapping of the diffusion process of molecules. Since diffusion reflects interactions with fibres and membranes, it can be used to test hypotheses about tissue architecture and plausible pathways in the brain (Le Bihan et al., 2001). Using DTI, Tamietto et al (2012) were able to reconstruct fibre bundles between the superior

colliculous, pulvinar and amygdala of healthy subjects, thereby providing critical evidence for a subcortical ‘low road’ in humans that was previously lacking. Interestingly, these pathways were more or less equivalent in density for each hemisphere of the control subjects, but were substantially denser in the blind hemisphere of the blindsight subject. The authors suggested that destruction of the visual cortex led to functional strengthening of connections that underlie unconscious emotion processing.

Several statistical and conceptual issues are worth discussing in relation to studying affective blindsight. Firstly, it is somewhat puzzling that although certain patients can discriminate facial expressions above chance level, performance is not anywhere near perfect (60-70% accuracy is representative in such studies), which calls into question the usefulness of this residual processing ability. Secondly, to take the original observations of deGelder et al (1999) as an example, above chance discrimination performance was observed for dynamic faces in a randomised presentation, but not for static faces in a blocked presentation. This could mean genuine affective blindsight exists, but is restricted to certain conditions. However, it may also mean that affective blindsight does not exist and this pattern of data simply reflects regression to the mean: following one extreme random event, the next is more likely to be less extreme. A third issue is that, conceptually, it is hard to reconcile an observer’s subjective lack of awareness with their above-chance discrimination ability. According to an objective criterion, these studies do not indicate evidence for unconscious discrimination of emotional stimuli- *the discrimination of emotion is evidence that the observer is aware of the stimuli.*

It is important to reflect on a number of other plausible mechanisms underlying affective blindsight. Firstly, destruction of the visual cortex does not necessarily imply that processing of stimuli presented to the blind portion of the visual field must be mediated by a sub-cortical pathway. There is evidence that the LGN, which receives direct retinal input (before the visual cortex) has direct projections to motion sensitive cortex, thereby bypassing V1 and the site of cortical damage (Schmid et al., 2010; Sincich, Park, Wohlgenuth, & Horton, 2004). Furthermore, there may be some massive unobserved post lesion or experience dependent cortical plasticity, meaning that functions previously undertaken by damaged substrates may be subserved by other cortical structures. Finally, at the most basic level, it is problematic to use observations from individuals with brain damage to make assumptions about how unconscious processes operate in healthy populations, particularly given the small number of patients in these studies.

1.9. Summary and Structure of Thesis.

In this chapter I have reviewed several independent lines of evidence that suggest perceptually invisible stimuli can elicit autonomic and behavioural responses consistent with emotional arousal. Additionally, functional neuroimaging suggests that this may be mediated by specific neural substrates that operate somewhat independently of conscious awareness. However, as outlined throughout, this evidence base is equivocal, since there are a number of controversial and inconsistent findings. As a result, there is a need for a more refined, quantitative analysis of the literature. Moreover, a major limitation of this literature is that i) insufficiently rigorous measures of awareness and ii) inadequate controls for low-level confounds may undermine the validity of many findings. The different paradigms for manipulating awareness also have unique mechanisms, which precludes researchers from making generalisations beyond the specific paradigm used. Facing these formidable challenges will ultimately depend on more quantitative and systematic measures of awareness and combining research knowledge across carefully controlled, novel experimental designs. This is the goal of the present work.

- Whereas *Chapter 1* has provided a narrative review, *Chapter 2* provides a timely, quantitative analysis of the extent to which emotionally salient (specifically threatening) stimuli receive a processing advantage in the absence of awareness. Meta analytic methods are employed to evaluate the existing evidence base from three widely-used behavioural paradigms: masked visual probe, binocular rivalry and breaking continuous flash suppression.
- *Chapter 3* investigates the extent to which unconsciously presented threatening stimuli modulate adaptive physiological and behavioral responses. Specifically, CFS is employed in combination with an attentional cuing paradigm and concurrent measurement of skin conductance. Awareness is assessed according to stringent objective and subjective criteria.
- *Chapter 4* tackles the intractable issue of low level confounds. Face stimuli, particularly fearful faces, are most commonly used to evidence unconscious processing of emotionally salient stimuli. The low-level salience of facial expressions is quantified by calculating their effective contrast: the relationship between their Fourier spectrum and the human contrast sensitivity function. We investigate whether the prioritized processing fear faces is better explained by low-level variability (effective contrast) or unconscious processes sensitive to threat. Two detection tasks are employed, using masking and CFS paradigms.

- In *Chapter 5*, an attentional cuing study is conducted, with face stimuli presented under normal (conscious) viewing, CFS and masking. Signal detection measures are employed to assess observer awareness. Controls for low level confounds are employed.

Are Visual Threats Prioritised Without Awareness?: A Critical Review and Meta Analysis Involving 3 Behavioural Paradigms and 2696 Observers.

2.a. Abstract.

Given capacity limits, only a subset of stimuli give rise to a conscious percept. Neurocognitive models suggest that humans have evolved mechanisms that operate without awareness and prioritize threatening stimuli over neutral stimuli in subsequent perception. In this meta analysis, we review evidence for this ‘standard hypothesis’ emanating from three widely used, but rather different experimental paradigms that have been used to manipulate awareness. We found a small pooled threat-bias effect in the masked visual probe paradigm, a medium effect in the binocular rivalry paradigm and highly inconsistent effects in the breaking continuous flash suppression paradigm. Substantial heterogeneity was explained by the stimulus type: the only threat stimuli that were robustly prioritized across all three paradigms were fearful faces. Meta regression revealed that anxiety may modulate threat-biases, but only under specific presentation conditions. We also found that insufficiently rigorous awareness measures, inadequate control of response biases and low level confounds may undermine claims of genuine unconscious threat processing. Considering the data together, we suggest that uncritical acceptance of the standard hypothesis is premature: current behavioral evidence for threat-sensitive visual processing that operates without awareness is weak.

2.b. Publication note.

This chapter is published as a research article: Hedger, N., Gray, K.L.H., Garner, M. Adams, W.J.A. (2016). Are visual threats priotitized without awareness? A critical review

and meta analysis involving 3 behavioral paradigms and 2696 observers . Psychological Bulletin, in press.

2.c. Contributions.

The search of the literature was conducted by Nicholas Hedger and Katie Gray. Analysis and write up were completed by Nicholas Hedger under the supervision of Wendy Adams and Matthew Garner.

2.1. Background.

Our eyes receive a vast array of visual information. However, due to capacity limits, only a sub-set of stimuli are consciously perceived at any one time (Dehaene & Changuex, 2011). The visual system must cope with these capacity constraints by guiding sensory processing towards the stimuli that are most important to our survival. Since it may take hundreds of milliseconds for visual stimulation to generate a conscious percept (Koch, 2004) it would be advantageous for threats to influence perception or behavior (e.g. by directing attention, or initiating physiological responses) before, or independently of their conscious registration. Such an advantage could, quite literally, be the difference between survival and death. The ‘standard hypothesis’ (Tamietto & deGelder, 2010) holds that humans have evolved a dedicated subcortical visual pathway that evaluates threat independently of conscious awareness and guides the selection of stimulus information for prioritized processing (Ohman, 2005; Ohman, Carlsson, Lundqvist, & Ingvar, 2007). However, despite the intuitive appeal of this notion, the extent to which threatening stimuli are genuinely processed in the absence of awareness remains strongly debated (Pessoa, 2005; Pessoa & Adolphs, 2010).

It is clear and uncontroversial that we are not aware of all aspects of visual processing; for instance, we cannot report the ‘low-level’ activity of individual retinal ganglion cells. Rather, the majority of research interest (and controversy) in unconscious perception is rooted in claims that the ‘meaning’ of a stimulus (such as whether it is threatening) can be registered without awareness and influence subsequent perceptual and cognitive operations (Goodale & Milner, 2004; Hannula, Simons, & Cohen, 2005; Hesselmann & Moors, 2015; Pessoa, 2005).

To investigate this issue, techniques that allow visual input to be dissociated from awareness such as backward masking, binocular rivalry and continuous flash suppression have become widely used (see Figure 2.1). These methods have the potential to reveal the answer to the question: ‘Are visual threats prioritized without awareness?’. We present a critical review and meta-analysis of this literature.

Figure 2.1. Number of PubMed citations that include the terms ‘backward masking’, ‘binocular rivalry’ and ‘continuous flash suppression’ in the title and / or abstract as a function of publication date. Please contact the author, or consult the published article for this figure.

2.2. The Present Review: Justification and Scope.

Understanding the division of labour between conscious and unconscious processes has attracted considerable research effort (Bargh & Morsella, 2008; Dubois & Faivre, 2014; Gayet, Van der Stigchel, & Paffen; Hassin, 2013). Moreover, the extent to which unconsciously presented threats modulate perception remains one of the most contested issues in psychology. Previous reviews of this topic (e.g. Hannula et al., 2005; Pessoa, 2005) have been narrative rather than

quantitative, and have drawn on a limited sample of publications. A comprehensive meta-analysis is particularly valuable in the area of unconscious perception, where findings are controversial and inconsistent. Null effects induced by unconsciously presented stimuli might be due to the genuine absence of an effect, or due to the necessary weakening of visual signals induced by the paradigms used to manipulate awareness. Effect sizes in the literature may therefore be small, which makes parameter estimation and power analysis particularly informative for future studies.

Our review also provides insight into the relationship between awareness and other visual selection processes, such as attention. There is some dispute as to whether awareness of and attending to a stimulus are equivalent constructs, or whether the conscious vs. unconscious processing distinction is entirely separate from attentional selection (Koch & Tsuchiya, 2007; Lamme, 2003). There is also a related, ongoing discussion about whether attention is necessary or sufficient for awareness and vice versa (Cohen, Cavanagh, Chun, & Nakayama, 2012, van Boxtel, Tsuchiya, & Koch, 2010). Quantifying the extent to which attentional selection occurs independently of awareness provides empirical data to inform this debate.

Our analyses are also important in the context of emotional disorders such as anxiety. Although threat sensitive mechanisms enable humans to respond effectively to danger, anxiety can be a maladaptive condition that is prototypically associated with hypersensitivity to threat, excessive fear and disruption to normal functioning (Eysenck, 1997). Prominent cognitive theories suggest that this hypersensitivity contributes to the etiology, maintenance or exacerbation of anxious disorders (Bishop, 2007; Matthews & Macleod, 2005). Specifically, this hypersensitivity is thought to arise from dysfunction in ‘automatic’ threat-sensitive mechanisms that operate without conscious awareness (Mogg & Bradley, 1998). A better understanding of mechanisms involved in unconscious emotion processing will inform cognitive-behavioral models of psychopathology, and help refine therapeutic interventions that systematically target discrete cognitive biases e.g. cognitive-behavioral therapies (Rapee, & Heimberg, 1997) or cognitive bias modification (Beard, 2011).

The standard hypothesis, which states that threats are prioritized in the absence of their conscious registration, continues to shape a large body of theoretical work, experimental research and clinical practice – our review provides a timely and comprehensive analysis of evidence in this area. It a) clarifies to what extent and under what conditions threatening stimuli are prioritized without awareness. b) It identifies important gaps and shortcomings in the literature and c) suggests new directions for future research, including improved methods of data acquisition, analysis and reporting.

2.3. Definitions.

Although most people have an intuitive grasp of what ‘threat’ and ‘conscious awareness’ mean, these abstract concepts are hard to define in a manner precise enough for scientific exploration. In fact, in the empirical literature they are often vaguely described and have long been a source of confusion (Pessoa, 2008; Le Doux, 2013; Wiens, 2007).

2.3.1. *What is a threatening stimulus?*

Ecological theories propose that there are three broad classes of threatening stimuli, which reflect the different mechanisms by which an organism associates a signal with the likely occurrence of a negative outcome (Adolphs, 2013; Boyer & Bergstrom, 2011). Firstly there may be an initial repertoire of ‘phylogenetic’ threat stimuli (see Ohman & Mineka, 2001, for a discussion) whose associations may have been set by evolution, such as an approaching predator (Ohman & Mineka, 2001), or heights (Poulton, Davies, Menzies, Langley, & Silva, 1998). Secondly, there are ‘ontogenetic’ threats that are learnt to be dangerous, such as weapons (Blanchette, 2006). Lastly, there are those stimuli that pose no immediate intrinsic threat themselves, but are symbolic, more abstract representations of the above two classes of stimuli (e.g. negative word stimuli, warning signs). The mechanisms through which these stimuli acquire threat value may vary: e.g. classical conditioning, vicarious conditioning/ modeling of others (Ollsson & Phelps, 2007) or through verbal pathways (Field, Lawson, & Banerjee, 2008). Across a range of species, these three categories of stimuli have been found to elicit a continuum of adaptive physiological, behavioural and cognitive responses that form part of a ‘defensive cascade’ (Blanchard & Blanchard, 1988). Moreover, despite the apparent diversity in these stimulus categories, they all elicit the subjective experience of negative affect in large samples of human observers (e.g. Bradley, Codispoti, Cuthbert, & Lang, 2001). At the evolutionary level, this may reflect the fact that diverse situations of predation, contamination, status loss, social exclusion and conspecific violence have all been legitimate and recurrent fitness threats for humans, the effects of which are all well documented in the archaeological record (Boyer & Bergstrom, 2011). At the psychological level, theories have reconciled the apparent diversity of threat stimuli with their subjective similarity by proposing that emotional evaluations are mostly based on an initial, primitive ‘core’ affective evaluation of whether stimuli are negative or positive (Barrett, 2006). These evaluations are termed ‘core’ because bivalent categorical distinctions between good and bad (appetitive and aversive) are made

by all humans and are present from birth (Barret, Mesquita, Ochsner, & Gross, 2007). Indeed, emotional evaluations of stimuli are mostly explained by the basic dimensions of valence and arousal (Greenwald, Cook, & Lang, 1989).

Based on the above literature, in this review, we define a threatening stimulus as any negatively valenced visual signal that is predictive of adverse affects to the physical or emotional well-being of the receiver. Examples of threat stimuli include fearful faces, images of animal attack, negative words and otherwise neutral stimuli that have been conditioned to predict a negative event (e.g. via pairing with an electric shock). Considerable evidence suggests that these stimuli trigger a broad pattern of defensive physiological responses (Bradley et al., 2001) and adaptive changes in perception, including their prioritized access to conscious awareness and attentional resources (e.g. Vuilleumier, 2005; Yang, Zald, & Blake, 2007).

2.3.2. What is ‘awareness’ and how is it manipulated and measured?

Various meanings of the term ‘awareness’ are conflated in cognitive psychology (Bargh & Morsella, 2008; Dehaene & Changeux, 2011), which are rooted in two, largely independent research domains. These are i) *subliminal perception*: which is concerned with the processing of ***stimuli of which one is unaware*** and ii) *unconscious cognition*: which is concerned with ***mental processes of which one is unaware*** (Hassin, 2013). In our review, the term ‘awareness’ is used to refer to the former definition, i.e. the awareness of a stimulus.

How has awareness of stimuli been measured? The simplest, but least conservative method is to use observers’ reports to index whether a stimulus is perceived. Historically, this *subjective* approach derives its motivation from the idea that only observers themselves have access to their inner states and that this is the only reliable source of information about conscious experience (James, 1890). However the development of signal detection theory (SDT: Green & Swets, 1996), raised concerns that subjective measures are prone to response bias or criterion effects, such as reluctance to report a signal if it is degraded or brief. According to SDT, due to internal neural noise, the absence of a signal may elicit a strong sensory state and the presence of a signal may elicit a weak sensory state. Reports of awareness are thus probabilistic statements based on an internal threshold that demarcates sufficient “strength of evidence” that a signal was present (Pastore, Crawley, Berens, & Skelly, 2003). If an observer sets this threshold too high, they may incorrectly reject their conscious perception (a type 2 error) and report they are ‘unaware’ of the stimulus. As a result of these issues associated with subjective report, *objective criteria* have also

been employed to determine awareness. The objective approach measures awareness according to an observer's ability to perform statistically above chance in discriminating alternative stimulus states (e.g. left or right location) in a forced-choice classification task (Macmillan & Creelman, 2005). For instance, observers might be given 100 trials in which they are asked to report whether a stimulus appeared left or right of fixation (where both are equally probable). Under the null hypothesis (observers are unaware of the stimulus), we would expect observers to respond correctly on approximately 50 of the trials (the most likely outcome given random responses). However, if an observer achieves 59 or more correct responses, the null hypothesis is rejected (performance is 'significantly' above chance performance, according to a binomial test) and that observer would be classified as 'objectively aware'. Objective awareness checks that probe stimulus detection via discrimination of a stimulus dimension (e.g. 'Was it on the left or right?') that is orthogonal to the critical dimension ('Was it visible?') are thought to be less prone to the response biases that can effect an observers' subjective report of the phenomena under investigation. A current view is that both objective and subjective measures have conceptual and practical limitations and so a range of measures should be used in combination to comprehensively characterize visual awareness (for extended discussion see Sandberg, Timmermans, Overgaard, & Cleeremans, 2010; Szecepanowski & Pessoa, 2007; Wiens, 2007).

2.4. Experimental Paradigms.

We applied three criteria when searching the evidence base for experimental paradigms to investigate our research question. Firstly, we reasoned that the paradigm must include an experimental manipulation that suppresses threatening and neutral stimuli from awareness. Secondly the paradigm must include a behavioral measure sensitive to enhanced perceptual selection of the threatening (relative to the neutral) stimulus to index its prioritization. Paradigms that manipulate awareness, but measure 'late' semantic congruency effects unrelated to perceptual selection (Algom, Chajut, & Lev, 2004), such as masked emotional Stroop and masked semantic priming, were not included (other meta-analyses on these subjects exist elsewhere, see Bar Haim et al., 2007; Van den Bussche, Van den Noortgrate, & Reynvoet, 2009). Third, we made an *a priori* decision that each paradigm must be represented by at least 10 independent studies to allow useful and informative analyses. A summary of other excluded paradigms is included in Appendix A. Implementing these criteria resulted in the inclusion of three experimental paradigms in the analyses: masked visual probe, binocular rivalry and breaking continuous flash suppression.

2.4.1. The masked visual probe paradigm.

Description.

In backward masking, a briefly presented target stimulus is quickly replaced by a salient, co-located ‘mask’ stimulus (typically before 40 ms). If the presentation parameters are manipulated appropriately, observers indicate being aware of the mask, but not the target stimulus, i.e. the target is masked from conscious perception. Backward masking appears to disrupt and replace visual processing of the target stimulus (Breitmeyer & Ogmen, 2000; Rolls, Tovee, & Panzeri, 1999). Theories suggest that masking weakens and abbreviates the target-related visual signal, eliminating re-entrant feedback from later stages of processing, which is critical for maintaining a representation in awareness (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006; Green et al., 2005). Masking is relatively simple to administer and continues to be widely used in studies that aim to manipulate visual awareness.

The masked visual probe (MVP) paradigm combines backward masking with an attentional cuing paradigm. The generic trial sequence is shown in Figure 2.2: (i) Observers view a central fixation point. (ii) A threat stimulus and a neutral stimulus are presented either side of fixation for a brief duration (typically <40ms), immediately followed by (iii) co-located mask stimuli. (iv) A probe stimulus is then presented at either the location preceded by the threat (valid) or the neutral stimulus (invalid). (v) Observers are asked to report an aspect of the probe (a two alternative forced choice discrimination) as quickly as possible.

Figure 2.2. Schematic of generic trial sequence from a masked visual probe task. Please contact the author, or consult the published article for this figure.

What can the MVP paradigm tell us and how is this evidenced?

The MVP paradigm provides an effective tool to probe the theoretical construct of the “orienting network” (Posner, 2012). Since our cognitive systems have limited capacity, they need mechanisms to selectively enhance perceptual processing of relevant, particularly threatening, stimuli. The orienting network is involved in this process by changing the distribution of processing resources across the visual field: attention is disengaged from an initial location (or locations) and engaged elsewhere. This re-distribution of attention is indexed by enhanced behavioural performance and increased neural activity at attended, versus unattended locations (Chica, Martin-Arevalo, Botta, & Lupianez, 2014). The MVP task was developed after initial reports that detection latencies to probe stimuli can be modulated by preceding visual cues (Posner, Snyder, & Davidson, 1980). It follows that spatial attention can be assessed by comparing response latencies to probes that appear in the location of the threat stimulus (often termed ‘valid cue trial’) to those from the neutral location (‘invalid cue trial’). Faster responses in valid (vs. invalid) cue trials suggest that attention is preferentially allocated at the location of the threat stimulus. Thus, by incorporating masking to manipulate stimulus awareness, researchers can determine the extent to which unconsciously presented threat stimuli are prioritized in spatial attention.

Example study: Mogg, Bradley, and Williams (1995).

The most frequently cited MVP study included in our analyses was conducted by Mogg, Bradley and Williams (1995). The authors examined the attentional biases towards subliminally presented negative and neutral stimuli in clinically anxious and healthy control participants. The observers completed an MVP task where they were presented masked pairs of negative and neutral words for 14ms. For anxious observers, but not normal controls, responses to the subsequent probes were significantly faster in valid trials - consistent with attention being preferentially drawn to masked threat stimuli.

To objectively assess awareness of stimuli, the observers completed a separate block of trials, where they discriminated between trials in which word stimuli were presented prior to the mask (50% of trials), or no stimulus was presented prior to the mask (50% of trials). Observers who performed significantly above chance (i.e. significantly above 50% accuracy) were removed from the MVP analyses (5 out of 32 participants). Thus, the data suggest that anxious observers exhibit attentional biases toward threatening stimuli that they are objectively unaware of. The authors interpreted their findings as evidence for an “automatic, preconscious processing bias in anxiety” (p. 31).

2.4.2. Binocular rivalry.

Description.

Under normal viewing conditions, our two eyes receive slightly different views of the world. The visual system is able to combine these similar images into a coherent percept via binocular fusion (Howard & Rogers, 1995). However, binocular rivalry (BR) may occur when our two eyes receive very different input at corresponding retinal locations, with images typically presented separately to each eye via a mirror stereoscope or as a coloured anaglyph (see Figure 2.3). In such cases, the visual system cannot combine the two eyes' images into a coherent percept and instead, perception alternates between them (Wheatstone, 1838). The extended and invariant visual stimulation in BR is thus rather different from backward masking, in which awareness is manipulated by rapidly changing the visual input.

At a neural level, BR has been attributed to reciprocal inhibition between neural populations representing the two eyes' stimuli at distributed stages of the visual processing hierarchy (Blake & Logothetis, 2002). The neural population exerting strongest inhibition achieves access to awareness. Subsequent neural adaptation of the dominant population progressively reduces inhibition of the suppressed stimulus, resulting in a perceptual switch - the previously suppressed stimulus reaches dominance and so on (Alais, Cass, O'Shea & Blake, 2012).

Figure 2.3. Schematic of typical stimuli, percepts and response options in a binocular rivalry paradigm. Please contact the author, or consult the published article for this figure.

What can binocular rivalry tell us? How is this evidenced?

Although BR has been investigated by vision scientists for more than 170 years, a landmark paper by Crick and Koch (1998) stimulated a renewed interest in BR research by popularising the idea that it allows investigation of the dynamics and neural concomitants of consciousness, owing to its capacity to dissociate visual input from awareness (Baker, 2010). At the theoretical level, the perceptual alternations in BR reflect a natural constraint: two different stimuli cannot occupy the same space at the same time, thus the visual system must 'choose' perception of one over the other. Researchers are interested in binocular rivalry because the ensuing 'choices' may be indicative of a variety of processes that the visual system uses to selectively process the retinal images evoked by the environment.

BR does have some similarity to natural vision, in that competition occurs between multiple sensory inputs, with only a subset reaching conscious perception. BR has thus been conceptualized as a means to mimic this selection process under more controlled conditions, by assessing which of two co-located stimuli are prioritized in the competition for awareness. This prioritization is often indexed by the proportion of time that one stimulus dominates (i.e. is perceived) over the other (Levelt, 1965). Whereas the MVP paradigm has been used to index attentional modulation by stimuli suppressed from awareness, the BR paradigm allows researchers to index unconscious processing via the speed or probability with which stimuli gain access to conscious perception. The logic is that if a threatening image is prioritized in the competition for awareness, it will be perceived for a larger proportion of a BR trial than a competing neutral stimulus. There is some evidence that dominance in BR is modulated by higher-level factors, such as object recognition (Yu & Blake, 1992) and surface organization (Graf & Adams, 2008). However, low-level stimulus properties, such as higher contrast and luminance, also robustly increase stimulus dominance in BR (Levelt, 1965). The stimulus properties that increase perceptual dominance in BR are often referred to as determining ‘stimulus strength’.

During prolonged viewing periods, both rivaling images are likely to be perceived multiple times, as perception alternates between the two. This limits the extent to which dominance in BR reflects a purely *unconscious* processing advantage, since prolonged perception of a stimulus could be driven by conscious processes acting on the dominant (visible) image. To address this issue, one can instead record which stimulus is the first to achieve perceptual dominance. This ‘first percept’ measure is considered more suited to investigating the early stages of perceptual selection, since only the initially dominant stimulus is reported (Carter & Cavanagh, 2007).

Example study: Anderson, Siegal, Bliss-Moreau and Feldman Barrett (2011).

The most cited BR study in our analyses was conducted by Anderson, Siegal, Bliss-Moreau and Feldman Barrett (2011). Via an affective learning procedure, Anderson et al. (2011) associated neutral faces with descriptions of social behaviors that were negative (e.g. “he threw a chair at his classmate”), positive (e.g. “he gave up his seat on the bus to a pregnant lady”), or neutral (e.g. “he rode the elevator with a coworker”). In the subsequent BR task, one of the conditioned face images was presented to one eye, and an image of a house was presented to the other eye. Participants continuously reported their percept (face or house) over the 10-second rivalrous trial. Faces paired with negative social behaviors were perceived for significantly longer than the faces paired with positive or neutral social behaviors, or novel faces. The authors concluded that “what we know about someone influences not only how we feel and think about them, but also whether or not we see them in the first place” (p.1448).

2.4.3. Breaking continuous flash suppression (bCFS) Paradigm.

Description.

Continuous flash suppression (CFS, Tsuchiya & Koch, 2005) is a variant of BR in which a stimulus presented to one eye is suppressed from awareness by a competing dynamic noise pattern presented to the other eye. Suppression during CFS is more potent than during traditional BR (as defined by contrast detection thresholds; Tsuchiya, Koch, Gilroy, & Blake, 2006). Temporally, the periods of suppression induced by CFS can last about 10 times longer than suppression induced by traditional BR (Tsuchiya & Koch, 2005). Another attractive property of CFS is that perceptual suppression of a target stimulus can reliably be induced from the onset of a trial. Thus, in comparison to traditional BR, CFS allows for more controlled, predictable and prolonged manipulations of awareness.

The relative strength of suppression induced by CFS may be due to a number of factors; the dynamic nature of the mask may reduce the neural adaptation that causes frequent perceptual switches in traditional BR (Shimaoka & Kaneko, 2011). Moreover, the spatiotemporal structure of the mask may exploit human sensory sensitivity; the mask can be selected to maximize human contrast and flicker sensitivity (Yang & Blake, 2012). It is currently disputed as to whether CFS constitutes a particularly robust form of binocular rivalry, or whether it results from distinct mechanisms (Shimaoka & Kaneko, 2011).

What can the bCFS paradigm tell us? How is this evidenced?

A popular application of CFS has been to use the length of the initial suppression period in CFS as a correlate of the unconscious salience of the suppressed image. This is referred to as the breaking continuous flash suppression, or bCFS paradigm (the ‘b’ refers to ‘breaking’ CFS- see Figure 2.4). Suppression duration is usually measured by the time it takes for an observer to report the presence or location of an initially suppressed stimulus whose contrast is increased over time. This is rooted in the similar assumption that is made about traditional BR: more salient stimuli gain access to awareness more quickly. Thus as with BR, researchers have capitalized on the bCFS paradigm since it may offer insight into the competitive dynamics that underlie prioritized access to conscious perception. For instance, to enable adaptive behaviour, it might be predicted that threatening images would gain faster access to awareness than neutral images. The bCFS paradigm offers a means of testing this prediction.

Figure 2.4. Schematic of typical stimuli, percepts and response options in a breaking continuous flash suppression paradigm. Please contact the author, or consult the published article for this figure.

This paradigm offers several advantages over a conventional rivalry task in which dominance durations are compared for stimuli that compete for resources at the same time and in the same space. Firstly, the likelihood of mixed percepts and associated response biases are reduced (albeit not eliminated) as the trial ends as soon as an observer detects the target stimulus. Secondly, when the duration of percepts are compared between stimuli engaged in BR, it is hard to determine whether increased dominance is due to the salience of the dominant stimulus or the ineffectiveness of the suppressed stimulus. Instead, in bCFS, response times are compared across different stimuli that compete against a common ‘baseline’ dynamic masking pattern, making differential suppression times easier to interpret.

Example study: Yang, Zald and Blake (2007).

The most cited bCFS paper included in our analyses was conducted by Yang, Zald and Blake (2007). Yang et al. presented happy, fearful and neutral faces under CFS and recorded the time it took for observers to detect a face. Each trial consisted of a face presented at a random quadrant in one eye, whilst the CFS mask was presented to the other eye updating at a rate of 10Hz. Results showed that observers were faster at detecting the location of fearful expressions than both

happy and neutral expressions. The findings were interpreted as evidence that “negatively charged facial expressions gain preferential access to awareness” (p.885).

2.5. Meta Analyses: Inclusion and Coding Decisions.

2.5.1. Inclusion criteria.

All studies included in our analyses met all of the following criteria:

1. The study used one of the following paradigms: masked visual probe, binocular rivalry, or breaking continuous flash suppression.
2. The study was published as a journal article in the English language on or before March 31, 2015.
3. A processing difference between threat-related and neutral stimuli could be assessed. Comparisons between neutral and “emotive” (a combination of positive and threatening) stimuli were excluded.
4. The study was conducted on human subjects.
5. The study was not a re-analysis of existing data.
6. Sufficient information was available for an effect size to be estimated (see “Meta Analysis: Methods”, section below).

2.5.2. Other coding and inclusion decisions.

1. Because anxiety has consistently been linked to increased processing biases for threatening stimuli (Bishop, 2007), we treated samples that were categorized as having high or low levels of self-reported anxiety as separate samples of observers. This allowed us to quantify the effects of anxiety as a moderator. When separate analyses were reported for two or more groups according to some other dimension or personality trait (e.g. carriers of a particular gene; Carlson, Mujica-Parodi, Harmon-Jones, & Hajcak, 2012), the data were pooled into one sample.
2. We excluded samples of patients that were reported to have a clinical diagnosis, unless this was an anxiety disorder. This was done to reduce unnecessary variance, as depression might be expected to modulate threat bias (Mogg et al., 1995; Mogg & Bradley, 2005), but there were insufficient data to reliably characterize effects of disorders other than anxiety. In

practice, only 12 studies included in the analyses reported a depression measure, and these varied across studies (Beck Depression Inventory: Beck, Ward, Mendelson, Mock, & Erbaugh, 1961; Montgomery-Ashberg Depression Rating Scale: Montgomery & Asberg, 1979; Depression Anxiety Stress Scale; Lovibond & Lovibond, 1996).

3. When studies involved a mood induction, therapeutic intervention or drug treatment expected to reduce or enhance threat-related biases (e.g. Maoz, Abend, Fox, Pine, & Bar Haim, 2013), we only included experimental data collected prior to the intervention (at baseline), or from a control group that did not receive an intervention. If there were no baseline data or control group, the study was excluded.
4. If the study manipulated levels of threat intensity (e.g. by conditioning a threat image with an aversive event, or neutral event: Beaver et al., 2005), our effect size reflects the processing difference between the highest level of threat (i.e. the threat image paired with the aversive event) and a neutral stimulus.
5. If studies used spatially inverted threat stimuli to control for low-level confounds (e.g. Yang et al., 2007), the corresponding data were excluded from our main analyses, since this manipulation reduces the recognizable emotional content of the stimulus (Gray et al., 2013). Instead, independent analyses were conducted to examine the effect of this manipulation on the magnitude of threat bias.
6. If the study included a manipulation or degradation of stimuli that was not pertinent to our research question (e.g. spatial filtering: Stein, Seymour, Hebart, & Sterzer, 2013), we included only data corresponding to the un-manipulated (e.g. unfiltered) stimuli, to reduce unnecessary variance.
7. If a study included a conditioning procedure, which assigned negative (CS+) and neutral valences (CS-) to stimuli, we excluded the data if the CS- was not intrinsically 'neutral' (e.g. if the CS+ and CS- were both angry faces; Raes, Koster, Van Damme, Fias, & De Raedt, 2010).

2.5.3. General search and coding strategies.

The search for relevant studies and their coding was conducted by two authors (NH, KHLG). First, we conducted PubMed database searches. Second, we examined the reference sections of all relevant literature reviews for additional studies. Third, we searched the reference sections of all qualifying articles and articles listed as citing the qualifying articles on Google

Scholar. Database search terms, and a summary of the excluded articles are presented according to the ‘Preferred Reporting Items for Systematic reviews and Meta Analysis’ guidelines (PRISMA: Moher, Liberati, Tetzlaff, Altman, & Altman, 2009). The search terms and associated PRISMA flowcharts can be found in the supplementary material (Appendix B).

Details of the coding / moderator variables used within each experimental paradigm are detailed in later sections. The inter-coder agreement between the two authors was high. We calculated the intra-class coefficients (ICCs) and kappa coefficients for the continuous and categorical moderators respectively. The ICCs were all 1.0 due to the straightforward nature of the continuous moderator data and the kappa coefficients ranged from 0.91 (for stimulus type) to 1.0 (for all other moderators). Rare disagreements were resolved via a discussion between the four authors.

2.6. Meta Analysis: Methods.

2.6.1. Effect size metric

The effect size index used for all outcome measures was Cohen’s d , the standardized difference between means (Cohen, 1977). In all cases, a positive value indicates a perceptual bias towards a threatening stimulus relative to a neutral stimulus.

2.6.2. Standardizers for d .

Our primary estimator of Cohen’s d was d_z – the difference between means standardized by the standard deviation of difference scores. The advantage of this metric is that it can be computed directly from just t , p or F values and the corresponding degrees of freedom (Lakens, 2013):

$$d_z = \frac{t}{\sqrt{N}}$$

As all our effects emanated from repeated measures designs, we also estimate an effect size estimate that corrects for the pre-post correlation (d_{RM}) wherever possible (see Appendix C).

In both cases (d_z , d_{RM}), the standard error was calculated via the generic formula:

$$SE = \sqrt{\frac{1}{\frac{N+d^2}{2N}}} \times \sqrt{2(1-r)}$$

If no exact t or p values were reported (e.g. ' $p < .05$ '), we either estimated the effect size from the available information, or, when necessary, excluded it from the analyses (see Appendix D). Additionally, we used two multiple imputation methods to estimate unreported values of moderator variables (see Appendix E). *Regression imputation* (RI) is 'optimistic' and uses the existing relationship between the reported moderator values and effect size to predict the unreported values. Conversely, *random-sample imputation* (RSI) is more conservative and assumes that missing values are random samples of the reported moderator values (i.e. the existing relationship is not predictive of the missing values).

2.6.3. Model and analysis decisions.

We made an *a priori* decision to analyze our effect size data in a random effects model, due to its tolerance of heterogeneous effect sizes and conservative nature of estimation (Cumming, 2011). The random effects model assumes that each study estimates different values from a distribution of population parameters, rather than assuming that studies are direct replications of each other (Schmidt, Oh, & Hayes, 2009).

We assessed heterogeneity across effect sizes by using Cochran's Q and I^2 statistics. Unless reported otherwise, parameter estimates were obtained via restricted maximum likelihood estimation, owing to its superior accuracy given a smaller number of studies (Lopez-Lopez, Marin-Martinez, Sanchez-Meca, Van den Noortgate, & Viechtbauer, 2014). Statistical tests of model coefficients were computed via Wald-type chi squared tests. We additionally used a pseudo- R^2 statistic (Raudenbush, 1994) to assess the extent of effect size heterogeneity that was explained by moderators included in the model (see Appendix F). Model comparisons were conducted via likelihood ratio tests. All analyses were conducted with the 'metafor' package (Viechtbauer, 2010) implemented in the *R* programming language.

2.6.4. Handling dependency among effect sizes.

For each paradigm, we explicitly coded the number of included conditions (nested within samples) and samples (independent groups of participants, nested within studies). Many of the

samples were exposed to multiple conditions, which generates multiple effect sizes for these samples. For instance, in some cases, samples were exposed to more than one type of threatening stimulus (e.g. to fear and angry faces; Gray et al., 2013), meaning that this important moderator occurs at the within sample level and information would be lost by aggregating these effects. Thus, to minimize this information loss and increase statistical power, we used *conditions, rather than samples* as the unit of analysis in our models (k = conditions).

When samples contribute multiple effect sizes in this way, the assumption of independence may be violated and bias the outcome of the meta-analysis, particularly if there is anything unrepresentative about these samples (Matt & Cook, 2009; Rosenthal, 1991). To examine the influence of dependency on our results, we used two strategies. Firstly, we created multi-level models (see Cheung, 2014) wherein conditions (level 2) were nested within their samples (level 3). Because a structural equation modelling approach is used to estimate these models, this allowed us to specify interesting constraints that are otherwise very difficult to test. Using this approach, we were able to partition the heterogeneity between effect sizes into that occurring at level 2 (between conditions) or level 3 (between samples) and also statistically examine whether there was a significant amount of effect size dependency (i.e. does a 3 level model provide a significantly better fit than a 2 level model?). Secondly, we examined the influence of dependency via sensitivity analyses: using random selection procedures, we created data sets where dependency was eliminated by selecting one effect size per independent sample (Greenhouse & Iyengar, 1994).

2.7. Meta Analysis: Results

2.7.1. The MVP paradigm

Summary of included data.

Our inclusion criteria resulted in 28 MVP studies being analyzed, comprising 1407 participants across 39 independent samples. We derived 44 estimates of the threat effect size. The coding system and summary of effects used in the analyses are shown in Tables 2.1 and 2.2. Detailed information about each included effect and demographic information can be found in Appendix G).

Dependent measures.

For the MVP paradigm, Cohen's d reflects the difference in response time between valid and invalid cue trials. Positive values indicate that attention is biased towards the spatial location of threat-related stimuli (faster responses in valid trials).

Table 2.1. Coding of Individual Effects in the MVP Paradigm.

Moderator	Type	Values	Description of variable / theoretical justification	Descriptive statistics	Missing cases
Stimulus type	Categorical	1=fearful face 2= angry face 3= disgust face 4=threatening word. 5=IAPS image 6 = fear relevant CS+	The type of threatening stimulus used in the experiment. Justification: The magnitude of threat bias may differ as a function of the semantic/physical properties of the stimulus. In addition, masking may not be equally effective for all stimulus types (Wiens & Ohman, 2007).	k=44 1=15 2=17 3=1 4=8 5=2 6=1	0
SOA	Continuous	12-34	Stimulus onset asynchrony (SOA) between presentation of the threat and mask stimulus. Justification: SOA is directly related to visibility/awareness of the target stimulus, which may modulate biases towards threat-stimuli.	k=44 M=20.89 SD=8.07 Range=12-34	0
Awareness measure	Dichotomous	0=none 1=objective	How awareness of stimuli was assessed (with an objective awareness check, or with no awareness check). Justification: Previous research suggests that threat-related biases to briefly presented stimuli may be a function of the criteria used to assess awareness (Pessoa, Japee, Sturman, & Ungerleider, 2006).	k=44 0=18 1=26	0
Probe response	Dichotomous	1 = 'where' 2 = 'what'	The response that an observer is required to make to the visual probe in the response phase: 'where' tasks involve reporting the location of the probe (left or right), whereas 'what' tasks involve discriminating the probes' orientation. Justification: Previous research has shown that these tasks are associated with different response times and error rates (Mogg & Bradley, 1999a).	k = 44 1 = 33 2 = 11	0
hpwr	Continuous	0.27-0.43	Metric that summarizes the statistical power of objective awareness checks (see text). Justification: Awareness checks with low power increase the probability that target stimuli were not fully / always suppressed from awareness, i.e. deviations from chance performance in the awareness check may not be detected. This increases the likelihood that threat-related biases could be driven by a small proportion of undetected trials where the observer was aware of stimuli (Hannula et al., 2007; Reingold, 2004).	k=26 M=0.38 SD=0.06 Range= 0.27-0.43	0
Stimulus size	Continuous	2.8-7.0	Diameter of the threat stimulus in degrees of visual angle. Justification: Affective processing of threat images may increase with stimulus size (Codispoti & Cesarei, 2007). Moreover, masking may not be as effective for larger stimuli (Wiens & Ohman, 2007).	k=22 M=5.3 SD=1.4 Range=2.8-7.0	22
STAI-T	Continuous	29.4-58.1	The samples' mean trait anxiety level, as assessed by the Spielberger state-trait anxiety scale (Spielberger, Gorsuch, Lushene, Vagg, & Jacobs, 1983). Justification: Anxiety is associated with enhanced processing biases towards threatening stimuli in visual probe tasks (Bar Haim et al., 2007).	k=15 M=41.34 SD=8.85 Range=29.4-58.10	29
Visual field	Dichotomous	1=left 2=right	Location of threat stimulus: left or right visual field. Evidence suggests that affective perception may be lateralized, such that the right hemisphere may be particularly sensitive to emotional stimuli (Gainotti, Caltagirone, & Zoccolotti, 1993; Mormann et al., 2011). For this reason, some studies have split analyses according to the visual field threatening stimuli were presented in.	k=10 1=5 2=5	0

Table 2.2. Summary of Effects Included in the MVP Analyses

Study/Effect	Sample code	Condition code	Exp	N	Stim	SOA	Aw meas	Probe resp	Hpwr	VF	Group	Stim size	STAI-T	Pool ID
1) Mogg et al., 1994 (i)	1	1	1	36	4	14	1	1	0.38	0	High Trait Anxious		42.4	0
Mogg et al., 1994 (ii)	2	2	1	30	4	14	1	1	0.38	0	Low Trait Anxious		29.1	0
2) Mogg et al., 1995 (i)	3	3	1	17	4	14	1	1	0.36	0	Clinical Anxiety		58.1	0
Mogg et al., 1995 (ii)	4	4	1	15	4	14	1	1	0.36	0	Normal Controls		39.1	0
3) Mogg & Bradley, 1999b (i)	5	5	1	33	2	17	1	1	0.41	1		2.9	42.0	0
Mogg & Bradley, 1999b (ii)	5	6	1	/	2	17	1	1	0.41	2		2.9	39.0	0
Mogg & Bradley, 1999b (iii)	6	7	3	22	2	17	1	1	0.43	1		2.9	39.0	0
Mogg & Bradley, 1999b (iv)	6	8	3	/	2	17	1	1	0.43	2		2.9	42.0	0
4) Mogg & Bradley, 2002 (i)	7	9	1	11	2	17	1	2	0.43	0	High Social Anxiety	4.5	43.7	0
Mogg & Bradley, 2002 (ii)	8	10	1	16	2	17	1	2	0.43	0	Low Social Anxiety	4.5	33.9	0
5) Fox, 2002 (i)	9	11	2	18	1	17	1	2	0.43	1	High Trait Anxious	5.7	50.4	0
Fox, 2002 (ii)	9	12	2	/	1	17	1	2	0.43	2	High Trait Anxious	5.7	50.4	0
Fox, 2002 (iii)	10	13	2	18	1	17	1	2	0.43	1	Low Trait Anxious	5.7	29.4	0
Fox, 2002 (iv)	10	14	2	/	1	17	1	2	0.43	2	Low Trait Anxious	5.7	29.4	0
6) Keogh et al., 2003 (i)	11	15	1	81	4	17	1	1	0.43	0				0
7) Beaver et al., 2005 (i)	12	16	2	10	6	17	1	2	0.43	0	High-aversive group			0
8) Hunt et al., 2006 (i)	13	17	1	55	4	17	1	1	0.43	0				0
9) Koster et al., 2007 (i)	14	18	1a	49	2	34	1	1		0		6.7		0
Koster et al., 2007 (ii)	15	19	2	24	2	34	1	1	0.27	0		6.7		0
Koster et al., 2007 (iii)	16	20	3	19	2	14	1	1	0.27	0		6.7		0
10) Murphy et al., 2007 (i)	17	21	1	12	1	17	0	2		0	Placebo control			0
11) Stone & Valentine, 2007 (i)	18	22	1	24	2	17	1	1	0.39	0				0
Stone & Valentine, 2007 (ii)	19	23	2	28	2	17	1	1	0.39	0				0
12) Wirth & Schultheiss, 2007 (i)	20	24	2	52	2	12	0	1		0				0
13) Schultheiss & Hale, 2007 (i)	21	25	1	52	2	12	0	1		0				0
Schultheiss & Hale, 2007 (ii)	22	26	2	60	2	12	0	1		0				0
14) Carlson & Reinke, 2008 (i)	23	27	1	30	1	33	0	1		0		6.0		0
Carlson & Reinke, 2008 (ii)	24	28	2	30	1	33	0	1		0		6.0		0
15) Monk et al., 2008 (i)	25	29	1	17	2	17	0	1		0	Generalized Anxiety Disorder			0
Monk et al., 2008 (ii)	26	30	1	12	2	17	0	1		0	Control Group			0
16) Carlson et al., 2009a (i)	27	31	1	12	1	33	1	1	0.31	1				0
Carlson et al., 2009a (ii)	27	32	2	/	1	33	1	1	0.31	2				0
17) Carlson et al., 2009b (i)	28	33	1	30	5	33	0	1		0		7.0		0
18) Helzer et al., 2009 (i)	29	34	1	112	4	20	1	1		0				0
19) Fox et al., 2010 (i)	30	35	1	104	5	14	0	2		0		4.0	40.1	0
20) Carlson & Reinke, 2010 (i)	31	36	1	12	1	33	1	1		0		6.0		0
21) Thomason et al., 2010 (i)	32	37	1	20	1	17	0	1		0	ll carriers			1
Thomason et al., 2010 (ii)	32	37	1	31	1	17	0	1		0	s- allele carriers			1
22) Sutton & Altarriba, 2011 (i)	33	38	2	64	4	30	0	2		0				0
23) Carlson et al., 2012 (i)	34	39	1	40	1	33	1	1	0.32	0	SS allele	6.0		2
Carlson et al., 2012 (ii)	34	39	1	10	1	33	1	1	0.32	0	L allele	6.0		2
24) Carlson et al., 2013a (i)	35	40	1	40	1	33	1	1	0.32	0		6.0		0
25) Carlson et al., 2013b (i)	36	41	1	15	1	33	1	1	0.32	0		6.0		0
26) Maoz et al., 2013 (i)	37	42	1	24	3	17	0	2		0	Treatment group			3
Maoz et al., 2013 (ii)	37	42	1	27	3	17	0	2		0	Control group			3
27) McCrory et al., 2013 (i)	38	43	1	40	2	17	0	1		0				0
28) Carlson et al., 2014 (i)	39	44	1	55	1	33	1	1	0.32	0		6.0		0

Note. Dashes indicate that the sample is the same as the preceding row. Pool ID is a coding variable that indicates the effects that are pooled together into one sample.

Overall effect size of threat-related bias.

Figure 2.5 depicts the outcome of the MVP meta-analysis. A small, pooled effect of threat bias was detected ($k=44$, $N= 1407$, $d_z = 0.28$, 95% CI [0.16 0.40], $p < .001$). The probability of superiority metric (Grissom & Kim, 2005) indicates that, after controlling for individual differences, the likelihood that a randomly sampled observer will respond faster to probes following threat relative to neutral stimuli is 58% [55% 61%]. The pooled effect remained significant when any single contributing effect was removed from the model (leave-one-out analysis, all $ps < .001$). Moreover, Rosenthal's 'fail safe N ' (Rosenthal, 1991) revealed that the number of effects averaging null results required to render the pooled effect non-significant was 1125¹. Non-parametric 'trim and fill' analyses (Duval, 2005), did not suggest that any effects had been suppressed by publication bias (see also funnel plot in Figure 5b).

Substantial heterogeneity was detected ($Q(43) = 151.24$, $p < .001$). The I^2 statistic indicated that 77% of the heterogeneity between studies could not be accounted for by sampling variability, justifying the use of the random effects model. Fifty-eight percent of heterogeneity was located at the between condition level and only 19% was located at the between sample level. Moreover, a 3-level, nested model did not provide a better fit to the data than a traditional 2 level model ($LRT= 0.249$, $p = .618$)², suggesting the influence of dependency was limited. To explain this heterogeneity across threat-related biases, we examined the influence of moderators, which are summarized in Table 2.1.

¹ As others have noted, this method is likely to be biased, because the choice of adding a *zero* effect size to the observed effects neglects the possibility of unpublished studies finding *negative* effects (Begg & Berlin, 1988), which would substantially reduce the fail-safe N . Moreover, this method also does not directly model the effect of (a) the heterogeneity of the observed effects and (b) the sample sizes of the added studies, meaning the effect of adding N studies with an averaged null effect would be the same regardless of whether they had sample sizes of 10 or 10,000 (Becker, 2005). These technical issues should be considered when interpreting fail-safe N values that appear to be of a considerable size.

² In this paradigm, there were relatively few cases where participants completed more than one condition. As a result, based on the currently available data, the statistical power to distinguish the between condition and between sample is likely to be low.

Figure 2.5. Outcomes of MVP analyses.

a): Forest plot of effects from the MVP analyses. Error bars are 95% confidence intervals. Dotted red line is the pooled summary effect, shaded region is the 95% confidence interval b) Funnel plot. Dotted line is pooled effect size. Coloured contours represent p values (Black: =1, Red=.05, Yellow=.01, Blue=.001, Green= .0001, Orange= .00001, Purple=.000001). Please contact the author, or consult the published article for this figure.

Regression models with one moderator.

A summary table of the one-moderator models and plots of all main effects can be found in Appendix H.

An effect of stimulus type was detected ($Q(5) = 13.78, p = .017$), and including this moderator in the model accounted for 24.34% of the total heterogeneity among effects. There was a large bias for fearful faces ($d_z = 0.58, [0.37\ 0.78], p < .001$) but significant pooled biases were not detected for any other stimulus types (see Figure 2.6a). Fearful faces yielded larger biases than angry faces, disgust faces and word stimuli ($ps < .05$). No other significant differences between stimulus types were detected.

The distribution of SOAs between target stimulus and mask was bimodal, so we dummy coded SOAs as either long (30, 33, or 34ms) or short (12, 14, or 17ms). A main effect of this factor was detected ($Q(1) = 9.23, p = .002$) and this moderator accounted for 29.10% of the heterogeneity in effects: threat-related biases were significantly larger at longer SOAs (see Figure 2.6b).

No difference was detected between studies that did vs. did not include an objective awareness check ($Q(1) = 0.04, p = .834$). To assess the statistical power of objective awareness checks, we used the effect size index Cohen's h (the arcsine transformed difference between chance performance and a target level of above-chance performance; Cohen, 1977). To summarize power in a single metric (h^{pwr}), we calculated the largest value of h that each awareness check would be underpowered to detect (by assuming power of 79% to detect at the $\alpha = .05$ level). In other words, this analysis asks, "what is the upper limit of discrimination performance that participants could attain in the awareness check, but still be classified as 'unaware'?" The mean value of h^{pwr} was 0.37 ($SD = 0.06$) - a small-to-medium effect size. In practice, this means that, on average, it is accepted that participants are objectively unaware of stimuli if 2AFC performance is less than 68%, i.e. up to 18% above chance level.

Meta-regression detected no evidence that h_{pwr} predicted the magnitude of threat related bias ($Q(1) = 0.32, p = .856$ ($RI: b = 0.02, [-0.08\ 0.12], p = .694$, $RSI: b = 0.02, [-0.09\ 0.12], p = .754$). Thus, although awareness checks were lacking in statistical power, and threat biases are larger with long SOAs, these data do not provide direct evidence that threat related biases can be attributed to undetected deviations from chance performance. However, given the low variability in h_{pwr} values across studies (range = 0.27-0.43), and the limited number of effects that had an associated objective awareness check ($k = 26$), limited power exists to detect this potential relationship.

Trait anxiety levels were reported for 15 effects. Anxiety was entered as a continuous predictor of the corresponding threat-biases via meta-regression (Figure 2.6c), revealing that elevated anxiety is associated with larger threat bias ($b=0.03$, $R^2=37.18\%$, $p = .008$; RI : $b = 0.02$, $[0.01\ 0.04]$, $p = .016$; RSI : $b = 0.01$, $[-0.01\ 0.03]$, $p = .234$). The model indicated that threat-related biases would reduce to statistical non-significance for samples with Spielberger trait anxiety scores below 40 (STAI-T, Spielberger, Gorsuch, Lushene, Vagg, & Jacobs, 1983). However, when restricting our analyses to samples whose anxiety levels were unreported, a small threat bias was still detected ($d_z=0.25$, 95% CI $[0.12, 0.39]$, $p<.001$).

Across 22 effects, we found no evidence that stimulus size modulated effect sizes ($b=0.04$, $R^2 = 0.00$, $p = .624$, RI : $b = 0.02$, $[-0.08\ 0.12]$, $p = .694$; RSI : $b = 0.02$, $[-0.09\ 0.12]$, $p = .754$). We also found no evidence that probe response modulated effect sizes ($Q(I) = 0.14$, $p = .708$). Pooled effect sizes were of similar magnitude in the ‘where’ ($d_z = 0.26$, $[0.12\ 0.40]$, $p<.001$) and ‘what’ versions of the task ($d_z=0.32$, $[0.07\ 0.57]$, $p=.013$). Five studies split their analyses by visual field, yielding 10 effects. No effect of visual field was detected ($Q(I) = 1.93$, $p = .165$). However, when left and right visual field were analyzed separately, threat-related biases were only statistically significant for stimuli presented in the left visual field (left: $d_z=0.68$, $[0.23\ 1.15]$, $p=.003$, right: $d_z=0.23$, $[-0.21\ 0.68]$, $p=.304$).

Models with two-way interactions.

Models with two-way interactions are summarized in the tables and figures in Appendix H. An interaction was detected between stimulus type and STAI-T ($Q(2)=15.13$, $p<.001$); the threat biases elicited by all stimuli had a positive association with STAI-T, but the slope was largest for angry faces, then fearful faces and words. The interaction between awareness measure and SOA was marginally significant ($Q(I)=3.73$, $p=.054$), such that the effect of SOA on threat bias was greater when no awareness check was conducted. We did not test for higher order interactions due to low numbers of observations and empty cells in some moderator categories.

Figure 2.6. Main effects from the MVP paradigm.

Random effects models with (a) stimulus type, (b) SOA and (c) STAI-T as the sole moderator. Error bars/ shaded grey regions are the 95% confidence intervals. Size of points is inversely proportional to the standard error of the effect (larger = more precision). Please contact the author, or consult the published article for this figure.

Multiple regression models.

We used multiple regression to determine the model that optimally explained the heterogeneity in effects. Only main effects were included since interactions were either non-significant, or involved a substantially reduced number of effects. This also enhanced the interpretability of our final model. We used a backward elimination strategy, starting with a model that contained all moderators, then eliminating moderators consecutively on the basis on their p

value. Since competing models differed in terms of the number of coefficients, we used maximum likelihood estimation to compare models via likelihood ratio tests (*LRT*).

Complete effects models.

We first analyzed models where moderators were reported for all effects ($k=44$: complete effects models). These moderators (the only ones with no missing values) were stimulus type, awareness measure, SOA and probe response. The backward elimination strategy revealed that the optimal complete effects model included only stimulus type and SOA as predictors, accounting for 31.72% of the heterogeneity in effects (see Figure 2.7).

Reduced effect models.

We next evaluated the influence of additional moderators that were only reported for a subset of effects, by including only effects for which these moderator values were reported (reduced effects models). The predictors h^{pwr} ($k=26$) stimulus size ($k=22$) and visual field ($k=10$) did not significantly improve the model fit, but STAI-T ($k=15$) did ($LRT=9.73$, $p=.002$).

Model comparisons with imputed data.

Using RI to estimate the missing data, the best fitting model included stimulus type, SOA and STAI-T, and accounted for 52.50% [31.23 74.11] of the heterogeneity among effects. However, with RSI, STAI-T did not significantly improve model fit.

Sensitivity analyses.

As noted earlier, some of the effects in our model shared a sample with another effect. The outcome of our analyses may therefore be biased if the samples contributing multiple effects were unrepresentative (Greenhouse & Iyengar, 1994). We therefore constructed two new data sets using random selection procedures such that no independent sample contributed more than one effect size to the model ($k=39$). The pooled effect sizes were $d_z = 0.29$ [0.16 0.42], $p<.001$ and $d_z = 0.23$ [0.11 0.34], $p<.001$ for the first and second random selections respectively, and these datasets both resulted in the same final model (including stimulus type and SOA), following multiple regression. This further suggests that the presence of shared samples / dependency did not substantially bias our analyses.

Figure 2.7. Predicted effect sizes from the final model, plotted as a function of actual (observed) effect sizes from the MVP paradigm.

Size of points is inversely proportional to the standard error of the effect (larger = more precision). The panel shows the relative importance of each predictor (normalized contribution to R^2 across all orderings of regressors). Please contact the author, or consult the published article for this figure.

2.7.2. Summary of MVP findings.

In the MVP paradigm, we detected a small threat bias when effect sizes were pooled. A substantial amount of heterogeneity was explained by the type of stimulus, the SOA between stimulus and mask and the observers' state anxiety.

i) Threat stimuli are not equally prioritized.

The threat related bias is predominantly attributable to fearful faces. Notably, we detected no threat related bias for any other individual stimulus type. Removing fearful faces from the analysis nearly halved the magnitude of the pooled effect $d_z=0.15$ [0.05 0.24], $p=.004$.

ii) Stimulus visibility may modulate threat related biases

The data provide indirect support for the idea that stimulus visibility moderates threat related biases: effects were substantially larger when the SOA between target and mask was >30 ms. Importantly, studies using stringent signal detection criteria show that the majority of observers can reliably detect stimuli when they are masked with an SOA of ~30ms (Pessoa, Japee, Sturman, & Ungeleider, 2006). Furthermore, the interaction between awareness measure and SOA approached statistical significance - the effect of SOA on threat-bias was greater when there was no awareness measure. This further suggests that inadequate awareness measures combined with partial stimulus visibility could have contributed to the observed threat effects in several studies. Another interpretation of the moderating effect of SOA is that a brief presentation may degrade processing of a masked stimulus in general, thereby reducing effect sizes, regardless of whether this results in visibility or not. However, irrespective of whether effect sizes are moderated by awareness of the stimuli, or simply by the strength of visual signals, either possibility illustrates the methodological issues associated with using brief presentations to manipulate awareness.

iii) Awareness was not carefully measured in all studies.

Eighteen effects were not associated with any awareness check to verify the efficacy of the masking procedure and so cannot make strong categorical claims about genuinely 'unconscious' processing. Furthermore, power analyses revealed that objective awareness checks were underpowered to detect small to medium deviations from chance performance. Thus, in many cases, type II errors (failure to detect awareness) may have occurred.

iv) Threat related biases are related to, but not dependent on high anxiety levels.

Our analyses generally support the proposed link between attentional bias to masked threat and anxiety. However, the data do not strongly suggest that preconscious threat-related biases *require* high anxiety levels - a statistically significant threat-related bias was observed in samples for which levels of anxiety were not reported, but are likely to converge around healthy population means.

2.7.3. *Binocular Rivalry*

Summary of included data.

Fourteen binocular rivalry (BR) studies (comprising 788 subjects in total) were included in our analyses. These studies reported data from 22 independent samples, providing 31 effect size estimates. The coding system and summary of the included effects are displayed in Tables 2.3 and 2.4 respectively. Detailed information on each effect size and demographic information can be found in Appendix I.

Dependent measures.

For the BR paradigm, a positive value of d reflects prioritized perceptual selection of threatening stimuli over neutral stimuli. The first dependent measure we refer to as *total dominance*, which is defined as the difference between threatening and neutral stimuli in terms of the proportion of total trial time (within rivalry trials) that each was perceptually dominant (e.g. Alpers & Gerdes, 2007). The second outcome measure is *initial dominance*, which is summarized by the difference between threat and neutral stimuli in terms of the proportion of rivalry trials on which each was reported as the first percept (e.g. Gray, Adams, & Garner, 2009).

Overall effect of threat-related bias.

Figure 2.8 displays the main meta-analytic results. A medium pooled effect of threat bias was detected ($k=31$, $N=788$, $d_z=0.47$, 95% CI [0.30 0.63], $p<.001$). After controlling for individual differences, this is consistent with a 63% [58% 67%] chance that a randomly sampled observer will perceive threatening stimuli longer/ more frequently than neutral stimuli. The effect remained statistically significant when any single effect was removed (leave-one-out analyses, all $ps<.001$). Rosenthal's fail-safe N indicated that 1559 effects averaging a null result would be required to reduce the pooled effect to non-significance. Trim and fill analyses did not suggest the suppression of null effects (see funnel plot, Figure 2.8b).

Substantial heterogeneity was detected ($Q(30)=165.33$, $p<.001$). The I^2 statistic indicated that 83% of the heterogeneity between effects could not be accounted for by sampling variability. The vast majority of heterogeneity (82%) was located at the between condition level, and only 1% was located at the between sample level. Moreover, a 3 level, nested model did not provide a better fit to the data than a traditional 2 level model ($LRT=.001$, $p=.972$), suggesting virtually no influence of dependency on effect sizes. We examined the influence of several moderators to explain this heterogeneity (Table 2.3).

Table 2.3. Coding System for Individual Effects in the BR Paradigm

Moderator	Type	Values	Description / theoretical justification	Descriptive statistics	Missing cases
Stimulus type	Categorical	1=fearful face 2=angry face 3=disgust 4=conditioned neutral face (CS+) 5=grating (CS+) 6= IAPS/ pictorial	Type of threatening stimulus presented in the rivalry trial. Justification: The magnitude of threat bias in rivalry may differ as a function of the semantic/physical properties of the stimulus categories	k=31 1=13 2=7 3=4 4=2 5=2 6=3	0
Dominance measure	Dichotomous	1= total dominance 2= initial dominance	Whether the effect reflects initial dominance (which stimulus is perceived first) or total dominance (which stimulus is perceived for the longest time over the course of a trial). Justification: These are thought to partially reflect separate processes. In initial dominance, the observer's only response is the first stimulus they perceive. This initial percept thus reflects only the 'bottom up' early stages of perceptual selection. However, with total dominance, both stimuli alternate in awareness, thus it is difficult to infer whether increased perception of threat stimuli is due to unconscious processes (i.e. processes acting on a suppressed threat stimulus), or contamination from periods of conscious evaluation during dominance periods (i.e. processes occurring when the threat stimulus is visible; Carter & Cavanagh, 2007).	k=31 1=20 2=11	0
Design	Dichotomous	0 = offline 1 = online	Whether threatening and neutral images competed in rivalry in the same trial (online) or across trials via a 'baseline' stimulus (offline). Justification: Differences between concurrently presented stimuli determine rivalry dominance. It is not known whether these effects are transitive, thus the two designs may produce different effect sizes.	k= 31 0= 20 1= 11	0
Stimulus size	Continuous	1-11.5	The diameter of the threat stimulus in degrees of visual angle. Justification: Affective processing of threat images may vary over stimulus size (Codispoti & Cesarei, 2007). Moreover, larger stimuli increase the likelihood of piecemeal rivalry (mixed percepts, in which elements of both rivalling stimuli are visible; Blake, O'Shea, & Mueller, 1992).	k=26 M=6.09 SD=4.34 Range= 1-11.5	5
Rivalry trial length	Continuous	8-60	Variable representing the length of the rivalrous period. Justification: Across shorter trials, dominance proportion will be more tightly correlated with first percepts. With longer trial lengths, each stimulus will have be perceived more times, given that the number of perceptual switches are proportional to the length of the rivalrous period.	k=20 M=34.80 SD=21.82 Range=8-60	0
STAI-T	Continuous	27.9-50.5	The sample's mean trait anxiety level, as assessed by the Spielberger state-trait anxiety scale (Spielberger, Gorsuch, Lushene, Vagg & Jacobs, 1983). Justification: Anxiety is consistently linked with processing biases towards threatening stimuli (Bar Haim et al., 2007).	k=17 M=40.49 SD=5.95 Range=27.9-50.5	14
Stimulus inversion	Dichotomous	1=upright 2=inverted	Whether the threat stimulus is presented upright, or spatially inverted. Justification: Spatial inversion can impair recognition of the emotional content of stimuli, but leave low-level properties such as contrast, luminance and spatial frequency unchanged (Gray et al., 2013). Thus, if threat, or emotion were the primary determinant of the processing biases, we would expect these to be reduced, or altered when stimuli are inverted. Conversely, if low-level properties are the primary determinant, we would expect equivalent threat related biases for both the upright and inverted configurations.	k=12 1=6 2=6	0

Table 2.4. Summary of Effects Included in the BR Analyses

Study/Effect	Sample code	Condition code	Exp	N	Stim	Dom meas	Stim size	Trial len	Online	Group	STAI-T	Stim inv
1) Alpers et al., 2005 (i)	1	1	1	31	4	1	1.31	8	0		42.60	1
Alpers et al., 2005 (ii)	2	2	2	30	4	1	3.00	14	0		40.70	1
2) Alpers & Pauli, 2006	3	3	1	46	5	1	9.50	30	1		40.58	1
3) Alpers & Gerdes, 2007 (i)	4	4	1	30	1	1	9.05	15	1		39.90	1
Alpers & Gerdes, 2007 (ii)	4	5	1	/	2	1	9.05	15	1			1
Alpers & Gerdes, 2007 (iii)	5	6	2	22	2	1	1.00	8	1			1
4) Bannerman et al., 2008 (i)	6	7	1	27	1	1	11.50	60	1			1
Bannerman et al., 2008 (ii)	7	8	3	20	1	1	11.50	60	0			1
Bannerman et al., 2008 (iii)	7	11	3	/	1	1	11.50	60	0			2
5) Yoon et al., 2009 (i)	8	9	1	38	3	1	1.90	60	1			1
Yoon et al., 2009 (ii)	9	10	2	78	3	1	1.90	60	1			1
Yoon et al., 2009 (iii)	9	11	2	/	3	2	1.90		1			1
6) Gray et al., 2009 (i)	10	12	1	19	1	2	2.15		1		42.95	1
Gray et al., 2009 (ii)	10	13	1	/	2	2	2.15		1			1
Gray et al., 2009 (iii)	10	12	1	/	1	2	2.15		1			2
Gray et al., 2009 (iv)	10	13	1	/	2	2	2.15		1			2
7) Amting et al., 2010 (i)	11	14	1	16	1	2			1			1
Amting et al., 2010 (ii)	11	15	1	/	3	2			1			1
8) Anderson et al., 2011 (i)	12	16	1	57	6	1	1.50	10	0			1
Anderson et al., 2011 (ii)	13	17	2	41	6	1	1.50	10	0			1
9) Bannerman et al., 2011 (i)	14	18	1	30	2	1	3.25	60	0	Younger adults		1
Bannerman et al., 2011 (ii)	14	14	1	/	2	1	3.25	60	0	Younger adults		2
Bannerman et al., 2011 (iii)	14	19	1	30	2	1	3.25	60	0	Older adults		1
Bannerman et al., 2011 (iv)	14	15	1	/	2	1	3.25	60	0	Older adults		2
10) Ritchie et al., 2012 (i)	15	20	1	18	1	1	5.95	60	0			1
Ritchie et al., 2012 (ii)	/	16	1	5	1	1	5.95	60	0			2
11) Lerner et al., 2012 (i)	16	21	1	11	1	1		36	0			1
12) Singer et al., 2012 (i)	17	22	1	16	1	2	11.50		0	Control group	27.90	1
Singer et al., 2012 (ii)	17	23	1	/	1	1	11.50	40	0			1
Singer et al., 2012 (iii)	18	24	1	16	1	2	11.50		0	Social anxiety group	50.50	1
Singer et al., 2012 (iv)	18	25	1	/	1	1	11.50	40	0			1
Singer et al., 2012 (v)	19	26	1	14	1	2	11.50		0	Panic disorder group	43.11	1
Singer et al., 2012 (vi)	19	27	1	/	1	1	11.50	40	0			1
13) Anderson et al., 2013 (i)	20	28	1	152	2	1		10	0		38.52	1
Anderson et al., 2013 (ii)	20	29	1	/	2	2			0			1
14) Gerdes & Alpers., 2014 (i)	21	30	1	20	5	2	4.00		0	Control group	37.35	1
Gerdes & Alpers., 2014 (ii)	22	31	1	21	5	2	4.00		0	Phobic group	41.47	1

Note. An ‘I’ in the condition code indicates that the effect emanates from a spatially inverted stimulus and is thus analysed separately from the main analyses.

Figure 2.8. Outcomes of BR analyses.

a): Forest plot of effects from the BR analyses, error bars are 95% confidence intervals. Red dotted line is the pooled summary effect, shaded red region is the 95% confidence interval b) Funnel plot. Coloured contours represent p values (Black: =1, red=.05, yellow=.01, Blue=.001, Green= .0001, Orange= .00001, Purple=.000001). Please contact the author, or consult the published article for this figure.

Regression models with one moderator.

A table and figure summary of all main effects can be found in Appendix J.

Stimulus type (including fearful, angry and disgust faces, international affective picture system (IAPS, Lang, Bradley & Cuthbert, 2008) images, conditioned neutral faces and conditioned gratings) was detected as a significant moderator of threat related biases ($Q(5) = 13.24, p = .021$), accounting for 29.92% of the total heterogeneity among effects (see Figure 2.9a). Moderate to large effects for fearful faces ($d_z = 0.73, [0.50\ 0.97]$, $p < .001$), disgust faces ($d_z = 0.47, [0.11\ 0.83]$, $p = .014$) and IAPS images ($d_z = 0.66, [0.20\ 1.12]$, $p = .005$) were detected. Fearful faces and IAPS images yielded larger threat-related biases than angry faces ($ps < .050$). No other differences between stimulus type were detected.

The dominance measure (total, initial) was a marginally significant moderator of threat-related bias ($Q(1) = 3.08, p = .079$, see Figure 2.9b) accounting for 6.86% of heterogeneity. A moderate effect for total dominance was detected ($d_z = 0.57, [0.37\ 0.77]$, $p < .001$), whereas initial dominance effects were small ($d_z = 0.27, [0.00\ 0.54]$, $p = .048$).

An effect of design was also detected ($Q(1) = 4.01, p = .045$, see Figure 2.9c), accounting for 12.96% of heterogeneity, such that online designs ($d_z = 0.68, [0.41\ 0.95]$, $p < .001$) yielded larger threat-related biases than offline designs ($d_z = 0.35, [0.15\ 0.54]$, $p = .001$).

We were able to determine stimulus size for 26 effects. This predictor was marginally significant (see Figure 2.9d): larger stimuli produced larger threat related biases (*observed*: $b = 0.039, R^2 = 12.81\%, p = .058$; *RI*: $b = 0.038 [-0.004\ 0.081]$, $p = .075$; *RSI*: $b = 0.032, [-0.011\ 0.076]$, $p = .140$).

Trait anxiety levels were available for 17 effects. No effect of anxiety on the magnitude of the threat bias was detected (*observed*: $b = 0.008, R^2 = 0.00, p = .657$; *RI*: $b = 0.007, [-0.033, 0.047]$, $p = .716$; *RSI*: $b = 0.006, [-0.031, 0.043]$, $p = .744$). Across the 20 total dominance effects, no effect of trial length on the magnitude of threat-related bias was detected ($b = 0.003, R^2 = 0.00, p = .525$).

Figure 2.9. Main effects for the BR paradigm.

Random effects models with (a) stimulus type (b) dominance measure (c) design (d) or stimulus size as the sole moderator. Size of points is inversely proportional to the standard error of the effect (larger = more precision). Error bars/ shaded regions are the 95% confidence intervals. Please contact the author, or consult the published article for this figure.

Models with two-way interactions.

Plots and tables summarizing all interactions can be found in Appendix J. An interaction between trait anxiety and design was detected, such that anxiety was more strongly associated with threat bias in offline designs ($b=-0.30$, $Q(1)=4.25$, $p=.039$). The interaction between stimulus size and dominance measure was marginally significant ($b=-$

0.07, $Q(1)=3.40$, $p=.065$) such that the positive association between stimulus size and threat related bias was larger in total than initial dominance effects.

To examine the effect of spatial inversion on threat related bias, in a separate model we combined data from conditions where threat-related biases were reported for both upright and spatially inverted stimuli ($k=12$, only available for fearful and anger stimulus types). Although inversion reduced the threat bias (upright: $d_z=0.32$, inverted: $d_z=0.13$), this was not a significant main effect ($Q(1)=0.68$, $p=.409$). Critically, however, we detected an interaction between stimulus type and inversion ($Q(1)=3.93$, $p=.047$); contrasts revealed that inversion significantly reduced biases for fearful faces ($Q(1)=4.55$, $p=.033$) but not angry faces ($Q(1)=0.31$, $p=.580$).

Multiple regression models.

Complete effects models.

Our full model contained three predictors: stimulus type, dominance measure and design, since these were the only moderators with no missing values. The backward elimination strategy and likelihood ratio tests indicated that this model was significantly better than models with any of these predictors removed and was thus retained as the final model (See Figure 2.10). The model accounted for 74.70% of the heterogeneity in effects.

Reduced effects models.

Reduced effect models that included stimulus size ($k=26$), trial length ($k=20$) or state anxiety ($k=17$) were not significantly better than the full model with three predictors.

Model comparisons with imputed data.

After using both RI and RSI to estimate the missing values for stimulus size, trial length and state anxiety, the best fitting model was unchanged.

Sensitivity analyses.

We constructed two new data sets using random selection procedures such that no sample contributed more than one effect size to the model ($k=22$). The pooled effect sizes were $d_z = 0.60$, $[0.41\ 0.78]$, $p < .001$ and $d_z = 0.56$, $[0.37\ 0.75]$, $p < .001$ for the first and second random selections respectively, and these resampled datasets resulted in the same final model following multiple regression. This suggests that the presence of shared samples did not substantially bias our analyses.

Figure 2.10. Predicted values from the final model, plotted as a function of actual (observed) effect sizes from the BR paradigm.

Size of points is inversely proportional to the standard error of the effect (larger = more precision). The panel shows the relative importance of each predictor (normalized contribution to R^2 across all orderings of regressors). Please contact the author, or consult the published article for this figure.

2.7.4. Summary of BR findings.

For the BR paradigm, we detected a moderately-sized overall threat bias that was larger than that found with the MVP paradigm. A model containing the type of stimulus, the dominance measure and design as moderators provided a good fit to the data.

i) Stimulus type.

Similarly to the MVP paradigm, the size of the threat bias depended on the type of stimulus; in both the MVP and BR paradigms, fearful faces produced a large and highly reliable effect. Strikingly, in both MVP and BR paradigms, the effect produced by angry faces was significantly smaller, and not significantly different from zero.

Fearful faces were the most widely used threat stimulus in the BR paradigm (contributing 42% of our analysed effects), whilst some other stimulus categories (e.g. IAPS) were sparsely represented, limiting the precision of their effect size estimates.

ii) Effects are smaller for initial than total dominance.

The initial dominance is thought to be a more objective measure of the unconscious perceptual selection of stimuli in the competition for awareness than total dominance (Berry, 1969; Gray et al., 2009; Ooi & He, 1999). Because total dominance is quantified from alternating perception of threatening and neutral stimuli, both conscious (during dominance of threat) and unconscious processes (during suppression of threat) could contribute to these effects. Our analyses revealed that total dominance effects were larger than initial dominance, suggesting that threat related biases in BR are strongly modulated by conscious processing.

iii) Stimulus size.

There is good evidence that rivalry occurs within spatially localized regions (e.g. Kovacs et al., 1996). This can lead to piecemeal rivalry for larger stimuli: perception is not exclusively of one stimulus or the other, but a mixed patchwork of the two. Piecemeal rivalry may actually become more prevalent than global rivalry when stimuli are large; the optimum stimulus size for coherent, whole-image rivalry is less than 1 degree of visual angle (DVA; Blake et al., 1992). Generally, the stimuli presented in the BR studies were considerably larger than this (the mean stimulus size was 6 DVA in diameter), suggesting that piecemeal rivalry may have occurred frequently. This, in turn, increases the risk of response biases and criterion effects, since these are more likely to come into play during the ambiguous, mixed perceptual states in piecemeal rivalry. For instance, a threatening stimulus may be reported as the dominant percept because it has more behavioral relevance and is more noticeable to an observer, when in fact local regions of both threat and neutral stimuli are visible. Our data provide some support for this - there was a marginally significant association between stimulus size and threat related bias. Moreover, stimulus size was more predictive of threat related biases in total dominance tasks than initial dominance tasks (stimulus size x dominance measure interaction). This further suggests

that mixed perception may play a role; mixed perception often occurs at the time of perceptual switches, which are lacking in paradigms that only measure the first percept (i.e. initial dominance).

iv) Experimental design

The design (offline vs. online) was predictive of threat related biases: online designs yielded larger effects than offline designs. Importantly, in online designs, when competing images (e.g. a fearful face and neutral face) are presented simultaneously, they may not satisfy a necessary condition of binocular rivalry: that the images presented to each eye are sufficiently different. In particular, if the faces are matched in terms of identity and, more importantly, orientation (as in, for example, Alpers & Gerdes, 2007, Amting et al., 2010) they may be binocularly fused, with the resultant percept differing from neutral. Thus, fusion may prompt an observer to report that a threatening stimulus is dominant, when in fact no rivalry occurred at all. Indeed, one experiment with an online design (Bannerman et al., 2008, Experiment 2a-which we excluded from our analyses) reported that when aligned fearful and neutral faces were presented dichoptically, observers did not experience any rivalry.

v) Binocular rivalry and anxiety.

Some studies included in our analyses have suggested a positive association between anxiety and threat bias in BR (Gray et al., 2009; Singer et al., 2012). This was not consistent across all studies that included this measure, and unlike the MVP analyses, our meta-regression did not detect a relationship between trait anxiety and threat dominance overall. Some studies reported no difference in threat bias between anxious and non-anxious populations (Alpers & Gerdes, 2007; Anderson et al., 2013). Another showed larger threat biases for anxious populations, relative to controls in initial dominance, but effects in both directions in total dominance, depending on the specific diagnosis (Singer et al., 2012). Although our analyses detected no main effect of anxiety, the relationship between anxiety and rivalry may be a function of the dominance measure, stimulus type, and specific diagnosis. More data will be needed to clarify this relationship.

2.7.5. Breaking continuous flash suppression

Summary of included data.

Fourteen bCFS studies (comprising 501 subjects) were included in the analyses. These studies reported data from 18 independent samples, providing 27 effect size estimates. The coding system and summary of the included effects are displayed in Tables 2.5 and 2.6 respectively. Details of how each effect size was computed can be found in Appendix K.

Dependent measures.

For the bCFS paradigm, a positive value of d indicates prioritized detection of threatening stimuli from CFS. In almost all cases, the dependent measure was response time, where a positive value indicates faster detection of threatening stimuli (e.g. Yang et al., 2007). In other cases, a positive value indicates more accurate localization of threat stimuli following shorter, fixed duration CFS trials (indexed by accuracy in forced choice responses- see Hedger, Adams, & Garner, 2015; Oliver, Mao, & Mitchell, 2014).

Overall effect of threat-related bias.

Figure 2.11 displays the main meta-analytic results for the bCFS paradigm. Across all effects, the pooled effect size was small, negative and non-significant ($k=27$, $N=501$, $d_z = -0.14$, $[-0.45\ 0.17]$, $p = .376$). After controlling for individual differences, this is consistent with a 46% chance that a randomly sampled observer will perceive threatening stimuli faster/ more efficiently than neutral stimuli. Trim and fill analyses did not suggest the suppression of any unpublished effects (see funnel plot, Figure 2.11b). The test for heterogeneity was significant ($Q(26) = 252.56$, $p < .001$ and the I^2 statistic indicated that nearly all the heterogeneity across effects (94%) was due to factors other than sampling variability. For heterogeneity, 33% was located at the between condition level and 61% was located at the between sample level. However, a 3 level nested model did not provide a significantly better fit than a traditional two level model ($LRT = 2.61$, $p = .187$). We examined the influence of several moderators to explain this heterogeneity (Table 2.5).

Table 2.5. Coding System for Individual Effects in the bCFS Paradigm

Moderator	Type	Values	Description / theoretical justification	Descriptive statistics	Missing cases
Stimulus type	Categorical	1= Fearful face	Type of threatening stimulus presented in the CFS trial. Justification: The magnitude of threat bias in CFS may differ as a function of the semantic/physical properties of the stimulus categories	k=27	0
		2= Angry face		1 = 10	
		3= Disgust face		2 = 3	
		4= Dominant face		3 = 3	
		5= Untrustworthy face		4 = 4	
		6 = Negative word		5 = 4	
		7 = Pictorial/ IAPS		6 = 2	
Stimulus size	Continuous	Range = 1.15- 5.20	The diameter of the threat stimulus in degrees of visual angle. Justification: Affective processing of threat images may vary over stimulus size (Codispoti & Cesarei, 2007). Moreover, a large stimulus size increases the likelihood of piecemeal rivalry (mixed percepts, in which elements of both stimulus and mask are visible; Blake et al., 1992).	k = 22	7
				M = 3.03	
				SD = 0.92	
				Range = 1.7 – 5.20	
Awareness measure	Dichotomous	1= Response time	The measure by which an observer's detection of stimuli from CFS is identified Justification: Response times may comprise multiple components- a motor component (i.e. the time taken to press a button), a perceptual component (the time it takes for a stimulus to reach awareness) and a decisional component (the time it takes to use the available information to determine that the stimulus is visible). Un-speeded forced-choice localization tasks are less affected by the motor and decisional components, since response latencies are not diagnostic.	k=27	0
		2 = Localization accuracy		1 = 24 2 = 3	
Stimulus inversion	Dichotomous	1= Upright	Whether the threat stimulus is presented upright, or spatially inverted.	k= 18	0
		2=Inverted	Justification: Spatial inversion can impair recognition of the emotional content of stimuli, but leave low-level characteristics such as contrast, luminance and spatial frequency unchanged (Gray et al., 2013). The logic of this manipulation is that if low-level properties were the cause of a threat bias, one might expect a similar sized bias for upright and inverted stimuli. Conversely, if threat sensitive processes were the cause of a threat bias, we would expect a threat bias in an upright configuration, but not in an inverted configuration.	1 = 9 2 = 9	

Table 2.6. Summary of Effects Included in the bCFS Analyses

Study/Effect	Sample code	Condition code	Exp	N	Stim	Stim Size	Aw meas	Group	Stim inv
1) Yang et al., 2007 (i)	1	1	1	12	1	1.9	1	Control group	1
Yang et al., 2007 (ii)	1	11	1	/	1	1.9	1		2
Yang et al., 2007 (iii)	2	2	2	12	1	1.9	1		1
Yang et al., 2007 (iv)	2	12	2	/	1	1.9	1		2
2) Sterzer et al., 2011 (i)	3	3	1	20	1	2.0	1		1
3) Sylvers et al., 2011 (i)	4	4	1	87	1	3.4	1		1
Sylvers et al., 2011 (ii)	4	5	1	/	3	3.4	1		1
4) Yang & Yeh, 2011 (i)	5	6	1	12	6	2.0	1		1
Yang & Yeh, 2011 (ii)	5	13	1	/	6	2.0	1		2
Yang & Yeh, 2011 (iii)	6	7	2	12	6		1		1
Yang & Yeh, 2011 (iv)	6	14	2	/	6		1		2
5) Chen & Yeh., 2012 (i)	7	8	1	30	1	5.2	1	Control group	1
6) Stein & Sterzer, 2012 (i)	8	9	1	16	2	2.0	1		1
7) Stewart et al., 2012 (i)	9	10	1	23	4	3.4	1		1
Stewart et al., 2012 (ii)	9	11	1	/	5	3.4	1		1
Stewart et al., 2012 (iii)	10	12	2	21	4	3.4	1		1
Stewart et al., 2012 (iv)	10	13	2	/	5	3.4	1		1
Stewart et al., 2012 (v)	11	14	3	28	4	3.4	1		1
Stewart et al., 2012 (vi)	11	15	3	/	5	3.4	1		1
8) Gray et al., 2013 (i)	12	16	3	41	2	2.5	1		1
Gray et al., 2013 (ii)	12	17	3	/	1	2.5	1		1
Gray et al., 2013 (iii)	12	15	3	/	2	2.5	1	Control Group	2
Gray et al., 2013 (iv)	12	16	3	/	1	2.5	1		2
9) Stein et al., 2014a (i)	13	18	1	12	1	3.5	1		1
Stein et al., 2014a (ii)	13	17	1	/	1	3.5	1		2
10) Capita et al., 2014 (i)	14	19	1	46	1	1.7	1		1
11) Oliver et al., 2014 (i)	15	20	1	40	1		2		1
Oliver et al., 2014 (ii)	15	21	1	/	3		2		1
Oliver et al., 2014 (iii)	11	19	2	39	1		2		2
Oliver et al., 2014 (iv)	12	110	2	/	3		2		2
12) Getov et al., 2014 (i)	16	22	1	36	4		1		1
Getov et al., 2014 (ii)	16	23	1	/	5		1		1
13) Jusyte et al., 2015 (i)	17	24	1	24	1	2.7	1	Control Group	1
Jusyte et al., 2015 (ii)	17	25	1	/	3	2.7	1		1
Jusyte et al., 2015 (iii)	17	26	1	/	2	2.7	1		1
14) Hedger et al., 2015a (i)	18	27	1	29	7	5.2	2		1

Figure 2.11. Outcomes of bCFS analyses.

(a) Forest plot of effects from the bCFS analyses, error bars are 95% confidence intervals. Red dotted line is the pooled summary effect, shaded red region is the 95% confidence interval b) Funnel plot. Coloured contours represent p values (Black: =1, red=.05, yellow=.01, Blue=.001, Green= .0001, Orange= .00001, Purple=.000001). Please contact the author, or consult the published article for this figure.

Regression models with one moderator.

A table and figure summary of all main effects can be found in Appendix I. Stimulus type (including fearful, angry and disgust, dominant and untrustworthy faces, negative words and IAPS images) was detected as a moderator of threat related biases ($Q(6) = 41.32, p < .001$), accounting for 65.38% of the total heterogeneity among effects (see Figure 2.12a). A moderate positive bias was detected for fearful faces ($d_z = 0.49, [0.17, 0.82], p < .001$), whereas large, negative biases were detected for dominant faces ($d_z = -0.96, [-1.47, -0.44], p < .001$), untrustworthy faces ($d_z = -0.68, [-1.18, -0.17], p = .008$), and negative words ($d_z = -1.69, [-2.58, -0.79], p < .001$). Fearful faces yielded larger biases than dominant and untrustworthy faces and negative words (all $ps < .001$). The pooled effect for negative words was significantly smaller than for angry faces, disgust faces and IAPS images (all $ps < .002$).

No effect of stimulus size (*observed*: $b = -0.180, R^2 = 0.00, p = .379$; *RI*: $b = 0.007, [-0.033, 0.047], p = .716$; *RSI*: $b = 0.006, [-0.031, 0.043], p = .744$) or awareness measure were detected ($Q(1) = 0.192, p = .661$).

Models with two-way interactions.

Plots and tables summarizing all interactions can be found in Appendix I. No interactions involving stimulus type, stimulus size or awareness measure were detected. To observe the effect of spatial inversion on threat related bias, in a separate model, we combined data from conditions where threat-related biases were reported for both upright and spatially inverted stimuli ($k = 18$). Inverted stimuli actually yielded larger threat related biases (upright: $d_z = 0.15$, inverted: $d_z = 0.34$) although the main effect of inversion was not significant ($Q(1) = 0.816, p = .367$). Critically, we detected an interaction between stimulus type and inversion ($Q(1) = 12.811, p = .005$, see Figure 2.12b), i.e. inversion had a differential effect on threat-related bias depending on the stimulus type. Contrasts revealed that threat-related biases for fear ($p = .837$) and anger ($p = .372$) faces did not differ significantly between upright and inverted configurations, but inversion was associated with significantly larger effect sizes for disgust faces ($p = .044$) and negative words ($p < .001$).

Models with multiple moderators.

Complete effects models.

Our full model contained two predictors: stimulus type and awareness measure, since these were the only moderators with no missing values. The backward elimination strategy eliminated awareness measure, meaning that the best fitting model included only stimulus type, as described above (Figure 2.12a).

Figure 2.12. Effects from the bCFS paradigm.

a) Random effects model with stimulus type as the sole moderator. b) Model depicting the interaction between stimulus type and stimulus inversion. Size of points is inversely proportional to the standard error of the effect (larger = more precision). Error bars are the 95% confidence intervals. Please contact the author, or consult the published article for this figure.

Reduced effects models.

A reduced effect model that included stimulus size ($k=22$) did not significantly improve the model fit.

Model comparisons with imputed data.

Using both imputation methods, the best-fitting model remained unchanged.

Sensitivity analyses.

We constructed two new data sets using random selection procedures such that no sample contributed more than one effect size to the model ($k=22$). The pooled effect sizes were $d_z=0.09$, $[-0.27\ 0.47]$, $p=.618$ and $d_z = 0.06$, $[-0.30\ 0.43]$, $p=.725$ for the first and second random selections respectively, and these resampled datasets resulted in the same final model (including just stimulus type), following multiple regression. This suggests that the presence of shared samples did not substantially bias our analyses.

2.7.6. Summary of bCFS findings.

In our analysis of the bCFS literature, we estimated a very small, negative, non-significant effect of threat related bias. Although many studies provided significant effects, there was substantial heterogeneity, with many effect sizes being strongly positive or strongly negative.

i) Evidence for reversed threat biases.

Some aspects of the data were similar to the MVP and BR paradigms. Again, fearful faces yielded threat-biases that were substantially larger than other stimulus categories. However, unlike these paradigms, a striking discrepancy was observed in that we found evidence for substantial reversed biases for some threat stimuli: negative word stimuli, and dominant and untrustworthy faces were *slower* to break suppression than their neutral counterparts. In addition to conflicting with the data from other paradigms, these findings conflict with the basic notion that unconscious threat processing is concerned with *expediting* the processing and perception of threatening stimuli to promote survival.

ii) Low-level confounds may explain some threat-related biases.

Contrary to our findings for the BR paradigm, we found that biases for fear and angry faces were indistinguishable between upright and inverted configurations. In fact, the pooled effect was slightly *larger* for inverted configurations. Given that inversion reduces the recognizable threat content of facial expressions, but maintains their low-level

image properties (Gray et al., 2013), this provides good evidence that detection advantages/disadvantages for these stimuli may be mainly attributed to low-level properties such as contrast and spatial frequency content - factors known to robustly affect rivalry dominance (Baker & Graf, 2009), rather than threat sensitive processes.

2.8. Discussion.

2.8.1. Summary of outcomes.

The primary goal of our meta-analysis was to examine the extent to which unconsciously presented threatening stimuli are prioritized in visual processing, relative to neutral stimuli. Our analyses revealed evidence for a small pooled threat-prioritization effect in the MVP paradigm, a medium effect in the BR paradigm and inconsistent effects in the bCFS paradigm.

2.8.2. Differences between paradigms.

The three paradigms we reviewed did not only yield pooled effects of different magnitude, they were also moderated by different variables, affirming our decision to analyze them separately. This is perhaps unsurprising, because the three paradigms differ with respect to how they disrupt normal visual processing (Breitmeyer, 2015). Research indicates that brief, masked presentations interfere with awareness by impeding the temporal integration of neural responses to successive stimuli (Kovacs, Vogels, & Orban, 1995). In contrast, BR is a complex multi-stage phenomenon, comprising of low-level, interocular inhibitory components (Tong & Engel, 2001; Tong, Meng, & Blake, 2006) and higher-level effects that increase the depth of suppression along the ventral processing stream (Nguyen, Freeman, & Alais, 2003). The strength of suppression induced by CFS is also more potent than BR, as demonstrated by sensitivity measurements (Yang & Blake, 2012). Moreover, masking and CFS may differ with respect to how they attenuate neural responses in the dorsal and ventral processing streams (Almeida, Mahon, Nakayama, & Caramazza, 2008). For instance, there is an ongoing discussion about whether CFS spares processing via the dorsal ‘vision for action’ pathway relative to masking (Hebart & Hesselman, 2012). Such a difference might provide an a priori expectation that stimuli

presented under CFS are more likely to elicit behavioral responses. These different suppression mechanisms should therefore be expected to differ with respect to how they restrict the neural representation of threat-relevant stimuli.

It is also important to consider that the MVP paradigm may reflect a different visual selection process to BR and bCFS. During BR, awareness alternates between two retinally co-located images presented to the two eyes. In contrast, in attentional cuing tasks, attending to a stimulus at one location impairs discrimination of a stimulus at another location, but does not, in itself, cause it to disappear from awareness. However, although selective attention and interocular suppression clearly have different perceptual consequences, they may engage common competitive mechanisms. For instance, Mitchell, Stoner and Reynolds (2004) found that cuing attention to a surface engaged in rivalry enhanced its dominance. Similarly, Ooi and He (1999) found that a stimulus is more likely to become dominant if accompanied by a salient ‘pop out’ cue. Human brain imaging also shows that the activation of regions involved in attentional switching and perceptual switching in BR are similar (Knapen, Brascamp, Pearson, van Ee, & Blake, 2011). Indeed, behavioral evidence shows that in the absence of attention there are no variations in consciousness that define binocular rivalry (Brascamp & Blake, 2012). Thus, despite apparent differences, biases observed in BR and MVP paradigms may be governed by a similar neural competition process that is prompted by rival stimulus representations. In this context, is notable that when controlling for the differences in stimuli that have been used in each paradigm, the data are broadly consistent. When considering only those stimuli that have been used in all three paradigms (fear faces, angry faces, disgust faces and IAPS images), effect sizes are not moderated by paradigm ($Q(2) = 2.37, p = .306$) and there is no interaction between stimulus and paradigm ($Q(6) = 2.48, p = .870$).

2.8.3. Which threat stimuli receive prioritized processing?

One other interesting finding was the existence of strong *reversed* biases for some threatening stimuli in the bCFS paradigm: neutral stimuli were consistently prioritized over negative words, untrustworthy and dominant faces. These findings conflict with the basic notion that when encountering threat, its privileged processing is beneficial (Nesse, 1999). Stewart and colleagues (2012) propose a framework to account for these discrepancies by suggesting that indirect threats (e.g. fearful faces) may induce fight or flight responses and

heighten cortical arousal to reduce suppression of threat stimuli, whereas direct threats (e.g. angry faces, dominant faces) may also induce passive responses, characterized by ‘freezing’ and reduced cortical arousal, which may prolong suppression of threat. However, this framework cannot accommodate the reversed bias for negative words, which are not direct threats, nor can it accommodate for the lack of reversed biases for directly threatening stimuli in the BR and MVP paradigm. At any rate, whereas freezing behaviors and physiological changes have adaptive properties in the context of threat (reducing detection by predators, conserving energy) these should not be conflated with *actively suppressing the perception* of threatening stimuli, which seems maladaptive. In fact, freeze responses in many mammals are associated with *hypervigilance* to threat cues that prime a subsequent fight or flight reaction (Campbell, Wood, & McBride, 1997).

One finding that was consistent across all three paradigms was that fearful faces elicited the largest, most reliable threat related biases (MVP: $d_z = 0.56$, BR: $d_z = 0.58$, bCFS: $d_z = 0.49$). In fact, it is worth noting that removing fearful faces from the analysis substantially reduced the pooled effect size in each paradigm (MVP: 0.28 to 0.15, BR: 0.47 to 0.31, bCFS: -0.04 to -0.50). This sensitivity to fear is consistent with a large body of neuroimaging literature which has demonstrated that fear faces elicit responses in threat sensitive brain regions, even when suppressed by masking (Whalen et al., 2004), BR (Pasely, Mayes, & Schultz, 2004) and CFS (Jiang & He, 2006). Another commonality worth noting is that in all three paradigms, angry faces produced substantially smaller, non-significant, and even negative effects (MVP: 0.11, BR: 0.08, bCFS: -0.07). This is somewhat surprising, given that angry faces signal a *direct* threat to an observer (‘I am angry’), whereas fearful faces only *indicate* the presence of a threat (‘I am afraid’). It is hard to explain why an effective threat detection system would have the capacity to prioritize an *indicator* of threat in the environment (a fearful face), without similar sensitivity to stimuli that are more *directly* threatening (an angry face). One possibility is that fearful faces are more salient on a purely sensory level, and that this is a better predictor of their enhanced processing than their effect on threat sensitive processes (Gray et al., 2013; Lee, Susskind, & Anderson, 2013). We discuss this possibility in the following section.

2.8.4. Low level confounds.

In our bCFS analyses, we found that biases for some stimulus categories (e.g. fearful faces) did not differ between upright and inverted configurations. Recent reports have shown that inverted facial expressions, while retaining luminance, contrast and spatial frequency profile, have vastly reduced recognizable emotional content, according to signal detection and implicit measures (Gray et al., 2013) and also valence, arousal and dominance ratings (Hedger, Adams, & Garner, 2015b). Therefore, the fact that the detection advantage for fearful over neutral faces is equivalent in magnitude between upright and inverted configurations suggests that simple low-level variability between expressions may drive this effect. This more parsimonious explanation negates the need to invoke unconscious threat sensitive processes. It is notable that very few MVP studies have attempted to control for low-level stimulus properties (the exceptions being Carlson & Reinke, 2008; and Fox, 2002). This is important, since if stimuli differ on some other dimension other than their perceived threat, it cannot unequivocally be claimed that perceived threat is the cause of the processing bias unless adequate controls are implemented.

Recently, it has been reported that the prioritized detection of fearful faces from backward masking and CFS is poorly explained by perceived threat (indexed by valence arousal and dominance ratings) and is better explained by low-level stimulus characteristics - the distribution of luminance contrast across spatial frequency in relation to the human contrast sensitivity function (Hedger et al., 2015b). In particular, several authors have noted that the increased luminance contrast associated with the greater exposure of iris and scleral field in the fear expression may be a good predictor of their prioritized detection over neutral faces (Gray et al., 2013; Hedger et al., 2015b; Lee, Susskind, & Anderson, 2013). Notably, this suggests a purely sensory detection advantage that can occur independently of threat, or emotion sensitive processes. Given that i) fearful faces were the most commonly used stimuli in conditions contributing to our analyses and ii) these conditions contributed the largest effect sizes to the pooled estimate, this is a non-trivial issue.

2.8.5. Assessment of awareness and response criteria.

In the MVP analyses, we found evidence that awareness moderates threat related biases: effects were substantially larger when the SOA between target and mask was

increased to > 30 ms. This is particularly important, given evidence that observers can reliably detect stimuli that are presented for this duration when stringent, signal detection criteria are used to assess awareness (Pessoa, Japee, & Ungerleider, 2005; Pessoa, Japee, Sturman, & Ungerleider, 2006). Furthermore, many MVP studies did not include any explicit awareness check to verify the efficacy of the masking procedure, which substantially limits the validity of strong conclusions about ‘unconscious’ processing on the basis of these observations.

Related, but separable concerns are applicable to the assessment of awareness in the BR and bCFS paradigms. In BR, the perceptual switches between stimuli are not always well defined and discrete, making it difficult to reliably measure which stimulus is dominant at any one time. Although some studies have included a ‘mixed-percept’ response option to address this issue (Alpers & Gerdes, 2007; Lerner et al., 2012), the boundary between perception of one image and another in rivalry is often graded and temporally uncertain (Knapen et al., 2011). Thus, regardless of the available response options, perceptual reports are still heavily reliant on an observer’s individual criteria in classifying when one image is (primarily) dominant or the percept is mixed (Pessoa, 2005). It is possible, for example, that response biases could inflate effect sizes, if a threatening stimulus is reported when elements of both threatening and neutral images are visible.

Similarly, response times in bCFS tasks reflect both a ‘pure’ suppression duration, during which none of the target stimulus is visible, but can also reflect the time taken, and criterion used, to report that a stimulus has become visible. This concern is particularly pertinent when one considers that several studies included in our analysis did not include a non-CFS control condition to verify that there were no inherent differences in detectability of threatening vs. non-threatening stimuli under suprathreshold conditions (Capitao et al., 2014; Chen & Yeh, 2012; Gray et al., 2013; Justyte et al., 2015; Stein et al., 2013; Sylvers et al., 2011). Furthermore, even in the cases where such a control condition has been included, this typically consists of presenting identical stimuli to both eyes (Sterzer et al., 2011; Stewart et al., 2012), which may not be perceptually comparable. For instance, response times are highly variable in a bCFS task, due to the stochastic temporal dynamics of BR (Lehky, 1995), whereas in a non-CFS control condition, the appearance of a binocularly presented target whose contrast is linearly increased is much more easily anticipated (Stein, Hebart, & Sterzer, 2011). Since the target stimulus and mask are simply superimposed in control tasks, there is also the absence of partial stimulus visibility that can occur during perceptual switches in binocular rivalry, including CFS. Thus, such

control tasks are not perceptually comparable and may not be equipped to rule out the influence of response biases. We must be cautious, therefore, in interpreting differential response times in *bCFS* studies as solely reflecting unconscious processing.

2.8.6. Threat-related biases and anxiety.

Evidence for a relationship between anxiety and threat-related bias varied across paradigms. We found strong evidence for a relationship between trait anxiety and threat bias in the MVP paradigm, but evidence for this association in the BR paradigm was more mixed, with both affirmative (Gray et al., 2009, Singer et al., 2012) and null findings (Alpers & Gerdes, 2007; Anderson et al., 2013). In the *bCFS* paradigm, we identified only one study that included an anxiety measure, which prevented meta-analytic examination. Measurement of threat biases in anxious populations could be complicated by the fact that anxious observers are less capable of discriminating between threatening and neutral signals (Lissek et al., 2009) and often interpret ambiguous stimuli as threatening (Clark & McManus, 2002) as a consequence of a lowered threshold for perceiving threat (Mogg & Bradley, 1998). Indeed, Lee, Kang, Kim, and An (2008) note that neutral faces may provide an inappropriate baseline in studies of emotion processing, since they may be evaluated as negative depending on the experimental context and the psychological state of the observer.

Another possibility is that anxiety is only associated with enhanced threat biases when stimulus presentation is brief; our BR analyses revealed that the association between anxiety and threat bias was stronger in the initial dominance measure (although the anxiety x dominance measure interaction did not reach significance). In keeping with this, evidence from the visual probe paradigm has suggested an anxiety-enhanced bias towards threat at short presentations, but this is less reliable at longer stimulus presentations (Mogg & Bradley, 2006; Mogg, Philppot, & Bradley, 2004). Moreover, in clinically anxious populations, threat related biases have been found to be larger in subliminal than supraliminal versions of the emotional Stroop task (Bar Haim et al., 2007). Eye movement data also indicate that threat biases in anxiety are typically observed during the initial phases of stimulus presentation (e.g. first fixations; Calvo & Avero, 2005; Mogg, Garner, & Bradley, 2007). These observations are broadly consistent with cognitive models of

anxiety, which posit that selective attention for threat is mediated by mechanisms operating early in information processing (Williams, Watts, Macleod, & Mathews, 1997).

2.8.7. Future directions.

Our review reveals a number of topics that, in our view, warrant further investigation.

Dissociating awareness, stimulus degradation and suppression.

Firstly, there is a need for a more refined, systematic investigation of the representation of subliminal stimuli. To optimally study unconscious threat processing, a paradigm should manipulate awareness and not any other aspect of visual processing. However, all known methods for rendering stimuli invisible do so by making them drastically different from a consciously viewed counterpart. Thus although suppression methods appear to allow experimenters to conveniently ‘switch awareness off’, they likely do so by attenuating the gain of neural responses and degrading the strength of visual signals, relative to consciously viewed stimuli (Yuval-Greenberg & Heeger, 2013). It has been argued, for instance, that binocular rivalry may not be optimally suited for studying visual consciousness, since it may have unique neural mechanisms that do not generalize to other stimulus conditions and perceptual phenomena (Blake, Brascamp, & Heeger, 2014). Therefore, it remains possible that other paradigms may eliminate awareness, but spare visual processing to the extent that threat responses remain effective. Testing a range of suppression paradigms that rely on different mechanisms, will allow more reliable dissociation of null effects resulting from the genuine absence of unconscious threat-sensitive process from those resulting from methodological limitations (e.g. Faivre, Berthet, & Koudier, 2012).

On a related note, the extent to which ‘dominance’ and ‘suppression’ in BR and CFS are functionally the same as ‘awareness’ and ‘unawareness’ remains an empirical question. Under some conditions, participants may retain some residual sensitivity to, or phenomenal awareness of stimuli in the suppression phase. For instance, colours of objects suppressed under rivalry can nonetheless appear as a diffuse “cloud” superimposed on the dominant image (Hong & Blake, 2009) and suppressed, drifting gratings can still give an impression of movement, when only the dominant image is visible (Zabood, Lee, & Blake, 2011). Moreover, when two flickering forms engage in rivalry, they can be temporally

integrated into ‘beats’, despite observers only being consciously aware of one form (Carlson & He, 2000). These examples of ‘stimulus fractionation’ are widespread in the rivalry literature and suggest that fusion and rivalry can co-occur, such that some aspects of a stimulus may be suppressed (form) but others may be fused (colour, motion, temporal information).

Are threat stimuli comparable?

A common criticism of meta analysis is that researchers combine different types of studies in a single analysis (i.e. a problem of “apples and oranges”, Bornstein, 2009). For instance, combining the data from individual studies that use either fearful faces or dominant faces (as in the case of the bCFS analyses) yields a threat related bias that is near zero, but this does not adequately characterize the effect elicited by each stimulus. However, meta analyses allow us to quantify these differences despite the fact that these stimuli were not directly compared within the same empirical study. Given the substantial heterogeneity explained by stimulus type in all three paradigms, we should question the extent to which all stimuli defined as threatening are truly comparable.

Though both fearful and dominant faces may be threatening, fearful faces may be perceived as a salient threat of physical harm in the nearby environment, whereas dominant faces may be perceived as more nuanced threat to social status. At the behavioural level, reacting to a fearful face may require a fast behavioural response, whereas responding to a dominant face may promote submissive withdrawal and behavioural adjustments related to longer term risk assessment. Indeed, at the neural level, researchers have differentiated between systems for responding to ‘potential threat’ and ‘imminent danger’ (Fiddick, 2011). As a result, a more refined characterization of threatening stimuli is required in future research. This could include a number of important dimensions that may modulate the threat response, such as the proximity (Mobbs et al., 2007), predictability (Whalen et al., 2007) or directness (Adams et al., 2011) of the threat and the psychological state of the observer (Bishop, 2007).

A related recommendation is that experimental methods should routinely test the crucial possibility that stimuli intended to be threatening or neutral may simply not be perceived as such by participants. This problem may arise because self-report ratings can be influenced by distortions such as social norms and the investigators’ expectations (Daglish & Power, 1999). Indeed, implicit measures of valence have been shown to be inconsistent with self report measures and may reveal that observers judge both ‘neutral’ and ‘threatening’ categories as being broadly similar in valence (e.g. Lee et al., 2008). As

implicit measures are relatively immune to response biases, they could be used in place of, or in conjunction with self report measures.

What kind of awareness matters?

The research literature that we have reviewed consists of paradigms that disrupt normal visual processing so that awareness can be studied. It has been argued that although this type of awareness is interesting, it is not particularly relevant to understanding the impact of threatening stimuli on behavior and clinical conditions such as anxiety (Pessoa, 2013). The primary reason that is often cited for this position is that “subliminal stimuli do not occur naturally” (Bargh & Morsella, 2008, p. 78). Whilst this claim seems unfalsifiable, it is clear that we do not, outside of the lab, often encounter a 10Hz stimulus presented to just one eye, or isolated faces images that are masked after only 17 milliseconds. For this reason, Bargh and Morsella propose that studying unawareness of *the influence of a stimulus* is more important to understanding human behavior than the *unawareness of a stimulus itself*. Although a reasonable concern, it is also true that stimuli can also be rendered invisible in more typical circumstances than those induced by BR and masking. For instance, the majority of traffic accidents can be attributed to inattention and forms of perceptual blindness (Chun & Marois, 2002). Visual crowding and motion-induced blindness are other instances where stimuli are rendered invisible, but under conditions that are likely to occur frequently in cluttered and dynamic natural scenes (Bonneh, Cooperman, & Sagi, 2001; Koudier, Berthet, & Faivre, 2011). Using these paradigms may prove informative and allow conclusions to more readily be generalized to typical viewing conditions.

Interocular suppression and anxiety

Although there have been studies into the efficacy of pharmacological and cognitive interventions to modulate threat related biases in anxious individuals using the MVP paradigm (Maoz et al., 2013; Murphy et al., 2008), this has not been attempted using BR and bCFS. This is somewhat surprising, since there is good evidence that perceptual switches in binocular rivalry are linked to the balance of inhibitory neurotransmitters (van Loon et al., 2013), prefrontal cortex activity (Amting et al., 2010) and attentional control (Carter et al., 2005; Paffen & Alais, 2011); all of which have been implicated in maladaptive perceptual biases in anxiety and considered therapeutic treatment targets (Bishop, 2009; Eysenck, Derakshan, Santos, & Calvo, 2007). There are therefore, potential gains from applying well controlled versions of these tasks to investigate maladaptive biases in threat processing in anxiety. For instance, early investigations with the bCFS

paradigm appear to show some sensitivity to self-reported anxiety (Capitao et al., 2014) and other social trait characteristics (Stewart et al., 2012).

What drives threat-related biases?

A major limitation of the ‘threat’ literature is that it is often unclear whether sensory or affective dimensions of stimuli drive prioritized processing (Adams, Gray, Garner, & Graf, 2011). The claim that perceived threat is the cause of a processing bias is crucial to the experimental logic of many studies, but is often based on assumption, rather than empirical evidence. It is important to control for, or at least explicitly characterize, the influence of low-level stimulus properties. One way to circumvent the issue of low-level confounds is to pair perceptually similar stimuli (pairs of neutral faces, simple gratings) with negative and neutral valences, via fear conditioning/ affective learning. Two studies in our meta analysis employed this technique (Alpers et al., 2005; Anderson et al., 2011) and observed evidence for a threat-related bias even when low-level confounds were eliminated via this method. Similarly, we would suggest that control stimuli (such as spatially and contrast inverted faces) provide a good means of dissociating sensory and affective factors (e.g. Gray et al., 2013). In addition to controlling for low-level stimulus properties, future studies could systematically measure affective dimensions of stimuli, via perceived valence, arousal and dominance ratings (Hedger et al., 2015b). This would allow a more detailed, precise and standardized examination of the relative contribution of low-level and affective factors.

Relating behavioral and neuroimaging measures

Neuroimaging techniques have the potential to reveal the neural signatures and brain regions underlying unconscious threat processing. For instance, there is considerable evidence that the amygdala is an important component of the neural circuitry involved in threat processing (for a review see Adolphs, 2008). However, patients with amygdala lesions nonetheless show prioritized processing of threat stimuli in a wide range of behavioral tasks (Tsuchiya, Moradi, Felsen, Yamazaki, & Adolphs, 2009; Piech et al., 2010; Piech et al., 2011). It therefore remains an interesting question as to whether amygdala activation in response to unconsciously presented faces has a *causal* role in driving threat responses, or whether it is simply *correlated* with the processing of threatening stimuli. Thus, whilst these neuroimaging findings are invaluable in many respects, they are more easily interpretable when combined with sensitive, well-controlled behavioral measures of enhanced threat processing.

When a procedure (e.g. masking) prevents conscious awareness of stimuli, but behavioral evidence for threat processing is detected, it is often concluded that threat processing temporally precedes awareness, or that it is ‘preconscious’ (Fox et al., 2010; Mogg et al., 1995, Sylvers et al., 2011). However, the rapidity of stimulus presentation is not related to the rapidity of processes under study (Vanrullen, 2011). Restricting presentation time directly affects the quality of visual input or equivalently, the signal to noise ratio. It is possible that conscious awareness requires more robust visual input than threat processing, but the two processes occur at similar latencies when the signal to noise ratio is sufficient. Thus, effects generated by subliminal stimuli do not shed light on the relative speeds of awareness and threat processing. This is important, given that one of the proposed advantages of unconscious threat processing is that it is faster than general purpose visual processing (Tamietto & deGelder, 2010), but the evidence for this component of the standard hypothesis is very mixed (Pessoa, 2010). This issue may be investigated with further studies using electrophysiological methods with fine temporal resolution, although it is currently unclear whether responses at ‘deep’ (subcortical) brain structures can be reliably estimated via EEG or MEG (Baumgartner, Pataia, Lindinger, & Deecke, 2000).

Awareness measures and response bias.

To clarify whether threat related biases are genuinely independent of awareness, future MVP studies could assess the relationship between stimulus visibility and threat bias by parametrically varying the SOA between target and mask. Recent work applying this logic to studying explicit and implicit measures of affective processing has revealed that these are strongly dependent on visibility/ the SOA (Lahteenmaki, Hyona, Koivisto, & Nummenmaa, 2015). Ideally, signal detection awareness measures would be employed concurrently with the visual probe trials in a manner that corrects for individual response bias. Researchers could also conduct *a priori* power calculations to determine the number of trials required for a sensitive awareness check. At a minimum, this could be calculated post hoc to assess the likelihood of type two errors.

Response biases were identified as an issue for BR and bCFS studies. To combat the issue of response biases, some researchers have implemented non-rivalrous ‘simulations’ that attempt to mimic piecemeal rivalry, by alternating the transparency of regions of superimposed images, with the temporal dynamics of these alternations drawn from rivalry data (Baker & Graf, 2009; Lee & Blake, 2004). Similar simulations could be used to characterize an observers’ tendency to report perception of a threatening stimulus

as a function of its physical, quantifiable visibility. This would provide an effective way to estimate the extent of response bias under rivalry-like conditions of partial visibility.

Do threat-related biases generalize?

All three paradigms that we reviewed revealed some evidence of threat-related bias under certain conditions. It is possible, however, that even the shared findings (e.g. a processing advantage of fearful over neutral faces) rely on the particular (and arguably unusual) stimulus conditions common to all tasks. All paradigms involve simple displays of no more than two isolated, static stimuli, with a high probability that a threatening stimulus will appear in one of a limited number of locations. More commonly, the human observer needs to select a subset of stimuli for further processing from a complex, dynamic scene, after, or in parallel with image segmentation, and often solving for partial occlusion. Novel behavioral paradigms that manipulate the number, predictability and location of threatening stimuli in more complex displays could be implemented to place more real-world demands on the capacity limits and physiological constraints of the visual system. After all, unconscious threat processing is of little functional benefit if it only operates in very simple environments.

2.8.8. Summary and Conclusions.

Our meta analysis and critical review makes a number of important empirical, methodological and theoretical contributions. At the empirical level, we have quantitatively combined data from a large and diverse research field, in which there was little overarching consensus. This has allowed us to identify and precisely quantify relationships between threat bias and stimulus, paradigm and observer parameters, in a way not possible by considering the results of individual studies alone. At the stimulus level, we have shown that fearful faces are the only stimuli that reliably elicit a threat effect across paradigms. However, the threat bias elicited by fearful faces may be attributable to low level confounds; fearful faces also reliably elicit equivalent bCFS effects (relative to neutral faces) in a spatially inverted configuration. At the paradigm level, we note that within the MVP paradigm, threat biases are strongly moderated by SOA. This effect of SOA was also stronger in studies where no awareness check was conducted. We should therefore be cautious in interpreting data from the MVP literature, since effects may be accounted for, to some extent, by partial visibility that was undetected due to insufficient

awareness measures. We also found evidence that prioritization of threat may be quite limited at the early stages of perceptual selection, as indexed by small initial dominance effects in the BR paradigm. At the observer level, our analyses suggest that anxiety may modulate threat biases, but only under specific presentation conditions.

We can think of our methodological contribution in two ways. Firstly, our analysis of the literature has direct implications for the design of future experiments and which methods may form the basis for interesting new research questions. Secondly, in terms of our meta analysis itself, we have applied rigorous methods to tackle important issues, for example by using a novel combination of recent approaches to tackle dependency between effects and missing data.

At the theoretical level, we have raised important questions about how awareness is measured and the ecological validity of different methods used to manipulate awareness. We have also evaluated evidence for the notion that anxious individuals have an unconscious bias for threat across several different paradigms. This novel analysis invites the field to revisit conclusions drawn from studies that have only employed masking to manipulate awareness (e.g. Bar Haim et al., 2007). Lastly, at the most basic level, our analyses may call for a re-definition of the scope and limits of visual processing that transpire without awareness, which has been discussed alongside some recent theoretical frameworks (e.g. Hassin, 2013; Tamietto & deGelder, 2013) and narrative reviews with no quantitative component (e.g. Axelrod, Bar & Rees, 2015).

Considering our meta-analyses and critical review together, we suggest that uncritical acceptance of the standard hypothesis, which states that threat stimuli can be identified and prioritized without awareness, is premature. We emphasize the significant methodological issues surrounding the assessment of awareness, response bias and low-level confounds. Tackling these substantial issues will require rigorous measures of awareness and combining evidence across carefully controlled, novel and ecologically valid experimental designs.

Autonomic Arousal and Attentional Orienting to Visual Threat are Predicted by Awareness.

3.a. Abstract.

The rapid detection and evaluation of threat is of fundamental importance for survival. Theories suggest that this evolutionary pressure has driven functional adaptations in a specialized visual pathway that evaluates threat independently of conscious awareness. This is supported by evidence that threat-relevant stimuli rendered invisible by backward masking can induce physiological fear responses and modulate spatial attention. The validity of these findings has since been questioned by research using stringent, objective measures of awareness. Here, we use a modified continuous flash suppression paradigm to ask whether threatening images induce adaptive changes in autonomic arousal, attention, or perception when presented outside of awareness. In trials where stimuli broke suppression to become visible, threatening stimuli induced a significantly larger skin conductance response than nonthreatening stimuli and attracted spatial attention over scrambled images. However, these effects were eliminated in trials where observers were unaware of the stimuli. In addition, concurrent behavioral data provided no evidence that threatening images gained prioritized access to awareness. Taken together, our data suggest that the evaluation and spatial detection of visual threat are predicted by awareness.

3.b. Publication Note.

This chapter is published as a research article: Hedger, N., Adams, W. J., & Garner, M. (2015). Autonomic arousal and attentional orienting to visual threat are predicted by awareness. Journal of Experimental Psychology. Human Perception and Performance, 41(3), 798–806. <http://doi.org/10.1037/xhp0000051>: <https://goo.gl/maU9l9>

3.c. Contributions.

Experimental design, data collection, analysis and write up were completed by Nicholas Hedger under the supervision of Wendy Adams and Matthew Garner. Jamie Findlay assisted with the data collection and a subset of the behavioral data was submitted for his BSc dissertation.

3.1. Introduction.

Given capacity limits, a critical function of vision is to direct resources in order to promote efficient detection and evaluation of threat in the environment. Additionally, it is widely held by neurocognitive theories that threatening stimuli are processed in the absence of conscious awareness, possibly via a subcortical visual pathway to the amygdala (for a review, see Tamietto & de Gelder, 2010), a structure implicated in mediating adaptive fear responses (Adolphs, 2013).

Empirical support for such theories comes predominantly from studies employing backward masking, a technique in which a briefly presented (typically ~30 ms) target stimulus is rendered invisible via masking by a subsequent, co-located stimulus (Esteves & Ohman, 1993). Threatening stimuli presented under these conditions can nonetheless elicit responses in both central and autonomic nervous systems consistent with fear arousal. For example, images of threatening stimuli increase amygdala activity (Morris, Ohman, & Dolan, 1999; Nomura et al., 2004) and autonomic skin conductance responses (SCRs) when rendered invisible by masking (Esteves, Dimberg, & Ohman, 1994). In addition, masked threat stimuli are suggested to selectively direct spatial attention to enhance threat appraisal and behavioral response selection. The masked visual probe paradigm provides evidence for this effect: an observer's detection of a peripheral target is enhanced when co-located with a preceding masked threatening image, relative to a neutral image (e.g. Carlson, Fee, & Reinke, 2009; Fox, 2002; Mogg & Bradley, 2002). Together, these findings suggest that signals of threat are registered independently of awareness and influence adaptive changes in behavior and physiology.

The debate concerning the most valid assessment of awareness has an extended history. In the main, this centers on a tension between measures that index subjective vs.

objective awareness (Cheesman, & Merikle, 1984; Merikle, Smilek, & Eastwood, 2001). One approach is to use an observer's subjective report to index whether a stimulus is perceived. This *subjective* approach is rooted in the intuitive idea that only the observer can give reliable information about their conscious experience, since awareness is subjective in nature (Wiens, 2006). The subjective approach has been challenged however, on the basis of signal detection models, which state that observers may be reluctant to report weak or brief percepts (Green & Swets, 1966). Accordingly, a more stringent approach relies on *objective* measures: observers are deemed unaware of a stimulus only when performance is at chance in a forced choice discrimination task (e.g. determining the presence or location of the stimulus, see Pessoa, Japee & Ungerleider, 2006). A current view is that both objective and subjective measures have conceptual and practical limitations and thus should be used in combination to comprehensively characterize visual awareness (Szcepanowski & Pessoa, 2007; Wiens, 2006).

It is in the context of this objective approach that more recent evidence has questioned the ability of the backward masking paradigm to reliably suppress stimuli from awareness. For example, studies using signal detection criteria to assess awareness revealed that the majority of observers can reliably detect stimuli masked after 33ms (Pessoa, Japee, Sturman & Ungerleider, 2006). Furthermore, exposure durations of 17-25ms may result in above chance discrimination of masked threat stimuli, even when subjective reports indicate unawareness (Szcepanowski & Pessoa, 2007). Such findings are important, given that physiological and behavioural threat responses observed in studies employing ~30 ms presentation durations (Carlson, Reinke, & Habib, 2009; Morris, Ohman, & Dolan, 1999; Whalen et al., 1998) are cited as concomitants of "unconscious" processing. Thus, it has been suggested that many previous studies reporting unconscious responses to threat may be explained by sub-optimal measures of awareness (Pessoa, 2005). Consistent with this stance, Codispoti, Mazzetti and Bradley (2009) showed that physiological responses to emotionally salient stimuli (including SCRs) require exposures of at least 80 ms to be reliably measurable (a duration at which stimuli are also reliably detected) and that exposure durations of 25 ms or less eliminate all physiological reactions. Furthermore, it has recently been suggested that other methodological limitations of masking studies, such as temporally inaccurate presentation methods (Hannula, Simons, & Cohen, 2005; Wiens et al., 2004) and trial order confounds (Wiens, Katkin, & Ohman, 2003) could explain the physiological responses to briefly presented stimuli observed in previous studies. Thus, although previous physiological studies suggest that threat can be

evaluated without awareness of the eliciting stimulus, the limitations of the masking paradigms employed require that such findings are re-examined. Similar concerns apply to the masked visual probe literature, which is characterised by a number of studies that do not include an explicit awareness check (Carlson & Reinke, 2008; Fox, Cahill & Zougkou, 2010; Sutton & Altarriba, 2011; Writh & Schultheiss, 2007). Instead assumptions are made about the efficacy of the masking procedure based on previous literature. Again, therefore, the associated claims about unconscious threat processing should be revisited.

Continuous flash suppression (CFS) is an increasingly popular technique that allows longer and more robust manipulations of awareness than backward masking (Lin & He, 2009). A stereoscope or anaglyph is used to present a dynamic masking pattern to one eye, which suppresses conscious perception of stimuli presented to the other. Despite continuous input to both retinae, the observer may perceive only the masking pattern for long durations, sometimes up to several minutes (Tsuchiya & Koch, 2005).

A rapidly growing body of studies employing CFS has demonstrated that emotionally salient stimuli suppressed from awareness can induce changes in neural activity indicative of threat processing. For instance, studies employing CFS in conjunction with functional magnetic resonance imaging (fMRI) have recorded differential amygdala responses to invisible fearful and neutral faces (Jiang & He, 2006; Vizueta, Patrick, Jiang, Thomas, & He, 2012). However, these studies did not report convergent changes in behaviour or autonomic physiological arousal that characterise functional threat responses. Although the amygdala is responsive to threat, contemporary models propose that it also sensitive to a broader range of stimulus properties (e.g. unpredictability, value, subjective preference) that modulate cognition and behaviour (Bar & Neta, 2007; Herry et al., 2007; Pessoa, 2010). Moreover, patients with amygdala lesions nonetheless exhibit selective prioritisation of threat in a range of behavioural tasks (Piech et al., 2011; Tsuchiya, Moradi, Felsen, Yamazaki, & Adolphs, 2009). Thus, it is unclear whether differential amygdala responses to suppressed stimuli can be directly related to functional threat responses, without measures of behavioural performance or autonomic arousal.

Psychophysical studies indicate that both spatial and feature-based attention may modulate the processing of low-level properties of simple stimuli (e.g. Gabor patch orientation) rendered invisible by CFS (Bahrami, Carmel, Walsh, Rees, & Lavie, 2008; Kanai, Tsuchiya, & Verstraten, 2006; Shin, Stolte, & Chong, 2009). In addition, the relative salience of invisible Gabor stimuli may drive eye movement responses (Rothkirch, Stein, Sekutowicz, & Sterzer, 2012). These observations have shaped the modal view that

that attention and awareness are dissociable processes (Koch & Tsuchiya, 2007). By extension, conscious perception of the threat stimuli may not be required for their attentional selection. However, the extent to which threatening stimuli modulate spatial attention under CFS has not yet been investigated.

A variant of CFS: breaking continuous flash suppression (bCFS) has also been used to assess visual processing without awareness. The initial suppression duration is used as a correlate of a stimulus' unconscious processing advantage. For example, fearful faces are detected more quickly than neutral faces from invisibility induced by CFS (Yang, Zald, & Blake, 2007), suggesting they are prioritised at unconscious stages of processing. However, recent research suggests that these differences in detection times may be better explained by low-level stimulus variations such as contrast (Gray, Adams, Hedger, Newton, & Garner, 2013) or by criterion effects (Stein, Hebart, & Sterzer, 2011) rather than by unconscious processes sensitive to threat. The extent to which signals of threat are evaluated under CFS, therefore, remains a matter of debate.

In the present study, we assessed the extent to which unconsciously presented visual threat evokes adaptive changes in physiology and behavioural performance. To investigate this, we recorded SCRs to stimuli rendered invisible by CFS. With two concurrent behavioural tasks, we additionally examined whether threatening stimuli presented outside of awareness i) gain prioritised access to awareness or ii) modulate spatial attention.

3.2. Method.

Stimulus images from the International Affective Picture System (IAPS; Lang, Bradley & Cuthbert, 2008) were chosen on the basis of their high emotional validity. Four images depicted biologically relevant threats (images of animal attack, e.g. a snake in a striking pose, IAPS numbers: 1050, 1300, 1726, 1930) and 4 depicted non-threatening animals (e.g. a deer, IAPS numbers: 1440, 1610, 1630, 1710) (mean valence and arousal ratings on a scale of 1-9 for images of animal attack: 3.90 and 6.59 respectively, non-threat images: 7.90 and 4.61 respectively – see Lang et al., 2008). We used images of animal attack as threat stimuli because they elicit the most robust SCRs of all image categories in the IAPS set (Bradley, Codispoti, Cuthbert, & Lang, 2001). All images were matched for root mean squared (RMS) contrast and mean luminance. Each image subtended 6.2 x 4.1

degrees of visual angle (DVA) at the viewing distance of 60 cm. Stimuli were prepared using MATLAB (Mathworks, Natick, MA) and presented using Inquisit 2 software (Millisecond Software) on a Sony Trinitron (1280 x 1024 pixel resolution, 60 Hz) monitor.

In the *experimental* block, we employed a CFS paradigm that closely mirrored that of Jiang, Costello, Fang, Huang and He (2006). Figure 3.1 shows a schematic of the trial sequence. (i) At the beginning of each trial, a central fixation cross was presented to each eye. (ii) Subsequently, observers viewed a CFS display via a mirror stereoscope, in which one eye viewed a pair of dynamic masks (random patterns of red-tinted, high contrast ellipses), which updated at 10Hz. The other eye viewed intact and block-scrambled (block size 0.17 x 0.17 DVA) versions of one IAPS image, which were presented simultaneously at 1.4 DVA to the left and right of a central fixation cross. Each eye's display was framed by a random dot surround (9.5 x 11.4 DVA) to assist fusion. The IAPS images (and their scrambled versions) were slowly introduced by linearly increasing contrast from 0-100% over the 800ms presentation period to reduce the risk of onset transients breaking suppression. (iii) This display was followed by an inter-stimulus interval (ISI) of 100ms during which only the fixation cross and random dot surround were presented. (iv) To measure the allocation of spatial attention, a Gabor patch (2.5 x 2.5 DVA, 4 cycles per degree) was briefly displayed (100ms) at the location preceded by the intact (valid) or scrambled image (invalid). The Gabor was oriented 1° clockwise or counter-clockwise from the vertical axis. The observer reported the orientation (clockwise or counter-clockwise) of the Gabor via key press "as quickly and accurately as possible". (v) At the end of each trial, observers were prompted to indicate (via key press) whether "*anything other than the mask*" had been visible during the CFS presentation, thereby providing a subjective awareness check. Importantly, this awareness check did not rely on recognition of the target stimulus, which could be influenced by criterion effects. Moreover, the uniform colour of the mask allowed observers to use any deviation from this colour as a cue to aid stimulus detection. The inter-trial interval of 8 seconds was sufficiently long to enable SCRs from consecutive stimuli to be differentiated (Boucsein et al., 2012; Breska, Maoz, & Ben Shakh, 2011).

Figure 3.1. Schematic of trial sequence from the experimental block.

Observers viewed a CFS display for 800 ms. The allocation of attention was measured by discrimination of a subsequent probe. Subjective awareness was measured by forced choice responses at the end of each trial. Please contact the author, or consult the published article for this figure.

The University of Southampton's local ethics committee approved the study and all observers gave informed consent. Twenty-nine naive undergraduate students (11 male) each completed 128 randomly ordered experimental trials (2 stimulus categories (threatening, non-threatening) x 4 images x 16 repetitions), counterbalanced across probe position (valid, invalid), location of the intact image (left, right) and eye of suppressed images (left, right). Observers also completed a separate block of 128 *control* trials as an objective awareness check. Control trials were identical to the experimental trials, except that after each stimulus presentation, rather than viewing a Gabor, observers reported whether the intact stimulus had appeared to the left or right of fixation (2AFC). Observers whose performance significantly exceeded chance level (binomial limit= 74 correct responses) were classified as being objectively aware of (at least some of) the stimuli; their data were not included in analyses of 'unaware' trials – see below. An *a priori* calculation confirmed that this task provided adequate statistical power (> 80 %) to detect even small (> Cohen's $h = 0.22$) deviations from chance performance.

Throughout the experiment, skin conductance responses were sampled at 1000Hz using a BioPac MP150 amplifier (Goleta, CA). SCRs were collected by applying Ag-AgCl electrodes with conductive gel to the medial phalanges of the second and third fingers of the observer's non-dominant hand. Data from 6 observers were removed from the SCR analyses due to recording failure.

SCR data were reduced by submitting the raw amplitudes to a low pass filter (cut-off frequency: > 25 Hz). Responses were defined as the maximum peak-to-peak amplitude in a pre-defined window 1-4 seconds after stimulus onset, measured in microsiemens (Cacioppo, Tassinary, & Berntson, 2007). We did not define a lower limit to demarcate a “non-response” (Dawson, Schell, & Filion, 2000), to avoid masking potentially small differences in trials where stimuli were successfully suppressed. To correct for individual differences in reactivity, these amplitudes were normalised by dividing by the individual’s mean SCR amplitude (Lyken & Vennables, 1971). SCR outliers were identified as lying beyond 3 standard deviations from the mean following a log transformation to correct for skew. Visual review of these outliers revealed large, abrupt changes in amplitude consistent with movement or recording artefacts. These accounted for 0.72% of the data and were excluded from further analyses. SCRs for each observer were summarised by a “threat effect” score: the differential response to threatening vs. non-threatening stimuli (*Normalised SCR to threatening images - Normalised SCR to non-threatening images*). A positive score represents a larger SCR to threatening than non-threatening images and indicates fear arousal.

3.3. Results.

Trials in the *experimental block* were separated into those in which observers were not aware of the target stimulus (*unaware*) vs. those in which at least some part of the stimulus was detected (*aware*). Trials were classified as *unaware* (1638 trials) when the observer (i) indicated that no stimulus was visible in the subjective report for that trial and (ii) the observer performed at chance level in the 2AFC objective control task ($N=14$). Conversely, *aware* trials (829 trials) were taken from all observers ($N=23$), and included all trials on which the subjective report indicated that the stimulus had broken suppression.³ Data in each trial category were pooled across observers and analysed with a standard bootstrap resampling procedure (Efron & Tibshirani, 1993; Howell, 2013).

Figure 3.2 depicts the SCR data from the *experimental block*, separated by trial category. Aware trials induced significantly larger SCRs to threatening than non-

³ We used conservative criteria to define unaware trials, excluding data from observers performing above chance in the awareness check. However, partitioning the data into aware and unaware trials based only on subjective reports (so that all data from all observers was used) produced results that were nearly indistinguishable from those reported in the main text.

threatening images ($M = 0.159$, 95% CI [0.004, 0.309], $p = .041$). However, this differential response was eliminated in unaware trials ($M = -0.070$, 95% CI [-0.159, 0.022], $p = .140$). Thus, when observers were unaware of stimuli, as determined by subjective and objective criteria, we found no evidence that SCRs were modulated by visual threat. Furthermore, the SCR threat effect was significantly larger in aware than unaware trials ($p = .010$).

Figure 3.2. SCR data from the CFS block, plotted as a function of trial category.

Error bars give 95% confidence intervals obtained via bootstrapping. Please contact the author, or consult the published article for this figure.

To assess the effect of threat on target visibility, we determined the percentage of trials in which stimuli broke suppression (bCFS) as a function of stimulus category, as indexed by forced choice decisions at the end of each trial. In contrast to previous findings with emotional face stimuli (Gray et al., 2013; Yang, et al., 2007), we found no evidence that threatening and non-threatening images differed in visibility, $t(28) = .736$, $p = .468$

(threatening images: $M = 32.17\%$, non-threatening images: $M = 30.50\%$, $d = 0.05$), suggesting no unconscious prioritisation of threat, relative to non-threat images.

To measure the effect of threat on spatial attention, we used the data from the attentional cueing task to compute an accuracy difference score. In this challenging probe discrimination task, modulation of spatial attention is reflected in differences in discrimination accuracy, rather than response time⁴ (Jiang et al., 2006). Attention was thus measured by subtracting the probe discrimination accuracy in invalid trials from accuracy in valid trials. A positive score reflects greater discrimination accuracy at the location of the target stimulus (vs. the scrambled image) suggesting that attention has been allocated to the target (Chica, Martin-Arevalo, Botta, & Lupiáñez, 2014).

Figure 3.3 shows substantial differences in attentional allocation between aware and unaware trials. We first compared the accuracy difference scores in each condition to zero (no accuracy difference between valid and invalidly cued trials). In aware trials, threatening images attracted spatial attention ($M = 8.95\%$, 95% CI [1.13, 16.80], $p = .027$), but non-threatening images did not: observers did not show significantly enhanced probe discrimination following non-threatening images ($M = 6.03\%$, 95% CI [-2.35, 14.25], $p = .153$). In contrast, in unaware trials, neither threatening ($M = 0.96\%$, 95% CI [-5.71, 7.74], $p = .783$), nor non-threatening images ($M = 1.14\%$, 95% CI [-5.57, 7.89], $p = .748$) attracted spatial attention. In other words, in the absence of awareness, there is no attentional preference for intact images of either category, relative to their scrambled counterparts. Next, we directly compared stimulus categories (threat vs. non-threat) in terms of attentional allocation. The difference in attention to threatening and non-threatening images was not significant in aware trials ($p = .640$) or unaware trials ($p = .944$). When collapsing accuracy difference scores across stimulus categories, overall attentional allocation to intact vs. scrambled images was significantly greater in aware than unaware trials ($p = .047$), consistent with an increased attentional preference for intact /

⁴ Enhanced probe discrimination in aware trials was not the product of a speed-accuracy trade-off; probe validity did not modulate reaction times in any condition ($ps > .18$). Raw accuracy scores for each of the eight conditions also were statistically above chance (binomial limit = 68%), suggesting null effects were not due to floor effects. In aware trials: threatening valid: 82%; threatening invalid: 73%; nonthreatening valid: 74%; nonthreatening invalid: 68%; and in unaware trials: threatening valid: 76%; threatening invalid: 75%; nonthreatening valid: 77%; nonthreatening invalid: 76%.

behaviourally relevant images in aware trials, irrespective of stimulus type. Finally, to decompose the effect of awareness on attentional orienting to intact images, we compared the accuracy difference scores for aware vs. unaware trials for each stimulus category (e.g. threat stimuli in aware trials compared to threat stimuli in unaware trials). This analysis revealed that the overall enhanced attending to intact images in aware trials, relative to unaware trials was 1.6 times greater for threatening (mean difference = 7.99, $p = .065$), than non-threatening stimuli (mean difference = 4.89, $p = .188$). In summary, consistent with the physiological measures, the data suggest that threat stimuli only modulate attention when they are available to awareness.

Figure 3.3. Data from the attentional cueing task, plotted as a function of stimulus and trial category.

Error bars are 95% confidence intervals obtained via bootstrapping. Please contact the author, or consult the published article for this figure.

3.4. Discussion.

In the present study, we used a CFS paradigm to assess physiological and behavioural concomitants of threat processing under conscious and unconscious viewing conditions. Under conscious processing, threat stimuli induced robust physiological responses and captured spatial attention over scrambled images. However, we found no evidence of threat evaluation without awareness: in unconscious presentations we found (i) no threat-related autonomic responses, (ii) no enhanced salience as measured by bCFS and (iii) no greater attentional allocation to threat (or non-threat) images relative to scrambled images. Notably, these null findings in unaware trials were coupled with robust positive effects under aware conditions, making it unlikely that a lack of sensitivity in our task or analyses can account for the null effects. Moreover, our own meta-analyses suggest that attentional effects for animal attack stimuli in subliminal cueing paradigms are medium ($d_z=0.51$) and that biases for threatening stimuli (fearful faces) in bCFS paradigms are large ($d_z=0.71$) (Hedger, Adams, & Garner, 2014). Consequently, our sample size provided good statistical power (.85 and .98 respectively) to detect similar effects of this magnitude. How can we reconcile our findings with previous studies that have reported evidence of emotion / threat processing without awareness?

Many previous demonstrations of selective responses to unconsciously presented threat stimuli have primarily used backward masking to prevent briefly presented stimuli from reaching awareness. However, some of these studies did not assess objective awareness on an individual basis (e.g. Carlson & Reinke, 2008; Dimberg, Thunberg, & Elmehed, 2000; Esteves et al., 1994; Fox et al., 2010; Morris et al., 1999; Writh & Schultheiss, 2007), or extrapolated objective 2AFC data from one sample of observers to another (e.g. Carlson et al., 2009; Nomura et al., 2004; Ohman & Soares, 1994; Phillips et al., 2004; Sutton & Altarriba, 2011). Moreover, some objective awareness checks may have lacked sufficient statistical power to establish above-chance detection performance, leaving open the possibility of awareness on at least some trials (e.g. Katkin, Wiens, & Ohman, 2001; Wiens et al., 2003). For these reasons, it is unclear whether these studies can unequivocally support claims of unconscious perception. Here, we assessed awareness at

the individual level by using both subjective and objective criteria and found no evidence for physiological or behavioural threat responses under unaware conditions. Our results are thus more consistent with other lines of evidence that suggest that responses to threatening stimuli are strongly modulated by their conscious detection (Hoffman, Lipka, Mothes-Lasch, Miltner, & Straube, 2012; Mayer et al., 1999; Peira, Golkar, Ohman, Anders & Wiens, 2012).

It is important to note that two recent studies with stringent awareness measures have reported SCR's to conditioned (Raio, Carmel, Carasco, & Phelps, 2012) and unconditioned fearful faces presented under CFS (Lapate, Bokers, Li, & Davidson, 2013). It is possible therefore, that fearful faces, but not animals in attack postures, are able to modulate autonomic responses independently of awareness (notably, Lapate and colleagues also observed that physiological responses to threatening animal stimuli were dependent on awareness). However, to maximise ecological benefit, an effective threat-detection system should mediate adaptive responses to a wide range of threat stimuli in a variety of viewing conditions. For instance, the fear expression is only an *indicator* of threat in the environment, rather than being a *direct* threat to survival and wellbeing. It seems unlikely that humans would have evolved unconscious sensitivity to an indicator of unspecified threat in the environment (fear), without sensitivity to the classes of directly threatening stimuli (e.g. animal attack). Indeed, single cell recordings indicate that the amygdala contains large populations of neurons that are highly selective to both face (Rutishauser et al., 2011) and animal stimuli (Mormann et al., 2011). Moreover, human neuroevolutionary models claim that unconscious vision remains particularly sensitive to snakes, due to the deadly threat they posed to early mammals (Cook & Mineka, 1991; Ohman & Mineka, 2001). Future studies should therefore directly compare autonomic and behavioural responses to a broader range of threat stimuli to investigate the source of this discrepancy.

Our data do not, of course, preclude unconscious processing of other types of threat stimuli. In addition, it is important to note that other dependent measures, such as event related potentials (ERPs) may reveal selective threat responses in the absence of the physiological and behavioural responses that we measured. This consideration is particularly important given that even when behavioural correlates of semantic processing are eliminated, electrophysiological correlates may still be present (Heil, Rolke, & Pecchinenda, 2004). Similarly, studies have reported differential amygdala activation to threatening and neutral stimuli under binocular rivalry (e.g. Jiang & He, 2006; Pasely,

Mayes, & Schultz, 2004; Vizueta et al., 2012). With both these ERP and fMRI measures it is important to consider that any pair of stimuli may induce differential activation at the neural level. However, what is the functional purpose of sensitivity to threat at the neural level if this does not cause adaptive changes in physiology and perception? Our study provides no evidence that threatening visual stimuli induce adaptive physiological or behavioural responses when reliably suppressed from awareness via CFS. Thus, combining our paradigm with neuroimaging methods would help further clarify the role of the amygdala in fear behaviour during conscious and unconscious threat processing.

Recently, concerns have been raised regarding the measurement of awareness in binocular rivalry paradigms (including CFS). The boundary between aware and unaware states in rivalry may be graded rather than discrete (Pessoa, 2005), and involve “mixed” states of “piecemeal” rivalry (Alais & Blake, 2005). For this reason, an observer’s perceptual report of awareness is likely to be influenced by their response criteria. These concerns apply to bCFS studies that use suppression duration (indexed by response time) as a measure of differential unconscious processing (as in Gray et al., 2013; Sylvers, Brennan, & Lilienfeld, 2011; Yang et al., 2007). This response time comprises both an unconscious component that reflects “pure” suppression duration, but is also confounded by the time observers take to use visible information to report a stimulus (Stein et al., 2011). It may therefore be inappropriate to conceptualise differential response times in these studies as genuine differences in unconscious processing. We obviated these issues by using fixed presentation durations and a more conservative forced choice task to evaluate awareness and found no advantage for threatening images in breaking suppression. In addition, work from our lab shows that the widely reported “fear advantage” in which fear faces emerge quickly from CFS (Yang et al., 2007) can be explained entirely by the low-level characteristics of fear faces; stimuli with the same low-level properties (i.e. spatial frequency content) but whose emotional expressions are unrecognizable are similarly detected faster following CFS (Gray et al., 2013; Hedger, Adams, & Garner, 2013). Critically, this finding is explained without reference to a sub-cortical pathway that prioritises threat during the suppression phase of CFS and is thus consistent with the present data.

One other important issue in the measurement of awareness is whether an observer’s attempts to monitor their detection of a target may interfere with any unconscious processing that may have otherwise taken place. As Koudier, de Gardelle,

Sackur and Dupoux (2010) note, this is conceptually similar to an “observer effect”, in which an observer’s attempt to monitor the state of a system may alter the processes of that system. For instance, in the masked priming literature, Carr and Dagenbach (1990) found that asking observers to make semantic judgements about a masked stimulus could disrupt priming effects. Is it possible that, in the current study, unconscious processing was similarly disrupted by observers’ concurrent monitoring of target visibility? We think this is unlikely for several reasons. Firstly, our awareness check required a simple, low-level perceptual judgement (i.e. “was anything other than the mask visible?”), which contrasts with the more demanding, semantic judgement employed by Carr and Dagenbach. Our dynamic mask was a uniform hue, which was easily differentiated from the target stimulus: observers were not required to identify the target and could detect it purely on the basis of hue. The cognitive and attentional resources required to monitor target visibility were thus minimised, making it unlikely that our concurrent awareness measures muted any unconscious processing. In fact, one recent eye tracking study has shown that, despite observers actively trying (and failing) to detect a suppressed Gabor patch, attention was nonetheless driven to its location (Rothkirch et al., 2012). Furthermore, in our study, effects in unaware trials were not simply reduced in magnitude, but were in the opposite direction predicted by the unconscious prioritisation of threat: in unaware trials SCRs and attentional cueing effects were greater to *non-threatening* images (this is also opposite to the pattern found in aware trials, where effects were consistent with enhanced threat processing). Lastly, in our study, CFS was used in all trials and thus observers monitored target visibility in all trials. Trials were classified as ‘aware’ or ‘unaware’ post hoc on the basis of target detection. Thus, the process of monitoring awareness would equally affect both aware and unaware trials.

Finally, on a cautionary note, we must consider the limitations of CFS and other paradigms used to manipulate awareness. To assess the role of awareness in threat processing, an optimal paradigm would modulate only awareness, and not any other aspect of visual processing. However, current methods for rendering sensory inputs invisible create stimulus conditions that may not be commonly encountered during everyday visual experience (Kim & Blake, 2005), or require conclusions to be generalized from small samples of individuals with brain damage (e.g. Tamietto et al., 2009). Indeed, neuroimaging data suggest that suppression via both backward masking and CFS disrupts neural activity representing the suppressed stimulus at early stages of visual processing

(Rolls, Tovee & Panzeri, 1999; Tong, 2003; Yuval-Greenberg & Heeger, 2013) and differentially affects activity in the ventral and dorsal processing stream (Almeida, Mahon, Nakayama, & Caramazza, 2008; Fang & He, 2005). It remains possible, therefore, that suppression paradigms with different underlying mechanisms, such as motion induced blindness (Bonneh, Cooperman, & Sagi, 2001), or visual crowding (Koudier, Berthet, & Faivre, 2011) may eliminate awareness but spare visual processing to the extent that threat-responses remain effective.

To summarise, using a CFS paradigm, we found no evidence for physiological or behavioural responses to threat when awareness is eliminated according to stringent subjective and objective criteria. Our data compliment other lines of evidence (see Pessoa, 2005; Pessoa & Adolphs, 2010) that challenge current claims that threat is identified unconsciously to direct processing resources.

Fearful Faces Have a Sensory Advantage in the Competition for Awareness

4.a. Abstract.

Only a subset of visual signals give rise to a conscious percept. Threat signals, such as fearful faces, are particularly salient to human vision. Research suggests that fearful faces are evaluated without awareness and preferentially promoted to conscious perception. This agrees with evolutionary theories that posit a dedicated pathway specialized in processing threat-relevant signals. We propose an alternative explanation for this "fear advantage." Using psychophysical data from continuous flash suppression (CFS) and masking experiments, we demonstrate that awareness of facial expressions is predicted by effective contrast: the relationship between their Fourier spectrum and the contrast sensitivity function. Fearful faces have higher effective contrast than neutral expressions and this, not threat content, predicts their enhanced access to awareness. Importantly, our findings do not support the existence of a specialized mechanism that promotes threatening stimuli to awareness. Rather, our data suggest that evolutionary or learned adaptations have molded the fearful expression to exploit our general-purpose sensory mechanisms.

4.b. Publication note.

This chapter is published as a research article in the Journal of Experimental Psychology: Human Perception and Performance: Hedger, N., Adams, W. J., & Garner, M. (2015). Fearful faces have a sensory advantage in the competition for awareness. Journal of Experimental Psychology. Human Perception and Performance, 41, 1748-1757: <https://goo.gl/bCpWEn>

4.c. Contributions.

Experimental design, data collection, analysis and write up were completed by Nicholas Hedger under the supervision of Wendy Adams and Matthew Garner.

4.1. Introduction.

An important predictor of survival is the ability to detect threat. However, given the capacity limits of our sensory systems, not all visual inputs give rise to a conscious percept - many stimuli within our field of view go undetected in the competition for neural resources (Dehaene & Changeux, 2011). How does a limited-capacity system selectively process those inputs of most significance for survival? A widely held view is that humans have a specialised, subcortical visual pathway that expedites the processing of threatening stimuli (Ohman, 2005). An important component of this proposal is that this pathway is thought to operate independently of conscious awareness (Tamietto & de Gelder, 2010). In the context of survival, it would be advantageous if threats in the environment could influence behaviour before, or without, an observer's awareness of them.

Evidence that threat can be processed preconsciously, or without awareness, comes from paradigms in which visual input is dissociated from awareness (Kim & Blake, 2005). In backward masking, awareness of a briefly presented image is restricted by the subsequent presentation of a co-located mask. Neuroimaging evidence indicates that masked fearful faces can increase amygdala activity, which is indicative of emotional arousal (Whalen et al., 1998; Whalen et al., 2004). Behaviourally, an observer's response to a peripheral "probe" stimulus is faster when preceded by a masked fearful face than a masked neutral face (Carlson & Reinke, 2008; Fox, 2002). Continuous flash suppression (CFS) is a technique in which a stable image shown to one eye is suppressed from perception by a dynamic stream of images presented to the other (Tsuchiya & Koch, 2005). Fearful faces break into awareness from CFS more quickly than neutral faces (Sylvers, Brennan, & Lilienfeld, 2011; Yang, Zald, & Blake, 2007). Together, these findings suggest that fear faces are evaluated without awareness and gain prioritised access to conscious vision. Interestingly, our own meta-analyses show that fearful faces are the only threat stimuli to be reliably prioritised over neutral stimuli across the masked visual probe, binocular rivalry and continuous flash suppression paradigms (Hedger, Adams, & Garner,

2014). Thus, there does seem to be something ‘special’ about the processing of subliminally presented fearful expressions that warrants careful investigation.

Ecological models distinguish between two components of visual signals: content and efficacy (Guilford & Dawkins, 1991). The former relates to the “message” of the signal, whereas the latter relates to the efficient transmission of the signal in relation to the sensory biases of an observer. It is often assumed that fearful faces are prioritised in the competition for awareness due to their *content*, since they signal important information about potential threats (Sylvers et al., 2011; Yang et al., 2007). However, fearful faces may also be prioritised due to their *efficacy*; fast detection could be mediated by purely low-level factors, such as the high contrast physical signal emanating from the eye region – i.e. the increased exposure of the iris and scleral field (Lee, Susskind, & Anderson, 2013). This latter position is in-line with the ‘sensory bias hypothesis’ (Horstmann & Ansorge, 2009; Horstmann & Bauland, 2006), which states salient facial expression signals are exploit the sensory biases of human observers. Thus, without characterising the sensory properties of facial signals, we risk attributing prioritised detection to threat-sensitive processes, when it may be better explained by the low-level physical salience of the expression.

Here, we consider this dilemma. There are two possible mechanisms via which fear faces, or other threat-relevant stimuli might gain prioritised access to conscious vision. First, humans might have evolved specialised mechanisms that evaluate the threat content of visual signals prior to their conscious registration. Second, the physical expression of fear might exploit the sensory tuning of early, general-purpose visual processing. In this latter case, the apparent “threat advantage” could be parsimoniously explained by sensory efficacy, without the involvement of pre-conscious mechanisms sensitive to threat.

Typically, these two accounts are conflated, since the low-level characteristics of facial expressions define the content of the communicated emotion (e.g. wide eyes signal fear). To resolve this issue, therefore, one must experimentally dissociate a stimulus’ sensory and affective properties. In the present study, we addressed this issue with a combination of image analyses and behavioural data. First, we use known properties of early visual processes to estimate the efficacy with which emotional expressions are received by human observers. Second, we use stimulus manipulations that modulate the threat content of our images, without affecting sensory efficacy. Third, we present behavioural data from CFS and masking paradigms that quantify the extent to which emotional expressions gain access to awareness. Lastly, we determine whether this is better

predicted by (i) the images' low-level, effective contrast (a quantity indifferent to threat), or (ii) their threat-content.

4.2. Image Analyses.

Stimulus detection is influenced by low-level properties such as luminance and contrast. Moreover, humans are more sensitive to contrast at certain spatial frequencies, as defined by their contrast sensitivity function (De Valois & De Valois, 1990). Thus, differences in detection between stimuli (e.g. fearful vs. neutral faces) that differ in these low-level properties cannot be considered a valid measure of threat-related processing. The contrast sensitivity function can be conceptualised as a modulation transfer function for spatial contrast energy at early visual stages. Measuring the extent to which stimuli exploit this sensitivity thus provides an estimate of their sensory efficacy. We asked whether fearful and neutral expressions differ in the extent to which they exploit the contrast sensitivity function (i.e. do they differ in “*effective contrast*”?).

4.2.1. *Effective contrast.*

We analysed the NimStim face set, a collection of face stimuli that is widely used in studies of emotion recognition and is subsequently employed in our behavioural experiments. The set includes 24 male and 19 female models from multiple ethnicities (Tottenham et al., 2009). First, we applied an opaque elliptical mask to eliminate external features before equating mean luminance and root mean squared contrast (RMS) of the images (following standard practice in psychophysical experiments). For our initial analyses, we mirrored the average size (13.5 cm bizygomatic diameter, see Katsikitis, 2003) and a typical distance (220cm) of a human face during social interactions.

To calculate effective contrast we followed the procedure of Baker and Graf (2009), implemented in MATLAB (Mathworks). This measure of effective contrast has previously been found to be a good predictor of stimulus salience in binocular rivalry competition (Baker & Graf, 2009). We obtained the amplitude spectrum (contrast energy as a function of spatial frequency) of each face image (figure 4.1a, left panel). We then fitted a second order polynomial to the contrast sensitivity data set “ModelFest” (Watson & Amuhada, 2005) to obtain a continuous contrast sensitivity distribution (figure 4.1a, middle panel,

normalised to the 0-1 range). By multiplying this distribution by the amplitude spectrum, we obtained effective contrast as a function of spatial frequency, for each stimulus (figure 4.1a, right panel).

Summing this contrast across spatial frequency produces an overall estimate of contrast energy after attenuation by the contrast sensitivity function, i.e. the image's effective contrast. Fear faces had higher effective contrast than neutral faces for 41 of the 43 models, and this difference was large in magnitude ($d = 0.76$, 95% CI [0.31 1.21], $p < .001$).

To confirm that this finding generalised beyond this particular image set, we extended our analysis to 641 images by including 4 other widely used face sets (fronto-parallel faces only): the Karolinska Directed Emotional Faces (KDEF: Lundqvist, Flykt, & Ohman, 1998), The Radboud Faces Database (RaFD: Langer et al., 2010), The Pictures of Facial Affect Dataset (Ekman & Friesen, 1976), The Montreal Set of Facial Displays of Emotion (MSFDE: Beaupre & Hess, 2005). The sensory advantage of fear was large and significant in all but the Ekman & Friesen set⁵ (figure 4.1b). The pooled effect size across face sets, calculated via a random effects model, was large and significant ($k=6$, $N= 641$, $d= 1.00$, 95% CI [0.69 1.31], $p < .001$). Based on this pooled estimate, the “probability of superiority” (Grissom & Kim, 2005), i.e. the likelihood that a randomly sampled fearful face will have a sensory advantage over a randomly sampled neutral face, is 84% (95% CI [75% 90%]).

4.2.2. Stability across viewing distances.

The effective contrast differences described above are not scale-invariant; they depend on the particular retinal size of the images. If our physical expression of fear is optimised to be salient in everyday social contexts (Gray, Adams, Hedger, Newton, & Garner, 2013) then this sensory advantage of fear should be robust over distances at which humans typically socialise and communicate. To test this possibility, we extended our

⁵ This outlying result may be attributed to the low control of head orientation, lighting, and lower image resolution compared with other, more recent sets. Moreover (a) the effect is directionally consistent and (b) statistical power is lower, given the significantly smaller number of images in the Ekman set. Thus, this discrepancy should not greatly impact the interpretation of our main findings.

analyses to simulated viewing distances of 50 - 500 cm. As shown in figure 4.1c, the sensory benefit of fear is largest within interpersonal proximities that characterise human social interactions (120 to 360 cm, region within dotted lines, see Argyle, 2013).

Figure 4.1. Effective contrast analyses and data.

(a) Image analysis for one example fearful face. (CSF = contrast sensitivity function). (b) Forest plot depicting the effect sizes for effective contrast differences between fearful and neutral faces (open = open mouthed, closed = closed mouthed). Error bars are 95% confidence intervals. The diamond depicts the pooled effect size. (c) The difference in effective contrast (arbitrary units) between fearful and neutral models as a function of viewing distance. Coloured symbols indicate the mean within each face set, shaded grey

region is the bootstrapped 95% confidence interval. The asterisk indicates the viewing distance used for the initial analyses. Dashed vertical lines span the distances that characterise typical human social interactions. Please contact the author, or consult the published article for this figure.

Importantly, our measure of effective contrast is derived from “classic” contrast sensitivity data (i.e. from challenging threshold conditions). It could be argued, therefore, that most normal (non CFS or unmasked) viewing conditions are *suprathreshold*, to which the shape of this *threshold* contrast sensitivity function may not apply. Indeed, contrast matching experiments have found that perceived suprathreshold contrast is largely invariant with spatial frequency (Georgeson & Sullivan, 1975). However, as De Valois and De Valois (1990) note, contrast matching is not a direct sensitivity measurement and as such, one cannot conclude what the suprathreshold contrast response function is for different spatial frequencies. In fact, other measures, such as magnitude estimation, show that the high and low frequency attenuation of the contrast sensitivity function is maintained at suprathreshold levels (Cannon, 1979). It is therefore inappropriate to conceptualise the contrast sensitivity function as an epiphenomenon restricted to threshold conditions. By extension, the detection of stimuli in natural viewing conditions can be understood, at least to a first approximation, in terms of the properties of the contrast sensitivity function.

Our analyses suggest that fearful expressions are optimised to excite the early visual processes of proximal observers: fear faces contain greater contrast energy at the spatial frequencies that humans are sensitive to, relative to neutral faces. This advantage is purely sensory, and generalises across gender and race. This sensory advantage could be either an evolutionary or learned adaptation.

The case for an unconscious processing advantage for threatening stimuli is most often evidenced by the prioritisation of fearful over neutral expressions. However, angry faces, although also signalling threat, typically yield smaller, inconsistent effect sizes (Hedger et al., 2014) and have even been reported to be *disadvantaged* relative to neutral faces in CFS paradigms (Gray, et al., 2013). Moreover, happy faces, although not signalling threat, have also been found to be prioritised over neutral stimuli in detection paradigms (Becker, Anderson, Mortensen Neufeld, & Neel, 2011) and there is some evidence that they are processed subliminally (Dimberg, Thunberg, & Elmehed, 2000; Schupp et al., 2004). Given the inconsistent and complex nature of these findings, it is

important to understand whether processing differences between these expressions are better explained by their sensory and affective characteristics. To this end, we used CFS (Experiment 1) and masking paradigms (Experiment 2) to investigate whether effective contrast can predict conscious perception of fearful, angry, happy and neutral faces.

4.3. Behavioural Experiment 1: Access to Awareness from CFS.

Under most viewing conditions, our two eyes receive slightly different views of the world and we perceive a single “fused” percept (Howard & Rogers, 1995). However, when dissimilar images are presented to our two eyes at corresponding retinal locations, conscious perception alternates between the two images as their neural representations compete for awareness (Blake & Logothetis, 2002)- a phenomenon termed binocular rivalry. In some respects, this is a controlled phenomenon that can be used to mimic aspects of natural vision, which involves selection amongst multiple sensory inputs, which are assigned to or omitted from conscious perception. In CFS, a dynamic masking pattern is presented to one eye, which can render a stimulus presented to the other eye invisible for seconds before it breaks suppression and enters conscious awareness (Tsuchiya & Koch, 2005). The length of this initial suppression has been used as an index of the unconscious salience of the suppressed image (the bCFS paradigm, Stein & Sterzer, 2014). Here, we use this bCFS paradigm to measure the extent to which stimuli gain access to conscious perception.

4.3.1. Methods.

We selected 4 NimStim models, on the basis of their high emotional validity (mean expression recognition accuracy was 87% - see Tottenham et al., 2009), portraying fearful, happy, angry and neutral expressions. Stimuli subtended 6.2 x 4.1 degrees of visual angle (DVA) at the viewing distance of 60 cm on a 1280 x 1024 pixel resolution, gamma corrected monitor. In order to decouple our images’ low-level, effective contrast from their affective properties, we presented the face stimuli in two different conditions (figure 4.2a). *Normal* faces were presented upright with veridical contrast polarity. *Control* faces were rotated 180 degrees with reversed contrast polarity, producing an image similar to a

photographic negative. Together, these manipulations severely disrupt the recognition and affective evaluation of facial expressions (Gray et al., 2013). Critically, however, they do not alter effective contrast⁶. Thus, if the threat or valence of face images is the critical factor in driving access to awareness, we would expect any threat advantage to be reduced or eliminated for the control images. Conversely, if effective contrast is the key predictor for a ‘threat advantage’, then a similar advantage for the fear expression should be observed within normal and control stimuli.

The trial sequence is shown in figure 4.2b. A central fixation cross was presented to each eye via a mirror stereoscope for 1 second. Subsequently, observers viewed a CFS display for 800 milliseconds, during which one eye viewed a pair of dynamic masking patterns and the other viewed a face stimulus at 1.4 DVA to the left or right of fixation. Face stimuli were introduced gradually by linearly increasing RMS contrast from 0-100% over the 800 millisecond period. Each eye’s display was framed by a random dot surround (9.5 x 11.4 DVA) to assist binocular alignment. Temporal frequency influences the strength of CFS suppression, with mask frequencies above 10Hz exerting weaker suppression than those below (Yang & Blake, 2012). We therefore used a 20Hz mask to ensure that faces broke suppression on a substantial proportion of trials. After the CFS presentation, observers were required to make a forced choice decision as to whether “anything other than the mask” was visible during the trial. This unspeeded measure does not measure response times, or recognition of the target stimulus, which are susceptible to criterion effects (Stein & Sterzer, 2014)

Twenty-two undergraduate students completed 256 experimental trials (4 expressions (anger, happy, fear, neutral) x 2 stimulus manipulations (normal, control) x 32 repetitions), balanced across face location (left or right of fixation). Our sample size provides in excess of 95% power to detect a large effect size (Cohen’s $d = 1.15$, the magnitude of difference in detection between fearful vs. neutral faces from a similar CFS paradigm- Yang et al., 2007).

⁶ Perceived contrast is affected more by low than high luminances (Haun & Peli, 2013). All normal faces had luminance histograms that were negatively skewed (third moment: $M = 0.10$, $SD = 0.21$). Thus, luminance profile reversal may have marginally increased the perceived contrast of control faces, relative to normal faces (which is in contrast to their decreased detection). Therefore, the effect of stimulus type (normal vs. control detection) cannot be explained by changes in the skew/luminance histogram. Importantly, all relationships between effective contrast and detection remained significant after controlling for skew.

4.3.2. Results.

The percentage of CFS trials in which face stimuli became visible is shown in figure 4.2c. Visibility was modulated by expression ($F(3, 63) = 5.33, p = .002$) with fear faces visible most often, followed by happy, neutral, then angry faces. It is notable that angry faces were detected least often, as this conflicts with the notion that threat is selectively prioritised. Pair-wise comparisons revealed fear and happy faces were both detected more frequently than angry faces ($ps < .05$). In addition, stimulus manipulation strongly modulated visibility ($F(1, 21) = 33.31, p < .001, d = 1.06, 95\% CI [0.57\ 1.54]$): normal faces ($M = 50.56, SE = 5.11$) were detected more frequently than control faces ($M = 26.85, SE = 4.34$). Critically, expression and manipulation did not interact in their effects on visibility ($F(3, 63) = 0.18, p = .905$): the main effect of expression was similar for both the normal ($F(3, 63) = 3.14, p = .031$) and control ($F(3, 63) = 3.00, p = .028$) stimuli, with fear detected most often, followed by happy, neutral and anger in both cases. Importantly, this means that the same modulation of visibility by expression and the same “fear advantage” was observed with control stimuli, whose emotional content *severely disrupted*. Thus, some physical property, that is unaffected by the stimulus manipulations, must be driving the effect of expression.

Does effective contrast predict visibility? Visibility was closely related to effective contrast (figure 4.2d) and a linear regression across the 16 facial models revealed that this was significant, $R^2 = .301, p = .026$.

Importantly, however, the main effect of stimulus manipulation (normal vs. control) cannot be explained by low-level properties, as the two stimulus categories have equivalent effective contrast. The mechanisms that govern visual awareness may therefore discriminate faces from non-faces (Stein, Sterzer, & Peelen, 2012), but we found no evidence that emotion or threat had an effect on detection beyond that explained by basic low-level variability between expressions.

Figure 4.2. Experiment 1 data.

(a) Examples of normal and control stimuli. (b) Schematic of CFS trial sequence. (c) Stimulus visibility in the CFS task, as a function of expression and stimulus manipulation. Error bars are ± 1 SEM. (d) Stimulus visibility as a function of effective contrast, collapsed across manipulation, shaded region is ± 1 SEM. Please contact the author, or consult the published article for this figure.

4.4 Behavioural Experiment 2: Access to Awareness from Visual Masking

The bCFS paradigm has been widely used to investigate the competition for visual awareness. However, we might question whether this represents a naturalistic example of how stimuli compete for awareness; binocular rivalry is infrequently encountered in daily life (Arnold, 2011). In contrast, backward masking involves conditions more typical of everyday vision; observers frequently encounter brief glimpses of stimuli when sampling dynamic scenes via saccades and fixations. In our second experiment, therefore, we investigated whether effective contrast can predict the detection of briefly presented, masked facial expressions. In addition, we asked observers to provide affective ratings of the face stimuli, allowing us to assess the contributions of (i) low-level contrast and (ii) affective factors in stimulus detection.

4.4.1. Methods.

Figure 4.3 shows the masking paradigm. At the beginning of each trial, observers viewed the fixation cross for one second. Next, two masks were presented either side of fixation for 200 ms, followed by a target (intact) and non-target (block-scrambled) face for a variable duration (13.3, 26.6, 40, 53.3, 66.7, 80, 93.3, 106.7 or 120ms). Subsequently, two new mask stimuli were presented for 200ms, immediately following the face presentation. Participants were asked to indicate, as accurately as possible, whether the intact face had appeared to the left or right of fixation.

Figure 4.3. Schematic of trial sequence in the masking task. Please contact the author, or consult the published article for this figure.

All stimulus dimensions matched those in Experiment 1. The scrambled face matched the amplitude spectrum averaged across all target face stimuli, ensuring the target could not be localised via non-specific differences in luminance or contrast between the two sides of the display. Mask stimuli also matched the averaged spectral slope of all face stimuli. This prevented interactions between the spatial frequency profile of the target and mask from influencing detection.

An independent sample of 11 participants completed 1152 randomly ordered trials (4 expressions (anger, fear, happy, neutral) x 2 stimulus manipulations (normal, control) x 9 stimulus onset asynchronies (SOAs: 13.3, 26.6, 40, 53.3, 66.7, 80, 93.3, 106.7 or 120ms) x 16 repetitions), balanced across the location of the face stimulus (left, right). Our sample size provided in excess of 95% power to detect the same target effect size as defined for experiment 1.

Observers also completed a Self-Assessment Manikin (SAM) to evaluate face stimuli on the three dimensions of emotional assessment: valence, arousal and dominance (see Bradley & Lang, 1994). On each trial, observers initiated face presentation, which was displayed (unmasked) for 120ms (maximum SOA in the masking task). Valence, arousal and dominance ratings (1-9 scale) were made in separate blocks, consisting of 32 randomly ordered trials.

4.4.2. Results.

Following standard practice, valence and arousal ratings are summarised in ‘affective space’ (figure 4.4a). For normal faces, the distribution of stimuli follows the expected “boomerang” shape (Bradley, Codispoti, Cuthbert, & Lang, 2001) with higher arousal levels reported for stimuli with large positive or negative valence. However, no such pattern is visible for the control stimuli. A cluster analyses confirms this – the distribution of normal stimuli is optimally explained (as determined by Bayesian Information Criterion) by a 3 cluster model that clearly differentiates between the positive (happy) negative (fear, anger) and neutral (neutral) expressions. In contrast, the distribution of control stimuli is optimally explained by a one-cluster model; expressions are not differentiated in affective space. This, consistent with previous work (Gray et al., 2013) confirms that our stimulus manipulations of spatial and contrast inversion severely alter the emotional evaluation of facial expressions. It is possible that increasing statistical power

may detect residual discrimination (Reingold & Merikle, 1988). Nonetheless, it is clear from figure 4a that control stimuli elicit a qualitatively different pattern of affective evaluations from observers. Thus, if affective dimensions are important, this difference should be expected to alter the effect of expression on detection for control stimuli relative to normal stimuli.

Figure 4.4b displays the 2AFC performance accuracy from the masking task. Data were fitted with cumulative normal distributions free to vary in position and slope. Detection thresholds were estimated from these fits for 75% correct performance (upper binomial limit). Thresholds were significantly and substantially modulated by stimulus manipulation, with observers requiring longer SOAs to detect control faces than normal faces (normal: $M = 60.6$ ms, control: $M = 84.6$ ms, $t(10) = 10.7$, $p < .001$, $d = 2.54$, 95% CI [1.38 3.69]). Notably, these detection thresholds occur at much briefer stimulus exposures than those at which observers made affective judgements in the SAM task (120 ms). Thus, discrimination of expressions would have been *even poorer* under conditions that are sufficient for any fear advantage. Similarly to Experiment 1 there was a main effect of expression on detection in both the normal ($F(3, 30) = 9.95$, $p < .001$) and control conditions ($F(3, 30) = 9.22$, $p < .001$). This effect was again similar in the two conditions, with no interaction between expression and stimulus manipulation ($F(3, 30) = 1.15$, $p = .345$). In other words, although spatial and contrast inversion inhibited emotional recognition of the control stimuli, this did not affect the ‘fear advantage’ for detection. Normal and control fearful faces were detected at shorter SOAs than both neutral and angry faces ($ps < .05$, pairwise comparisons). Figure 4.4c illustrates the relationship between effective contrast and detection threshold, for all stimuli. Effective contrast was a similarly good predictor of detection thresholds in both the normal ($R^2 = 0.36$, $p = .014$) and control ($R^2 = 0.41$, $p = .004$) configurations.

Figure 4.4. Experiment 2 data.

(a) Distribution of stimuli in affective space, according to valence and arousal ratings. Valence is normalised to a -4 to +4 range such that 0 indicates neutral. Symbol colour represents facial expression, symbol shape represents clusters obtained via Bayes criteria (i.e. normal = 3 clusters, control = 1 cluster). (b) The proportion of correct face localisation responses is plotted as a function of SOA and expression, with cumulative normal fits. The dashed red lines indicate the mean thresholds for normal and control stimuli (c) Detection threshold as a function of effective contrast. Shaded region indicates ± 1 SEM. Please contact the author, or consult the published article for this figure.

Interestingly, adding expression as a second predictor significantly increased the variance explained by this regression model. However, as we have discussed, it is

important not to conflate ‘expression’ (which is a physical signal) with ‘emotion’ (which implies an affective evaluation of the signal). Thus, the fact that expression adds to the model fit simply means that effective contrast cannot *entirely* account for the effect of expression; it does not, in itself, entail the influence of emotion sensitive processes.

To strengthen this conviction, we performed regression analyses with valence, arousal and dominance as predictors of thresholds. For both the normal and control configurations, none of these variables significantly predicted detection thresholds (all $ps > .05$). Notably, the same was true when these affective ratings were used as predictors of the bCFS visibility data from Experiment 1. Moreover, tests for zero partial association revealed that the relationship between effective contrast and detection thresholds remained significant after controlling for the influence of these variables (normal: $t(11) = -2.57, p = .026$, control: $t(11) = -3.47, p = .005$). In summary, we found that low-level effective contrast predicts stimulus visibility, but found no evidence for any influence of emotion sensitive processes.

4.5. Local Image Analyses.

Our data establish that global differences in effective contrast can predict the prioritisation of faces in the competition for awareness. However, we can refine our analyses further to ask whether this is driven by particular image regions. These regions were defined by symmetric Gaussian windows whose size and standard deviation (2 DVA, 0.5 DVA respectively) matched the stimuli used to derive the Modelfest data (Watson & Amuhada, 2005). These windows were applied to 130 uniformly distributed, overlapping locations within each image and effective contrast was calculated for each region. The relative (z scored) effective contrast, averaged across the models used in our experiments is shown in figure 4.5. Consistent with previous suggestions (Gray et al., 2013; Lee et al., 2013) it is clear that the eye region is highly salient within all expressions, but particularly so for fear faces. This can be attributed to the increased exposure of the white scleral field and dark iris. These features are unique amongst primates and probably co-evolved with human social communication to enhance detectability of gaze (Kobayashi & Kohshima, 1997). Expressing fear amplifies this sensory benefit by increasing the vertical dimension of the scleral field (see Hedger, 2014 for a demonstration).

Figure 4.5. Local variations in effective contrast.

Image colour/luminance represents the Z-scored effective contrast. Please contact the author, or consult the published article for this figure.

4.6. Discussion.

Shaped by sociobiological pressures, human signals are designed to reliably convey information to observers. In the context of threat, a sender may express fear to warn others of danger, or to signal appeasement. However, before a signal can be acted upon, it must be detected. Theories of enhanced signal function by design (Dukas, 1998) thus predict that the facial expression of fear would converge on a form that exploits the sensory processes of a proximal observer. By analysing the Fourier content of faces in the context of human interactions, we found that facial expressions differ in the extent to which they stimulate low-level visual processes. This mechanism, effective contrast, provides a parsimonious explanation for the prioritisation of fearful faces in the competition for awareness, across rather different paradigms: CFS and masking. Critically, this ‘threat advantage’ was independent of perceived emotion; it was unchanged for stimuli with the same effective contrast, but altered emotional content. Moreover, face detection was not predicted by observers’ affective ratings. Our data are inconsistent with the notion that the threat value of fear faces is evaluated outside of awareness and determines access to conscious vision. Instead, our data suggest that access is determined by the tuning of very early visual processes, i.e. the contrast sensitivity function.

Previous work has speculated that the prioritised detection of threat relevant stimuli (including fearful faces) may be linked to simple, low-level stimulus properties (Bar, & Neta, 2006; Coelho, Cloete, & Wallis, 2010; Gray et al., 2013; Stein & Sterzer, 2012; Lee et al., 2013; Yang et al., 2007). However, these studies have not quantified these properties

with respect to underlying human sensory processes. Moreover, to our knowledge, our study is the first to explicitly quantify both sensory (effective contrast) and affective properties (SAM) of facial expressions as predictors in a detection paradigm. We found that low-level sensory properties were by far the best predictor of stimulus detection and found no evidence that detection was modulated by threat-sensitive processes. The data thus suggest that the fear advantage is most parsimoniously explained by low-level properties of the fear expression, negating the need to invoke the role of threat, or emotion sensitive processes.

Detectability is unlikely to be the only factor that has shaped the fear expression. Expressing fear enhances the expresser's field of view, eye movement velocity and nasal airflow- linking it to functional advantages in the context of threat (Susskind et al., 2008). Fear may also enhance the expresser's stimulus detection by shaping how light is cast onto their retina (Lee, Mirza, Flanagan, & Anderson, 2014). In other words, the expression of fear appears to be adaptive for both senders and receivers, in terms of efficient transmission and reception of visual information.

We observed a robust "face advantage" in both experiments. Normal faces were better detected than control faces, despite being equivalent in effective contrast. This suggests that the visual system is sensitive to stimuli that are specifically face-like, and this sensitivity is not yoked to awareness of the stimuli. It has been found in bCFS studies that stimulus inversion has a detrimental effect on the detection of human faces but no effect on detection of chimpanzee faces (Stein, et al., 2012). This suggests that pre-conscious visual processes selectively promote stimuli that resemble conspecific faces to conscious perception, presumably because of their social relevance. However, our data suggest that this sensitivity does not extend to the evaluation of facial emotion; emotional evaluations had no effect on stimulus detection beyond that explained by low-level image properties.

How can we reconcile a robust face advantage with the absence of emotion-sensitive processes? Determining whether a stimulus is a face represents a coarser-level judgement than identifying its emotional expression. Visual masking studies have shown that identifying a specific object requires substantially more processing time than identifying its general category, whilst determining an object's category co-occurs with its detection (Grill-Spector & Kanwisher, 2005). The present data are thus consistent with a framework in which the detection and categorisation of faces (i.e. face vs. non face), but not the evaluation of facial expression, is performed at an early processing stage by the same perceptual mechanisms. In contrast, the data are poorly explained by models

suggesting that emotional evaluation precedes and drives face detection (Palermo & Rhodes, 2007).

It is important to consider well-documented phenomena that appear to conflict with our “low-level” account of the fear advantage. One relevant example is that anxious populations exhibit enhanced processing of fear faces, which is commonly attributed to dysfunction in threat-sensitive mechanisms that operate without awareness (Bar Haim, Lamy, Pergamin, Bakermans-Kranenberg, & van IJzendoorn, 2007). However, differences in perceptual selection observed in anxious populations are not limited to threat relevant contexts. For instance, anxiety is associated with enhanced attentional capture by perceptually salient neutral stimuli (Moser, Becker, & Moran, 2012, Moran & Moser, 2014). Correspondingly, enhanced biases for fear faces in anxious populations could be a function of either the perceptual or emotional properties of the stimuli. Thus, processing differences displayed by anxious populations may not be inconsistent with our account.

Another interesting phenomenon is that eye gaze direction can modulate detection of fearful faces, such that averted fearful gazes are prioritised over directed fearful gazes (Milders, Hietan, Leppanen, & Braun, 2011). This makes good ecological sense in terms of perceived threat, since the presence of an unknown threat in the environment (averted gaze) may be more dangerous than a threat directly from the target (directed gaze). Importantly, however, Chen and Yeh (2012) found directly contradictory results using schematic faces, in which low-level variability is reduced. In a bCFS paradigm, Chen and Yeh found that schematic fearful faces with directed gaze were detected faster than those with averted gaze. Notably, the removal of the salient eye white in schematic stimuli also resulted in a lack of an overall “fear advantage” for detection. These opposing findings, likely due to simple physical variations between the particular stimulus sets employed, pose problems for accounts that posit specialised threat detection mechanisms as the cause of processing biases (see Becker, et al., 2011 for a related discussion).

Several studies have observed differential amygdala responses to fearful and neutral faces rendered invisible by masking and CFS (Jiang & He, 2006; Whalen et al., 2004), which has been interpreted as evidence that fearful faces are evaluated without awareness via a pathway involving the amygdala. However, whether this neural activity is linked to adaptive changes in perception is hard to determine without convergent behavioural measures. Importantly, patients with amygdala lesions nonetheless show prioritised detection of fear in bCFS and visual search tasks (Piech et al., 2010; Tsuchiya, Moradi, Felsen, Yamakazi, & Adolphs, 2009). Moreover, recent work using a CFS

paradigm suggests that attentional orienting to threat stimuli may be dependent on their conscious detection (Hedger, Adams, & Garner, 2015). Whether amygdala activity to unconsciously presented threat stimuli (in response to either low-level or affective properties) has a functional role in promoting their detection, therefore, remains an interesting question.

How do our data fit with suggestions that processing of threatening stimuli is driven by evaluation of content in the low spatial frequencies? (Mermillod, Droit-Volet, Devaux, Schaefer, & Vermeulen, 2010; Willenbockel, Leopre, Nguyen, Bouthillier, & Gosselin, 2012). Such observations are thought to support the notion that coarse, rapid, magnocellular input to the amygdala is sufficient for the evaluation and subsequent detection of threat stimuli (Tamietto & de Gelder, 2010; Vuileumier, Armony, Driver, & Dolan, 2003). Importantly, human contrast sensitivity is greater for low spatial frequencies, meaning that they are weighted more heavily in our effective contrast calculations. Thus, our data also suggest that low spatial frequencies are important, but that this relates to the distribution of contrast at these spatial scales, rather than the *evaluation* of the content at these scales.

In summary, our data suggest that, through evolutionary or learned adaptations, fearful faces are optimised to stimulate human sensory biases. This mechanism may provide a parsimonious explanation of the “fear advantage” in the competition for awareness that negates the need to invoke preconscious processes sensitive to threat.

Facial Emotion and Implicit Attention

5.a. Abstract.

The visual probe (VP) paradigm provides behavioural evidence that emotionally salient stimuli receive prioritised processing relative to neutral stimuli. Moreover, masked versions of the VP task suggest that this prioritisation may occur at unconscious stages of processing. Due to inadequate controls, it remains fundamentally unclear whether the prioritised processing of emotionally salient stimuli is driven by their affective salience, or simple low-level variability between emotive and neutral stimuli. Moreover, recent work has questioned the extent to which emotional face stimuli genuinely receive prioritised processing in the absence of awareness. In the present study, we presented stimuli under masked, continuous flash suppression (CFS) and normal viewing conditions, with awareness assessed according to sensitive signal detection criteria. Under normal viewing conditions, attention was allocated more to emotional than neutral face stimuli, although this was entirely explained by their low-level properties, rather than recognisable emotional content. Under CFS and masked presentations, we found no evidence for increased attentional allocation to emotional face stimuli. Moreover, an observer's awareness of the stimuli (as assessed by d prime) predicted the magnitude of attentional cuing. Together, the data suggest that the prioritised processing of emotional face stimuli is restricted to conditions of awareness, and may be parsimoniously explained by simple low-level variability between emotional and neutral face stimuli.

5.b. Contributions.

Experimental design, data collection, analysis and write up were completed by Nicholas Hedger under the supervision of Wendy Adams and Matthew Garner.

5.1. Introduction.

Human visual perception has limited capacity: at any given moment we are consciously aware of only a reduced set of sensory events; only a subset of stimuli presented to our retinae achieve this privileged status of “fame in the brain” (Dennett, 1993). To optimise behaviour, the visual system must select important stimuli for further processing. The visual probe paradigm provides evidence of this effect: An observer’s detection of a peripheral target is facilitated when preceded by an emotionally salient stimulus relative to a neutral stimulus (Bar Haim, Lavy, Pergamin, Bakermans-Kranenburg, & van Ijzendoorn, 2007). There is tremendous interest in understanding the mechanisms of this selection process - how does the visual system prioritize the stimuli that are most important to its survival?

One proposed mechanism, entrenched in evolutionary theories, is that humans possess an independent, sub-cortical visual pathway that operates without awareness and rapidly directs processing resources towards threatening stimuli (Garrido, 2012; Tamietto & de Gelder, 2010). This theory has intuitive appeal - it may take hundreds of milliseconds for retinal stimulation to generate a conscious percept (Koch, 2004). If threats could modulate an observer's behaviour rapidly and independently of their conscious registration, survival odds would be increased (Morris, Öhman, & Dolan, 1999). This notion is intriguing, because it suggests that there are specialised ways of (and independent neural substrates for) prioritising affective stimuli. Moreover, the idea has also influenced thinking about clinical disorders. For instance, dysfunction in the systems involved with preconscious threat detection are thought to underlie the hypersensitivity to threat and maladaptive perceptual biases exhibited by individuals with anxiety disorders (Mogg & Bradley, 1998; Ohman & Mineka, 2001).

Evidence for the unconscious prioritisation of threat has typically relied on measuring responses to stimuli that are presented to observers outside of awareness (Kim & Blake, 2005). A long history of observations from paradigms such as backward masking, binocular rivalry and continuous flash suppression (CFS) has revealed that threat stimuli suppressed from awareness can nonetheless elicit adaptive changes in neural activity (Jiang & He, 2006; Whalen et al., 2004; Williams, Morris, McGlone, Abbott, & Mattingley, 2004) and physiological arousal (Lapate, Rokers, Li, & Davidson, 2013; Ohman & Soares, 1994).

Behaviourally, the masked visual probe (MVP) has provided evidence that threat stimuli receive prioritized processing in the absence of awareness. On a typical trial, threatening and neutral target stimuli are presented on either side of a central fixation cross, before being rapidly replaced by salient mask stimuli. The brief stimulus onset asynchrony (SOA) between the target stimuli and mask stimuli (usually ~17 or ~33 milliseconds) means that observers typically report perceiving the mask stimuli, but not the preceding target stimuli (Wiens & Ohman, 2007): visual input of the target stimuli is dissociated from awareness. After presentation of the mask, a ‘probe’ (usually a small dot or arrow) is presented in either the valid location (the location preceded by the threat stimulus) or the invalid location (the location preceded by the neutral stimulus) and the observer is asked to report its location. If response times (RTs) to valid probes are faster than responses to invalid probes, it can be inferred that spatial attention has been preferentially drawn to the location of the threat stimulus (Chica, Martín-Arévalo, Botta, & Lupiáñez, 2014). Thus, the MVP paradigm can be used to measure the extent to which threatening stimuli receive prioritised processing in the absence of their conscious registration.

In a recent meta-analysis of the MVP paradigm (Hedger, Gray, Garner, & Adams, 2016) we found that the magnitude of threat-related bias (i.e. the valid vs. invalid RT difference) tends to be small (Cohen’s $d_z = 0.28$). Our analyses also suggested that the effect size was strongly modulated by stimulus visibility: the threat related bias was significantly larger if the SOA between stimuli and masks was >30 ms than if it was < 30ms. This is important, given that many observers achieve above-chance detection of 33ms targets when stringent signal detection measures of awareness are employed (Pessoa, Japee, Sturman, & Ungerleider, 2006; Szczepanowski & Pessoa, 2007) and suggests that unintended stimulus visibility increases threat-related biases. Moreover, we also found that this SOA-bias relationship was greater within paradigms that did not include an awareness check to verify that masking successfully eliminated stimulus visibility. Interestingly, this suggests threat related biases in the MVP paradigm could be modulated by, or perhaps even driven by residual awareness of the masked stimuli.

In the studies that have included an explicit awareness check to assess awareness of masked stimuli, this is usually implemented via a separate block of trials wherein observers complete an alternative forced choice (AFC) task such as discriminating between different masked stimuli (Carlson, Reinke, & Habib, 2009; Fox, 2002; Mogg, Bradley, & Hallowell, 1994). In general, if observers’ performance is not significantly better than

chance in this control task, it is concluded that any threat biases obtained during the experimental trials can be attributed to processes that occur independently of awareness of the threat stimuli.

Establishing null sensitivity to stimuli via a forced choice task in this way is associated with formidable practical and conceptual issues (Wiens, 2008). For instance, awareness checks in the MVP paradigm have typically lacked statistical power, suggesting that the likelihood of type 2 errors (failure to detect an observers residual discrimination of target stimuli) may have been at unacceptable levels. Our meta analysis found that, on average, across MVP experiments, observers were classed as unaware of stimuli if 2AFC performance was less than 68%. Importantly, this permits deviations from chance performance that are moderate in magnitude (Cohen's $h = 0.38$, see Cohen, 1977), which invalidates strong statements about truly 'unconscious' processing of the masked stimuli. Moreover, it is important to note that only one study employed a signal detection measure (d' - d prime) that corrected for individual response bias (Koster, Verschuere, Burssens, Custers, & Crombez, 2007). Taken together, these limitations suggest that more rigorous methods are needed to assess awareness⁷.

Another interesting question, which is increasingly receiving attention, is whether any effects induced by stimuli presented outside of awareness depend on the method used to manipulate awareness. For instance, it is possible that threat stimuli can modulate attention independently of awareness, but that the brevity of masked presentations degrades all processing of the target stimuli such that any attentional modulation is too small to be detected. Masking necessitates presentation times that are substantially briefer (< 40 ms) than those chosen to optimise attentional cueing effects in standard, supraliminal versions of the visual probe task (usually around 500 ms; (Bar-Haim, et al., 2007). Since the presentation of stimuli in the masked version of the visual probe paradigm is an order of magnitude briefer than in the standard version, this confounds any comparison between aware and unaware processing. A more direct comparison would require that subliminal

⁷ Notably, not all authors have claimed that observers were completely unaware of the masked stimuli, and have instead claimed that awareness has been "restricted" (Carlson & Reinke, 2008; Mogg & Bradley, 1999). Nonetheless, it remains a matter of contention, with theoretical importance, to determine whether threat stimuli attract attention under genuine conditions of unawareness.

stimuli are not so degraded, or temporally disadvantaged, relative to a supraliminal counterpart. Continuous flash suppression (CFS), which has been heralded as a “game changer” in the study of unconscious processing (Sklar et al., 2012) may provide one solution to this problem. In CFS, stimuli presented to one eye can be suppressed from awareness by presenting a dynamic noise pattern to the other eye (Tsuchiya & Koch, 2005). With appropriate presentation parameters, CFS can render stimuli invisible for several seconds, allowing ample time for unconscious processes to engage with the suppressed stimuli. The use of CFS in a visual probe paradigm may therefore provide a more suitable comparison between unaware and aware states.

Finally, a critical conceptual issue concerns the stimulus attributes that drive the prioritisation of threat stimuli across all paradigms: standard visual probe, masked visual probe and CFS. Although it has been demonstrated that certain classes of threat stimuli, such as fearful faces, are reliably prioritized, one idea gaining traction is that this prioritization may be better explained by their low-level properties than by threat-sensitive processes (Gray, Adams, Hedger, Newton, & Garner, 2013; Hedger, Adams, & Garner, 2015b; Stein & Sterzer, 2012). For instance, Lee and colleagues (2013) found that the increased luminance contrast resulting from the greater exposure of the scleral field in fearful faces (relative to neutral faces) was a good predictor of enhanced performance in an attentional cuing task. More recently, work from our own lab revealed that the relationship between a face’s amplitude spectrum and the human contrast sensitivity function was a better predictor of the prioritised detection of fear faces in masking and CFS tasks than perceived valence or arousal (Hedger et al., 2015b). This sensory advantage of the fear expression is particularly important, since fearful faces give rise to the largest, most reliable threat-related biases in the MVP paradigm of all stimulus types (Hedger et al., 2016). As highlighted in our meta analyses, it is critical that researchers provide adequate stimulus controls such that threat-related biases driven by the semantic content of stimuli (i.e. their affective content) are distinguished from effects due to simple low-level variability across stimuli.

The current study aims to address issues identified as contentious or unresolved in our analyses of previous literature. Specifically, we ask: i) Do emotionally salient stimuli modulate attention in normal viewing conditions? (ii) Do unconsciously presented emotional stimuli modulate attention when principled criteria are imposed to assess awareness? (iii) Are these biases explained by affective, or low-level variability

across stimuli? iv) Are these effects modulated by the method used to render stimuli perceptually invisible?

5.2. Method.

5.2.1. Participants.

From our meta analyses, we determined that 41 participants would be required to attain sufficient power (95%) to detect the attentional effects observed when fear and neutral faces compete in the MVP paradigm ($d_z = 0.58$). For this reason, data collection was terminated when 41 participants (9 male) had completed the experiment. All had normal or normal-to-corrected vision.

5.2.2. Stimuli.

Stimuli were four facial models, taken from the NimStim face set (Tottenham et al., 2009), depicting neutral, fearful and happy expressions. All stimuli were placed within an opaque elliptical mask to eliminate any external features and were equated in luminance and root mean squared (RMS) contrast. Face stimuli were presented in two configurations. *Normal* faces were presented upright with veridical contrast polarity. *Control* faces were rotated 180 degrees with reversed contrast polarity, producing an image similar to a photographic negative (see Figure 5.1). These manipulations severely disrupt the recognition

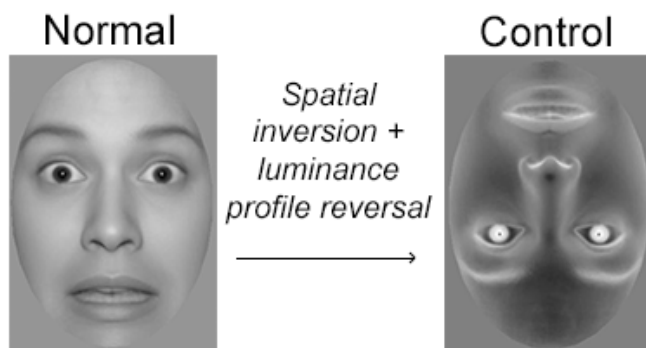


Figure 5.1. Face stimuli presented in the normal and control configuration.

and affective evaluation of facial expressions (Gray et al., 2013; Hedger et al., 2015b). Critically, however, they do not alter low-level stimulus properties (RMS contrast, mean luminance, amplitude spectra). Thus, if the valence of face images is critical in directing spatial attention, we would expect any effect of expression to be reduced or eliminated for the control images, relative to the normal images (i.e. an interaction between expression and stimulus manipulation). Conversely, if low-level properties of the stimuli explain the effect of expression, we would anticipate a similar main effect of expression for both normal and control stimuli (i.e. no interaction between expression and stimulus manipulation). All stimuli subtended 6.2 x 4.1 degrees of visual angle (DVA) at the viewing distance of 70 cm on a 1280 x 1024 pixel resolution, gamma corrected monitor. In all trials, observers viewed the display via a mirror stereoscope, and each eye's image was framed by a random dot surround (9.5 x 11.4 DVA) to control vergence.

5.2.3. Questionnaire measures.

Previous work suggests that attentional biases towards emotional stimuli are modulated by anxiety and related trait characteristics (Fox, 2002; Mogg & Bradley, 1999). Before the visual probe experiment, all observers completed the following measures of general and social anxiety: Trait Anxiety Inventory (STAI-T, Spielberger et al., 1983), Social Interaction Anxiety Scale (SIAS, Heimberg, Mueller, Holt, Hope, & Liebowitz, 1992) and Social Phobia Scale (SPS, Heimberg et al., 1992).

5.2.4. Procedure.

Each trial began with the presentation of a central fixation cross whose duration was randomly sampled from the range of 300 to 1000 ms to avoid anticipatory responses. Observers completed 560 trials in total. On '*signal*' trials (336 trials), pairs of face stimuli were presented to observers. On '*noise*' trials (224 trials), no face stimuli were presented to observers; intermingling these noise trials enabled concurrent evaluation of stimulus awareness (see '*noise trials*' section). There were three presentation conditions (Figure 5.2).

5.2.5. Presentation conditions.

Masked presentation.

In the *masked* presentation condition (Figure 5.2a - 112 trials), our trial sequence mirrored that of previous literature (Fox, 2002). Two face stimuli appeared either side of fixation for 17ms before being immediately replaced by two masks (patterns of high contrast ellipses) for 17ms. A 17ms SOA between face and mask has been commonly employed in previous MVP studies (Beaver, Mogg, & Bradley, 2005; Fox, 2002; Koster, Verschuere, Burssens, Custers, & Crombez, 2007; Mogg & Bradley, 1999, 2002), due to the refresh rate of standard cathode ray tube (CRT) monitors. Immediately after presentation of the mask, a dot appeared at the location preceded by the left or right face and observers were required to report its location as quickly and accurately as possible.

CFS presentation.

In the *CFS* presentation condition (Figure 5.2b - 112 trials), two faces were presented monocularly (counterbalanced across eyes) on either side of fixation for 500ms, whilst the other eye viewed a pair of dynamic masking patterns either side of fixation that updated at a rate of 10Hz. Immediately after, a dot appeared at the location preceded by the left or right face and observers were required to report its location as quickly and accurately as possible.

Standard presentation.

In the *standard* presentation condition (Figure 5.2c - 112 trials), two faces were presented monocularly (counterbalanced across eyes) on either side of fixation for 500 ms, whilst only the fixation cross and surround was presented to the other eye. We chose monocular presentation of face stimuli to allow a straightforward comparison with the CFS condition. Immediately after the face presentation, a dot appeared at the location preceded by the left or right face and observers were required to report its location as quickly and accurately as possible.

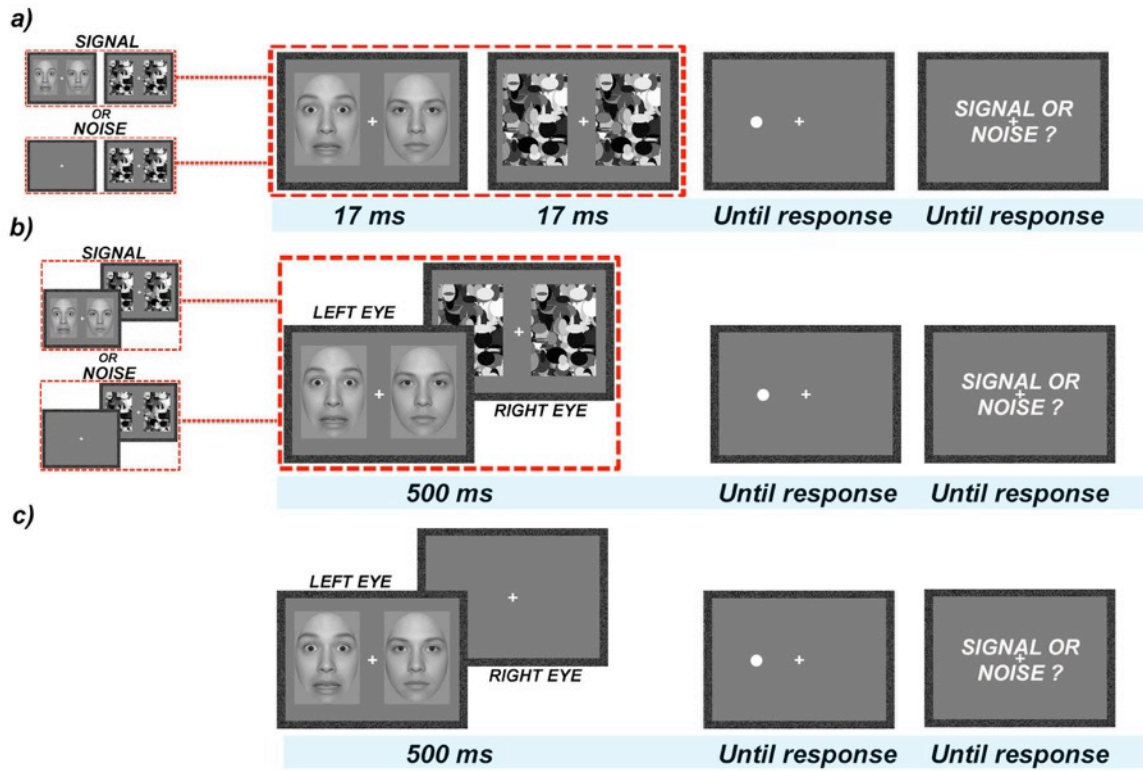


Figure 5.2. Schematic of trial sequences for the three presentation conditions.

a) Masked presentation b) CFS presentation c) Standard presentation. Masked and CFS trials had an equal number of signal trials (trials where face stimuli were presented) and noise trials (trials where no face stimuli were presented) – these are shown in the leftmost panels.

Within each presentation condition (*masked*, *CFS*, *standard*) there were two stimulus pairing conditions (see Figure 5.3).

5.2.6. Stimulus pairing conditions.

Emotion-neutral trials: emotional biases.


Mirroring conventional visual probe studies, *emotion - neutral trials* (64 trials), were designed to measure selective processing under conditions where a neutral and emotional stimulus compete for resources. The face presentation consisted of an emotional face (32 fear, 32 happy) presented to one side of fixation and a neutral face presented to the other. Within each emotion-neutral pair, half of the trials were valid (subsequent probe appeared in the location of the emotional face) and half were invalid (probe appeared in the

location of neutral face). These trials were repeated with face stimuli presented in both normal (16 trials) and control manipulations (16 trials).

Normal-control trials: face biases.

Normal-control trials (48 trials) were designed to measure selective attention when normal and control face stimuli (with matching emotional expression) compete for resources. In normal-control trials (48 trials), a normal face (16 neutral, 16 fear, 16 happy) was presented on one side of fixation and a face with the same emotion, but in a control configuration was presented at the other. Within each normal-control pair, half of the trials were valid (subsequent probe appeared in the location of the normal face) and half were invalid (probe appeared in the location of control face).

**a) EMOTION - NEUTRAL
BIAS**



**b) NORMAL - CONTROL
BIAS**




Figure 5.3. Schematic of the two stimulus pairing conditions.

a) Emotion-neutral trials - an emotional (fear happy) and neutral face were presented either side of fixation b) Normal-control trials - a normal (fear, happy, neutral) and control face (same expression) stimulus were presented either side of fixation.

5.2.7. Noise trials.

50% of the trials within the CFS and masked presentation conditions (112 masked, 112 CFS) were ‘noise’ trials. These trials were identical to signal trials, except no face stimuli were presented prior to the mask (for masked presentations) or to the opposite eye to the mask (for CFS presentations). If observers are unaware of the stimuli, they should perform at chance in discriminating signal trials from noise trials (Wiens, 2008). Thus, on each trial, after the observer reported the location of the probe, they were asked to indicate whether the preceding presentation had been a ‘noise’ trial or a ‘signal’ trial. It was clearly

explained to the participants that on 50% of the trials where a mask was presented, faces were also presented and that they had to discriminate these cases from those in which no faces were presented. Participants were also informed that there were no time constraints for this response and that they should prioritise accuracy over speed. The 224 trials for each presentation condition meant that the forced choice task had adequate (80%) power to detect very small deviations from chance performance (Cohens h of 0.16 or larger).

5.2.8. *Summary.*

All 41 observers completed 336 signal trials (3 presentation conditions (masked, CFS, normal) x 2 stimulus pairing conditions (i) emotion-neutral trials: 2 emotions (fear vs. neutral, happy vs. neutral) x 2 face manipulations (normal, control), (ii) normal-control trials: 3 emotions (neutral, fear, happy) x 16 repetitions) and 224 noise trials (112 masked, 112 CFS). The side of the emotional /control face, the eye of face presentation, the location of the probe and the validity of the probe were counterbalanced. Trial order was randomized for each participant.

5.3. Results.

5.3.1. *Sample characteristics.*

Sample characteristics are found in table 5.1.

Table 5.1. Summary of Sample Characteristics

Characteristic	$M (SD)$
Age	21.18 (3.54)
STAI-T	41.46 (10.20)
SIAS	20.24 (10.55)
SPI	19.87 (12.28)

5.3.2. *Assessing observers' awareness of the face stimuli.*

Following standard practice, d' values were computed from the difference between the z-transformed hit rates (proportion of signal trials that were correctly identified) and false alarm rates (proportion of noise trials that were incorrectly classified as signal trials). For masked presentations, d' was consistent with poor discrimination between signal and noise trials - at the group level, performance was not significantly better than chance ($M = 0.04$, $t(40) = 1.54$, $p = .130$). Importantly, no individual observer significantly exceeded chance performance in correctly classifying signal and noise trials (assessed via binomial test, upper binomial limit = 127 correct responses). For CFS presentations, sensitivity was slightly higher and significantly different from zero at the group level ($M = 0.06$, $t(40) = 2.55$, $p = .015$). At the individual level, two observers performed significantly above chance in distinguishing signal and noise trials. These two observers were excluded from all further analyses (with the exception of the correlation analysis shown in Figure 5.6). After removal of these observers, the group d' was not significantly different from zero for either masked ($M = 0.04$, $t(38) = 1.44$, $p = .158$) or CFS ($M = 0.04$, $t(38) = 2.01$, $p = .051$) presentations.

5.3.4. *Visual probe data.*

Data reduction.

Response times (RTs) corresponding to incorrect responses were removed (1.57% of RT data) and a log transform was applied to correct for skew. The mean log RT was calculated for each observer for each presentation condition and cue validity. Values that were more than 3 standard deviations from these means were defined as outliers and removed (1.93%). The analyses reported below were conducted on the remaining 96.51% of the RT data.

Emotion-neutral trials.

We calculated an emotional face bias score for each condition (invalid RT - valid RT) such that positive values indicate that attention is drawn to the location of the emotional expression. The resultant attentional bias scores are summarised in figure 5.4a. Figure 5.4b shows the attentional bias effect size (collapsed across expression and

manipulation) in each presentation condition. Although there were no main effects of expression or inversion on bias scores, and no interactions involving these variables, the effect of presentation condition approached conventional significance levels ($F(2,74) = 3.096, p = .051$). We detected a modest attentional bias towards emotional faces in standard presentations ($d_z = 0.36, 95\% \text{ CI } [0.04 \text{ } 0.69], p = .021$). However, in the masking and CFS conditions, the attentional bias effects were trivially small and did not reach significance (masking: $d_z = -0.14, [-0.46 \text{ } 0.17], p = .462$; CFS: $d_z = 0.07, [-0.24 \text{ } 0.39], p = .793$). Thus, we detected an attentional cuing effect towards emotional stimuli in standard presentations, where stimuli were fully visible, but we did not detect this in masking or CFS trials. Although we detected no effects involving expression, or manipulation on attentional bias scores, we examined the planned 2 way interactions between stimulus manipulation and expression in each presentation condition. This revealed that there was no significant interaction between expression and stimulus configuration in any presentation condition (all p s $> .67$). This suggests that facial emotion had no effect on attentional bias beyond that explained by basic low-level variability between expressions. In fact, the attentional bias effect in the standard presentation condition (widely reported in previous literature: Bar-Haim et al., 2007) was *larger* for control face stimuli (normal: $d_z = 0.26, 95\% \text{ CI } [-0.05 \text{ } 0.58], p = .114$, control: $d_z = 0.35, 95\% \text{ CI } [0.03 \text{ } 0.67], p = .033$).

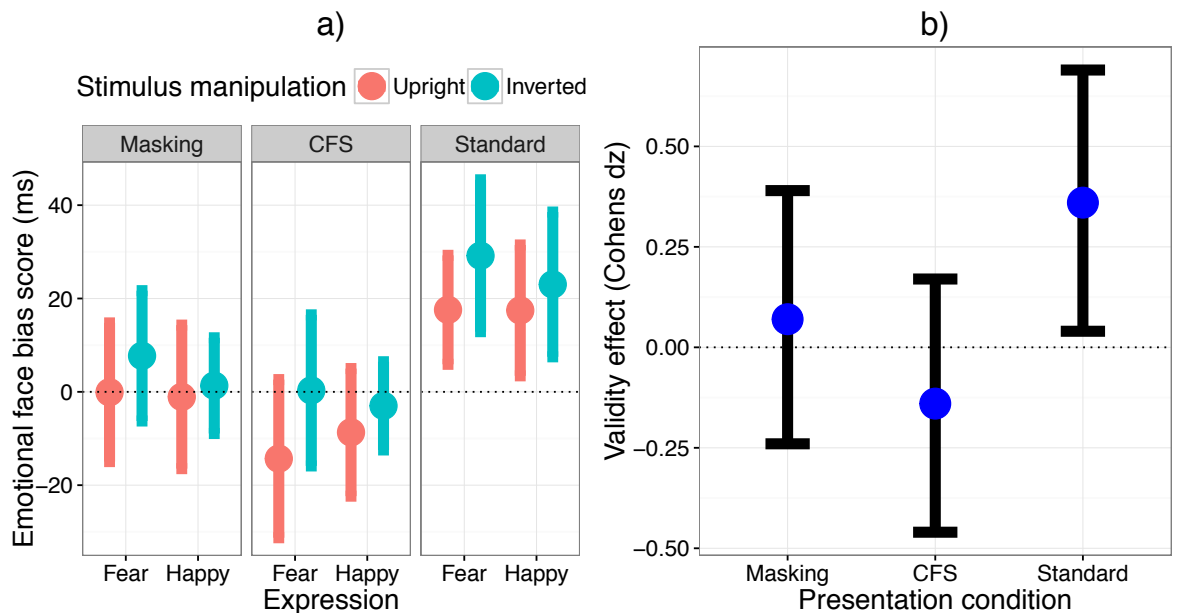


Figure 5.4. Attentional biases in emotion-neutral bias trials.

a) Attentional bias scores (invalid RT - valid RT) plotted as a function of expression, stimulus manipulation and presentation condition. Error bars are +/- 1 SE. b) The main

effect of validity in each presentation condition, expressed as Cohen's d_z is plotted as a function of presentation condition. Error bars are 95% confidence intervals

Normal – control trials

We calculated a normal face bias score for each condition (invalid RT - valid RT) such that positive values indicate that attention is drawn to the location of the normal face. The resultant attentional bias scores are summarised in figure 5.4a. Figure 5.5b shows the attentional bias effect size in each presentation condition (pooled across all stimuli). Mirroring the trend in emotion- neutral trials, there was a significant main effect of presentation condition ($F(2,74) = 6.08, p = .004$). The bias score was significantly different from zero in standard trials $d_z = 0.33, [0.01\ 0.65], p = .045$ and in CFS trials $d_z = 0.50, [0.17\ 0.83], p = .003$ but not in masking trials $d_z = -0.26, [-0.58\ 0.06], p = .125$. No further main effects or interactive effects on bias scores were detected.

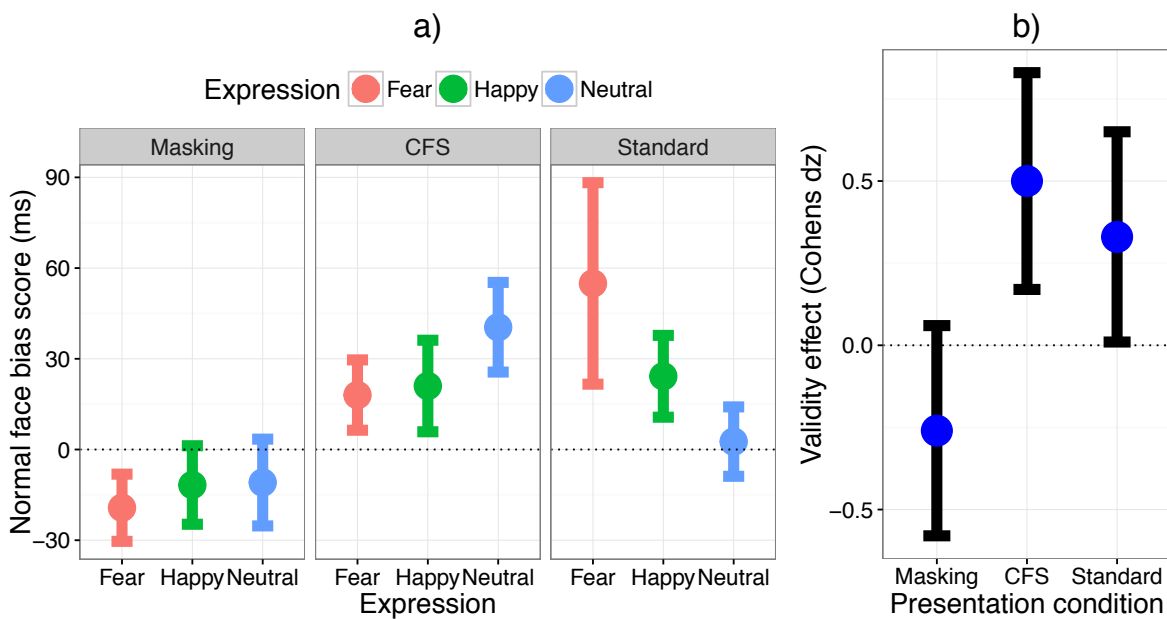


Figure 5.5. Attentional biases in normal – control bias trials.

a) Attentional bias scores (invalid RT - valid RT) plotted as a function of expression, stimulus manipulation and presentation condition. Error bars are ± 1 SE. b) The main effect of validity in each presentation condition, expressed as Cohen's d_z is plotted as a function of presentation condition. Error bars are 95% confidence intervals.

Association with awareness and anxiety measures

Our previous work (Hedger et al., 2016) suggests that some emotion related biases found in backward masking and CFS paradigms may be due to, or modulated by observers' awareness of the stimuli. To investigate this in the current study, we examined the relationship between awareness of the stimuli (as indexed by performance in the 'signal' vs. 'noise' discrimination task) and attentional bias effects in the masked and CFS paradigm conditions. All observers, including the two who performed at above chance level in the awareness task were included in these correlational analyses. For each observer, we computed a single attentional bias score, collapsed across all conditions, and a single d' score for the awareness task collapsed across CFS and masked presentations. Performance in the awareness task (i.e. the ability to distinguish 'signal' from 'noise' trials) was significantly and positively correlated with attentional bias ($F(1,39) = 8.452, R^2 = .178, p = .006$), as shown in Figure 5.6, suggesting that attentional biases are inflated when observers have some awareness of the stimuli. Notably, the best-fit line passes very close to (0,0), suggesting that awareness of the stimuli not only increases attentional bias, but may be required for attentional bias effects to occur. When data were split by paradigm, there was no significant difference between the CFS and masking paradigms in terms of the size of the correlation between awareness and attentional bias, and neither correlation reached significance when analysed alone (masking: $F(1,39) = 0.18, R^2 = .010, p = .675$, CFS: $F(1,39) = 2.53, R^2 = .056, p = .133$). When the data were split by stimulus pairing condition, performance in the awareness task predicted attentional biases in both emotion-neutral bias and normal control bias trials (emotion -neutral: $F(1,39) = 6.38, R^2 = .141, p = .015$, normal - control: $F(1,39) = 4.28, R^2 = .098, p = .045$).

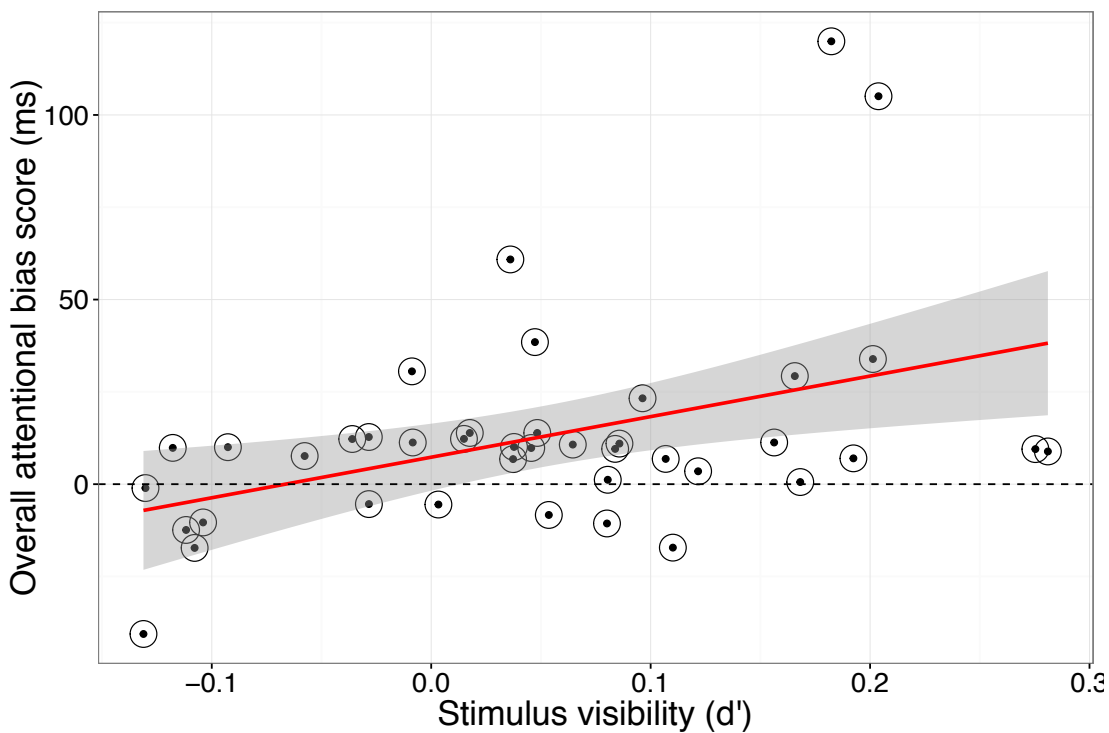


Figure 5.6. Association between overall d' and overall attentional bias score.

Red line is the least squares fit to the data, shaded region is ± 1 SE.

STAI-T did not predict attentional biases for any stimulus under any presentation condition. The only detected association involving the questionnaire measures was a negative relationship between SIAS and biases for normal fearful faces presented under standard viewing conditions.

5.4. Discussion.

Our experimental design allowed us to assess the extent of attentional orienting to various face stimuli under different conditions of awareness. A few key findings emerged:

- (i) In the standard, supraliminal paradigm, we found evidence that emotional faces attracted attention, when competing with neutral faces, and this was not significantly modulated by stimulus manipulation – it was actually larger for faces in which emotion is hard to identify – the control faces. No such effect of emotion on attention was found under masked or CFS conditions.
- (ii) Normal faces attracted attention over control faces

within the standard and CFS conditions, suggesting a preference for more natural, face-like stimuli. (iii) Importantly, within the masked and CFS paradigms, attentional effects were predicted by observers' awareness of the stimuli, suggesting that attentional biases are modulated, or even driven by awareness.

5.4.1. Attentional capture by emotionally salient stimuli.

Attentional biases towards emotionally salient stimuli were modulated by presentation condition. In standard, 500ms trials, we observed attentional biases towards emotionally salient stimuli, of a similar magnitude to those reported in previous, well-cited studies (Fox, 2002). Importantly, these effects were not reduced when the stimuli were presented in the manipulated condition (in fact, they increased slightly). This suggests that the apparent effect of emotion on attentional selection within the standard (conscious) visual probe paradigm is driven by low-level factors, and previous results may have been mistakenly attributed attentional effects to emotion-sensitive mechanisms.

When stimuli were suppressed from awareness via masking and CFS, emotion-related attentional effects were small and insignificant. However, as with any null result, it is worth discussing possible sources of a type 2 error. On a statistical level, it is important to reiterate that statistical power was high and the sample size calculations were based on a large body of previous literature (see 'participants' section). Secondly, in relation to the sample characteristics, the mean state anxiety level was relatively high ($M = 41.46$) and above the value expected to produce detectable biases towards threatening stimuli under subliminal presentations (Hedger et al., 2016). Unlike studies that have solely used masking to manipulate awareness, it is unlikely that null effects in unconscious presentations can be explained by simple restrictions on presentation time, since this was equated in normal and CFS trials (to 500 ms). Importantly, we detected significant attentional modulation effects under normal presentation conditions, suggesting that the task, in itself, was sensitive to modulations of attention. Our results are consistent with a number of studies that have failed to detect evidence for the emotional modulation of attention under masking (Fox, Cahill, & Zougkou, 2010; Koster et al., 2007) and CFS (Hedger, Adams, & Garner, 2015a).

5.4.2. Attentional preference for 'face like' stimuli.

We observed evidence for attentional biases for normal faces over manipulated, control faces in both standard and CFS trials. The latter finding is consistent with a large body of work from the breaking continuous flash suppression (bCFS) literature, which has consistently demonstrated that upright faces break CFS suppression faster / more frequently than inverted faces (Jiang, Costello, & He, 2007; Stein, Sterzer, & Peelen, 2012). Moreover, fMRI studies indicate that face selective regions of the temporal cortex still exhibit a measurable response that differentiates upright from inverted faces when presented under CFS (Jiang & He, 2006). Despite this previous evidence for enhanced sensitivity to face like configurations presented under CFS, ours is the first to provide behavioral evidence that they modulate spatial attention.

Why is it that CFS seems to spare selective attention to face like configurations (in face trials), but not emotional expressions (in emotion bias trials)? One possibility is that discriminating a face from a non-face (a coarse, basic-level classification) is an easier task than discriminating different expressions (a finer, sub-ordinate classification) and may thus be less affected by degradation associated with CFS suppression. Recent work using bCFS and backward masking paradigms suggests that any differential processing of subliminal facial expressions is better explained by their low level variability than evaluation of their emotional content (Hedger et al., 2015b). On the other hand, these concerns surrounding low-level confounds are thought not to apply to the differential detection of upright and inverted faces: low-level stimulus properties are preserved after spatial and contrast inversion. Instead, the prioritised detection of upright faces is assumed to reflect some higher-level ‘face- sensitive’ process. An alternative explanation, is that this preferential processing of upright faces does not reflect face-sensitive processes, but rather the fact that ‘top heavy’ patterns in general are more easily detectable, since humans have a robust upper hemifield advantage in basic visual sensitivity (Skrandies, 1987). In fact, recently it has been shown that upright ‘protofacial’ stimuli (a simple triangular configuration of dots, resembling the position of the eyes and mouth) break CFS more rapidly than their inverted counterparts (Akechi et al., 2015). Future work should aim to dissociate effects driven by face sensitive processes from effects driven by simple differences in sensitivity in the upper and lower hemifield. Clearly, these two possibilities have drastically different implications for the level and complexity of visual processing that transpires without awareness. This, in turn, suggests caution when inferring high level processing based on a preference for upright face configurations.

5.4.3. Implications for paradigms used to manipulate awareness.

A recent concern, that has been raised by many, is whether the perceptual suppression induced by techniques such as masking and CFS are functionally similar to those that may occur under natural viewing conditions (Blake, Brascamp, & Heeger, 2014; Hesselmann & Moors, 2015). If they are not, then studies employing these techniques may tell us about the peculiarities of the techniques used, rather than revealing any conclusions about unconscious processing that generalises to natural viewing conditions. In addition, a related concern is that conclusions emanating from different paradigms used to manipulate awareness may not generalise to one another (Dubois & Faivre, 2014) meaning that a null effect in one paradigm does not necessarily entail the absence of unconscious processing, since affirmative findings may be found with a different paradigm. Our findings strengthen these concerns. For instance, based on the data from the standard and masked presentations alone, one could conclude that an attentional preference for face-like configurations depends on their conscious registration, whereas when taken together with the CFS data, the conclusion is that the absence of effects may be due to the methodological limitations of the masking paradigm. Similarly, the absence of an emotion bias in CFS or masked presentations does not necessarily imply that emotional stimuli fail to modulate attention under all conditions of unawareness. For instance, Faivre, Berthet and Koudier (2012) found that affective priming was eliminated when primes were presented under CFS, but robust priming effects were observed when primes were rendered indiscriminable by crowding. The study of unconscious processing is thus highly susceptible to the error of ‘denying the antecedent’ when interpreting null effects.

5.4.4. Implications for assessment of awareness.

If attentional cuing operated independently of awareness of the cuing stimuli, we should expect no association between discrimination of stimulus presence and the magnitude of the attentional cuing effect. Instead, our data revealed evidence that increased stimulus awareness (as assessed by d') predicted attentional biases, despite the limited range of d' values and the relatively low level of sensitivity. Recent research employing stringent signal detection measures of awareness have revealed that observers are more capable than previously thought at detecting brief, masked signals. In fact, one study has shown that the majority of observers can reliably detect images of fearful faces that are

masked after 25, or even 17 ms (Szczepanowski & Pessoa, 2007). Although these deviations from chance performance were small, they are non-trivial in the context of the attentional effects emanating from the masked visual probe paradigm, which are also very small. This, taken together with our own data, illustrates the importance of providing sensitive, well-powered and objective awareness measures.

5.4.5. Implications for control of low level confounds.

We observed no interactions between expression and stimulus manipulation in emotion -neutral trials. This suggests that recognisable emotion had no further effect on attention beyond that explained by low-level variability across expressions. In fact, in emotion –neutral trials, attentional cuing effects were *larger* for control faces than for their normal counterparts. This conflicts with the idea that emotional valence drives attention towards facial expressions, since control faces have vastly reduced recognisable emotional content (Gray et al., 2013). Considering this in the light of our previous work, we suggest routine use of control stimuli to verify that apparently ‘emotional’ modulations of attention are not driven by simple, low-level variability that exists between faces of different expressions. Fearful and happy facial expressions in particular tend to have more contrast energy at the spatial scales humans are sensitive to, relative to neutral faces (Hedger et al., 2015). Thus, a processing advantage for these expressions is predicted on the basis of simple sensory factors alone. To strengthen the conviction that emotion sensitive processes genuinely influence attentional selection, one needs to demonstrate that effects of expression are stronger in the normal than control configuration, or that they are associated with affective ratings of the stimuli (e.g. via valence and arousal reports).

5.4.5. Conclusion.

In conclusion, our data suggest that attentional biases to emotionally salient stimuli are predicted by awareness. We detected attentional cuing effects under normal viewing conditions, but not under two different conditions of unawareness. Moreover, we provide direct evidence that an observer’s awareness of stimuli predicts the magnitude of attentional cuing effects. Finally, even under full awareness, we found that attentional cuing effects towards emotionally salient stimuli was fully accounted for by low-level

stimulus confounds. When considered alongside our meta-analysis, these findings could motivate a reinterpretation of previous literature and stimulate further well-controlled studies on the relationship between attention and awareness.

Discussion

6.1. Motivation for Thesis.

A fundamental constraint on human visual processing is its limited capacity. To enable adaptive functioning, the visual system must prioritise processing resources to promote perception of stimuli that will facilitate successful interactions with the environment. Although this selection process is a critical feature of the visual system, it remains poorly understood. This subject has broad appeal, as it raises challenging questions about how ‘value’ or ‘meaning’ are assigned to sensory signals and how this shapes our perception of the world. The topic is also intriguing because it is aligned with some of the most entrenched themes in cognitive psychology, such as modularity – the idea that there are specialised ways of (and neural pathways for) processing affective stimuli (Ohman & Mineka, 2001). Over the past two decades, since the scientific study of emotion has become ‘mainstream’ (e.g. Fox, 2008; Panksepp, 2004), there has been a proliferation of studies in the domain of ‘affective vision’; there is tremendous interest in uncovering how emotionally salient stimuli are prioritised by the visual system. *Chapter 1* reviewed the long history of observations that have shaped the pervasive ‘standard hypothesis’: that emotionally significant stimuli are processed via an independent subcortical pathway that operates rapidly, automatically and independently of conscious awareness.

Although the standard hypothesis has long been criticized, sometimes vigorously (Pessoa, 2005; Pessoa & Adolphs, 2010), its basic tenets, i) that there is a specialised sub-cortical pathway for processing affective stimuli, and ii) that affective information can be processed unconsciously, still continue to shape contemporary theory, experimental work and clinical practice. On a surface level, this standard hypothesis seems plausible for a number of reasons:

1. The notion of a subcortical pathway is appealing because it is assumed to be faster than a cortical one, and rapid processing of affective stimuli is adaptive.
2. Conscious perception has a limited capacity (Marois & Ivanoff, 2005). In an optimal system, the processing of affective stimuli would be less constrained by these capacity limits.
3. By extension, conscious awareness is probably a relatively recent evolutionary development (Reber, 1992) and it is unclear why it is required for efficient functioning (Chalmers, 1995). It therefore seems possible that brain function in phylogenetically older, sub-cortical structures is less associated with conscious experience.
4. Many lower species exhibit adaptive responses to predators and prey, but have entirely decentralised nervous systems: they are nothing more than simple visuo-motor stimulus response systems that seem very unlikely to support subjective experience (Barron & Klein, 2016). Hence, in many species, adaptive responses to salient stimuli seem ‘automatic’ and can occur without any central integration of sensory inputs. It is possible that similar mechanisms have been conserved by evolution.

Despite these arguments regarding the plausibility of the standard hypothesis, one needs to provide evidence. To investigate the scope of affective visual processing that transpires unconsciously, psychologists have developed several techniques for dissociating visual input from awareness. *Chapter 1* provides a review of this literature, asking the question: *to what extent are emotional visual stimuli processed without conscious awareness?* The conclusions stemming from many of these studies is that emotionally salient stimuli can elicit adaptive behavioural, physiological and neural responses when presented outside of awareness. However, the reviewed evidence for the standard hypothesis was undermined by three pervasive issues, which were addressed in subsequent chapters:

1. **Inadequate assessment of awareness.** In a very real sense, the literature is inconsistent. Much of inconsistency between findings that support or challenge the standard hypothesis may be explained by different ways that

awareness has been characterised. *Chapter 1* revealed that many studies supporting the standard hypothesis used insufficiently rigorous measures of awareness. This was explored *quantitatively* via the meta analysis reported in *Chapter 2*, which revealed that threat related biases may be predicted by unintended stimulus visibility. Using a modified CFS paradigm, the extent to which threat stimuli modulate adaptive changes in behaviour and physiology was assessed experimentally using stringent and well-powered objective and subjective measures of awareness (*Chapter 3*). Finally, attentional cuing by emotional stimuli presented under masking and CFS was explored using signal detection criteria to assess awareness (*Chapter 5*).

2. **Poor control of low-level confounds.** Many of the behavioural studies reviewed in *Chapters 1 & 2* indicate that affective (particularly threatening) stimuli presented outside of awareness receive a processing advantage. Notwithstanding issues surrounding awareness, in many studies, the stimulus characteristics that lead to an unconscious processing advantage often remain unclear. Specifically, it is possible that a processing advantage for emotional stimuli could be driven by simple low level confounds, such as variability in luminance and contrast (Gray, Adams, Hedger, Newton, & Garner, 2013). Critically, if this were the case, the data are parsimoniously explained without the need to invoke unconscious processes sensitive to threat. *Chapters 4 and 5* tested whether the widely reported ‘fear advantage’ is better explained by the low-level or affective properties of the stimuli, using a combination of image analyses and behavioural data from a CFS and masking paradigm.
3. **Differences between paradigms used to manipulate awareness.** Two of the most prominent methods for manipulating awareness: masking and CFS, rely on separable mechanisms, making it likely that they also differ with respect to how they affect the representation of emotionally salient stimuli. Indeed, analyses presented in *Chapter 2* demonstrated that they have different consequences for the prioritisation of threat stimuli. Despite this, there are relatively few behavioural studies that have directly compared effects emanating from both paradigms (Faivre, Berthet, & Koudier, 2012). Such a comparison is important, as it reveals whether one can take the results from one paradigm, and make valid, general conclusions about

unconscious processing, or whether the scope of unconscious processing is dependent upon the peculiarities of each paradigm. In *Chapters 4 and 5* we compare the prioritised processing of threat stimuli when presented under masking and CFS.

6.2. Key Findings, Implications and Future Research.

6.2.1. Inadequate awareness measures in previous work.

Although the notion of threat sensitive processing without awareness has been persuasive to many (Tamietto & de Gelder, 2010), others have argued that the evidence base may be undermined by inadequate assessment of awareness (Hannula, Simons, & Cohen, 2005; Pessoa, 2005). The meta analysis, reported in *Chapter 2*, represents the first attempt to quantitatively characterise the extent of the issue. In the masked visual probe (MVP) paradigm, it was found that a substantial proportion of reported effects emanated from conditions where no awareness check was conducted to verify that observers were genuinely unaware of the stimuli. Awareness checks were also found to be underpowered, and insensitive to medium deviations from chance performance. Moreover, threat related biases were larger when stimuli were presented for >30 ms than < 30 ms and this effect of presentation time was larger under conditions when no awareness check was conducted. This strongly suggests that threat related biases are associated with unintended stimulus visibility.

In the binocular rivalry (BR) paradigm, it was found that threat related biases were quite limited at the initial, unconscious stages of perceptual selection – initial dominance effects were substantially smaller than total dominance effects. Analyses also suggested that threat related biases were greater under conditions associated with incoherent rivalry periods (i.e. piecemeal rivalry) and response bias. In the breaking continuous flash suppression (bCFS) paradigm, many studies failed to decouple differential response times driven by unconscious processing from those driven by processes that occur when images are visible to the observer.

In sum, the inadequate assessment of awareness was found to be a recurrent issue within every paradigm under study, which seriously undermines strong claims about ‘unconscious processing’. This analysis is timely and important, because the notion of

unconscious emotion-sensitive processing has become accepted as ‘common knowledge’ in recent textbooks (Goodale & Milner, 2013) and narrative reviews (Axelrod, Bar, & Rees, 2015). Moreover, current theoretical frameworks suggest that essentially any high-level process that can be performed consciously can also be performed unconsciously (Hassin, 2013). The analyses reported in *Chapter 2* caution against these simplifying assumptions and emphasise the growing body of recent null results in the literature (Hedger et al., 2015a; Hesselman & Moors, 2015; Lahteenmaki et al., 2015) as well as the significant issues associated with establishing conditions of unawareness.

6.2.2. No evidence for adaptive responses to emotionally salient stimuli when stringent, objective awareness measures are employed.

The above limitations concerning the assessment of awareness, as described in *Chapter 2*, required that such findings were re-examined with more rigorous and sensitive awareness measures. This shaped the approach to *Chapters 3 and 5*.

In *Chapter 3*, threat stimuli were presented under CFS and employed both subjective and objective awareness measures. This paradigm allowed measurement of physiological arousal, spatial attention and access to awareness. In trials where observers reported awareness of stimuli, threat images modulated spatial attention and elicited a larger skin conductance response than non-threat images. However, these effects were eliminated when observers were unaware of stimuli according to subjective and objective criteria. Moreover, no evidence was found to support the idea that threat images broke suppression and entered awareness more readily than non-threatening images.

In *Chapter 5*, emotional face stimuli were presented under both masking and CFS and concurrent signal detection measures were employed to assess awareness. No evidence was found for an attentional bias towards face stimuli presented under CFS or masking but a significant attentional effect was detected when observers were fully aware of stimuli in standard presentations. Moreover, it was found that an observer’s awareness of stimuli, as assessed by *d* prime, predicted the magnitude of attentional cuing towards stimuli presented in the masked and CFS conditions.

In sum, the data from *Chapters 3 and 5* are consistent with the idea that adaptive responses to emotionally salient stimuli are predicted by, rather than independent of awareness. How can these results be reconciled with previous, apparently contradictory

literature? As reviewed in *Chapter 1*, much of the evidence base has relied on responses to backward masked stimuli, with awareness assessed via subjective measures. Some authors have argued that subjective awareness measures may permit ‘weakly conscious’ processes – cases where observers are aware of stimuli, but it is below their criterion for subjective report (Snodgrass & Shevrin, 2006). Observers can be under-confident about their perceptual experiences (Bjorkman, Juslin, & Winnman, 1993) and may well interpret this uncertainty as a lack of perception – particularly if the visual signal is brief – as is the case with masking. As such, many behavioural researchers agree that subjective measures provide weak evidence for unawareness (Pessoa, 2005), if any evidence at all (Holender, 1986).

A possible interpretation then, is that the objective measures employed in *Chapters 3 & 5* simply probe a more elementary form of awareness that is less amenable to the response strategies associated with subjective reports. Indeed, subjective measures have been empirically demonstrated to be less sensitive than objective measures in the context of masking. Szecepanowski and Pessoa 2007 found evidence for a ‘dissociation zone’, whereby if masked fear faces are presented for durations between 17 and 25 ms, observers were above chance at discriminating between presence and absence of fear faces (objective awareness) but failed to discriminate between their correct and incorrect responses, as assessed by confidence ratings (subjective unawareness).

The discrepancy between our data and previous studies may not be as simple as a distinction between subjective and objective thresholds. Indeed, several studies employing objective measures of awareness have found evidence for unaware emotional perception (Jiang & He, 2006; Raio, Carmel, Carrasco, & Phelps, 2012). Instead, the difference may be explained by the *sensitivity* of our objective awareness measures. Both the objective awareness checks conducted in *Chapter 3 and Chapter 5* are sensitive enough to detect deviations from chance performance that are an order of magnitude smaller than those detectable by any previous visual probe studies (*as reviewed in Chapter 2*). This means that observers classified as ‘objectively aware’ in our studies may have been characterised as ‘objectively unaware’ in previous studies due to a lack of statistical power. Such false negatives are particularly dangerous in this literature, since null results are critical for claims that observers are unaware of stimuli (Vadillo, Konstantinidis, & Shanks, 2016).

Regarding false negatives, it is also important to note that the null findings reported for the unaware presentations of *Chapter 3 and 5* do not simply reflect a lack of sensitivity within the tasks to detect enhanced processing of emotionally salient stimuli. Both tasks

were well powered, according to analyses of previous literature (*see Chapter 2*) and both detected prioritisation of emotionally salient stimuli under normal viewing conditions. Instead, our data are consistent with the idea that adaptive responses to affective stimuli are eliminated under objectively unaware conditions.

6.2.3. The processing bias for fearful faces is best explained by low-level properties.

The meta analysis in *Chapter 2* revealed that the only stimuli to elicit a reliable, detectable threat-related bias, across the MVP, BR and bCFS paradigms were fearful faces. Removing fearful faces from the analyses substantially reduced the pooled threat related bias. By contrast, angry faces yielded substantially smaller, non-significant and even negative effects. Although the selective processing advantage for the fearful faces is widely reported in the literature as evidence for the standard hypothesis (Mendez Bertolo et al., 2016; Ohman & Mineka, 2001), the broader pattern of results makes little sense in the context of survival. There is no intrinsic survival value in narrowly tuned, exclusive sensitivity to fearful faces; fear faces signal a conspecific response to the presence of threat in the environment – they are not directly threatening themselves. It is arguably more important to be sensitive to the classes of *directly* threatening stimuli that elicited the fear expression in the first place (animal attack, angry faces). Thus the large bias for fear faces, relative to other threat stimuli, is not well explained by unconscious processes that have evolved to promote survival.

In the analyses of the bCFS literature, it was found that the large processing advantage for fear expressions was entirely replicated (and of the same magnitude) when faces were presented in a spatially inverted configuration. Given that the recognisable threat content of fear expressions is significantly reduced when presented in an inverted configuration (Gray et al., 2013; Hedger et al., 2015b), this suggests simple low-level variability between expressions may drive this fear advantage. In *Chapter 4*, the source of this low-level variability was investigated. In both a masking and CFS paradigm, it was found that the prioritised awareness of fear faces is predicted by their effective contrast: the relationship between their Fourier spectrum and the contrast sensitivity function. Fear faces have higher effective contrast than neutral faces and this, not their threat content predicts their prioritised access to awareness. The fear advantage was poorly explained by

perceived emotion; it was unchanged for control stimuli with the same effective contrast, but reduced recognisable emotional content. Importantly, the fear advantage was also not predicted by observers' affective evaluations of the stimuli (valence, arousal, dominance). Critically, this completely negates the need to invoke unconscious processes sensitive to threat. Instead, the data suggest that the fear expression has evolved to be salient to our general-purpose sensory mechanisms.

This purely 'low level' account of the fear advantage was again supported by work presented in *Chapter 5*. Attentional biases for emotional expressions were not altered when stimuli were presented in the control configuration. In fact, attentional biases were marginally *larger* for control stimuli than normal stimuli, suggesting that emotion has no effect attentional allocation beyond that explained by basic low-level variability across stimuli.

Taken together, the findings from *Chapters 2, 4 and 5* have major implications for the interpretation of previous literature. Importantly, by far the most commonly used stimulus class, when investigating the prioritised processing of threat, is fearful faces (accounting for 37% of the included effect sizes). Despite the abundant interest in this phenomenon, relatively few studies have included control stimuli that dissociate the contributions of low-level and affective properties of the stimuli to this processing advantage. Therefore, given that fear faces have a robust, purely sensory advantage over neutral expressions (*Chapter 4*), these studies cannot unequivocally claim evidence for unconscious threat perception. This is a substantial concern, given the volume of evidence for unconscious prioritization of threat stimuli that involves fearful faces.

6.2.4. Implications for paradigms used to manipulate awareness.

John Bargh and Ezequiel Morsella (2008) have previously argued that "assessing the unconscious in terms of processing subliminal stimuli is analogous to evaluating the intelligence of a fish based on its behaviour out of water". The principle reason for this argument is that common methods for manipulating awareness do so via stimulus conditions that rarely occur in natural viewing. Indeed, binocular rivalry and CFS occur *precisely because they violate a natural constraint*- that two different stimuli cannot occupy the same time and space. Any paradigm that renders a stimulus perceptually

invisible does so by impeding the normal processing of a stimulus in some way (Fogelson et al., 2014; Yuval-Greenberg & Heeger, 2013).

One way of tackling this issue is to employ a number of suppression paradigms, with different mechanisms. If findings generalise across paradigms, this strengthens the conviction that the research is illuminating fundamental features of perception. In *Chapter 4* we found that fear faces were most efficiently detected from both masking and CFS, however, this simply reflects the fact that suppression in both paradigms is strongly contrast dependent. In *Chapter 5* we found that both masking and CFS eliminated a processing advantage for emotional over neutral expressions. One paradigm-related difference was observed, however: observers showed an attentional preference for normal faces over control faces suppressed via CFS, but not via masking. This may be because it is a necessary, but undesirable requirement of masking that presentation time is restricted, which may have selectively muted the attentional effects in masked presentations. This reflects the broader point that, regardless of how robustly a paradigm manipulates awareness, it will always have unintended consequences and disadvantages. For instance, binocular disparity cues provide important information about the proximity of threats and have ‘camouflage breaking’ qualities that are important in segregation and object detection (Wardle, Cass, Brooks, & Alais, 2010). These ecologically germane aspects of threat detection cannot be investigated with CFS or rivalry given the necessity of presenting threat stimuli to only one eye.

A fruitful avenue for future research may be to investigate paradigms that generalise more readily to natural viewing, such as crowding – which occurs frequently in cluttered, naturalistic environments. Alternatively, motion induced blindness clearly relies on separable mechanisms to masking and CFS, since the dependence on luminance contrast is reversed – higher contrast targets actually experience more robust suppression (Bonneh, Cooperman, & Sagi, 2001).

6.2.5. Implications for anxiety and other disorders.

This thesis has important implications for models of anxiety, which state that preconscious or ‘automatic’ processing biases for threat stimuli are involved in the aetiology/ maintenance of anxiety disorders (Mogg & Bradley, 1998). These models derive their evidence base almost entirely from findings employing backward masking, which

have revealed that anxious individuals exhibit a larger attentional bias to masked threat, as indexed by MVP and masked Stroop tasks (Bar Haim et al., 2007). Indeed in *Chapter 2*, via meta-regression, it was found that previous MVP studies supported an association between the magnitude of threat bias and an observer's trait anxiety level. However, relying on a small number of well controlled, but highly circumscribed paradigms in this way precludes any firm claims about real world significance. By contrast, if the phenomenon under study is found to hold across diverse stimuli and perceptual phenomena, these claims become more reasonable. However, this effect did not generalise, since no such association was observed in the binocular rivalry paradigm. This novel analysis invites the field to revisit conclusions based on masking studies alone, particularly considering the issues associated with the assessment of awareness discussed in section 6.2.1 and 6.2.2.

Chapters 4 and 5, which emphasise the importance of low level confounds, also have implications for the anxiety literature. For instance, it is routinely reported that anxiety is associated with an elevated processing bias for fear faces (Capitao et al., 2014; Fox, 2002; Gray et al., 2009). This is often interpreted as evidence that anxious individuals are hyper-sensitive to threat. However, unless adequate controls for low level confounds are implemented, all these studies may actually be showing is that anxious individuals have difficulty ignoring salient stimuli, *regardless of their threat relevance*. This alternative explanation is aligned with recent proposals that attentional dysfunction in anxious individuals is not limited to threat relevant contexts but reflects a generalized deficit in attentional control (Moser, Becker, & Moran, 2012).

The dangers of low-level confounds may be much more serious than typically appreciated when studying populations who might be expected to show *diminished* threat processing, such as patients with a recent brain injury (Tsuchiya et al., 2009), or individuals who have received an intervention to alleviate anxiety symptoms (Murphy, Downham, Cowen, & Harmer, 2008). A processing bias for fear faces in these populations may be wrongly as interpreted as diagnostic of 'unimpaired threat processing' or a 'failed intervention'. In reality, all these observations may tell us is that these observers have a normal contrast sensitivity function and are thus more sensitive to fear faces. For this reason, it is essential that adequate controls are provided when studying clinically relevant differences in the processing of emotional stimuli.

6.2.6. *Implications for neuroimaging.*

Although this thesis did not employ neuroimaging methods, its findings have important implications for the interpretation of neuroimaging literature. As reviewed in *Chapter 2*, a mounting literature suggests that emotionally salient stimuli suppressed from awareness via masking, binocular rivalry and CFS elicit responses in the amygdala (Jiang & He, 2006; Whalen et al., 2004). Any two stimuli that differ in one or more dimension may elicit differential activity at the neural level. However, what is the purpose of neural sensitivity to these stimuli if it does not lead to adaptive changes in physiology or perception? *Chapters 3 and 5* provide no evidence for adaptive physiological and behavioural responses when stimuli are reliably suppressed from awareness via masking and CFS. It is possible therefore, that these neural signatures may be unrelated to prioritized processing of emotionally salient stimuli. Indeed, behavioural data from lesion studies often conflict with fMRI data, such that a brain region thought to be critical for a process is shown to be entirely dispensable. For instance, observers with lesions to the amygdala nonetheless show prioritized perception of threat stimuli in a number of behavioural tasks (Tsuchiya et al., 2009; Piech et al., 2010). Concomitant, well-controlled behavioural tasks are therefore essential in interpreting neuroimaging data.

A second, related possibility, raised in *Chapter 4* is that differential responses in the amygdala may be modulated by low-level properties of the stimuli. As with the behavioural literature, the most commonly studied comparison in the neuroimaging literature is between fear and neutral faces. However, almost all of these studies have failed to include control stimuli to rule out the influence of low level confounds. As evidence suggests that the amygdala is involved in ‘high level’ abstract functions such as processing ‘value’ or ‘relevance’ (Adolphs, 2013), the assumption is that it is less tuned to ‘low level’, sensory properties of stimuli, such as contrast. However, there is good evidence that crude sensory properties, such as the abruptness of auditory tones (Herry et al., 2007), flicker (Straube et al., 2010), and the sharpness of contours (Bar & Neta, 2007) can modulate amygdala activity. Thus, without adequate controls, it is impossible to claim that differential amygdala responses are driven by affective, rather than low-level sensory properties of stimuli. Thus, neuroimaging techniques may be beneficial in uncovering the neural mechanisms that drive unconscious prioritization of emotionally salient stimuli, but only when combined with appropriately sensitive measures of awareness, controls for low

level confounds and concomitant behavioural methods that are sensitive to enhanced perceptual selection.

6.3. Concluding Remarks.

The primary function of the visual system is to promote successful interactions with the environment. One of the most critical of these interactions is to avoid threats to survival and wellbeing. Despite the obvious, real-world importance of threat detection, this is an area that remains poorly understood. Scientists and lay-people alike have long been fascinated by the notion that affective visual stimuli hold a ‘special’ status in the visual system and can be registered without awareness to influence our thoughts and behaviour. This ‘standard hypothesis’ remains prominent - all but accepted as fact - despite the major methodological challenges involved in testing it, and contradictory experimental findings.

Chapter 2 represents the first systematic, quantitative review of empirical work that has explored the extent to which unconsciously presented threats modulate behaviour. This and the subsequent empirical chapters converge on two central ideas. The prioritisation of affective stimuli was found to be i) either associated with, or restricted to conditions of awareness (*Chapter 3, Chapter 5*) or ii) parsimoniously explained by simple low-level confounds (*Chapter 4, Chapter 5*). In conclusion, the work presented in this thesis seriously undermines the standard hypothesis and should caution researchers against uncritical acceptance of previous, apparently supporting evidence.

Appendices

Appendix A: Chapter 2. Summary of Excluded Paradigms.

Table A1. Summary of Excluded Paradigms

Paradigm used to manipulate awareness.	Reasons for exclusion
Visual crowding (e.g. Koudier, Berthet & Faivre, 2011)	Research employing this technique has typically investigated semantic priming, rather than changes in perceptual selection in response to stimuli rendered invisible by crowding (e.g. Faivre, Berthet, & Koudier, 2012; Koudier, Berthet, & Faivre, 2011).
Motion induced blindness (MIB: e.g. Bonne, Cooperman, & Sagi, 2001)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by MIB. This is probably due to practical constraints such as the upper limit on the retinal size of stimuli (~1DVA) that can be rendered invisible by this method (see Bonne, Cooperman, & Sagi, 2001, figure 2b).
Suppression by transcranial magnetic stimulation (TMS: e.g. Jacobs, de Graaf, Goebel, & Sack, 2012)	This research has primarily focused on subjective awareness during simple discrimination tasks (e.g. Corthout, Utti, Ziemann, Cowey & Hallett, 1999; Jacobs, de Graaf, Goebel, & Sack, 2012). There have been some attempts to disrupt processing of emotional stimuli (faces/ bodily postures) by TMS, but these did not assess changes in perception induced by the suppressed stimuli (Filmer & Monsell, 2013; Jolij & Lamme, 2005).
Chromatic flicker fusion/ dichoptic colour masking (e.g. Jiang, Zhou & He, 2007)	We are not aware of any studies that have measured a threat-neutral processing difference to stimuli rendered invisible by chromatic flicker fusion. This may be due to practical constraints, such that stimuli have to be monochromatic, low-contrast and low spatial frequency (Schurger, Pereira, Treisman, & Cohen, 2010).
(single) Flash suppression (e.g. Wolfe, 1984).	We are aware of one study that has recorded behavioral responses to stimuli rendered invisible by flash suppression, but this did not allow a comparison between threatening and neutral stimuli (Krieman, Fried, & Koch, 2002).
Generalised flash suppression (GFS: e.g. Wilke, Logothetis, & Leopold, 2003)	We are not aware of any studies that have measured a threat-neutral processing difference to stimuli rendered invisible by GFS. This method may not be optimal for subliminal perception research given that invisibility depends on several seconds of prior adaptation to a visible target stimulus (Wilke, Logothetis, & Leopold, 2003).
Attentional blink (e.g. Maratos, 2011)	In the attentional blink paradigm, the behavioral performance measure is usually whether a first stimulus (T1) suppresses perception of a second stimulus (T2) (e.g. Maratos, 2011; Vermeulen, Godefroid, & Mermillod, 2009). Relatively few studies have investigated how stimuli rendered invisible by the attentional blink (T2) impact on subsequent perceptual selection. (Giesbrecht, Bischof, & Kingstone, 2004; Qian, Meng, Chen, & Zhou, 2012). At any rate, given the rapid serial visual presentation associated with this task, it would be difficult to ascertain whether changes in perceptual selection of stimuli presented after the T2 were actually induced by the (invisible) T2 itself or by the (visible) T1.
CFS with visual probe (e.g. Hedger, Adams, & Garner, 2015)	To our knowledge, only two studies have investigated attentional cuing effects in response to threatening stimuli rendered invisible by CFS (Hedger, Adams, & Garner, 2015; Tan, Ma, Gao, Wu, & Fang, 2011).
Load induced blindness (e.g. Macdonald & Lavie, 2008)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by LIB.
Surprise induced blindness (Asplund, Todd, Snyder, Gilbert & Marois, 2010)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by SIB.
Adaptation induced blindness (e.g. Motoyoshi & Hayakawa, 2010)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by AIB.
Change blindness (CB: e.g. Simons & Rensinck, 2005)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by CB.
Distractor induced blindness (DIB: e.g. Michael, Hesselmann, Kiefer, & Niedeggen, 2011)	To our knowledge, no study has investigated changes in perceptual selection in response to threat stimuli rendered invisible by DIB.

Appendix B: Chapter 2. Search Terms and Omissions.

MVP paradigm.

PUBMED search terms:

((((((((((((((((sub-threshold[Title/Abstract]) AND dot-probe[Title/Abstract])) OR ((subliminal*[Title/Abstract]) AND attentional bias[Title/Abstract])) OR ((awareness[Title/Abstract]) AND attentional orientation[Title/Abstract])) OR ((masking[Title/Abstract]) AND dot-probe[Title/Abstract])) OR ((attention-orienting[Title/Abstract]) AND masked[Title/Abstract])) OR ((probe detection[Title/Abstract]) AND masked[Title/Abstract])) OR ((pre-attentive[Title/Abstract]) AND visual probe[Title/Abstract])) OR ((probe detection[Title/Abstract]) AND preconscious*[Title/Abstract])) OR ((subliminal[Title/Abstract]) AND bias[Title/Abstract])) OR ((masked[Title/Abstract]) AND cueing[Title/Abstract])) OR ((subliminal[Title/Abstract]) AND attentional bias*[Title/Abstract])) OR ((preconscious[Title/Abstract]) AND attention[Title/Abstract])) OR ((unaware[Title/Abstract]) AND attentional bias[Title/Abstract])) OR ((nonconscious[Title/Abstract]) AND dot probe[Title/Abstract])) OR ((masked[Title/Abstract]) AND dot-probe[Title/Abstract])) OR ((masked[Title/Abstract]) AND dot probe[Title/Abstract])) OR ((visual probe[Title/Abstract]) AND masked[Title/Abstract]))

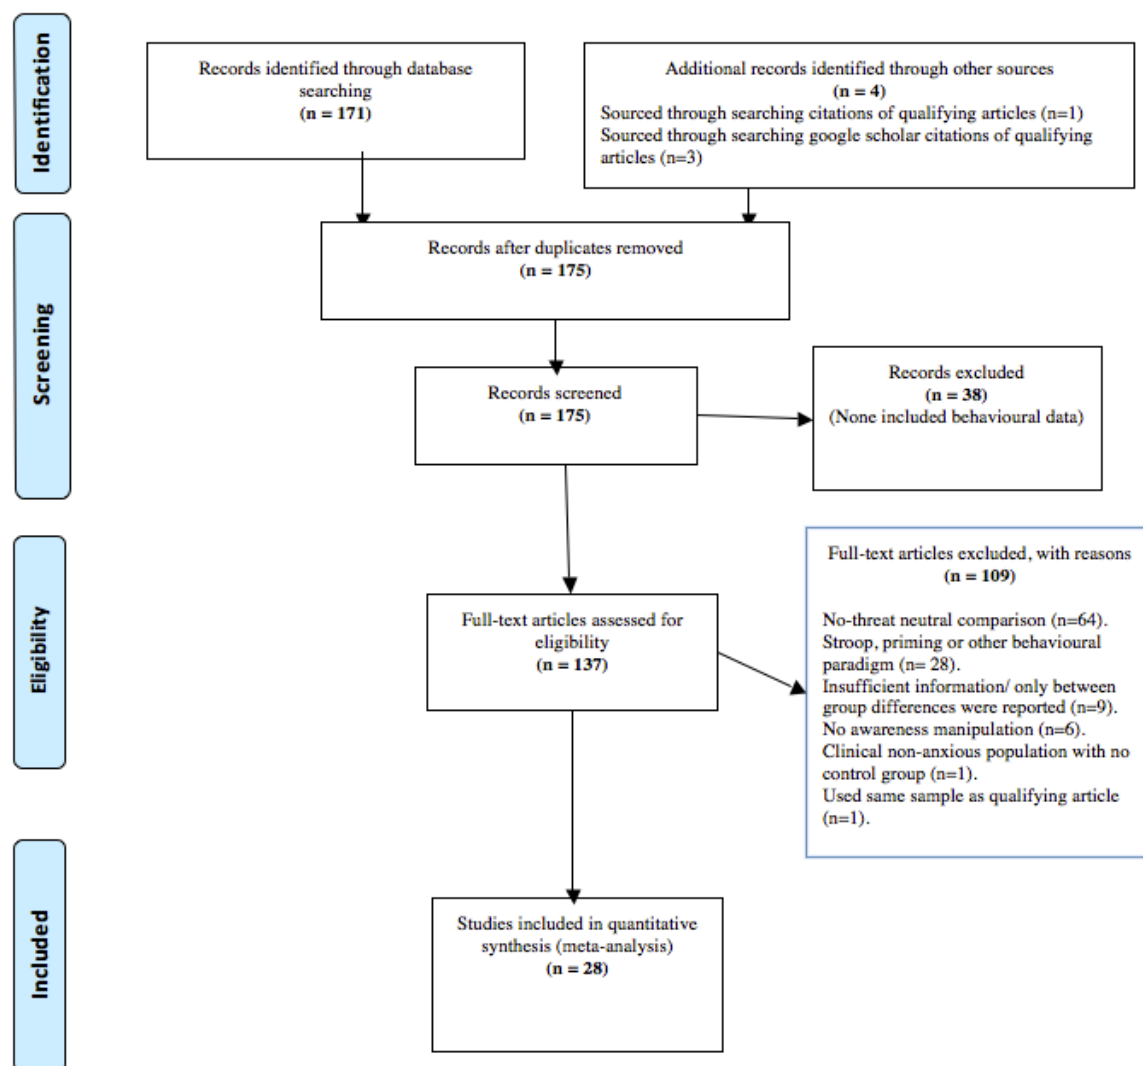


Figure A1. PRISMA flow diagram for the MVP search.

Note that this shows only the most important reasons for excluding a study (several studies were excluded for multiple reasons).

BR paradigm

PUBMED search terms:

(((((binocular rivalry[Title/Abstract]) AND emotion*[Title/Abstract])) OR ((negative*[Title/Abstract]) AND binocular rivalry[Title/Abstract])) OR ((threat*[Title/Abstract]) AND binocular rivalry[Title/Abstract])) OR ((emotion*[Title/Abstract]) AND interocular suppression[Title/Abstract])) OR ((negative*[Title/Abstract]) AND interocular suppression*[Title/Abstract])) OR ((phobic*[Title/Abstract]) AND interocular suppression[Title/Abstract]))

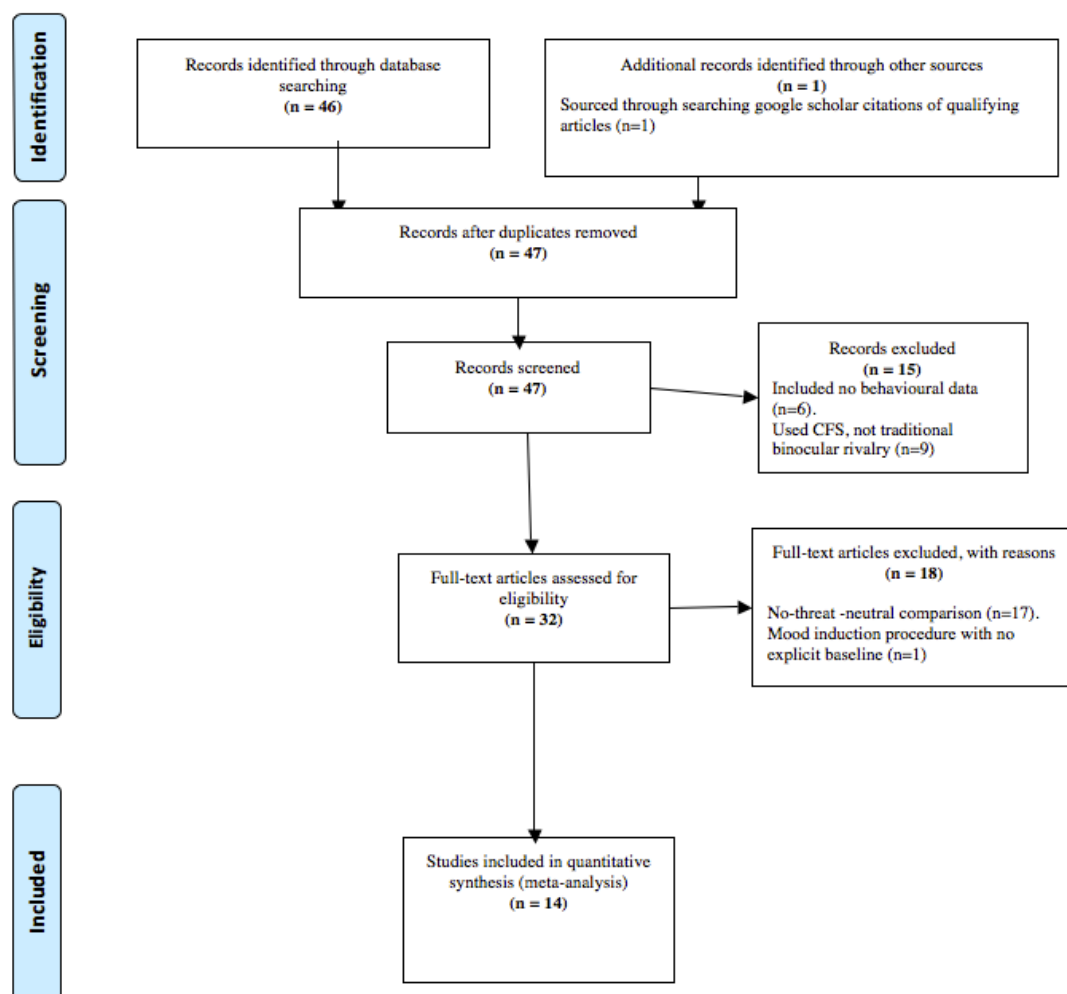


Figure A2. PRISMA flow diagram for the BR search.

bCFS paradigm***PUBMED search terms:***

((((((((((continuous flash suppression[Title/Abstract]) AND emotion*[Title/Abstract])) OR ((negative*[Title/Abstract]) AND continuous flash suppression[Title/Abstract])) OR ((threat*[Title/Abstract]) AND continuous flash suppression[Title/Abstract])) OR ((emotion*[Title/Abstract]) AND interocular suppression[Title/Abstract])) OR ((negative*[Title/Abstract]) AND interocular suppression*[Title/Abstract])) OR ((phobic*[Title/Abstract]) AND interocular suppression[Title/Abstract])) OR ((dynamic noise*[Title/Abstract]) AND suppression[Title/Abstract]))OR ((dynamic mask*[Title/Abstract]) AND suppression[Title/Abstract]))))

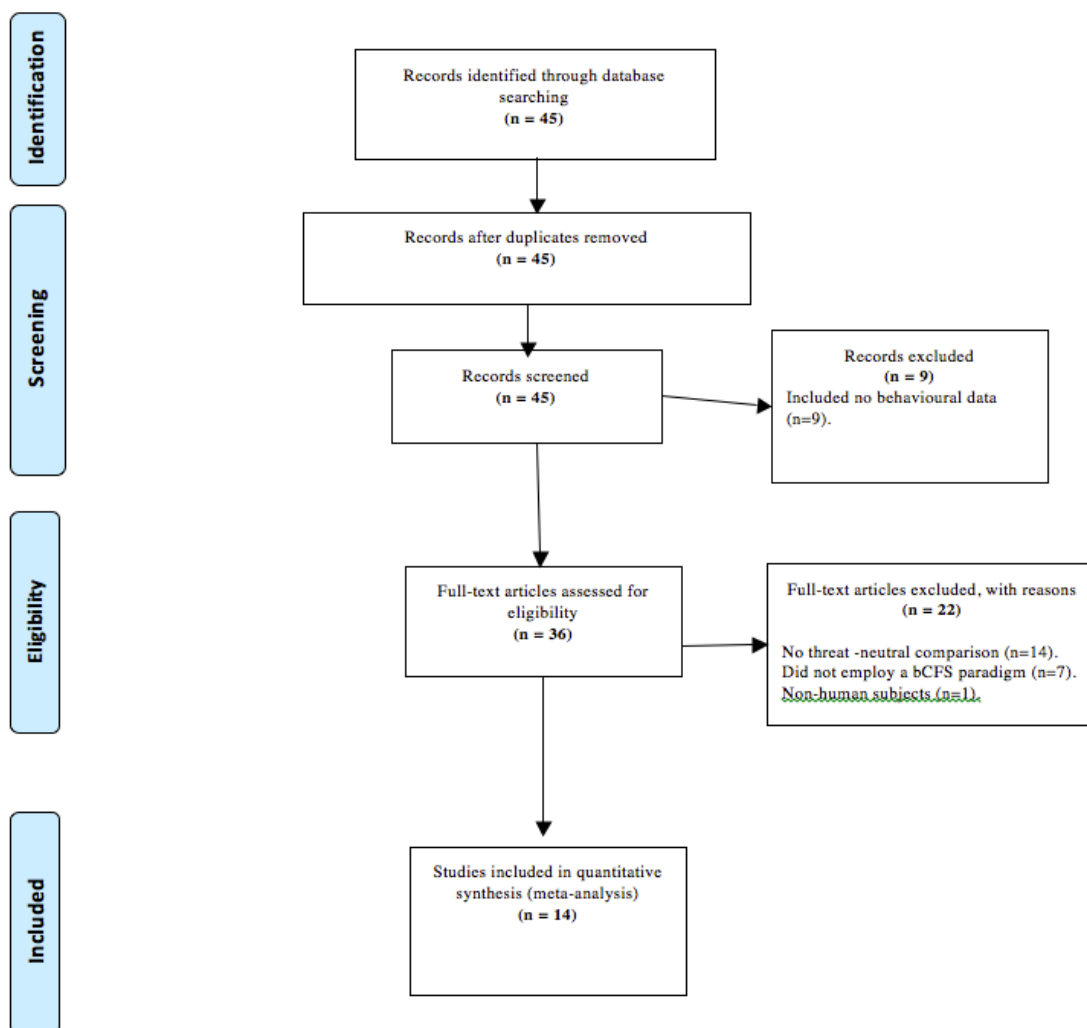


Figure A3. PRISMA flow diagram for the bCFS search.

Appendix C: Chapter 2. Decisions Regarding Standardisers for d

As Dunlap and Cortina (1996) note, d_z calculated from a repeated measures design may lead to inflated estimates (relative to an equivalent independent samples effect) as it does not correct for the pre-post correlation (r). In paired designs, if r is greater than .5, then effect sizes based on d_z will be inflated by a factor of $\sqrt{2(1-r)}$. We therefore also report the correlation-adjusted effect size (d_{RM}):

$$d_{RM} = \frac{t}{\sqrt{N}} \times \sqrt{2(1-r)}$$

There are several reasons why it is desirable to report d_{RM} as well as d_z . Firstly we note that i) this issue may have been ignored in other meta-analyses; the methods used to compute d for the paired design are highly inconsistent (Lakens, 2013), ii) independent designs are often the ‘default’ in meta-analyses and so most existing ‘benchmark’ effect sizes are based on independent samples iii) the paired correlation gives useful information about the precision of the experimental design (Cumming, 2012). Thus, wherever possible, we report the available estimate of r for each paradigm.

As the paired correlation was never reported, we used the equations reported in Morris & Deshon (2002) to recover these statistics from the available means, standard deviations and paired t value. Further information on how to interpret d_z and d_{RM} can be found in Lakens (2013).

The MVP paradigm

The paired correlation could be estimated for 13 effects ($M = .92$, $SD = .05$). The high correlation entails a highly sensitive design and precise estimation of the effect size. Thus the correlation-adjusted pooled effect size for the MVP paradigm is 0.11.

The BR paradigm

We do not report the estimated correlation-adjusted effect size for the BR paradigm, as this is distorted by the nature of the task. For instance, in a trial where a threatening and neutral stimulus are engaged in rivalry, perceiving one of the rivalling stimuli for more time will generally entail seeing the other for less time. A negative correlation therefore reflects the nature of the task, rather than a lack of sensitivity in the paired design. In this context, we do not believe a conversion from d_z to d_{RM} makes conceptual sense.

The bCFS paradigm

The paired correlation could be estimated for 19 effects ($M = .94$, $SD = .03$). The high correlation entails a highly sensitive design. Thus the correlation-adjusted pooled effect size for the MVP paradigm is -0.05 .

Appendix D: Chapter 2. Decisions Regarding Missing Information.

In cases where no relevant statistics were reported, we were sometimes able to calculate d_z by estimating the mean and standard deviation of difference scores from published figures, using GraphClick software (Version 3.0; Arizona Software). If a paper contained no information to compute an effect size, we contacted the corresponding author to obtain the necessary data. If this was unsuccessful, we adopted two conservative approaches to estimate the effect, as summary effects calculated from meta-analyses are often modestly inflated (Button et al., 2013): i) Where an effect was reported to be non-significant, but no exact statistics were reported, we estimated the effect size assuming that $p = .50$ to obtain a representative sample of outcomes (Cooper & Hedges, 1994). ii) In cases where no exact p value was reported, but the paper reported that an effect was significant (e.g. “ $p < .050$ ”), we assumed the upper bound (e.g. $p = .050$) to provide a conservative estimate of the effect size. These two procedures were only used when it was possible to determine the direction of an effect (from condition means/figures) and in practice accounted for just 21% of the included effects. If insufficient information was available to use the above methods to estimate an effect size, the effect was excluded from the analysis.

Appendix E. Chapter 2. Methods for Imputing Missing Moderator Values.

To estimate unreported values of continuous moderators, we used two multiple imputation methods (Van Buuren & Groothuis-Oudshoorn, 2011). With “*regression imputation*” we performed regression on a bootstrapped sample of observed data. From the resulting regression coefficients, we imputed values for the unreported data via predictive mean matching (van Buuren, 2012). These imputed values were then combined with the observed data and a meta regression was performed on the full data set. We then repeated this process 1000 times and pooled across analyses. For “*random-sample imputation*” we simply imputed unreported data from random samples of the observed data 1000 times, combined this with the observed data and performed the analyses on all full, imputed

datasets and pooled across analyses. Thus, the regression imputation is an ‘optimistic’ estimate, based on the assumption that the observed relationship is predictive of the missing values (taking into account the uncertainty in the data), whereas random sample imputation is conservative as it assumes that missing values are randomly distributed.

Appendix F. Chapter 2. Explaining Sources of Heterogeneity.

The pseudo R^2 statistic (López-López, Marín-Martínez, Sánchez-Meca, Van den Noortgate, & Viechtbauer, 2014) is computed as:

$$R^2 = \frac{\tau_{RE}^2 - \tau_{ME}^2}{\tau_{RE}^2}$$

Where τ_{RE} is the total heterogeneity as estimated by the empty (no moderator) random effects model and τ_{ME} is the amount of residual heterogeneity as estimated by the mixed effects (moderator) model. The pseudo R^2 thus estimates the proportional reduction in heterogeneity after including moderators. Note that it does not involve sampling variability at all. Hence, it is possible to get very large R^2 values, even when there are still discrepancies between the regression line and the observed effect sizes (when those discrepancies are not much larger than what one would expect based on sampling variability alone). In fact, when $\tau_{ME}^2 = 0$, then $R^2 = 1$. However, this *does not* imply that the points all fall perfectly on the regression line - the residuals are just not larger than expected based on sampling variability. This statistic should thus be interpreted with caution, particularly for analyses with a small k .

Appendix G. Chapter 2. Summary of included effects: The MVP paradigm.

Demographic information.

The mean age of participants included in the MVP effects was 21.89 (range: 12 – 40). The gender ratio (females: males) was 1.45 (range: 0.55- 1.45). 24 effects were defined as being from undergraduate populations, 14 were defined as “consenting adults” (usually a mixture of university students and staff) 4 were child populations (all less than 14 years of age) and 2 had a clinical anxiety disorder. None of these demographic variables (age, gender ratio, population) were significant moderators of effect size.

Table A2. Summary of Included Effects in the MVP Paradigm.

Study/Effect	Method	Source	Notes
1) Mogg et al., 1994 (i)	t & N	p 856	Data is taken from the “no-stress condition” (see ‘other coding and inclusion decisions no 3 in the main text’).
Mogg et al., 1994 (ii)	t & N	p 856	As above
2) Mogg et al., 1995 (i)	t & N	p 26	
Mogg et al., 1995 (ii)	t & N	p 26	
3) Mogg & Bradley, 1999 (i)	t & N	p 722	
Mogg & Bradley, 1999 (ii)	$p=.5$	p 722	p value of .5 is assumed (see Appendix D)
Mogg & Bradley, 1999 (iii)	t & N	p 729	
Mogg & Bradley, 1999 (iv)	t & N	p 731	Experiment 2 was excluded, since it contains no explicit comparison between threat and neutral stimuli (happy and threatening faces were used) - see ‘inclusion criteria’ no 3 in the main text. We did not split into high and low anxiety groups here, since an effect size cannot be computed for each sample separately from the reported information.
4) Mogg & Bradley, 2002 (i)	t & N	p 1408	We split the data by social anxiety, rather than trait anxiety, since the analyses are more detailed for these groups.
Mogg & Bradley, 2002 (ii)	t & N	p 1408	
5) Fox, 2002 (i)	t & N	p 57	
Fox, 2002 (ii)	t & N	p 57	
Fox, 2002 (iii)	$p=.5$	p 57	p value of .5 is assumed, since the effect is indicated to be non-significant but no statistics were reported (see Appendix D).
Fox, 2002 (iv)	$p=.5$	p 57	As above
6) Keogh et al 2003 (i)	$Mdiff$ and $SDdiff$	p 88	Mean differences and standard deviation of differences were computed from figures on page 88- using GraphClick software. The effect represents the pooled effect across all groups and word types.
7) Beaver et al., 2005 (i)	t & N	p 74	An effect size can be computed for the “high aversive” group only. This group could be considered the group for which the conditioning procedure was most effective. Only data from experiment 2 is included, since in experiment 1, the stimuli are not masked during the MVP trials.
8) Hunt et al., 2006 (i)	$Mdiff$ and $SDdiff$	p 423	Mean differences and standard deviation of differences were computed from figures on page 423- using GraphClick software. The effect represents the pooled effect across all groups and word types.
9) Koster et al., 2007 (i)	$p=.5$	p 288	Interaction between cue validity and cue valence is non significant, so p value of 0.5 assumed (see Appendix D - direction of effect is inferred from table).
Koster et al., 2007 (ii)	$p=.5$	p 290	As above

Koster et al., 2007 (iii)	$p=.5$	p 291	As above. Experiment 1b does not use masking and so data were excluded.
10) Murphy et al., 2007 (i)	$p=.5$	p 508	Placebo group only. No exact p value reported so .5 assumed.
11) Stone & Valentine, 2007 (i)	$t \& N$	Provided by author	
Stone & Valentine, 2007 (ii)	$t \& N$	Provided by author	
12) Wirth & Schultheiss, 2007 (i)	$t \& N$	Provided by author	
13) Schultheiss & Hale, 2007 (i)	$t \& N$	Provided by author	
Schultheiss & Hale, 2007 (ii)	$t \& N$	Provided by author	
14) Carlson & Reinke, 2008 (i)	$t \& N$	p 524	Although many components of attention were analysed (orienting, disengagement), our effect reflects the RT difference between congruent and incongruent trials, in line with the comparison most commonly made in the other MVP studies.
Carlson & Reinke, 2008 (ii)	$t \& N$	p 526	As above
15) Monk et al., 2008 (i)	$t \& N$	Provided by author	
Monk et al., 2008 (ii)	$t \& N$	Provided by author	
16) Carlson et al., 2009a (i)	$p \& N$	p 1387	
Carlson et al., 2009a (ii)	$p \& N$	p 1387	
17) Carlson et al., 2009b (i)	$F \& N$	p 538	
18) Helzer et al., 2009 (i)	$p=.5$	p. 6	Effect reported to be non-significant with no exact stats so $p=.5$ assumed (see Appendix D).
19) Fox et al., 2010 (i)	$p=.5$	p 5	Effect reported to be non-significant with no exact stats so $p=.5$ assumed (see Appendix D). Direction of effect is inferred from table.
20) Carlson & Reinke, 2010 (i)	$F \& N$	p 22	
21) Thomason et al., 2010 (i)	$p=.5$	p. 6	
Thomason et al., 2010 (ii)	$t \& N$	p.6	Angry faces were included, but no stats are available for this comparison.
22) Sutton & Altarriba, 2011 (i)	$t \& N$	p 743	Experiment 1 was excluded (it did not use a masking procedure).
23) Carlson et al., 2012 (i)	$p \& N$	p 205	
Carlson et al., 2012 (ii)	$p \& N$	p 205	
24) Carlson et al., 2013a (i)	$t \& N$	p 4	
25) Carlson et al., 2013b (i)	$t \& N$	p 2597	
26) Maoz et al., 2013 (i)	$t \& N$	Provided by author	
Maoz et al., 2013 (ii)	$t \& N$	Provided by author	
27) McCroy et al., 2013 (i)	$t \& N$	p 5	Both groups analysed together
28) Carlson et al., 2014 (i)	$p \& N$	p 5	

Appendix H. Chapter 2. Main Effects and Interactions: The MVP paradigm.

Table A3. Moderators of Threat-related Bias in the MVP Paradigm: Main Effects.

Moderator	k	df	Q^O	p	R^2
Stimulus type	44	5	21.12	.017*	24.34
SOA	44	1	9.23	.002**	29.10
Awareness measure	44	1	0.04	.835	0.00
Probe response	44	1	0.14	.708	0.00
Retinal size	22	1	0.24	.624	0.00
STAI-T	15	1	6.90	.008**	37.18
H ^{per}	26	1	0.03	.858	0.00
Visual field	10	1	1.93	.165	11.99

^O Omnibus test for comparison between levels of a moderator

Plots of all main effects can be found at the following URLs:

http://figshare.com/articles/MVP_Main_effects_1/1466750

http://figshare.com/articles/MVP_Main_effects_2/1466751

Table A4. Moderators of Threat-related Bias in the MVP Paradigm: Two-way Interactions.

Moderators	<i>k</i>	Model matrix full?	<i>df</i>	<i>Q^I</i>	<i>p</i>	<i>R²</i>
Stimulus type x SOA	44	No	3	1.44	.697	24.32
Stimulus x awareness measure	44	No	2	1.87	.391	19.88
Stimulus x probe response	44	No	3	3.23	.358	25.63
Stimulus x retinal size	22	No	2	1.30	.521	6.35
Stimulus x STAI-T	15	No	2	8.78	.012*	71.01
Stimulus x H ^{per}	26	No	2	0.79	.673	0.00
SOA x awareness measure	44	Yes	1	3.73	.054 ^M	30.86
SOA x probe response	44	Yes	1	0.04	.834	20.84
SOA x retinal size	22	Yes	1	1.54	.215	0.00
SOA x H ^{per}	26	Yes	1	0.21	.646	0.00
Awareness measure x probe response	44	Yes	1	1.10	.295	0.00
Awareness measure x retinal size	22	Yes	1	1.44	.229	0.00
Probe response x retinal size	22	Yes	1	1.29	.254	0.00
Probe response x STAI-T	15	Yes	1	4.69	.030*	50.61
Retinal size x H ^{per}	16	Yes	1	1.44	.230	0.00
STAI-T x retinal size	11	Yes	1	1.56	.211	65.49
STAI-T x H ^{per}	14	Yes	1	2.43	.112	50.99

Note: If the model matrix is not full, this indicates that redundant predictors were removed (e.g. there were no STAI-T data for effects that used disgust faces- thus this coefficient was removed from the model).

^I test of the interaction coefficient

^M non-significant trend

Plots of all interactions can be found at the following URLs:

http://figshare.com/articles/MVP_Interactions_1/1466752

Appendix I. Chapter 2. Summary of Included Effects: the BR Paradigm.

Demographic information.

The mean age of participants included in the BR effects was 27.67 (range: 20 – 71). The gender ratio (females: males) was 2.64 (range: 1- 9.5). 24 effects were defined as being from undergraduate populations, 14 were defined as “consenting adults” (usually a mixture of university students and staff) 1 was an elderly population (mean age 71) and 5 had a clinical anxiety disorder. None of these demographic variables were significant moderators of effect size.

Table A5. Summary of Included Effects in the BR paradigm

Study/Effect	Method	Source	Notes
1) Alpers et al., 2005 (i)	$t\&N$	p.29	The difference between the predominance ratio for CS+ and the CS- at baseline (before conditioning) versus the same predominance ratio after block 3 (after conditioning). Calculated from t value (Experiment 1: p 29). No t statistics/ degrees of freedom are reported for the initial percept.
Alpers et al., 2005 (ii)	$t\&N$	p.30	The difference between the predominance ratio for CS+ and the CS- at baseline (before conditioning) versus the same predominance ratio after block 4 (after conditioning). Calculated from t value. No statistics are reported for the initial percept In both of the experiments in this study we think a comparison of the final block to baseline is optimal, since it compares the most threatening (i.e. most conditioned) stimulus to a baseline.
2) Alpers & Pauli, 2006	$t\&N$	p.603	The initial percept data does not contain an explicit comparison between threat and neutral- the data are collapsed across all emotional stimuli for this analysis- thus this can not be included (see 'inclusion criteria' no 3 in the main text).
3) Alpers & Gerdes, 2007 (i)	$t\&N$	p.499	
Alpers & Gerdes, 2007 (ii)	$t\&N$	p.499	
Alpers & Gerdes, 2007 (iii)	$t\&N$	p.500	The initial percept data (for all experiments: 1, 2 and 3) does not contain an explicit comparison between threat and neutral- the data are collapsed across all emotional stimuli for this analysis. Experiment 3 also collapses all analyses across all emotional stimulus types. Thus none of this data can be included (see 'inclusion criteria' no 3 in the main text).
4) Bannerman et al., 2008 (i)	$t\&N$	p 320	
Bannerman et al., 2008 (ii)	$p=.001$	p 324	No exact p value is reported so upper bound is assumed.
Bannerman et al., 2008 (iii)	$p=.5$	p.324	Calculated by assuming a p value of 0.5 to achieve a representative sample of outcomes (see Appendix D - direction of effect inferred from figure on p 323).
5) Yoon et al., 2009 (i)	$t\&N$	Provided by author	
Yoon et al., 2009 (ii)	$t\&N$	Provided by author	
Yoon et al., 2009 (iii)	$t\&N$	Provided by author	
6) Gray et al., 2009 (i)	$t\&N$	Provided by author	
Gray et al., 2009 (ii)	$t\&N$	Provided by author	
Gray et al., 2009 (iii)	$t\&N$	Provided by author	
Gray et al., 2009 (iv)	$t\&N$	Provided by author	
7) Arnting et al., 2010 (i)	$t\&N$	p 10041	
Arnting et al., 2010 (ii)	$p=.5$	p 10041	Figure on p 10041 indicates non-significance. Thus effect size calculated by assuming a p value of 0.5 to achieve a representative sample of outcomes (see Appendix D . Direction of effect inferred from figure on same page).
8) Anderson et al., 2011 (i)	$t\&N$	p. 1447	
Anderson et al., 2011 (ii)	$t\&N$	p. 1448	
9) Bannerman et al., 2011 (i)	$t\&N$	p.375	
Bannerman et al., 2011 (ii)	$p=.6$	p.375	Calculated by assuming the minimum p value indicated in text. Direction of effect inferred from figure on same page.
Bannerman et al., 2011 (iii)	$t\&N$	p.375	
Bannerman et al., 2011 (iv)	$p=.7$	p.375	Calculated by assuming the minimum p value indicated in text. Direction of effect inferred from figure on same page.
10) Ritchie et al., 2012	$t\&N$	Provided by	

Chapter 7

(i)	Ritchie et al., 2012	$t \& N$	author	
(ii)			Provided by author	For both experiments, we took the data from the foveal condition only, to ensure that the stimulus conditions were most similar to the other studies in the analysis (which generally present stimuli foveally).
11)	Lerner et al., 2012	$p = .05$		No exact p value is reported so upper bound is assumed (see Appendix D).
(i)				
12)	Singer et al., 2012	$t \& N$	Provided by author	
(i)				
	Singer et al., 2012	$t \& N$	Provided by author	
(ii)				
	Singer et al., 2012	$t \& N$	Provided by author	
(iii)				
	Singer et al., 2012	$t \& N$	Provided by author	
(iv)				
	Singer et al., 2012	$t \& N$	Provided by author	
(v)				
	Singer et al., 2012	$t \& N$	Provided by author	
(vi)				
13)	Anderson et al.,	$p = .5$	p.623	
2013 (i)				
	Anderson et al.,	$p = .5$	p.624	“Scowling” is an expression of anger, so the stimuli used in this study are coded as angry. The focus of this study is differences in threat-bias between groups and no within groups threat-biases are included. However, we collapsed across groups and refer to the “main effect of face type” being non-significant and assume a p value of .5 (see Appendix D). Directions of effects are inferred from the table on page 622).
2013 (ii)				
14)	Gerdes & Alpers.,	$p = .5$	p. 19	
2014 (i)				
	Gerdes & Alpers., 2014	$t \& N$	p. 19	Both effects from this study are from the initial percept measure. Only between group differences are reported for continuous rivalry.
(ii)				

Appendix J. Chapter2. Main Effects and Interactions: The BR paradigm

Table A6. Moderators of Threat-related Bias in the BR Paradigm: Main Effects.

Moderator	k	df	Q^o	p	R^2
Stimulus type	31	5	13.24	.021*	29.92
Dominance measure	31	1	3.08	.079 ^M	6.86
Design	31	1	4.01	.045*	12.96
Retinal size	26	1	3.59	.058 ^M	12.81
Rivalry trial length	20	1	0.33	.568	0.00
STAI-T	17	1	0.19	.657	0.00

Plots of all main effects can be found at the following URL:

http://figshare.com/articles/BR_Main_effects_1/1466755

Table A7. Moderators of Threat-related Bias in the BR Paradigm: Two-way Interactions.

Moderators	<i>k</i>	Model matrix full?	<i>df</i>	<i>Q'</i>	<i>p</i>	<i>R</i> ²
Stimulus type x dominance measure	31	No	2	2.50	.286	57.60
Stimulus x design	31	No	2	0.47	.791	51.98
Stimulus type x retinal size	26	No	2	0.18	.915	0.00
Stimulus x rivalry trial length	20	No	2	0.22	.897	48.60
Stimulus x STAI-T	17	No	2	0.04	.981	0.00
Dominance measure x Design	31	Yes	1	1.62	.203	27.20
Dominance measure x retinal size	26	Yes	1	3.40	.065 st	31.07
Dominance measure x STAI-T	17	Yes	1	1.50	.221	7.29
Design x retinal size	26	Yes	1	2.65	.103	32.18
Design x rivalry trial length	20	Yes	1	0.12	.729	17.51
Design x STAI-T	17	Yes	1	4.25	.039 ^t	30.81
Retinal size x STAI-T	15	Yes	1	0.01	.972	0.00

Plots of all interactions can be found at the following URLs:

http://figshare.com/articles/BR_Interactions_1/1466756

http://figshare.com/articles/BR_Interactions_2/1466757

Appendix K. Chapter 2. Summary of Included Effects: the bCFS

Paradigm.

Demographic information.

The mean age of participants included in the bCFS effects was 21.94 (range: 9 – 39). The gender ratio (females: males) was 2.64 (range: 0.33- 88). 29 effects were defined as being “consenting adults” (usually a mixture of university students and staff) and 4 were defined as undergraduate students. 2 were child populations (mean age 9). None of these demographic variables were significant moderators of effect size.

Table A8. Summary of Included Effects for the bCFS paradigm.

Study/Effect	Method	Source	Notes
1) Yang, et al., 2007 (i)	$t \& N$	p.884	
Yang, et al., 2007 (ii)	$t \& N$	p.884	
Yang, et al., 2007 (iii)	$t \& N$	p.884	
Yang, et al., 2007 (iv)	$t \& N$	p.884	We do not include data from experiment 3, since only eye stimuli are used (see 'other coding and inclusion decisions' no 6 in the main text).
2) Sterzer et al., 2011	$p = .01$	p 1620	No exact p value is reported so upper bound is assumed (direction of effect is inferred from figure on p 1620). Data is taken from the control group.
3) Sylvers et al., 2011 (i)	$t \& N$	p.1283	
Sylvers et al., 2011 (ii)	$t \& N$	p.1283	
4) Yang & Yeh., 2011 (i)	$F \& N$	p. 225	
Yang & Yeh., 2011 (ii)	$F \& N$	p. 225	
Yang & Yeh., 2011 (iii)	$F \& N$	p. 228	
Yang & Yeh., 2011 (iv)	$F \& N$	p. 228	
5) Chen & Yeh, 2012 (i)	$t \& N$	Provided by author	This effect represents the difference in response time between detecting fearful and neutral faces, collapsed across the direction of eye gaze (direct/ averted).
6) Stein & Sterzer, 2012	$t \& N$	Provided by author	The remaining experiments (2-4) have no explicit comparison between threatening and neutral stimuli and so are not included.
7) Stewart et al (i)	$t \& N$	p. 719	Here, and for the rest of the effects reported in this study, we calculate our effect based on the difference in response time between the "most" untrustworthy or dominant face and the neutral face (see 'other coding and inclusion decisions' no 4 in the main text).
Stewart et al., 2012 (ii)	$t \& N$	p. 719	
Stewart et al., 2012 (iii)	$t \& N$	p. 719	
Stewart et al., 2012 (iv)	$t \& N$	p. 719	
Stewart et al., 2012 (v)	$t \& N$	p. 721	
Stewart et al., 2012 (vi)	$t \& N$	p. 721	
8) Gray et al., 2013 (i)	$t \& N$	Provided by author	
Gray et al., 2013 (ii)	$t \& N$	Provided by author	
Gray et al., 2013 (iii)	$t \& N$	Provided by author	
Gray et al., 2013 (iv)	$t \& N$	Provided by author	
9) Stein et al., 2014 (i)	$t \& N$	Provided by author	Both effects for this study represent the difference in response time between fearful faces and neutral faces in a broadband spatial frequency (see 'other coding and inclusion decisions' no 6 in the main text).
Stein et al., 2014 (ii)	$t \& N$	Provided by author	
10) Capita et al., 2014 (i)	$p \& N$	p. 1030	Direction is inferred from figure
11) Oliver et al., 2014 (i)	$p = .001$	p. 7	Data is taken from the "subjective awareness" measure- since this is most analogous to that from other studies.
Oliver et al., 2014 (ii)	$p \& N$	p. 7	Direction is inferred from figure.
Oliver et al., 2014 (iii)	$p = .001$	p. 13	
Oliver et al., 2014 (iv)	$p \& N$	p. 13	Direction is inferred from figure.
12) Getov et al., 2014 (i)	$t \& N$	p.4	As with Stewart et al., (2012), here, and for the rest of the effects reported in this study, we calculate our effect based on the difference in response time between the "most" untrustworthy or dominant face and the neutral face (see 'other coding and inclusion decisions' no 4 in the main text).
Getov et al., 2014 (ii)	$t \& N$	p.4	
13) Jusyte et al., 2015 (i)	$M \text{ and } SD_{diff}$	p.290	For all the effects in this study, we used GraphClick software to calculate the mean differences and standard deviation of the difference scores to compute d. This information is displayed in figure 2 in this paper.
Jusyte et al., 2015 (ii)	$M \text{ and } SD_{diff}$	p.290	
Jusyte et al., 2015 (iii)	$M \text{ and } SD_{diff}$	p.290	
14) Hedger et al., 2015 (i)	$t \& N$	Provided by author	Provided by author

Appendix I: Main effects and Interactions: The bCFS Paradigm

Table A9. Moderators of Threat-related Bias in the bCFS Paradigm: Main Effects.

Moderator	<i>k</i>	<i>df</i>	<i>Q^p</i>	<i>p</i>	<i>R²</i>
Stimulus type	27	6	41.32	<.001***	65.38
Awareness measure	27	1	0.19	.661	0.00
Retinal size	22	1	0.77	.379	0.00

Plots of all main effects can be found at the following URL:

http://figshare.com/articles/bCFS_Main_effects_1/1466762

Table A10. Moderators of Threat-related Bias in the bCFS Paradigm: Two Way Interactions.

Moderators	<i>k</i>	Model matrix full?	<i>df</i>	<i>Qⁱ</i>	<i>p</i>	<i>R²</i>
Stimulus type x awareness measure	27	No	1	0.24	.625	59.97
Stimulus type x retinal size	22	No	2	3.90	.142	61.19

Note: Interaction between Awareness measure and retinal size could not be calculated due to empty cells.

Plots of all interactions can be found at the following URL:

http://figshare.com/articles/bCFS_Interactions/1466763

References

- Adams, R. B., Franklin, R. G., Kveraga, K., Ambady, N., Kleck, R. E., Whalen, P. J., ... Nelson, A. J. (2012). Amygdala responses to averted vs direct gaze fear vary as a function of presentation speed. *Social Cognitive and Affective Neuroscience*, 7(5), 568–577. <http://doi.org/10.1093/scan/nsr038>
- Adams, W. J., Gray, K. L. H., Garner, M., & Graf, E. W. (2011). On the ‘special’ status of emotional faces... Comment on Yang, Hong, and Blake (2010). *Journal of Vision*, 11(3), 10. <http://doi.org/10.1167/11.3.10>
- Adolphs, R. (2013). The Biology of Fear. *Current Biology*, 23(2), R79–R93. <http://doi.org/10.1016/j.cub.2012.11.055>
- Akechi, H., Stein, T., Kikuchi, Y., Tojo, Y., Osanai, H., & Hasegawa, T. (2015). Preferential awareness of protofacial stimuli in autism. *Cognition*, 143, 129–134. <http://doi.org/10.1016/j.cognition.2015.06.016>
- Alais, D., & Blake, R. (2005). *Binocular rivalry*. MIT press.
- Alais, D., Cass, J., O’Shea, R. P., & Blake, R. (2010). Visual Sensitivity Underlying Changes in Visual Consciousness. *Current Biology*, 20(15), 1362–1367. <http://doi.org/10.1016/j.cub.2010.06.015>
- Algom, D., Chajut, E., & Lev, S. (2004). A rational look at the emotional stroop phenomenon: a generic slowdown, not a stroop effect. *Journal of Experimental Psychology. General*, 133(3), 323–338. <http://doi.org/10.1037/0096-3445.133.3.323>
- Almeida, J., Mahon, B. Z., Nakayama, K., & Caramazza, A. (2008). Unconscious processing dissociates along categorical lines. *Proceedings of the National Academy of Sciences of the United States of America*, 105(39), 15214–15218. <http://doi.org/10.1073/pnas.0805867105>
- Alpers, G., & Pauli, P. (2006). Emotional pictures predominate in binocular rivalry. *Cognition and Emotion*, 20(5), 596–607. <http://doi.org/10.1080/02699930500282249>
- Alpers, G. W., & Gerdes, A. B. M. (2007). Here is looking at you: emotional faces predominate in binocular rivalry. *Emotion (Washington, D.C.)*, 7(3), 495–506. <http://doi.org/10.1037/1528-3542.7.3.495>
- Alpers, G. W., Ruhleder, M., Walz, N., Mühlberger, A., & Pauli, P. (2005). Binocular rivalry between emotional and neutral stimuli: A validation using fear conditioning and EEG. *International Journal of Psychophysiology*, 57(1), 25–32. <http://doi.org/10.1016/j.ijpsycho.2005.01.008>
- Amaral, D. G., Behniea, H., & Kelly, J. L. (2003). Topographic organization of projections from the amygdala to the visual cortex in the macaque monkey. *Neuroscience*, 118(4), 1099–1120.
- Amting, J. M., Greening, S. G., & Mitchell, D. G. V. (2010). Multiple mechanisms of consciousness: the neural correlates of emotional awareness. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 30(30), 10039–10047. <http://doi.org/10.1523/JNEUROSCI.6434-09.2010>
- Anderson, E. C., Dryman, M. T., Worthington, J., Hoge, E. A., Fischer, L. E., Pollack, M. H., ... Simon, N. M. (2013). Smiles may go unseen in generalized social anxiety disorder: evidence from binocular

- rivalry for reduced visual consciousness of positive facial expressions. *Journal of Anxiety Disorders*, 27(7), 619–626. <http://doi.org/10.1016/j.janxdis.2013.07.004>
- Anderson, E., Siegel, E. H., Bliss-Moreau, E., & Barrett, L. F. (2011). The visual impact of gossip. *Science (New York, N.Y.)*, 332(6036), 1446–1448. <http://doi.org/10.1126/science.1201574>
- Argyle, M. (2013). *Bodily Communication*. London: Routledge.
- Arnold, D. H. (2011). Why is Binocular Rivalry Uncommon? Discrepant Monocular Images in the Real World. *Frontiers in Human Neuroscience*, 5. <http://doi.org/10.3389/fnhum.2011.00116>
- Ashwin, C., Wheelwright, S., & Baron-Cohen, S. (2006). Finding a face in the crowd: Testing the anger superiority effect in Asperger Syndrome. *Brain and Cognition*, 61(1), 78-95.
- Asplund, C. L., Todd, J. J., Snyder, A. P., Gilbert, C. M., & Marois, R. (2010). Surprise-induced blindness: a stimulus-driven attentional limit to conscious perception. *Journal of Experimental Psychology. Human Perception and Performance*, 36(6), 1372–1381. <http://doi.org/10.1037/a0020551>
- Axelrod, V., Bar, M., & Rees, G. (2015). Exploring the unconscious using faces. *Trends in Cognitive Sciences*, 19(1), 35–45. <http://doi.org/10.1016/j.tics.2014.11.003>
- Bahrami, B., Carmel, D., Walsh, V., Rees, G., & Lavie, N. (2008). Spatial attention can modulate unconscious orientation processing. *Perception*, 37(10), 1520–1528.
- Baker, D. H. (2010). Visual Consciousness: The Binocular Rivalry Explosion. *Current Biology*, 20(15), R644–R646. <http://doi.org/10.1016/j.cub.2010.06.010>
- Baker, D. H., & Graf, E. W. (n.d.). Natural images dominate in binocular rivalry. Retrieved 14 September 2015, from <http://www.pnas.org>
- Bannerman, R. L., Milders, M., De Gelder, B., & Sahraie, A. (2008). Influence of emotional facial expressions on binocular rivalry. *Ophthalmic & Physiological Optics: The Journal of the British College of Ophthalmic Opticians (Optometrists)*, 28(4), 317–326. <http://doi.org/10.1111/j.1475-1313.2008.00568.x>
- Bannerman, R. L., Regener, P., & Sahraie, A. (2011). Binocular rivalry: A window into emotional processing in aging. *Psychology and Aging*, 26(2), 372–380. <http://doi.org/10.1037/a0022029>
- Bar, M., & Neta, M. (2007). Visual elements of subjective preference modulate amygdala activation. *Neuropsychologia*, 45(10), 2191–2200. <http://doi.org/10.1016/j.neuropsychologia.2007.03.008>
- Bargh, J. A., & Morsella, E. (2008). The Unconscious Mind. *Perspectives on Psychological Science : A Journal of the Association for Psychological Science*, 3(1), 73–79.
- Barrett, L. F. (2006). Solving the emotion paradox: categorization and the experience of emotion. *Personality and Social Psychology Review: An Official Journal of the Society for Personality and Social Psychology, Inc*, 10(1), 20–46. http://doi.org/10.1207/s15327957pspr1001_2
- Barron, A. B., & Klein, C. (2016). What insects can tell us about the origins of consciousness. *Proceedings of the National Academy of Sciences*, 113(18), 4900–4908. <http://doi.org/10.1073/pnas.1520084113>
- Baumgartner, C., Pataia, E., Lindinger, G., & Deecke, L. (2000). Neuromagnetic recordings in temporal lobe epilepsy. *Journal of Clinical Neurophysiology: Official Publication of the American Electroencephalographic Society*, 17, 177–189. <http://dx.doi.org/10.1097/00004691-200003000-00007>

- Beard, C. (2011). Cognitive bias modification for anxiety: current evidence and future directions. *Expert Review of Neurotherapeutics*, 11(2), 299–311. <http://doi.org/10.1586/ern.10.194>
- Beaver, J. D., Mogg, K., & Bradley, B. P. (2005). Emotional conditioning to masked stimuli and modulation of visuospatial attention. *Emotion (Washington, D.C.)*, 5(1), 67–79. <http://doi.org/10.1037/1528-3542.5.1.67>
- Beck, A. T., Ward, C. H., Mendelson, M., Mock, J., & Erbaugh, J. (1961). An inventory for measuring depression. *Archives of General Psychiatry*, 4, 561–571. <http://dx.doi.org/10.1001/archpsyc.1961.01710120031004>
- Becker, B. J. (2005). Fail-safe N or file-drawer number. In H. R. Rothstein, A. J. Sutton, & M. Borenstein (Eds.), *Publication bias in meta-analysis: Prevention, assessment and adjustments* (pp. 111–125). Chichester, England: Wiley.
- Becker, D. V., Anderson, U. S., Mortensen, C. R., Neufeld, S. L., & Neel, R. (2011). The face in the crowd effect unconfounded: happy faces, not angry faces, are more efficiently detected in single- and multiple-target visual search tasks. *Journal of Experimental Psychology. General*, 140(4), 637–659. <http://doi.org/10.1037/a0024060>
- Begg, C. B., & Berlin, J. A. (1988). Publication bias: A problem in interpreting medical data. *Journal of the Royal Statistical Society Series A (General)*, 151, 419–463. <http://dx.doi.org/10.2307/2982993>
- Beaupré, M. G., & Hess, U. (2005). Cross-cultural emotion recognition among Canadian ethnic groups. *Journal of Cross-Cultural Psychology*, 36, 355–370.
- Bishop, S. J. (2007). Neurocognitive mechanisms of anxiety: an integrative account. *Trends in Cognitive Sciences*, 11(7), 307–316. <http://doi.org/10.1016/j.tics.2007.05.008>
- Björkman, M., Juslin, P., & Winman, A. (1993). Realism of confidence in sensory discrimination: the underconfidence phenomenon. *Perception & Psychophysics*, 54(1), 75–81.
- Blake, R., Brascamp, J., & Heeger, D. J. (2014). Can binocular rivalry reveal neural correlates of consciousness? *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 369(1641), 20130211. <http://doi.org/10.1098/rstb.2013.0211>
- Blake, R., & Logothetis, N. K. (2002). Visual competition. *Nature Reviews. Neuroscience*, 3(1), 13–21. <http://doi.org/10.1038/nrn701>
- Blake, R., O'Shea, R. P., & Mueller, T. J. (1992). Spatial zones of binocular rivalry in central and peripheral vision. *Visual Neuroscience*, 8(5), 469–478. <http://doi.org/10.1017/S0952523800004971>
- Blanchard, D. C., & Blanchard, R. J. (1988). Ethoexperimental approaches to the biology of emotion. *Annual Review of Psychology*, 39, 43–68. <http://doi.org/10.1146/annurev.ps.39.020188.000355>
- Blanchette, I. (2006). Snakes, spiders, guns, and syringes: how specific are evolutionary constraints on the detection of threatening stimuli? *Quarterly Journal of Experimental Psychology (2006)*, 59(8), 1484–1504. <http://doi.org/10.1080/02724980543000204>
- Bonneh, Y. S., Cooperman, A., & Sagi, D. (2001). Motion-induced blindness in normal observers. *Nature*, 411(6839), 798–801. <http://doi.org/10.1038/35081073>
- Boucsein, W., Fowles, D. C., Grimnes, S., Ben-Shakhar, G., Roth, W. T., Dawson, M. E., ... Society for Psychophysiological Research Ad Hoc Committee on Electrodermal Measures. (2012). Publication

- recommendations for electrodermal measurements. *Psychophysiology*, 49(8), 1017–1034.
<http://doi.org/10.1111/j.1469-8986.2012.01384.x>
- Boyer, P., & Bergstrom, B. (2011). Threat-detection in child development: an evolutionary perspective. *Neuroscience and Biobehavioral Reviews*, 35(4), 1034–1041.
<http://doi.org/10.1016/j.neubiorev.2010.08.010>
- Bradley, M. M., Codispoti, M., Cuthbert, B. N., & Lang, P. J. (2001). Emotion and motivation I: defensive and appetitive reactions in picture processing. *Emotion (Washington, D.C.)*, 1(3), 276–298.
- Bradley, M. M., & Lang, P. J. (1994). Measuring emotion: the Self-Assessment Manikin and the Semantic Differential. *Journal of Behavior Therapy and Experimental Psychiatry*, 25(1), 49–59.
- Bradley, M. M., & Lang, P. J. (2000). Emotion and motivation. *Handbook of psychophysiology*, 2, 602–642.
- Brascamp, J. W., & Blake, R. (2012). Inattention abolishes binocular rivalry: perceptual evidence. *Psychological Science*, 23(10), 1159–1167. <http://doi.org/10.1177/0956797612440100>
- Brascamp, J. W., van Ee, R., Noest, A. J., Jacobs, R. H. A. H., & van den Berg, A. V. (2006). The time course of binocular rivalry reveals a fundamental role of noise. *Journal of Vision*, 6(11), 1244–1256.
<http://doi.org/10.1167/6.11.8>
- Breitmeyer, B. G. (2015). Psychophysical ‘blinding’ methods reveal a functional hierarchy of unconscious visual processing. *Consciousness and Cognition*, 35, 234–250.
<http://doi.org/10.1016/j.concog.2015.01.012>
- Breitmeyer, B. G., & Ogmen, H. (2000). Recent models and findings in visual backward masking: a comparison, review, and update. *Perception & Psychophysics*, 62(8), 1572–1595.
- Breska, A., Maoz, K., & Ben-Shakhar, G. (2011). Interstimulus intervals for skin conductance response measurement. *Psychophysiology*, 48(4), 437–440. <http://doi.org/10.1111/j.1469-8986.2010.01084.x>
- Button, K. S., Ioannidis, J. P. A., Mokrysz, C., Nosek, B. A., Flint, J., Robinson, E. S. J., & Munafò, M. R. (2013). Power failure: why small sample size undermines the reliability of neuroscience. *Nature Reviews Neuroscience*, 14(5), 365–376. doi:10.1038/nrn3475
- Cacioppo, J. T., Tassinary, L. G., & Berntson, G. G. (2007). *Handbook of Psychophysiology* (3rd ed.). New York: Cambridge University Press.
- Calvo, M. G., & Avero, P. (2005). Time course of attentional bias to emotional scenes in anxiety: Gaze direction and duration. *Cognition & Emotion*, 19(3), 433–451.
<http://doi.org/10.1080/02699930441000157>
- Campbell, B. A., Wood, G., & McBride, T. (1997). Origins of orienting and defensive responses: An evolutionary perspective. In P. J. Lang, R. F. Simons, & M. T. Balaban (Eds.), *Attention and orienting: Sensory and motivational processes* (pp. 41–67). Hillsdale, NJ: Erlbaum, Inc.
- Cannon, W. B. (1927). The James–Lange theory of emotions: a critical examination and an alternative theory. *American Journal of Psychology*, 39, 106–124
- Cannon, M.W. (1979). Contrast sensation: A linear function of stimulus contrast. *Vision Research*, 19, 1045–1052
- Capitão, L. P., Underdown, S. J. V., Vile, S., Yang, E., Harmer, C. J., & Murphy, S. E. (2014). Anxiety increases breakthrough of threat stimuli in continuous flash suppression. *Emotion (Washington, D.C.)*, 14(6), 1027–1036. <http://doi.org/10.1037/a0037801>

- Carlson, J. M., Beacher, F., Reinke, K. S., Habib, R., Harmon-Jones, E., Mujica-Parodi, L. R., & Hajcak, G. (2012). Nonconscious attention bias to threat is correlated with anterior cingulate cortex gray matter volume: a voxel-based morphometry result and replication. *NeuroImage*, 59(2), 1713–1718. <http://doi.org/10.1016/j.neuroimage.2011.09.040>
- Carlson, J. M., Cha, J., Harmon-Jones, E., Mujica-Parodi, L. R., & Hajcak, G. (2014). Influence of the BDNF genotype on amygdalo-prefrontal white matter microstructure is linked to nonconscious attention bias to threat. *Cerebral Cortex (New York, N.Y.: 1991)*, 24(9), 2249–2257. <http://doi.org/10.1093/cercor/bht089>
- Carlson, J. M., Cha, J., & Mujica-Parodi, L. R. (2013). Functional and structural amygdala – Anterior cingulate connectivity correlates with attentional bias to masked fearful faces. *Cortex*, 49(9), 2595–2600. <http://doi.org/10.1016/j.cortex.2013.07.008>
- Carlson, J. M., Fee, A. L., & Reinke, K. S. (2009). Backward masked snakes and guns modulate spatial attention. *Evolutionary Psychology*, 7(4), 147470490900700404.
- Carlson, J. M., & Mujica-Parodi, L. R. (2014). Facilitated Attentional Orienting and Delayed Disengagement to Conscious and Nonconscious Fearful Faces. *Journal of Nonverbal Behavior*, 39(1), 69–77. <http://doi.org/10.1007/s10919-014-0185-1>
- Carlson, J. M., Mujica-Parodi, L. R., Harmon-Jones, E., & Hajcak, G. (2012). The orienting of spatial attention to backward masked fearful faces is associated with variation in the serotonin transporter gene. *Emotion*, 12(2), 203–207. <http://doi.org/10.1037/a0025170>
- Carlson, J. M., & Reinke, K. S. (2008). Masked fearful faces modulate the orienting of covert spatial attention. *Emotion (Washington, D.C.)*, 8(4), 522–529. <http://doi.org/10.1037/a0012653>
- Carlson, J. M., & Reinke, K. S. (2010). Spatial attention-related modulation of the N170 by backward masked fearful faces. *Brain and Cognition*, 73(1), 20–27. <http://doi.org/10.1016/j.bandc.2010.01.007>
- Carlson, J. M., Reinke, K. S., & Habib, R. (2009). A left amygdala mediated network for rapid orienting to masked fearful faces. *Neuropsychologia*, 47(5), 1386–1389. <http://doi.org/10.1016/j.neuropsychologia.2009.01.026>
- Carlson, T. A., & He, S. (2000). Visible binocular beats from invisible monocular stimuli during binocular rivalry. *Current Biology*, 10(17), 1055–1058. [http://doi.org/10.1016/S0960-9822\(00\)00672-2](http://doi.org/10.1016/S0960-9822(00)00672-2)
- Carr, T. H., & Dagenbach, D. (1990). Semantic priming and repetition priming from masked words: evidence for a center-surround attentional mechanism in perceptual recognition. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 16(2), 341–350.
- Carter, O., & Cavanagh, P. (2007). Onset rivalry: brief presentation isolates an early independent phase of perceptual competition. *PloS One*, 2(4), e343. <http://doi.org/10.1371/journal.pone.0000343>
- Carter, O. L., Presti, D. E., Callistemon, C., Ungerer, Y., Liu, G. B., & Pettigrew, J. D. (2005). Meditation alters perceptual rivalry in Tibetan Buddhist monks. *Current Biology: CB*, 15(11), R412–413. <http://doi.org/10.1016/j.cub.2005.05.043>
- Chalmers, D. J. (1995). Facing up to the problem of consciousness. *Journal of Consciousness Studies*, 2(3), 200–219.
- Cheesman, J., & Merikle, P. M. (1986). Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 40(4), 343–367. <http://doi.org/10.1037/h0080103>

- Chen, Y.-C., & Yeh, S.-L. (2012). Look into my eyes and I will see you: unconscious processing of human gaze. *Consciousness and Cognition*, 21(4), 1703–1710. <http://doi.org/10.1016/j.concog.2012.10.001>
- Cheung, M. W. L. (2014). Modeling dependent effect sizes with three-level meta-analyses: a structural equation modeling approach. *Psychological Methods*, 19(2), 211.
- Chica, A. B., Martín-Arévalo, E., Botta, F., & Lupiáñez, J. (2014). The Spatial Orienting paradigm: How to design and interpret spatial attention experiments. *Neuroscience & Biobehavioral Reviews*, 40, 35–51. <http://doi.org/10.1016/j.neubiorev.2014.01.002>
- Chun, M. M., & Marois, R. (2002). The dark side of visual attention. *Current Opinion in Neurobiology*, 12(2), 184–189.
- Codispoti, M., & De Cesare, A. (2007). Arousal and attention: picture size and emotional reactions. *Psychophysiology*, 44(5), 680–686. <http://doi.org/10.1111/j.1469-8986.2007.00545.x>
- Codispoti, M., Mazzetti, M., & Bradley, M. M. (2009). Unmasking emotion: exposure duration and emotional engagement. *Psychophysiology*, 46(4), 731–738. <http://doi.org/10.1111/j.1469-8986.2009.00804.x>
- Coelho, C. M., Cloete, S., & Wallis, G. (2010). The face-in-the-crowd effect: when angry faces are just cross(es). *Journal of Vision*, 10(1), 7.1-14. <http://doi.org/10.1167/10.1.7>
- Cohen, J. (1977). *Statistical power analysis for the behavioral sciences*. New York, NY: Routledge.
- Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, 16(8), 411–417. <http://doi.org/10.1016/j.tics.2012.06.013>
- Cook, M., & Mineka, S. (1991). Selective associations in the origins of phobic fears and their implications for behavior therapy. In *Handbook of behavior therapy and psychological science: An integrative approach* (pp. 413–434). Elmsford, NY, US: Pergamon Press.
- Cooper, H., & Hedges, L. V. (1994). *The handbook of research synthesis*. New York: Russell Sage Foundation.
- Coren, S., & Russell, J. A. (1992). The Relative Dominance of Different Facial Expressions of Emotion under Conditions of Perceptual Ambiguity. *Cognition and Emotion*, 6(5), 339–356. <http://doi.org/10.1080/02699939208409690>
- Corthout, E., Uttl, B., Ziemann, U., Cowey, A., & Hallett, M. (1999). Two periods of processing in the (circum)striate visual cortex as revealed by transcranial magnetic stimulation. *Neuropsychologia*, 37(2), 137–145.
- Crick, F., & Koch, C. (1998). Consciousness and neuroscience. *Cerebral Cortex (New York, N.Y.: 1991)*, 8(2), 97–107.
- Cumming, G. (2011). *Understanding the new statistics: Effect sizes, confidence intervals, and meta-analysis*. New York, NY: Routledge.
- Dalgleish, T. (2004). The emotional brain. *Nature Reviews Neuroscience*, 5(7), 583–589. <http://doi.org/10.1038/nrn1432>
- Dalgleish, T., & Power, M. (1999). *The handbook of cognition and emotion*. Sussex, UK: Wiley and Sons, Ltd. <http://dx.doi.org/10.1002/0470013494>

- Dawson, M. E., Schell, A. M., & Filion, D. L. (2000). The electrodermal system. In J. T. Cacioppo, L. J. Tassinary, & G. G. Berntson (Eds.), *Handbook of Psychophysiology* (2nd ed, pp. 200–223). Cambridge, UK: Cambridge 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(18), 3759–3763.
- Dehaene, S., & Changeux, J.-P. (2011). Experimental and theoretical approaches to conscious processing. *Neuron*, 70(2), 200–227. <http://doi.org/10.1016/j.neuron.2011.03.018>
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211. <http://doi.org/10.1016/j.tics.2006.03.007>
- Demos, K. E., Kelley, W. M., Ryan, S. L., Davis, F. C., & Whalen, P. J. (2008). Human amygdala sensitivity to the pupil size of others. *Cerebral Cortex (New York, N.Y.: 1991)*, 18(12), 2729–2734. <http://doi.org/10.1093/cercor/bhn034>
- Dennett, D. C. (1993). *Consciousness Explained*. Penguin UK.
- De Valois, K.K. (1977). Spatial frequency adaptation can enhance contrast sensitivity. *Vision Research*, 17, 1057–1065.
- Dimberg, U. (1982). Facial Reactions to Facial Expressions. *Psychophysiology*, 19(6), 643–647. <http://doi.org/10.1111/j.1469-8986.1982.tb02516.x>
- Dukas, R. (1998). *Cognitive Ecology: The Evolutionary Ecology of Information Processing and Decision Making*. Chicago: University of Chicago.
- Dunlap, W. P., Cortina, J. M., Vaslow, J. B., & Burke, M. J. (1996). Meta-analysis of experiments with matched groups or repeated measures designs. *Psychological Methods*, 1(2), 170–177. doi:10.1037/1082-989X.1.2.170
- Duval, S. J. (2005). The trim and fill method. In H. R. Rothstein, A. J. Sutton, & M. Borenstein (Eds.), *Publication bias in meta-analysis: Prevention, assessment, and adjustments* (pp. 127–144). Chichester, England: Wiley.
- Ekman, P., & Friesen, W. V. (1976). *Pictures of facial affect*. Palo Alto, CA; Consulting Psychologists.
- Efron, B., & Tibshirani, R. J. (1993). *An Introduction to the Bootstrap*. New York: Chapman & Hall.
- Filmer, H. L., & Monsell, S. (2013). TMS to V1 spares discrimination of emotive relative to neutral body postures. *Neuropsychologia*, 51(13), 2485–2491. <http://doi.org/10.1016/j.neuropsychologia.2013.09.029>
- Eriksen, C. W. (1960). Discrimination and learning without awareness: A methodological survey and evaluation. *Psychological Review*, 67, 279–300.
- Esteves, F., Parra, C., Dimberg, U., & Ohman, A. (1994). Nonconscious associative learning: Pavlovian conditioning of skin conductance responses to masked fear-relevant facial stimuli. *Psychophysiology*, 31(4), 375–385.
- Eysenck, M. W., Derakshan, N., Santos, R., & Calvo, M. G. (2007). Anxiety and cognitive performance: attentional control theory. *Emotion (Washington, D.C.)*, 7(2), 336–353. <http://doi.org/10.1037/1528-3542.7.2.336>

- Faivre, N., Berthet, V., & Kouider, S. (2012). Nonconscious influences from emotional faces: a comparison of visual crowding, masking, and continuous flash suppression. *Consciousness Research*, 3, 129. <http://doi.org/10.3389/fpsyg.2012.00129>
- Fang, F., & He, S. (2005). Cortical responses to invisible objects in the human dorsal and ventral pathways. *Nature Neuroscience*, 8(10), 1380–1385. <http://doi.org/10.1038/nn1537>
- Feinstein, J. S., Adolphs, R., Damasio, A., & Tranel, D. (2011). The human amygdala and the induction and experience of fear. *Current Biology: CB*, 21(1), 34–38. <http://doi.org/10.1016/j.cub.2010.11.042>
- Fiddick, L. (2011). There is more than the amygdala: Potential threat assessment in the cingulate cortex. *Neuroscience & Biobehavioral Reviews*, 35(4), 1007–1018. <http://doi.org/10.1016/j.neubiorev.2010.09.014>
- Field, A. P., Lawson, J., & Banerjee, R. (2008). The verbal threat information pathway to fear in children: the longitudinal effects on fear cognitions and the immediate effects on avoidance behavior. *Journal of Abnormal Psychology*, 117(1), 214–224. <http://doi.org/10.1037/0021-843X.117.1.214>
- Flykt, A., Esteves, F., & Öhman, A. (2007). Skin conductance responses to masked conditioned stimuli: Phylogenetic/ontogenetic factors versus direction of threat? *Biological Psychology*, 74(3), 328–336. <http://doi.org/10.1016/j.biopsycho.2006.08.004>
- Fox, E. (2002). Processing emotional facial expressions: The role of anxiety and awareness. *Cognitive, Affective & Behavioral Neuroscience*, 2(1), 52–63.
- Fox, E. (2008). *Emotion science cognitive and neuroscientific approaches to understanding human emotions*. Palgrave Macmillan.
- Fox, E., Cahill, S., & Zougkou, K. (2010). Preconscious processing biases predict emotional reactivity to stress. *Biological Psychiatry*, 67(4), 371–377. <http://doi.org/10.1016/j.biopsych.2009.11.018>
- Gainotti, G., Caltagirone, C., & Zoccolotti, P. (1993). Left/right and cortical/subcortical dichotomies in the neuropsychological study of human emotions. *Cognition and Emotion*, 7(1), 71–93. <http://doi.org/10.1080/02699939308409178>
- Garrido, M. I. (2012). Brain Connectivity: The Feel of Blindsight. *Current Biology*, 22(15), R599–R600. <http://doi.org/10.1016/j.cub.2012.06.012>
- Gayet, S., Van der Stigchel, S., & Paffen, C. L. E. (2014). Breaking continuous flash suppression: competing for consciousness on the pre-semantic battlefield. *Frontiers in Psychology*, 5, 460. <http://doi.org/10.3389/fpsyg.2014.00460>
- Georgeson, M.A., & Sullivan, G.D. (1975). Contrast constancy: Deblurring in human vision by spatial frequency channels. *Journal of Physiology*, 252, 627–656.
- Getov, S., Kanai, R., Bahrami, B., & Rees, G. (2015). Human brain structure predicts individual differences in preconscious evaluation of facial dominance and trustworthiness. *Social Cognitive and Affective Neuroscience*, 10(5), 690–699. <http://doi.org/10.1093/scan/nsu103>
- Gerdes, A., & Alpers, G. W. (2014). You see what you fear: Spiders gain preferential access to conscious perception in spider-phobic patients. *Journal of Experimental Psychopathology*, 5, 14–28. <http://dx.doi.org/10.5127/jep.033212>

- Giesbrecht, B., Bischof, W. F., & Kingstone, A. (2004). Seeing the light: Adapting luminance reveals low-level visual processes in the attentional blink. *Brain and Cognition*, 55(2), 307–309.
<http://doi.org/10.1016/j.bandc.2004.02.027>
- Goodale, M., & Milner, D. (2013). *Sight unseen: An exploration of conscious and unconscious vision*. OUP Oxford.
- Graf, E. W., & Adams, W. J. (2008). Surface organization influences bistable vision. *Journal of Experimental Psychology: Human Perception and Performance*, 34(2), 502–508.
<http://doi.org/10.1037/0096-1523.34.2.502>
- Gray, K. L. H., Adams, W. J., & Garner, M. (2009). The influence of anxiety on the initial selection of emotional faces presented in binocular rivalry. *Cognition*, 113(1), 105–110.
<http://doi.org/10.1016/j.cognition.2009.06.009>
- Gray, K. L. H., Adams, W. J., Hedger, N., Newton, K. E., & Garner, M. (2013). Faces and awareness: low-level, not emotional factors determine perceptual dominance. *Emotion (Washington, D.C.)*, 13(3), 537–544. <http://doi.org/10.1037/a0031403>
- Green, M. F., Glahn, D., Engel, S. A., Nuechterlein, K. H., Sabb, F., Strojwas, M., & Cohen, M. S. (2005). Regional brain activity associated with visual backward masking. *Journal of Cognitive Neuroscience*, 17, 13–23.
- Green, D. M. Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: Wiley.
- Greenhouse, J. B., & Iyengar, S. (1994). Sensitivity analysis and diagnostics. In H. Cooper & L. V. Hedges (Eds.), *The handbook of research synthesis*. New York, NY: Russell Sage Foundation Publications.
- Greenwald, M. K., Cook, E. W., & Lang, P. J. (1989). Affective judgment and psychophysiological response: Dimensional covariation in the evaluation of pictorial stimuli. *Journal of Psychophysiology*, 3(1), 51–64.
- Grieve, K. L., Acuña, C., & Cudeiro, J. (2000). The primate pulvinar nuclei: vision and action. *Trends in Neurosciences*, 23(1), 35–39.
- Grissom, R. J., & Kim, J. J. (2005). *Effect sizes for research: A broad practical approach*. Mahwah, NJ: Erlbaum.
- Grillon, C., & Cornwell, B. (2007). Comments on ‘Fearful and Sexual Pictures Not Consciously Seen Modulate the Startle Reflex in Human Beings’, by Ruiz-Padial and Vila, Online on the Biological Psychiatry web site. *Biological Psychiatry*, 62(5), 541–542.
<http://doi.org/10.1016/j.biopsych.2006.12.026>
- Grill-Spector, K., & Kanwisher, N. (2005). Visual recognition: as soon as you know it is there, you know what it is. *Psychological Science*, 16(2), 152–160. <http://doi.org/10.1111/j.0956-7976.2005.00796.x>
- Hannula, D. E., Simons, D. J., & Cohen, N. J. (2005). Imaging implicit perception: promise and pitfalls. *Nature Reviews Neuroscience*, 6(3), 247–255. <http://doi.org/10.1038/nrn1630>
- Guilford, T., & Dawkins, M. S. (1991). Receiver psychology and the evolution of animal signals. *Animal behaviour*, 42(1), 1–14.
- Hassin, R. R. (2013). Yes It Can: On the Functional Abilities of the Human Unconscious. *Perspectives on Psychological Science: A Journal of the Association for Psychological Science*, 8(2), 195–207.
<http://doi.org/10.1177/1745691612460684>
- Haun, A. M., & Peli, E. (2013). Perceived contrast in complex images. *Journal of vision*, 13(13), 3–3.

- Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, 293(5539), 2425–2430. <http://doi.org/10.1126/science.1063736>
- Hebart, M. N., & Hesselmann, G. (2012). What Visual Information Is Processed in the Human Dorsal Stream? *The Journal of Neuroscience*, 32(24), 8107–8109. <http://doi.org/10.1523/JNEUROSCI.1462-12.2012>
- Hedger, N. (2014). Fearfulness and effective contrast. *figshare*.
<http://dx.doi.org/10.6084/m9.figshare.1104425>
- Hedger, N., Adams, W. J., & Garner, M. (2013). Fearful faces: no emotion-based processing without awareness under continuous flash suppression. *Journal of Vision*, 13, 586–586. doi: 10.1167/13.9.586
- Hedger, N., Adams, W. J., & Garner, M. (2015a). Autonomic arousal and attentional orienting to visual threat are predicted by awareness. *Journal of Experimental Psychology. Human Perception and Performance*, 41(3), 798–806. <http://doi.org/10.1037/xhp0000051>
- Hedger, N., Adams, W. J., & Garner, M. (2015b). Fearful faces have a sensory advantage in the competition for awareness. *Journal of Experimental Psychology. Human Perception and Performance*, 41(6), 1748–1757. <http://doi.org/10.1037/xhp0000127>
- Hedger, N., Gray, K. L. H., Garner, M., & Adams, W. J. (2016). Are Visual Threats Prioritized Without Awareness? A Critical Review and Meta-Analysis Involving 3 Behavioral Paradigms and 2696 Observers. *Psychological Bulletin*. <http://doi.org/10.1037/bul0000054>
- Heil, M., Rolke, B., & Pecchinenda, A. (2004). Automatic Semantic Activation Is No Myth Semantic Context Effects on the N400 in the Letter-Search Task in the Absence of Response Time Effects. *Psychological Science*, 15(12), 852–857. <http://doi.org/10.1111/j.0956-7976.2004.00766.x>
- Helzer, E. G., Connor-Smith, J. K., & Reed, M. A. (2009). Traits, States, and Attentional Gates: Temperament and Threat Relevance as Predictors of Attentional Bias to Social Threat. *Anxiety, Stress, and Coping*, 22(1), 57–76. <http://doi.org/10.1080/10615800802272244>
- Herry, C., Bach, D. R., Esposito, F., Di Salle, F., Perrig, W. J., Scheffler, K., ... Seifritz, E. (2007). Processing of temporal unpredictability in human and animal amygdala. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 27(22), 5958–5966. <http://doi.org/10.1523/JNEUROSCI.5218-06.2007>
- Hesselmann, G., & Moors, P. (2015). Definitely maybe: can unconscious processes perform the same functions as conscious processes? *Consciousness Research*, 584. <http://doi.org/10.3389/fpsyg.2015.00584>
- Hoffmann, M., Lipka, J., Mothes-Lasch, M., Miltner, W. H. R., & Straube, T. (2012). Awareness modulates responses of the amygdala and the visual cortex to highly arousing visual threat. *NeuroImage*, 62(3), 1439–1444. <http://doi.org/10.1016/j.neuroimage.2012.05.064>
- Holender, D. (1986). Semantic activation without conscious identification in dichotic listening, parafoveal vision, and visual masking: A survey and appraisal. *Behavioral and Brain Sciences*, 9, 1-23.
- Hong, S. W., & Blake, R. (2009). Interocular suppression differentially affects achromatic and chromatic mechanisms. *Attention, Perception & Psychophysics*, 71(2), 403–411. <http://doi.org/10.3758/APP.71.2.403>

-
- Horstmann, G., & Ansorge, U. (2009). Visual search for facial expressions of emotions: A comparison of dynamic and static faces. *Emotion*, 9(1), 29.
- Horstmann, G., & Bauland, A. (2006). Search asymmetries with real faces: testing the anger-superiority effect. *Emotion*, 6(2), 193.
- Howard, I. P., & Rogers, B. J. (1996). *Binocular Vision and Stereopsis*. New York: Oxford University Press.
- Howell, D.C. (2013). *Statistical methods for Psychology (8th ed.)* Belmont, CA: Wadsworth.
- Hunt, C., Keogh, E., & French, C. C. (2007). Anxiety sensitivity, conscious awareness and selective attentional biases in children. *Behaviour Research and Therapy*, 45(3), 497–509.
<http://doi.org/10.1016/j.brat.2006.04.001>
- Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews. Neuroscience*, 2(3), 194–203. <http://doi.org/10.1038/35058500>
- Jacobs, C., de Graaf, T. A., Goebel, R., & Sack, A. T. (2012). The Temporal Dynamics of Early Visual Cortex Involvement in Behavioral Priming. *PLoS ONE*, 7(11), e48808.
<http://doi.org/10.1371/journal.pone.0048808>
- James, W. (1890). *Principles of psychology*. New York: Holt.
- Jiang, Y., & He, S. (2006). Cortical responses to invisible faces: dissociating subsystems for facial-information processing. *Current Biology: CB*, 16(20), 2023–2029.
<http://doi.org/10.1016/j.cub.2006.08.084>
- Jiang, Y., Costello, P., & He, S. (2007). Processing of invisible stimuli: advantage of upright faces and recognizable words in overcoming interocular suppression. *Psychological Science*, 18(4), 349–355.
<http://doi.org/10.1111/j.1467-9280.2007.01902.x>
- Jiang, Y., Shannon, R. W., Vizueta, N., Bernat, E. M., Patrick, C. J., & He, S. (2009). Dynamics of processing invisible faces in the brain: automatic neural encoding of facial expression information. *NeuroImage*, 44(3), 1171–1177. <http://doi.org/10.1016/j.neuroimage.2008.09.038>
- Jiang, Y., Zhou, K., & He, S. (2007). Human visual cortex responds to invisible chromatic flicker. *Nature Neuroscience*, 10(5), 657–662. <http://doi.org/10.1038/nn1879>
- Jolij, J., & Lamme, V. A. F. (2005). Repression of unconscious information by conscious processing: evidence from affective blindsight induced by transcranial magnetic stimulation. *Proceedings of the National Academy of Sciences of the United States of America*, 102(30), 10747–10751.
<http://doi.org/10.1073/pnas.0500834102>
- Jusyte, A., Mayer, S. V., Künzel, E., Hautzinger, M., & Schönenberg, M. (2015). Unemotional traits predict early processing deficit for fearful expressions in young violent offenders: an investigation using continuous flash suppression. *Psychological Medicine*, 45(2), 285–297.
<http://doi.org/10.1017/S0033291714001287>
- Kanai, R., & Tsuchiya, N. (2012). Qualia. *Current Biology*, 22(10), R392–R396.
<http://doi.org/10.1016/j.cub.2012.03.033>
- Kanai, R., Tsuchiya, N., & Verstraten, F. A. J. (2006). The scope and limits of top-down attention in unconscious visual processing. *Current Biology: CB*, 16(23), 2332–2336.
<http://doi.org/10.1016/j.cub.2006.10.001>

- Kanwisher, N., McDermott, J., & Chun, M. M. (1997). The fusiform face area: a module in human extrastriate cortex specialized for face perception. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 17(11), 4302–4311.
- Katkin, E. S., Wiens, S., & Ohman, A. (2001). Nonconscious fear conditioning, visceral perception, and the development of gut feelings. *Psychological Science*, 12(5), 366–370.
- Katsikitis, M. (2003). *The Human Face: Measurement and Meaning*. Dordrecht: Kluwer.
- Keogh, E., Thompson, T., & Hannent, I. (2003). Selective attentional bias, conscious awareness and the fear of pain. *Pain*, 104(1–2), 85–91.
- Kim, C.-Y., & Blake, R. (2005). Psychophysical magic: rendering the visible ‘invisible’. *Trends in Cognitive Sciences*, 9(8), 381–388. <http://doi.org/10.1016/j.tics.2005.06.012>
- Knapen, T., Brascamp, J., Pearson, J., Ee, R. van, & Blake, R. (2011). The Role of Frontal and Parietal Brain Areas in Bistable Perception. *The Journal of Neuroscience*, 31(28), 10293–10301. <http://doi.org/10.1523/JNEUROSCI.1727-11.2011>
- Kobayashi, H., & Kohshima, S. (1997). Unique morphology of the human eye. *Nature*, 387, 767–768. doi:10.1038/42842
- Koch, C. (2004). *The quest for consciousness: A neurobiological approach* (1st ed.). Denver, CO: Roberts & Company Publishers.
- Koch, C., & Reid, R. C. (2012). Neuroscience: Observatories of the mind. *Nature*, 483(7390), 397–398. <http://doi.org/10.1038/483397a>
- Koster, E. H. W., Verschuere, B., Burssens, B., Custers, R., & Crombez, G. (2007). Attention for emotional faces under restricted awareness revisited: do emotional faces automatically attract attention? *Emotion (Washington, D.C.)*, 7(2), 285–295. <http://doi.org/10.1037/1528-3542.7.2.285>
- Kouider, S., Berthet, V., & Faivre, N. (2011). Preference Is Biased by Crowded Facial Expressions. *Psychological Science*. <http://doi.org/10.1177/0956797610396226>
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14(7), 301–307. <http://doi.org/10.1016/j.tics.2010.04.006>
- Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 362(1481), 857–875. <http://doi.org/10.1098/rstb.2007.2093>
- Kreiman, G., Fried, I., & Koch, C. (2002). Single-neuron correlates of subjective vision in the human medial temporal lobe. *Proceedings of the National Academy of Sciences of the United States of America*, 99(12), 8378–8383. <http://doi.org/10.1073/pnas.072194099>
- Krolak-Salmon, P., Hénaff, M.-A., Vighetto, A., Bertrand, O., & Mauguière, F. (2004). Early Amygdala Reaction to Fear Spreading in Occipital, Temporal, and Frontal Cortex: A Depth Electrode ERP Study in Human. *Neuron*, 42(4), 665–676. [http://doi.org/10.1016/S0896-6273\(04\)00264-8](http://doi.org/10.1016/S0896-6273(04)00264-8)
- Lähteenmäki, M., Hyönä, J., Koivisto, M., & Nummenmaa, L. (2015). Affective processing requires awareness. *Journal of Experimental Psychology. General*, 144(2), 339–365. <http://doi.org/10.1037/xge0000040>

- Lakens, D. (2013). Calculating and reporting effect sizes to facilitate cumulative science: a practical primer for t-tests and ANOVAs. *Frontiers in Psychology*, 4, 863. <http://doi.org/10.3389/fpsyg.2013.00863>
- Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1), 12–18. [http://doi.org/10.1016/S1364-6613\(02\)00013-X](http://doi.org/10.1016/S1364-6613(02)00013-X)
- Lamme, V. A., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23(11), 571–579.
- Lang, P. J., Bradley, M. M., & Cuthbert, B. N. (2008). *International affective picture system (IAPS): Affective ratings of pictures and instruction manual* (Tech. Rep. No. A–8). Gainesville, FL: University of Florida.
- Lange, C., 1885. *The mechanism of the emotions*. In: Dunlap, E. (Ed.), *The Emotions*. Williams & Wilkins, Baltimore, Maryland, pp. 33–92.
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H. J., Hawk, S.T., & van Knippenberg, A. (2010). Presentation and validation of the Radboud Faces Database. *Cognition & Emotion*, 24, 1377–1388. doi: 10.1080/02699930903485076
- Lapate, R. C., Rokers, B., Li, T., & Davidson, R. J. (2013). Nonconscious Emotional Activation Colors First Impressions A Regulatory Role for Conscious Awareness. *Psychological Science*, 956797613503175. <http://doi.org/10.1177/0956797613503175>
- LeDoux, J. (1998). *The emotional brain: The mysterious underpinnings of emotional life*. Simon and Schuster.
- LeDoux, J. E. (2013). The slippery slope of fear. *Trends in Cognitive Sciences*, 17(4), 155–156. <http://doi.org/10.1016/j.tics.2013.02.004>
- LeDoux, J. E., Sakaguchi, A., & Reis, D. J. (1984). Subcortical efferent projections of the medial geniculate nucleus mediate emotional responses conditioned to acoustic stimuli. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 4(3), 683–698.
- Lee, D. H., Mirza, R., Flanagan, J. G., & Anderson, A. K. (2014). Optical origins of opposing facial expression actions. *Psychological science*, 25(3), 745–752.
- Lee, D. H., Susskind, J. M., & Anderson, A. K. (2013). Social transmission of the sensory benefits of eye widening in fear expressions. *Psychological science*, 24(6), 957–965. Lee, S.-H., & Blake, R. (2004). A fresh look at interocular grouping during binocular rivalry. *Vision Research*, 44(10), 983–991. <http://doi.org/10.1016/j.visres.2003.12.007>
- Lehky, S. R. (1995). Binocular rivalry is not chaotic. *Proceedings. Biological Sciences / The Royal Society*, 259(1354), 71–76. <http://doi.org/10.1098/rspb.1995.0011>
- Lerner, Y., Singer, N., Gonen, T., Weintraub, Y., Cohen, O., Rubin, N., ... Hendler, T. (2012). Feeling without seeing? Engagement of ventral, but not dorsal, amygdala during unaware exposure to emotional faces. *Journal of Cognitive Neuroscience*, 24(3), 531–542. http://doi.org/10.1162/jocn_a_00165
- Levelt, W. J. (1965). *On binocular rivalry* (Doctoral dissertation, Van Gorcum Assen).
- Levi, D. M. (2008). Crowding—an essential bottleneck for object recognition: a mini-review. *Vision Research*, 48(5), 635–654. <http://doi.org/10.1016/j.visres.2007.12.009>

- Liddell, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., ... Williams, L. M. (2005). A direct brainstem–amygdala–cortical ‘alarm’ system for subliminal signals of fear. *NeuroImage*, 24(1), 235–243. <http://doi.org/10.1016/j.neuroimage.2004.08.016>
- Lin, Z., & He, S. (2009). Seeing the invisible: the scope and limits of unconscious processing in binocular rivalry. *Progress in Neurobiology*, 87(4), 195–211. <http://doi.org/10.1016/j.pneurobio.2008.09.002>
- Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: anatomy, physiology, and perception. *Science (New York, N.Y.)*, 240(4853), 740–749.
- López-López, J. A., Marín-Martínez, F., Sánchez-Meca, J., Van den Noortgate, W., & Viechtbauer, W. (2013). Estimation of the predictive power of the model in mixed-effects meta-regression: A simulation study. *The British Journal of Mathematical and Statistical Psychology*. <http://doi.org/10.1111/bmsp.12002>
- Lovibond, S. H., & Lovibond, P. F. (1996). *Manual for the Depression Anxiety Stress Scales*. Psychology Foundation of Australia.
- Luo, Q., Holroyd, T., Majestic, C., Cheng, X., Schechter, J., & Blair, R. J. (2010). Emotional Automaticity Is a Matter of Timing. *The Journal of Neuroscience*, 30(17), 5825–5829. <http://doi.org/10.1523/JNEUROSCI.BC-5668-09.2010>
- Lykken, D. T., & Venables, P. H. (1971). Direct measurement of skin conductance: a proposal for standardization. *Psychophysiology*, 8(5), 656–672.
- Macdonald, J. S. P., & Lavie, N. (2008). Load Induced Blindness. *Journal of Experimental Psychology: Human Perception and Performance*, 34(5), 1078–1091. <http://doi.org/10.1037/0096-1523.34.5.1078>
- Macmillan, N. A. Creelman, C. D. (2005). *Detection theory: A user's guide*. New York: Cambridge University Press.
- MacLean, P. D. (1949). Psychosomatic disease and the ‘visceral brain’: recent developments bearing on the Papez theory of emotion. *Psychosomatic Medicine*, 11, 338–353.
- Maoz, K., Abend, R., Fox, N. A., Pine, D. S., & Bar-Haim, Y. (2013). Subliminal attention bias modification training in socially anxious individuals. *Frontiers in Human Neuroscience*, 7, 389. <http://doi.org/10.3389/fnhum.2013.00389>
- Maratos, F. A. (2011). Temporal processing of emotional stimuli: The capture and release of attention by angry faces. *Emotion*, 11(5), 1242–1247. <http://doi.org/10.1037/a0024279>
- Marois, R., & Ivanoff, J. (2005). Capacity limits of information processing in the brain. *Trends in Cognitive Sciences*, 9(6), 296–305. <http://doi.org/10.1016/j.tics.2005.04.010>
- Mather, M., & Knight, M. R. (2006). Angry faces get noticed quickly: Threat detection is not impaired among older adults. *The Journals of Gerontology Series B: Psychological Sciences and Social Sciences*, 61(1), P54–P57.
- Mathews, A., & MacLeod, C. (2005). Cognitive vulnerability to emotional disorders. *Annual Review of Clinical Psychology*, 1, 167–195. <http://doi.org/10.1146/annurev.clinpsy.1.102803.143916>
- Mayer, B., Merckelbach, H., de Jong, P. J., & Leeuw, I. (1999). Skin conductance responses of spider phobics to backwardly masked phobic cues. *Journal of Psychophysiology*, 13(3), 152–159. <http://doi.org/10.1027//0269-8803.13.3.152>

- McCrory, E. J., Brito, S. A. D., Kelly, P. A., Bird, G., Sebastian, C. L., Mechelli, A., ... Viding, E. (2013). Amygdala activation in maltreated children during pre-attentive emotional processing. *The British Journal of Psychiatry*, 202(4), 269–276. <http://doi.org/10.1192/bjp.bp.112.116624>
- Méndez-Bértolo, C., Moratti, S., Toledano, R., Lopez-Sosa, F., Martínez-Alvarez, R., Mah, Y. H., ... & Strange, B. A. (2016). A fast pathway for fear in human amygdala. *Nature neuroscience*.
- Merikle, P. M., Smilek, D., & Eastwood, J. D. (2001). Perception without awareness: perspectives from cognitive psychology. *Cognition*, 79(1–2), 115–134.
- Mermillod, M., Droit-Volet, S., Devaux, D., Schaefer, A., & Vermeulen, N. (2010). Are coarse scales sufficient for fast detection of visual threat? *Psychological Science*, 21(10), 1429–1437. <http://doi.org/10.1177/0956797610381503>
- Michael, L., Hesselmann, G., Kiefer, M., & Niedeggen, M. (2011). Distractor-induced blindness for orientation changes and coherent motion. *Vision Research*, 51(15), 1781–1787. <http://doi.org/10.1016/j.visres.2011.06.007>
- Milders, M., Hietanen, J. K., Leppänen, J. M., & Braun, M. (2011). Detection of emotional faces is modulated by the direction of eye gaze. *Emotion (Washington, D.C.)*, 11(6), 1456–1461. <http://doi.org/10.1037/a0022901>
- Mitchell, J. F., Stoner, G. R., & Reynolds, J. H. (2004). Object-based attention determines dominance in binocular rivalry. *Nature*, 429(6990), 410–413. <http://doi.org/10.1038/nature02584>
- Mobbs, D., Petrovic, P., Marchant, J. L., Hassabis, D., Weiskopf, N., Seymour, B., ... Frith, C. D. (2007). When Fear Is Near. *Science (New York, N.Y.)*, 317(5841), 1079–1083. <http://doi.org/10.1126/science.1144298>
- Mogg, K., & Bradley, B. P. (1998). A cognitive-motivational analysis of anxiety. *Behaviour Research and Therapy*, 36(9), 809–848. [http://doi.org/10.1016/S0005-7967\(98\)00063-1](http://doi.org/10.1016/S0005-7967(98)00063-1)
- Mogg, K., & Bradley, B. P. (1999). Orienting of Attention to Threatening Facial Expressions Presented under Conditions of Restricted Awareness. *Cognition and Emotion*, 13(6), 713–740. <http://doi.org/10.1080/026999399379050>
- Mogg, K., & Bradley, B. P. (2002). Selective orienting of attention to masked threat faces in social anxiety. *Behaviour Research and Therapy*, 40(12), 1403–1414.
- Mogg, K., & Bradley, B. P. (2005). Attentional bias in generalized anxiety disorder versus depressive disorder. *Cognitive therapy and research*, 29(1), 29–45.
- Mogg, K., & Bradley, B. P. (2006). Time course of attentional bias for fear-relevant pictures in spider-fearful individuals. *Behaviour Research and Therapy*, 44(9), 1241–1250. <http://doi.org/10.1016/j.brat.2006.05.003>
- Mogg, K., Bradley, B. P., & Hallowell, N. (1994). Attentional bias to threat: roles of trait anxiety, stressful events, and awareness. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, 47(4), 841–864.
- Mogg, K., Bradley, B. P., & Williams, R. (1995). Attentional bias in anxiety and depression: the role of awareness. *The British Journal of Clinical Psychology / the British Psychological Society*, 34 (Pt 1), 17–36.

- Mogg, K., Garner, M., & Bradley, B. P. (2007). Anxiety and orienting of gaze to angry and fearful faces. *Biological Psychology*, 76(3), 163–169. <http://doi.org/10.1016/j.biopsycho.2007.07.005>
- Mogg, K., Philippot, P., & Bradley, B. P. (2004). Selective attention to angry faces in clinical social phobia. *Journal of Abnormal Psychology*, 113(1), 160–165. <http://doi.org/10.1037/0021-843X.113.1.160>
- Moher, D., Liberati, A., Tetzlaff, J., Altman, D. G., & The PRISMA Group. (2009). Preferred Reporting Items for Systematic Reviews and Meta-Analyses: The PRISMA Statement. *PLoS Med*, 6(7), e1000097. <http://doi.org/10.1371/journal.pmed.1000097>
- Monk CS, Telzer EH, Mogg K, & et al. (2008). AMygdala and ventrolateral prefrontal cortex activation to masked angry faces in children and adolescents with generalized anxiety disorder. *Archives of General Psychiatry*, 65(5), 568–576. <http://doi.org/10.1001/archpsyc.65.5.568>
- Montgomery, S. A., & Asberg, M. (1979). A new depression scale designed to be sensitive to change. *The British Journal of Psychiatry*, 134, 382–389. <http://dx.doi.org/10.1192/bjp.134.4.382>
- Moore, D. A., & Healy, P. J. (2008). The trouble with overconfidence. *Psychological Review*, 115(2), 502–517. <http://doi.org/10.1037/0033-295X.115.2.502>
- Moran, T. P., & Moser, J. S. (2015). The color of anxiety: neurobehavioral evidence for distraction by perceptually salient stimuli in anxiety. *Cognitive, Affective & Behavioral Neuroscience*, 15(1), 169–179. <http://doi.org/10.3758/s13415-014-0314-7>
- Moreno, N., & González, A. (2007). Evolution of the amygdaloid complex in vertebrates, with special reference to the anamnio-amniotic transition. *Journal of Anatomy*, 211(2), 151–163. <http://doi.org/10.1111/j.1469-7580.2007.00780.x>
- Mormann, F., Dubois, J., Kornblith, S., Milosavljevic, M., Cerf, M., Ison, M., ... Koch, C. (2011). A category-specific response to animals in the right human amygdala. *Nature Neuroscience*, 14(10), 1247–1249. <http://doi.org/10.1038/nn.2899>
- Mormann, F., Kornblith, S., Quiñ Quiroga, R., Kraskov, A., Cerf, M., Fried, I., & Koch, C. (2008). Latency and Selectivity of Single Neurons Indicate Hierarchical Processing in the Human Medial Temporal Lobe. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 28(36), 8865–8872. <http://doi.org/10.1523/JNEUROSCI.1640-08.2008>
- Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*, 7(1), 105–125. doi:10.1037/1082-989X.7.1.105
- Morris, J. S., Öhman, A., & Dolan, R. J. (1999). A subcortical pathway to the right amygdala mediating ‘unseen’ fear. *Proceedings of the National Academy of Sciences of the United States of America*, 96(4), 1680–1685.
- Moser, J. S., Becker, M. W., & Moran, T. P. (2012). Enhanced attentional capture in trait anxiety. *Emotion (Washington, D.C.)*, 12(2), 213–216. <http://doi.org/10.1037/a0026156>
- Motoyoshi, I., & Hayakawa, S. (2010). Adaptation-induced blindness to sluggish stimuli. *Journal of Vision*, 10(2), 16.1–8. <http://doi.org/10.1167/10.2.16>
- Murphy, S. E., Downham, C., Cowen, P. J., & Harmer, C. J. (2008). Direct effects of diazepam on emotional processing in healthy volunteers. *Psychopharmacology*, 199(4), 503–513. <http://doi.org/10.1007/s00213-008-1082-2>

- Nesse, R. M. (1999). Proximate and evolutionary studies of anxiety, stress and depression: Synergy at the interface. *Neuroscience and Biobehavioral Reviews*, 23, 895–903. [http://dx.doi.org/10.1016/S0149-7634\(99\)00023-8](http://dx.doi.org/10.1016/S0149-7634(99)00023-8)
- Nguyen, V. A., Freeman, A. W., & Alais, D. (2003). Increasing depth of binocular rivalry suppression along two visual pathways. *Vision Research*, 43(19), 2003–2008.
- Nomura, M., Ohira, H., Haneda, K., Iidaka, T., Sadato, N., Okada, T., & Yonekura, Y. (2004). Functional association of the amygdala and ventral prefrontal cortex during cognitive evaluation of facial expressions primed by masked angry faces: an event-related fMRI study. *NeuroImage*, 21(1), 352–363. <http://doi.org/10.1016/j.neuroimage.2003.09.021>
- Ohman, A. (2005). The role of the amygdala in human fear: automatic detection of threat. *Psychoneuroendocrinology*, 30(10), 953–958. <http://doi.org/10.1016/j.psychneuen.2005.03.019>
- Ohman, A., Carlsson, K., Lundqvist, D., & Ingvar, M. (2007). On the unconscious subcortical origin of human fear. *Physiology & Behavior*, 92(1–2), 180–185. <http://doi.org/10.1016/j.physbeh.2007.05.057>
- Ohman, A., & Mineka, S. (2001). Fears, phobias, and preparedness: toward an evolved module of fear and fear learning. *Psychological Review*, 108(3), 483–522.
- Ohman, A., & Soares, J. J. (1994). ‘Unconscious anxiety’: phobic responses to masked stimuli. *Journal of Abnormal Psychology*, 103(2), 231–240.
- Ohman, A., & Soares, J. J. (1998). Emotional conditioning to masked stimuli: expectancies for aversive outcomes following nonrecognized fear-relevant stimuli. *Journal of Experimental Psychology. General*, 127(1), 69–82.
- Olsson, A., & Phelps, E. A. (2007). Social learning of fear. *Nature Neuroscience*, 10(9), 1095–1102. <http://doi.org/10.1038/nn1968>
- Ooi, T. L., & He, Z. J. (1999). Binocular rivalry and visual awareness: the role of attention. *Perception*, 28(5), 551–574.
- Padmala, S., & Pessoa, L. (2008). Affective Learning Enhances Visual Detection and Responses in Primary Visual Cortex. *The Journal of Neuroscience*, 28(24), 6202–6210. <http://doi.org/10.1523/JNEUROSCI.1233-08.2008>
- Panksepp, J. (1998). *Affective neuroscience: The foundations of human and animal emotions*. Oxford university press.
- Papez, J. W. (1937) A proposed mechanism of emotion. *Archives of Neurology & Psychiatry*, 38, 725–743.
- Pasley, B. N., Mayes, L. C., & Schultz, R. T. (2004). Subcortical discrimination of unperceived objects during binocular rivalry. *Neuron*, 42(1), 163–172.
- Pastore, R. E., Crawley, E. J., Berens, M. S., & Skelly, M. A. (2003). ‘Nonparametric’ A’ and other modern misconceptions about signal detection theory. *Psychonomic Bulletin & Review*, 10(3), 556–569.
- Pegna, A. J., Khateb, A., Lazeyras, F., & Seghier, M. L. (2005). Discriminating emotional faces without primary visual cortices involves the right amygdala. *Nature Neuroscience*, 8(1), 24–25. <http://doi.org/10.1038/nn1364>
- Peira, N., Golkar, A., Ohman, A., Anders, S., & Wiens, S. (2012). Emotional responses in spider fear are closely related to picture awareness. *Cognition & Emotion*, 26(2), 252–260. <http://doi.org/10.1080/02699931.2011.579087>

- Pessoa, L. (2005). To what extent are emotional visual stimuli processed without attention and awareness? *Current Opinion in Neurobiology*, 15(2), 188–196. <http://doi.org/10.1016/j.conb.2005.03.002>
- Pessoa, L. (2008). On the relationship between emotion and cognition. *Nature Reviews Neuroscience*, 9(2), 148–158. <http://doi.org/10.1038/nrn2317>
- Pessoa, L. (2010). Emotion and attention effects: is it all a matter of timing? Not yet. *Frontiers in Human Neuroscience*, 4, 172. <http://doi.org/10.3389/fnhum.2010.00172>
- Pessoa, L., & Adolphs, R. (2010). Emotion processing and the amygdala: from a ‘low road’ to ‘many roads’ of evaluating biological significance. *Nature Reviews Neuroscience*, 11(11), 773–783. <http://doi.org/10.1038/nrn2920>
- Pessoa, L., Japee, S., Sturman, D., & Ungerleider, L. G. (2006). Target visibility and visual awareness modulate amygdala responses to fearful faces. *Cerebral Cortex (New York, N.Y.: 1991)*, 16(3), 366–375. <http://doi.org/10.1093/cercor/bhi115>
- Pessoa, L., Japee, S., & Ungerleider, L. G. (2005). Visual awareness and the detection of fearful faces. *Emotion (Washington, D.C.)*, 5(2), 243–247. <http://doi.org/10.1037/1528-3542.5.2.243>
- Phillips, M. L., Williams, L. M., Heining, M., Herba, C. M., Russell, T., Andrew, C., ... Gray, J. A. (2004). Differential neural responses to overt and covert presentations of facial expressions of fear and disgust. *NeuroImage*, 21(4), 1484–1496. <http://doi.org/10.1016/j.neuroimage.2003.12.013>
- Piech, R. M., McHugo, M., Smith, S. D., Dukic, M. S., Van Der Meer, J., Abou-Khalil, B., & Zald, D. H. (2010). Fear-enhanced visual search persists after amygdala lesions. *Neuropsychologia*, 48(12), 3430–3435. <http://doi.org/10.1016/j.neuropsychologia.2010.07.009>
- Piech, R. M., McHugo, M., Smith, S. D., Dukic, M. S., Van Der Meer, J., Abou-Khalil, B., ... Zald, D. H. (2011). Attentional capture by emotional stimuli is preserved in patients with amygdala lesions. *Neuropsychologia*, 49(12), 3314–3319. <http://doi.org/10.1016/j.neuropsychologia.2011.08.004>
- Posner, M. I. (2012). *Attention in a social world*. New York, NY: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199791217.001.0001>
- Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology*, 109(2), 160–174.
- Poulton, R., Davies, S., Menzies, R. G., Langley, J. D., & Silva, P. A. (1998). Evidence for a non-associative model of the acquisition of a fear of heights. *Behaviour Research and Therapy*, 36(5), 537–544.
- Qian, W., Meng, Q., Chen, L., & Zhou, K. (2012). Emotional Modulation of the Attentional Blink Is Awareness-Dependent. *PLoS ONE*, 7(9), e46394. <http://doi.org/10.1371/journal.pone.0046394>
- Raes, A. K., Koster, E. H. W., Van Damme, S., Fias, W., & De Raedt, R. (2010). Aversive conditioning under conditions of restricted awareness: effects on spatial cueing. *Quarterly Journal of Experimental Psychology (2006)*, 63(12), 2336–2358. <http://doi.org/10.1080/17470218.2010.492995>
- Raio, C. M., Carmel, D., Carrasco, M., & Phelps, E. A. (2012). Nonconscious fear is quickly acquired but swiftly forgotten. *Current Biology*, 22(12), R477–R479. <http://doi.org/10.1016/j.cub.2012.04.023>
- R L De Valois, & Valois, K. K. D. (1980). Spatial Vision. *Annual Review of Psychology*, 31(1), 309–341. <http://doi.org/10.1146/annurev.ps.31.020180.001521>
- Rapee, R. M., & Heimberg, R. G. (1997). A cognitive-behavioral model of anxiety in social phobia. *Behaviour Research and Therapy*, 35(8), 741–756. [http://doi.org/10.1016/S0005-7967\(97\)00022-3](http://doi.org/10.1016/S0005-7967(97)00022-3)

- Raudenbush, S. W. (1994). Random effects models. In H. Cooper & L. V. Hedges (Eds.), *The handbook of research synthesis* (pp. 301–321). New York, NY: Russell Sage Foundation.
- Reagh, Z. M., & Knight, D. C. (2013). Negative, but not positive emotional images modulate the startle response independent of conscious awareness. *Emotion (Washington, D.C.)*, 13(4), 782–791. <http://doi.org/10.1037/a0032286>
- Reber, A. S. (1992). The cognitive unconscious: An evolutionary perspective. *Consciousness and Cognition*, 1(2), 93–133. [http://doi.org/10.1016/1053-8100\(92\)90051-B](http://doi.org/10.1016/1053-8100(92)90051-B)
- Reingold, E. M. (2004). Unconscious perception: assumptions and interpretive difficulties. *Consciousness and Cognition*, 13(1), 117–122. <http://doi.org/10.1016/j.concog.2003.11.002>
- Ritchie, K., Bannerman, R. L., & Sahraie, A. (2011). The effect of fear in the periphery in binocular rivalry. *Perception*, 40(12), 1395–1401.
- Rolls, E. T., Tovée, M. J., & Panzeri, S. (1999). The neurophysiology of backward visual masking: information analysis. *Journal of Cognitive Neuroscience*, 11(3), 300–311.
- Rosenthal, R. (1991). *Meta-analytic procedures for social research* (rev. ed.). Newbury Park, CA: SAGE Publications. <http://dx.doi.org/10.4135/9781412984997>
- Rothkirch, M., Stein, T., Sekutowicz, M., & Sterzer, P. (2012). A direct oculomotor correlate of unconscious visual processing. *Current Biology: CB*, 22(13), R514–515. <http://doi.org/10.1016/j.cub.2012.04.046>
- Ruiz-Padial, E., & Vila, J. (2007). Fearful and sexual pictures not consciously seen modulate the startle reflex in human beings. *Biological Psychiatry*, 61(8), 996–1001. <http://doi.org/10.1016/j.biopsych.2006.08.046>
- Rutishauser, U., Tudusciuc, O., Neumann, D., Mamelak, A. N., Heller, A. C., Ross, I. B., ... Adolphs, R. (2011). Single-Unit Responses Selective for Whole Faces in the Human Amygdala. *Current Biology*, 21(19), 1654–1660. <http://doi.org/10.1016/j.cub.2011.08.035>
- Sandberg, K., Timmermans, B., Overgaard, M., & Cleeremans, A. (2010). Measuring consciousness: Is one measure better than the other? *Consciousness and Cognition*, 19(4), 1069–1078. <http://doi.org/10.1016/j.concog.2009.12.013>
- Schmid, M. C., Mrowka, S. W., Turchi, J., Saunders, R. C., Wilke, M., Peters, A. J., ... Leopold, D. A. (2010). Blindsight depends on the lateral geniculate nucleus. *Nature*, 466(7304), 373–377. <http://doi.org/10.1038/nature09179>
- Schmidt, F. L., OH, I.-S., & Hayes, T. L. (2009). Fixed- versus random- effects models in meta-analysis: Model properties and an empirical comparison of differences in results. *British Journal of Mathematical and Statistical Psychology*, 62, 97–128. <http://dx.doi.org/10.1348/000711007X255327>
- Schmidt, T., & Vorberg, D. (2006). Criteria for unconscious cognition: three types of dissociation. *Perception & Psychophysics*, 68(3), 489–504.
- Schroeder, C. E., Mehta, A. D., & Givre, S. J. (1998). A spatiotemporal profile of visual system activation revealed by current source density analysis in the awake macaque. *Cerebral Cortex*, 8(7), 575–592. <http://doi.org/10.1093/cercor/8.7.575>
- Schultheiss, O. C., & Hale, J. A. (2006). Implicit Motives Modulate Attentional Orienting to Facial Expressions of Emotion. *Motivation and Emotion*, 31(1), 13–24. <http://doi.org/10.1007/s11031-006-9042-9>

- Schupp, H. T., Ohman, A., Junghöfer, M., Weike, A. I., Stockburger, J., & Hamm, A. O. (2004). The facilitated processing of threatening faces: an ERP analysis. *Emotion (Washington, D.C.)*, 4(2), 189–200. <http://doi.org/10.1037/1528-3542.4.2.189>
- Schurger, A., Pereira, F., Treisman, A., & Cohen, J. D. (2010). Reproducibility distinguishes conscious from nonconscious neural representations. *Science (New York, N.Y.)*, 327(5961), 97–99. <http://doi.org/10.1126/science.1180029>
- Seligman, M. E. (1971). Phobias and preparedness. *Behavior Therapy*, 2(3), 307–320. [http://doi.org/10.1016/S0005-7894\(71\)80064-3](http://doi.org/10.1016/S0005-7894(71)80064-3)
- Shimaoka, D., & Kaneko, K. (2011). Dynamical systems modeling of Continuous Flash Suppression. *Vision Research*, 51(6), 521–528. <http://doi.org/10.1016/j.visres.2011.01.009>
- Shin, K., Stolte, M., & Chong, S. C. (2009). The effect of spatial attention on invisible stimuli. *Attention, Perception & Psychophysics*, 71(7), 1507–1513. <http://doi.org/10.3758/APP.71.7.1507>
- Simons, D. J., & Rensink, R. A. (2005). Change blindness: past, present, and future. *Trends in Cognitive Sciences*, 9(1), 16–20. <http://doi.org/10.1016/j.tics.2004.11.006>
- Sincich, L. C., Park, K. F., Wohlgemuth, M. J., & Horton, J. C. (2004). Bypassing V1: a direct geniculate input to area MT. *Nature Neuroscience*, 7(10), 1123–1128. <http://doi.org/10.1038/nn1318>
- Singer, N., Eapen, M., Grillon, C., Ungerleider, L. G., & Hendler, T. (2012). Through the eyes of anxiety: Dissecting threat bias via emotional-binocular rivalry. *Emotion (Washington, D.C.)*, 12(5), 960–969. <http://doi.org/10.1037/a0027070>
- Sklar, A. Y., Levy, N., Goldstein, A., Mandel, R., Maril, A., & Hassin, R. R. (2012). Reading and doing arithmetic nonconsciously. *Proceedings of the National Academy of Sciences*, 109(48), 19614–19619. <http://doi.org/10.1073/pnas.1211645109>
- Skrandies, W. (1987). The upper and lower visual field of man: Electrophysiological and functional differences. In *Progress in sensory physiology* (pp. 1-93). Springer Berlin Heidelberg.
- Smith, A. T., Singh, K. D., Williams, A. L., & Greenlee, M. W. (2001). Estimating receptive field size from fMRI data in human striate and extrastriate visual cortex. *Cerebral Cortex (New York, N.Y.: 1991)*, 11(12), 1182–1190.
- Snodgrass, M., & Shevrin, H. (2006). Unconscious inhibition and facilitation at the objective detection threshold: replicable and qualitatively different unconscious perceptual effects. *Cognition*, 101(1), 43–79. <http://doi.org/10.1016/j.cognition.2005.06.006>
- Stein, T., Seymour, K., Hebart, M. N., & Sterzer, P. (2014). Rapid fear detection relies on high spatial frequencies. *Psychological Science*, 25(2), 566–574. <http://doi.org/10.1177/0956797613512509>
- Stein, T., & Sterzer, P. (2012). Not just another face in the crowd: detecting emotional schematic faces during continuous flash suppression. *Emotion (Washington, D.C.)*, 12(5), 988–996. <http://doi.org/10.1037/a0026944>
- Stein, T., & Sterzer, P. (2014). Unconscious processing under interocular suppression: getting the right measure. *Frontiers in Psychology*, 5. <http://doi.org/10.3389/fpsyg.2014.00387>
- Stein, T., Sterzer, P., & Peelen, M. V. (2012). Privileged detection of conspecifics: evidence from inversion effects during continuous flash suppression. *Cognition*, 125(1), 64–79. <http://doi.org/10.1016/j.cognition.2012.06.005>

- Stewart, L. H., Ajina, S., Getov, S., Bahrami, B., Todorov, A., & Rees, G. (2012). Unconscious Evaluation of Faces on Social Dimensions. *Journal of Experimental Psychology. General*, 141(4), 715–727. <http://doi.org/10.1037/a0027950>
- Straube, T., Dietrich, C., Mothes-Lasch, M., Mentzel, H.-J., & Miltner, W. H. R. (2010). The volatility of the amygdala response to masked fearful eyes. *Human Brain Mapping*, 31(10), 1601–1608. <http://doi.org/10.1002/hbm.20960>
- Susskind, J. M., Lee, D. H., Cusi, A., Feiman, R., Grabski, W., & Anderson, A. K. (2008). Expressing fear enhances sensory acquisition. *Nature Neuroscience*, 11(7), 843–850. <http://doi.org/10.1038/nn.2138>
- Sutton, T. M., & Altarriba, J. (2011). The automatic activation and perception of emotion in word processing: Evidence from a modified dot probe paradigm. *Journal of Cognitive Psychology*, 23(6), 736–747.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs: General and Applied*, 74(11), 1–29. <http://doi.org/10.1037/h0093759>
- Spielberger, C. D., Gorsuch, R. L., Lushene, R., Vagg, P. R., & Jacobs, G. A. (1983). *Manual for the State-Trait Anxiety Inventory*. Palo Alto, CA: Consulting Psychologists Press.
- Sylvers, P. D., Brennan, P. A., & Lilienfeld, S. O. (2011). Psychopathic traits and preattentive threat processing in children: a novel test of the fearlessness hypothesis. *Psychological Science*, 22(10), 1280–1287. <http://doi.org/10.1177/0956797611420730>
- Szczepanowski, R., & Pessoa, L. (2007). Fear perception: Can objective and subjective awareness measures be dissociated? *Journal of Vision*, 7(4), 10–10. <http://doi.org/10.1167/7.4.10>
- Tamietto, M., Castelli, L., Vighetti, S., Perozzo, P., Geminiani, G., Weiskrantz, L., & Gelder, B. de. (2009). Unseen facial and bodily expressions trigger fast emotional reactions. *Proceedings of the National Academy of Sciences*, 106(42), 17661–17666. <http://doi.org/10.1073/pnas.0908994106>
- Tamietto, M., & de Gelder, B. (2010). Neural bases of the non-conscious perception of emotional signals. *Nature Reviews Neuroscience*, 11(10), 697–709. <http://doi.org/10.1038/nrn2889>
- Tamietto, M., Pullens, P., de Gelder, B., Weiskrantz, L., & Goebel, R. (2012). Subcortical Connections to Human Amygdala and Changes following Destruction of the Visual Cortex. *Current Biology*, 22(15), 1449–1455. <http://doi.org/10.1016/j.cub.2012.06.006>
- Tan, J., Ma, Z., Gao, X., Wu, Y., & Fang, F. (2011). Gender Difference of Unconscious Attentional Bias in High Trait Anxiety Individuals. *PLoS ONE*, 6(5), e20305. <http://doi.org/10.1371/journal.pone.0020305>
- Thomason, M. E., Henry, M. L., Paul Hamilton, J., Joormann, J., Pine, D. S., Ernst, M., ... Gotlib, I. H. (2010). Neural and behavioral responses to threatening emotion faces in children as a function of the short allele of the serotonin transporter gene. *Biological Psychology*, 85(1), 38–44. <http://doi.org/10.1016/j.biopsycho.2010.04.009>
- Tong, F. (2003). Primary visual cortex and visual awareness. *Nature Reviews. Neuroscience*, 4(3), 219–229. <http://doi.org/10.1038/nrn1055>
- Tong, F., & Engel, S. A. (2001). Interocular rivalry revealed in the human cortical blind-spot representation. *Nature*, 411(6834), 195–199. <http://doi.org/10.1038/35075583>
- Tong, F., Meng, M., & Blake, R. (2006). Neural bases of binocular rivalry. *Trends in Cognitive Sciences*, 10(11), 502–511. <http://doi.org/10.1016/j.tics.2006.09.003>

- Tottenham, N., Tanaka, J. W., Leon, A. C., McCarry, T., Nurse, M., Hare, T. A., ... Nelson, C. (2009). The NimStim set of facial expressions: Judgments from untrained research participants. *Psychiatry Research*, 168(3), 242–249. <http://doi.org/10.1016/j.psychres.2008.05.006>
- Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8(8), 1096–1101. <http://doi.org/10.1038/nn1500>
- Tsuchiya, N., Koch, C., Gilroy, L. A., & Blake, R. (2006). Depth of interocular suppression associated with continuous flash suppression, flash suppression, and binocular rivalry. *Journal of Vision*, 6(10), 1068–1078. <http://doi.org/10.1167/6.10.6>
- Tsuchiya, N., Moradi, F., Felsen, C., Yamazaki, M., & Adolphs, R. (2009). Intact rapid detection of fearful faces in the absence of the amygdala. *Nature Neuroscience*, 12(10), 1224–1225. <http://doi.org/10.1038/nn.2380>
- Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2016). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, 23(1), 87–102. <http://doi.org/10.3758/s13423-015-0892-6>
- Van Boxtel, J. J. A., Tsuchiya, N., & Koch, C. (2010). Consciousness and attention: on sufficiency and necessity. *Consciousness Research*, 1, 217. <http://doi.org/10.3389/fpsyg.2010.00217>
- Van Buuren, S. (2012). *Flexible Imputation of Missing Data*. CRC Press.
- Van Buuren, S., Groothuis-Oudshoorn, K. (2011). mice: Multivariate Imputation by Chained Equations in R. *Journal of Statistical Software*, 45, 1–67.
- Van den Bussche, E., Van den Noortgate, W., & Reynvoet, B. (2009). Mechanisms of masked priming: a meta-analysis. *Psychological Bulletin*, 135(3), 452–477. <http://doi.org/10.1037/a0015329>
- van den Hout, M. A., De Jong, P., & Kindt, M. (2000). Masked fear words produce increased SCRs: an anomaly for Ohman's theory of pre-attentive processing in anxiety. *Psychophysiology*, 37(3), 283–288.
- van Loon, A. M., Knapen, T., Scholte, H. S., St. John-Saaltink, E., Donner, T. H., & Lamme, V. A. F. (2013). GABA Shapes the Dynamics of Bistable Perception. *Current Biology*, 23(9), 823–827. <http://doi.org/10.1016/j.cub.2013.03.067>
- Vanrullen, R. (2011). Four common conceptual fallacies in mapping the time course of recognition. *Frontiers in Psychology*, 2, 365. <http://doi.org/10.3389/fpsyg.2011.00365>
- Vermeulen, N., Godefroid, J., & Mermillod, M. (2009). Emotional Modulation of Attention: Fear Increases but Disgust Reduces the Attentional Blink. *PLoS ONE*, 4(11), e7924. <http://doi.org/10.1371/journal.pone.0007924>
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36, 1–48. <http://dx.doi.org/10.18637/jss.v036.i03>
- Vizueta, N., Patrick, C. J., Jiang, Y., Thomas, K. M., & He, S. (2012). Dispositional fear, negative affectivity, and neuroimaging response to visually suppressed emotional faces. *NeuroImage*, 59(1), 761–771. <http://doi.org/10.1016/j.neuroimage.2011.07.015>
- Vuilleumier, P. (2005). How brains beware: neural mechanisms of emotional attention. *Trends in Cognitive Sciences*, 9(12), 585–594. <http://doi.org/10.1016/j.tics.2005.10.011>

- Vuilleumier, P., Armony, J. L., Driver, J., & Dolan, R. J. (2003). Distinct spatial frequency sensitivities for processing faces and emotional expressions. *Nature Neuroscience*, 6(6), 624–631.
<http://doi.org/10.1038/nn1057>
- Wardle, S. G., Cass, J., Brooks, K. R., & Alais, D. (2010). Breaking camouflage: Binocular disparity reduces contrast masking in natural images. *Journal of vision*, 10(14), 38–38.
- Whalen, P. J. (2007). The uncertainty of it all. *Trends in Cognitive Sciences*, 11(12), 499–500.
<http://doi.org/10.1016/j.tics.2007.08.016>
- Whalen, P. J., Kagan, J., Cook, R. G., Davis, F. C., Kim, H., Polis, S., ... Johnstone, T. (2004). Human Amygdala Responsivity to Masked Fearful Eye Whites. *Science*, 306(5704), 2061–2061.
<http://doi.org/10.1126/science.1103617>
- 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. *The Journal of Neuroscience*, 18(1), 411–418.
- Wheatstone, C. (1838) On some remarkable, and hitherto unobserved, phenomena of binocular vision. *Philosophical Transactions of the Royal Society London*, 128, 371–394.
- Whitney, D., & Levi, D. M. (2011). Visual Crowding: a fundamental limit on conscious perception and object recognition. *Trends in Cognitive Sciences*, 15(4), 160–168.
<http://doi.org/10.1016/j.tics.2011.02.005>
- Wiens, S. (2006). Remain aware of awareness. *Nature Reviews Neuroscience*, 6(3).
<http://doi.org/10.1038/nrn1630-c1>
- Wiens, S. (2008). Concepts of visual consciousness and their measurement. *Advances in Cognitive Psychology*, 3(1–2), 349–359. <http://doi.org/10.2478/v10053-008-0035-y>
- Wiens, S., Fransson, P., Dietrich, T., Lohmann, P., Ingvar, M., & Arne, Ö. (2004). Keeping It Short A Comparison of Methods for Brief Picture Presentation. *Psychological Science*, 15(4), 282–285.
<http://doi.org/10.1111/j.0956-7976.2004.00667.x>
- Wiens, S., Katkin, E. S., & Ohman, A. (2003). Effects of trial order and differential conditioning on acquisition of differential shock expectancy and skin conductance conditioning to masked stimuli. *Psychophysiology*, 40(6), 989–997.
- Wiens, S., & Ohman, A. (2002). Unawareness is more than a chance event: comment on Lovibond and Shanks (2002). *Journal of Experimental Psychology. Animal Behavior Processes*, 28(1), 27–31.
- Wiens, S., & Öhman, A. (2007). Probing unconscious emotional processes: On becoming a successful masketeer. In J. A. Coan & J. J. B. Allen (Eds.) *Handbook of emotion elicitation and assessment*. New York: Oxford University Press, pp. 65–90.
- Willenbockel, V., Lepore, F., Nguyen, D. K., Bouthillier, A., & Gosselin, F. (2012). Spatial Frequency Tuning during the Conscious and Non-Conscious Perception of Emotional Facial Expressions – An Intracranial ERP Study. *Frontiers in Psychology*, 3. <http://doi.org/10.3389/fpsyg.2012.00237>
- Wilke, M., Logothetis, N. K., & Leopold, D. A. (2003). Generalized Flash Suppression of Salient Visual Targets. *Neuron*, 39(6), 1043–1052. <http://doi.org/10.1016/j.neuron.2003.08.003>
- Williams, M. A., Morris, A. P., McGlone, F., Abbott, D. F., & Mattingley, J. B. (2004). Amygdala responses to fearful and happy facial expressions under conditions of binocular suppression. *The Journal of*

- Neuroscience: The Official Journal of the Society for Neuroscience*, 24(12), 2898–2904.
<http://doi.org/10.1523/JNEUROSCI.4977-03.2004>
- Williams, J. M. G., Watts, F. N., MacLeod, C., & Mathews, A. (1988). *Cognitive psychology and emotional disorders*. John Wiley & Sons.
- Wirth, M. M., & Schultheiss, O. C. (2007). Basal testosterone moderates responses to anger faces in humans. *Physiology & Behavior*, 90(2–3), 496–505. <http://doi.org/10.1016/j.physbeh.2006.10.016>
- Wolfe, J. M. (1984). Reversing ocular dominance and suppression in a single flash. *Vision Research*, 24(5), 471–478. [http://doi.org/10.1016/0042-6989\(84\)90044-0](http://doi.org/10.1016/0042-6989(84)90044-0)
- Yang, E., & Blake, R. (2012). Deconstructing continuous flash suppression. *Journal of Vision*, 12(3), 8. <http://doi.org/10.1167/12.3.8>
- Yang, E., Zald, D. H., & Blake, R. (2007). Fearful expressions gain preferential access to awareness during continuous flash suppression. *Emotion (Washington, D.C.)*, 7(4), 882–886. <http://doi.org/10.1037/1528-3542.7.4.882>
- Yang, Y.-H., & Yeh, S.-L. (2011). Accessing the meaning of invisible words. *Consciousness and Cognition*, 20(2), 223–233. <http://doi.org/10.1016/j.concog.2010.07.005>
- Yoon, K. L., Hong, S. W., Joormann, J., & Kang, P. (2009). Perception of facial expressions of emotion during binocular rivalry. *Emotion (Washington, D.C.)*, 9(2), 172–182. <http://doi.org/10.1037/a0014714>
- Yu, K., & Blake, R. (1992). Do recognizable figures enjoy an advantage in binocular rivalry? *Journal of Experimental Psychology. Human Perception and Performance*, 18(4), 1158–1173.
- Yuval-Greenberg, S., & Heeger, D. J. (2013). Continuous flash suppression modulates cortical activity in early visual cortex. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, 33(23), 9635–9643. <http://doi.org/10.1523/JNEUROSCI.4612-12.2013>
- Zadbood, A., Lee, S.-H., & Blake, R. (2011). Stimulus fractionation by interocular suppression. *Frontiers in Human Neuroscience*, 5, 135. <http://doi.org/10.3389/fnhum.2011.00135>